# Electronic Design®

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

NOVEMBER/DECEMBER 2018
electronicdesign.com

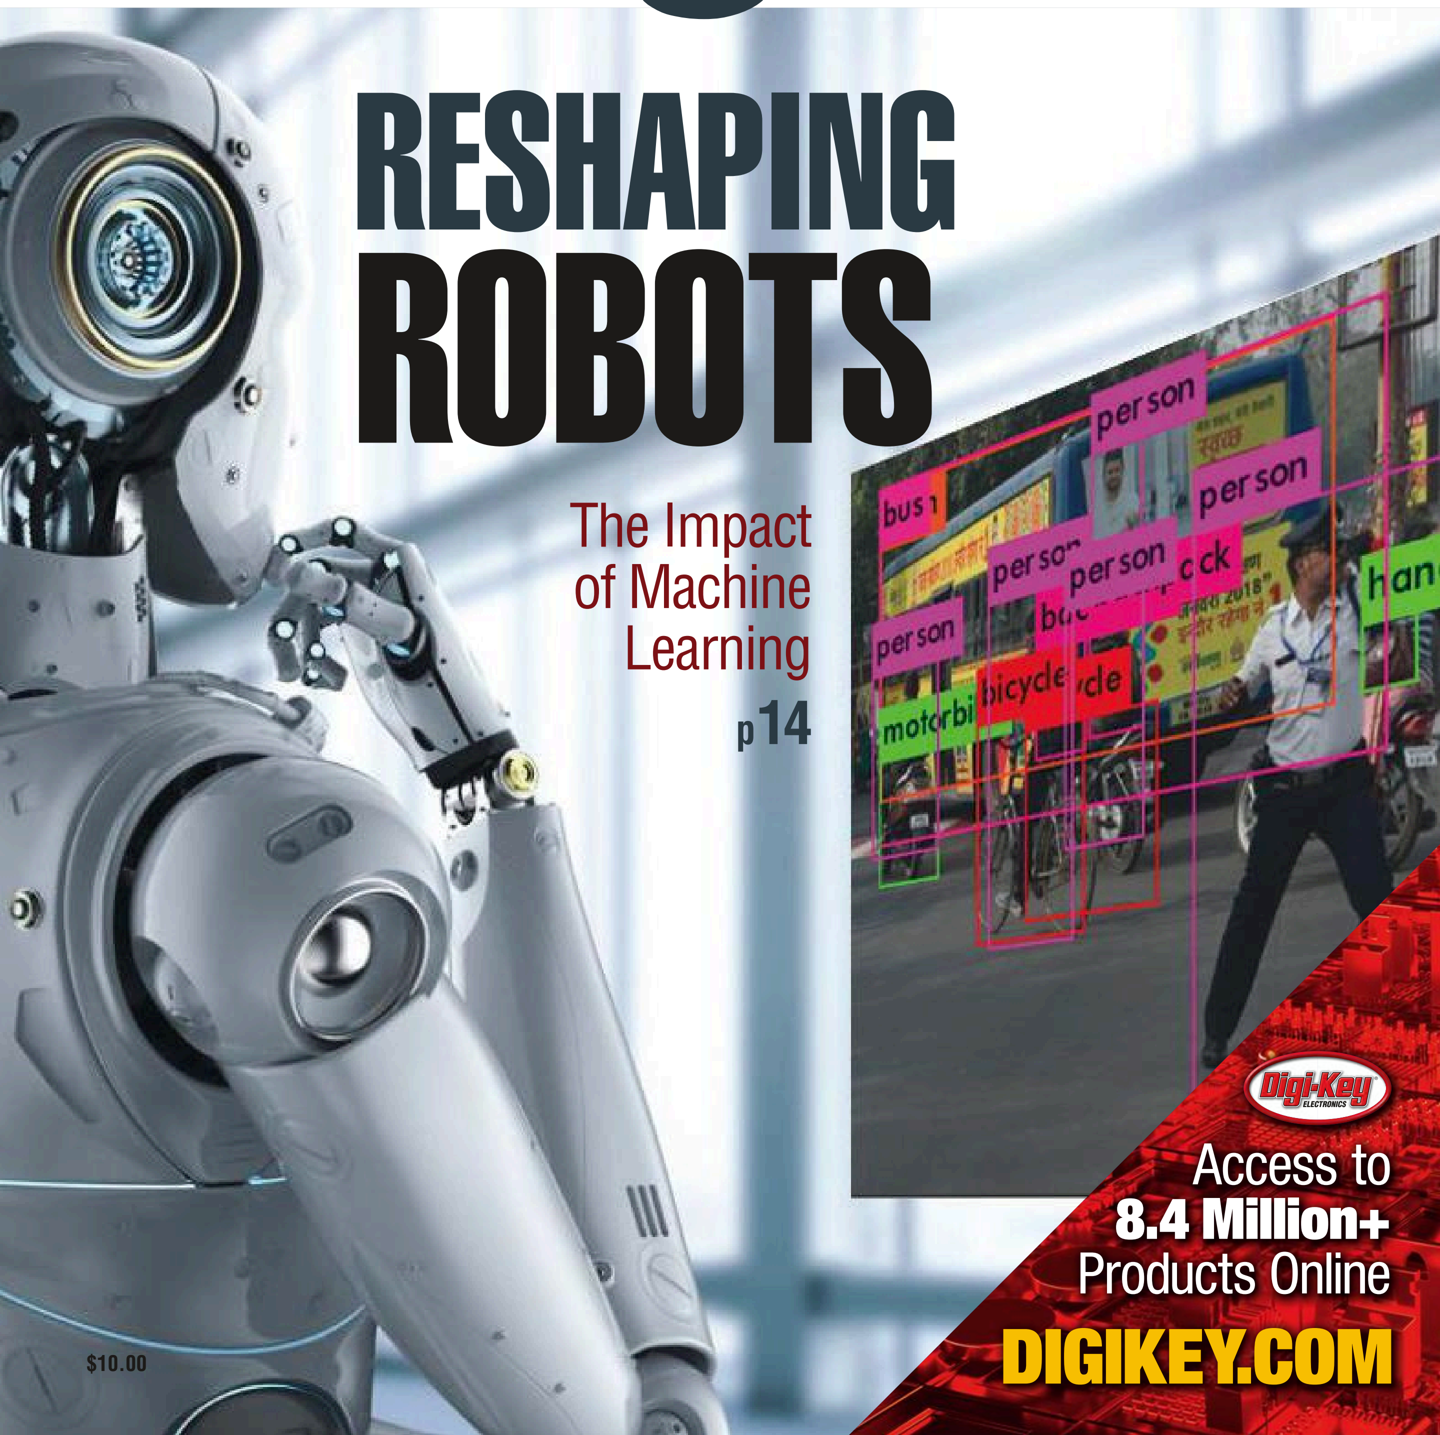


# RESHAPING ROBOTS

## The Impact of Machine Learning

p14



$10.00

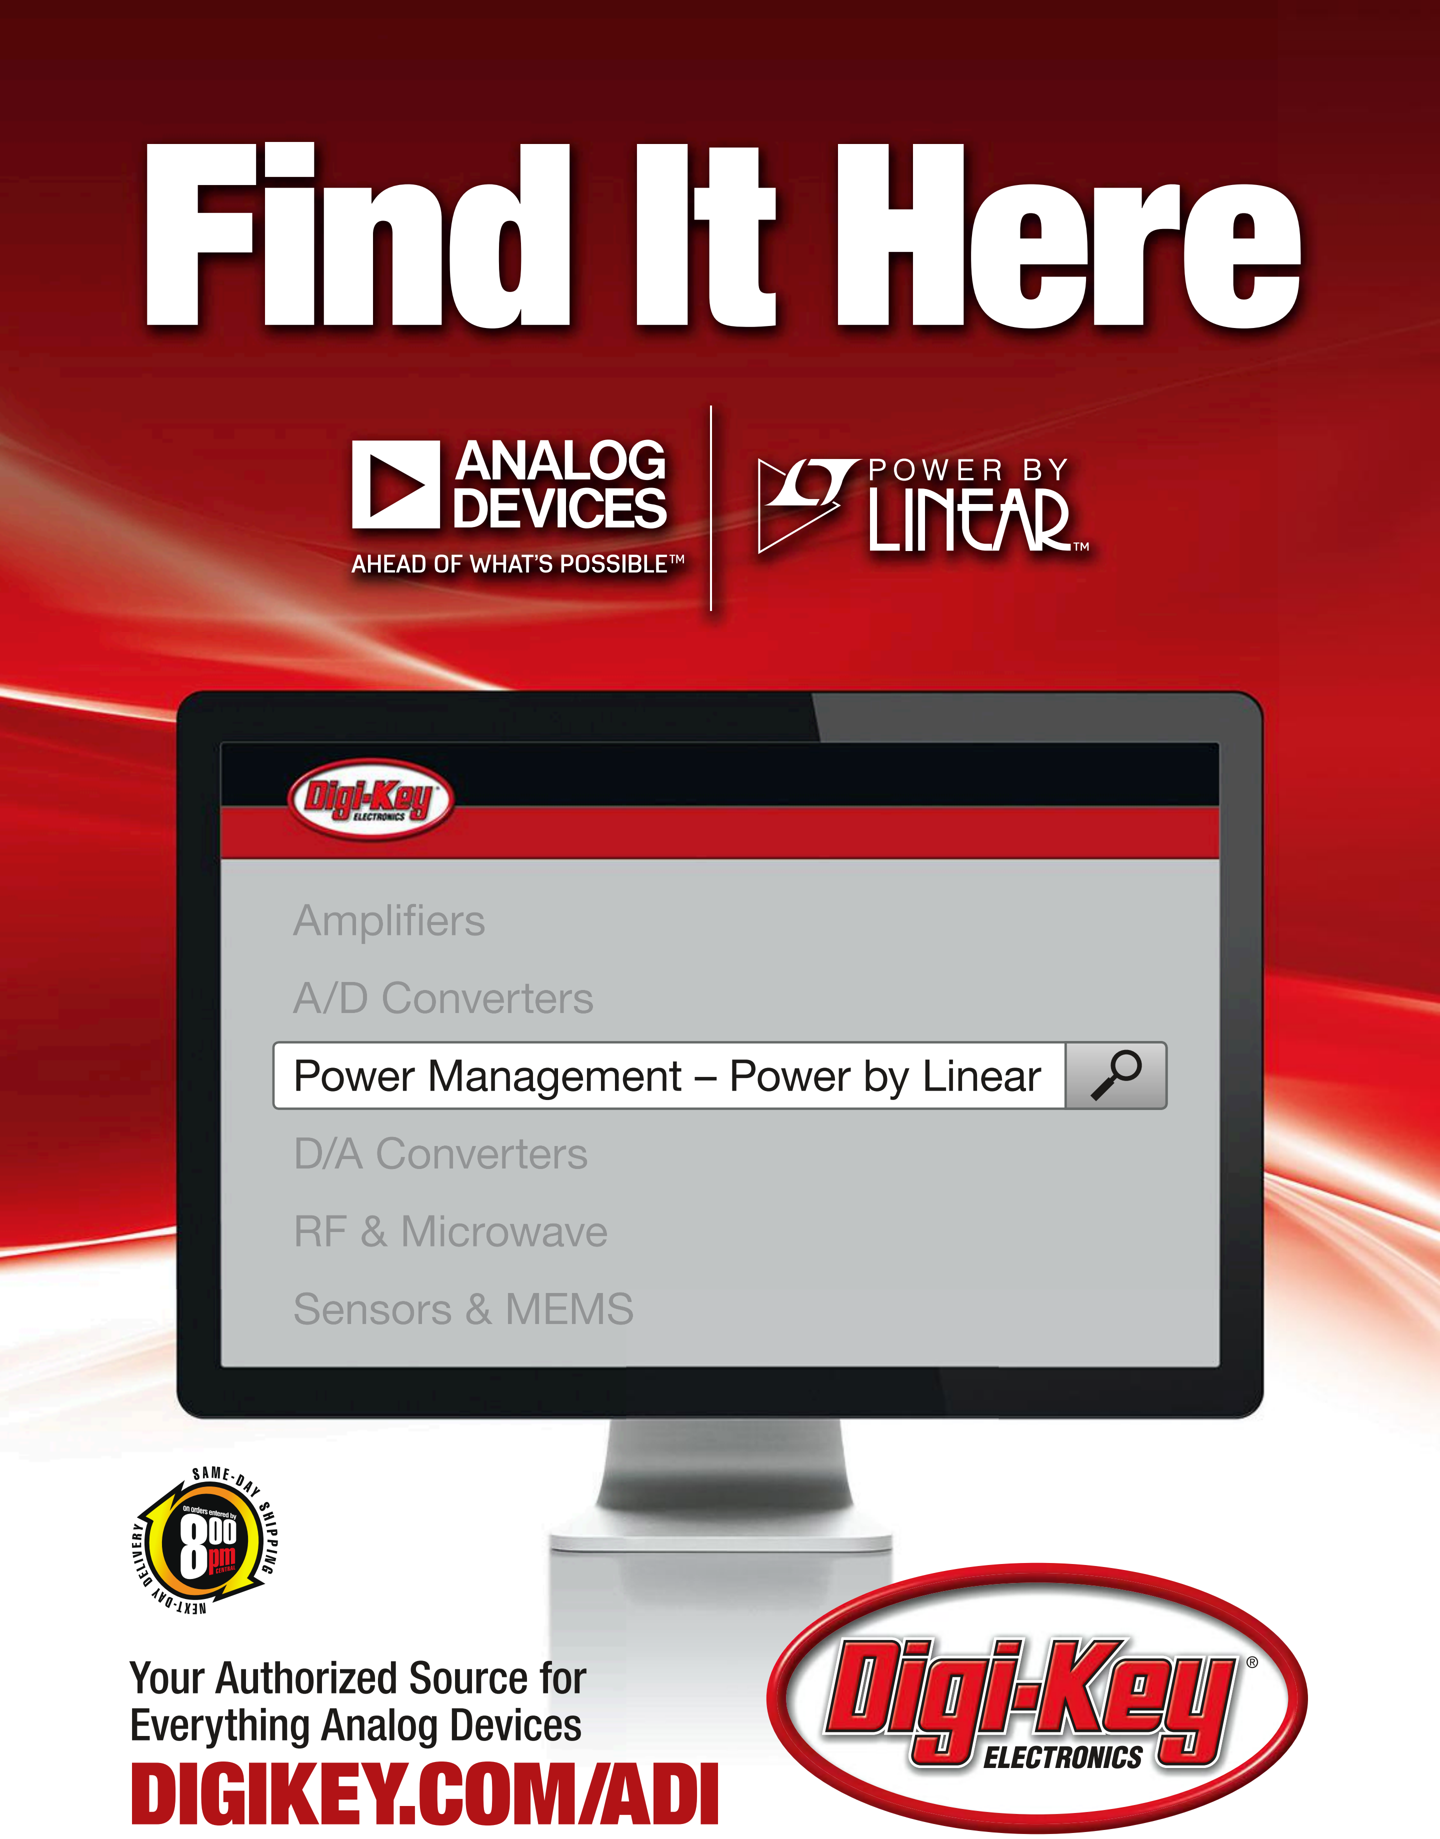
Find It Here
ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™
POWER BY
LINEAR™
Digi-Key ELECTRONICS
Amplifiers
A/D Converters
Power Management – Power by Linear
D/A Converters
RF & Microwave
Sensors & MEMS
SAME-DAY SHIPPING
on orders entered by
8:00 pm
NEXT-DAY DELIVERY
Your Authorized Source for
Everything Analog Devices
DIGIKEY.COM/ADI
Digi-Key®
ELECTRONICS

NOVEMBER/DECEMBER 2018 VOLUME 66, ISSUE 6

ElectronicDesign®

# IN THIS ISSUE

## FEATURES



14

48

52

## 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


## What's All This HP Museum Stuff, Anyhow?

No matter how old, Keysight test equipment is a treasure worth preserving. And so are the memories of those who crossed paths with legends in the industry. Paul Rako takes us on a tour.

*https://www.electronicdesign.com/test-measurement/what-s-all-hp-museum-stuff-anyhow*

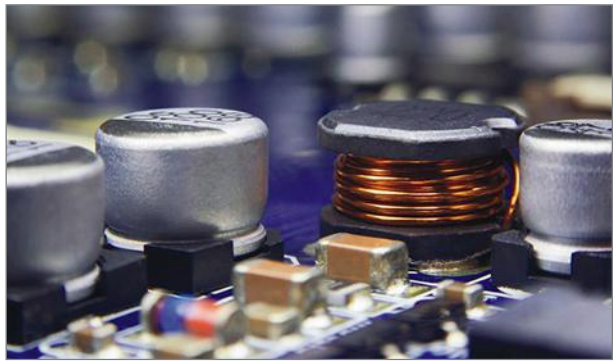

## Designing a Discontinuous-Conduction-Mode Flyback Transformer

CET Technology's Craig Lombard presents a step-by-step approach on how to design a compact DCM flyback transformer that can meet your needs.

*https://www.electronicdesign.com/power/designing-discontinuous-conduction-mode-flyback-transformer*

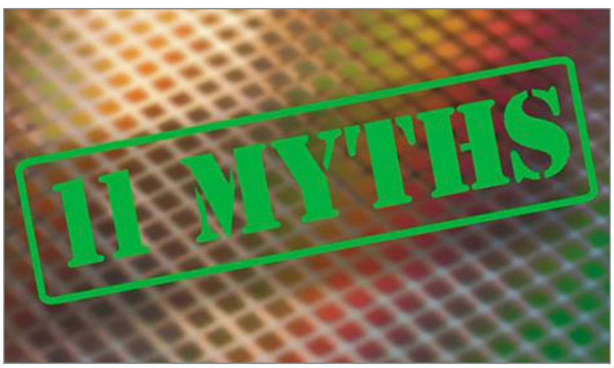


## 11 Myths About Custom Silicon

Preconceptions about custom chip design are that it's costly and takes up too much time. But what's the real story in today's semiconductor industry climate? S3 Semiconductors' Edel Griffith breaks it down.

*https://www.electronicdesign.com/analog/11-myths-about-custom-silicon*

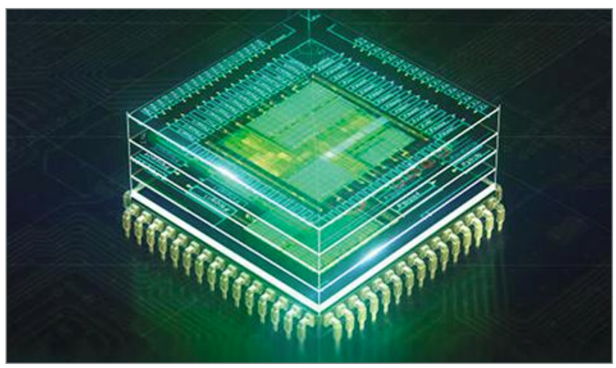

## Path to Systems: Why a SiP?

Second in a series, this article presents key advantages and challenges ahead for system-in-package (SiP) technology in the grand scheme of semiconductor integration and specifically, embedded systems.

*https://www.electronicdesign.com/embedded-revolution/path-systems-why-sip*

# From Concept to Reality

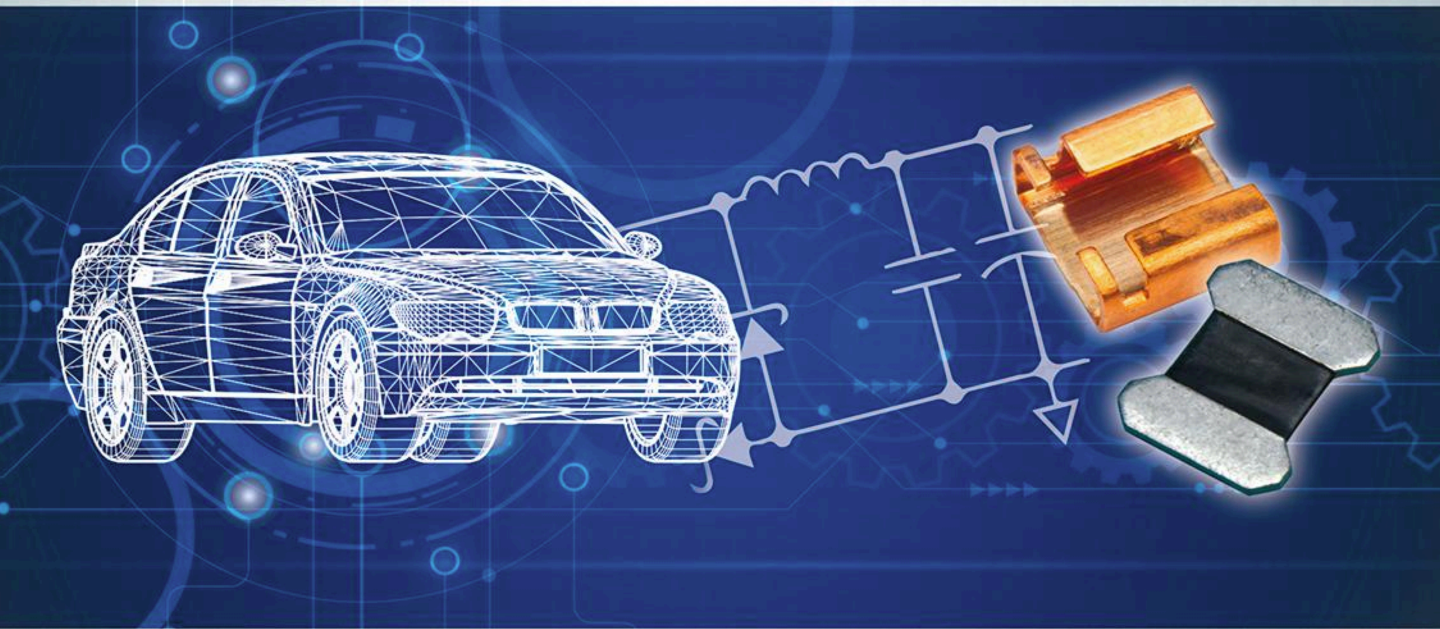

## Our new passive component solutions will help you make the leap!

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

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

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

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

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

### NEW! From KOA Speer


**Precision Thick Film Resistors (RK73G)**

**Metal Plate Current Sense Resistors (TLR, TLRZ & TLRH)**


**Wide Terminal Current Sense Resistors (WK73S & WU73)**

KOA®
KOA SPEER ELECTRONICS, INC.

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

ElectronicDesign.

**NOVEMBER/DECEMBER 2018**

## EDITORIAL

EXECUTIVE DIRECTOR, CONTENT: **KAREN FIELD** karen.field@informa.com

ASSOCIATE EDITOR/COMMUNITY MANAGER: **ROGER ENGELKE** roger.engelke@informa.com

ASSOCIATE EDITOR/COMMUNITY MANAGER: **JEREMY COHEN** jeremy.cohen@informa.com

EMBEDDED/SYSTEMS/SOFTWARE: **WILLIAM WONG** bill.wong@informa.com

ANALOG/POWER: **MARIA GUERRA** maria.guerra@informa.com

ASSOCIATE CONTENT PRODUCER: **JAMES MORRA** james.morra@informa.com

CONTRIBUTING EDITOR: **LOUIS E. FRENZEL** lou.frenzel@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: **JULIE GILPIN** julie.gilpin@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.com

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

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

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

PLEASE SEND INSERTION ORDERS TO: orders@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

informa

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


BatteryHolders.com
WISHING YOU ALL A HAPPY HOLIDAY!

Editorial

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

# Memory… It's Complicated

One constant throughout the CPU multicore battles, the rise of GPGPUs, and FPGAs in the cloud is memory. It remains a key element of every system and is often found in many forms, from cache to DRAM to NAND flash. The latter garners the most attention these days at events like the Flash Memory Summit (FMS), where the latest storage technologies were out in force.

NAND flash continues to dominate the non-volatile (NV) mass-storage market in light of dramatically dropping prices for multi-layer QLC NAND. New interfaces like M.2, U.2, and SD Express are bringing higher throughput interfaces using PCI Express and NVMe, which are pushing SATA and SAS out of its entrenched areas close to the processor. SATA and SAS are managing ever-growing arrays, though.

One theme at FMS that stood out this year was persistent memory and its associated hardware, NVDIMMs. These DIMMs are on the same bus as DRAM, but with non-volatile memory that can range from NAND flash to Intel's Optane. Putting this memory next to a processor increases throughput and reduces latency compared to even NVMe storage.

New memory technologies like MRAM and RRAM are playing in this space. Embedded MRAM is also getting more play, with the major foundries now supporting it. Many ASIC developers are using it, too, but don't ask who. It tends to be the secret sauce that no one wants to talk about.

There are multiple challenges for these emerging memory technologies, with differences in price, availability, and capacity being major considerations when competing with NAND flash and DRAM. Objective Analysis' Jim Handy notes, "Spot prices, a leading indicator of where contract prices are likely to head, have been on a steady decline through all of 2018. Objective Analysis expects this trend to continue with NAND flash, then migrate to DRAM, until prices reach cost, resulting in a profitless 2019 for memory chips. This will slow the migration to the new memory technologies profiled in our new report 'Emerging Memories Poised to Explode.'"

Memory appetite among developers is insatiable. Likewise, memory hierarchies are becoming even more complex. For example, the memory bus may now have different DIMMs installed. Hard drives could be part of the mix, but they're way out on the periphery where capacity and price trumps performance.

Still, memory technology is one area that continues to push the limits with long-term growth almost assured, especially when compared to CPUs and even GPUs that are frequency-challenged. Increasing and improving cores is great, but memory technologies lend themselves to this type of scalability as well.

As Bob Hope used to say, "Thanks for the memory." ed

**ED. NOTE:** In our Sept./Oct. issue of *Electronic Design*, several technical errors appeared in the "How Does Power Factor Correction Impact Your Utility Bill?" article. We apologize for those errors and any confusion they may have caused. Please go to *www.electronicdesign.com* for the corrected version of the article.

# News

## SiTime Targets the High End of TIMING-DEVICE MARKET

To limit interference and support higher throughput, 5G networks need a different architecture than 4G communications systems. Instead of being located in remote cell towers, radios will be mounted on street lamps, traffic lights, rooftops and parking garages closer to smartphones and other connected devices. But to establish connections with each other, 5G radios needs to have precising timing—around 10 times more precise than 4G equipment.

That not only means boosting the precision of timing devices in wireless infrastructure. It also means protecting their signals from vibrations, humidity, sudden temperature changes, and other disturbances that can throw off frequency stability, according to Piyush Sevalia, vice president of marketing for Santa Clara, Calif.-based SiTime. The company is trying to check off both requirements to expand its share of the $1.5 billion networking and telecommunications timing market.

The company recently announced its first line of oven-controlled oscillators—more commonly called OCXOs—which have the lowest noise, highest precision, and most stable performance over sudden changes in temperature of any timing component, making them critical to the operation of communications networks. Traditional OCXOs are used to keep time in the core network but could move into radios as timing requirements increase.

The Emerald timing devices are based on microelectromechanical systems (MEMS) instead of the tiny crystals traditionally used to generate timing signals that enable the operation of chips and circuit boards in almost every electronic device. That results in components with 10 times the stability over temperature than quartz OXCOs, according to SiTime. Sevalia said that they are also smaller and consume less power than alternatives.

"These components have been the realm of crystal oscillators for so long," said Sevalia. Traditional OCXOs contain a sliver of transparent quartz that vibrates at a specific frequency when run through with electricity. The crystal is surrounded by a temperature-controlled oven. Limiting variations in ambient temperature serves to prevent unwanted changes in frequency. They generally range in cost from around $10 to $100 each, according to SiTime.

But with the need for more accurate timing, 5G radio equipment could benefit from the use of OCXOs. The prob-

lem is that being mounted to street lamp poles or traffic lights can introduce the OCXOs to vibrations that throw off the resonance of the crystal, altering their frequency. Because SiTime's MEMS components weigh only 1/3000th of crystal resonators, less energy couples to them, resulting in higher vibration resistance, said Sevalia.

"It was just a matter of time before we got here, and now here we are," he said. In addition to alternatives for oven-controlled crystal oscillators, SiTime is also targeting customers that typically use voltage-controlled crystal oscillators (VCXOs) that output different frequencies based on the input voltage, temperature-compensated crystal oscillators (TCXOs) that can better tolerate heat fluctuations, and standard crystal oscillators (XOs).

Business has been booming. SiTime's products are stealing market share in factories, cars, and other systems subject to extreme temperatures, violent vibrations and other stressors, even though they generally cost more than crystal devices. Last year, under chief executive Rajesh Vashist, the company increased unit shipments around 85% to 425 million. That gave the company more than a billion units shipped since it was founded in 2005.

SiTime's Emerald product line is based on a programmable analog architecture, allowing the company to program the timing devices to any output frequency between 1 and 220 MHz based on customer demands. The devices support temperature operation from −40 to 85°C. Before the end of the first quarter of 2019, SiTime plans to give customers the ability to program the devices themselves and broaden the temperature range to 40 to +105°C.

Since SiTime's Emerald products are less vulnerable to temperature and vibration, customers have more freedom in the layout of circuit boards and can avoid using additional shielding for thermal isolation. The timing devices are also smaller than alternatives, so the company is planning to sell products that match the dimensions of common OCXOs. That way, telecommunications companies can use them like drop-in replacements, said Sevalia.

"With 5G, every radio may need an OCXO," he told *Electronic Design*. Potential customers could include Intel, which has announced that it would combine SiTime's timing devices with its latest wireless modems for 5G infrastructure. Last month, Intel said that both Ericsson and Nokia were planning to build radio access equipment based on its multimode wireless modems for global 5G deployments. ■

## **POWER-EFFICIENT NOVUTENSOR** Takes on Inference at the Edge

**NOVUMIND IS ANOTHER** startup with its eyes on inference at the edge, leveraging dedicated technology designed to deliver high inference throughput with minimal power. It's not alone in this space and many GPU and FPGA machine-learning (ML) platforms are available, but fewer dedicated platforms are shipping at this point. Many of these target specific applications, such as Intel's Movius series. The advantage of the newer chips is that their power requirements are a few watts instead of the hundreds needed for high-end GPU boards.

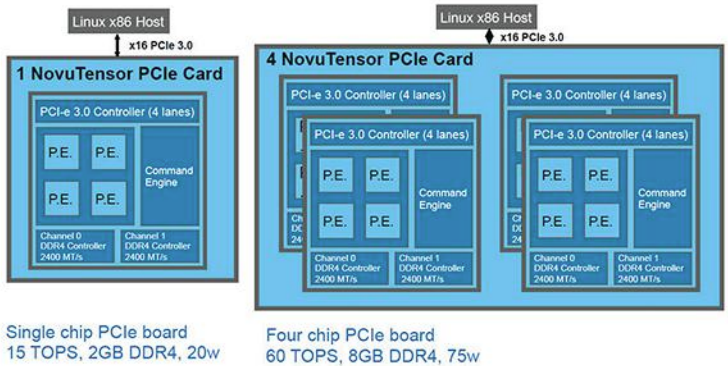


**1. The 400-MHz NovuTensor delivers 15 TOPS while the quad-chip PCI Express card pushes 60 TOPS.**

Dr. Ren Wu, founder and CEO of NovuMind, says, "Until now, GPUs have powered the advances in AI, particularly around the training of deep-neural-network models from large sets of data. Once models are trained, however, the challenge is to deploy them at scale. GPUs and other processors are expensive and consume large amounts of power. Their architectures are optimized for two-dimensional matrix computation. While they perform well when processing large batches of data, these chips are not suited for real-time applications that require low latency. They also lack power efficiency and they tend to be very expensive. With the arrival of our NovuTensor chip, we are breaking these barriers and ushering in a new era where AI can be deployed at scale."

NovuMind's 400-MHz NovuTensor is designed to deliver 15 TOPS while using under 5 W for the neural engine. The chip actually needs 15 W. It's available as a chip or on a short PCI Express card *(Fig. 1).* The card includes four chips delivering a combined 60 TOPS.

*(Continued on page 29)*

WILLIAM WONG | Senior Technology Editor

# Entering a New Robotics Age with Machine Learning

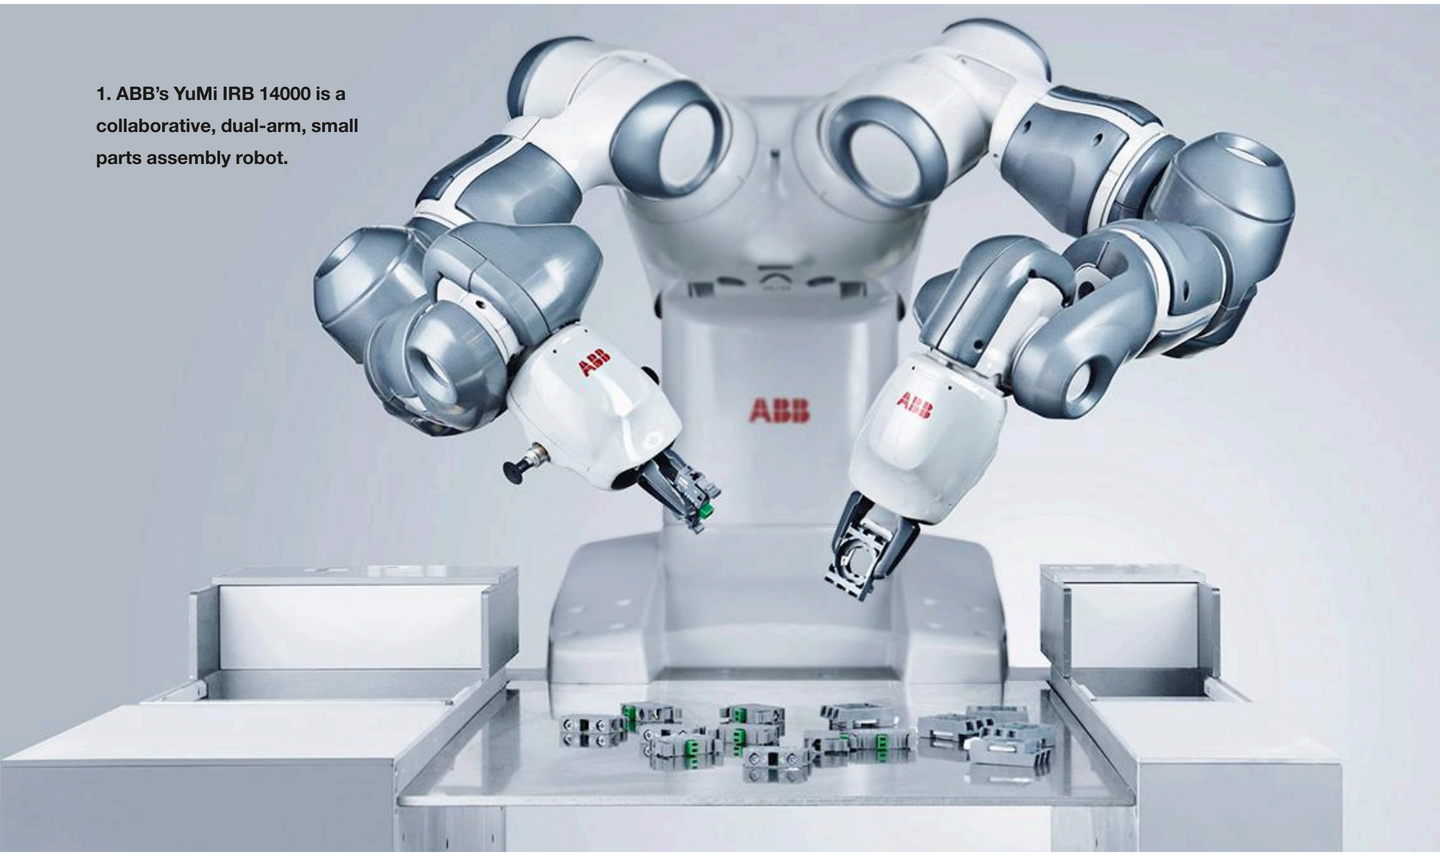


1. ABB's YuMi IRB 14000 is a collaborative, dual-arm, small parts assembly robot.

**From the flood of sensors being incorporated into robots to the rise of neural-net inference, robotics technology is experiencing a seismic shift in performance.**

Machine-learning (ML) technology is radically changing how robots work and dramatically extending their capabilities. The latest crop of ML technologies is still in its infancy, but it looks like we're at the end of the beginning with respect to robots. Much more looms on the horizon.

ML is just one aspect of improved robotics. Robotics has demanding computational requirements, and that's being helped by improvements in multicore processing power. Changes in sensor technology and even motor control have made an impact, too. Robots are being ringed with sensors to provide a robot with more context of its environment.

Thanks to all of these improvements, robots can tackle new tasks, work more closely with people, and become practical in a wider range of environments. For example, ABB's YuMi "you and me" IRB 14000 robot *(Fig. 1)* is a collaborative, dual-arm, small parts assembly robot. It has flexible hands and a camera-based part-location system. Cameras are embedded in the grippers.

Safety functionality built into YuMi makes it possible to remove barriers to collaboration—there's no longer a need for fencing and cages (Underwriters Laboratory ISO 13849-1 PL, or performance level, b Cat b). Its lightweight arms, designed to have no pinch points, have a floating plastic casing that can also be padded. The software can detect collisions and stop movement within milliseconds. All of the wiring for the self-contained system is internal.

Locus Robotics' LocusBot *(Fig. 2)* is another example of a mobile robot designed to work collaboratively with humans. The wheeled robot can move autonomously through a warehouse. Its tablet interface runs the LocusApp that allows users to interact with the robot. The tablet has a built-in scanner to easily identify an item's bar code. Robots are linked wirelessly to the LocusServer.

The LocusBot weighs in at 100 pounds, has a payload up to 100 pounds, and can move at speeds up to 1.4 m/s. The battery charges in less than an hour, providing a 14-hour runtime.

**2. Locus Robotics' LocusBot is designed to work collaboratively with humans.**

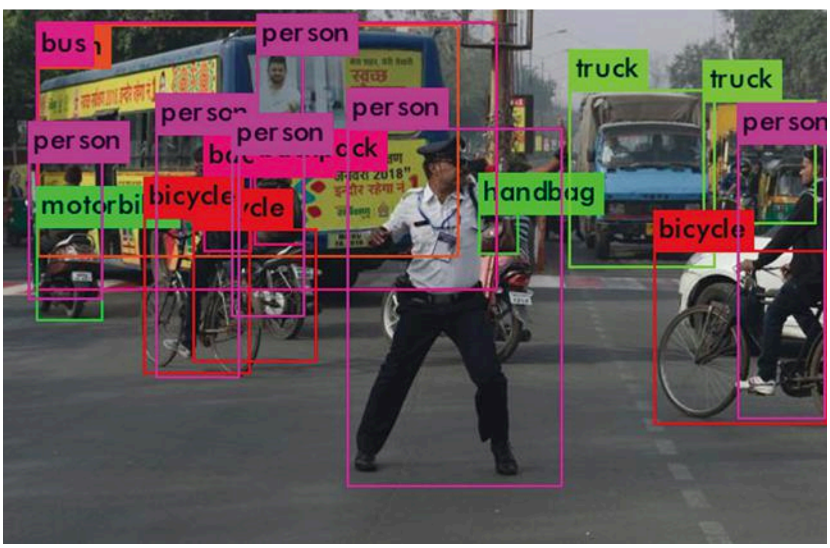


**3. Machine learning is allowing robots to track many different types of objects using cameras.**

**MORE SENSORS**

Robots are using more sensors and different types of sensors to evaluate their environment. The YuMi includes motor feedback and cameras in addition to movement sensors. The LocusBot incorporates range sensors in its base.

The ability to track many different types of objects using cameras is also becoming important, not only to robots, but autonomous vehicles that need to coexist safely with people *(Fig. 3)*. Smaller, lower-cost sensors enable developers to place better sensors in more locations to provide a robot with a wider field of view than people.

High-resolution camera sensors like OmniVision Technologies' 24-Mpixel OV24B series chip *(Fig. 4)* help bring visual information to machine-learning tools. The OV23B is built on OmniVision's PureCel Plus stacked-die technology. It can deliver full resolution at 30 frames per second (fps). It can

Multiple sensors of the same type are becoming the norm, such as the four-camera systems that provide surround view in many cars. Dual cameras can provide parallel support to do 3D imaging. Another approach is to use 3D cameras like Intel's RealSense *(Fig. 5)*. These cameras actually include a 2D sensor in addition to the 3D sensor. This allows one system to provide range information as well as a video stream that can be used by machine-learning systems to identify objects.

also do 4-cell binning with a 6-Mpixel resolution at 60 fps, 4K2K video at 60 fps, 1080p video at 120 fps, or 720p video at 240 fps.

The OV24B sensor family has on-chip re-mosaic, 2x2 mirrorless phase-detection auto-focus (PDAF) support. It also offers very good low-light performance. Just what a high-end smartphone or high-performance robot needs.

Multiple sensors of the same type are becoming the norm, such as the four-camera systems that provide surround view in many cars. Dual cameras can provide parallel support to do 3D imaging. Another approach is to use 3D cameras like Intel's RealSense *(Fig. 5)*. These cameras actually include a 2D sensor in addition to the 3D sensor. This allows one system to provide range information as well as a video stream that can be used by machine-learning systems to identify objects.

The RealSense depth modules use infrared, time-of-flight (ToF) sensing to provide 3D information. They also have 1080p video camera sensors. The standard versions use USB 3.0 connections; modules without the case are available separately. They work best for short- to mid-range indoor applications.

LiDAR and radar work well indoors and outdoors. The need for such sensors in cars and autonomous vehicles has driven down their cost and led to enhancements in performance. Autonomous vehicles are essentially robots, so one might say these sensors are being designed for robotic applications. In any case, they're equally applicable to other robotic applications as well. Likewise, they're getting smaller.

Texas Instruments' mmWave sensors *(Fig. 6)* are available for short- and long-range operation. These radar systems aren't affected by lighting and other optical issues. The short-range

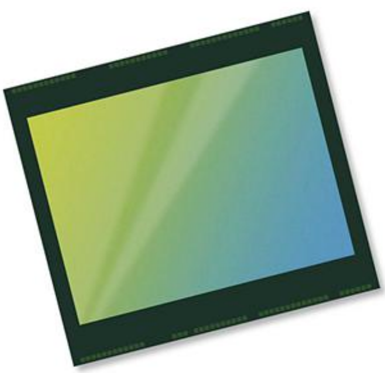

**4. OmniVision's OV24B 24-Mpixel camera sensor brings visual information to machine-learning tools.**

**5. Intel's RealSense cameras provide 3D position information in addition to conventional 2D image streams.**

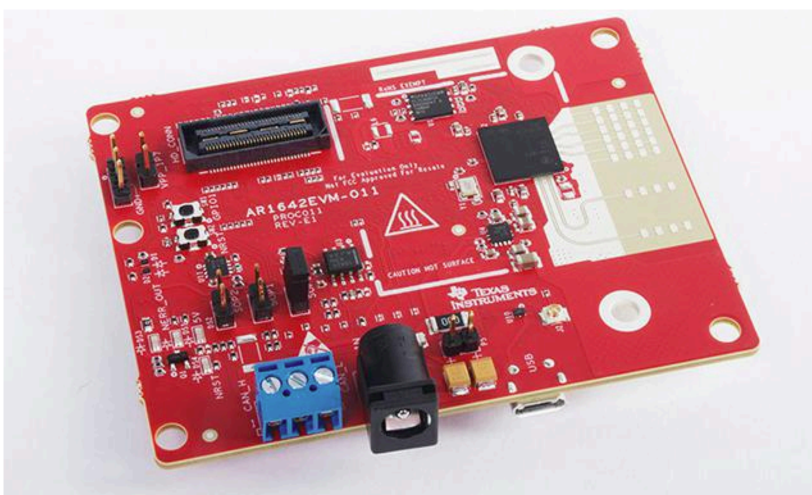

**6. The AWR1642 EVM module is Texas Instruments' mmWave radar reference design. The tiny array on the right is the radar antenna matrix.**

versions are useful for applications like mobile robots or even detecting when people are in a room. They're superior to many motion detectors that only indicate the presence or distance to an object rather than movement direction. This ability is useful for a self-opening door in knowing whether a person is coming toward the door or passing in front of it. Robots can also use this type of information.

### SEEING MACHINE-LEARNING IMPROVEMENTS

One area of research that's having a major impact on robotics is convolutional-neural-network (CNN) and deep-neural-network (DNN) inference. It allows an application to utilize inputs from the new cameras and 3D sensors to identify objects. These tasks are computationally heavy, and there's lots of data involved in both training and inference. For most robots these days, ML's main usage will involve inference.

Companies like Habana with its Goya chip and Flex Logix with its NMAX architecture *(Fig. 7)* are working to reduce the cost and power requirements of using high-performance ML inferencing. The Goya HL-1000 processor can process 15,000 images/s with a 1.3-ms latency, while dissipating only 100 W of power, for the ResNet-50 with a small batch size.

Small batch sizes are important for inferencing—changing the weights and activation information that's necessary for large ML models is time-consuming. Making this change quickly and keeping this information local will be the key to success for ML chips. The NMAX architecture has L1 RAM close to the computing elements plus an L2 memory between that and the off-chip DRAM. A high-speed network fabric links everything together.

FPGA companies are also tailoring their solutions for machine learning. Xilinx's Adaptive Compute Acceleration Platform (ACAP) reworks the FPGA into a more customized compute engine surrounded by accelerators and mated with real-time and application processors. ACAP also has multiple high-speed interconnects linking various components within the system, including neural-network accelerators.

DSPs and GPUs also continue to crunch ML inference chores. For instance, NVIDIA's Jetson Xavier can deliver 32 TOPS of performance *(Fig. 8)*. This latest Jetson shares the same basic architecture as the NVIDIA T4 that's being used for training as well as cloud inference tasks. The Xavier module packs in 512 of the latest NVIDIA cores. It also has its own 64-bit ARMv8 processor complex, a 7-way VLIW vision accelerator, and a pair of NVIDIA Deep Learning Accelerators (NVDLAs), in addition to 16 CSI-2 video input lanes plus eight SLVS-EC lanes. It can deliver 20 INT8 TOPS using only 30 W.

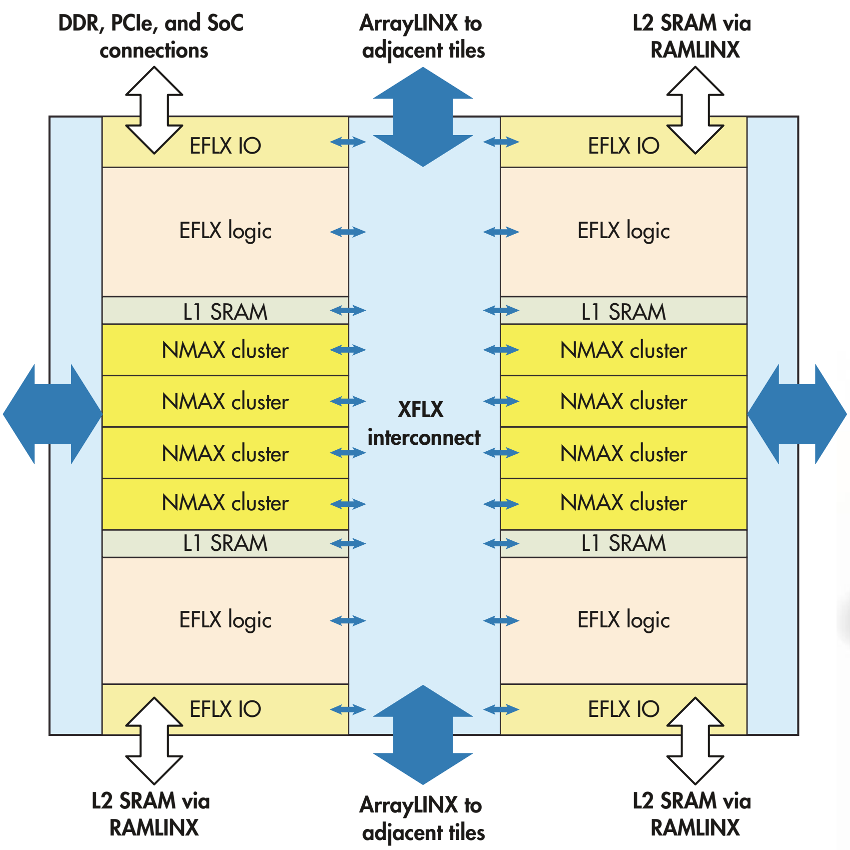


**7. Flex Logix's NMAX is one machine-learning hardware platform designed to allow robots to analyze sensor information from a variety of source.**

**8. NVIDIA's Jetson Xavier can deliver 32 TOPS of performance for machine-learning chores.**

Another emerging technology worth mentioning is CCIX (pronounced "see 6"). The Cache Coherent Interconnect for Accelerators is built on PCI Express hardware with its own low-latency, cache-coherent protocol designed to link processors with ML hardware accelerators as well as other accelerators. CCIX is so new that nothing except FPGAs can implement it, but that will be changing over the next couple years.

**MACHINE LEARNING FROM EXPERIENCE AND EXPERIMENTATION**

Using inference engines trained in the cloud or on local servers is where it's at now, but ML research continues. On that front, self-learning is currently a hot topic, and yet another change in how robots will be built and programmed. Such research is often based in neural networks as well, though it looks to be a way to address some of the harder robotic challenges, from handling fabrics to selecting and picking up objects—much along the lines of what Google Research has been doing *(Fig. 9)*.

Google Research turned a hoard of robotic arms loose on a tray of unsuspecting gadgets, foam doohickeys, and Legos to see if they could figure out what to do with them. The 3D cameras were overlooking the tray and arm. It's only one of hundreds of research projects around the world that have taken off because of the availability of hardware, sensors, and ML software.

Still, physical devices are expensive and take up space. They can also be slow. On the other hand, simulation doesn't take up space and scales to available

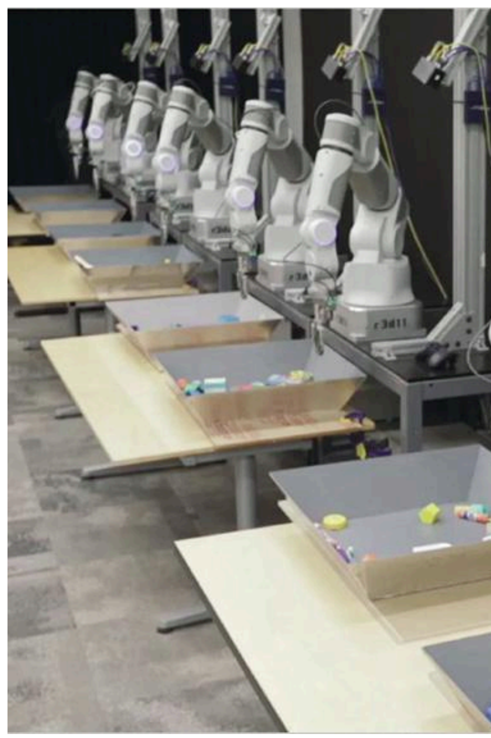

**9. Google let these robotic arms learn to pick up objects.**

computing resources. The challenge is the supporting framework. This is where NVIDIA's Isaac robot simulation environment comes into play. Isaac is built around an enhanced version of Epic Games' Unreal Engine 4 and the Physx physics engine.

Isaac is tuned for cobots and conventional robots, but simulation is also playing a big part in that other major robotic application—self-driving cars. These typically simulate a much larger environment, with many objects requiring heavy-duty computing and storage environments typically found in cloud computing.

What's changing the way robotics works are all of these technological advances. Individually they're useful, but combining them has radically changed how robots are designed, trained, and employed. We're just at the starting point and that's rather exciting. However, we should not forget issues such as security and safety in the rush to build the next great robot. ed

**Physical devices are expensive and take up space. They can also be slow. On the other hand, simulation doesn't take up space and scales to available computing resources. The challenge is the supporting framework.**

BOB WARREN | Memory and Storage Architect, Spin Transfer Technologies
www.spintransfer.com

# MRAM Means New Possibilities for NVDIMMs

Replacing DRAM with MRAM in new NVDIMM designs will result in higher performance, lower power, easier management, and new ways of using NVDIMM.

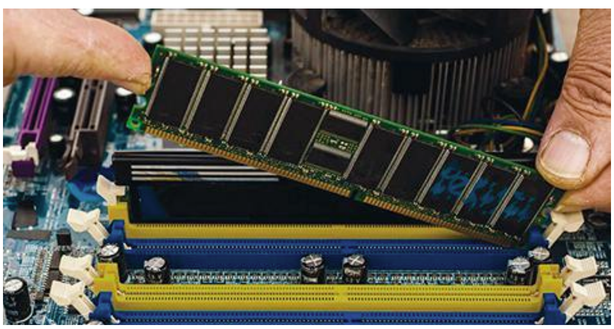

NVDIMM (non-volatile dual in-line memory module) is a new class of memory designed to allow non-volatile memory to be connected to the DIMM bus traditionally populated by DRAM. Being a new memory, we're early into its lifecycle, with two variants in production and a third on the way. But proponents of NVDIMM are also in the early days of understanding the diverse ways that they can integrate the technology into their systems. As user ideas proliferate, they will spawn new NVDIMM architectures, further broadening NVDIMM's application scope.

The newest NVDIMM combines flash memory for storage with a DRAM cache for faster access. While DRAM clearly has extremely broad adoption in computer systems, it's fraught with management and power challenges. An emerging memory technology, magnetoresistive RAM (MRAM), is simpler

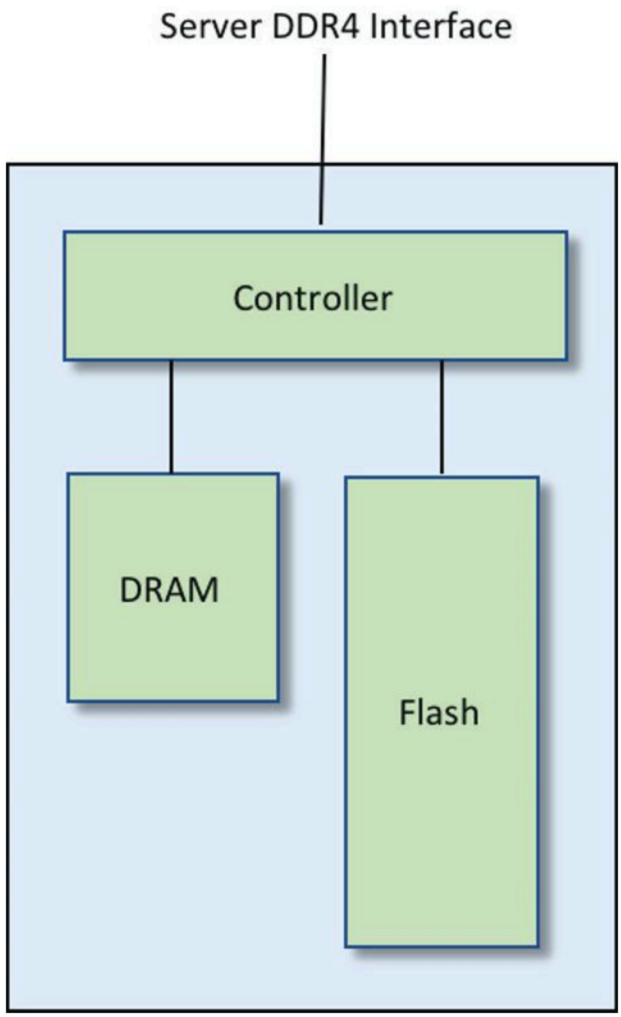


**1. A typical NVDIMM-N configuration combines flash and DRAM managed by a controller that moves data between the two during power transitions.**

to manage than DRAM, is targeted to use less energy, and is itself non-volatile. Replacing the DRAM in an NVDIMM with MRAM in new NVDIMM designs will result in higher performance, lower power, easier management, and new ways of using NVDIMM. It could also inspire new architectures for future versions.

**OVERVIEWS: NVDIMM AND MRAM**

There are two existing versions of NVDIMM (with one more in the standardization process):

- NVDIMM-F is the original, and it consists only of flash memory wrapped in a DDR4 interface.
- NVDIMM-N adds the DRAM cache, keeping the DDR4 interface. Reads and writes go through the DRAM, with cache changes being written back to flash per system policy. Because DRAM is volatile, its contents must be stored in the flash when power is shut off or fails. Some form of backup battery is needed to provide the energy for that transfer.

NVDIMM-P is a newer standard being formalized through the JEDEC organization. Intended to accommodate new types of memory, its first configuration will be monolithic, with subsequent variant versions to follow *(Fig. 1)*.

MRAM, meanwhile, is a newer non-volatile memory technology that uses very small magnetic elements as the storage cell. Unlike flash, they're byte-addressable and require no special management like wear-leveling and refresh. And, as companies like Spin Transfer Technologies innovate to bring this technology into the broad commercial sphere, MRAM will be able to compete with DRAM on cost.

**HANDLING POWER FAILURES**

One of the major challenges of NVDIMM-N is the need to save the contents of the DRAM when the power fails or is shut off. Under normal operating conditions, the cache must be saved when the power is turned off, and then, depending on how that save operation is done, the DRAM will need to be reloaded on power up to restore the state of the system when power was shut off.

Powering down a system is part of normal operation and can be planned for. But power failures come unpredictably, and multiple power on-off cycles may occur before full power is ultimately restored. This is typically handled by detecting the failure and then saving the cache to flash. During that save

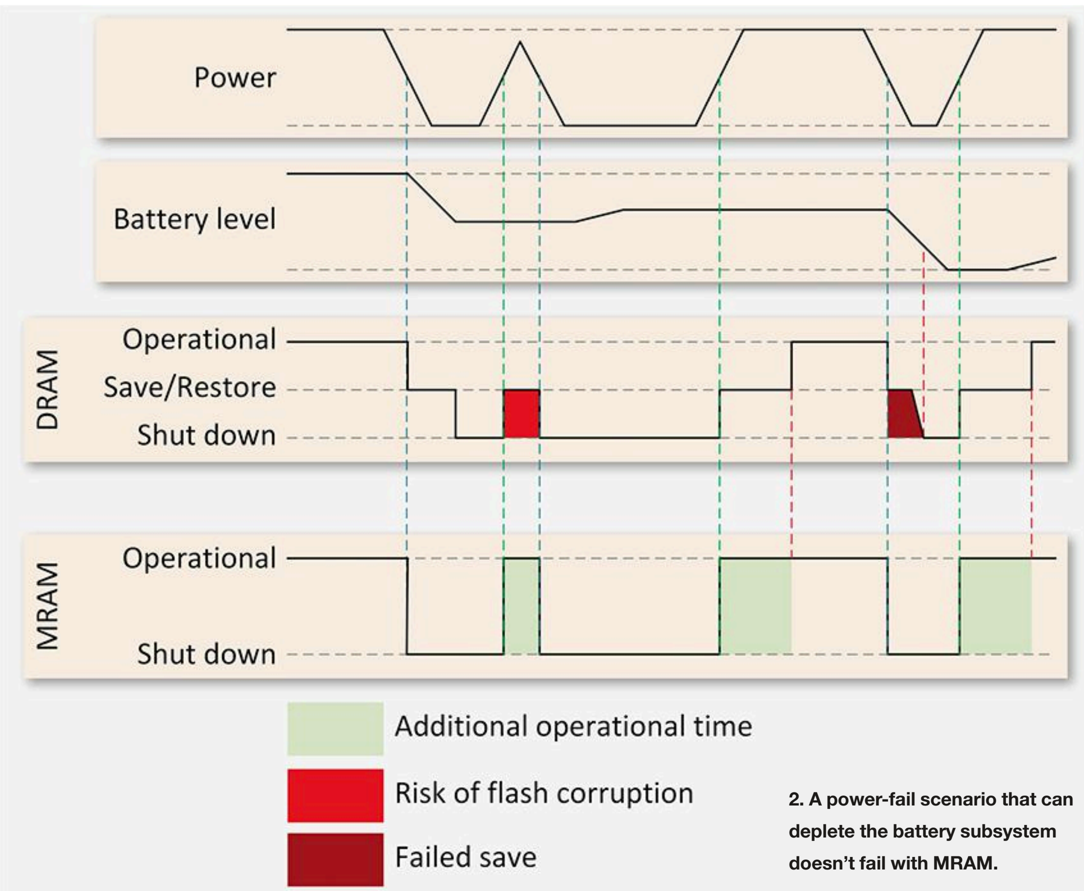


**2. A power-fail scenario that can deplete the battery subsystem doesn't fail with MRAM.**

## Transient Immunity Testers from AVTECH

**The Avtech AVRQ series of high-voltage, high-speed pulse generators is ideal for testing the common-mode transient immunity (CMTI) of next-generation optocouplers, isolated gate drivers, and other semiconductors.**

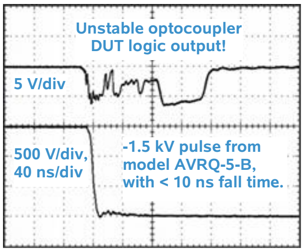


- **Kilovolt amplitudes (±1, ±1.5, -2 kV)**
- **Transition times down to 10 ns, dV/dt rates up to 120 kV/us**
- **Switchable daughterboards to handle a variety of DUT package styles**
- **GPIB, RS-232 ports standard**
- **Ethernet / VXI optional**

AVTECH

Nanosecond Electronics
Since 1975

**Pricing, manuals, datasheets and test results at:**
**http://www.avtechpulse.com/semiconductor**

## 30, 50 and 100 Volt Lab Pulsers

**Avtech offers an extensive series of user-friendly 30, 50 & 100 Volt general-purpose lab pulsers. We can provide an alternative for the discontinued Agilent 8114A or HP214!**

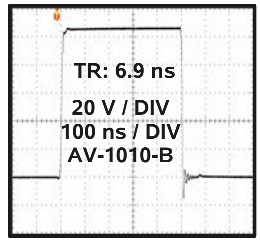


| | |
|---|---|
| **Model AV-1015-B:** | **50 Volts, 10 MHz** |
| **Model AV-1010-B:** | **100 Volts, 1 MHz, 25 ns to 10 ms, 10 ns rise time** |
| **Model AV-1011B1-B:** | **100 Volts, 2 ns rise time** |
| **Model AV-1011B3-B:** | **30 Volts, 0.5 ns rise time** |

**Variable baseline and burst mode options and output currents to 8 Amps with accessory transformers.**

AVTECH

Nanosecond Electronics
Since 1975

**Pricing, manuals, datasheets and test results at:**
**http://www.avtechpulse.com/general**

## Nanosecond Laser Diode Drivers With Butterfly Diode Sockets

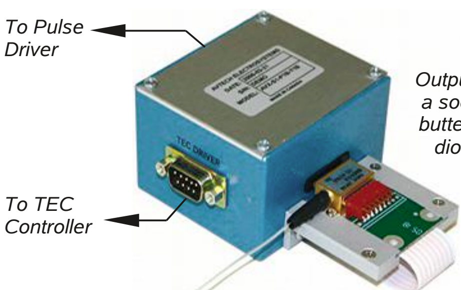


*Output module with a socket-mounted butterfly-packaged diode installed.*

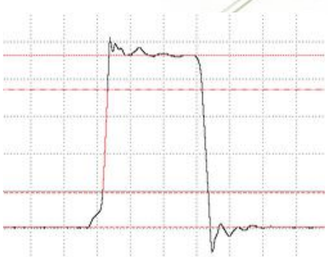

**Model AVO-9A-B**
**40 mA/DIV**
**1 ns/DIV**

Each of the 19 models in the Avtech AVO-9 series of pulsed laser diode drivers includes a replaceable output module with an ultra-high-speed socket suitable for use with sub-nanosecond rise time pulses. Models with maximum currents of 0.1A to 10A are available with pulse widths from 400 ps to 1 us. GPIB, RS-232, and Ethernet control available.

AVTECH

Nanosecond Electronics
Since 1975

**Pricing, manuals, datasheets and test results at:**
**http://www.avtechpulse.com/laser**

## 2 to 200 Amp Current Pulsers

Avtech has pioneered the design of user-friendly 2 to 200 Amp constant current pulsers. We offer over 35 models in five series for laser diode, air bag, squib, fuse and other applications. The basic specifications are as follows:

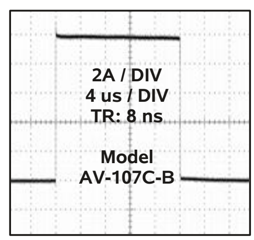


| Series | I, V | PW | TR |
|---|---|---|---|
| AV-107 | 2 - 20 A, 60 V | 0.2 - 200 us | 10 - 30 ns |
| AV-106 | 5 - 100 A, 100 V | 0.5 us - 1 ms | 50 ns - 1 us |
| AV-108 | 12.5 - 200 A, 100V | 2 us - 1 ms | 5 - 15 us |
| AV-109 | 10 - 100 A, 5 V | 10 us - 1 s | 10 us |
| AV-156 | 2 - 30 A, 30 V | 1 us -100 ms | 0.2 - 50 us |

Avtech has a long history of producing one-of-a-kind custom units.

AVTECH

Nanosecond Electronics
Since 1975

**Pricing, manuals, datasheets and test results at:**
**http://www.avtechpulse.com/current**

operation—which happens *after* system power is gone—a backup battery supplies the necessary energy. This adds some system-level requirements:

- Software is needed to manage the power transitions (including normal shutdown).
- Space in the flash may need to be allocated for cache storage.
- The battery itself takes precious space, has a limited lifetime, and introduces materials into the system that are covered by the Restriction of Hazardous Substances (RoHS) standard.
- If a primary battery is used, then replacement is needed after some predictable time. If the battery is rechargeable, then recharging circuitry is required, and both need to be periodically tested.
- When capacitors or other energy storage is used, operation can become unpredictable as the components age.

Repeated power cycles create yet other challenges. If the power comes back up, causing the cache contents to be read back out of flash, and then fails again during that load process, the flash contents can be corrupted. Worse yet, each power cycle further drains the backup battery, with no time for complete recharging if the battery is rechargeable. After too many cycles, the battery will have been exhausted.

Because of this, some NVDIMMs include a power-cycling session in their testing or validation. Power is turned off and on every few milliseconds for a couple of hours to ensure that the system can handle sustained power-cycling events.

### MRAM AS CACHE

Replacing the DRAM cache with an MRAM cache eliminates the challenges associated with power failures. MRAM's non-volatile character means that, if power goes off, MRAM maintains its contents. There's no need to save the cache contents to flash when power goes away, nor does the cache need to be reloaded when power returns, making power-up faster and safe.

This means that systems using MRAM in NVDIMMs will not require the backup battery to support the MRAM. And by not having to reload the cache contents on power-up, there's no risk of flash corruption on a subsequent failure. Finally, the software stack no longer needs code for managing the cache-saving process, and the byte-addressability of the MRAM further simplifies software *(Fig. 2)*.

Note that if system design specifications call for writing the cache to flash on power-down for any reason, it's still possible

> Replacing the DRAM cache with an MRAM cache eliminates the challenges associated with power failures. MRAM's non-volatile character means that, if power goes off, MRAM maintains its contents. There's no need to save the cache contents to flash when power goes away, nor does the cache need to be reloaded when power returns, making power-up faster and safe.

to do so even with MRAM. However, it's no longer required as a fundamental characteristic of NVDIMM-N use.

### NVDIMMs FOR SOFTWARE-DEFINED STORAGE

Software-defined storage adds a layer of abstraction above traditional memory-mapping schemes. Instead of defining a memory region by its physical address boundaries, a region is given a name; the region is referred to as a namespace.

In a software program, the memory region is referred to by the name instead of by memory bounds. An intervening controller manages where the namespace is allocated. In fact, for security purposes, a namespace could be relocated on each boot-up. This abstraction decouples the memory region as referred to in the program from the physical storage, allowing the storage to be changed or even interleaved without any source program changes.

These memory regions can be encrypted, which adds key management to the system requirements. The keys are typically stored in their own encrypted space. System robustness is enhanced if those keys are stored in a memory separate from the memory where the data itself is stored.

The keys are themselves encrypted, so if the key storage is compromised, the attacker will not gain access to the data. The root key for encrypting the keys comes from the system. Therefore, gaining access to the key storage—say, by removing the NVDIMM module from the system—keeps the data secure.

In such a system, the key storage can be defined as one namespace; the data as a separate namespace. Data writes involve encryption; data reads involve two levels of decryption (key and data). These details are abstracted from the system through a controller. When the system needs to read data, for example, it simply requests a read. The controller then recovers and decrypts the required data key from the key storage, reads and decrypts the data, and then places it on the bus. The system is unaware of the details of how the data was obtained.

An NVDIMM with MRAM as the cache works very well for this application—particularly if a new controller allows for independent access to the MRAM and the flash (rather than requiring all flash accesses to go through the cache). The MRAM can be used as the key-store namespace; the flash becomes the data-storage namespace. Because the MRAM and flash are separate chips, the failure of one doesn't affect the other.

### MRAM USES LESS ENERGY

Using MRAM enhances energy benefits:

- DRAMs utilize a pre-charge mechanism for reading data. That pre-charging operation uses energy that's not required for reading an MRAM.
- DRAMs need to be constantly refreshed. The refresh operation is performed by dedicated circuitry that keeps

track of which memory sectors need to be refreshed at each refresh event. That circuitry uses energy, as do the memory cells, during refresh. No such circuitry is required for MRAM, so that energy is saved.

- DRAMs must be read one line at a time. When multiple threads are served out of a single memory, accesses will appear to be relatively random. Therefore, most of the contents of each retrieved memory line won't be used, since successive reads will most likely come from different regions in memory. MRAM, by contrast, is byte-addressable. Only those bytes needed are retrieved, making access faster and reducing energy requirements.
- Cache storage and recovery in response to power events (up, down, fail, and recover) take time and use energy—some of which is supplied by the backup battery. MRAM requires no such handling, which saves that energy.

### MRAMs CREATE NEW NVDIMM OPPORTUNITIES

NVDIMMs as defined today leverage DRAMs as cache, mediating access to the full flash storage. Even for today's applications, the current architectures can benefit from a swap of an MRAM in place of the DRAM. Power events no longer need to be handled; backup energy requirements drop; performance increases; and management software becomes simpler.

Meanwhile, new MRAM-based designs that permit direct access to both the MRAM and flash can open up new applications like key-encrypted key storage for accessing namespaces in the flash memory. This is but one variation; other architectures are likely to present benefits for different applications.

NVDIMM architects are faced with the opportunity to innovate as system designers envision yet new ways to use NVDIMM memories. Spin Transfer Technologies remains committed to working with NVDIMM creators to explore how this new memory class can further reduce system cost, decrease system power, and improve system performance.

BOB WARREN is a storage industry expert with over 33 years of experience. Early in his career, he began developing storage controller subsystems to house fingerprint identification data. While working for Micro Technology Inc. (MTI) in 1996, Bob gained an appreciation for the impact of differing memory types and their value to storage. During this tenure at Western Digital and Seagate, Bob did initial development of the world's first hybrid hard-disk drive, as well as looking at applications for emerging memory types. He is currently working with Spin Transfer Technology to advance MRAM, one of these exciting new technologies.

(Continued from page 13)

Details about the chip are a bit sparse at this point. In general, one of its advantages is to 3D tensor calculations without unfolding the data into 2D matrices. According to its patent, “The contraction engine calculates the tensor contraction by executing calculations from equivalent matrix multiplications, as the tensors were unfolded into matrices, but avoiding the overhead of expressly unfolding the tensors. The contraction engine includes a plurality of outer product units that calculate matrix multiplications by a sum of outer products. By using outer products, the equivalent matrix multiplications can be partitioned into smaller matrix multiplications, each of which is localized with respect to which tensor elements are required.”

Part of this approach is to minimize the amount of data movement. This is also something that Flex Logix does with its NMAX approach to neural-net processing. Moving data around takes time and power, but it’s necessary to keep the matrix multipliers flowing. Most systems can’t keep these calculators running all of the time and are often waiting for data to arrive.

NovuTensor can be used for most inferencing chores. It’s able to handle challenging applications like scaling streaming media to 4K video using a single chip or 8K video using four chips *(Fig. 2).*

**2. A NovuTensor chip takes on tasks such as scaling streaming media to 4K video or 8K video using four chips.**

The challenge for developers will be benchmarking this and other chips with their applications and neural-network models. Most benchmarks these days don’t address real-world applications all that well. Likewise, the scale and implementation of a model can have a significant impact on how it’s partitioned and implemented on a particular system. This will be especially critical for embedded systems where using the smallest, lowest-power chip can make the difference between a good, economical product and an expensive one that doesn’t work. ■

M. ANTHONY AIELLO | Technology Strategist, AdaCore

# SSI: Continued Assurance from Requirements to Code

**Maintaining systems-to-software integrity (SSI) throughout the development process improves software quality, reduces development cost, and increases safety and security.**

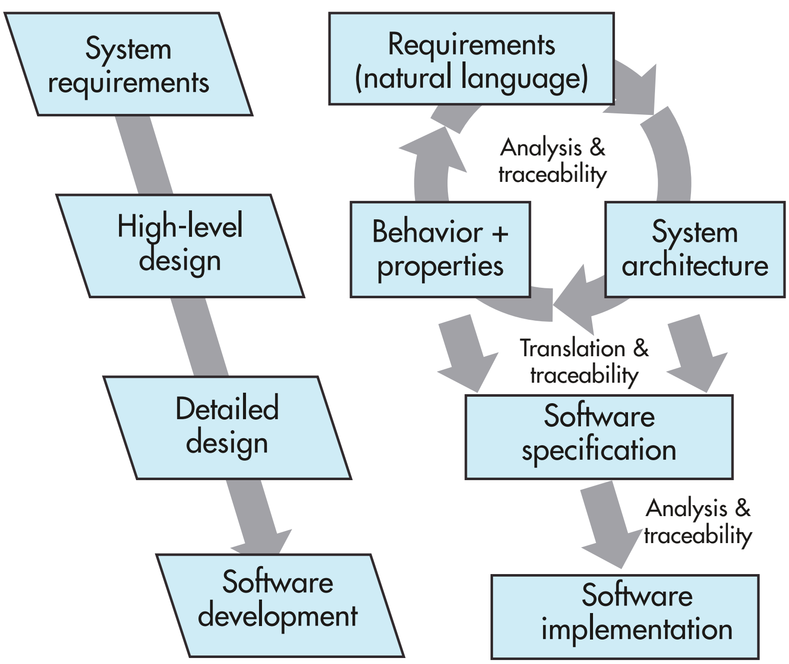


This high-level workflow supports SSI.

The development of systems, whether high-assurance or mission-critical, faces a perennial problem: We can envision and therefore want to build systems that are much more capable and complex than we can assure with traditional approaches. The result is predictable: System errors resulting in system failures that result in losses—sometimes including loss of life.

Many system errors can be traced to erroneous behavior of critical system software. Software, even more than systems, is prone to escape from the envelope of manageable complexity. Software has no physical extent, no physical weight, and no physical power demand per se. The pressures that constrain the features we want to add to systems don't constrain software. If we can imagine a feature, we can add the feature to software. The result is again predictable: Software that's far too complex to assure; that fails; and that causes system failure.

### GOING FORMAL?

Software assurance has traditionally been gathered through extensive testing. Unfortunately, testing isn't exhaustive and thus can only reveal the presence—not absence—of errors. But we can do better than just testing. At the system level, we apply analysis as well as testing. For example, when designing an aircraft, aerospace engineers conduct extensive aerodynamic analysis using computational fluid dynamics before going into the wind tunnel. The wind-tunnel tests validate the analytical model, upon which the fundamental assurance is based.

For software, we can apply formal methods. These are techniques based on mathematical representations of software behavior that allow for a truly exhaustive examination of software state to prove the absence of errors. Software testing then validates this comprehensive analysis.

# FREE WEBINAR

## November 29, 2018 | 2:00pm ET

**Registration URL:** http://comsol.com/c/7qw2

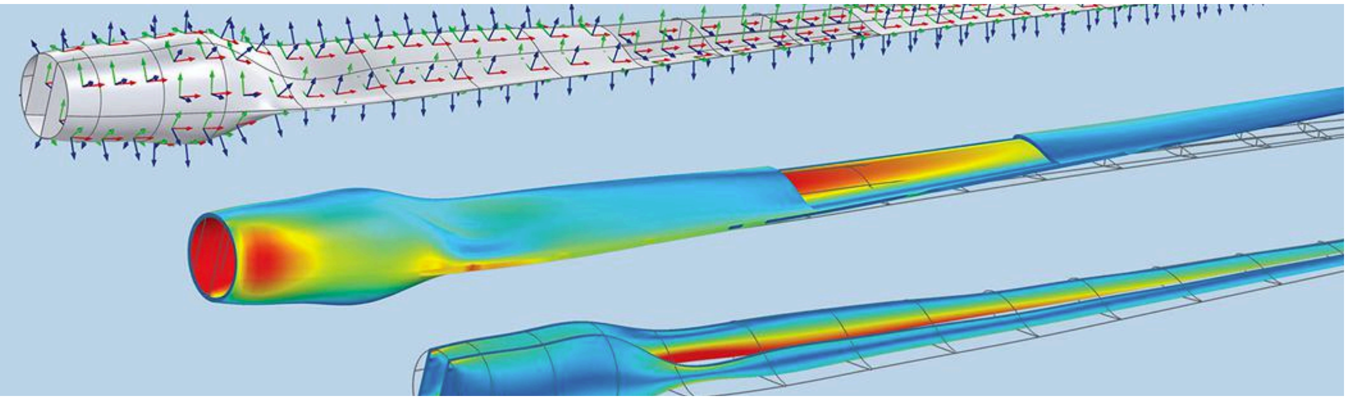

Wind turbine blade. Visualization of the shell local coordinate system (top) and the von Mises stress in the skin (middle) and spars (bottom).

# COMPOSITE MATERIALS MODELING

**In this webinar,** we will show how to model layered composite structures. Common examples of composite materials include fiber-reinforced plastic, laminated plates, and sandwich panels, all of which are widely used in the manufacturing of aircraft and spacecraft components, wind turbine blades, and other structures.

We will discuss two approaches that accurately model composite shells: layerwise theory and equivalent single layer theory. In addition to standard structural analyses (such as stationary, eigenfrequency, transient, and frequency domain), the presentation will go over micromechanical, first-ply failure, and linear buckling analyses. We will also demonstrate how to include various physical phenomena, such as heat transfer, electromagnetics, acoustics, and fluid flow, in a composite laminate model. A live demo in the COMSOL Multiphysics® software will show simulation results for the through-thickness and ply-wise stress distribution in the composite laminate. The webinar will conclude with a Q&A session.

**SPEAKER:**
**Pawan Soami, *Product Manager, COMSOL***

***Pawan Soami,*** Product Manager, COMSOL

Pawan Soami is the product manager for the Multibody Dynamics Module. He has been working with COMSOL in the Bangalore office since 2010. He received his master's degree in aerospace engineering at the Indian Institute of Science, Bangalore.

*Hosted with*
ElectronicDesign.

COMSOL

Despite its power, the use of formal methods isn't a silver bullet. In particular, software engineers must ask the right questions when conducting analysis. Asking the right questions is hard when critical system-level properties haven't been explicitly traced to software. Software engineers often don't know which properties need to be proved.

When focusing on SSI, systems engineers identify critical system-level properties and trace those properties through decomposition across system architecture and through translation from one artifact to the next. This attention to the evolution of system-level properties enables systems engineers to effectively communicate to software engineers the critical properties that they should prove. Focusing on SSI enables engineers to take full advantage of formal methods.

As model-based systems engineering has become more popular, tools supporting this approach have included increasingly powerful analyses. These tools go beyond simple checks for latency or space, weight, and power consumption, and they seek to demonstrate functional correctness of the architecture. Effectively, these tools apply formal methods at a higher level of abstraction. To take full advantage of these tools, systems engineers need clearly identified properties that they should prove. Focusing on SSI enables systems engineers to take full advantage of these analysis tools.

Talking to systems and software engineers in industry, we often hear that systems-engineering artifacts like requirements documents and design documents are "tossed over the wall" from one team to the next. This statement is somewhat astonishing. Surely all engineers involved in development—from systems engineers to software engineers—should be a single, integrated team.

Different people have different specialties, and software standards such as DO-178B/C require, at the highest Design Assurance Level, independence between those who develop a component and those who verify it. However, integrity is certain to be lost if there's not constant and managed communication amongst team members. Focusing on SSI enables systems and software engineers to bridge the walls that separate them, by focusing communication on the critical properties along with the artifacts that serve as the interface between team members.

## HOW IS SSI SUPPORTED?

SSI is a trait of high-assurance systems engineering—not a tool or technology. However, to facilitate SSI in systems engineering, tool support is essential. There are four integral parts to this tool support, plus an optional fifth: translation, traceability, analysis, argument, and (optionally) test-case generation.

### *Translation*

Translation assists in preserving integrity by automating the representation of properties expressed in one artifact into the next. For example, in the SysML modeling language, a constraint associated with a requirement that's written in OCL (the Object Constraint Language) and attached to an output port on an Internal Block Diagram might be translated into a synchronous observer in a Simulink model that's written using Simulink blocks and attached to an output port of a subsystem. Similarly, the constraint might be translated to a postcondition in a programming language like SPARK that supports software contracts.

Translation bridges the gap separating the architecture (in SysML) from the design (the Simulink model) or the implementation (in SPARK). It also reduces the likelihood of error by not requiring that an engineer manually recode the property in the target language. Of course, this requires that we have confidence in the correctness of the translation, but techniques are available (such as tool qualification under DO-178B/C) that can offer this confidence.

### *Traceability*

Traceability assists in preserving integrity by allowing engineers to follow the evolution of properties through system development. Continuing the example above, the Simulink synchronous observer would be traced to the OCL constraint so that if either was changed, engineers would be notified that the other must be reexamined or regenerated.

### *Analysis*

Analysis assists in preserving integrity by allowing engineers to prove property satisfaction across decompositions. At each phase of system development, greater abstraction is replaced by greater precision, until the system has been fully and precisely described. Decomposition plays a critical role in this ever-increasing precision, as larger abstract components are broken into multiple, smaller components.

With analysis, engineers can prove that the smaller components, together, satisfy the intended functionality of the larger, abstract component. Continuing the example above, the Simulink subsystem would be broken up into smaller subsystems, each with its own contracts that, together, must satisfy the observer for the top-level subsystem.

### *Argument*

Argument assists in preserving integrity by allowing engineers to record their rationale for property continuity when traceability or analysis are insufficient or unable to provide sufficient support. Continuing the example above, one of the subsystems may be COTS and might not have a sufficiently precise contract to enable proof that the observer for the top-level subsystem is satisfied. Engineers may know, however, that the way in which the COTS subsystem is used will, together with the other subsystems, always satisfy the observer for the top-level subsystem. The argument documents this rationale.

A key requirement of tools supporting SSI is that they meet systems engineers where they currently are: Rather than ask engineers to adopt entirely new languages and entirely new development methods, these tools should support existing languages and existing development methods.

***Test-Case Generation***

Test-case generation assists in confirming integrity by demonstrating that critical properties are satisfied by the system. Continuing the example above, the observer for the subsystem could form the basis of a test case or set of test cases to be included in that phase of integration testing.

A key requirement of tools supporting SSI is that they meet systems engineers where they currently are: Rather than ask engineers to adopt entirely new languages and entirely new development methods, these tools should support existing languages and existing development methods. Tools supporting SSI don't currently exist, but AdaCore is working on building them. When finished, they will support a workflow that might look something like what's illustrated in the *figure*.

**WHAT ABOUT HARDWARE?**

This article has focused on SSI in terms of software: system-to-software integrity. But SSI is equally applicable to hardware. The languages used in the final phases of system development will differ for hardware, but the overall approach can still significantly benefit hardware development as well as software development.

**THE BOTTOM LINE**

Taking an integrated approach to system development and verification—by identifying key properties early and ensuring that they're preserved as a system progresses throughout its lifecycle from requirements to implementation—can save time and effort. It also prevents mismatches between what a system is supposed to do and what it actually does. With automated support from existing and upcoming technology, SSI can offer the assurance that's needed for the most critical systems that society depends on. ed

DAVE HUGHES | CEO, HCC Embedded
www.hcc-embedded.com


**Transport Layer Security is the most widely used approach for security on the internet—the majority of secure transactions depend on it. But myths about TLS still abound.**

Security issues are persistently front and center when it comes to the internet, and Transport Layer Security (TLS) often is the go-to solution. Nonetheless, myths surround the technology. HCC Embedded CEO Dave Hughes looks to dispel some of these misconceptions.

**1. TLS is broken and can't provide adequate protection against hackers.**

Hearing about widely publicized security breaches, you would think that those designing security are incompetent. This is simply not the case. The truth is, there are no known hacks of TLS 1. Rather, these hackers were successful not due to faulty TLS, but because of a lack of software-quality processes.

For example, a well-designed static-analysis tool would have detected Apple's 2017 TLS vulnerability before it was released. And the Heartbleed Bug, which resulted from an implementation defect in some OpenSSL versions, was caused by software that did not check the scope of a protocol variable and then processed it blindly.

Software-quality processes that include unit testing and boundary case analysis/testing would have instantly alerted developers to the issue, and the detection would have been reinforced by other requirements of the lifecycle process.

**2. TLS 1.2 is perfect and will always protect you.**

This is never going to be true, but the main criticism facing TLS (and all attempts at security) is that it can be difficult to use safely in real-world environments. This has been demonstrated by the stream of security failures. However, as stated in Myth 1, these protocols can only be effective if they're implemented properly, using proven software-quality processes.

**3. TLS 1.3 will fix problems with TLS 1.2.**

TLS 1.3 is really an enhancement of 1.2, reducing information leakage and cleaning things up rather than fixing TLS 1.2. While it may not fix all problems, the main things that TLS 1.3 enhances are security and connection speed. Security is improved both through the algorithms used as well as the TLS protocol exchanges, which give fewer attack vectors. The speed of TLS has also been improved

by simplifying certain key protocol exchanges. Nothing was actually broken in TLS 1.2—it's just a continuous struggle to keep TLS ahead of the bad guys and to improve the user experience.

**4. TLS and network security are all about cryptography.**

Most recent network security failures have been caused by either the leak of key information by humans, badly written code, or poor integration of the security layer. While some cryptographic algorithms have been weakened in lab conditions, practical attacks rarely exploit these weaknesses.

**5. Cryptographic algorithms can't be broken.**

When using currently available computing power, this is probably true. However, computing power, particularly that available to nation-states, is increasing rapidly and consequently must be kept under observation. Embedded developers first need to assess their risks. If you're really concerned about being hacked by a nation-state, then you need to take different measures than if you're trying to protect your data from a company trying to extract competitor information. Many would argue that if a nation-state wants to hack you, you have bigger problems to address.

**6. Adding compression makes encryption more secure.**

Surprisingly, the opposite is the case. The reason is that compression reduces the size by keeping dictionaries of common strings. Therefore, by injecting strings and checking the change in size, you can find out how much that data is used. That's not an easy thing to process, but it's a very useful attack vector for the determined hacker.

**7. Properly implemented TLS has no information leakage.**

This isn't true, and although desirable, "zero information leakage" was not a goal of TLS. The types of leakage that occur can include the size of messages being conveyed and the fact that a conversation between two people has occurred. Developers need to decide if these are real risks for their system and take additional measures as appropriate. The use of VPNs, for example, can provide additional levels of security.

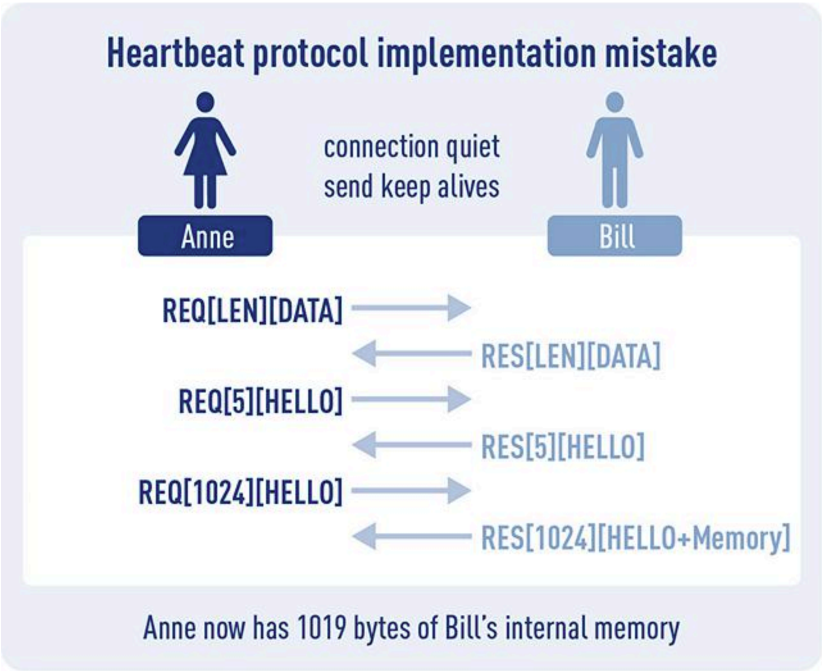


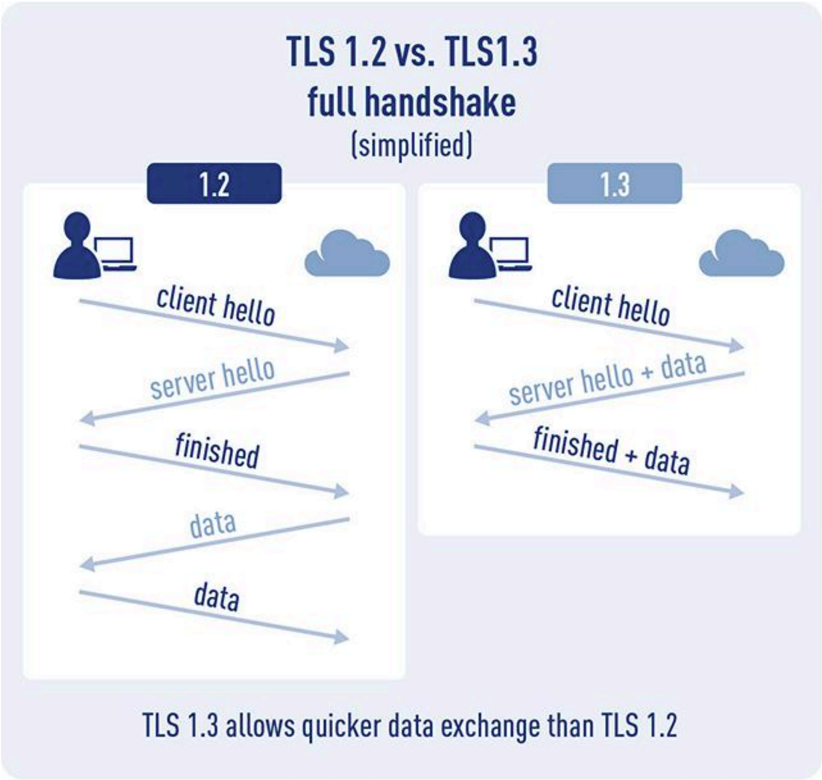

**State-of-the-art algorithms are available for most browsers, but they often negotiate down to older versions where required. Web browsers all fight for market share, so they want backwards compatibility with as many servers as possible. On the other hand, they want to be very secure, so they have a difficult balance to strike.**

**8. Web browsers mostly now use state-of-the-art security.**

State-of-the-art algorithms are available for most browsers, but they often negotiate down to older versions where required. Web browsers all fight for market share, so they want backwards compatibility with as many servers as possible. On the other hand, they all want to be very secure, so they have a difficult balance to strike. But as a matter of course, developers should all use the latest versions of TLS (1.2 and/or 1.3) and not allow this to be negotiated down. All modern browsers support at least TLS 1.2.

**9. Side-channel attacks are commonplace.**

Side-channel attacks have been used to break crypto-algorithms. To execute them, though, you normally need to be physically local to your target and have the ability to execute millions of test cases without anyone noticing. At best, it's useful for targeting a single specified target.

**10. TLS will never be broken.**

This seems unlikely. But since both computing power and the imagination of hackers continues to rise, it's possible that a point will be reached where trying to break TLS isn't worth it. This seems unlikely. Rather, it seems inevitable that breaches of TLS will occur and improvements will be required.

However, as the protocol and algorithms are improved, there may come a point where hacking the protocol is too difficult for most to attempt—and that other methods of attacking systems will

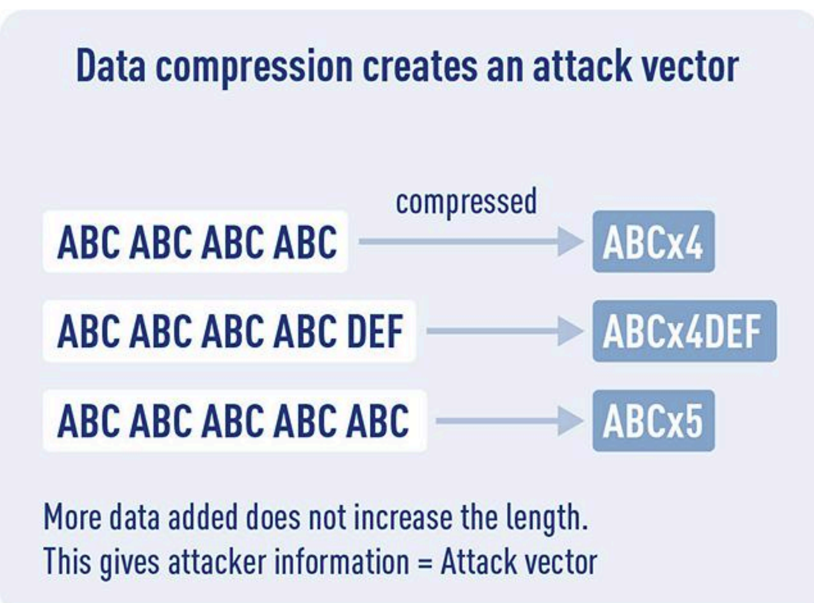

The expectation that there's a foolproof solution to security is naive. Most high-profile security breaches have come from three main sources: insiders divulging secrets, poor system management, and badly or inappropriately written software.

be preferred. The most obvious example of this is social engineering, where an attacker works in other ways to find system attack vectors. An example of this might involve befriending someone and getting likely password possibilities from him or her. The point is that the users and system managers will have become the weakest point in the system.

**11. Security is a perfected art.**

The expectation that there's a foolproof solution to security is naive. Most high-profile security breaches have come from three main sources: insiders divulging secrets, poor system management, and badly or inappropriately written software. The first two sources can only be dealt with by the organizations responsible for the security of the information. Clearly, there are no easy solutions where humans are involved.

But the third source—software quality—is often neglected, especially with rapidly increasing networking requirements and new situations that create new risks. This is where developers should focus their efforts: Sourcing and creating high-quality software developed through the use of proven lifecycle processes. ed

DAVE HUGHES founded HCC Embedded, a leading developer of reusable embedded software components, in 2000. In that time, the company has grown substantially, supplying many of the industry's major technology providers. Dave is a "hands-on" embedded specialist, who still actively contributes to the strategy and direction of HCC's core technologies. His extensive experience has made him one of the industry's leading authorities on fail-safe embedded systems, flash memory, and process-driven software methodologies. Dave is a graduate of the University of Sussex in England.

**REFERENCE:**

Verifiable Embedded TLS/DTLS

the **difference** between

THE DIFFERENCE BETWEEN:

# Screw Clamp and Spring Clamp DIN Rail Terminal Blocks

**TERMINAL BLOCKS** are a necessary part of almost every industry, supplying a way to hardwire components, subsystems, and systems together. As a staple of the electrical and electronics control systems used in a wide variety of applications, including utilities, automobile plants, chemical plants, and all types of automation and manufacturing that requires hardwired interconnection, the breadth of available products indicates how important application is to secure the right terminal block components. Most often, terminal blocks are found inside control panels to interconnect the panel wiring internally as well as to field wiring. Terminal blocks have a crucial role for providing good wire management in any project.

Because of trends in automation and manufacturing, the look and feel of terminal blocks has evolved and expanded into their own trends, including compact sizing, the use of a variety of materials for different environments, and the adjustment and expansion of clamp technology from those using screw clamps to spring clamps of various types. But whatever your application, you want to be sure to use the right product for the right job, whether the situation calls for DIN rail-mounted terminal blocks with front-entry wiring, side-entry wiring, or top entry wiring, or panel-mounted terminal blocks, there are products that provide the perfect solution for any challenge. Ultimately, your choice of terminal block must maintain a long life of reliable service, be as low maintenance as possible, and be able to be installed quickly and efficiently under any conditions.

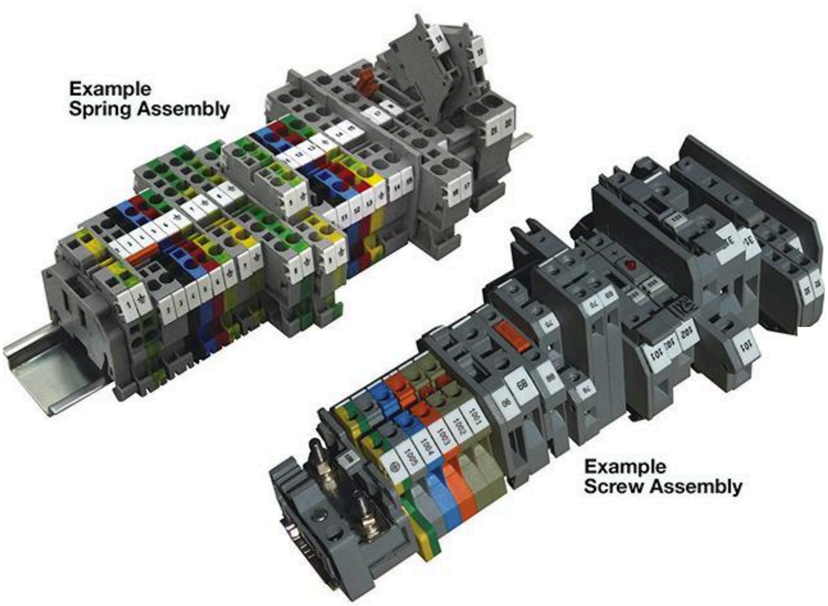


This article will explore the more common options for your terminal block requirements, looking at the two most standard products, the screw clamp and the spring clamp terminal block.

## Screw-Type DIN Rail Terminal Blocks

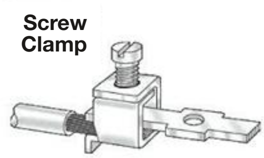


Screw terminals are used extensively in the distribution of electricity in homes, apartments,

**Screw Clamp Instructions**

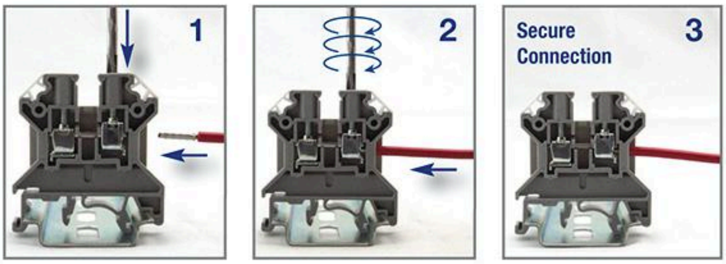


1 Insert Wire into open clamp

2 Tighten screw to specified torque

3 Remove Screwdriver, screw terminal closed with locking feature for a reliable/safe connection

and stores, as well as in utility plants and inside all types of controller and junction boxes for automation and manufacturing systems and in multiple industries.

Operation of this type of a screw terminal block is simple: you insert the wire into the open clamp, tighten the screw to a specified torque, remove your screwdriver, and close the locking feature over the terminal for a reliable and safe connection.

One of the benefits of using a screw-type terminal block is that they can be manufactured for very high-voltage applications and offer the maximum in wire efficiency and reliability to the user. These terminal blocks are supplied with open clamps, ready to accept wires. No wire preparations are necessary because the terminal can be accessed easily whether you are using a ferrule or not.

The terminal block must be able to handle the physical stress of multiple, large (sizes up to 250 kcmil) wires mounted in the same general location as well as handle the large currents going through the wires. This typically means that the terminal block must be sturdy and tough and made from high-impact materials such as polyamide PA66, self-extinguishing plastics whether mounted on a DIN rail or on the side of a metal enclosure. Once completed, these terminal blocks provide vibration-resistant, gas-tight connections designed for the most stringent requirements.

Most terminal blocks designed for industry are color coded so that the user can easily determine ground wires from hot wires and control wires. Other options are also available, such as push-in jumpers, partition plats for separation (used in wire management), and large center marking areas for important information.


## Spring-Clamp Terminal Blocks

There are many benefits to using the latest spring-clamp terminal blocks, the most important being the time it takes for installation. With the use of a tool or narrow screwdriver, push into an insertion space to open a pre-loaded stainless steel spring (see figure below). Your wire can then be inserted into an opening in the spring leg. By removing the tool, there is a spring force that presses the conductor against a copper tin-plated current bar for a highly reliable connection. These terminal blocks are used with solid or stranded wire and create a very secure connection—unlike screws that can loosen over time or in high-vibration applications. Spring-clamp terminal blocks are often designed into vehicles because of their natural resistance to vibration and shock. Under the high-vibration movement of vehicles and aircraft, the spring-clamp terminal block maintains a secure hold on the wire long after when a screw-type terminal block would have loosened and caused a potential failure.

Another type of spring-clamp terminal block is designed specifically for single-wire connections allows the wire to be pushed directly into the spring block without using a screwdriver or tool to complete the operation. Push-in spring-clamp terminal blocks are being used in more and more applications as their benefits are realized.

With spring-clamp terminal blocks in general, conductors can be terminated securely in mere seconds, independently of one another and with absolute reliability. This easy installation reduces wiring time by over 50 percent, while delivering a reliable, maintenance-free connection independent of operator skill. Not only are these terminal blocks versatile, offering a reliable gas-tight connection but they are vibration-proof, having no screw to loosen during operation.

Spring-clamp terminal blocks feature a push-in jumper system with various available pole configurations. They allow easy and reliable shorting connections, and tabs can be easily broken off for across terminal block shorting applications. Two jumper channels allow the user maximum flexibility.

These types of terminal blocks can also be used with all sizes of wire and are actually easier to use on small wires than a screw-type block. They provide a fail-proof and

**Spring Clamp Instructions**

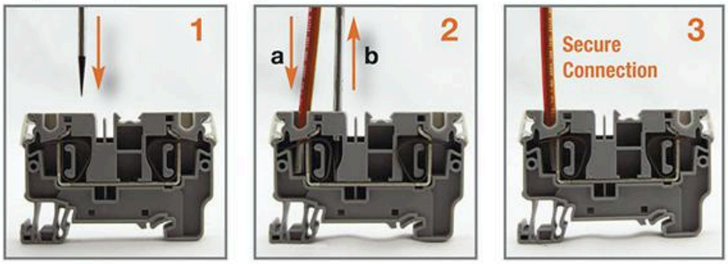


1. Insert Flat-Head Screw Driver into insertion space to open pre-loaded spring.
2. (a) Insert Wire into opening in spring leg.
   (b) Remove screwdriver.
3. The spring force presses conductor against the current bar for a reliable/safe connection.

safe connection with no torque requirements. Like other terminal blocks, spring-clamp devices accept wires with or without ferrules to provide a completely maintenance free operation.

This technology also allows DIN rail-mount terminal blocks to be some of the industry's most compact to date. An added feature to DIN rail devices is that the ground connection has been added to the carrier rail, eliminating the need to provide additional wiring steps, reducing the many work steps involved to a minimum.

Terminal block design and manufacturing has come a long way in the past few years due to the advancements in materials and the needs of specific applications. Many manufacturers offer a full line of components that fit the broadest uses in the industry. Contacting and reviewing data sheets on these components will familiarize you with what is available and what will best suit the criteria of your particular application. ■

FAQs

# Polycarbonate Enclosures for Harsh and Unprotected Environments

## What makes using a polycarbonate enclosure safe for harsh environments?

Polycarbonate is a very durable material which offers high impact resistance, low scratch resistance, and shatterproof safety to users. The material also provides temperature, corrosion, UV, and weather resistance for a wide variety of unprotected environments.

## What do I need to look for, besides ruggedness, when it comes to the design of an enclosure for outside use?

Overall design has a lot to do with the right enclosure for your application. For example, if located outside you'll want to be sure that it's rainproof and dustproof, even when it has air ventilation features as part of the unit. You will want it to have protection ratings such as IP 66/IP 67. You may also wish for the unit to be lightweight, which polycarbonate generally is.

## What type of sealing technology might I look for?

There are a lot of ways to seal an enclosure. For example, the Altech GEOS uses a unique enclosure design that actually diverts moisture to the back side of the enclosure using an overlapping cover and drain system. This new sealing principle is called "Drain Protect" and discharges moisture by providing a widely overlapping cover with a built-in drainage channel along the top edge of the enclosure base to divert moisture to the back side.

Additionally, an elastomer seal provides for optimum protection from the environment. An elastomer seal provides additional protection against concerns such as dust and dirt.

## What about impact protection?

Good question. Knowing your application is key, but most outdoor applications do need a greater amount of impact protection. The right enclosure will most likely offer a high degree of impact protection up to IK09 impact rating in order to assure your equipment doesn't get damaged.

## What difficulties might I have when configuring an enclosure for my components?

That depends on the enclosure you purchase, of course. What you'll want to look for is an enclosure that offers some sort of extendable, functional installation system with mounting brackets. Make sure the unit you select is usable for mounting plates and standard rails, ideally at different installation heights. You might also look for tool-less installation of your mounting bracket for easy and time-saving mounting.

## What else might I consider for outside use?

It is always important to make sure the enclosure not only has the technical ratings you need, but also the required certifications for your application. Watch for proper certifications you might need, including VDE, UL 50, IEC 62208, IEC 61439, and CSA 22.2.

## Do I need to consider how to brand my company on a purchased enclosure?

Many companies use enclosures as part of their branding. When making a purchase, be certain that the enclosure company offers options for applying branding elements, such as hot embossing, screen, pad, digital printing, or laser marking.

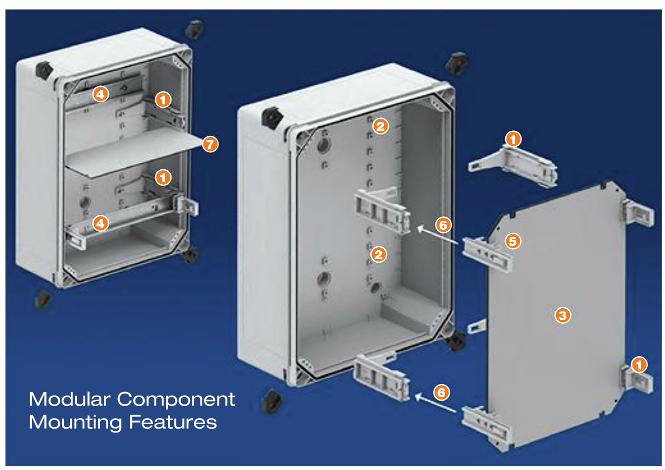


1 Functional installation system with mounting bracket. 2 Tool-less installation of the mounting bracket in 25 mm location points. 3 Universally usable for mounting plates and 4 standard rails. 5 Different installation heights possible in the assembly bracket. 6 Subsequent insertion of pre-assembled modules. 7 Partition walls.

### Are there adequate ventilation systems available?

Yes, there are a number of ventilation methods used. For example, Altech GEOS uses a patented air ventilation element that prevents condensation while maintaining an IP65 protection rating. Through the use of a continuous and high rate of air exchange, condensation is kept out of the enclosure. This air exchange takes place even in environments with almost constant air humidity and temperature changes.

The interior of the enclosure heats up from component use and, due to the permanent pressure equalization, all enclosure parts including the seal are permanently relieved of pressure. The system also provides 100% protection against foreign bodies entering the enclosure for maintenance-free operation.

### Can I find enclosures that I can customize to my needs?

Yes. Every application can have its nuances, which means that you'll want to be able to customize your enclosures much of the time. If you choose to work with a company that offers customization opportunities, you'll save a lot of time and effort. Look for a company that can customize the enclosures they sell by offering milling and printing services.

But it's not just the outside of the enclosure that typically needs to be specialized: Some applications lend themselves to external mounting, while others will need to be mounted from the inside. You may also want free choice between screw closure and quick-release closure. Finally, look for a company that provides a wide variety of enclosure sizes, with different heights, base sizes, and opaque or transparent cover options.

### What might some primary applications be?

Enclosures made for harsh environments and outdoor use can provide protection for a wide variety of applications. Some notable ones might be for any type of electrical installation, such as those found in the solar and wind power industry; for large plant or building heating and air conditioning controls; and industrial applications such as the food and beverage industry, where washdown might be expected.



# design FAQs

## Understanding and Using E-Stops

FREQUENTLY ASKED QUESTIONS

**Q: What is an e-stop and how is it used?**

A: E-stops, or emergency stop switches, are used to ensure machine as well as personnel safety. They are used to provide a consistent and predictable failsafe response on a wide range of electrical machinery and must stop the machine without creating additional hazards. The devices can be highly specialized for emergency shutdown of equipment and meet workplace and machine safety standards established by international and U.S. regulatory commissions.

**Q: Is there a difference between e-stops and regular stop switches?**

A: E-stops provide what can be considered foolproof equipment shutdown, and always require a human action for resetting. Often, the switch requires an additional step—a twist, pull, or key—in order to release the electrical contacts prior to the machine being in a position to be restarted. As a general standard, e-stops must be a red operator with a yellow background. Companies often offer additional product options for particular applications. For example, some companies offer the same operators in various specialty colors. These non-red operators do not qualify as "Emergency Stops," but can be applied in a similar way to stop applications. Black operators are used as a Machine Stop, similar in function to e-stops, but are simply a different color. Typical application for these devices is when the machine's "OFF" button is required to be manually reset before restarting the machine. Blue operators are an accepted designated color for stopping water or sprinkler systems, and yellow operators are an accepted designated color for shutting off gas lines. Different color operators were designed for customers with specific needs.


**Q: What standards must e-stops adhere to in order to be consid ered properly certified?**

A: Always check with your supplier to assure their devices are tested and approved by appropriate institutions relevant to your application. Some important standards include: IEC60947-5-1 and EN60947-5-5; VDE0660; UL508; CSA: C22.2 No. 14-95; and NEMA Type 4X, 12. These are some of the most used in the U.S. while other compliance and rating bodies also exist for other countries as well.

**Q: What types of applications are required to have e-stops installed?**

A: All industry segments mandate e-stops for safe operation, including, but not limited to industries involving vehicles and transportation, medical treatment and diagnostics, industrial machinery, oil and gas, food and beverage, water and waste water, and instrumentation. Therefore, designers will want to have a knowledge of e-stop fundamentals, and switch characteristics and capabilities, as well as the international and U.S. standards and compliance requirements that need to be met for their application.

**Q: How can I begin to select the right e-stop for my application?**

A: The first step is to determine where the e-stop fits within your machine control system and what category of emergency shutdown is needed according to the standard you are adhering to. The intended application often determines the electrical and mechanical specifications, as well as the size and placement. So, a thorough understanding of the machinery and associated control systems is key to making the right e-stop choice. Requirements vary by industry segment, therefore standards for e-stops used in the transportation industry may differ significantly from those used on process machinery or medical equipment, and will be governed by different regulatory bodies, provides online access to information that allows designers to select the panel opening size, type of actuator, type and number of contact blocks, and all the electrical ratings you'll need to adhere to.

**Q: Are there specific changes in any of the standards that I should be aware of?**

A: There are continual adjustments being made to the standards as issues arise. One of particular interest recently is concern about generators — a key component in many industries. The NEC 2017 standard, for instance, has recently changed for article 445 on disconnecting means and shutdown of prime mover, and should be read thoroughly before finishing your design. This new change in the NEC standard requires the generator installer to include a lockable disconnecting means, see picture, that would disable all prime mover start control circuits—which would then require a mechanical reset. This additional shutdown e-stop switch is to be located outside the equipment room or enclosure and should not be installed with an e-stop lockout or a lockable shroud.

**Q: What is the difference between e-stop lockout shrouds and e-stop guards?**

A: E-stop shrouds are usually lockable and used where safety is a requirement for the particular application's specific design and installation parameters. E-stop guards, on the other hand, are not lockable and can be used in any application where protection of the operator is required but the operator has a full access to the e-stop operator. This guard is meant to prevent the machine operator from accidentally bumping into the e-stop while working, but does not provide a lockout function. ■

Altech Corp.®

Sponsored by Altech Corp.

# design FAQs

## Alternative Solutions for Backup Power

FREQUENTLY ASKED QUESTIONS

**Q. What current options are available for backup power solutions?**

A. Systems that require backup power solutions usually have small battery backups or generators. Small backup systems traditionally rely on batteries for energy storage while larger systems may use generators or more uncommon systems such as flywheels, superconducting magnetic energy storage, or fuel cells for larger installations. While there can be a need for a battery based solution, batteries have a limited life span. Their life is based on the number of discharges and charges they undergo. Battery installations are difficult due to the long list of regulations and standards they must abide. Their installations must comply with local and national health and safety standards. Environmental regulations, which mainly affect larger installations, can include requirements for ventilation, handling, and disposal of batteries. Flywheels, for example, require an extreme amount of maintenance. Also, systems that use generator sets and fuel cells have poor turn-on response and are not very reliable.

**Q. What alternatives provide a more secure and reliable power source?**

A. Two alternatives that can be used instead of generators are an all-in-one power solution system and an ultra-capacitor system.

An all-in-one power solution combines a power supply unit, battery charger, battery care module, or backup module. The all-in-one battery solution is based on the battery operating as a storage device (batteries are not included). The available power is automatically distributed among load and battery, while supplying power to the load always is the first priority. The maximum available current of the load output is two times the value of the device's rated current. If the device is disconnected from the main power source, the battery will supply the load until the battery voltage reaches 1.5 V per cell. This prevents the battery from deep discharge. The battery charging is controlled via a microprocessor. Using algorithms, the battery's condition will be detected and based on that, an appropriate charging mode is chosen. The real-time diagnostics system will continuously monitor the charging progress and indicate possibly occurring faults such as elements in short circuit, accidental reverse polarity connection or disconnection of the battery by the battery fault LED and a flashing code of the diagnosis LED. The all-in-one solution can provide a variety of charging voltages: 12VDC, 24VDC, and 48VDC. The solution is highly efficient; up to 91% via switching technology.

The ultra-capacitor solution stores energy within a compact design that does not use a battery but supper capacitors to store energy. The load is energized from the buffer module until it is depleted. The capacitors are either a 12 or 24V system design. It operates between temperatures of -40°F to 140°F and can provide up to 10,000 watts of energy plus extension modules. The system can be customized up to 600A. Backup times depend on the state of the charge of the ultra-capacitors and the load in which they are supplying. Compared to standard buffer modules, ultra-capacitor units are capable of prolonged backup times (up to 55 minutes) and fast discharges. Capacitors excel at controlled shutdown functions and allowing for the protection of computer systems. Backup times are dependent on the load and amount

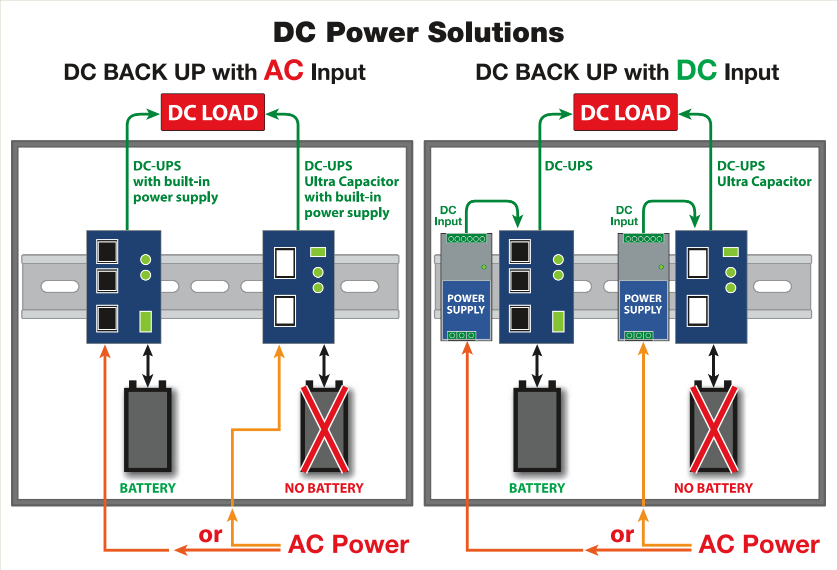

of capacitors within the unit. Backup times can be calculated to better serve the requirements of required functions.

**Q. In which situations would one prefer to use an ultra-capacitor versus an all-in-one multi-power solution?**

A. An all-in-one solution is beneficial for cleaner power and versatility of use. The solution is suitable for most common battery types and the device itself is mounted in a compact design case for protection. A higher efficiency of the battery is achieved thanks to continuous control over time. They offer more monitoring in the main connection nodes: input, output load, battery. The solution provides adjustable current if that is a requirement and ease of battery diagnosis. The all-in-one solution has several output protection features as against troubleshoots like short circuiting, overload, deep battery discharge, etc. The ease of diagnosis is possible through the event logging of: number of battery charging cycles, charge cycles completed, aborted charge cycles, Ah charged, charging time, total number of transitions standby / backup, etc. Examples of the event management are: checking the load output, shutdown management of PCs (UPS function), RESET management of a generic equipment.

The all-in-one solution offers DIN rail mounting. Also, the load output will not be affected by battery conditions. The solution insures continuous power supply to the load even in conditions of completely discharged batteries. The automatic multi-stage operation optimizes and adapts to the battery status. It can recharge deeply discharged batteries even when their voltage is close to zero, thus allowing recharge and complete recovery of flat batteries.

The downside of the all-in-one solution is that it uses a battery. They are used more widely in industry and are easier to implement. Battery has a shorter life cycle and more maintenance required. Also, they have to be replaced more often; standard lifetime is three to five years. Charging batteries also takes a significant long time. The ultra-capacitor solution is beneficial for customizable systems but could incur high cost on initial-purchase hardware and installation. They are resistant to shocks and vibration, environmentally safe, have no toxic chemicals, and are virtually maintenance-free. The ultra-capacitor modules are built with supper capacitors that last up to 20 years. The use of ultra-capacitors also reduces wiring time due to integrated energy storage. The wiring is also vibration-secured via spring-loaded plugs. There is no limitation on their use; they can be incorporated into any design. Ultra-cap modules are also chargeable in a short time and able to provide large amount of energy when needed. They have microcontroller-based charging and discharging of the ultra-capacitors. The ultra-capacitors have seamless switchover and a long operational life.

**Q. What are the industry benefits to using such systems?**

A. All-in-one solutions can be found in several industries. They are used in infrastructure industries like industrial water pumping, fire protection systems, power supply continuity, and remote measurement stations. They are also used in commercial industries like audio, lighting, electric vehicles, off-highway equipment, and wireless control.

Ultra-capacitor solutions can be found in manufacturing industries, data collection services, and energy and construction industries. Among manufacturing industries using them are textile manufacturing, molding machinery, automotive, automation, packaging, and steel production. They are also used as power backups for data centers, feeding systems, wind turbines, and tunneling machines. ■

Sponsored by Altech

MIKE HORTON | CTO, ACEINNA

# Choosing an IMU for Your Autonomous-Vehicle Project

**When developing in application spaces that are unforgiving of inadequacy, particularly autonomous vehicles, pinpointing the right device for critical navigation subsystems becomes essential.**

Sensing technologies for electronic situational awareness has seen lots of progress, driven by all the attention currently being given to autonomous vehicles (AV). The advances in materials, topology, and processes have resulted in a new generation of sensor systems, which in turn enable the creation of new, cost-effective solutions.

This explosion in targeted automated solutions has been enabled by the significant reduction in cost in advanced subsystems. The need for accurate, reliable, and efficient systems to support AV location and movement has thrown the spotlight on a core infrastructure technology, the inertial measurement unit (IMU).

Formerly the realm of the military and other "cost-is-no-object" application spaces like aerospace, IMUs are now available to be used in the most cost-sensitive applications. One example is the European VineRobot project, which is developing a novel unmanned agricultural robot equipped with several non-invasive sensing technologies to monitor vineyards. Until now, it has required direct human action to determine critical vineyard yield aspects like vegetative growth, water status, and grape composition.

Multiple issues are involved in developing a robot that can operate in the real world, and they're exacerbated by application-specific issues. In the case of a vineyard robot, both technical and business challenges need to be addressed.


**1. The AGV developed by the European VineRobot project uses non-invasive sensing technologies to monitor vineyards for vegetative growth, water status, and grape composition.**

Vineyards differ in a significant way from most other agricultural situations. The most notable aspect is terrain, as most vineyards in Europe are on hillsides and other, often steep, landscapes. This creates several challenges to an agricultural robot for vineyards, as tire profile, vehicle suspension, and motor torque become very important when driving off the vertical.

Another serious issue arising from operating (alternating up and downslope) on a steep incline is the higher importance on vehicle presence. What angle are the tires to the face? How steeply and in what direction is the vehicle tilted? There are lots of considerations when going from a 2D to a 3D surface for navigation purposes. Things get stickier quickly when you add issues like vehicle orientation for proper sensor operation.

# Design Note

## High Efficiency 20A Monolithic Silent Switcher®2 Regulator for SoC and µP Applications

Zhongming Ye

Power budgets continually rise for advanced SoC (system on chip) solutions used in industrial and automotive systems. Each successive SoC generation adds power hungry devices and increases data processing speed. These devices require reliable power, including 0.8V for cores, 1.2V and 1.1V for DDR3 and LPDDR4, and 5V, 3.3V and 1.8V for peripheral and auxiliary components. Moreover, advanced SoCs require higher performance than traditional PWM controllers and MOSFETs can provide. As a result, the solutions necessary must be more compact, with higher current capability, higher efficiency, and more importantly, superior EMI performance. This is where our Power by Linear™ Monolithic Silent Switcher 2 buck regulators can satisfy advanced SoC power budgets while meeting SoC size and thermal constraints.

### 20A Solution from 20V Input for an SoC

The LTC®7150S raises the bar for high performance in industrial and automotive power supplies. It features high efficiency, small form factor and low EMI. Integrated high performance MOSFETs and thermal management features enable reliable and continuous delivery of currents up to 20A from input voltages up to 20V without heat sinking or airflow, making it ideal for SoCs, FPGA, DSP, GPU and µPs in industrial, transportation and automotive applications.

Figure 1 shows a 1.2V output at 20A solution for SoC and CPU power using the LTC7150S switching at 1MHz. This circuit can be easily modified to accommodate other output combinations, including 3.3V, 1.8V, 1.1V and 0.6V to take advantage of the wide input range of the LTC7150S. The LTC7150S has the output current capability to act as a first stage 5V supply, which can be followed by a number of downstream second-stage switching or LDO regulators at various outputs.

### Silent Switcher 2 with Excellent EMI Performance

Passing EMI regulations at high currents typically involves a complicated design and test challenge, including numerous trade-offs in solution size, efficiency, reliability and complexity. Traditional approaches control EMI by slowing down the MOSFET switching edge rates and/or lowering switching frequency. Both of these strategies come with trade-offs, such as reduced efficiency, increased minimum on- and off-times and larger solution size. Alternative mitigation techniques, such as a complicated bulky EMI filter or metal shielding, add significant costs in board space, components and assembly, while complicating thermal management and testing.

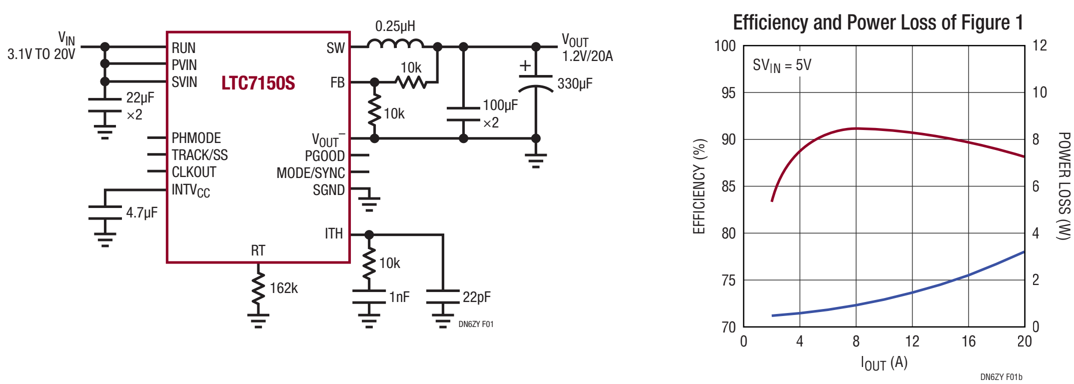


Figure 1. Schematic and Efficiency of the Buck Converter: $12V_{IN}$ to $1.2V_{OUT}$ at 20A

Analog Devices' proprietary Silent Switcher 2 architecture self-cancels EMI via integrated hot loop capacitors, minimizing noisy antenna size. This, combined with integrated MOSFETs, significantly reduces switching node ringing and associated energy stored in the hot loop, even with very fast switching edges. The result is exceptional EMI performance while minimizing the AC switching losses. Silent Switcher 2 is incorporated in the LTC7150S to minimize EMI and deliver high efficiency, greatly simplifing EMI filter design and layout, ideal for noise-sensitive environments. LTC7150S passes the CISPR22/32 conducted and radiated EMI peak limits with only a simple EMI filter in front. Figure 2b shows the radiated EMI CISPR22 test result.

## High Frequency, High Efficiency Fits Tight Space

Integrated MOSFETs, integrated hot-loop decoupling capacitors, built-in compensation circuit—all take the design complexity out of the system and minimize total solution size with circuitry simplicity and Silent Switcher architecture. Thanks to high-performance power conversion, LTC7150S delivers high current without the need for additional heat sinks or airflow. Unlike most solutions, both low EMI and high efficiency can be achieved at high frequency operation, ensuring small passive component size. Figure 3 shows a 2MHz solution, which uses a small 72nH inductor and all ceramic capacitors in a very low profile solution for FPGA and µP applications.

## Conclusion

The demand for more intelligence, automation, and sensing in industrial and automotive environments is resulting in a proliferation of electronic systems that require increasingly high performance power supplies. Low EMI has risen from an afterthought to a key power supply requirement—in addition to solution size, high efficiency, thermal efficiency, robustness, and ease-of-use. The LTC7150S meets stringent EMI demands in a compact footprint by incorporating Silent Switcher 2 technology. Integrated MOSFETs and thermal management features enable robust and reliable delivery of currents up to 20A continuously from input ranges up to 20V, with switching frequency ranges as high as 3MHz.

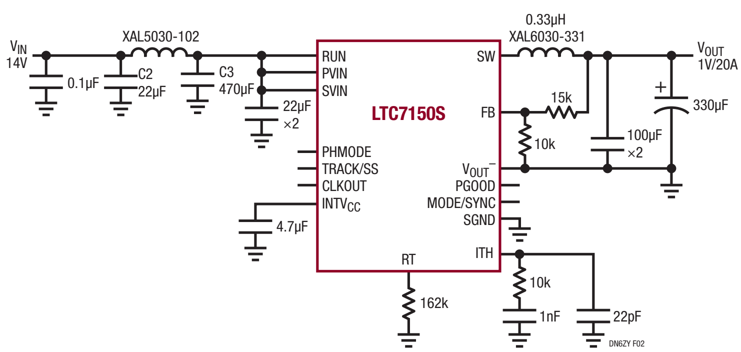


Figure 2(a) $V_{IN}$ = 14V, $V_{OUT}$ = 1V, 20A. $f_{SW}$ = 400kHz

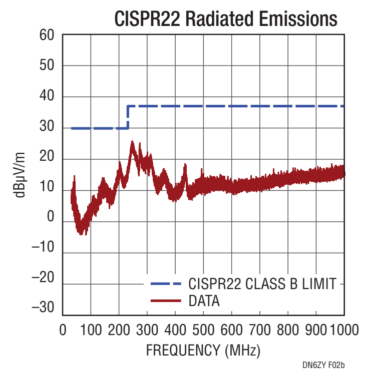


Figure 2(b) Radiated EMI Performance of Figure 2a

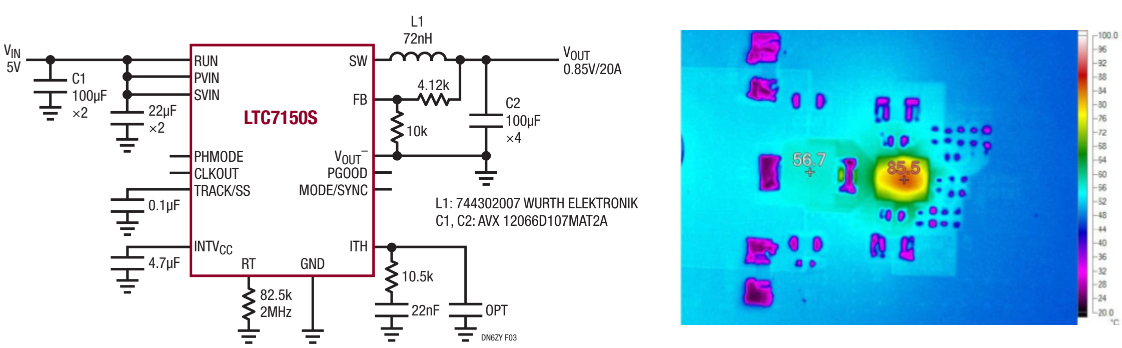


Figure 3. LTC7150S Schematic and Thermal Image for 5V Input to 0.85V/20A Output with $f_{SW}$ = 2MHz

**Data Sheet Download**
www.analog.com/LTC7150S

DN575-7/18


ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™
POWER BY LINEAR™


GEOS
Tough • Flexible • Safe
from Altech
by els spelsberg
Tough
UV and weather resistant
Corrosion resistant
Temperature resistant -35° to 80°C
Rugged
High impact resistance - IK 09
Flexible
Modular component mounting system
Lightweight
Easily machined
Multiple mounting and closure options
Safe
Drain Protect seal system
Air ventilation element option
High protection rating IP 67
Insulated protection
Safety agency approvals (UL, CSA)
High quality materials and an innovative design make the new GEOS series a universal solution and long lasting enclosure for indoor and outdoor applications.
c UL us
PROTECTION IP67
5 sizes (up to 15.75 x 19.69 x 8.9 in.)
Gray or Transparent cover
2 different depths
Features
All around protection through "Drain Protect"
Optimized box dimensions
Multiple mounting options
Modular component mounting system
Varied closure systems
Mounting standoffs in the cover
Optional Drain Knockout
Prevents Condensation
Optional Air Ventilation System
ENCLOSURE FINDER
AltechCorp.com/EnclosureFinder
Altech Corp.®
Your Source for Automation & Control Components
AltechCorp.com or 908.806.9400 for all the details and catalogs

That's where the IMU comes in. It provides backup position information for vineyard navigation. In addition, an IMU handles all other aspects of managing a robotic vehicle in an environment where vehicle position and angle are not only critical for safe maneuvering, but also for proper sensor suite functionality for the intended application of vine growth and health monitoring.

On top of that, this custom solution must work in the smallest and steepest of European vineyards—and it has to be cheap. Winemakers operate on some of the thinnest margins, and like many small farmers, are very vulnerable to catastrophic issues that can render a season unprofitable. That's why this new accessibility and cost-effectiveness in the latest generation of IMUs is critical. It enables solutions to be developed for spaces that previously couldn't afford such a sophisticated solution *(Fig. 1)*.

Such a useful solution previously could not be created, because navigational systems like the IMU were too expensive. With a good IMU now available for well under $100, the scope of application spaces that can develop their own targeted robotic vehicle solution has expanded to apply to anyone with a reasonable amount of resources.


**2. Modern IMU units can be very compact—as small as 24 × 37 × 9.5mm—for easy integration into a wide range of design projects.**

An IMU determines an AV's speed, angle, and direction, often adding magnetic-field information for added directional input. This task is done by using an integrated suite of accelerometers, gyroscopes, and sometimes magnetometers. An important task for the IMU is to provide baseline vehicle information to supplement GPS signals, to create a completely accurate snapshot of all the AV's motion and position characteristics.

But how do you know you have the right tool for your application? Many products are out there, and frankly you can probably pick almost any of them and have something that will work in your system for your application. However, an integral part of engineering is optimizing performance—using a solution that isn't the most efficient means, of course, you will never achieve the best performance.

To paraphrase Charles Dickens, the difference between the right solution and the almost-right solution is the difference between lightning and a lightning bug.

### APPLICATION DRIVEN

First things first. What kind of AV is going to use the IMU? Will it operate in a closed facility, or does it have to operate outdoors? What kind of accuracy do you need?

Beyond knowing orientation in three dimensions (three degrees of freedom), you should also know the additional dimensions of rotation, for six degrees of freedom. Then, once you add magnetometers, a device is considered by the industry to have nine degrees of freedom.

These requirements can all be easily addressed by reading the spec sheet and checking off the level of functionality you want. Most vendors and manufacturers offer several levels of product to enable scaling.

Here's where you decide on primary device characteristics—this is the critical stage. A good example would be magnetic sensing, a feature not available in all IMUs. If you buy an IMU with one and your project turns out to not need magnetic information, it's wasted time and money. It's almost as bad (but definitely not nearly as big a nuisance) as initially buying an IMU without magnetic sensing, and then later determining (usually due to changing customer requirements) that you really do need that sensing capability.

That's not to say that just because you can relatively easily satisfy this part of this selection process, it doesn't mean it shouldn't have the same level of diligence as for any other aspect of the design process. Here's where you also can weed out first-echelon needs like sensitivity and accuracy, culling the obviously inapplicable and creating a short list for further research.

### BODY ISSUES

Beyond the application, among the first concerns is the device footprint. Needless to say, if the device is too large for the space provided (and sometimes ironically, if it's too small), you will have significant issues integrating it into your product.

If you don't need a lot of robustness or protection from the elements, you can get IMUs with nine degrees of freedom in packages smaller than two square inches, and with protective packaging and coatings smaller than 16 square inches. By properly establishing all environmental and operational parameters, you can optimize size vs. performance and get the best solution *(Fig. 2)*.

Not only do you need to consider the device itself, but also the clearances for any cables for power and signal required. Something as relatively trivial to performance as whether the connectors are on the top, bottom, or side can make or break the selection decision.

Other than footprint, packaging issues involve materials (cost and weight), any coatings for ruggedness or water resistance (cost, weight and size), and mounting requirements. Accessibility may or may not be an issue, based on desired operational lifetime and maintenance requirements.

## POWER PLAY

The other primary consideration involves the issue of power. Whether the IMU is going into a large wheeled or winged AV with big batteries, or a small quadcopter with barely enough power to keep itself in the air, does make a critical difference.

The good news here is that the latest IMUs integrate very efficient (and small) microelectromechanical systems (MEMS)-based sensors. That means they will be both powerful and accurate, while consuming less power than legacy gyroscopes and electromechanical systems.

When selecting an IMU, you're going to want one that draws very little power, while accepting as wide a voltage as possible to ease system integration. Depending on your battery configuration and other subsystem needs, some IMUs may not be able to cost-effectively interface with your power bus.

## THE INTERNALS

Of course, an IMU is only as good as the sensors inside of it. How good is the gyro performance? The magnetic sensors? The latest IMUs have advanced MEMS-based inertial sensors, and 3-axis magnetic capability, while some devices may either use legacy sensing technology or low-grade MEMS *(Fig. 3)*.

Software is an often-overlooked aspect of the IMU selection process. If the communication protocols aren't up-to-date, or the system interfaces aren't effectively implemented, you will never achieve your solution's potential. The right software infrastructure should provide a simple and clean graphic interface to display, record, playback, and analyze all of an IMU's system parameters.

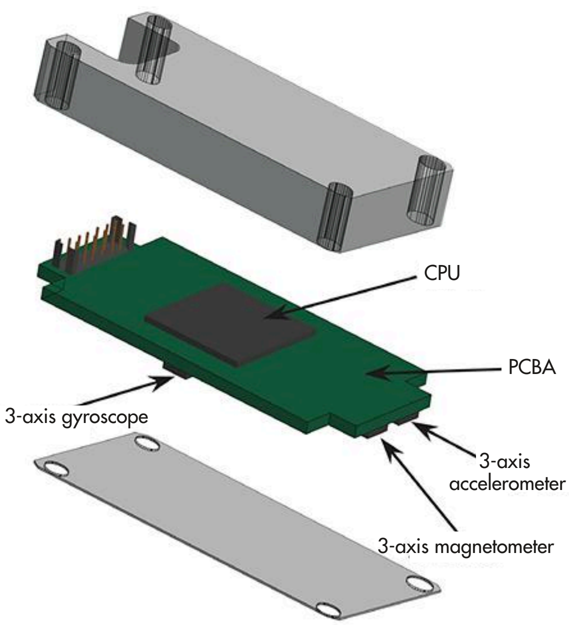


**3. Designed for complex autonomous applications, the latest IMUs have advanced MEMS-based inertial sensors (3-axis accelerometer and 3-axis gyroscope), as well as 3-axis magnetic capability.**

Having good software also expands the capabilities of the system. For example, an advanced solution with user-configurable settings to optimize system performance can address highly dynamic applications more effectively than a system with a limited number of parameters.

## UNDER A CLOUD

This could be considered a part of the first step, but once you have your communication protocols set up and your signal bus established, the question shifts to what level of integration and functionality you desire in the system. That can be dealt with after you determine your form factor and figure out how it will be powered.

Part of specifying the desired IMU is also to determine what peripheral functionality it will provide. At the minimum, the IMU will be a subsystem that tells the AV its location, where it's pointed, and how fast it's moving that way (and which direction on the compass, if you have magnetic functionality).

The IMU can also provide the vehicle systems with sensor data for basic operation. Many AVs are designed for use in upgraded consumer or commercial vehicles, which often have their own sensors systems for stability and roll detection, so the IMU serves the navigation system. If you're creating a vehicle from scratch, however, you can use the IMU to replace many of the sensors previously required for that job, simplifying your system and reducing your bill of materials (BOM).

Speaking of the BOM, the recent significant drop in the cost of an IMU and other advanced subsystems has also created a shopper's opportunity for design engineers. There's lots of competition in the space right now, so it's as almost as important to evaluate costs as it is to evaluate performance.

This is especially critical when looking to replace legacy devices—it's easy to be pleasantly surprised at the cost of the new stuff to the point you don't perform due diligence in pricing. This also leaves you open to being second-guessed by procurement when they swap your selection out for a less costly unit themselves.

## LOOKING FORWARD

Whenever you design a complex system, success or failure is often determined by how well you can integrate the various subsystems involved. Specifying the right device for the task, especially for critical navigation subsystems, will help ensure development of an optimal design. This is especially critical when addressing an application space as unforgiving of inadequacy as autonomous vehicles. ed

MIKE HORTON is the CTO of ACEINNA, which provides MEMS-based sensing solutions. Horton is responsible for corporate technology strategy and inertial-navigation-related technology development. Prior to ACEINNA, he founded Crossbow Technology, a leader in MEMS-based inertial navigation systems and wireless-sensor networks.

BAKUL DAMLE | Business Director of Battery Management in the Mobile Group
NAZZARENO ROSSETTI | PhD EE, Maxim Integrated Products Inc.
www.maximintegrated.com

# Advanced Battery Management for the Rest of Us

**A disruptive approach to battery modeling is critical for an effective fuel-gauge system to make it approachable, cost-effective and accurate.**

Battery management, like military radars and supersonic jets, is a sophisticated technology that's been out of reach for engineers who don't have special clearance or deep pockets—until now.

Charging and fuel gauging—at the heart of every battery-management system—are critical components of any mobile or IoT application. Battery performance relies on a high-quality battery model to drive the fuel-gauging algorithm. Extraction of the right model for the chosen battery is difficult, expensive work. In fact, only a few large manufacturers have the resources to develop such models.

The accessibility of accurate models can be a huge barrier to the proliferation of low-volume portable applications. This article presents a disruptive approach to cracking the battery-management barrier with a quick, cost-effective, high-performance solution.

### FUEL GAUGING FOR THE FEW

The release of energy from a battery adds up to nothing less than a controlled explosion. The energy stored in the battery (capacity) depends on load and temperature. Hence, developers recognize the importance of characterizing the battery under various conditions. Once a model tuned to the battery behavior is extracted, it's loaded into the fuel-gauge chip. This closely supervised process results in accurate gauging and safe battery charging and discharging.

IC vendors have traditionally focused on high-volume applications, since a few weeks of lab work are necessary for model extraction. This time-consuming, customized work yields maximum battery performance, in cases such as minimizing the state-of-charge (SOC) error and correctly predicting if the battery is nearly empty. *Figure 1* illustrates a battery SOC.

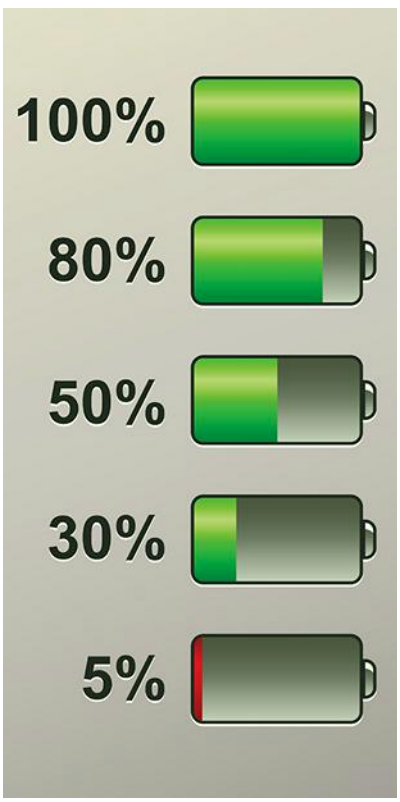


**1. It can be time-consuming to maximize battery performance and minimize the state-of-charge (SOC) error while correctly predicting if the battery is nearly empty.**

## FUEL GAUGING FOR THE MANY

By studying the characteristics of common lithium batteries, it's possible to develop an algorithm capable of handling the majority of these batteries. Such an algorithm can use a battery model tuned to the specific application and be embedded into the fuel-gauge ICs. Subsequently, designers are able to generate the battery models themselves using a simple configuration wizard included in the evaluation kit software. The system designer needs to provide only three pieces of information:

1. Capacity (often found on the label or datasheet of the battery)
2. The voltage per cell, considered the empty point for the battery (application-dependent)
3. Whether the battery-charge voltage is above 4.275 V (per cell, in case of multiple series cells)

With such an algorithm, the system designer's characterization work, already performed by the manufacturer, essentially disappears. Assuming a system error budget of 3% in the prediction of the SOC, the algorithm's model should cover 97% of the test cases.

In addition, the algorithm should include several adaptive mechanisms that help the fuel gauge learn about the battery characteristics to improve its accuracy even more. One such mechanism guarantees that the fuel-gauge output converges to 0% as the cell voltage approaches the empty voltage, thus reporting 0% SOC at the exact time that the cell voltage reaches empty.

For many users, just knowing SOC or remaining capacity (mAhr) isn't sufficient. They really want to know how much run-time they can get out of the residual charge. Simplistic methods such as dividing the remaining capacity by the present or future load can lead to unrealistically optimistic answers. The proposed algorithm should come to the rescue, providing an accurate time-to-empty register based on battery parameters, temperature, and load effects, as well as the empty voltage of the application.

The proposed algorithm's advantages are obvious. High-volume manufacturers can use it as a starting point to get development going even before selecting the right battery for the application. They can then turn to a finely tuned battery model at a more mature stage of the development. The small-volume manufacturer can use it to define the model of the battery and run with it, having the confidence that most batteries out there will be compatible.

Such an algorithm is built into ModelGauge m5 EZ ultra-low power, standalone fuel-gauge ICs developed by Maxim Integrated.

## FUEL GAUGING FOR ALL

A logical next step to simplify a battery system's design is to put it on an Arduino form factor shield *(Fig. 2)*. Maxim's MAXREFDES96 IoT Power Supply is an Arduino form factor shield, powered by a 660-mAh Li-ion battery for running an untethered Arduino system *(Fig. 3)*. The design features both of Maxim's battery-management technologies: a highly integrated battery charger and a ModelGauge m5 EZ fuel gauge. Additional Maxim devices provide system-management and power-supply functions.

These technologies boost charging and improve the accuracy of battery-capacity feedback to help maximize performance. The MAXREFDES96 may be charged through the host board USB connector, the host board power adapter, or via the on-board power connector. Furthermore, the system accommodates any single cell Li-ion battery, allowing for development with batteries of different sizes from different manufacturers.

**An untethered Arduino board is appealing to an extended range of applications and users, including hobbyists and enthusiasts.**

Free firmware supports operation with Arduino and mbed.org platform boards.

With the MAXREFDES96, the battery model can be directly stored in nonvolatile memory on the MAX17201, or on the Arduino board. In the latter case, the models are transferred from the board to the shield at power up. This arrangement allows for the use of a variety of batteries. And with the Arduino platform, many available software drivers can be employed.

An untethered Arduino board is appealing to an extended range of applications and users, including hobbyists and enthusiasts. The system can be leveraged for its portability and deployed in a remote location for data gathering; for its transportability by moving a running test between locations without interruption; or as a backup system in case the main power fails at a critical time of the development. In all cases, the system facilitates battery management, leading to quicker charging, more accurate fuel gauging, and longer battery life.

## CONCLUSION

We have highlighted the critical importance of battery modeling for an effective fuel gauge system, and discussed the barriers to obtain these models that impedes the proliferation of low-volume battery applications. A disruptive approach, based on the ModelGauge m5 EZ algorithm and the MAXREFDES96 Arduino shield, helps simplify and lower the cost of battery-system implementation—and is available to all. ed

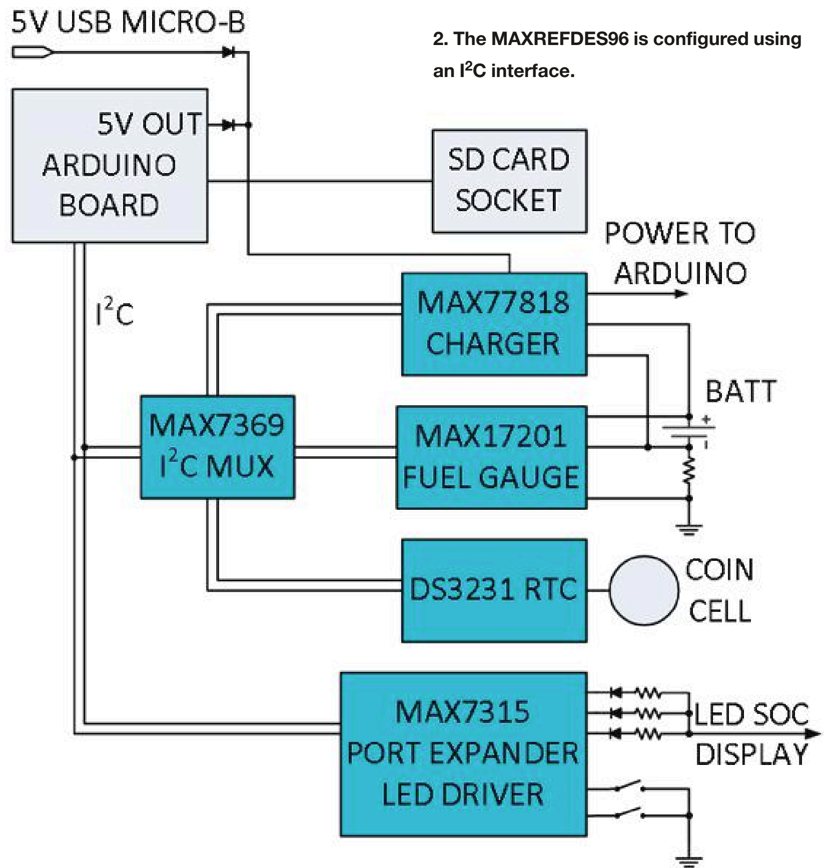


2. The MAXREFDES96 is configured using an I²C interface.

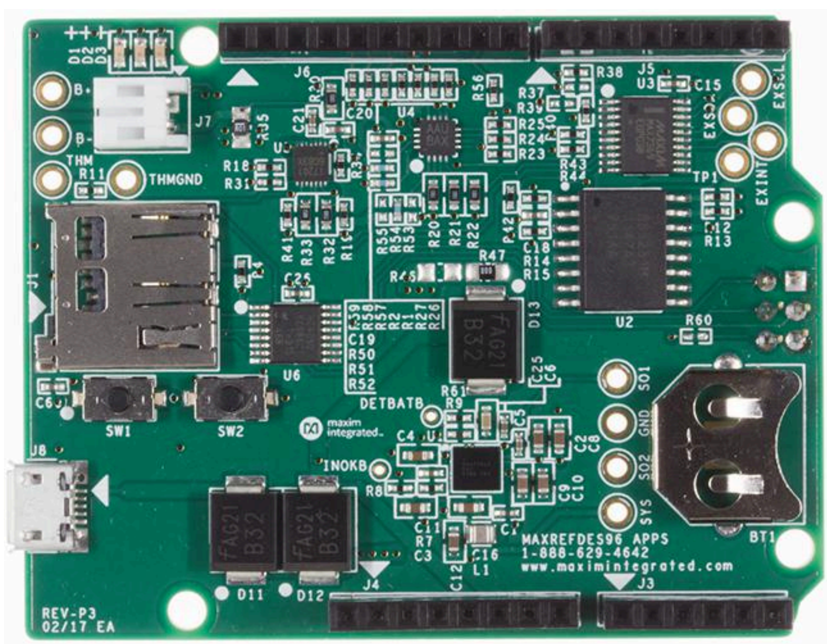


3. The MAXREFDES96 is an Arduino Uno R3-compatible board that provides battery charge, boost and data-logging capabilities.

AARON SHULTZ | Applications Engineering Manager, Analog Devices Inc.
PETER HAAK | Consultant

# Pocket-Sized White-Noise Generator Quickly Tests Circuit's Response

**Using a single op amp and a properly selected resistor, this white-noise generator provides flat output versus frequency and speeds testing of circuit performance.**

The behavior of analog circuits can often be characterized by sweeping their inputs with a signal going across a range of frequencies and then observing the output. These sweeps can be composed of discrete input frequencies or a swept sine.

For the first option, extremely low-frequency sine waves (below 10 Hz) are difficult to produce cleanly, and usually requires a processor, digital-to-analog converter (DAC), and some complex, precise filtering to produce relatively clean sine waves. Further, for each frequency step, the system must settle, making the sequential full sweep with many frequencies into a slow process. The faster option is to test at fewer discrete frequencies, but increases the risk of skipping over critical frequencies where high-Q phenomena reside.

In contrast, a white-noise generator is simpler and faster than a swept sine wave because it effectively produces all frequencies at the same time with the same amplitude. Imposing white noise at the input of a device under test (DUT) can quickly produce an overview of the frequency response over an entire frequency range. In this case, there's no need for expensive or complex swept sine-wave generator—simply connect the DUT output to a spectrum analyzer and watch. Using more averaging and longer acquisition times produces a more accurate output response across the frequency range of interest.

The expected response of the DUT to white noise is frequency-shaped noise. Using white noise in this fashion can quickly expose unexpected behavior such as weird frequency spurs, strange harmonics, and undesirable frequency-response artifacts. Furthermore, a white-noise generator allows a careful engineer to test a tester. Lab equipment that measures frequency response should produce a flat noise profile when measuring a known, flat white-noise generator.

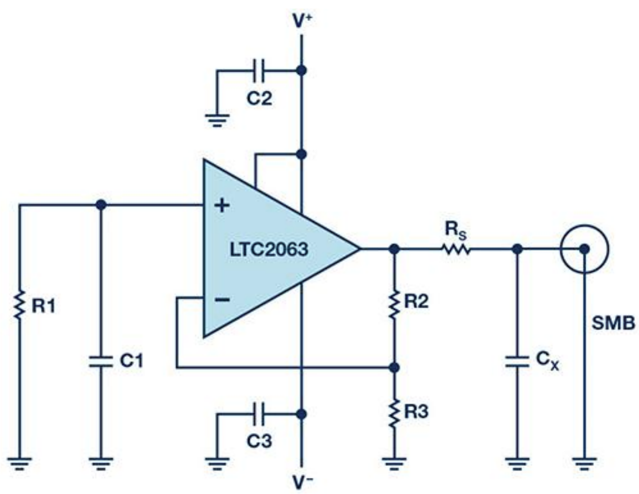


**1. In this white-noise generator, the low-drift micropower LTC2063 op amp amplifies the Johnson (thermal) noise of R1.**

On the practical side, a white-noise generator is easy to use, small enough for compact lab setups, portable for field measurements, and inexpensive. Quality signal generators with myriad settings are attractively versatile. However, versatility can hamper quick frequency-response measurements. A well-designed white-noise generator requires no controls, yet produces a fully predictable output.

## NOISY DISCUSSION

Resistor thermal noise, sometimes called Johnson noise or Nyquist noise, arises from thermal agitation of charge carriers inside a resistor. This noise is approximately white, with nearly Gaussian distribution. In electrical terms, the noise voltage density $V_{NOISE}$ is given by:

$V_{NOISE} = \sqrt{(4k_BTR)}$

where $k_B$ is the Boltzmann's constant, T is the temperature in Kelvin, and R is the resistance. Noise voltage arises from the random movement of charges flowing through the basic resistance, a sort of R × $I_{NOISE}$. *Table 1* shows examples at 20°C.

A 10-MΩ resistor, then, represents a 402 nV/√Hz wideband voltage-noise source in series with the nominal resistance. A gained up, resistor-derived noise source is fairly stable as a lab test-noise source, as R and T variations affect noise only by the square root. For instance, a change of 6°C from 20°C is a change of 293 kΩ to 299 kΩ. Because noise density is directly proportional to the square root of temperature, a change of 6°C temperature leads to a relatively small 1% noise density change. Similarly, with resistance, a 2% resistance change leads to a 1% noise density change.

In *Figure 1*, a 10-MΩ resistor R1 generates white, Gaussian noise at the positive terminal of an op amp, while resistors R2 and R3 increase the gain of the noise voltage to the output. Capacitor C1 filters out chopper-amplifier charge glitches. The output is a 10-µV/√Hz white-noise signal.

The gain (1 + R2/R3) is high at 21 V/V in this example. Even if R2 is high (1 MΩ), the noise from R2 compared to the gained-up R1 noise is inconsequential.

An amplifier for the circuit must have sufficiently low input-referred voltage noise so as to let R1 dominate as the noise source. The reason is that the resistor noise should dominate the overall accuracy of the circuit, not the amplifier. An amplifier for the circuit must have sufficiently low input-referred current noise to avoid ($I_{IN}$ × R2) to approach (R1 noise × gain) for the same reason.

### HOW MUCH AMPLIFIER VOLTAGE NOISE IS ACCEPTABLE IN THE WHITE-NOISE GENERATOR?

*Table 2* shows the increase in noise from adding independent sources. A change from 402 nV/√Hz to 502 nV/√Hz is only 1.9 dB in log volts, or 0.96 power dB. With op-amp noise at ~50% of the resistor noise, a 5% uncertainty in op-amp $V_{NOISE}$ changes the output noise density by only 1%.

A white-noise generator could use only an op amp without a noise-generating resistor. Such an op amp must exhibit a flat noise profile at its input. However, the noise voltage is often not accurately defined and has a large spread over production, voltage, and temperature. Other white-noise circuits may operate based on a Zener diode with far-less predictable characteristics. Finding an optimal Zener diode for stable noise with microamps of current can be difficult, however, particularly at low voltage (<5 V).

Some high-end white-noise generators are based on a long pseudorandom binary sequence (PRBS) and special filters. Using a small controller and DAC may be adequate; however, making sure that the DAC doesn't produce settling glitches, harmonics, or intermodulation products is something for experienced engineers. Moreover, choosing the most appropriate PRBS sequence adds complexity and uncertainty.

### A LOW-POWER ZERO-DRIFT SOLUTION

Two design goals dominated this project:

- An easy-to-use white-noise generator must be portable and most likely battery-powered, which means micropower electronics.
- The generator must provide uniform noise output even at low frequencies, below 0.1 Hz and beyond.

The prototype *(Fig. 2)* and the layout *(Fig. 3)* show the compact generator implementation that was achieved.

**2. The prototype of the pocket-size white-noise generator shows its compact layout, both a convenience as well as a necessity for effective operation.**

**TABLE 1: NOISE VOLTAGE DENSITY OF VARIOUS RESISTORS**

| Resistor | Noise voltage density |
|---|---|
| 10 Ω | 0.402 nV/√Hz |
| 100 Ω | 1.27 nV/√Hz |
| 1 kΩ | 4.02 nV/√Hz |
| 10 kΩ | 12.7 nV/√Hz |
| 100 kΩ | 40.2 nV/√Hz |
| 1 MΩ | 127 nV/√Hz |
| 10 MΩ | 402 nV/√Hz |

**TABLE 2: CONTRIBUTION OF OP-AMP NOISE**

| $R_{NOISE}$ (nV/√Hz) | Amp $e_n$ | Total input referred |
|---|---|---|
| 402 nV/√Hz | 300 | 501.6 nV/√Hz |
| 402 nV/√Hz | 250 | 473.4 nV/√Hz |
| 402 nV/√Hz | 200 | 449.0 nV/√Hz |
| 402 nV/√Hz | 150 | 429.1 nV/√Hz |
| 402 nV/√Hz | 100 | 414.3 nV/√Hz |

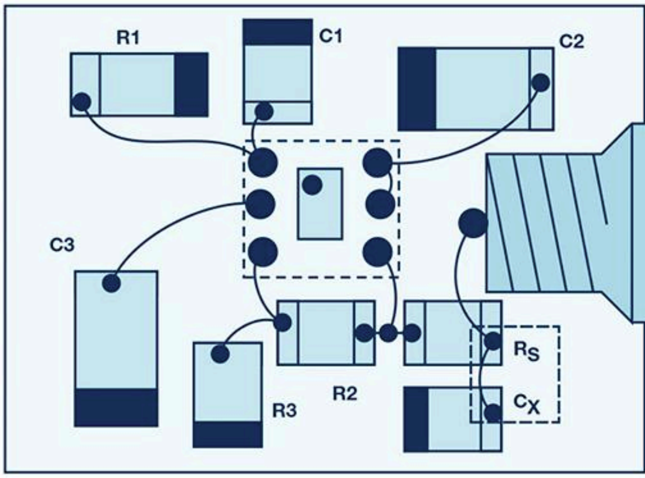


Note: Blue Fill SMD Terminals Are Grounded
R1 = 10 MΩ MMA0204 (MiniMELF)
Vishay/Beyschlag 1% TC50 (= Thin Film)
R2 = 1 M, R3 = 49900 Ω; $R_S$ = 10 kΩ; All 1% TC100 Thick Film
C1 = 22 pF C0G 5%; C2/C3 = 0.1 µF C0G
$C_X$ = 47 nF C0G 5% (See Text: "Optional Tuning")

**3. The schematic tells only part of the generator story, as the layout affects—and is affected by—EMI considerations.**

The LTC2063, a low-power, zero-drift op amp, met these requirements and other critical constraints. The noise voltage of a 10-MΩ resistor is 402 nV/√Hz; that of the LTC2063 is roughly half. The noise current of a 10-MΩ resistor is 40 fA/√Hz; the LTC2063's is less than half. The LTC2063 fits neatly into a battery application due to its typical supply current of 1.4 µA, while the supply can go down to 1.7 V (rated at 1.8 V). This is important since low-frequency measurements by definition require long settling times; therefore, this generator must remain powered by a battery for extended periods of time. The noise density of the LTC2063 input is roughly 200 nV/√Hz, and noise is predictable and flat over the frequency range (within ±0.5 dB). Assuming that the LTC2063's noise is 50% of thermal noise and op-amp voltage noise changes 5%, output noise density changes only 1%.

Zero-drift op amps don't have zero 1/f noise by design. Some are better than others and, especially for current noise, it's more common that the wideband specification is wrong or that 1/f noise is much higher than suggested in the datasheet. For some zero-drift op amps, the datasheet noise plot doesn't go down to the MHz-frequency region, possibly masking 1/f noise.

A chopper-stabilized op amp could be a solution to keep the noise flat at very low frequency. However, the high-frequency "noise bump" and switching noise must not spoil the performance. The data shown here supports the use of LTC2063 in the face of these challenges.

## CIRCUIT DESCRIPTION AND IMPLEMENTATION DETAILS

Thin-film resistor R1 (Vishay/Beyschlag MMA0204 10 MΩ) generates most of the noise. The MMA0204 is one of few 10-MΩ options to combine high quality with low cost. In principle, R1 could be any 10-MΩ device as signal current is very small, so 1/f noise can be neglected. It's best to avoid low-cost thick-film chips of questionable accuracy or stability for the primary element of this generator. For best accuracy and long-term stability, R2, R3, or $R_S$ should be 0.1% thin film—for example, TE CPF0603. Capacitors C2 and C3 can have almost any dielectric; C0G (NPO) dielectric devices can be used to guarantee low leakage current.

For best EMI rejection, the loop area formed by R1/C1/R3 should be minimized. In addition, R1/C1 should be very well shielded from electrical fields (discussed further below in the *EMI Considerations* section). Although not critical, R1 should be shielded from large temperature changes. With good EMI shielding, thermal shielding is often adequate.

The LTC2063 rail-to-rail input-voltage transition region of the common-mode voltage $V_{CM}$ range should be avoided, as crossover may result in higher, less-stable noise. For best results, use at least 1.1 V for V+ with the input at 0-V common mode.

Note that a value for $R_S$ of 10 kΩ may seem high, but the micropower LTC2063 presents a high output impedance; even 10 kΩ doesn't fully decouple the LTC2063 from load capacitance at its output. For this white-noise generator circuit, some output capacitance that leads to peaking can be a design feature rather than a hazard. The output sees 10-kΩ $R_S$ and a 50-nF $C_X$ to ground. This capacitor $C_X$ will interact with the LTC2063 circuit, resulting in some peaking in the frequency response. Such peaking can be used to extend the flat bandwidth of the generator, in much the same way that port holes in loudspeakers attempt to expand the low end. A high-impedance load is assumed (>100 kΩ), as a lower-Z load would significantly reduce the output level and may affect peaking.

## OPTIONAL TUNING

Several IC parameters (for example, $R_{OUT}$ and gain-bandwidth) affect flatness at the high-frequency limit. Without access to a signal analyzer, the recommended value for $C_X$ is 47 nF, which typically yields 200 Hz to 300 Hz (–1 dB) bandwidth. Nevertheless, $C_X$ can be optimized for either flatness or bandwidth, with $C_X$ = 30 nF to 50 nF as typical values. For wider bandwidth and more peaking, use a smaller $C_X$; for a more damped response, use a larger value. Critical IC parameters are related to op-amp supply current. Parts with low supply current may require a somewhat larger $C_X$, while parts with high supply current most likely require less than 30 nF while achieving wider flat bandwidth.

Plots highlight how $C_X$ values affect closed-loop frequency response. Output noise density vs. $C_X$ (at $R_S$ = 10-kΩ, ±2.5-V supply) is shown in *Figure 4*. The output RC filter is effective in eliminating clock noise. The plot shows output vs. frequency for $C_X$ = 0 and $C_X$ = 2.2 nF, 10 nF, 47 nF, and 68 nF.

Using $C_X$ = 2.2 nF results in mild peaking, while peaking is strongest for $C_X$ = 10 nF, gradually decreasing for larger $C_X$. The trace for $C_X$ = 68 nF shows no peaking, but it has visibly lower flat bandwidth. The best result is for $C_X$ ~ 47 nF; clock noise is three orders of magnitude below signal level. Due to limited vertical resolution, it's impossible to judge with fine precision the flatness of output amplitude vs. frequency. This plot was produced using ±2.5-V battery supply, though the design allows the use of two coin cells (about ±1.5 V).

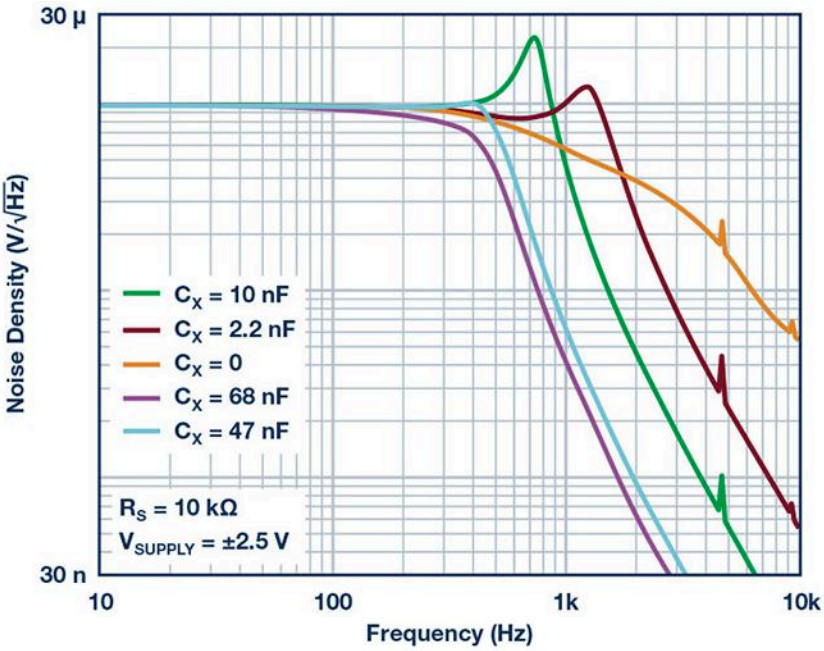


**4. Output noise density of the basic design shows peaking with various values of $C_X$.**

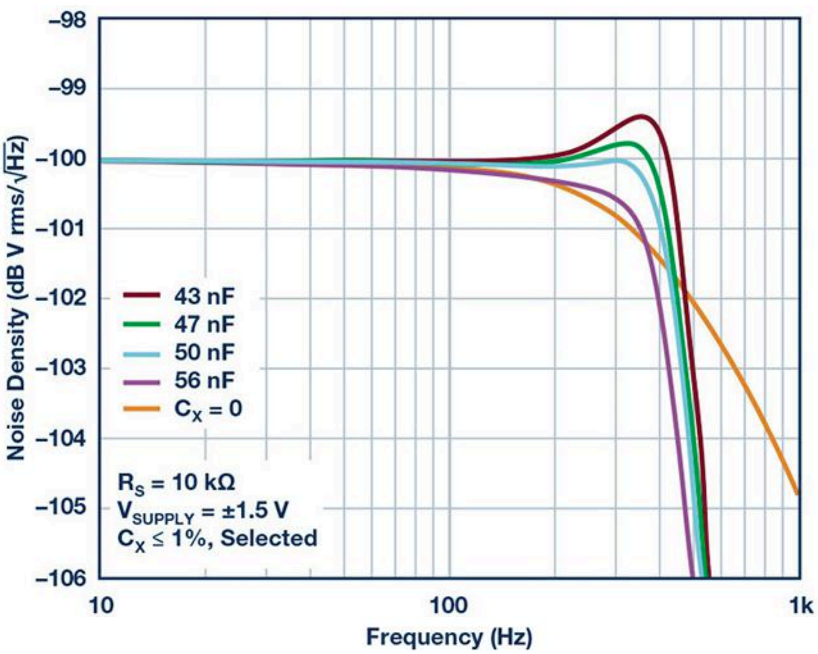


**5. This "zoomed-in" view of output noise density highlights the achieved flatness.**

*Figure 5* shows flatness magnified on the Y-axis. For many applications, flatness within 1 dB is enough to be useful and <0.5 dB is exemplary. Here, $C_X$ = 50 nF is best ($R_S$ = 10 kΩ, $V_{SUPPLY}$ ±1.5 V); $C_X$ = 45 nF, although 55 nF is acceptable.

High-resolution flatness measurements take time; this plot from 10 Hz to 1 kHz with 1000 averages requires about 20 minutes per trace. The standard solution uses $C_X$ = 50 nF. The traces shown for 43 nF, 47 nF, and 56 nF, all $C_X$ < 0.1% tolerance, show a small but visible deviation from best flatness. The orange trace for $C_X$ = 0 was added to show that peaking increases flat bandwidth (for Δ = 0.5 dB, from 230 Hz to 380 Hz).

Using two 0.1-µF C0G capacitors in series is probably the simplest solution for an accurate 50-nF value. A 0.1-µF C0G 5% 1206 component is easy to procure from Murata, TDK, and Kemet. Another option is a 47-nF C0G (1206 or 0805); this part is smaller but may not be as commonly available. As stated prior, optimum $C_X$ varies with actual IC parameters.

Flatness was also checked versus supply voltage *(Fig. 6)*. The standard circuit is ±1.5 V; changing the supply voltage to ±1.0 V or ±2.5 V shows a small change in peaking as well as a small change in the flat level (due to $V_N$ changing vs. supply, with thermal noise dominant). Both peaking and flat level change ~0.2 dB over the full range of supply voltage. The plot suggests good amplitude stability and flatness when the circuit is powered from two small batteries.

For this prototype with a ±1.5-V supply, flatness was within 0.5 dB up to approximately 380 Hz. At ±1.0-V supply, flat level and peaking slightly increase.

For ±1.5- to ±2.5-V supply voltage, the output level doesn't visibly change. Total V p-p (or $V_{RMS}$) output level depends on the fixed 10 µV/√Hz density, as well as on bandwidth. For this prototype, the output signal is ~1.5 mV p-p.

At some very low frequency (MHz range), noise density may increase beyond the specified 10 µV/√Hz. For this prototype, it was verified that at 0.1 Hz, noise density is still flat at 10 µV/√Hz. In stability versus temperature, thermal noise dominates. Thus, for T = 22(±6)°C, the amplitude change is ±1%, a change that would barely be visible on a plot.

**EMI CONSIDERATIONS**

The prototype uses a small copper foil with Kapton insulation as a shield. This foil, or flap, is wrapped around the input components (10 MΩ + 22 pF) and soldered to ground on the PCB backside. Changing the position of the flap has a significant effect on sensitivity to EMI and risk of low-frequency (LF) spurs. Experimentation suggests that LF spurs that occasionally show are due to EMI, and spurs can be prevented with very good shielding. With the flap, the prototype gives a clean response in the lab, without any additional mu-metal shielding. No mains noise or other spurs appear on a spectrum analyzer. If excess noise is visible on the signal, additional EMI shielding might be needed.

When an external power supply is used instead of batteries, common-mode current can easily add to the signal. It's recommended to connect the instrument grounds with a solid wire and use a CM choke in the supply wires to the generator.

**LIMITATIONS OF THE DESIGN**

There are always applications that require more bandwidth, such as the full audio range or ultrasound range. More bandwidth isn't realistic on a few microamps of supply current. With approximately 300 to 400 Hz of flat bandwidth, the LTC2063 resistor noise-based circuit could be useful to test some instruments for 50-/60-Hz mains frequency, perhaps geophone applications. The range is suitable for testing various VLF applications (for example, sensor systems), as the frequency range extends down to <0.1 Hz.

The output signal level is low (<2 mV p-p). A follow-on LTC2063 configured as a noninverting amplifier with a gain of five and further RC-output filter can provide a similarly well-controlled flat, wideband-noise output to 300 Hz with larger amplitude. If the closed-loop frequency range isn't maximized, a capacitor across the feedback resistor can lower the overall bandwidth. In this case, the effects of $R_S$ and $C_X$ will have less, or even negligible, effect at the edge of the closed-loop response.

The white-noise generator described here is a small but essential tool. With long measurement times—the norm for low-frequency applications—a simple, reliable, pocketable device that can produce near instantaneous circuit characterization is a welcome addition to the engineer's toolbox. Unlike complex instruments with numerous settings, this generator requires no user manual.

This particular design features low supply current, essential for battery-powered operation in long-duration VLF application measurements. When supply current is very low, there's no need for on/off switches. A generator that works on batteries also prevents common-mode currents.

The LTC2063 low-power, zero-drift op amp used in this design is the key to meeting the constraints of the project. Its features enable use of a noise-generating resistor gained up by a simple, noninverting op-amp circuit.

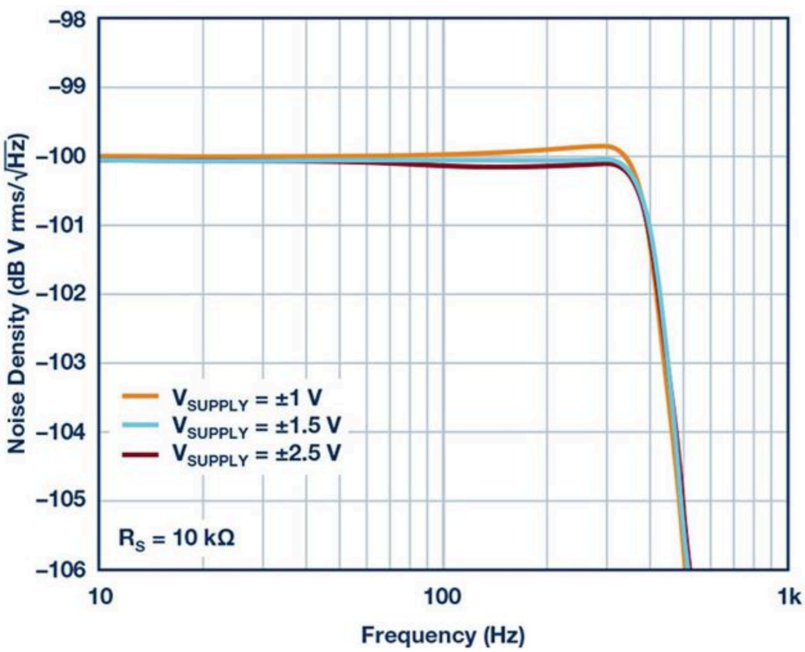


**6. Output noise density shows only a slight change in peaking for various supply voltages.**


AARON SCHULTZ is an applications engineering manager whose multiple system engineering roles in both design and applications have exposed him to topics ranging across battery management, photovoltaics, dimmable LED drive circuits, low voltage and high current dc-to-dc conversion, high-speed fiber-optic communication, advanced DDR3 memory R&D, custom tool development, validation, and basic analog circuits, while over half of his career has been spent in power conversion. PETER HAAK started in electronic development in 1986 and since 1993, he has worked as an independent consultant specializing in sensors and instrumentation. Peter lives in the center of 's-Hertogenbosch in the Netherlands next to St. Jan cathedral.

# Gated Counter Determines Which of Two Pulses Came First

DAVE CONRAD | Technical Consultant

**IN MANY APPLICATIONS,** there's a need to determine the difference in arrival time of two pulses, such as those derived from acoustic signals emanating from a single moving source. In this circuit, logic signals indicate the timing of the acoustic signals with good fidelity. (An analog front end, not shown here, shapes the analog signals and converts them to digital-logic signals.)

This circuit uses a dual-D flip-flop and an OR gate *(Fig. 1)* that provides a positive-polarity gate signal as well as an indication of which pulse came first, to be used by an MCU counter. Note that the OR gate could be eliminated if two MCU counters were used, one for each of the U1A Q and U1B Q outputs. In that case, only one of the two counters would accumulate counts, depending on which pulse arrived first and the time of arrival of the second pulse. Then, with additional coding, it would be possible to determine which pulse came first, and by how much.

This particular project required taking left, right, up, and down timing measurements, which would require four counters with gates—that's two more than the PIC MCU designated for the project could provide. Therefore, using the OR gate saves MCU resources.

$V_{left}$ and $V_{right}$ resting levels may be 1 or 0; but if 1, the timing between left and right pulses will be measured on the trailing edge of the pulses. This may induce errors if the pulse width of the trigger inputs is inconsistent. As the rising edge of $V_{left}$ clocks a logic 1 into U1A Q (L_R_Select), the U1B flip-flop is reset, thus forcing U1B Q to logic 0. A later pulse from the $V_{right}$ source reverses this condition, ending the counter interval. *Figure 2* shows the waveform timing for "right came first," while *Figure 3* shows the opposite case of "left came first."

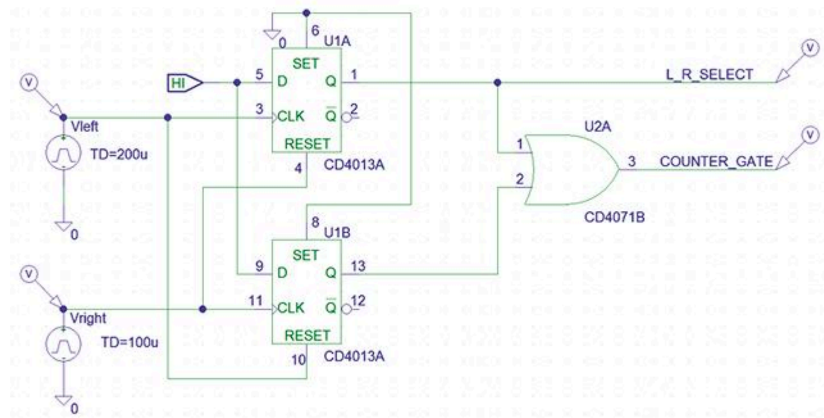


**1. By controlling timer count, this dual flip-flop circuit can determine which of two pulses arrived first.**

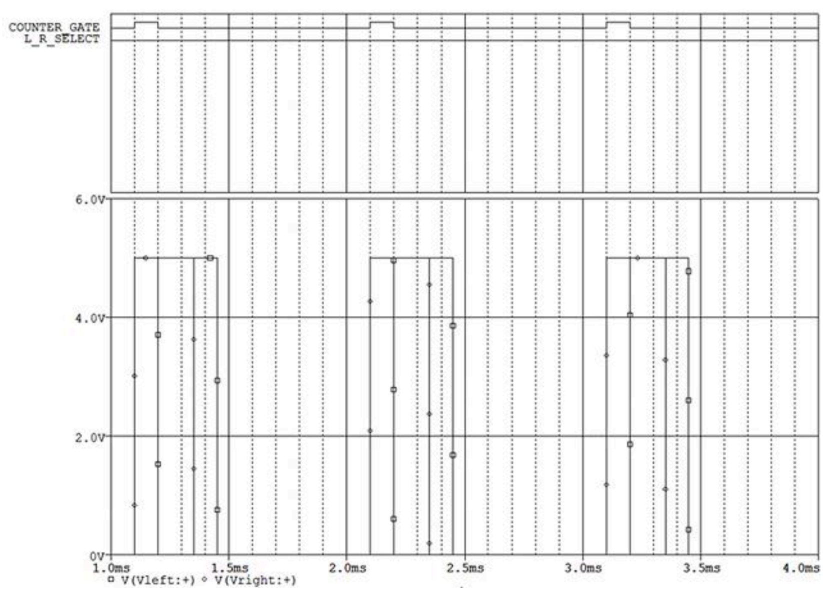


**2. The $V_{right}$ pulse arrives 100 μs before the $V_{left}$ pulse. Note that L_R_Select is low. If this signal is sampled by the MCU just after the gate signal goes high, this indicates right-before-left.**

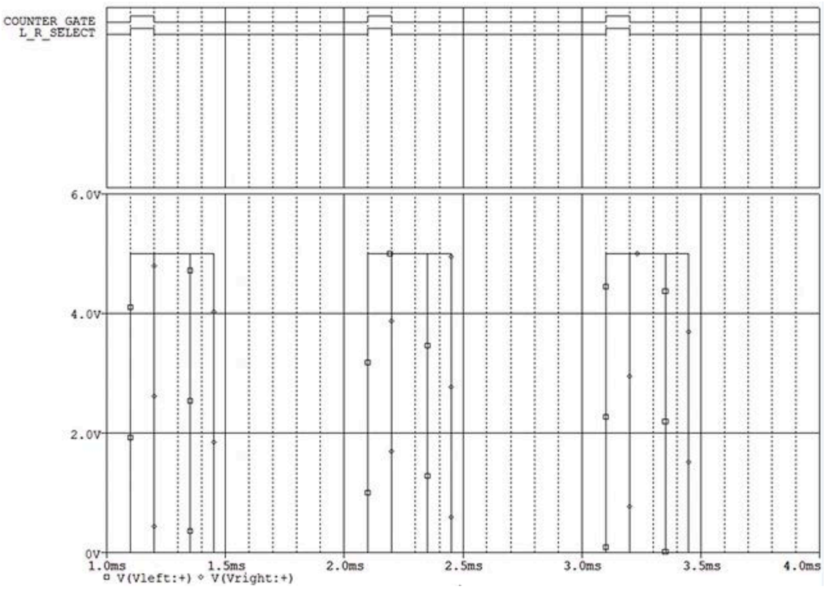


**◀ 3. The $V_{left}$ pulse arrives 100 µs before the $V_{right}$ pulse. Because the L_R_Select is high, if this signal is sampled by the MCU just after the gate signal goes high, this indicates left-before-right.**

In the case of simultaneously arriving pulse edges *(Fig. 4)*, there will be no gate output, which will result in zero counts accumulated by the PIC counter. The resolution around zero is limited by the rise and fall times of the flip-flop logic family used. Though the 4000-series logic family is fairly slow, it was adequate for this project. Faster logic families will allow for more resolution around zero, if that's a requirement.

The clock frequency can also limit resolution; the maximum count at 500 kHz is 110. The overall span is double that as the sound source moves from left to right, relative to the center line of the two sensors and the display. Higher clock frequencies may require a 16-bit or greater counter, depending on the maximum difference in time of arrival; this project uses a 500-kHz clock for 2-µs resolution. The maximum time difference to be measured is defined by the time of flight of the acoustic pulse in air over the distance from one sensor to another. Three inches at a wave velocity of 1135 feet per second (fps) limits the difference in arrival times to 220 µs.

Some pseudocode *(Fig. 5)* illustrates the method for using the circuit output to indicate position or motion, where PortC.0 is the Counter_Gate signal and PortC.2 is the L_R_Select signal, also known as LR_FLAG.

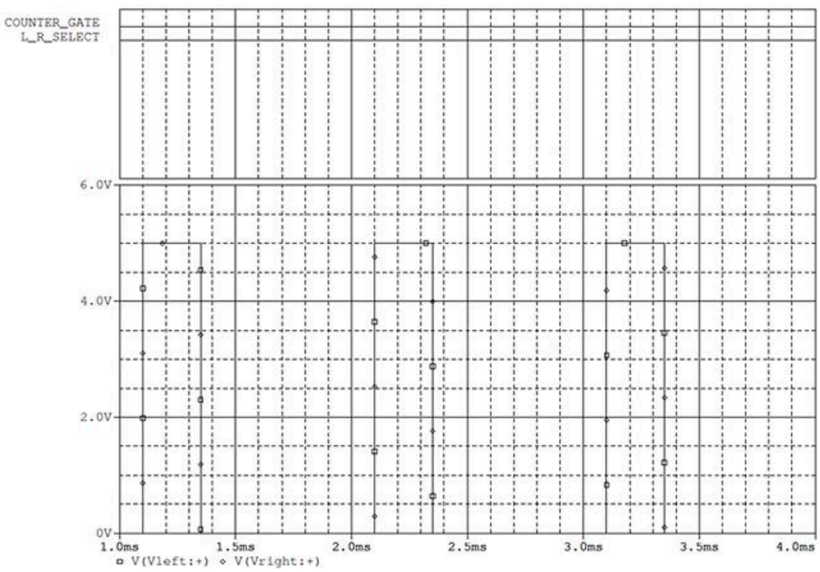


**4. If both signals arrive with the "simultaneous" window, zero counts are accumulated by the PIC counter.**

```
DO WHILE PORTC.0 = 0  'WAIT FOR PORTC.0 TO GO HIGH, THEN READ LEFT GATE STATUS
LOOP

   LR_FLAG = PORTC.2 'READ THE LEFT GATE, IF 1, IS LEFT, IF 0 IS RIGHT

DO WHILE PORTC.0 = 1  'POLL C.0 UNTIL GATE IS LOW (COUNT IS CAPTURED)
LOOP              'GATE IS CLOSED, READ TIMER THEN ZERO IT.

   X = TMR1L       'READ TIMER1 LOW BYTE REGISTER, 120 IS MAX COUNT
   TMR1L = 0       'ZERO TMR1 REGISTER
   TMR1H = 0
   IF LR_FLAG = 0 THEN X = X + 119 'OFFSET FOR RIGHT MOTION, 119 IS THE CENTER OF THE DISPLAY.
   IF LR_FLAG = 1 THEN X = 119 - X 'OFFSET FOR LEFT MOTION
      ENDIF
```

**◀ 5. This code shows how the dual FFT timing signals can be used to indicate left or right motion, relative to the center of a display.**

DAVE CONRAD, a retired electronics engineer, creates innovative light-show systems using LED or video displays, and also does some contract consultation. He can be reached at dconrad992@gmail.com.

## 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

# Maxim's Wearable Health Sensor Platform Adds a New Dimension

**ECG, heart rate, and skin temperature all can be monitored with the latest version of the company's Health Sensor Platform.**

Maxim Integrated Products' first Health Sensor Platform (HSP) made it much easier to build wearable devices. However, the latest version—HSP 2.0 *(Fig. 1)*—lets developers monitor electrocardiogram (ECG), heart rate, and body temperature. The modular system is designed to handle upgrades or enhancements, which in the long run can save developers months of design and implementation work.


**1. Maxim's Health Sensor Platform 2.0 (left) can monitor ECG, heart rate, and skin temperature. The device can be taken apart (right) to enable changes or upgrades to the sensor system.**

The $399 HSP 2.0 contains a number of Maxim Integrated's components, starting with the Darwin low-power MAX32630 ARM Cortex-M4F microcontrollers. The platform is tied into a MAX32664 sensor hub *(Fig. 2)* that handles a range of sensors, including the MAX86141—an optical sensor that tracks heart rate as well as being a pulse oximeter. The MAX30001 analog front end (AFE) handles the ECG sensors, while the MAX30205 handles temperature to within 0.1°C. Maxim also covers power management with its MAX20303 PMIC, which is critical for such mobile devices.

In addition, the HSP 2.0 uses third-party chips and modules, including the Bluetooth support. The system is charged and managed with a USB Type-C connection. The compact module has an expansion header and display interface and contains the processor and Bluetooth support *(Fig. 3)*.

ARM Mbed supports the MAX32630 microcontroller. This provides access to a large, open-source library of tools as well as a development and IoT communication framework. System software support for over-the-air updates is in the works. The MAX32630 also includes a secure element that can provide encryption and secure key support. An Android application is also provided that allows for remote viewing of data.

The platform doesn't have FDA approval as a medical device, but a developer could use it as the basis for such a product. It's a matter of getting the appropriate approvals so that data could be utilized by medical professionals. Of course, the device can be used to acquire and provide information to users.

Maxim Integrated has put together an amazing package that opens opportunities for professional developers and makers. There's no comparable platform available elsewhere. Its modular approach enables future hardware upgrades as well as custom sensor arrays without incurring major redesign costs. ed

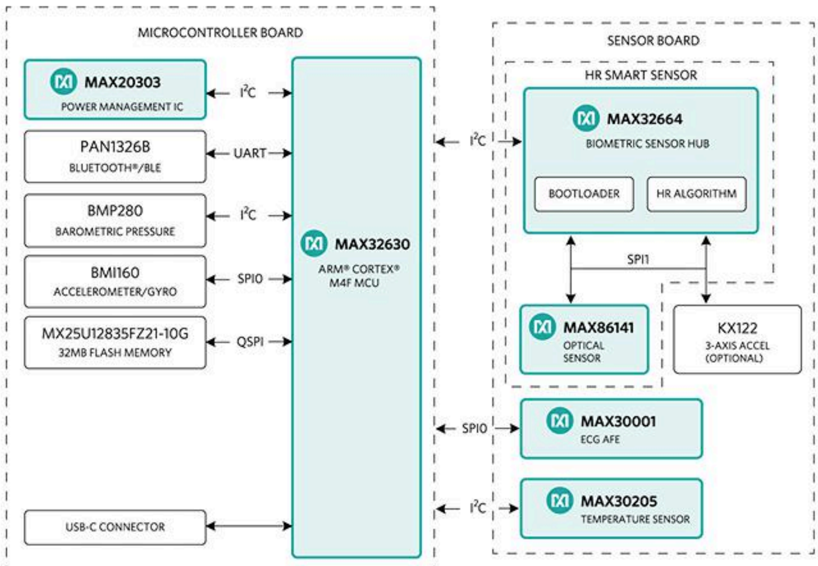



**2. The HSP 2.0 is built around Maxim's MAX32630 ARM Cortex-M4F tied to a MAX32664 sensor hub that handles a range of sensors.**

**3. The processor module is compact with an expansion header.**

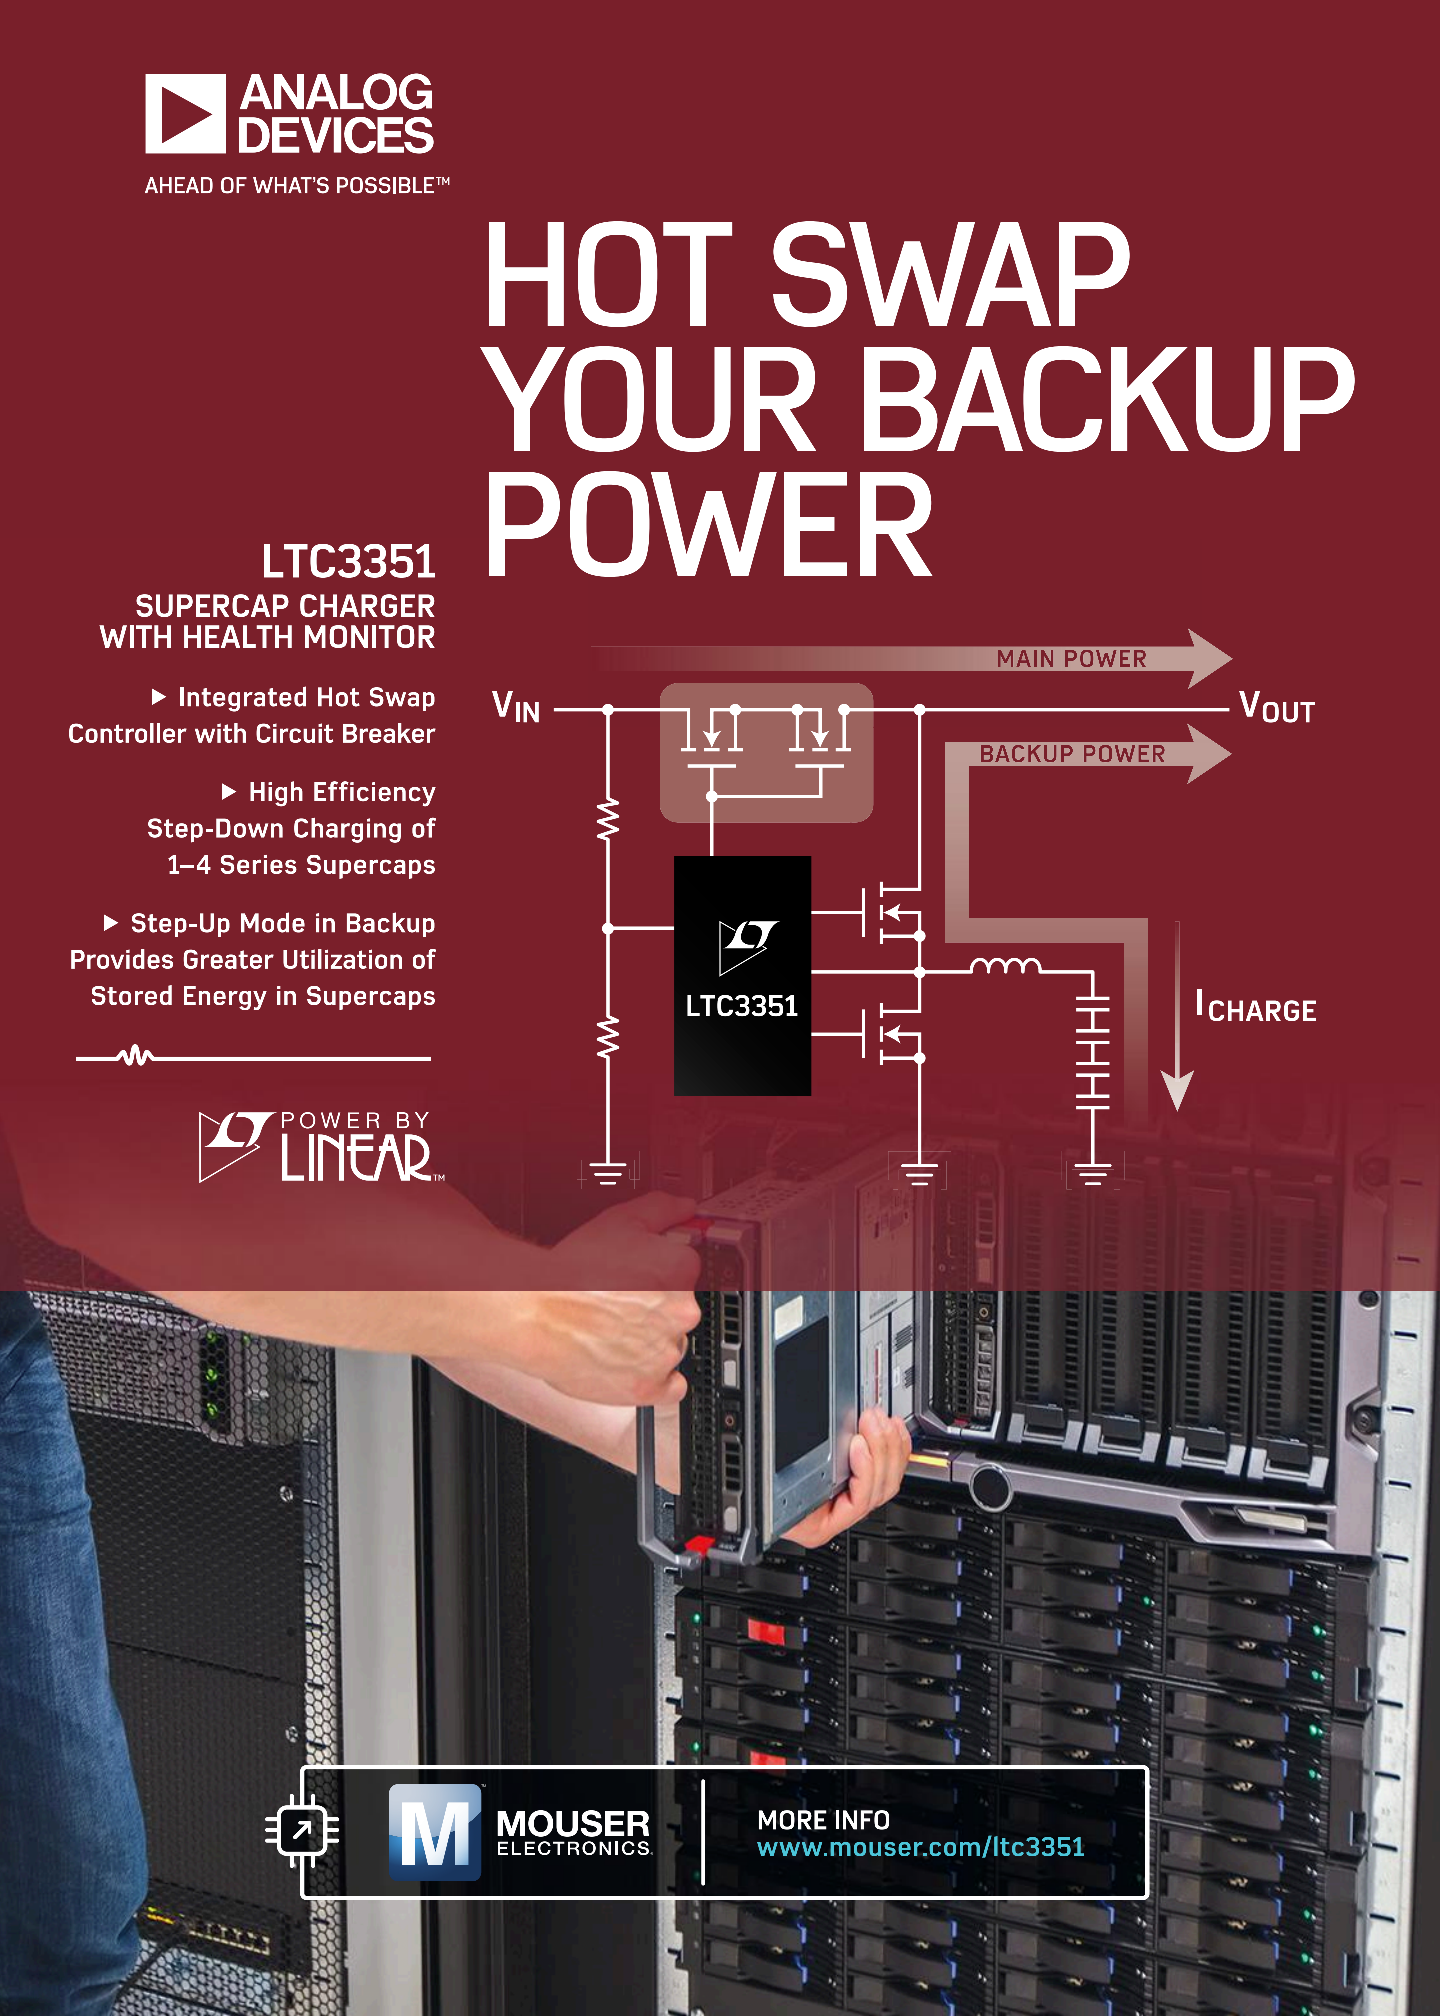
ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™
HOT SWAP YOUR BACKUP POWER
LTC3351
SUPERCAP CHARGER WITH HEALTH MONITOR
▶ Integrated Hot Swap Controller with Circuit Breaker
▶ High Efficiency Step-Down Charging of 1–4 Series Supercaps
▶ Step-Up Mode in Backup Provides Greater Utilization of Stored Energy in Supercaps
POWER BY LINEAR™
MAIN POWER
VIN
VOUT
BACKUP POWER
LTC3351
ICHARGE
MOUSER ELECTRONICS
MORE INFO
www.mouser.com/ltc3351