Electronic Design®

THE AUTHORITY ON EMERGING TECHNOLOGIES FOR DESIGN SOLUTIONS

MAY/JUNE 2019
electronicdesign.com

ROVING ROBOTS

From firefighting to supermarkets, robots are popping up everywhere p12

$10.00

Altech® PUSH-IN WIRE

Printed Circuit Board Terminal Blocks

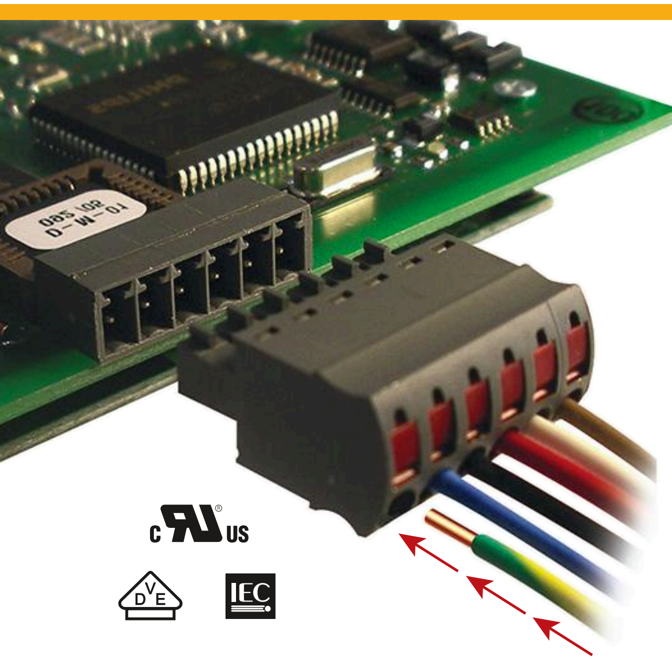

Time Saving
Push-In
TECHNOLOGY
Easy wire insertion
No tool required !

PUSH-IN PLUGS

- Unique ergonomic grip design for quick and easy mounting
- 3.50, 3.81, 5.00 and 5.08mm pin spacing
- Screw flange version for stability
- Imprinting and coding available

PUSH-IN POWER BLOCKS

- Rated current:
  20A **(AK4200)**
  35A **(AK4210)**
- Double row pins
- Many colors allow easy identification

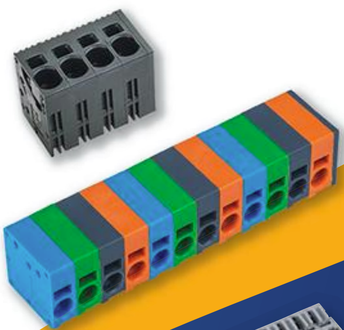

PUSH-IN PCB TERMINAL BLOCKS

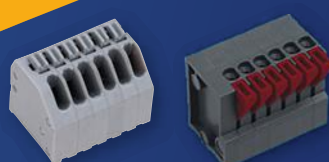

- 2.50 / 2.54mm micro terminals
- Vertical / horizontal types
- Lever version available
- Double row pins for high stability

Complete line of Plug-In Connectors

- All pin spacings
- Proven screw clamp design
- Compatible headers

Excellent Quality at Competitive Prices

FREE Demo Kits available

Altech Corp.®

Your Source for Automation & Control Components

AltechCorp.com/PCB
908.806.9400 for all the details & catalog

IN THIS ISSUE

FEATURES

COLUMNS & DEPARTMENTS

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 bi-monthly in Jan/Feb, Mar/Apr, May/June, July/Aug, Sept/Oct and Nov/Dec by Informa Media Inc., 9800 Metcalf Ave., Overland Park, KS 66212-2216. Paid rates for a one-year subscription are as follows: $60 U.S., $90 Canada, $120 International. Periodicals postage paid at Kansas City, MO, and additional mailing offices. Editorial and advertising addresses: ELECTRONIC DESIGN, 605 Third Avenue, New York, NY 10158. Telephone (212) 204-4200. Printed in U.S.A. Title registered in U.S. Patent Office. For subscriber services or to order single copies, write to Electronic Design, PO Box 2100, Skokie, IL 60076. POSTMASTER: Send change of address to Electronic Design, PO Box 2100, Skokie, IL 60076. Canadian Post Publications Mail agreement No. 40612608. Canada return address: IMEX Global Solutions, P.O. Box 25542, London, ON N6C 6B2.

ON ELECTRONICDESIGN.COM

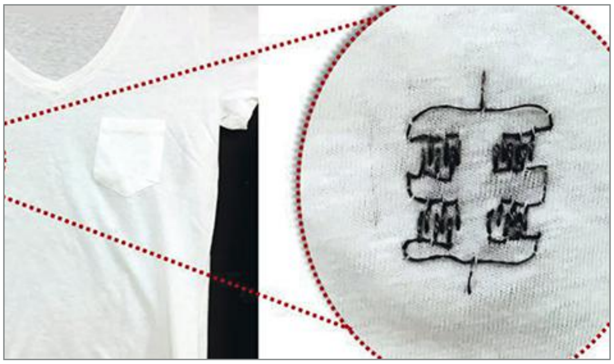

Charge-Storing Threads Turn Fabric into Supercapacitor

Researchers combined conductive threads with a polymer film to create a flexible mesh of electrodes on a textile backing that functions as a supercapacitor.

https://www.electronicdesign.com/power/charge-storing-threads-turn-fabric-supercapacitor

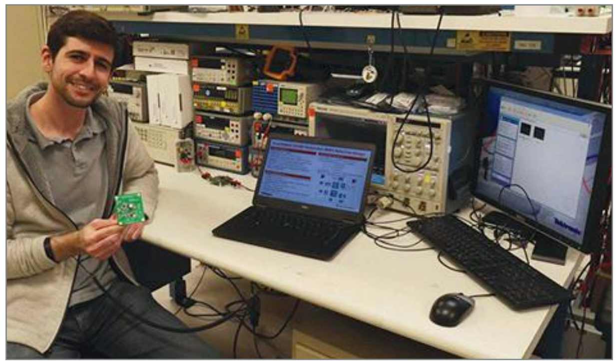

What Goes Into a "Simple" Buck-Regulator Chip?

With a good company, you're buying much more than the chip. You're also obtaining the confidence that the chip was in good hands, because it's usually designed by experts in the industry.

https://www.electronicdesign.com/power/what-goes-simple-buck-regulator-chip

How Waterproofing Will Change the Smartphone Design Paradigm

Waterproofing issues, which are garnering more attention than ever, can affect the design of embedded devices such as smartphones.

https://www.electronicdesign.com/industrial-automation/how-waterproofing-will-change-smartphone-design-paradigm

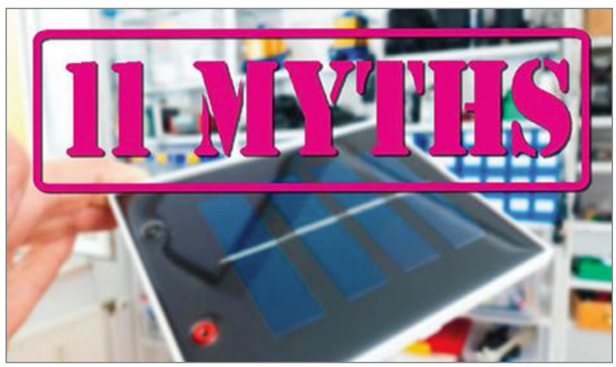

11 Myths About Energy Harvesting

Is energy harvesting expensive? Hard to implement? We'll dispel some of the myths surrounding the technology.

https://www.electronicdesign.com/power/11-myths-about-energy-harvesting


THINK™
HIGH PERFORMANCE
Color Keyed Automotive Blade Fuse Holders
KEYSTONE
06
• Available in colors to match fuse amperage colors for simplified fuse location • Available for Mini, Low Profile Mini and Standard automotive blade fuses • Fully insulated contacts • Available in Thru Hole Mount (THM) rated for 30 Amps and Surface Mount (SMT) rated for 20 Amps • Available in Bulk or on Tape and Reel
IT'S WHAT'S ON THE INSIDE THAT COUNTS
KEYSTONE®
ELECTRONICS CORP.
www.keyelco.com • (516) 328-7500 • (800) 221-5510
REQUEST CATALOG!

ElectronicDesign.

MAY/JUNE 2019

EDITORIAL

GROUP CONTENT DIRECTOR: **KAREN FIELD** karen.field@informa.com
EDITOR/CONTENT DIRECTOR: **BILL WONG** bill.wong@informa.com
SENIOR STAFF WRITER: **JAMES MORRA** james.morra@informa.com
TECHNICAL EDITOR: **CHRIS DeMARTINO** chris.demartino@informa.com
CONTRIBUTING EDITOR: **LOUIS E. FRENZEL** lou.frenzel@informa.com
ASSOCIATE EDITOR/COMMUNITY MANAGER: **ROGER ENGELKE** roger.engelke@informa.com

ART DEPARTMENT

GROUP DESIGN DIRECTOR: **ANTHONY VITOLO** tony.vitolo@informa.com
CONTENT DESIGN SPECIALIST: **JOCELYN HARTZOG** jocelyn.hartzog@informa.com
CONTENT & DESIGN PRODUCTION MANAGER: **JULIE JANTZER-WARD** julie.jantzer-ward@informa.com

PRODUCTION

GROUP PRODUCTION MANAGER: **GREG ARAUJO** greg.araujo@informa.com
PRODUCTION MANAGER: **SUSAN POSKIN** susan.poskin@informa.com

AUDIENCE MARKETING

USER MARKETING MANAGER: **DEBBIE BRADY** debbie.brady@informa.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@informa.com

REGIONAL SALES REPRESENTATIVES:

AZ, NM, TX: **GREGORY MONTGOMERY** **T** | 480.254.5540 gregory.montgomery@informa.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@informa.com

CT, DE, MA, MD, ME, NH, NJ, NY, PA, RI, VT, EASTERN CANADA:
ELIZABETH ELDRIDGE **T** | 917-789-3012 elizabeth.eldridge@informa.com

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

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@informa.com

INFORMA REPRINTS: **WRIGHT'S MEDIA** **T** | 877.652.5295

LIST RENTALS/ SMARTREACH CLIENT SERVICES MANAGER: **MARY RALICKI** **T** | 212.204.4284 mary.ralicki@informa.com

DIGITAL

GROUP DIGITAL DIRECTOR: **RYAN MALEC** ryan.malec@informa.com

DESIGN ENGINEERING & SOURCING GROUP

EXECUTIVE DIRECTOR, CONTENT: **KAREN FIELD** karen.field@informa.com
VP OF MARKETING: **JACQUIE NIEMIEC** jacquie.niemiec@informa.com

INFORMA MEDIA INC.

605 THIRD AVENUE
NEW YORK, NY 10158 USA **T** | 212.204.4200

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

2 GHz Clock Generator

- ***Square wave clocks from DC to 2.05 GHz***
- ***Random jitter <1 ps (rms)***
- ***80 ps rise and fall times***
- ***16-digit frequency resolution***
- ***CMOS, LVDS, ECL, PECL, RS-485***
- ***Phase adjustment & time modulation***

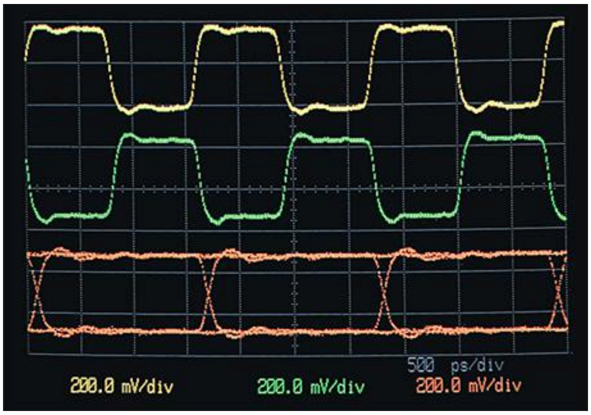

Plot shows complementary clocks and PRBS (opt. 01) outputs at 622.08 Mb/s with LVDS levels. Traces have transition times of 80 ps and jitter less than 1 ps (rms).

The CG635 generates clock signals—flawlessly. The clock signals are fast, clean and accurate, and can be set to standard logic levels.

How fast? Frequency to 2.05 GHz with rise and fall times as short as 80 ps.

How clean? Jitter is less than 1 ps and phase noise is better than –90 dBc/Hz (100 Hz offset) at 622.08 MHz.

How accurate? Using the optional rubidium timebase, aging is better than 0.0005 ppm/yr, and temperature stability is better than 0.0001 ppm.

You would expect an instrument this good to be expensive, but it isn't. You no longer have to buy an rf synthesizer to generate clock signals. The CG635 does the job better—at a fraction of the cost.

SRS
Stanford Research Systems
Phone: (408)744-9040 • info@thinkSRS.com • **www.thinkSRS.com**

Editorial

WILLIAM WONG | Senior Technology Editor
bill.wong@informa.com

Bumping into ROBOTS

When developing your next wireless-sensor-network application, it pays to rightsize the network with the optimal number of sensors and consider variables such as the network medium, mains, or battery power sources, wireless protocol options, and network configuration.

Have you bumped into a robot in a store lately? If you don't buy all your produce and products online, then you may run into one while shopping at a brick-and-mortar store like Walmart *(see figure).*

A number of factors have hastened this invasion of cobots. They range from frameworks like the Robot Operating System (ROS) to machine-learning (ML) inferencing to a host of smaller, low-cost sensors. Robots that used to have a couple sensors and maybe a camera or expensive LiDAR are enveloped in a mesh of sophisticated sensors these days.

Higher-resolution video, infrared, and thermal cameras are now available as are compact radar and LiDAR systems, often courtesy of the push for self-driving cars. The sensor net feeds deep-neural-network (DNN) models that provide planning software with ever-more detailed information about the local environment, including those pesky humans that tend to take offense when bumping into self-directed robots.

Walmart is only one of many companies with stores that will have robots roaming the aisles looking at inventory.

Neural networks are playing a rapidly increasing role in robotic systems, but they're a far cry from Asimov's fictional positronic brain—his three laws of robotics are still a bit in the future. Hardware acceleration is helping the adoption of ML technology, which is crucial because many applications, including robotics, require multiple ML models to process the sensor data streams. ML models are also being used for planning.

Neural networks are being employed in many components within cooperative robots, also known as cobots. Peaceful coexistence between cobots and humans will depend on how well designers can employ these tools and hardware to create cobots that perform their tasks while interacting with us and the rest of the world.

Cobots wandering around stores performing chores like inventory are just the start. There are smart suitcases that will follow you through the airport. Small delivery robots are being tested now. Some are rolling while others fly through the air. Not all cobots will be successful, but they will continue to improve and become more functional as developers take advantage of the latest hardware and software. ed

News

INTEL AGILEX FPGA

Brings CXL to Embedded, Data-Center Apps

The new line of Agilex FPGAs developed by Intel delivers an impressive array of features, including support for the new Compute Express Link (CXL) standard design that puts devices on an equal footing with host processors. The family takes advantage of the company's die-to-die interconnect, Embedded Multi-die Interconnect Bridge (EMIB), which was originally developed by Altera and is now part of Intel. EMIB allows Intel to mix EMIB-linked chiplets or tiles with support for high-speed transceivers, PCI Express (PCIe), and a quad-core ARM Cortex-A53 complex for the SoC edition *(Fig. 1)*.

The Agilex is based on Intel's 10-nm technology. It provides a 40% performance boost compared to the Intel Stratix 10 while reducing power by up to 40%. The top-end systems can deliver up to 40 TFLOPS (FP16 performance) using the new DSP blocks, and are designed to handle DDR5 and provide on-chip high bandwidth memory (HBM) along with support for off-chip Intel Optane DC persistent memory. The top end will also support 112-Gb/s transceivers as well as PCI Express Gen 5.

The EMIB support is used to provide connections from the FPGA fabric to these interfaces, but it can also be used for custom tiles. Intel's eASIC effort will allow companies to incorporate their tiles into custom Agilex FPGAs *(Fig. 2)*. The eASIC is an Agilex FPGA with customer-supplied tiles linked to the system using EMIB. This provides the flexibility of an FPGA with many of the advantages of an ASIC, but with much lower startup costs. The approach can be used to integrate many chips into one, configurable platform.

The CXL support *(Fig. 3)* is one of the major advances for Intel's FPGAs. Though it requires matching support in the host processor, it provides a common, cache-coherent, memory architecture that's built on PCI PCIe. It essentially puts all CXL devices on the same level, allowing the FGPA to gain access to shared memory and peripherals.

CXL is a new platform and Agilex will be one of the first out there to include it. Intel will have matching support in a future Xeon family, at which point CXL becomes very interesting. From a programming point of view, there's no difference between a processor and an FPGA working on the data in shared memory. It's even possible that one instance of an application will use processor-based support while another employs an FPGA implementation. The operating system could transparently utilize what's needed or available at the time.

Intel FPGAs are already used in the data center in the Alveo family, which employs PCIe interfaces. CXL also is built on PCIe, but it provides a lower-latency implementation because of the cache-coherent nature. ■

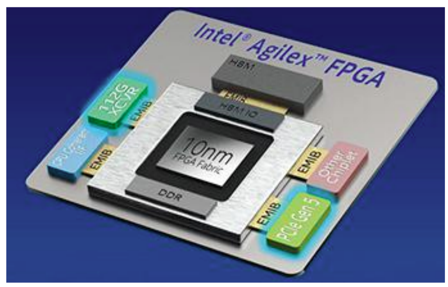

1. The Agilex FPGA uses EMIB connectivity to blend its FPGA fabric with other devices, including processors, high-speed interfaces, and memory.

2. Intel's eASIC provides a bridge between FPGA solutions and custom ASICs.

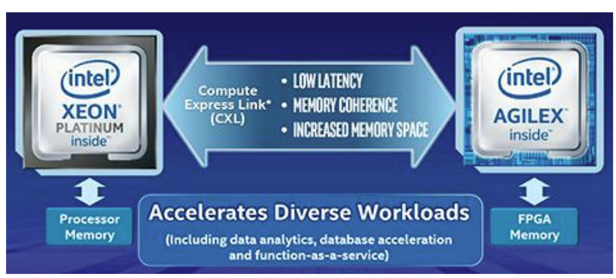

3. Agilex supports the new Compute Express Link (CXL) standard for linking devices to host processors in a cache-coherent way.

WILLIAM WONG | Senior Technology Editor

ROBOTS: Coming to a Sidewalk Near You

From sidewalks to firefighting and supermarkets to self-driving cars, robots are becoming ubiquitous.

Robots used to be hidden in assembly lines or research labs. Now you can find them in supermarkets, warehouses, and rolling down the sidewalk to deliver treats. They're also taking on dangerous chores like fighting fires. Some are controlled remotely while others are fully autonomous. New sensors, improved motors, and machine-learning (ML) support are all helping to improve what robots can do and how closely they can interact with people.

Robots are appearing more frequently in assembly lines, with robotic arms and systems from companies like those from FANUC *(Fig. 1)* providing a more flexible and functional setup. They allow assembly lines to be reconfigured quickly, given the advanced programming available as well as improved safety around people.

These days, robots aren't always sitting in fixed locations. Autonomous, mobile robots like those from KUKA *(Fig. 2)* move throughout factories often equipped with robotic arms or other actuators to do more than just roll around. Like many mobile robots, including self-driving cars, these moving platforms include a range of sensors to determine their position as well as objects in its surroundings. Technologies ranging from radar and LiDAR to thermal, 3D cameras, and ultrasonics are being used.

1. All sorts of robots, like these from FANUC, are providing more flexible assembly lines and working in closer proximity with people.

2. KUKA can combine mobile robotic platforms with robotic arms to support mobile pick-and-place operations.

SECURE MEASURES

Keeping autonomous robots in a self-contained area like a warehouse can simplify the use of robots, since they would only bump into other robots. Allowing a person to move into that space could be hazardous unless the robots were equipped with more sensors and software. Amazon is using a slightly different approach by having employees wear a special vest when moving into these more restricted areas. The vest provides location information about a person so that robots know to steer clear.

Robots operating in a secured area allows them to be simpler and move faster. Still, robots are being used in close quarters with people doing jobs that tend to be monotonous. For example, Giant Food Stores are now home to a host of Marty robots *(Fig. 3)* that roam around the stores looking for spills. The robot has a number of obstacle detection sensors, but its slow movement and the lack of any arms or protrusions makes it a relatively safe device.

Similar-sized robots are being developed to handle inventory chores. Robots that will stock shelves are generally limited to warehouses at this point, but could be used during the time a store is closed and with a limited numbers of humans around. People could also wear something like Amazon's vest and have additional controls to make it safer to work around robots.

Brain Corp.'s line of robotic solutions

3. Marty is a slow-moving robot on the lookout for spills in Giant Food stores.

includes some giant Roombas designed to handle industrial-sized cleaning chores like supermarkets *(Fig. 4)*. These don't bump into walls and furniture like tiny vacuum-cleaning robots. Instead, these floor scrubbers are designed to detect local obstacles in order to perform their function with limited or no human intervention. They typically operate when the store or area is closed so that there are few people to get in the way.

ROBOTS THAT DELIVER

Delivery robots approximately the size of shopping carts are taking to the sidewalks. Robots like Postmates' Serve *(Fig. 5)* are designed to handle local deliveries. Postmates already uses people driving cars to handle deliveries, but many are often within walking distance, perhaps enabling the use of Serve in such cases.

Under the hood of Serve is NVIDIA's Jetson AGX Xavier. It has a laser range finder on top along with other sensors spaced around the robot. Its large wheels are designed to go over curbs. They have already applied for a commercial permit in San Francisco, but it won't be long before this type of robot becomes more common.

TAKING ON DANGEROUS SITUATIONS

The terrible fire at Notre Dame Cathedral saw a number of robots in use, including Colossus *(Fig. 6)* from Shark Robotics. It has a 360-degree HD camera, thermal cameras, and NRBC sensors for detecting hazards. Colossus is designed for remote control as well as autonomous operation for hours. It can be equipped with a fire house, allowing it to get closer to a fire than would be recommended for human firefighters.

Other robots used to help control the fire included DJI Mavic Pro and Matrice M210 drones. These didn't deliver water like the Colossus, but they did fly over Notre Dame cathedral to provide infor-

4. Brain Corp. offers a number of robots that are industrial-sized Roombas for handling large cleaning chores.

5. Postmates' Serve delivery robot is controlled by NVIDIA's Jetson AGX Xavier.

6. Colossus, developed by Shark Robotics, was used to help fight the fire that gutted much of Notre Dame cathedral in Paris.

mation about the status of the fire. These stock drones, which utilized conventional cameras, were remotely piloted. Drones equipped with thermal cameras are being used in similar situations. Firefighters actually had to borrow the drones, since they didn't have any at this point.

One challenge with drones is that geofencing is being used more often. DJI's drones take this into account—the area around the cathedral was fenced, so DJI cooperated to disable that feature, allowing the drones to be used in this instance.

Sensor, ML, and robot platforms are making it possible to put robots in close proximity with people. However, tools like NVIDIA's Isaac Robot Engine simulator *(Fig. 7)* actually help develop and test robots. This type of simulation environment is also being used for self-driving car development, providing a much safer and cheaper way to test systems.

One of the challenges with these simulation environments is incorporating the sensors so that their characteristics match that of the real-world systems. Another challenge is to handle a large number of robots.

Such systems are becoming more sophisticated, too, so that they can be part of hardware-in-the-loop testing, where a real robot is being utilized with these same tools. Likewise, the simulators also support running the robotic software on a real platform like the Jetson Nano, while sensors operate in the simulated environment. ed

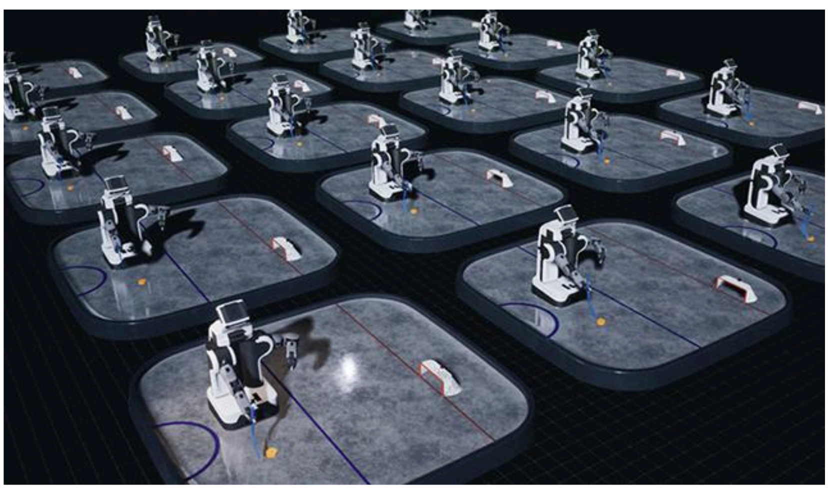

7. NVIDIA's Isaac Robot Engine simulation system can support HIL robotic testing.

LOU FRENZEL | Contributing Editor

Is Wireless Killing Our Wired Networks?

While wireless grabs all of the headlines, wired networks continue to reliably do their jobs. But, despite their reliability and relative ubiquity, some foresee the demise of the wired side.

Does the relentless expansion of wireless technologies and services signal the end of wired networks? Does the growth of the Internet of Things, forthcoming 5G, and ever-faster Wi-Fi foretell the eventual death of copper cabling? Some are predicting just this. Here are my thoughts.

THE WIRED WORLD

How many times have we wired the world? The first wired network was the telegraph, circa 1850. I'm not making this up. Telegraph lines actually came before the telephone and electrical power transmission. These didn't show up until the late 1870s through the early 1900s. All three—telegraph, telephone and electrical power businesses—eventually created massive networks and grids that remain today. The telephone system grew into a massive enterprise including fiber and DSL wiring.

But that's not all. In the 1960s, we got community TV antenna systems that morphed into the huge cable TV networks. Today these hybrid-fiber-cable (HFC) systems cover most of the U.S. In the 1970s and later, local-area networks (LANs) began connecting millions of PCs with twisted pair.

Somewhere along the way, factories, process plants, and other industrial facilities got wired with fieldbuses like HART, Modbus, PROFINET, Ethernet, and others. Office buildings, homes and even complete cities were wired with twisted-pair Ethernet. And let's certainly not forget the Internet and all the related fiber connecting the world. Now, wired networks have taken hold in cars, SUVs, and trucks. New vehicles have LIN, CAN, MOST, even Ethernet. What have I left out?

Given the huge wired infrastructure, you would think that the answer to the burning question is a solid "NO, wireless is not killing the wired world."

WIRELESS IS WINNING

However, all of the most recent networking hype, development, and application seems to be wireless-related. Wi-Fi everywhere. Massive expansion of the 5G cellular network and billions of smartphones. The public switched telephone network is fading away as subscribers abandon it for cellular service only. Bluetooth cable replacement, wireless headsets and speakers, and many others. Satellite TV distribution competing with cable TV.

Then there's the Internet of Things that's mostly wireless with its billions of sensors. Wireless utility-meter reading. And no telling how 5G will affect things. Initially, 5G broadband wireless services will challenge the DSL operators and cable companies.

With all those wireless signals floating around, it's amazing that the EMI hasn't compromised at least some of those systems. It's only going to get worse, so we should all be worried.

THE DEATH OF WIRED SYSTEMS IS PREMATURE

Mark Twain is acclaimed to have said "The reports of my death are greatly exaggerated." That's probably the case here. When you think about it, aren't many if not most of those wireless systems tied together to other networks with cables?

The cellular system can't exist without all of those attendant fiber interconnects. Server interconnects in data centers use exotic cabling like Infiniband. And isn't it wired Ethernet that ties all those Wi-Fi access points together to a server in a closet that is connected to outside fiber. Home wireless still needs that cable or DSL wiring to get internet access. Don't all wireless networks actually have a wired connection to some other system?

(Continued on page 47)

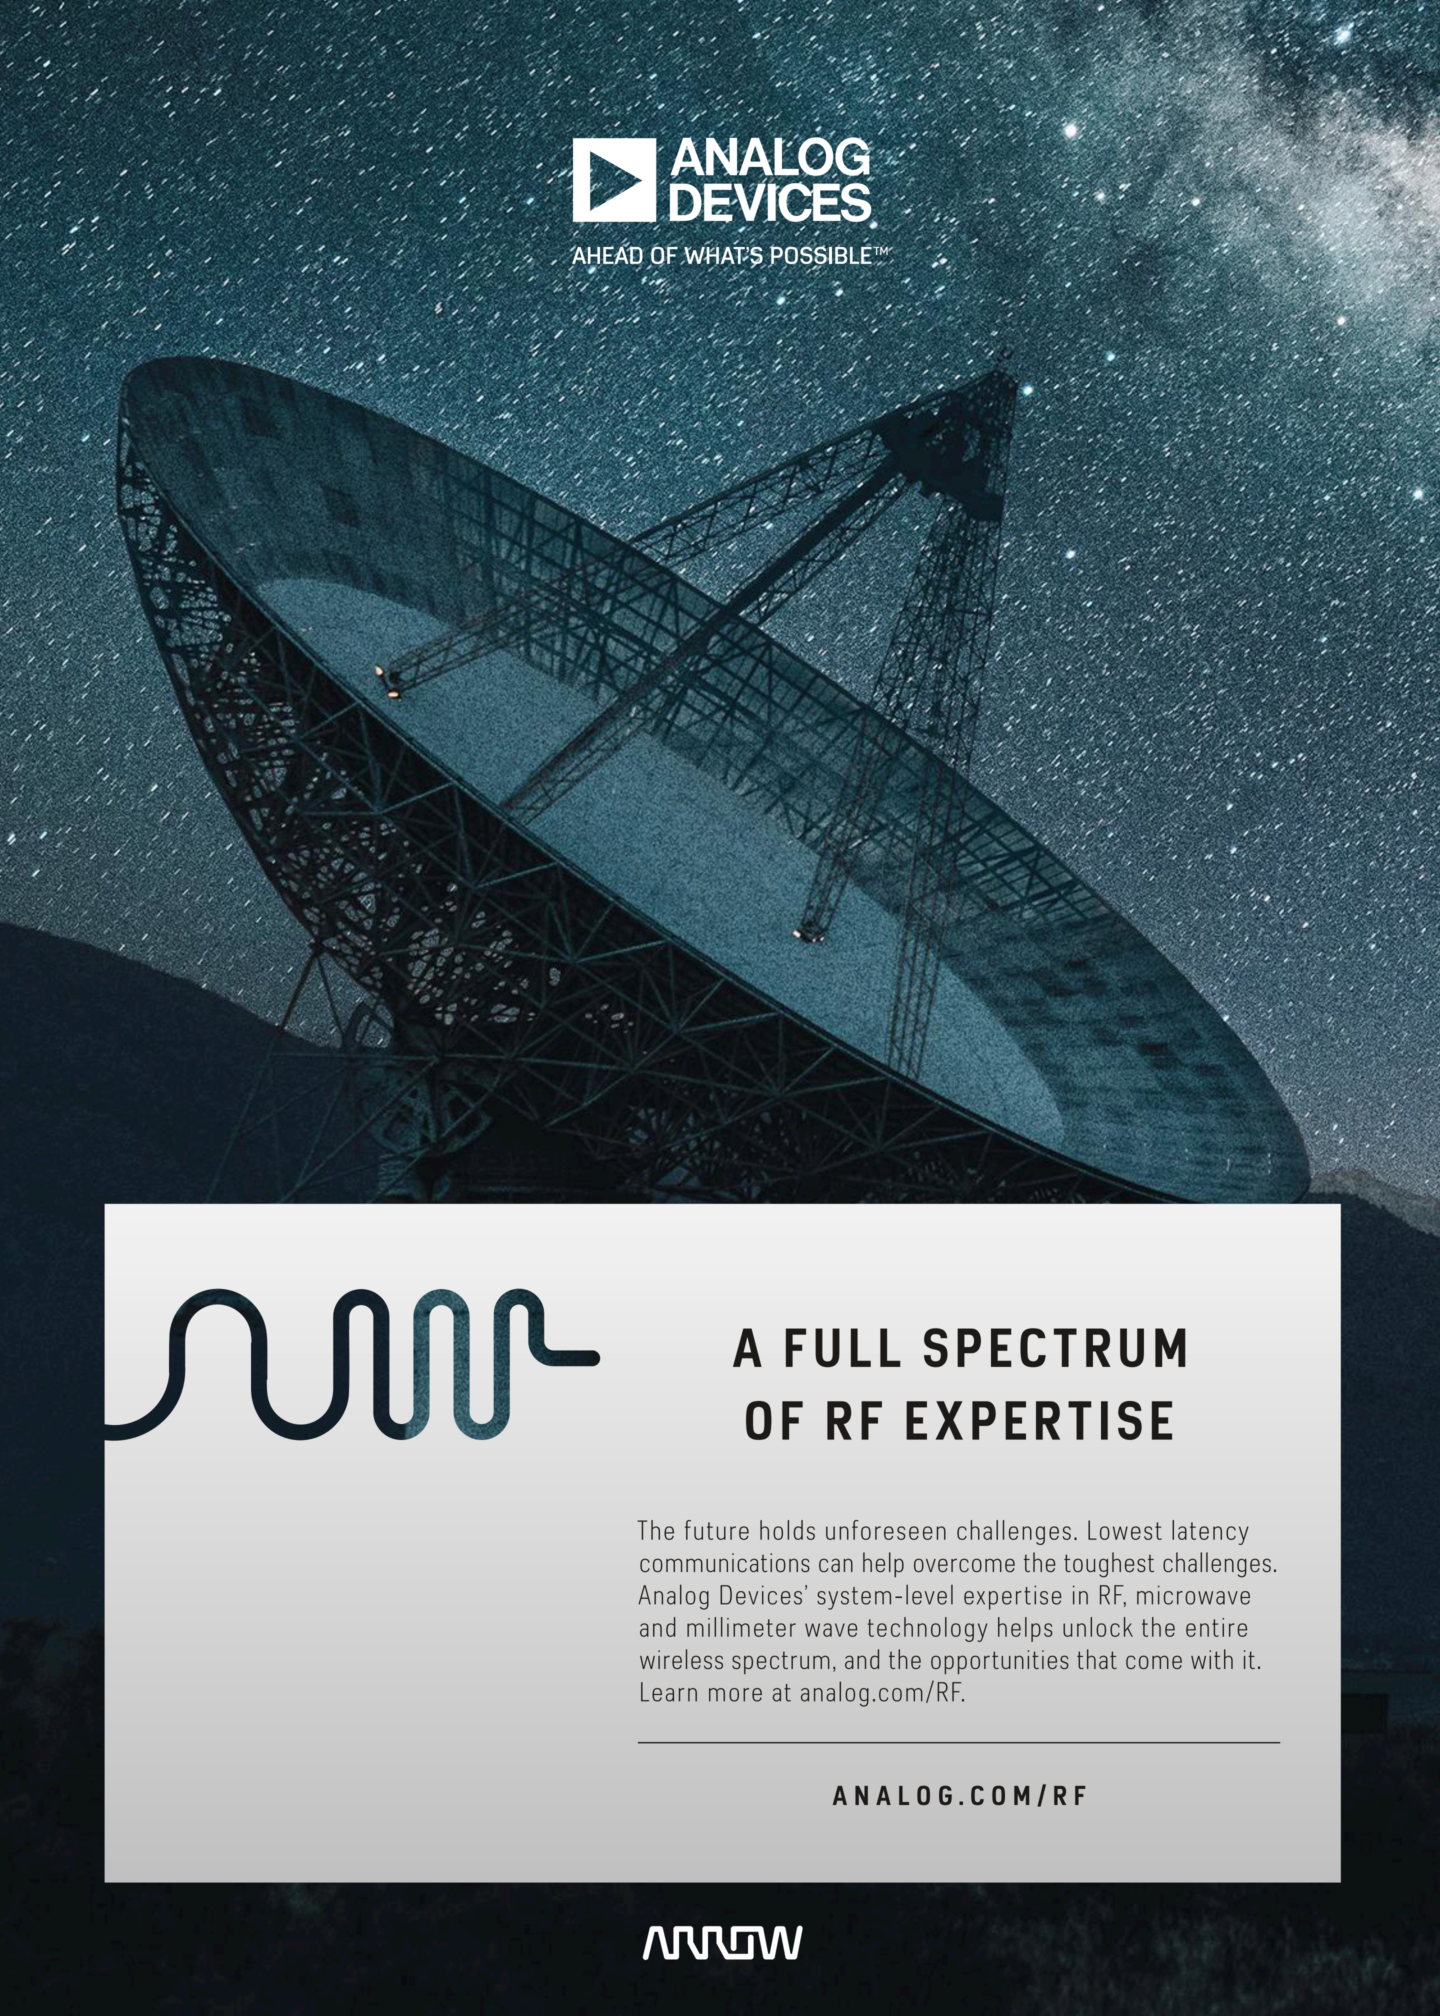
ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™
A FULL SPECTRUM OF RF EXPERTISE
The future holds unforeseen challenges. Lowest latency communications can help overcome the toughest challenges. Analog Devices' system-level expertise in RF, microwave and millimeter wave technology helps unlock the entire wireless spectrum, and the opportunities that come with it. Learn more at analog.com/RF.
ANALOG.COM/RF
ARROW

WILLIAM WONG | Senior Technology Editor

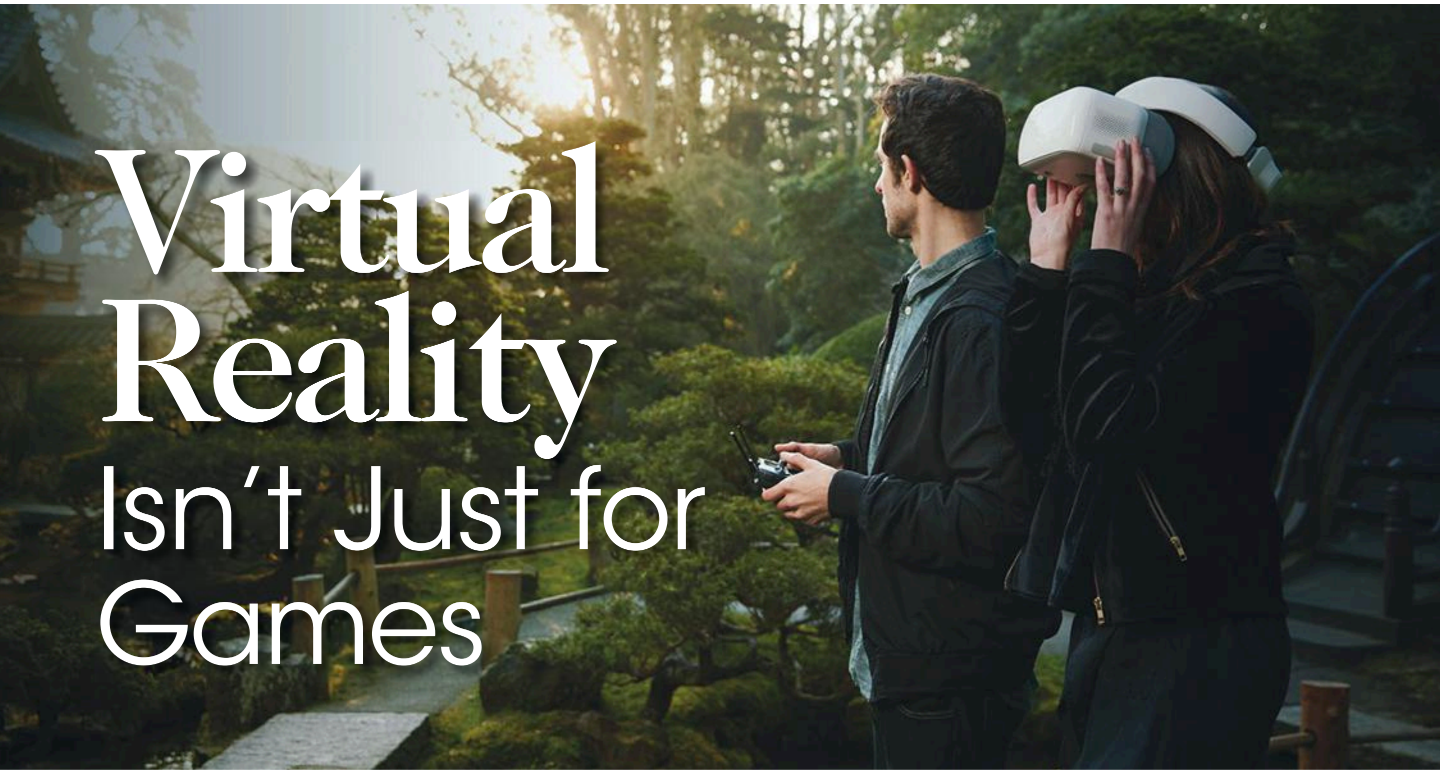

Virtual Reality Isn't Just for Games

Virtual reality is finding its way into a range of applications, from controlling drones to training workers about safety zones.

Virtual reality (VR) has made a dent in gaming, but its real value is in commercial and industrial environments. Improved graphics, lighter-weight goggles, and better software along with lower costs will make VR solutions ubiquitous across many arenas.

One area where VR is coming into its own is in training. VR allows training to be more interactive as well as putting users in situations that might otherwise be unsafe or hard to replicate in the real world. Oftentimes, training has consisted of books and Powerpoint presentations with minimal on-the-job training.

Royal Innovative Solutions (RIS) is one company that's working to bring VR to the highways. Their VR training solutions targeting highway safety and construction help reduce lost time from injuries and minimize litigation costs related to working in hazardous environments *(Fig. 1)*. The training environment uses a combination of VR and 3D animation. The VR approach boosts learning and retention rates.

The ability to rehearse operations in a realistic environment can lead to increased operational efficiency and production as well. Furthermore, RIS is

1. Royal Innovation Solutions is leveraging VR for training purposes.

2. The DJI Goggles can be used for flight control and/or camera control. Multiple users are able to view the camera video at the same time.

bringing training to many states where highway safety training was minimal or non-existent due to the costs and the ineffectiveness of book learning. VR also allows training to be done almost anywhere.

Another area where VR is having a major impact is in drone technology. Drone racing, like the International Drone Racing Association competitions, has participants piloting drones at very high speeds through flight mazes.

A more practical tool is the DJI Goggles from DJI for their complement of consumer and commercial drones *(Fig. 2)*. The VR headset is integrated with the flight control and camera system. It provides 1080p resolution to both eyes with an 85-deg. field of view. The drones can transmit at that resolution when the headsets are nearby and at 720p over longer distances. Up to four users are able to share the experience from a single drone.

A touchpad on the right side of the goggles can be used for command and control. It's also possible to use head tracking for control of the camera or the flight. Flight control and camera control are split; therefore, a pilot can handle flight control while other users view or control the camera. This can make tracking an object easier since a pilot doesn't have to worry about the drone movement as well as the camera movement. Head tracking makes camera movement very simple and natural.

The DJI drones have a number of flight control modes that blend well with VR. The simulated fixed wing mode allows a drone to act like a plane that flies at a constant speed. The VR headset movement is used to adjust yaw and pitch. There's also terrain following and a tripod mode that allows users to concentrate on camera movement.

The DJI Goggles are designed for drones compared to the conventional VR glasses. Antennas are built into the headband and it's balanced for head movement. The headset can flip up or be easily removed to provide viewing of the outside world. The touchpad is a useful addition and can be easily disabled. The system also has an HDMI input for external video feeds.

VR is being used in other applications, from car showrooms to present the latest model cars and trucks to virtual vacations for exploring venues without leaving home. The 1080p resolution is currently the norm, with 4K and higher pushing the envelope in cost and performance. These are also pushing the GPUs driving the systems because of the higher bandwidth. ed

MARK BEECHAM | MCU and Sensor Product Manager, Silicon Labs
www.silabs.com

Rightsizing Your Sensor Network Design for the IoT

When developing your next wireless-sensor-network application, it pays to rightsize the network with the optimal number of sensors and consider variables such as the network medium, mains, or battery power sources, wireless protocol options, and network configuration.

With the rise of the Internet of Things (IoT), we've seen a growing trend toward collecting and aggregating sensor data for a wide array of smart-home, industrial, green-energy, transportation, and smart-city applications. The general belief across the industry was the more data sent to the cloud or to a local control system, the better.

Much of this sensor data was in the form of physical sensing for human occupation, object detection, temperature, humidity, light, sound, and vibration. Over time, developers realized the difficulty of deploying large sensor networks, and many cloud companies concluded that access to vast amounts of data doesn't necessarily add value if you can't act on that data for decision-making.

In today's cost-sensitive environment, every deployed sensing device or network must balance the complexity of development and deployment with the value delivered to IoT service providers, cloud companies, and end users. We're now seeing a trend in deployed sensor networks of simplifying the network and measuring only what can be translated into cost savings and/or a more positive end-user experience.

SENSOR NETWORK TRENDS AND USE CASES

Let's consider two use cases that exhibit the emerging trend of deploying the optimal number of sensors required by the IoT application: a retail store and a commercial office building.

A store's most valuable real estate is its merchandise shelves, and a key metric is customer conversion. Most large retailers would benefit from a sensor network measuring occupancy, customer and item location, inventory level, and more. Ideally the system would use sensor data to increase the likelihood of customer purchases. The sensors could use new Bluetooth tracking technology such as angle of arrival (AoA) and angle of departure (AoD) as well as received signal strength indicator (RSSI) or phase shift for distance.

Such AoX *(Fig. 1)* and distance technologies enable the network to triangulate the position of a customer's smartphone or a device embedded in the shopping basket/cart to detect the customer's location and where shoppers spend their time in the store. This data combined with facial-recognition technology or smartphone identification could assign a portfolio to each shopper anticipating items they're most likely to buy and either offer in-store sales or bundles to incentivize them or deliver

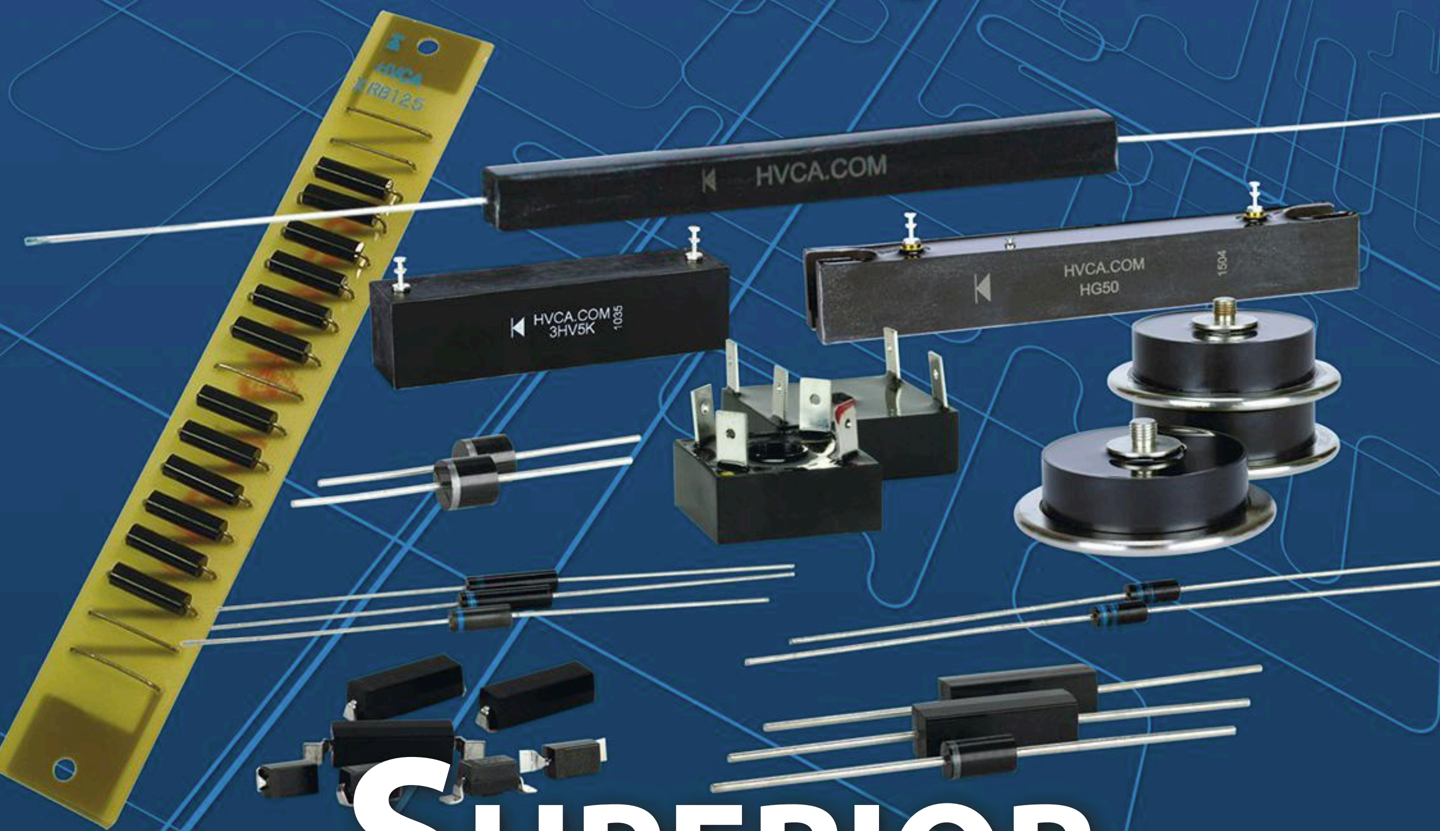
HVCA.COM
HVCA.COM
3HV5K
1035
HVCA.COM
HG50
1504

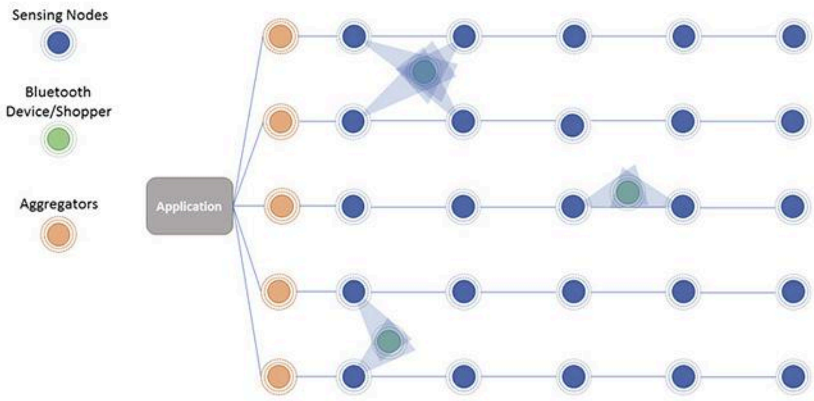

1. Angle of Arrival and Angle of Departure (AoX) technology, supported by Bluetooth 5.1, enhances the development of indoor positioning systems.

targeted marketing campaigns after in-store visits.

The optimized sensor network could also provide store managers with insights into buying decisions for various product types. For example, the sensor network might reveal that certain customers spend a lot of time comparing toothpaste brands while acting quickly with shampoos, choosing a brand they know. This information enables the store to intelligently stock their shelves with items customers are more likely to purchase, increasing the store's efficiency and profitability.

Converting an office building into a sensor-laden smart building will reduce maintenance and operating costs and enhance safety, comfort, and convenience for its occupants. Climate control is a major cost factor for facilities management. It can cost more than $50,000 a month to adequately cool or heat a typical 100,000-square-foot building. Cutting this cost by 15% by adding a $300,000 smart sensor network would be a wise investment, paying for itself in a few years.

Another valuable addition would be a sensor network optimized for damage avoidance. The ability to sense threats such as water leaks, flooding, smoke and fire and then react quickly can save significant repair cost while protecting occupants and even saving lives. Threat reduction through sensor networks could also reduce insurance premiums.

While these two sensor network examples can directly enhance the deployer's top or bottom line revenue, there are two key challenges to consider when implementing a sensor network: the physical medium for deployed sensors and how the sensor data is handled.

THE PHYSICAL MEDIUM

Every sensor network is configured differently, depending on the application requirements and physical environment. For example, some networks may be deployed in retail environments with uniform coverage while others may be implemented in corporate offices with varying room sizes. Some networks may be used 24/7 while others are active partially on a given day. A factory plagued with electromagnetic noise from motors can be a challenging environment for sensor networks, making some types of sensing or communication techniques impossible. Each of these scenarios requires an optimal sensing solution.

Consider the example of a commercial building. A common goal is to reduce energy consumption and cost for lighting and HVAC systems. There are two options: minimize lighting and HVAC usage in areas with little to no occupancy or implement power-efficient modes of lighting and climate control. With either option, deploy sensor types that are absolutely necessary, such as occupancy, light level, and temperature sensors. Additional sensors for humidity or pressure can add more confidence by ensuring that occupants aren't negatively impacted by optimizing lighting and climate controls.

To deploy an optimal sensor network, the developer must make a few decisions upfront:

"Will the sensing nodes be battery powered or wired?"

A battery-powered wireless sensor node is easier to install as there are no power lines to run. But battery replacement every few years can negatively impact the ROI equation. Running power lines (mains or power over Ethernet) can also be an expensive upfront cost, and local code compliances vary by region. Deploying a low-voltage system like PoE can be more cost-effective than running 110 V or 220 V to every sensor node. Another economical option is to install low-voltage LED lighting along with the sensor network for easier access to energy savings.

"Will the sensing node use wired or wireless communication?"

The solution depends on the tradeoffs of using wired power and wired data communications versus battery power and wireless data. Wireless can be cost-effective since there's no need to run data lines throughout a building, but other upfront engineering costs and complexities must be considered when implementing a wireless network. For new construction, it's very easy to run CAT5, which can pass data and power to every node.

With CAT5 wire, the data rate can be lower than the typical 1 Gb/s, meaning the CAT5 cable can run more than 100 meters. However, for some building renovations, accessing hard-to-reach places to run cable such as dropped ceilings can be nearly impossible. In such cases, wireless connectivity may be the best option.

When implementing a wireless network, developers must consider which medium and protocol to use. Wi-Fi is attractive from its ease-of-use due to the pervasive infrastructure and familiarity of IP packets, but Wi-Fi can only support a limited number of end nodes and access points and consume lots of power compared to other short-range options. While Zigbee solves the power problem, developing a Zigbee network may be challenging for developers who lack mesh expertise. The deployer will need to gauge the strengths and drawbacks of each protocol choice *(Fig. 2)* to determine the best path. There's no perfect solution. The best wireless choice is one where tradeoffs can be tolerated to achieve optimal application goals.

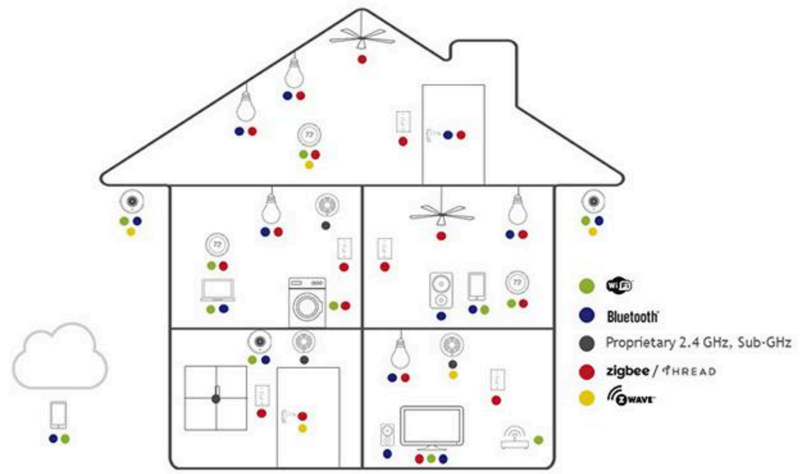

2. Developers have a wide choice of short-range protocols for wireless-sensing applications used in building automation.

"Where will decisions based on sensor data be made?"

Depending on the requirements of a commercial building or home, the decision to turn an HVAC system on or off can be made locally within the building by an HVAC controller. Alternatively, the sensor data can be sent to the cloud

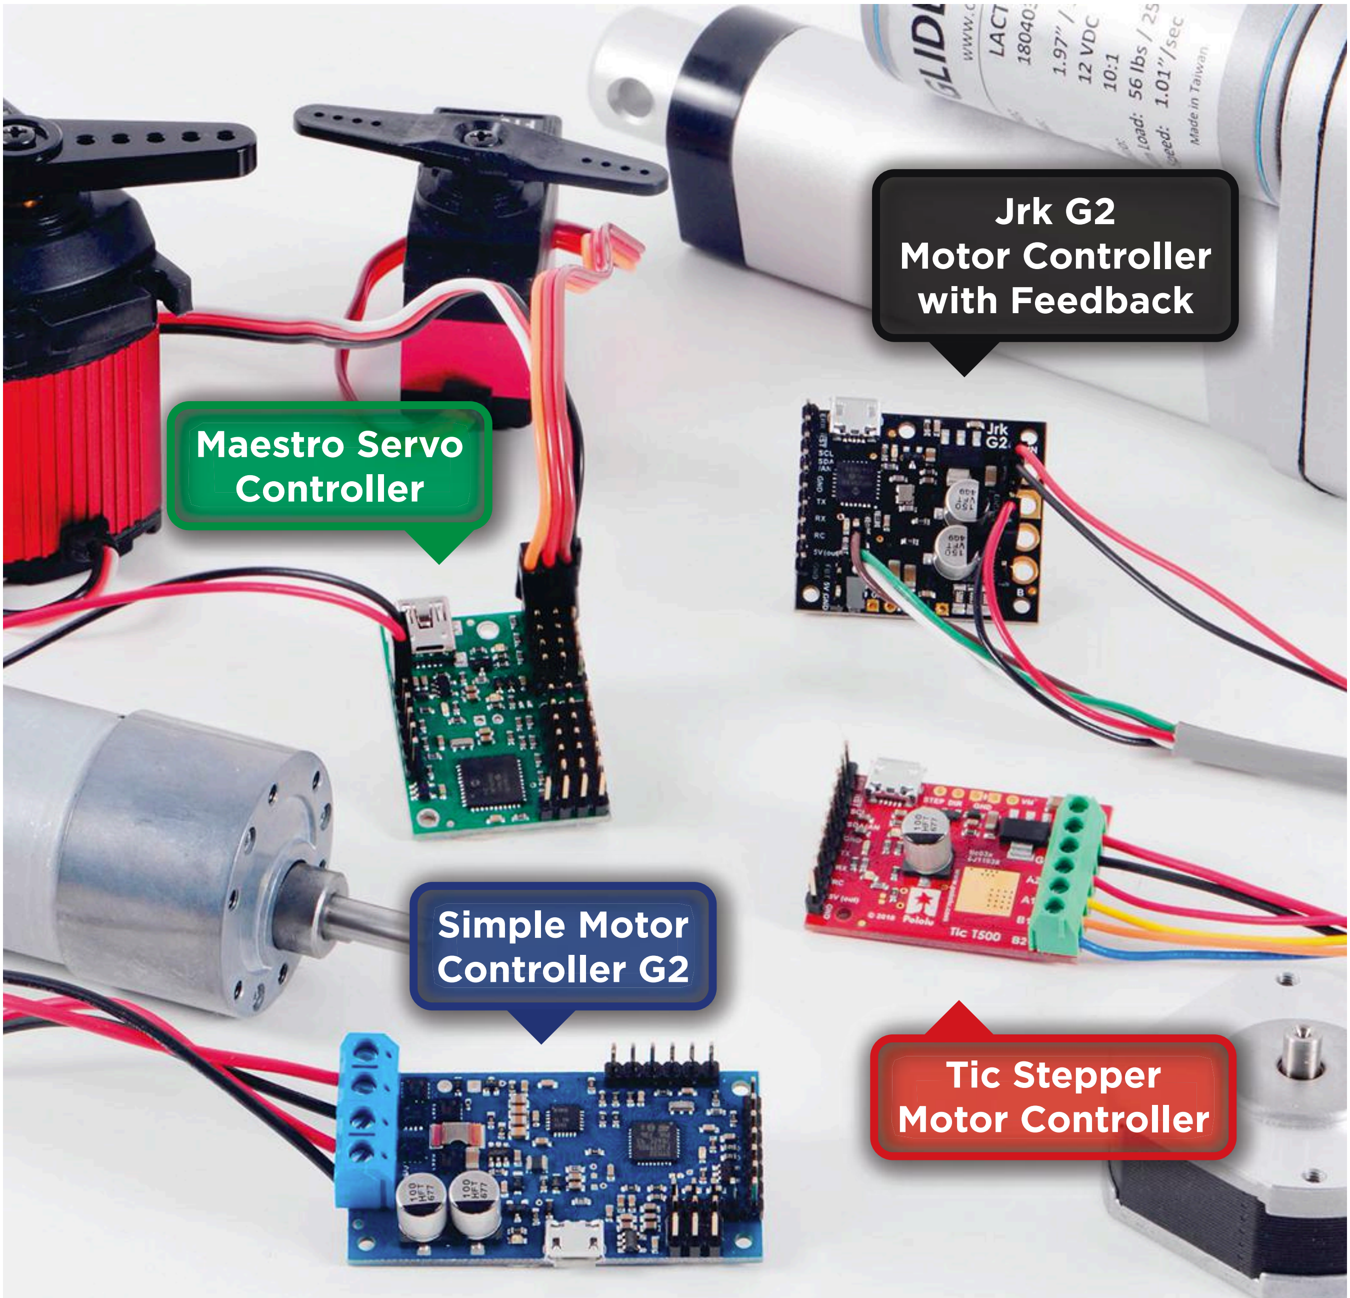

CONTROL INTERFACES

	USB	TTL serial	I²C	Analog voltage	RC	Rotary encoder	Scripting
Jrk G2 Motor Controllers with Feedback	✓	✓	✓	✓	✓		
Maestro Servo Controllers	✓	✓		✓*			✓
Tic Stepper Motor Controllers	✓	✓	✓	✓	✓	✓	
Simple Motor Controllers G2	✓	✓	✓	✓	✓		

* through scripting

Pololu
Robotics & Electronics

Get your projects moving with motor controllers from Pololu. Our controllers for RC servos, stepper motors, and brushed DC motors all include free software, so you can easily configure your motor's movement based on your selected control input.

Find out more at **www.pololu.com**

for processing, enabling remote control of the HVAC system.

This choice of local or remote data processing and control depends on many factors, including the complexity of the computations required for decision-making, how much data must be processed, and if post data processing or remote monitoring are needed. Local processing and control may be feasible for a smart home with limited data requirements. However, managing an HVAC system for a building with many rooms may require a gateway to transmit data to the cloud.

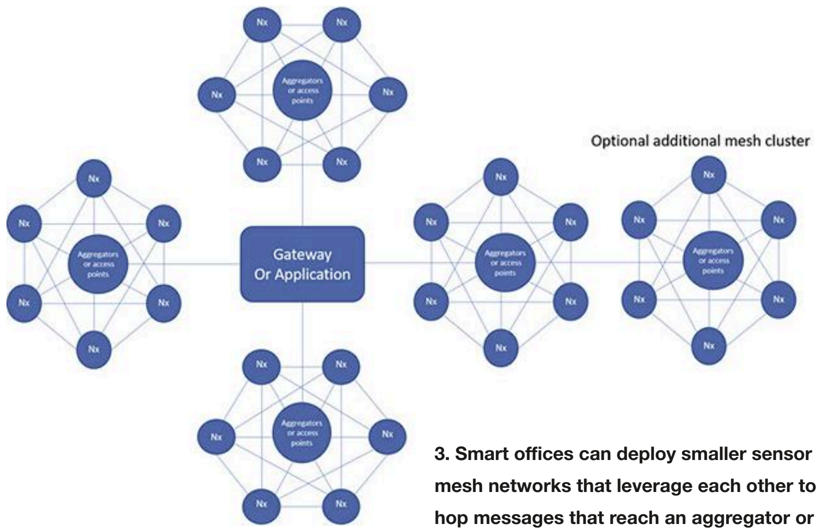

3. Smart offices can deploy smaller sensor mesh networks that leverage each other to hop messages that reach an aggregator or access point.

HANDLING SENSOR DATA

Handling the data generated by hundreds of sensor nodes is a showstopper for many developers attempting to deploy large-scale sensor networks. Setting up a few nodes in a lab and getting them to communicate to a host or the cloud is relatively simple, but the network design becomes exponentially difficult as the number of nodes increases.

The application that collects local or cloud-based data must identify and time-stamp all data to build an accurate model of the building. For a smart office, a host system will receive data from hundreds of sensors. This can be an unrealistic, expensive, and difficult problem. While a sensor on each end node may cost only $0.50, the requirements for a mesh network and cloud computation overhead may far outweigh the cost of the sensors. Experienced sensor network deployers are now adding only what's absolutely required to each end node. This rightsizing approach cuts data overhead and reduces the power consumption of end nodes, which can save battery-replacement cost.

Network topology is a multivariable problem. A trending solution is to use a hybrid topology that combines a mesh network with a bus or tree topology. For example, in a smart office, it makes sense to deploy many smaller sensor mesh networks that can leverage each other to hop messages that reach an aggregator or an access point *(Fig. 3)*. This access point can be part of another mesh network that's one step away from the host or includes the host and other aggregators. Using these "middle" nodes simplifies a difficult problem by essentially splitting the work between the aggregators and the host.

By using a hybrid topology, the cloud can receive organized data, easing the burden on the cloud application and enabling a more scalable network. The cloud application will most likely be customized for the specific task and must be created or licensed to provide monitoring capabilities for the building.

The current trend is to implement such an application through "software as a service" (SaaS). For example, an insurance company wishing to cut HVAC cost usually doesn't have the core competence to develop this application on its own. The application provides feedback and configurations that facility managers can use to custom-control their building environments to their liking.

CONCLUSION

Sensor networks can be a challenge to deploy. No one sensor network layout or measurement portfolio fits every IoT application. One network may focus on tracking assets or people, while another prioritizes sensing environmental changes to cut cost or avoid disasters. Whichever sensor network approach is chosen, it must bring value. This value must offset the cost and complexity of the physical medium in which it's deployed as well as the system that's deployed to handle the data and ultimately make decisions.

Deploying a sensor network must be carefully planned from battery-powered end nodes to large-scale cloud-based applications. When deploying or designing a sensor network, consider all aspects—even the slightest oversight can cause insurmountable tradeoffs. Consider how each sensor node will be powered, how data will be transferred and its path to the host application, and how the network will be serviced or scaled up over time. Sensor networks are beginning to achieve optimal value-to-cost ratios, and more businesses will deploy them to reduce cost, gain intelligence, and differentiate from the competition. ed

MARK BEECHAM is a product marketing engineer supporting and defining Silicon Labs' 8- and 32-bit microcontroller and sensor products. Specializing in embedded systems and sensors, Mark joined Silicon Labs in 2015; previously he worked for IBM and Texas Instruments. He holds a BSEE from Texas A&M University.

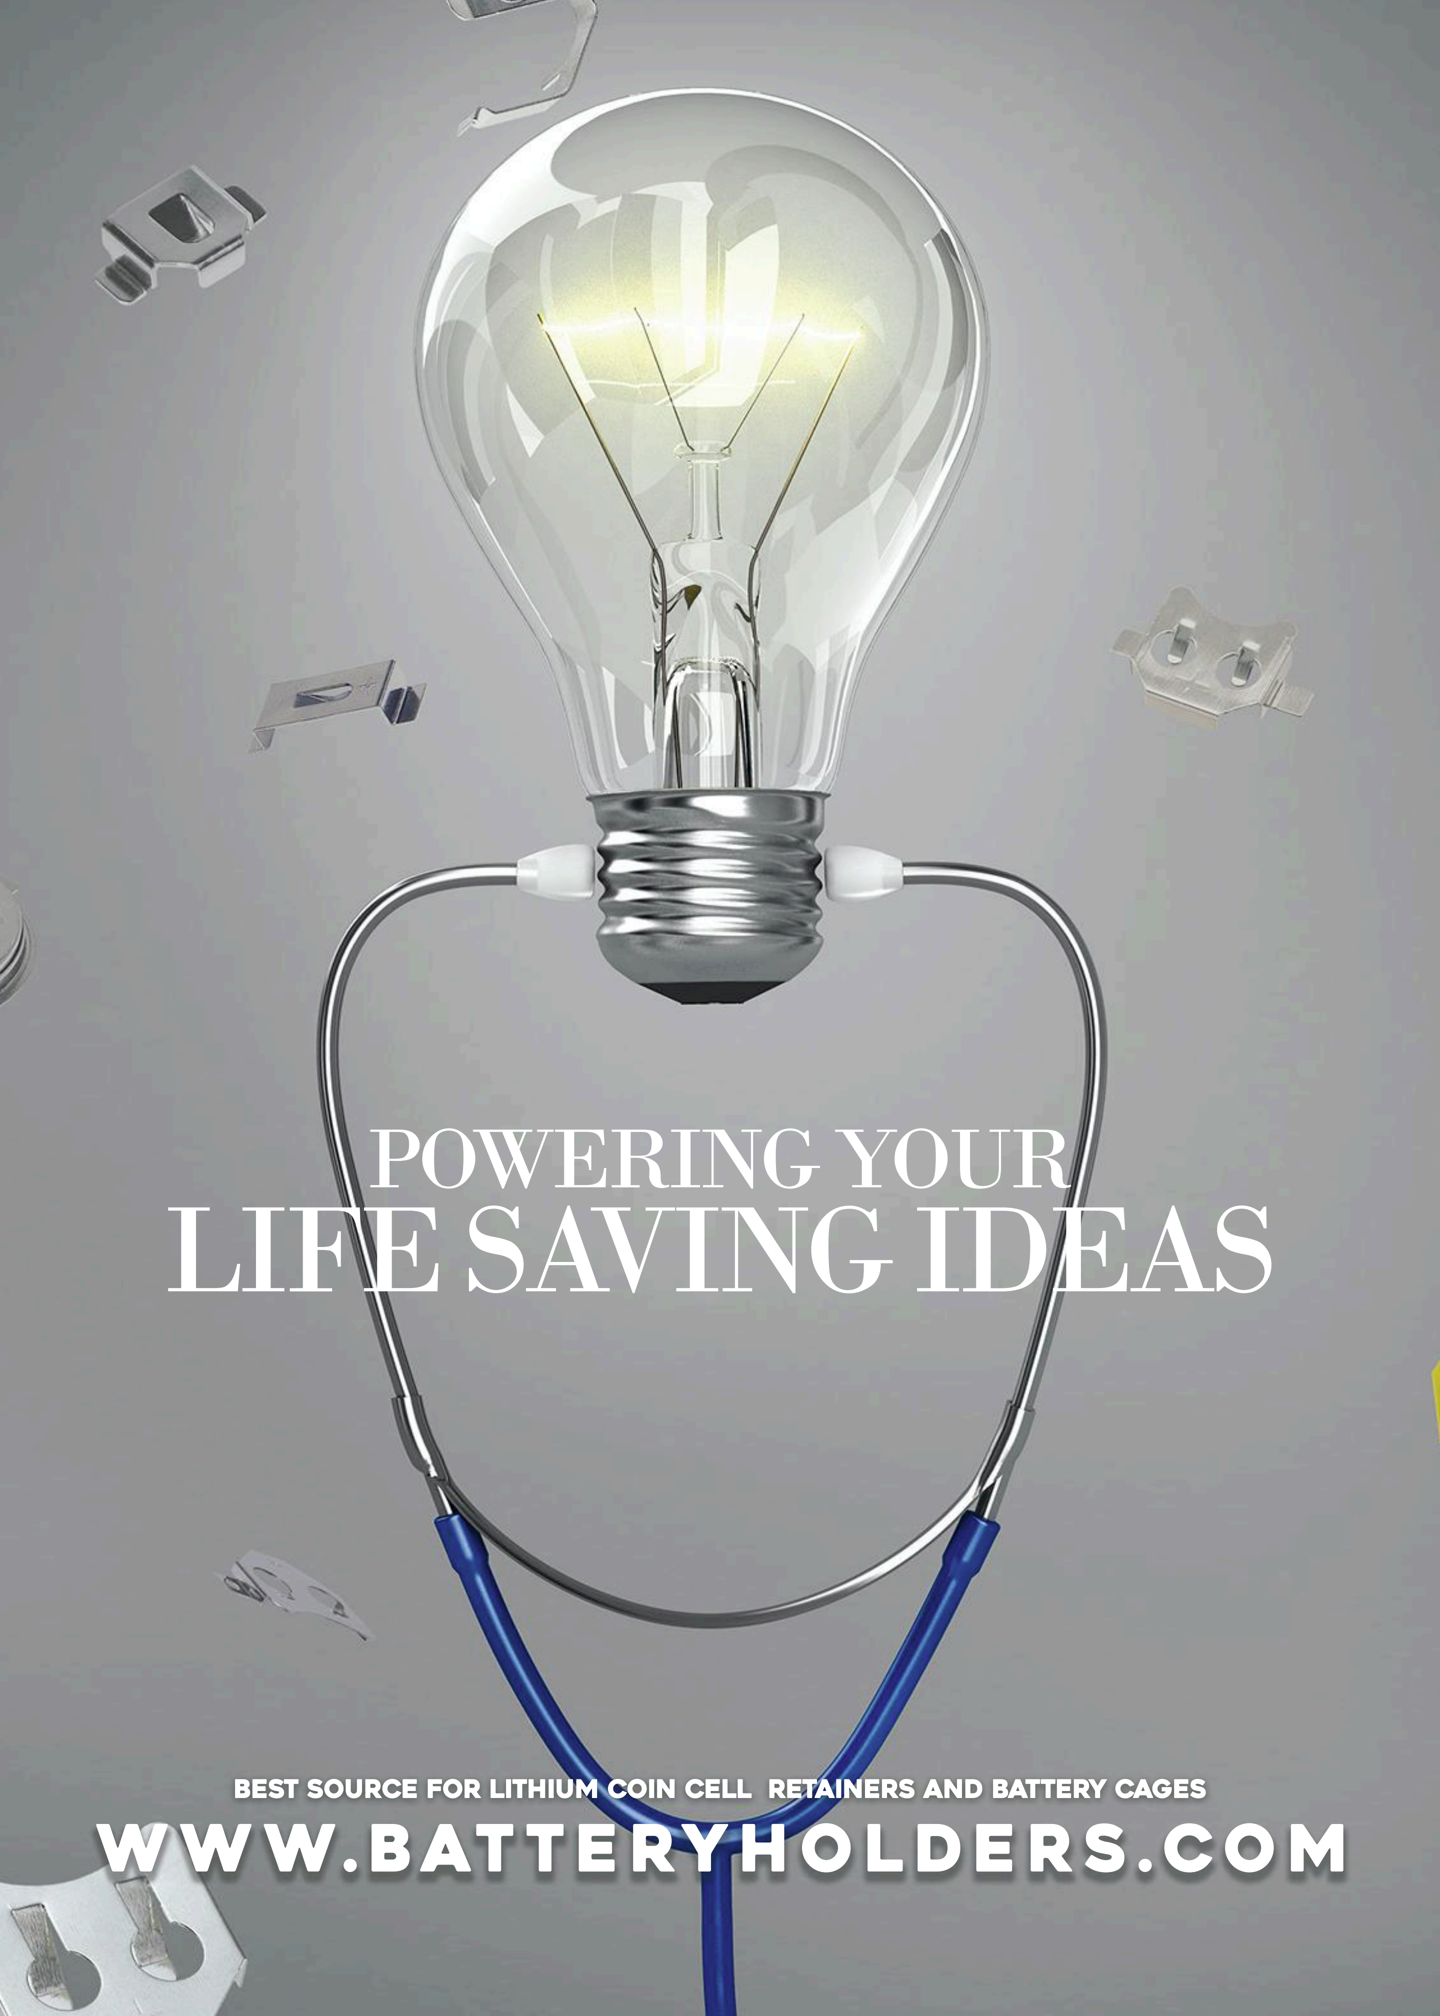
POWERING YOUR
LIFE SAVING IDEAS
BEST SOURCE FOR LITHIUM COIN CELL RETAINERS AND BATTERY CAGES
WWW.BATTERYHOLDERS.COM

HARSHA JAGADISH | Product Marketing Manager, MCU16, Microchip Technology
www.microchip.com

Design Separately, Integrate Seamlessly with Dual-Core Controllers

As the level of software integration becomes greater in embedded applications, designers are looking for ways to simplify the process. A dual-core digital signal controller offers one solution.

Embedded applications are becoming complex and sophisticated to meet several objectives. First, applications need to improve efficiency, which requires significant controller performance to run sophisticated algorithms. Next, the ubiquitous internet availability is enabling embedded applications to become "smarter" and more "connected." The third objective is to reduce cost by integrating several functions like sensor interfacing, connectivity, motor control, digital power conversion, security, and safety in a single controller.

Such a high level of integration requires respective domain experts to handle specific functional areas or modules; then, multiple functions need to be integrated into an end application. Often, multinational companies having their teams spread across the globe makes it even more important that various modules be able to be designed separately and integrated seamlessly with ease to reduce the development risk and efforts.

1. **Today's applications are increasingly "smart" and "connected."** *(Courtesy of Microchip)*

EFFICIENCY IMPROVEMENT

First, let's look at how the objective of improving energy efficiency requires increased controller performance. Consider the example of a motor-control application. The industry has been moving away from brushed dc motors, which offer 75-80% efficiency, and toward brushless dc (BLDC) motors or the newer permanent-magnet synchronous motors (PMSMs). These motors offer up to 85-90% better efficiency, reduced acoustical noise, and increased product lifespan.

A typical brushed dc motor control requires very simple direction and speed-control techniques that can be accomplished using an entry-level 8-bit microcontroller. In comparison, controlling a sensorless BLDC or PMSM motor with "field-oriented control" (FOC) is more sophisticated and computationally intensive. It allows for close control of energy used by the motor over a wide range of load or speed, and helps significantly improve the efficiency. Additional control algorithms may also be implemented based on application requirements, like "Rotor Stall Detection and Recovery," "Windmilling," "PI Loop Saturation and Antiwindup," "Flux Weakening," and "Maximum Torque per Ampere," which help improve the performance and response to a dynamic load, plus increase the overall efficiency.

All of these advanced control techniques are computationally intensive, involving math operations such as divide, multiply, square-root, and trigonometric operations, requiring significant central processing unit (CPU) bandwidth. Because these control functions need to be executed periodically at a high frequency, it's necessary that the CPU gets allocated at a specific time interval.

Such tight control-loop execution can take up most of the CPU bandwidth and impact other time-critical functions in a complex application. An embedded developer will have limited flexibility to add any additional functionalities like communication, safety monitoring, and system and housekeeping functions that could interfere with the time-critical control of the motor. The challenge increases in digital power applications, where the time-critical control-loop functions need to be executed at an even higher frequency.

COMPLEX SOFTWARE

Now, let's consider the next objective driven by internet or cloud connectivity. The latest industry trend is for applications to be "smart" and "connected," offering intelligence and accessibility from anywhere. These requirements demand embedded applications to include multiple software stacks such as:

- The main application function software. In our example, this function implements motor control, housekeeping, and user-interface operations that are commonly required in most of the applications.
- The communication software running the necessary network application protocols for connectivity.
- The security software for IP protection, privacy, data integrity, authenticity, and access control, and for thwarting any hacking possibilities.
- If applications involve human operations and may cause bodily injuries due to malfunctioning, then even the functional safety software needs to be part of such safety-critical applications.
- Some of the end applications may also have customization requirements where certain features will be unique to specific variants targeted for different market segments.

All of these function requirements necessitate various domain expert teams to be involved in developing respective software stacks and be able to optimally and quickly integrate them into an end application. Experts from multiple domains will need to coordinate very closely to design and implement an end application. This scenario gets further complicated in multinational companies where the expert teams will be spread across the globe.

COST REDUCTION

Finally, cost optimization is an important objective that's common to all end applications. Often, embedded engineers will not have the budget to consider a multi-microcontroller design, where an individual software stack can be executed on different microcontrollers with very little coordination. Going for a single microcontroller design with very high integration will be the most optimal solution. This further enables cost reduction due to a compact PCB design and a reduced number of external components like crystal oscillators and passive components.

WHAT ARE THE DEVELOPMENT CHALLENGES?

To implement sophisticated algorithms and execute multiple software stacks, embedded designers often choose a higher-performance microcontroller. However, this may not be the best choice due to the challenges associated with time-critical execution, multiple software stack development, integration, and testing.

A simple scheduler or a real-time operating system (RTOS) may serve the purpose of scheduling and executing multiple tasks from different stacks on a high-performance CPU in a time-spliced manner. But, a scheduler or an RTOS adds overhead that consumes CPU bandwidth, memory, and other microcontroller resources. The time-splicing also increases switching overhead, reducing the effective CPU utilization.

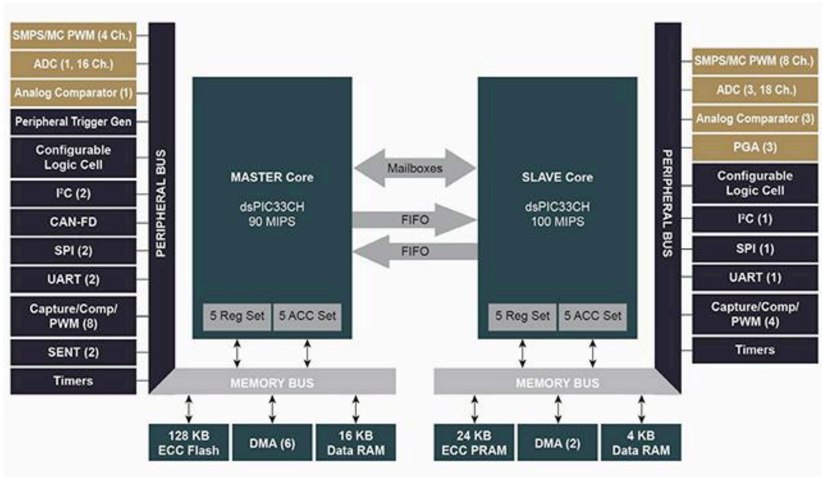

2. A typical dual independent core controller block diagram. *(Courtesy of Microchip)*

The scenario gets further complicated when two time-critical complex control loops need to be executed periodically at a precise and overlapping time interval or when two asynchronous safety-critical functions must be executed simultaneously in real time. In such cases, considering an even higher-performance microcontroller will not always serve the system requirements.

Even if a high-performance single core microcontroller has enough CPU bandwidth to accommodate multiple software stacks, maybe together with an RTOS, there are many other design complications to consider. Developing, integrating, and testing multiple software stacks needs a considerable amount of coordination among subject matter experts. It requires designing a compatible and modular software architecture that dynamically shares resources and exchanges information. The complications further increase if any of the legacy stacks doesn't have compatible architecture:

- Legacy stacks may have different architectures based on either polling mode or interrupt mode.
- Legacy stacks may be using the same microcontroller's resources, which now need to be shared without any conflicts to avoid hazards like race condition and deadlock.
- Stacks may have several common global variables and functions with the same names.
- Each stack may function perfectly when executed individually but may malfunction on integration. Debugging such an integrated solution will be a nightmare that increases the development time.

An already available standalone stack may not always help in reducing the development time when implemented on a single-core microcontroller. All of these challenges pose significant development risk and increase the time to market.

DUAL-CORE CONTROLLER

A dual-core controller helps to improve efficiency, simplify development efforts, and reduce cost with the following offerings:

- It offers higher performance than a similar single-core controller operating at twice the speed and is ideal for applications with two or more time-critical functions.
- It simplifies software development with dual independent cores that enable geographically dispersed software development; seamless integration with very minimal coordination; and easy feature customization across multiple variants of a product line.

ADI's High Power Silicon Switches Save Bias Power and External Components in Massive MIMO RF Front-End Design

Bilge Bayrakci
Analog Devices, Inc.

Multiple input, multiple output (MIMO) transceiver architectures are widely used in the design of high power RF wireless communications systems. As a step into the 5G era, massive MIMO systems covering cellular bands are now being deployed in urban areas to meet the emerging demand for high data throughput and a new variety of services. The availability of highly integrated single-chip baseband transceiver solutions, such as ADI's new ADRV9008/ADRV9009 family of products, made this achievement possible. Similar integration toward lower power consumption for thermal management and smaller size for cost is still needed at the RF front-end section of these systems in order to accommodate more MIMO channels.

MIMO architectures allow the RF power requirements on building blocks, such as amplifiers and switches, to be relaxed. However, as the number of parallel transceiver channels increase, the complexity of the peripheral circuits and power consumption scales accordingly. ADI's new high power switches in silicon technology are designed to simplify the RF front-end designs by eliminating the need for peripheral circuits and reducing power consumption to negligible levels. ADI's new high power switches in silicon offer RF designers and system architects the flexibility to increase their system complexity without RF front ends becoming a bottleneck in their designs.

In time division duplex (TDD) systems, there is a switch function incorporated at the antenna interface to isolate and protect the receiver input from transmitted signal power. This switch function can either be used directly at the antenna interface in relatively lower power systems, as shown in Figure 1, or used on the receive path, as shown in Figure 2, for higher power applications to ensure proper termination to the duplexer. Having a shunt arm on the switch outputs will help to improve the isolation.

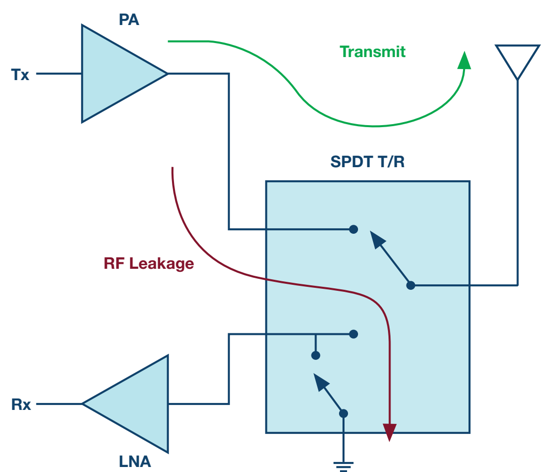

Figure 1. Antenna switch.

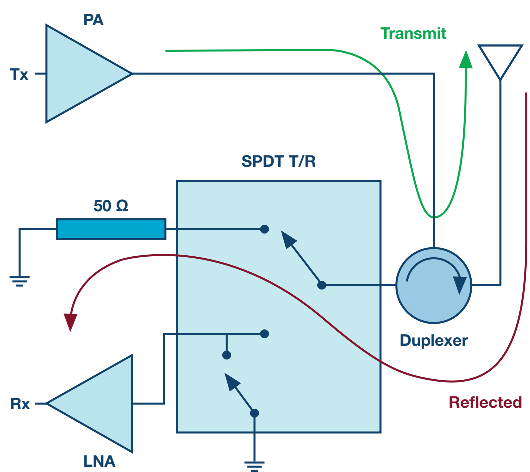

Figure 2. LNA protection switch.

Table 1. ADI's New High Power Silicon Switch Family

Part Number	Frequency	Insertion Loss	Isolation	P_{AVRG}	P_{PEAK}	Package
ADRF5130	0.7 GHz to 3.5 GHz	0.6 dB, 2.7 GHz 0.7 dB, 3.8 GHz	45 dB, 3.8 GHz	20 W	44 W	4 mm × 4 mm
ADRF5132	0.7 GHz to 5.0 GHz	0.60 dB, 2.7 GHz 0.65 dB, 3.8 GHz 0.90 dB, 5.0 GHz	45 dB, 3.8 GHz 45 dB, 5.0 GHz	3.2 W	20 W, 3.8 GHz 10 W, 5.0 GHz	3 mm × 3 mm
ADRF5160	0.7 GHz to 4.0 GHz	0.8 dB, 2.7 GHz 0.9 dB, 3.8 GHz	48 dB, 3.8 GHz	40 W	88 W	5 mm × 5 mm

PIN diode-based switches have been the preferred solutions due to their low insertion loss characteristics and high power handling capabilities. However, their need for high bias voltages to reverse bias for isolation and their need for high current to forward bias for low insertion loss are the shortcomings in the design of massive MIMO systems. In Figure 3, a typical application circuit is shown for a PIN diode-based switch and its peripherals. Three discrete PIN diodes are biased through their bias-tee circuits and are controlled through a high voltage interface circuit.

ADI's new high power silicon switches are better suited to massive MIMO designs. They run on single 5 V supply with less than 1 mA bias current and do not need external components or interface circuits. In Figure 4, internal circuit architecture is shown. A FET-based circuit operates on low bias currents and low supply voltages, bringing down the power consumption to negligible levels and helping the thermal management at the system level. Besides the ease of use, the device architecture yields better isolation as more shunt arms are incorporated on the RF signal paths.

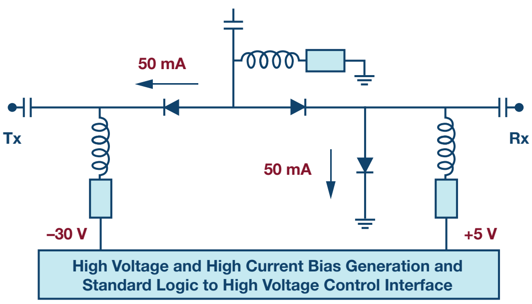

Figure 3. PIN diode switch.

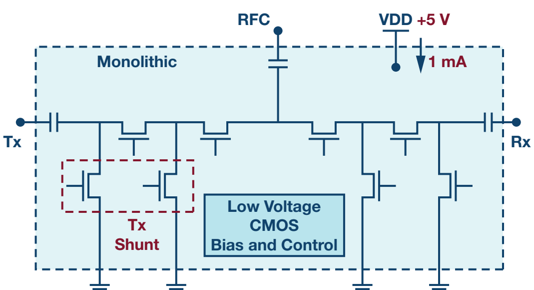

Figure 4. ADRV9008/ADRV9009 Silicon switch.

Figure 5 shows a side-by-side comparison of printed circuit board (PCB) artwork between PIN diode-based switch and new silicon switch on a single-layer PCB design. The silicon switch can fit into more than 10× smaller PCB area. It has simplified power supply requirements and does not need high power resistors.

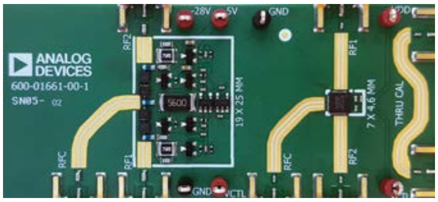

Figure 5. Side-by-side comparison of a PIN diode-based switch design to a silicon switch.

ADI's high power silicon switches can handle an RF peak power of up to 80 W—that is enough to cover, with margin, the peak-to-average power ratio requirements for massive MIMO systems. Table 1 shows ADI's family of high power silicon switches that are optimized for different power levels and various package options. These devices inherit the intrinsic advantages of silicon technology and yield better ESD robustness and part-to-part variation compared to its alternative solutions.

Massive MIMO systems will continue to evolve and there will be further need for even higher levels of integration. ADI's new high power silicon switch technology is well suited for multichip module (MCM) designs to integrate with LNAs to offer a complete, single-chip solution for TDD receiver front ends. ADI will also scale new designs toward higher frequencies and will lead similar solutions for millimeter wave 5G systems. Circuit designers and system architects will also benefit from the advantages of ADI's new silicon switches in other applications, such as phased array systems, as ADI expands its product portfolio with high power silicon switches toward X-band frequencies and higher commonly used frequency bands.

About the Author

Bilge Bayrakci is a marketing and product manager for RF and MW control products at ADI. He has an M.S.E.E. degree from Istanbul Technical University and more than 20 years of experience in the semiconductor industry. He has been with ADI since 2009. He can be reached at *bilge.bayrakci@analog.com*.

THE TOP 50
Electronics Distributors

2019 Platinum Sponsor MOUSER ELECTRONICS

	Company	2018 Global Revenue
1.	Arrow Electronics, Inc.	$29.677 billion
2.	Avnet, Inc. *	$19.040 billion
3.	WPG Holdings LTD	$18.070 billion
4.	Future Electronics Inc. **	$5 billion (EST)
5.	Digi-Key Electronics	$3.160 billion
6.	TTI Inc.	$2.800 billion
7.	Electrocomponents plc/Allied Electronics & Automation ***	$2.030 billion
8.	Mouser Electronics	$1.9 billion
9.	Smith & Associates	$1.66 billion
10.	Newark	$1.5 billion
11.	Rutronik Elektronische Bauelemente GmbH	$1.22 billion
12.	Fusion Worldwide	$1.008 billion
13.	DAC	$959.1 million
14.	Sager Electronics	$308 million
15.	America II Electronics	$283 million
16.	Master Electronics	$280 million
17.	NewPower Worldwide	$252.2 million
18.	PEI-Genesis	$250 million
19.	Bisco Industries Inc.	$202.86 million
20.	Rand Technology	$187 million
21.	The Powell Electronics Group	$173.9 million
22.	Richardson Electronics, Ltd.	$163.2 million
23.	Sourceability	$162 million
24.	RFMW Ltd.	$156 million
25.	Phoenics Electronics	$155 million
26.	Electro Enterprises	$151.462 million
27.	Classic Components	$120 million
28.	Hughes-Peters	$110.3 million
29.	Flame Enterprises	$95 million
30.	CoreStaff Co., Ltd.	$89 million
31.	Steven Engineering, Inc.	$85.32 million
32.	Chip 1 Exchange	$75 million
33.	Crestwood Technology Group	$75 million
34.	Symmetry Electronics, Corp.	$65 million
35.	Marsh Electronics, Inc.	$62.52 million
36.	Edge Electronics, Inc.	$57.9 million
37.	IBS Electronics, Inc	$52 million
38.	SMD Inc.	$45.8 million
39.	NRC Electronics, Inc.	$41.6 million
40.	Air Electro Inc.	$31.2 million
41.	Area51-ESG, Inc.	$28.41 million
42.	Diverse Electronics	$26.7 million
43.	Powertech Controls Co. Inc	$24.8 million
44.	March Electronics, Inc.	$23.98 million
45.	PUI (Projections Unlimited, Inc.)	$22.1 million
46.	Kensington Electronics, Inc.	$20.6 million
47.	Cumberland Electronics Strategic Supply Solutions Inc (CE3S)	$19.09 million
48.	Gopher Electronics Company	$18.3 million
49.	NASCO Aerospace & Electronics	$16.62 million
50.	Marine Air Supply	$16.2 million

* *SourceToday* adjusted Avnet Inc. revenue to 2018 calendar year
** *SourceToday* estimated Future Electronics global revenue
*** *SourceToday* adjusted Electrocomponents/Allied Electronics & Automation to 2018 calendar year

DUAL-CORE CONTROLLER: BETTER PERFORMANCE

A dual-core controller facilitates higher software integration by allowing different functions to execute on two independent cores. It's particularly helpful if an application requires execution of two time-critical functions periodically at a precise time or as a response to asynchronous events. With each time-critical function executing on two different independent cores, there will be no contention between the functions. This improves the overall CPU utilization due to reduced or no context switching overheads between the functions.

Many dual-core controllers come with dedicated resources that further reduce the switching and arbitration overheads. Some of the dual-core controllers also feature dedicated fast Program RAM (PRAM) coupled to one of the cores, typically to the slave core, which further improves the performance. Thus, a dual-core controller offers higher performance than a similar single-core controller operating at twice the speed.

DUAL-CORE CONTROLLER: SIMPLIFIED DEVELOPMENT

Many dual-core controllers offer dedicated memory, peripherals, and debug support with each core. A flexible resource-management scheme further allows shared resources allocation to either of the cores as per an application's requirement.

Such a microcontroller architecture enables independent software development with very minimal coordination between domain experts and facilitates easy integration. A dual-core controller particularly simplifies integrating two software stacks that are based on different architectures or require similar microcontrollers' resources, which can now run on two independent cores.

It's similar to developing the stacks to execute on two different controllers, but with the benefits of improved performance, optimal resource utilization, and reduced cost. This eliminates any complications associated with the stack integration, time-spliced resource sharing, and the associated hazard conditions.

A dual-core controller also enables easy post-integration debugging as each core comes with its own debug interfaces. Due to minimized dependencies between the stacks, the debugging gets extremely simplified to isolate issues and rectify them. Offering so many advantages, a dual-core controller significantly reduces the development risk and time to market.

To add to the list of benefits, a dual-core controller enables easy customization without modifying the main functionality. By architecting the main functionality to run on one core, the custom features can be implemented on another core. All of these offerings of a dual-core controller simplify software design even are multiple teams are involved across the globe, and they enable seamless integration with very minimal coordination efforts.

DUAL-CORE CONTROLLER: COST REDUCTION

By offering higher performance, a dual-core controller enables an embedded designer to realize complex applications using a single microcontroller. By simplifying the development, a dual-core controller drastically reduces design time and risk and enables competitive designs with reduced cost and time-to-market.

To practically realize all of the above benefits of a dual-core controller, a little experiment was conducted. In this demo, one of the cores (typically the slave core) implements motor control running a FOC algorithm to control a BLDC motor. To offer a graphical user interface, the other core (the master core) executes a graphics stack to interface an OLED display and implements a system function to interface with the potentiometer and buttons that control the speed and state of the motor.

To demonstrate the design simplicity offered by a dual-core device, the graphics stack and the motor-control software were developed by two different, geographically separated teams. With the flexibility to maintain independent software architecture, very little coordination was required between the two teams. One team with the expertise in motor control could very quickly implement the FOC algorithm to control a BLDC motor. And with the other team having expertise in graphical user interface development, both teams could leverage their experience in respective areas and quickly complete the project. There was very minimal coordination required to establish an agreement that would convey the buttons and potentiometer status between the two cores.

As an extended experiment, both the teams used already available software libraries to implement motor control and the graphics interface. As a result, the project was completed in no time, with very little effort spent on integrating two different legacy stacks. Because of the high performance of the core, a lot of CPU bandwidth was still available on both cores. To push the limits, an OLED display interface was also added on the slave core to showcase dynamic motor parameters without affecting the motor performance. Here's the live demonstration of the motor-control application with the GUI running on a dual-core digital signal controller (DSC).

An example of a dual-core controller that has all of these benefits is Microchip's dsPIC33CH128MP508 dual-core DSC. The dual-core dsPIC33CH offers high performance with dedicated memory and application-specific peripherals, suiting it for high-performance embedded, motor-control and digital power-conversion applications. The dual-core in this family enables designers to develop firmware for different system functions separately and then be brought together seamlessly without code blocks interfering with each other. ed

HARSHA JAGADISH is a product marketing manager in Microchip Technology's 16-bit microcontroller business unit. He has eight years of experience in the semiconductor industry and currently focuses on high-performance embedded solutions.

3. Shown is a multi-function complex application block diagram. *(Courtesy of Microchip)*

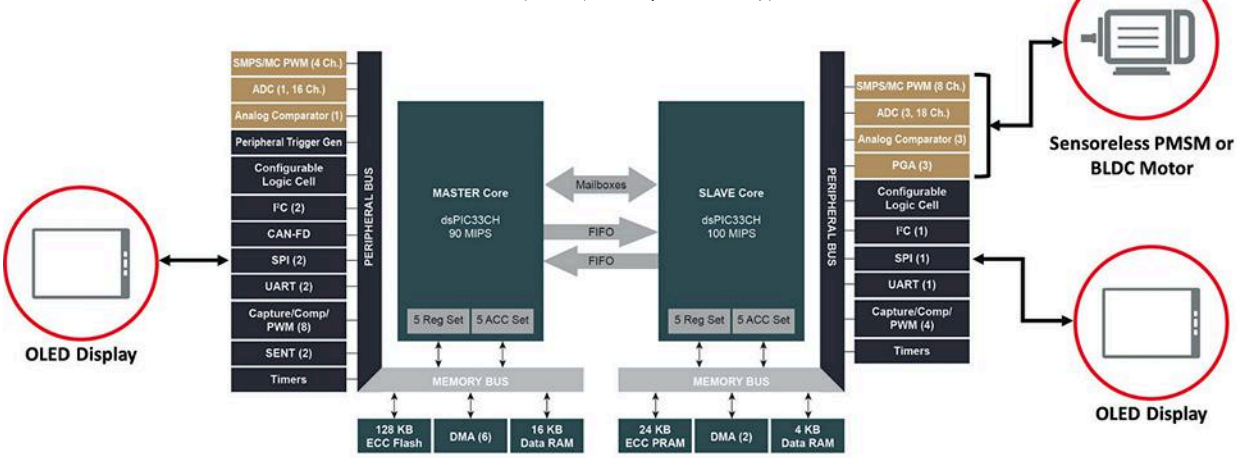

Using Simulation to Design Cooling Systems for Electronics Webinar

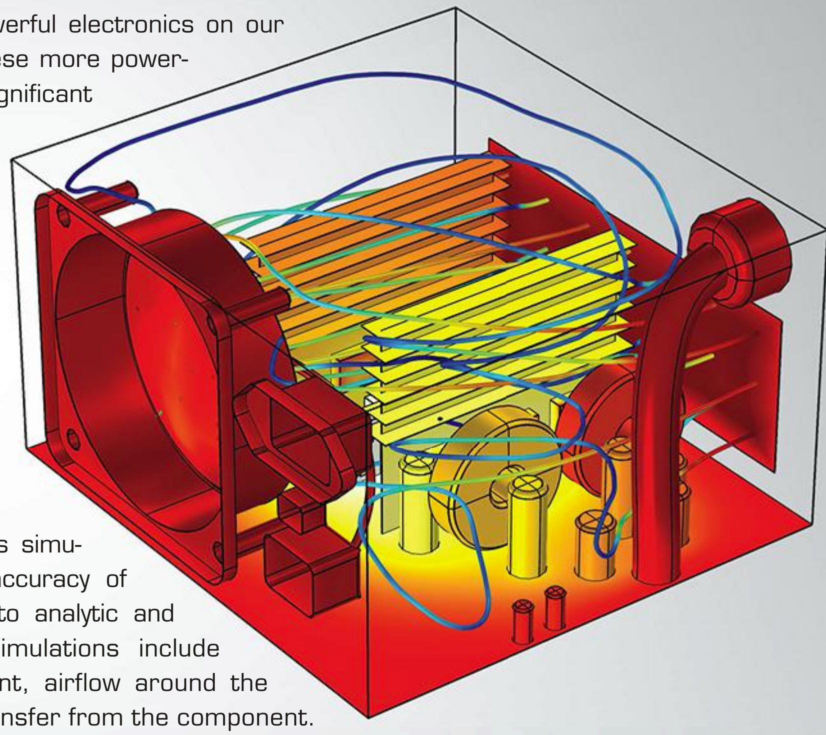

The impact of increasingly powerful electronics on our society cannot be overstated. These more powerful electronic systems produce significant heat that must be dissipated to prevent premature component failure. Engineers that design electronics face a significant thermal management challenge. This task becomes even more challenging when the cooling system relies on natural convection instead of forced convection from fans due to the relatively short life expectancy of fans.

Engineers can use multiphysics simulation software to improve the accuracy of their calculations in comparison to analytic and single-physics solutions. These simulations include heat generated by the component, airflow around the component, and radiative heat transfer from the component. Heat generation due to resistive heating in the board can be included with heat generated from components to determine the heat generated within the system. Airflow through the system due to either forced or natural convection can also be analyzed. For many systems, radiation must be considered for accurate temperature predictions due to the large amount of heat transfer that occurs via this mechanism in many electronics cooling problems.

In this presentation, Kyle Koppenhoefer from AltaSim Technologies will discuss the analysis of an electronics cooling problem subjected to a complex thermal environment. The webinar will also include a live demo in the COMSOL Multiphysics® software and a Q&A session.

REGISTER ONLINE AT:
www.electronicdesign.com/webcasts/using-simulation-design-cooling-systems-electronics-webinar

Speaker

Kyle Koppenhoefer, Principal, AltaSim Technologies

Kyle Koppenhoefer has been one of the principals and leaders at AltaSim Technologies for 15 years. He works with customers to identify how computational analysis can be used to further develop their products and manufacturing processes. Prior to cofounding AltaSim, Kyle worked for the Department of Defense and the Edison Welding Institute. He holds a PhD in civil engineering from the University of Illinois.

RON SQUIERS | Solutions Architect, Mentor, a Siemens Business

Virtual Emulation Knows No Bounds for Networking Design Verification

Virtual emulation is necessary to scale verification technologies with the growing speeds, capacities, and port counts of Ethernet designs, which is where the VirtuaLAB Ethernet App steps in.

With the remarkable surge in demand for connectivity, the internet has become the communications highway for billions of users. According to "Internet Live Stats," 4.2 billion people traveled the internet in 2019. That's about 55% of the world's population, and the amount of traffic generated by this connected population is truly staggering.

The vast web of computers and mobile devices that transmit, receive, and deliver this data wouldn't even be possible without advances in networking technologies. Ethernet is an iconic example—one that continues to evolve with astonishing rapidity.

Since the 1980s, when a seven-port Ethernet switch was introduced, port count has scaled dramatically. Today's Ethernet switches and routers now reach 1000 ports, and by the end of 2020, will accommodate 4000 ports. Also, in contrast to the throughput of 10 Mb/s of the seven-port Ethernet switches, they can now handle throughputs of 1/10/25/40/50/100/200/400 Gb/s.

While the Ethernet Alliance can foresee the future growth in number of ports, the increase of bandwidth to 1600 Gb/s is less obvious due to limits in the transmission medium. Dense wavelength division multiplexing and protocol analysis module signaling have been used to overcome these limits to date. We may anticipate a move to adopt more parallelism to expand the bandwidth. Latency in networking switches has constantly decreased, and today the lowest latencies have dropped below 1 µs.

A larger number of ports, expanding throughput, decreasing latency, and overall improvement in security and ease of use are making today's network switches and routers among the largest IC designs ever developed, reaching well beyond a billion gates and now on par with the largest processor and graphics chips. Verifying such complex IC designs, before silicon availability, is a daunting task. Fortunately, help has already arrived in the form of hardware emulation and, specific to Ethernet, with the VirtuaLAB Ethernet solution.

This article explains some of the technology behind why virtual emulation is the only way to keep up with the growing speeds and capacities of Ethernet designs.

IN-CIRCUIT EMULATION LIMITS

For example, a design of a modern system-on-chip (SoC) with a 256-port Ethernet interface and a variable bandwidth of 1/25/40/50/100/200/400 Gb/s surpasses the capabilities of traditional digital simulators. While hardware description language simulation can be used at the block level, verification of an entire design of a billion gates with simulated traffic is unrealistic and must be ruled out. This is a primary case for using hardware emulation in what's called in-circuit-emulation (ICE) mode.

Unique to the ICE verification mode is the ability to test a design with real traffic. An ICE configuration requires one Ethernet tester per port. Since a direct connection isn't possible due to the different speed domains between the tester and the emulated design under test (DUT), a speed rate adapter is inserted between the two. This reconciles the fast speed of the tester to the relatively low speed of the emulated DUT.

But when we tackle a modern switch design that can support a port configuration of 256 ports and variable bandwidths of 1/10/25/40/50/100 Gb/s, we start to see some cracks in the ice. Because the design has 256 ports, the ICE setup requires 256 Ethernet test ports, 256 Ethernet speed adapter ports, as well as CAT5 and IO cables to the emulator system (1K cable channels) *(Fig. 1)*. An ICE approach becomes difficult, not only from a cabling perspective, but also from a traffic-shape accuracy perspective when exercising a DUT through speed adapter ports.

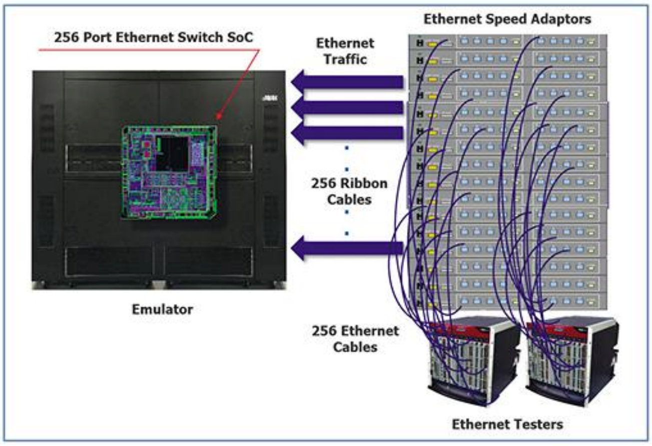

1. A 256-port Ethernet switch is verified using ICE.

Apart from the spaghetti-cable arrangement, the potential unreliability of the hardware, overall cost, and loss of traffic-shape accuracy, a frustrating limitation is that the entire setup supports only a single user located in the proximity of the emulation lab. Time-sharing using modern parallel and batch mode techniques is extremely difficult.

UNTANGLING ETHERNET DESIGN VERIFICATION

On the other hand, a virtual-emulation environment frees users from the limitations and restrictions of ICE. Let's compare the ICE setup to an equivalent one using a virtual approach. In the virtual scenario, Ethernet testers are modeled in software running under Linux and/or virtual machines on a workstation connected to an emulator. This model is an accurate representation of the actual physical tester, based on proven-in-implementation intellectual property.

Virtual testers typically include Ethernet Packet Generator and Monitor (EPGM) functions that do exactly that—generate, transmit, and monitor Ethernet packets exchanged with the DUT. It should have the ability to configure any GMII, XGMII, 25GMII, XLGMII, 50GMII, CGMII, 200GMII, 400GMII, or Physical Coding Sub-layer (PCS) interfaces for 1G, 10G, 25G, 40G, 50G, 100G, 200G, 400G, respectively.

The interface between the virtual tester and the DUT can consist of one too many instances of the data programming interface (DPI) to scale up the emulation test system to support terabit throughputs for the largest networking systems. Packet stimulus for transmission and packet reception for analysis flows bidirectionally through an EPGM tool to/from a virtual Ethernet transactor and Null-PHY connected to the DUT. One xRTL transactor is required for each port of any Media Independent Interface (xMII) or PCS-supported type *(Fig. 2)*.

For example, VirtuaLAB Ethernet can provide up to 64 ports per workstation for any speed type. Multiple VirtuaLAB applications are able to be bundled together across multiple workstations to support large port count configurations. High-speed optical link cards are used to co-model host workstations to the emulator.

This tightly integrated transport mechanism is tuned for maximum wall-clock performance rates and is completely transparent to the testbench. VirtuaLAB Ethernet's parallel runtime architecture has been implemented such that data-plane emulation throughput scales linearly with co-model port count *(Fig. 3)*. For example, emulation throughput performance with VirtuaLAB Ethernet for a 256-port SoC is in excess of 1T packets per day.

Apart from enabling high data-plane transport, several other benefits come from using this approach. First, reconfiguring the virtual tester to perform various functions is fast via remote access. Second, the workstation is a stable and reliable piece of equipment acquired at a fraction of the cost of a complex Ethernet tester of equivalent functionality. Third, these workstations fall neatly into standard data-center-grid engine requirements.

Even more important is the ability to handle multiple concurrent users and concurrency across distributed software models that are essential in supporting large software development teams. Emulation hardware and software architecture views must mind this scalability aspect from the ground up.

A virtual Ethernet device should support a directed test mode of operation to define and control the streaming of specific packet types, traffic shape, scheduling and error injection into the emulated DUT, as well as trace and analysis of packets returned from the DUT. VirtuaLAB Ethernet is an example of how Tcl, Python, and GUI interfaces can ease generation, analysis, and scoreboarding through rules-based user definitions.

VirtuaLAB Ethernet can be configured for multiple co-model hosts supported by one instance of software running on up to 16 co-model hosts, with up to 1K Virtual Ethernet xRTL Transactors connected to Null-PHYs and to the DUT ports on the emulator. These 1K VirtuaLAB instances can be managed through a centralized "controller" software application invoked from one of the workstations.

COMPLEX TESTING AND ANALYSIS

A virtual Ethernet solution is well-suited for complex test scenario generation, monitoring, and automation. For example, an interactive and batch-mode Tcl/Python programming command line interface (CLI) can be incorporated to control all media access controllers, and to generate myriad protocol and traffic flow scenarios. VirtuaLAB Ethernet supports traffic construction CLIs over non-homogenous packet types. Percentages of traffic per protocol type can be mixed, with differing packet sizes or random size per flow or stream. It's also possible to control packet transmission and arbitration.

Driven by the high rate of networking technology convergence, virtual Ethernet is part of a class of virtual networking solutions that includes other related protocols, such as Flexible Ethernet, 5G Enhanced Common Packet Radio Interface, and Optical Transport Networking. Virtual Ethernet solutions also perform offline traffic analysis, per port/flow statistics, and many other complex analysis functions. OSI Level 1 – Level 7 stack testing should also be provisioned.

Moreover, VirtuaLAB Ethernet supports stress testing and error injection for complex switching topology using the mutable (dynamic) port

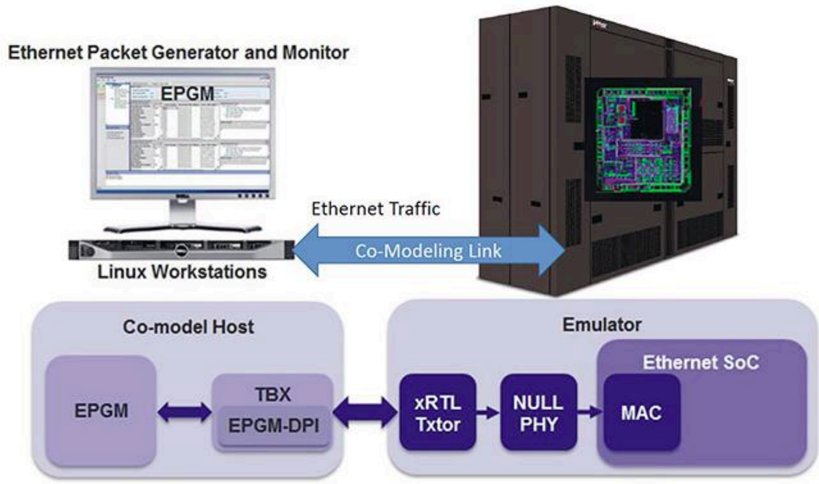

2. One instance of EPGM-DPI communicating to a Virtual Ethernet xRTL is all that's required.

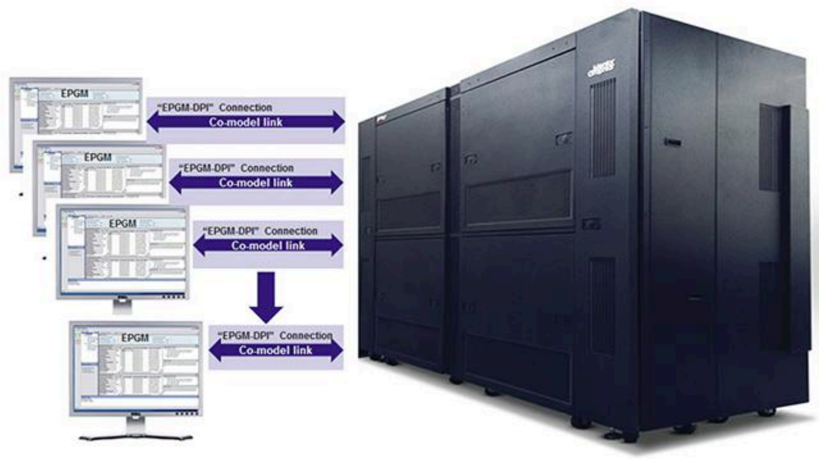

3. This is a high-level view of the multi-co-model topology.

group reconfiguration features for 1/10/25/40/50/100/200/400-Gb/s full-duplex speeds. For example, xMII/PCS widths, link speeds, link up/down, and fault states can all be dynamically configured during emulation runtime, and need no recompile steps to support testing numerous port group configurations.

Many protocol and performance violations are automatically reported, such as cyclic redundancy check, preamble, inter-frame gap (IFG), line rate, and forward error correction. Packets can be reviewed in interactive or batch sessions to examine packet statistics per flow/port, ingress/egress trace, meta-data such as signatures, time stamps, latency and content all on the wire.

One of the hallmarks of virtual Ethernet solutions in emulation is their ability to do complex performance analysis of large and deeply state-full SoC and multichip systems. Virtual Ethernet in emulation has been used in validating packet classification, filtering, bandwidth rates, policing, traffic shaping, class-of-service, congestion, memory management, packet sequence drops, IFG analysis, dynamic load balancing, group multipath and network fabrics using signature analysis features available in virtual Ethernet systems.

Virtual Ethernet solutions reside entirely within software, so anything can be imagined and implemented quickly for capture and analysis with 100% visibility. Examples include custom and vendor-specific protocols, support for higher function analysis add on and even non-standard packet interface types. Virtual Ethernet solutions easily accommodate architectural exploration, integration with third-party tools, and bypass SoC logic that has little value in being re-verified for systems seeking greater test throughput for subsystems.

For some complex test and analysis, millions of packets may be needed to hit steady state or transient analysis points of interest in the design. Take, for example, measuring bit rates for traffic flows in a terabyte Ethernet switch with a 32-MB packet buffer. Virtual Ethernet signature generation and packet time stamps can be used to measure cut-through and store and forward latency conditions, RFC2544, RFC2889, etc.

TESTING WORLD-CLASS STREAMING DEVICES

Virtual Ethernet provides a software-controlled environment for generating, transmitting, and analyzing Ethernet packets to test Ethernet SoCs with an emulation platform. Because the emulator runs on enterprise servers, multiple users located anywhere in the world can test over 20 million SoC packets per day concurrently. VirtuaLAB Ethernet brings complex Ethernet SoC designs to market on schedule with significantly increased productivity.

RON SQUIERS is a solutions architect who is passionate about innovation, strategic product development, and positioning in the networking domain. He has developed next-generation networking, communications, and encryption products for Mentor Graphics, Broadcom, Hewlett Packard, SGI, AMD, and several startups. Ron received his BSEE from the University of California at Berkeley and graduate studies at Stanford University. He has several patents in networking, encryption, and emulation. Ron is a member of IEEE, Ethernet/WiFi Alliance, and Tau Beta Pi organizations, and lives and works in the San Francisco Bay Area.

BOB ZOLLO | Keysight Technologies, Inc.

The Effect of Wire-Path Resistance on Battery Measurements

Wire and contact resistance, which can lead to voltage drops, will wreak havoc when testing cells. Here are some tips that can help you alleviate those headaches.

Proper wiring and probing/contacting are critical to making successful measurements. When testing cells or batteries, using good engineering practices is important to avoid noise pickup, signal degradation, and unwanted voltage drops through wiring and contact resistances.

If you're testing cells or batteries at high currents, resistance will be a critical consideration. As test currents rise, the resistance in the wires and in the contacts will cause significant voltage drops.

Discharging	
V_drop =	I_in_wire * (R1 + R2 + R3 + R4)
V_input =	V_cell - V_drop

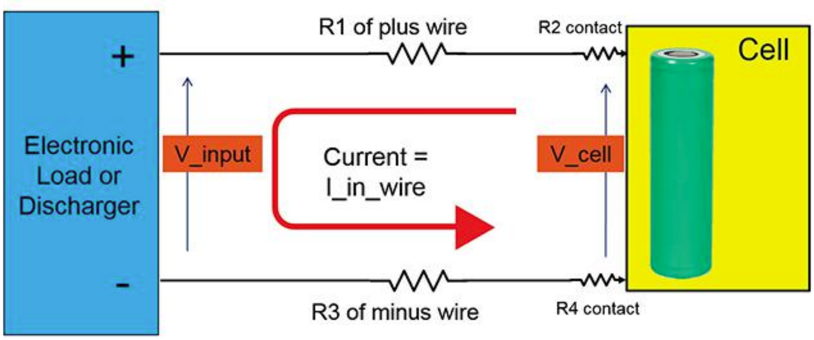

Impact of wire and contact resistance of voltage at input to electronic load or battery discharger.

WIRE-PATH RESISTANCE VOLTAGE DROPS ARE PROBLEMATIC DURING BATTERY DISCHARGE

These voltage drops can cause problems during the discharge of a battery into an electronic load, as electronic loads have a minimum operating voltage. Note when I say electronic load, this could be a benchtop electronic load or the discharge circuits in a battery tester (sometimes called a battery cycler). As the battery voltage drops during discharge, the battery voltage may approach the minimum operating voltage of the electronic load.

If large voltage drops occur in the wiring, then the voltage at the input to the electronic load may drop below its rated minimum operating voltage. When the electronic load is operated below its minimum operating voltage, the available current will be derated and you may not be able to draw the required current out of the battery.

COMPARING TWO DISCHARGING-CELL CASES

Let's look at the impact of wire and contact resistance for two cases of discharging a high capacity lithium-ion cell (*see table*). The *figure* shows a schematic diagram of the discharge setup.

So, the keys to successful discharge operation are:

- Use the shortest wire possible.
- Use thick wires to minimize the voltage drop in the wires.
- Use good quality probes to minimize contact resistance.

These are necessary to keep the total resistance low so that the voltage drop is minimal at high current during discharge.

OTHER CONSIDERATIONS

Is the cable well-constructed or does flexing the cable change the resistance? Inconsistent wire resistance from channel to channel can make some channels work properly because they meet wire resistance criteria to keep the voltage above the electronic load minimum input voltage when drawing high current. Meanwhile, other channels have too much resistance and thus don't work properly, causing channel-to-channel behavior variation.

Is the contact resistance changing over time? Contacts can become worn out and pin-type probes can go out of alignment, causing inconsistent pressure and inconsistent resistance. Contacts also may become dirty from dust and chemicals. Monitoring the number of cycles of use and performing preventative maintenance (cleaning, alignment, and replacement) can reduce these effects.

Furthermore, the batteries themselves could be a cause of inconsistent contact resistance because the battery contacts may be oxidized. Contacts are often made of materials that easily build up an oxide coating over time.

	Case 1 Long wires Insufficient wire gauge High contact R	Case 2 Short wires Proper wire gauge Low contact R
Lowest voltage on lithium-ion cell during discharge; this is the end-of-discharge voltage (EODV)	3.0 V	3.0 V
Discharge current	75 A	75 A
Minimum voltage input specification of the electronic load	2 V	2 V
Maximum current rating of the electronic load	100 A	100 A
Length of wire = the full wiring path from the electronic load to the (+) of the cell and from the (−) of the cell back to the electronic load	20 feet	8 feet
Size of the wire	8 AWG	0 AWG
Resistance of the wire per foot	0.0006285 Ω/foot	0.0000983 Ω/foot
Wire resistance	0.01257 Ω	0.0007864 Ω
Contact resistance of the test probes to the cell	0.02 Ω	0.005 Ω
Total resistance of the wiring path = wire R + contact R	0.03257 Ω	0.0057864 Ω
Impact of wire R and contact R	The wire R and the contact R both contribute significantly to the total R	The contact R is the dominant factor in the total R
Total voltage dropped in the wiring path = Total R * discharge current	2.44 V dropped in wire path	0.43 V dropped in wire path
Resultant voltage at the electronic load input = Discharge voltage − voltage drop in wire path	0.56 V	2.57 V
Electronic load condition	The electronic load is operating **far below** minimum voltage so 75-A loading is not possible	The electronic load is operating above minimum voltage so 75-A loading is possible

CONTACT RESISTANCE IMPACT ON VOLTAGE MEASUREMENTS

Now let's look at a measurement of voltage. Is contact resistance important there? Can measurements of cell voltage be distorted or impacted by contact resistance?

The short answer is no—contact resistance will not impact the measurement of voltage. Voltage measurements are made with the DMM configured with high impedance input. For a typical DMM measuring voltage, the input impedance is hundreds of megohms or even gigaohms. Because the input impedance is so high, virtually no current flows into the DMM measurement terminals. If no current is flowing, then resistance in the path has no effect on measurement.

But what about a battery tester? Is the input impedance as high on a battery tester as found on DMM?

Achieving very high input impedance on a measurement instrument is challenging. Most battery testers measure voltage in the same manner as a DMM, although the input impedance is not usually gigaohms because battery testers aren't trying to measure the wide range of devices that a standard DMM is expected to measure.

For example, the Keysight BT2152A is a battery test instrument focused on measuring self-discharge of lithium-ion rechargeable cells. (For more information about measuring self-discharge of lithium-ion cells, visit *www.keysight.com/find/bt2152a.)* When measuring cell voltage with the BT2152A, the input impedance is 10 MΩ. At the maximum cell voltage of 4.2 V (typical) and 10 MΩ of input impedance, the current flow into the BT2152A's input is 420 nA. This value is found using Ohm's law of V = I*R or I = V/R, so for a 4.2-V cell connected to a measurement system with 10-MΩ input impedance, current is calculated as 4.2 V/10 MΩ, or 420 nA.

Since the input impedance is 10 MΩ, only 420 nA of current will flow. With total wire-path resistances of far under 1 Ω (*see table*), the voltage drop across the wire is under 420 nV. Even if wire and contact resistances were to become thousands of times bigger than what's shown in the table and climb to 10 Ω, the voltage drop in the wire resistance would jump to 4.2 μV, which is insignificant measurement error compared to the 4.2 V being measured on the cell. As you can see, thanks to the high impedance input, the resistance of the wires and contacts don't impact the voltage measurement.

However, whenever you're making voltage measurements, noise pickup needs to be minimized by proper grounding and shielding, but that's the subject for another article. ed

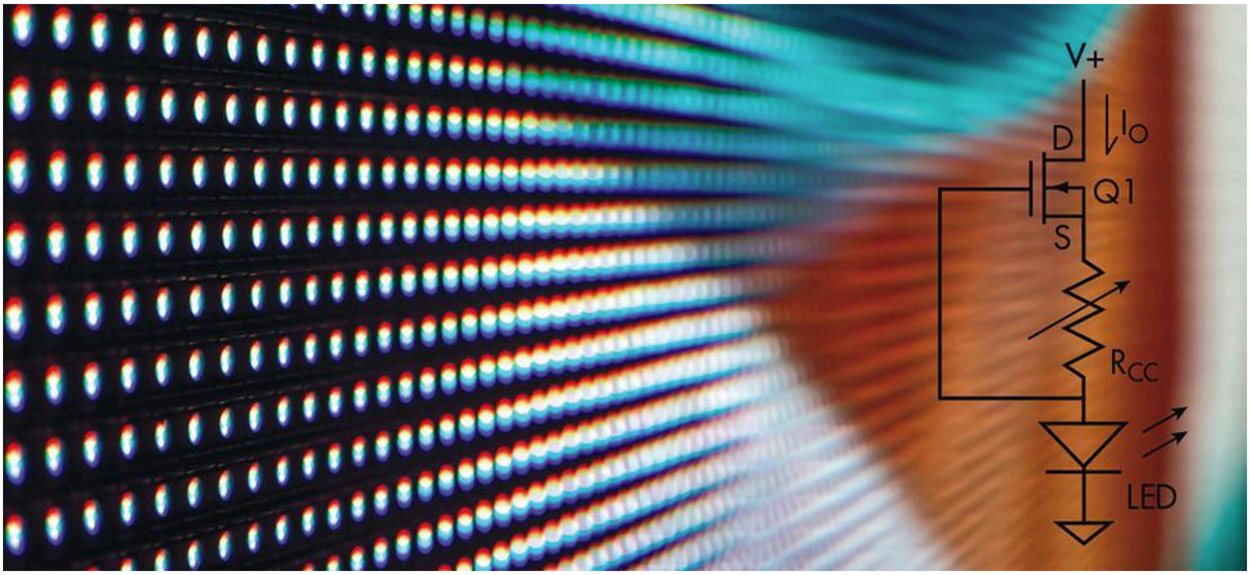

Enhanced VLED Driver Updates JFET Constant-Current Source

DAN AWTREY | Engineer

HIGH-BRIGHTNESS VISIBLE-LIGHT LEDS (VLEDs) require a constant-current source. This simple, RFI-free circuit takes advantage of the normally-on properties of the new, high-power depletion-mode (Dmode) MOSFETs to provide this source.

Historically, the simplest method for driving a load such as an LED uses a fixed or variable resistance between the power source and the load *(Fig. 1)*. While low cost, the current doesn't remain constant, but varies with supply voltage and temperature rise with current load. The poor efficiency can be improved by replacing the passive resistor with an active component or circuit whose impedance automatically adjusts to maintain a specific value of current, even with varying supply voltage and/or heavy loads.

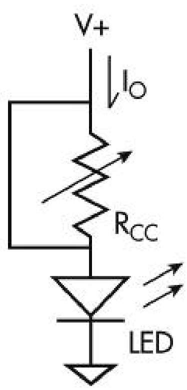

1. Using a current-limiting resistor is the simplest to control the current through an LED.

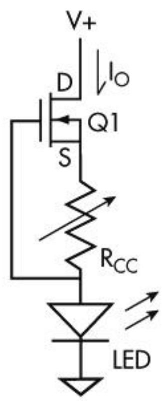

2. A better approach is to use a FET as a constant-current source.

Constant-current sources constructed using the normally-on properties of a JFET or Dmode FETs are restricted to very low power levels. These devices are controlled by voltage rather than current as in the older bipolar junction transistors. While small-signal JFETs are available in both p and n types, the new Dmode FETs are able to manage the heavy currents required by the new blue and white VLEDs for proper operation, but are currently restricted to n-types.

In a traditional Dmode FET constant-current circuit *(Fig. 2)*, only the FET and a resistor (fixed or variable) is required. Current flowing through the FET generates a voltage across the resistor with a value chosen to generate the level of pinch-off voltage ($V_{GS(off)}$) necessary to provide the desired constant current *(Io)*. The two-terminal CCS allows the load to be arbitrarily connected to either the drain or, most commonly, the gate terminal as shown.

The IXTP3N50D2 Dmode FET from IXYS Corp. powers a 5-W white VLED such as the LZI-00NW05 from LEDengin (190 lumens at 1 A); the design requires a 1-to-2 W resistor (R_{cc}) between 2 to 10 Ω. For maximum efficiency, the supply voltage should be as near the forward voltage (V_F) required by the VLED(s) for minimum supply voltage across the circuit.

Though the traditional CCS is compact and efficient, it provides only a limited range of current control with a variable R_{cc}, and usually requires an individually selected, fixed-value power resistor for each FET. This is because the pinch-off voltage ($V_{GS(off)}$) required to control the FET varies greatly from device to device, as does the V_F of the LED. This tends to make it impractical for volume production.

The new design allows for control of the pinch-off voltage ($V_{GS\ (off)}$) of the Dmode FET. It replaces the high-wattage series R_{cc} rheostat of the standard circuit with a series VLED while connecting a high-value, wide-tolerance potentiometer (voltage divider) in parallel with the VLED *(Fig. 3)*.

In the complementary combination of a Dmode FET and a VLED, the FET supplies a current controlled by voltage while the LED generates a voltage dependent on current, and the VLED is driven directly from the drain of Q1. Thus, maximum current flow is supplied to the VLED when the potentiometer is set at the source end of R_{cc}, with decreasing values selected by R_{cc} until the minimum current value (determined by the forward voltage of the VLED and the $V_{GS(off)}$ of the FET) is reached at the ground terminal of the potentiometer.

To avoid possibly damaging current

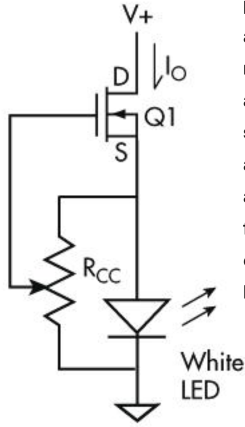

3. Using a potentiometer across the LED, rather than a rheostat in series, creates a more-efficient architecture for constant-current, controllable LED drive.

flow through the LED, it may be necessary to insert a small resistance between the gate and R_{cc}, as directly connecting the source and gate together turns the FET fully on. The V_F of a single white VLED ranges from 3.6 to 3.8 V, and the $V_{GS(off)}$ of the IXTP3N50D2 FET ranges from −2 to −4 V. This design provides a convenient method of varying the brightness of the VLED over a wide range without regard to the characteristics of a particular FET or VLED, which may be of any color or wavelength. These values also show that the V_F of a single white VLED is insufficient to allow the potentiometer to turn the control FET completely off *(Fig. 4)*.

If a full-on to full-off range is required, the simple solution is to use two LEDs in series. This works over the author's unit with a 2.8- to 12-V supply range and would work as well up to the voltage limit of the FET; however, it requires a heat sink. A better approach would be to add more LEDs in series to match the supply voltage for heating and efficiency reasons. For example, the 500-V, 3-A rating of the IXTP3N50D2 could control three parallel strings of 100 5-W VLEDS in an offline application.

Rather than the actual value or tolerance of the potentiometer, it's the ratio of the resistance value above the wiper to that below the wiper, then multiplied by the V_F of the VLED, that's important in setting the voltage on the Dmode gate.

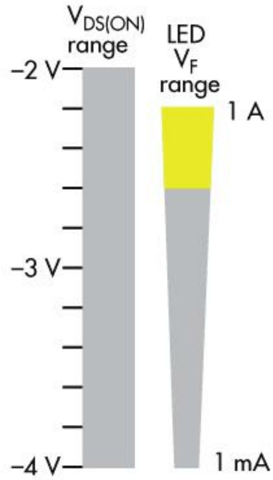

4. Overlap exists between FET control voltage with VLED intensity; yellow shows normal LED operating range.

The actual value of the potentiometer only determines the current it consumes and is equal to n times the VLED's forward voltage, then divided by the R_{cc} value across the VLED (where n is the number of LEDs in series).

For a single VLED with an arbitrary 100-kΩ potentiometer across the VLED, the current is only 0.037 mA (0.14 mW), so any low-wattage trim pot is suitable, which is orders of magnitude better than using a series rheostat. Also, since a Dmode FET is voltage-controlled, the value of the potentiometer could as easily be several times this value and still stay within the transistor's specifications, further reducing control-circuit power loss to a negligible value. When two (or more) VLEDs are connected in series, the potentiometer value may be increased proportionally to 200 kΩ (or more) to maintain the same level of power consumption.

Further, as the circuit contains no reactive or switching components, it has a unity power factor. Also, it doesn't create EMI, so Part 15 of FCC regulations concerning emissions don't apply.

DAN AWTREY (now retired) was in engineering and engineering management, and is a patent holder and author of over 20 articles in print and on the web. He last appeared in *Electronic Design* (May 1996) with "A Bootstrap Voltage Converter Using A Digital Gate."

(Continued from page 18)

Clearly, we have taken our reliable wired networks for granted. They're considered old technology and get no recognition or appreciation. However, despite this inattention, even disrespect, wired networks continue to improve. The cable TV HFC systems keep gaining more speed thanks to the latest upgrade to the DOCSIS standard. It can easily handle 1-Gb/s data, and even faster in some instances.

DSL has also ramped up its speed over the years as companies invented technologies that let you put 1 Gb/s of data on a twisted pair. And, after all these years, the telecom companies discovered how to minimize 60-Hz induced voltage and current interference to the PSTN wiring and DSLAMs by using the mysterious Induction Neutralizing Transformer (INT). The INT has been around for a while, but its effectiveness hasn't been fully exploited. Fast data can be achieved over twisted pair by using the latest standards such as VDSL2, G.fast, and others thanks to the INT.

As for other examples of a healthy wired infrastructure, consider these. Many of the small cells needed for 5G will be fiber. IoT nodes are aggregated in gateways wired to some other wired LAN or internet connection. DSL and cable internet connections to homes and businesses aren't going away. New 5G broadband wireless may take some business, but not all.

If you want an in-depth look at this issue, get a copy of the new book, *Computer Networking Breakthroughs You've Always Wanted, Without Needing Fiber Optic Cable…Even in the Age of the Internet of Things.* It's written by an acquaintance of mine, Russ Gundrum, a veteran of the telecommunications and cable businesses. He's also the guy to thank for promoting and deploying the INT.

My conclusion is that wired networks still serve a real purpose. We actually can't get along without them. Yes, wireless is getting all of the attention right now, but we should pay some respect to the wired networks we use every day. ed

INDEX

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

WILLIAM WONG | Senior Technology Editor
bill.wong@informa.com

The Ever-Improving Inference at the Edge

Using machine-learning inference on the edge has never been easier with platforms like NVIDIA's Jetson Nano.

Not all applications can utilize machine-learning (ML) inference, but it's possible with most. Doing it at the information source instead of in the cloud is becoming easier thanks to improved artificial-intelligence (AI) software support plus hardware acceleration that targets deep neural networks (DNNs). Platforms like Renesas' e-AI, STMicroelectronics' STM32CubeMx.AI, and NXP's eIQ all support ML and target hardware from conventional microcontrollers to systems with hardware acceleration.

ML hardware acceleration can significantly improve the performance of inference applications on the edge, opening up new application opportunities that would not be possible on stock hardware. GPGPUs and multicore CPUs led the charge, but ML-specific hardware has the edge. Even the latest version of these platforms have been enhanced to address the inference chores. For example, Intel's latest Xeons include instructions targeting ML and its Movidius video processing unit (VPU) zeros in on specific ML application spaces.

NVIDIA's Jetson Nano *(see figure)* brings a full SoC to the ML table. The 128-CUDA-core Maxwell GPGPU handles processing of most of the DNN models assisted by the 64-bit, quad-core Cortex-A57 CPU cluster. The compact DIMM module also includes 4 GB of DRAM and runs Linux. Its hardware encode and decode support can process a 4K or eight 1080p video streams while running ML models on each stream. Convection cooling easily handles the 5 to 10 W of power, allowing the Jetson Nano to work in compact, low-power AI applications on the edge.

Coprocessors are also answering the call for more efficient inference and identification chores in embedded systems, where a batch size of one is important. Servers typically handle large batch sizes more efficiently, but they're also working with larger datasets versus embedded systems that might have a single camera delivering data for analysis.

Chips like Flex Logix's InferX X1 target this space. The chip incorporates multiple nnMAX processing tiles specifically designed to handle each layer within a DNN model that's been trained on a server using boards like NVIDIA's latest Tesla T4 or FPGA boards such as Xilinx's Alveo or Intel's Programmable Acceleration Cards (PACs). The InferX X1 is optimized to implement Winograd acceleration, which can improve accuracy and performance of INT8 layers by 2.25. The system transforms the 3-by-3 convolution to a 4-by-4 with dynamic translation of weights to 12 bits. The support also handles input and output translation on-the-fly, minimizing the loading of weights within the system.

Figuring out whether AI will benefit an application is a chore in and of itself. However, once that determination is made, lots of options are available to implement these systems. Of course, one may have to apply AI just to wade through the options. ed

The Jetson Nano from NVIDIA is an SoC that supports machine-learning inference chores in embedded systems.