Electronic Design®

THE AUTHORITY ON EMERGING TECHNOLOGIES FOR DESIGN SOLUTIONS

JULY/AUGUST 2018
electronicdesign.com

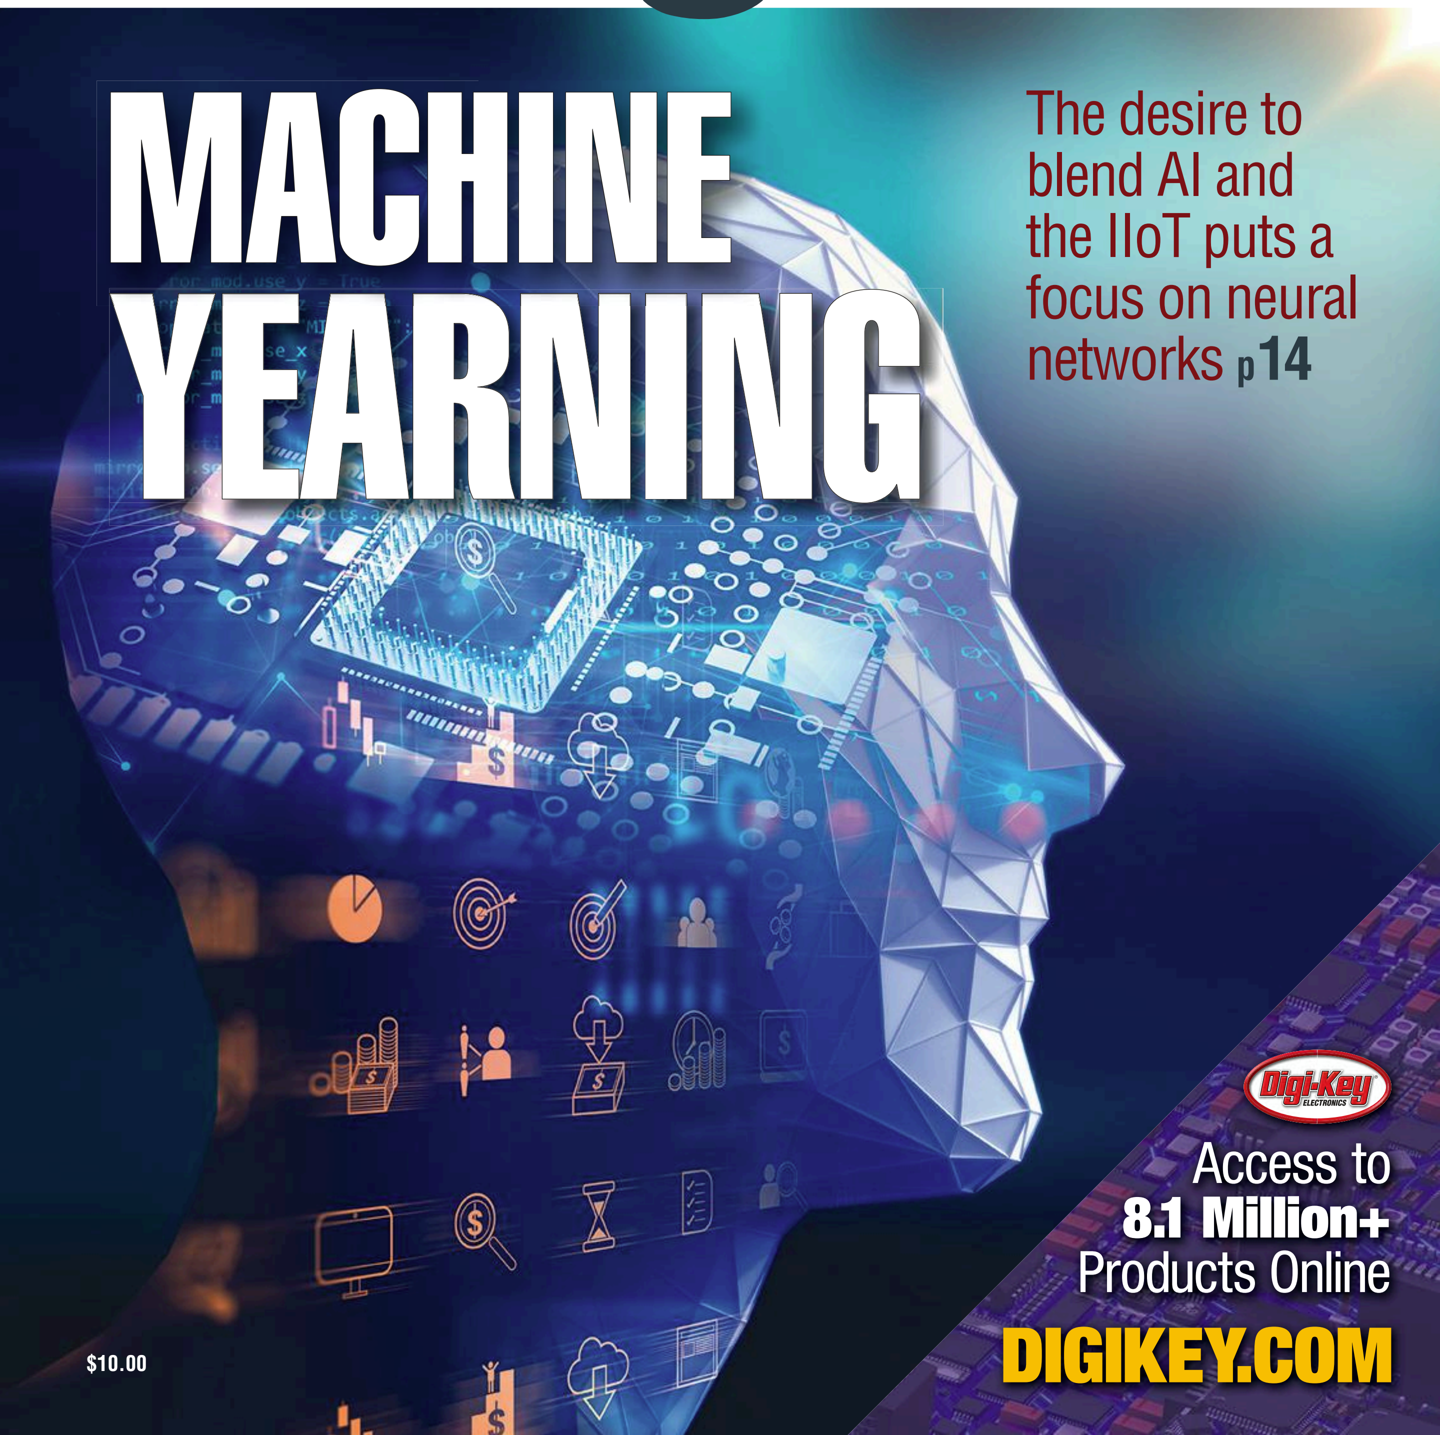

MACHINE YEARNING

The desire to blend AI and the IIoT puts a focus on neural networks p14

$10.00

the accuracy of *Rubidium*...

PRS10 Rubidium Oscillator (10 MHz)

- Less than 5×10^{-11} aging per month
- Ultra low phase noise (-130 dBc/Hz @ 10 Hz)
- 20 year lamp life
- 1 pps input and output
- RS-232 computer interface

$1495 (U.S. list)

FS725 Benchtop Rubidium Frequency Standard

- 5 MHz and 10 MHz outputs
- 0.005 ppm aging over 20 years
- Built-in distribution amplifier (up to 22 outputs)
- 1 pps input and output
- RS-232 computer interface

$2695 (U.S. list)

SRS rubidium frequency standards have excellent aging characteristics, extremely low phase noise and outstanding reliability.

The PRS10 component rubidium oscillator is designed for easy system integration. It has a 1 pps input for phase-locking to an external reference (like GPS) and provides 72 hour Stratum 1 level holdover.

The FS725 benchtop instrument is ideal for the metrology laboratory as well as the R&D facility – anywhere precision frequency is required. It generates 5 MHz and 10 MHz signals and has a built-in distribution amplifier with up to 22 outputs.

SRS **Stanford Research Systems**
1290-D Reamwood Ave. Sunnyvale, CA 94089 · email: info@thinkSRS.com
Phone (408) 744-9040 · Fax (408) 744-9049 · **www.thinkSRS.com**

IN THIS ISSUE

FEATURES

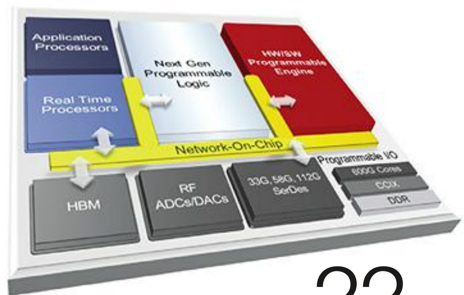

22

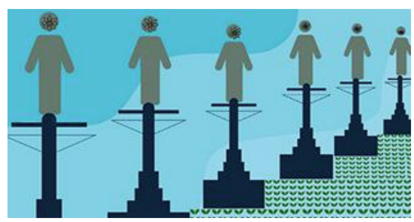

32

COLUMNS & DEPARTMENTS

48

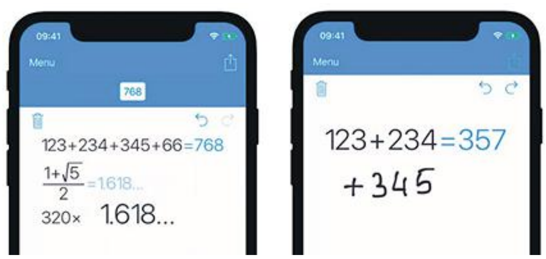

EDITORIAL MISSION:
To provide the most current, accurate, and in-depth technical coverage of the key emerging technologies that engineers need to design tomorrow's products today.

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

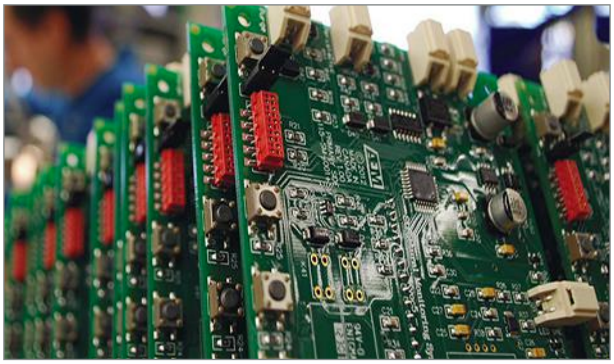

The Chronicles of GND (Part 2): Single Supply = Multiple Problems

The latest installment of "The Front End" from Kendall Castor-Perry, "The Filter Wizard" looks at how the ancestral "split-supply" PCB layouts have mostly ceded to "single-supply" configurations, and the GND implications that often ensue.

http://www.electronicdesign.com/analog/chronicles-gnd-part-2-single-supply-multiple-problems

How Modular Design Is Transforming Enterprise Connectivity

As transformative technologies such as 5G change the connectivity landscape, IT continues to turn more to modular approaches for enterprise system solutions.

http://www.electronicdesign.com/analog/how-modular-design-transforming-enterprise-connectivity

Bitcoin Mining: A Thermal Perspective

Why does Bitcoin mining consume humongous amounts of electricity, and how does it relate to thermal engineering? Mentor's John Wilson, Technical Marketing Engineer and Electronics Product Specialist, provides the answers.

http://www.electronicdesign.com/industrial-automation/bitcoin-mining-thermal-perspective

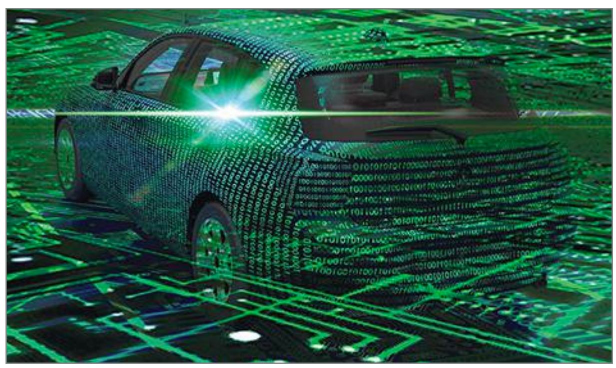

AI Poses a Tough Road Ahead for Autonomous Car Makers

Following recent crashes, reliable performance remains the biggest question mark, but cost remains another barrier to break through to achieve consumer viability.

http://www.electronicdesign.com/automotive/ai-poses-tough-road-ahead-autonomous-car-makers

ElectronicDesign.

JULY/AUGUST 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:
SHANNON ALO-MENDOSA **T** | 978.501.9116 shannon.alo-mendosa@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


CHARGE UP
YOUR SUMMER
MPD
MEMORY PROTECTION DEVICES
DL2032
DURACELL
WELCOME
BATTERYHOLDERS.COM
MPD
MEMORY PROTECTION DEVICES

The Unique AKIII Enclosure ....

The patented Air Ventilation elements avoid condensation while preserving the high protection class IP65. Occurring condensation is transported out of the enclosure immediately via air exchange.

The continuous and high air exchange allows the interior air to mix constantly with the environmental air and moves the moisture outward.

Even environments with nearly constant humidity and temperature possess air exchange, since the air on the interior of the enclosure warms due to the fittings. The individual components of the ventilation elements mainly consist of the rotating cover (rear side of box) and the 10μm filter element integrated on the inside of the box.

E254144
UL TYPE 3R, 12, 12K

EU

PROTECTION IP65

WINNER iF DESIGN AWARD

Advantages of Air Ventilation Elements

- Prevention of condensate water while preserving the high protection class IP65
- High air exchange rate
- 100% pressure compensation
- Foreign object protection
- Maintenance-free

Combination flange / flange end walls

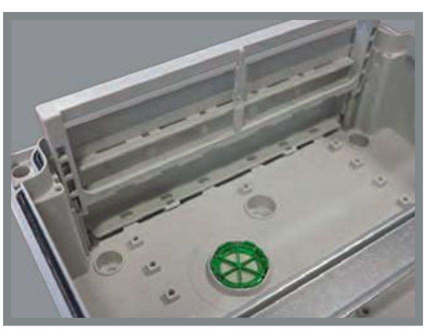

Closed wall

Knockout flange

Knockout flange

Combination flange

...with Patented Air Ventilation System

Patent No.
WO 2015/063057 A1

Accessories available
including air ventilation element
for use in other enclosures.

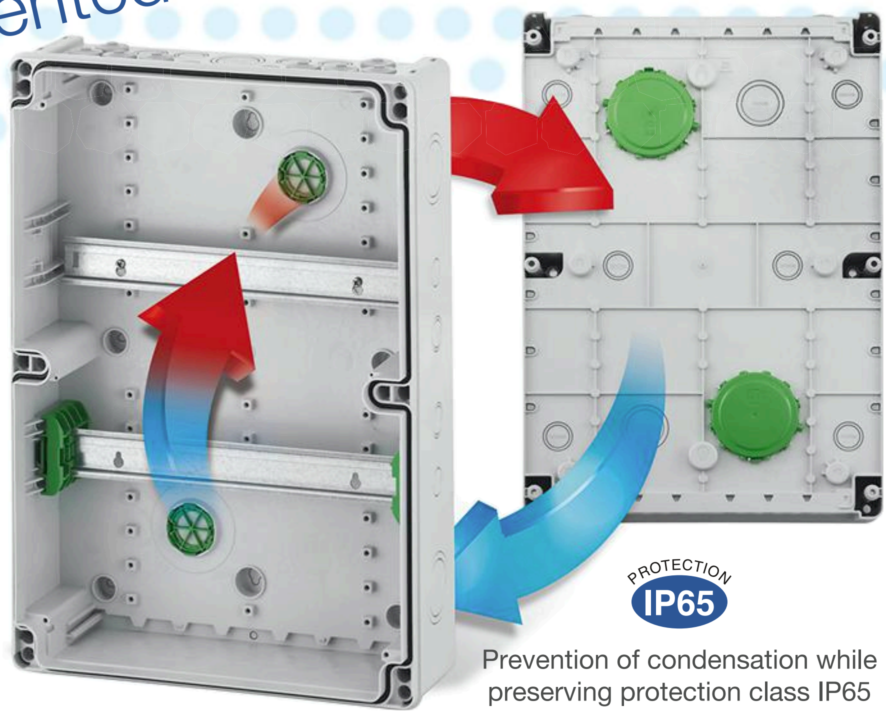

Prevention of condensation while
preserving protection class IP65

Wide Range of Features:

- Integrated air ventilation
- Top/bottom interchangeable flange endwalls
- Height-adjustable mounting rail
- 150 mm rail spacing
- Optimized hinged window
- Combinable
- Multiple installation options
- DLG focus test: Ammonia resistance
- Indoor and outdoor installation
- Free of halogen, PVC and silicone

The AKIII Enclosure Family

Enclosure Size Range: 315 x 300 x 155 mm (12.4 x 11.8 x 6.1 in) to 315 x 750 x 155 mm (12.4 x 29.5 x 6.1 in)
Knockouts Sizes: M40/50 - M20/25 - M20 - M20/25 -M32/40

High-performance power

1.8kW

400 - 700V_{IN}

98% efficiency

Isolated fixed-ratio

Ultra-high voltage DC-DC converter

4.4 x 1.4 x 0.3in
110.5 x 35.5 x 9.4mm

HIGH POWER

500W

200 - 420V_{IN}

93% peak efficiency

Isolated regulated

DC-DC converter

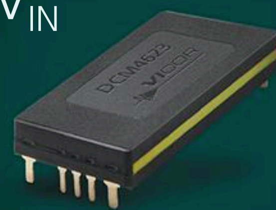

1.8 x 0.8 x 0.2in
47.9 x 22.8 x 7.2mm

NEW RELEASE

98% efficiency

750W, 12/48V bidirectional

Non-isolated bus converter

0.9 x 0.6 x 0.2in
23 x 17 x 7.4mm

Up to 96% efficiency

8 – 60V_{IN}

10 – 50V_{OUT}, 140W

Cool-Power ZVS buck-boost switching regulator

10 x 14 x 2.6mm

Qualify for a free evaluation product at **vicorpower.com**

VICOR

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

Fill Those Vacant Tech Jobs with Co-op Students

Finding young, qualified talent to work in the tech industry can be as challenging as trying to find your first job. Having any experience tends to improve the odds, but getting it isn't always easy. There are essentially three ways to fill out the resume: find a part-time job in industry, find an internship, or go to a university with a co-op program.

The first option is usually like trying to find that job without any experience. Some are lucky to nab a paid position, while others may have to work for free. Programmers can often participate in open-source projects, but hardware designers and engineers may find it more difficult to get involved in projects.

Many universities encourage internships, and some have programs designed to match students with companies that have internships. Part of the challenge with internships is that they're often unpaid, vary in duration, may be part time, and so on. They can provide budding engineers with a taste of industry and even some good training. However, unless an internship is repeated, the experience is more limited.

Co-op programs are designed for full-time positions that repeat over the course of earning a degree, and often extend the typical four-year program to five years. There are no breaks, as students wind up alternating between taking courses and working, usually alternating semesters. This makes it possible for a company to fill a position throughout the year with different students who are on alternating schedules.

Companies benefit by having a known, future graduate who will have experience with the company and its processes. The typical half-hour interview will now be a year or more of experience. Students benefit from exposure to industry and typically gain as much or more education in the real world that can also help direct their studies in college.

I can endorse Georgia Tech's co-op programing as a graduate from the program with a Bachelor of Electrical Engineering. I was able to pay for Tech with my co-op earnings, although these days, expenses tend to a bit higher. Other well-known schools have well-established co-op programs, too, including Drexel, Northeastern, Rochester Institute of Technology, Clarkson University, and many more.

Co-op programs within a company require some effort for both human resources as well as the groups where co-op students work. Though the programs are generally easier to implement in larger companies, small- and medium-sized organizations can participate as well. Co-op programs in universities are designed to work with all types of companies.

Co-op students can be useful for projects that might not otherwise be languishing, simply because the full-time staff is too busy. Of course, co-op graduates can be a good source for new staff, although there's no requirement to offer a co-op graduate a job.

I learned a lot at Georgia Tech and at Burroughs where I co-op'ed. I actually wound up at RCA's Sarnoff Research Labs when I graduated, and I know my experience at Burroughs was a good reason for that happening. You may want to hire a co-op graduate, so consider looking into starting a co-op program if you don't have one already. ed

News

MICRON BARRED FROM Selling Some Memory Chips in China

Micron has been barred from selling certain memory chips in China after a preliminary ruling came down in patent infringement cases brought against it by United Microelectronics Corporation and Fujian Jinhua Integrated Circuit. The company called the order was unfair because the patents in question were not used in its products. But it also said that restrictions will only hurt its business slightly.

The order undercuts the sale of DRAM modules and solid-state drives sold under Micron's Crucial and Ballistix brands. The company said on Thursday that it expects to lose only one percent of its revenue because of the order, which was handed down by Fuzhou Intermediate People's Court in Fujian Province. Micron still estimates to generate between $8 and $8.4 billion of revenue in the fourth quarter ending in August.

In December, Micron introduced a lawsuit in United States District Court for the Northern District of California that accused UMC of stealing trade secrets and giving the technology to Jinhua. The Taiwanese chip manufacturer responded in January with a lawsuit alleging that Micron had infringed on its patent rights in China. Micron said that its claims "rely on distorted interpretations of the patents and improper evidence."

"The central government of China has often stated that the rights of foreign companies are fairly and equally protected in China," the company said, adding that it believed the ruling

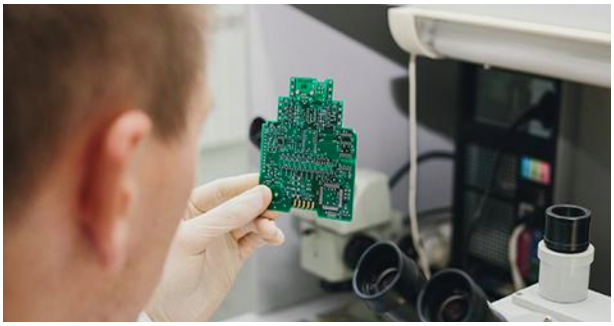

to be "inconsistent with this proclaimed policy." Joel Poppen, Micron's general counsel, added that it "will continue to aggressively defend against these unfounded patent infringement claims while continuing to work closely with its customers and partners."

Boise, Idaho-based Micron has been hoovering up profit as demand for memory chips, used in everything from smartphones and servers to thermostats and cars, outstrips supply. In turn, that has raised prices for its products. The company—which only two years ago was losing money, eliminating jobs and weighing an acquisition offer from China's Tsinghua Unigroup —nearly doubled its year-over-year profit last quarter, to $3.82 billion. ■

CROSSBAR EYES EMBEDDED SPACE with Resistive Random-Access Memory

GROWING REQUIREMENTS FOR lower power and high performance in embedded devices could push electronics engineers toward more specialty memory chips. That could include chips based on Crossbar's resistive random-access memory, or ReRAM, which attempts to bridge the gaps between nonvolatile flash memory technology and high-speed RAM.

Crossbar recently announced an agreement with Microsemi Corporation, one of the largest American makers of military and aerospace semiconductors. Microsemi said that it would license Crossbar's intellectual property and collaborate on the development of chips integrated with ReRAM technology, which Crossbar said is scalable down to 10 nanometers. The terms of the deal were not disclosed.

Crossbar's ReRam represents bits using the resistance of the circuits, which can be set and reset by a pulse of voltage. The memory is composed of silicon switching material sandwiched between two electrodes. When a voltage is applied, a metallic filament forms in the switching material, connecting the electrodes and changing resistance between them. Reversing the voltage causes the filament to recede, cutting off the conductance.

The company says that chips embedded with ReRAM technology devour less power, access information with lower latency, and have more endurance than flash memory devices, such as NOR and NAND. In addition, Crossbar's memory devices are nonvolatile, meaning that they preserve what's stored inside them after the system powers down. It sets the technology apart from DRAM devices that forfeit information when powered down.

Crossbar is also targeting the embedded space, with an eye toward edge computing, communications infrastructure, and automotive applications. Another focus is artificial-intelligence tasks, which are constantly bobbing for information stored in memory. ■

WILLIAM WONG | Senior Technology Editor

Neural-Network Hardware

Drives the Latest Machine-Learning Craze

From self-driving cars to the industrial Internet of Things, neural networks are reshaping the problem-solving methods of developers.

Artificial-intelligence (AI) research covers a number of topics, including machine learning (ML). ML covers a lot of ground as well, from rule-based expert systems to the latest hot trend—neural networks. Neural networks are changing how developers solve problems, whether it be self-driving cars or the industrial Internet of Things (IIoT).

Neural networks come in many forms, but deep neural networks (DNNs) are the most important at this point. A DNN consists of multiple layers, including input and output layers plus multiple hidden layers *(Fig. 1)*. The number of nodes depends on the application; each node has a weight associated with it. An input, such as an image, is supplied at one end and the outputs provide information about the inputs based on the weights.

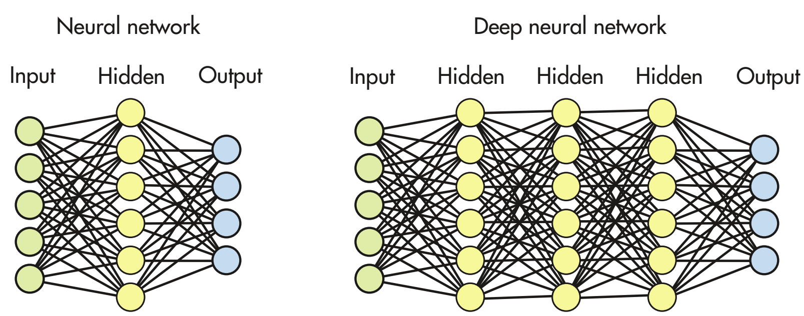

1. Deep neural networks (right) have multiple hidden layers.

The magnitude of the inputs, weights, and calculations done at each node are of no importance in a logical sense. However, they're crucial in terms of implementation, because this affects the amount of computational performance necessary for a system, as well as the amount of power required to perform the calculations. The fewer bits involved, the lower the performance and power requirements; hence, a more efficient implementation.

As it turns out, the number of bits used to encode weights and calculations done can often be significantly reduced—sometimes to a single bit—although eight bits is typically sufficient. Some implementations even employ small floating-point encodings, since the number of significant digits is less important than a value's potential range.

TRAINING AND INFERENCE

There are different ways to create and implement DNNs. The typical method is to design a network and then train it. Training consists of presenting input data and matching results so that the weights in the nodes within the DNN matrix are adjusted with each new piece of information.

The architecture of a DNN is simple, but not so much when it comes to design and implementation. The overall system results can be affected by a number of issues, from the format, size, and quality of the input to the type of output desired. Likewise, the training process of a successful design assumes that the model will converge over time as more examples are presented. In essence, the system comes to recognize the desired characteristics.

The resulting weights and network model are then used as an inference engine. In this case, there's no change to the network while in use. Essentially all of the machine learning (ML) for this type of system is done in the training session. A better implementation with more training or a different model can replace an existing implementation if additional input is available.

For example, an image-recognition system in a car could also record images that would subsequently be sent to the cloud for additional training. This would allow thousands of cars to provide information to a cloud-based training system, which in turn would generate an improved model that could be distributed to those cars in the future after testing. It's probably not a good idea to have a self-direct learning system for each car that would change on its own, because it would be difficult to determine if the training was improving the system or not.

Trained models can often be optimized so that the inference engine employs smaller weights or uses fewer resources. For example, the number of bits utilized for weights and calculations may be reduced but provide the same or similar results. Keep in mind that a DNN is designed to deliver probabilities as results, and it's possible that false-positive and false-negative results can be generated. The idea of a good design is to minimize or effectively eliminate these, especially in safety-related applications like self-driving cars, but these consequences in other applications may be more than acceptable.

DIFFERENT NEURAL NETWORKS

Convolutional neural networks (CNNs) are a specific type of DNN. CNNs are normally used for image recognition. There's also recurrent neural networks (RNN) designed to address temporal behavior, allowing the system to address problems such as speech and handwriting recognition. RNNs can be used for a variety of audio and video processing applications, too.

Another type of neural network is the spiking neural network (SNN). Unlike a DNN, where information always propagates through the network, an SNN has nodes that trigger when a threshold is reached. The output can actually be a train of signals from the system, rather than values, as results for all outputs in a DNN. This is more akin to how the brain operates, but it's still an approximation.

"Spiking neural networks, considered by many to be the next generation of artificial intelligence, is based on an explicit incorporation of time in the form of events or spikes for representation of data," says Nara Srinivasa, CTO of Eta Compute. "When this representation is sparse and the hardware that runs these models is fully asynchronous, these event-driven networks can learn to solve machine-learning-type problems at a fraction of the power consumption of machine-earning approaches today."

SNNs can often address applications that DNNs and their ilk are less adept at performing. For example, a DNN trained to recognize flowers in images can now be used to scan new images to find things that look like a flower, but it may be difficult to selectively identify a particular type of flower. An SNN can often do the latter; however, it may be less effective identifying anything that looks like a flower.

The challenge for developers new to this space is that high-level ideas like DNNs and machine learning are easy to appreciate. The sticking point is that no one solution that works best all of the time. Likewise, the tradeoffs in designing, implementing, and delivering a system can be difficult to address because of the number of options as well as the implications of choosing among them. In addition, the tradeoffs in training, accuracy, and speed may make some applications possible, while others become out of reach.

Implementation issues are further complicated because hardware optimization can significantly affect the performance of a system. As noted, different possible implementation criteria and hardware optimized for one approach may be unsuitable or even useless for another.

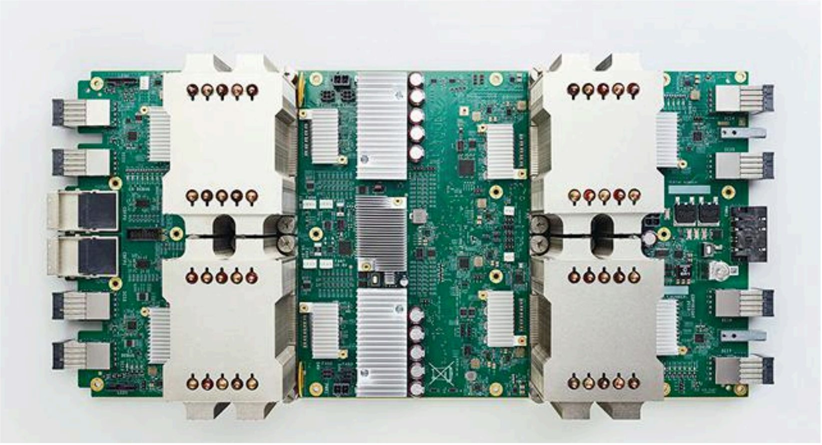

2. Google's TensorFlow Processing Unit (TPU2) delivers 45 TFLOPS. The newer TPU3 is eight times more powerful.

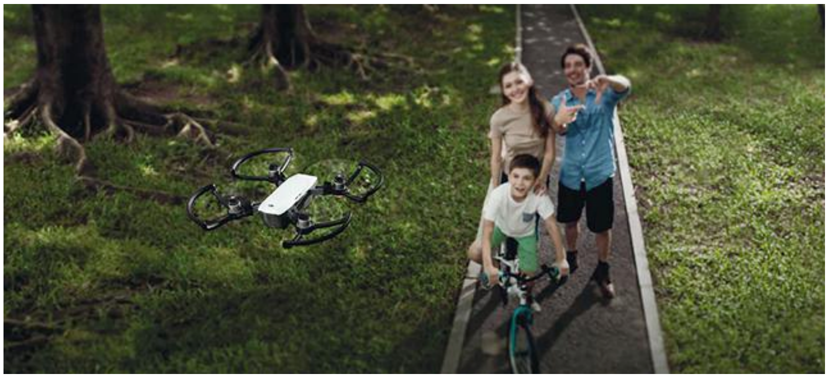

3. DJI's SPARK drone uses Intel's Myriad 2 visual processing unit (VPU) with ML support to recognize gestures such as framing and taking a picture.

SPECIALIZED ML HARDWARE

Hardware acceleration can significantly improve machine-learning performance across the board. Some hardware is often targeted at training, while other hardware may be optimized for inference chores. Sometimes a system can do double duty. In certain cases, the hardware can reduce the power requirements so that a system can run off of very little power, including battery-operated environments.

At one end of the spectrum is Google's TensorFlow Processing Unit (TPU2). TensorFlow is one of the more popular ML frameworks. The TPU2 *(Fig. 2)* is designed for the cloud where dozens or hundreds of TPU2 boards would be used to handle a range of training and inference chores.

Unlike many ML accelerators, the TPU2 targets a subset of the computing used in a DNN model. Of course, it targets the area that's the most computationally intensive, allowing the host processor to fill in the blanks. Each of the TPU2 chips has 16 GB of high bandwidth memory (HBM) supporting a pair of cores. Each core has a scalar floating-point unit and 128-×-128, 32-bit floating-point mixed multiply unit (MXU). It can deliver 45 TFLOPS of processing power.

Google has already announced the TPU3. It's eight times more powerful than the TPU2, and is water-cooled.

At the other end of the spectrum is Intel's Myriad X visual processing unit. The Myriad X includes 16, 128-bit VLIW SHAVE (Streaming Hybrid Architecture Vector Engine) processors. It also has a pair of 32-bit RISC processors along with hardware accelerators and neural-network hardware. Its 16 MIPI lanes support up to eight HD cameras. The earlier Myriad 2 was used in DJI's Spark drone *(Fig. 3)*. It allows a drone's camera to recognize gestures in order to perform functions such as framing and taking a photo.

These are just two examples of dedicated ML hardware.

GPGPUs AND FPGAs DO AI

Dedicated hardware isn't the only way to support machine learning. In fact, conventional CPUs were used initially, but alternatives typically provide much better performance, making it possible to employ much larger and more complex models in applications. Likewise, many applications require a range of models to operate on different data. Self-driving cars are good example—multiple models would be employed for different functions, from analyzing an engine's performance to recognizing people and cars around the vehicle.

Two frequently used platforms in ML applications are GPGPUs and FPGAs. As with dedicated hardware, developers need to consider many options and tradeoffs. These types of platforms provide more flexibility since they depend more on the software. This tends to make it easier to support a wider range of models and methodologies as well as support new ones. Dedicated hardware is generally more limited.

These days, the GPGPUs from Nvidia, AMD, Arm, and Intel can be used for computation, not just for delivery of flashy graphics on high-resolution displays. In fact, GPGPUs dedicated to computation with no video outputs are common and target ML applications in the cloud. This works equally well for GPGPUs in desktops, laptops, or even system-on-chip (SoC) platforms.

Part of the challenge for GPGPUs and FPGAs is that they were tuned for different applications. For example, single- and double-precision floating-point matrix operations are quite useful for everything from simulating weather to finding gas and oil from sensor information. Higher precision is usually an advantage.

Unfortunately, DNNs and the like may underutilize this type of precision, as noted earlier. Still, GPGPUs and FPGAs often handle small integers well, and the latest generation of these platforms is being tailored to handle the types of computations needed for ML applications. On the plus side, many applications will require additional processing that these platforms already support to massage incoming data or process results. Overall, such platforms are more adaptable to a wider range of programming chores.

FPGAs have an edge when it comes to adaptability. This can be especially useful where methodologies are still changing, and new approaches may require different hardware acceleration. GPGPUs will work well if the application fits the hardware.

Andy Walsh, Senior Director, Strategic Marketing at Xilinx, notes "Artificial intelligence is a hotspot of innovation in computing today. But to stay best-in-class, both network models and acceleration algorithms must be constantly improved. To keep pace with the design cycles in AI, developers are looking for more than just a chip. They need a computing platform that is highly performant and adaptable.

Recognizing this need, we've recently unveiled a new product category that goes far beyond the capabilities of an FPGA. The breakthrough adaptive compute acceleration platform from Xilinx is a highly integrated, multi-core, heterogeneous platform that can be programmed at the hardware and software levels. Using tools familiar to developers, an adaptive compute acceleration platform (ACAP) can be tuned to deliver performance and power efficiency for AI inference that outshines any other processing platform."

Embedded FPGAs (eFGPAs) are being quick to embrace machine learning. For example, Flex Logix Technologies' EFLX 4K AI *(Fig. 4)* has DSP blocks that can deliver 10 times the MACs compared to the EFLX 4K DSP that employs typical FPGA DSP blocks. The AI DSP block eliminates logic not required for ML matrix operations and is optimized for 8-bit configurations common in ML applications. The AI DSP blocks can also be configured as 16-bit MACs, 16-×-8 MACs, or 8-×-16 MACs.

AI eFPGAs allow a designer to incorporate just the amount of ML computational power needed for an application. A smaller, more power-efficient version might be used with a microcontroller, while larger arrays could be used where more performance is needed.

SOFTWARE MAKES AI WORK

Regardless of the hardware used, ML and AI applications need software to take advantage of the hardware. On the plus side, vendors of general hardware such as CPUs, GPUs, and FPGAs are delivering software that works with the standard machine-learning frameworks. This means that applications developed on the frameworks can be ported to different hardware, giving developers a choice of what hardware will be used in a particular application.

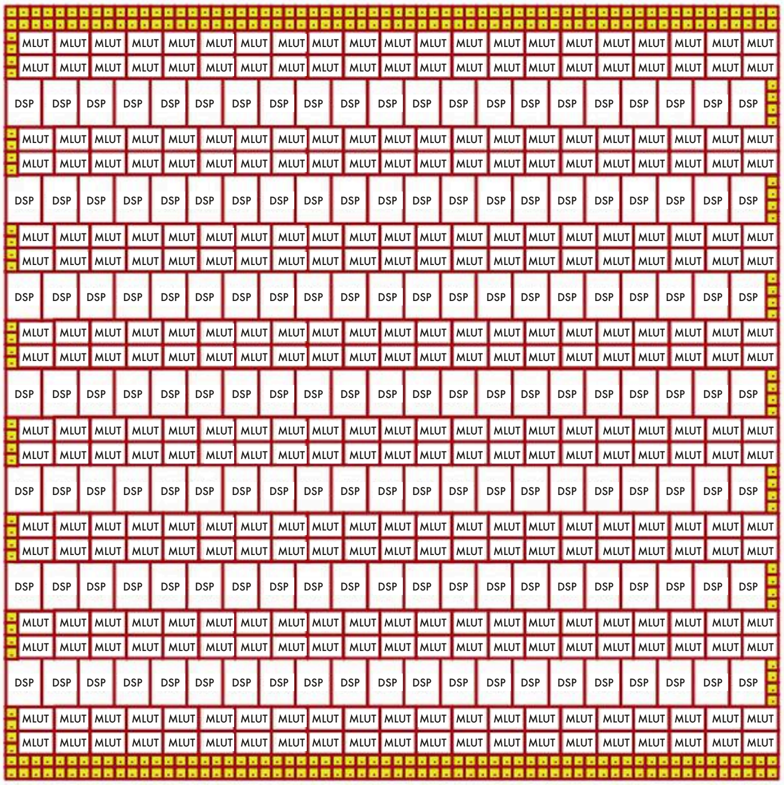

Flex Logix Technologies' EFLX 4K AI has DSP blocks that can deliver 10 times the MACs compared to the EFLX 4K DSP that employs typical FPGA DSP blocks.

Choosing which hardware platform to use for an application can be difficult, since different approaches may offer performance and efficiency levels that vary by an order of magnitude. Likewise, the AI portion of an application may only be a fraction of the code involved, even if it requires a significant portion of the computational support. ed

4. Embedded FPGAs (eFPGAs) like Flex Logix Technologies' EFLX 4K AI are delivering building blocks optimized for today's ML frameworks.

WILLIAM WONG | Senior Technology Editor

IoT Sensor Micros Need to Sip Power

Minimal power consumption is the watchword for IoT sensors, and multiple pathways can be taken to reach this goal.

Many devices in the Internet of Things (IoT) are battery-operated or they may even use power-scavenging systems. The bottom line is that the process tends to be the part of the system using the most power, so optimizing how it uses power is critical.

The typical way to minimize power requirements is to run at full speed for as little time as possible and power down at other times. These days, most micros have a range of power-down modes with different power-utilization features that developers can take advantage of depending on the application. Some may only need to run the real-time clock (RTC), while others may have to maintain some or all of the RAM contents.

These methods usually help with almost any microcontroller, but what other options do developers have when choosing a low power platform?

CHIP DESIGN

Microcontroller chip designers generally choose technologies that will minimize power requirements while providing performance. Each iteration tends to improve power utilization, but some techniques offer significant advantages over the conventional method.

Eta Compute takes an asynchronous approach versus the conventional clocked design. The challenge with a clocked design is that it's always cycling, and thus consumes power while in this mode. Eta's design, with its Cortex-M3 and spiking neural-network support, takes advantage of this approach. The machine-learning (ML) ASIC consumes less power than most Cortex-M3 platforms.

Another approach is applied by Ambiq Micro for its Apollo 2. Its Subthreshold Power Optimized Technology (SPOT) platform uses transistors biased in the subthreshold region of operation. This allows the Cortex-M4F to consume less power than most Cortex-M3 implementations. Floating-point support is included. A Bluetooth version of the chip is available, too.

The challenge with both approaches is that the design process is more difficult than the conventional chip design. New design libraries help, but at this point, the technologies have been applied to a limited number of chips. Nonetheless, both hit the sweet spot for IoT devices.

MEMORY DESIGN

Texas Instruments' (TI) 16-bit MSP430 family has been known for its low-power operation. The family includes a wide variety of chips that target specific sensor and control applications.

One feature found in in the TI MSP430 FRAM (ferroelectric RAM) group of chips is FRAM memory. This non-volatile memory technology doesn't have the speed and lifetime limitations that bog down flash memory. It was

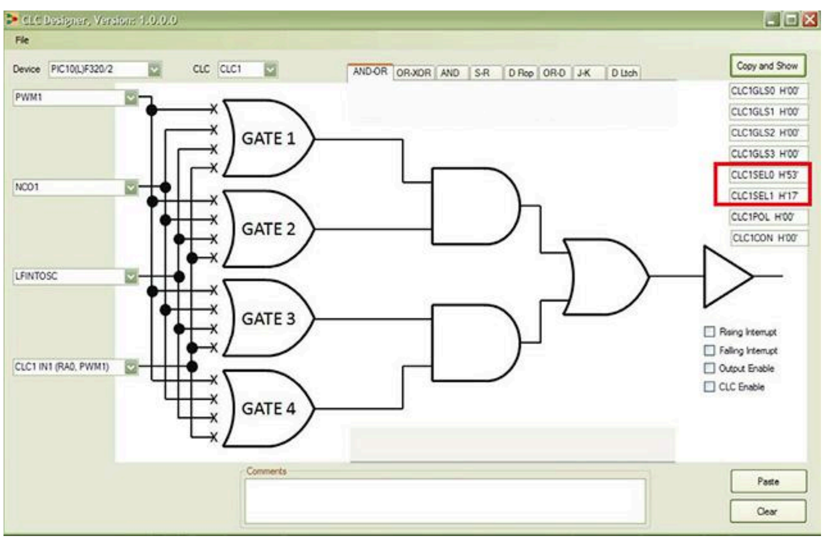

Microchip's Configurable Logic Cell design tool helps designers link smart peripherals together.

originally used for data storage, but now TI has chips with a unified memory for code and data, which creates significant flexibility in memory allocation.

Non-volatile data storage has other advantages for IoT environments where the processor runs only part of the time because data needn't be stored in flash memory, or RAM doesn't require battery backup. This allows the MSP430 to use very low power modes because the FRAM doesn't need power to maintain its contents. Not having to move data to and from flash storage means start-up and shutdown is faster, which also reduces power requirements.

FRAM is exclusive to TI's MSP430, but a competitor looms on the horizon. Magnetoresistive RAM (MRAM) shares FRAM's features compared to flash memory. Standard MRAM microcontrollers aren't available yet, but the foundries used by fabless micro vendors are now including MRAM as part of their standard fare.

PROCESSOR OFFLOAD

Smart peripherals are another way to limit power consumption by allowing data acquisition and transfer while the processor is in sleep mode. A smart I^2C slave device may have built-in address recognition, allowing it to wake up the processor when the device is selected.

The other way to build smart peripherals is to link multiple peripherals together. For example, Microchip's Configurable Logic Cell (CLC) can tie the outputs from peripherals to the input of another peripheral or interrupt. With the CLC configuration tool *(see figure)*, designers can generate designs quickly and easily.

Most microcontroller vendors have a variation of CLC for their chips. STMicroelectronics calls its implementation "autonomous peripherals." NXP's term for intelligent peripheral management is FlexIO. And Silicon Labs has its Peripheral Reflex System (PRS). NXP has its own set of configurable peripherals.

Cirrus Logic's PSoC takes configuration to the extreme with near FPGA flexibility. Not only can designers configure the linkage between peripherals, but also the peripherals themselves. PSoC chips come with a set of digital and analog blocks that can be configured into a range of standard and custom peripherals.

Power efficiency is affected by numerous aspects of a design, including the processor and its memory, peripherals, and their configuration. Developers will need to consider all of these aspects when selecting a device to meet the cost and lifetime requirements of a system.

WILLIAM WONG | Technical Editor

Not Your Father's FPGAs Anymore

FPGAs continue to evolve from the original collection of gates and routing to handle tasks from communications to artificial intelligence.

FPGA's, like GPUs, have changed significantly from their initial inception that took a narrower view of the solution space. They've morphed from a collection of gates and routing to taking on jobs that range from communications to artificial intelligence (AI).

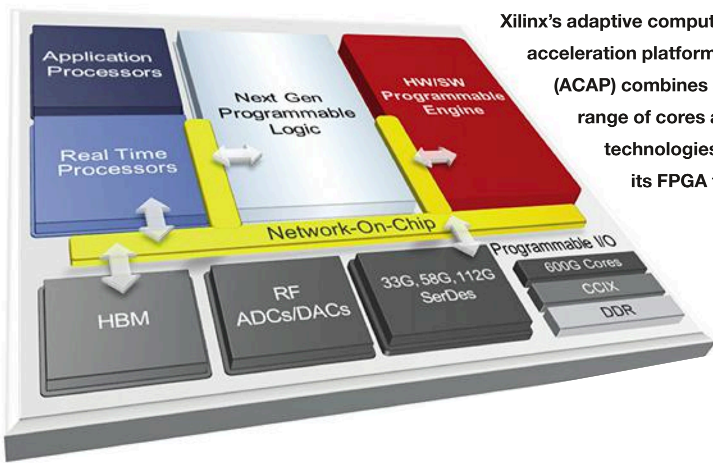

Xilinx's adaptive compute acceleration platform (ACAP) combines a range of cores and technologies with its FPGA fabric.

CHANGING FPGA LANDSCAPE

FPGAs used to be single chips like most devices. Though they've grown in size in terms of transistors, the underlying architecture is morphing as well. FPGAs, large and small, are still available and useful in their original context, but the overall FPGA solution landscape is much broader.

At the high end, interposer technology is allowing types of devices to be combined on a single die. Xilinx's adaptive compute acceleration platform (ACAP) pushes the envelope with multiple hard cores employing different silicon node processes and analog support, such as mixed signal and SERDES onto a single chip (*see figure*).

Interposer technology is available in a number of forms like Intel's embedded multi-die interconnect bridge (EMIB). EMIB uses through-silicon vias (TSVs) like other interposer technologies, but smaller, individual interposers are utilized between die. This is more like using connectors than a backplane. TSV technology is also applied in stacking dies, which are common in high bandwidth memory that's also being tied to FPGAs.

Mixing technologies brings significant advantages when it comes to analog technology. It allows for high-speed SERDES to be optimized, since they utilize transistor technology different from that needed for FPGA fabrics.

Another change in the FPGA space is the rise of embedded FPGAs. Usually, FPGA intellectual property (IP) was only used in chips from FPGA vendors. Now a number of embedded FPGA vendors provide IP so that FPGA fabrics are part of a system-on-chip (SoC) solution. This is the flip side of FPGA SoC designs from FPGA vendors that incorporate hard-core processors and peripherals with an FPGA fabric. In both instances, the approach is to use compatible FPGA technology that meshes with the other IP versus using interposer technology to link dissimilar silicon components.

Also worth mentioning is the impact of flash-based FPGAs. Their low power and instant-on capabilities allows them to be used in applications such as mobile or in IoT devices where RAM-based FPGAs may be impractical.

CHANGING ARCHITECTURES

The use of hard-core processors and peripherals in FPGAs has been commonplace for a while, as are building blocks like DSP and floating-point units, but even these are changing as application demands shift.

AI is one tool that's reinventing how FPGAs work. These days, AI's machine learning (ML) means implementing deep-neural-network (DNN) support. DNN inference support can benefit from small floating-point or integer support, whereas most DSP blocks lean toward larger values amenable to the problems addressed by earlier FPGAs.

FPGA vendors have been eyeing the lookup table (LUT) and interconnect fabric as well, to see how to improve implementations in the face of wider FPGA usage.

Another way FPGAs are changing is the use of custom blocks. It's most common in embedded FPGAs, but some vendors like Achronix will incorporate custom or customer IP into a conventional FPGA fabric. This is different than the aforementioned interposer support in that the IP uses the same silicon technology employed by the FPGA logic. Moreover, it's designed to work within the FPGA interconnect fabric versus adding hard-core blocks. The custom IP support is integrated into the FPGA software tools so that the custom IP blocks can be connected and utilized like the usual FPGA LUTs.

Custom IP of this type is normally included in columns like LUTs and other FPGA building blocks such as DSP units. However, custom IP can have other advantages; for example, an additional connection fabric that's independent of the FPGA connectivity fabric. This can be handy for high-speed applications (e.g., networking), where it's possible to improve the flow of data significantly compared to a conventional FPGA implementation.

SOFT CORES AND RISC-V

Soft-core processors have been common since FPGAs appeared. Some vendors had their own optimized soft-core processors, such as Intel and its NIOS II and Xilinx with its MicroBlaze. There are a number of third-party cores for processors, like the venerable 8051 and Arm's Cortex-M1. The latter tended to provide more mobility between FPGAs compared to the proprietary soft cores from FPGA vendors, which were optimized for their FPGA hardware.

Enter RISC-V. RISC-V is actually an open-source instruction set standard. However, it leads to hardware implementation, including those on FPGAs. This approach has been spearheaded by Microsemi's Mi-V FPGA ecosystem. RISC-V is available for most FPGA platforms, in addition to showing up in RISC-V-based SoCs as well as GreenWaves Technologies' multicore GAP 8 SoC for machine-learning applications.

Even as FPGA hardware continues to morph, one of the greatest changes for FPGA technology is the software used to create the IP that runs on the FGPA. But that will have to wait for a future article. ed

PicoScope® 5000D Series PC Oscilloscopes

pico® Technology

The engineer's mixed-signal, flexible-resolution solution.

MSOs available with 16 digital channels

NEW

For engineers in mixed-signal environments, capturing both high-speed digital measurements and high-precision analog waveforms often requires two different oscilloscopes. A PicoScope® 5000D is the solution. Its innovative FlexRes® architecture combines multiple, high-resolution ADCs so users can selectively optimize sampling rate or resolution as needed. With less distortion and greater memory depth than typical 8-bit scopes. Plus, 5000D MSO models allow users to accurately time-correlate digital and analog channels and apply advanced triggers across both. Add a robust set of powerful software features, all standard, and a 5000D scope is both a great solution and an unmatched value.

KEY FEATURES

Flexible 8- to 16-bit hardware resolution • Up to 200 MHz analog bandwidth • Up to 512 MS capture memory • 130,000 waveforms per second • Arbitrary waveform generator • 200 MHz spectrum analyzer • Serial bus decoding (18 protocols) • Math channels & filters • Mask limit testing • Advanced digital triggers • DeepMeasure™ waveform analysis tool • Software development kit

US$1155 - 3025 **To learn more: visit www.picotech.com/A128 or call 800-591-2796**

MARK CARLSON | Co-Chair of the SNIA Technical Council and the SNIA Cloud Storage Technical Work Group
www.snia.org

The Persistence of MEMORY

A new storage media offers a blend of volatile memory speed and non-volatile persistence. Welcome to the world of Persistent Memory.

Typically, when you run a software program on your computer, mobile device or via a server in the cloud, you're using volatile memory, or memory that doesn't retain its contents when the power is off. When the code is finished with its tasks, the resulting data is usually handed off to non-volatile media, like a hard drive or solid-state drive (SSD). The data can live there until it needs to be pulled again for a given software process.

Imagine, though, that you could cut down on the lag, or latency, that happens with this hand off? You'd have the perfect blend of volatile memory speed and non-volatile persistence. In fact, the new nascent storage media is called just that: Persistent Memory.

A number of competing technologies exist on the market. Companies are all trying to get the lowest latency possible out of the underlying media. The ultimate goal with persistent memory is to achieve the same or better latency than dynamic random-access memory (DRAM); then it can be a DRAM replacement.

Persistent media already exists, but it uses a workaround by placing it behind DRAM, which has the lowest latency. What the host sees is fast DRAM latency. It makes that DRAM persistent by using Persistent Memory media as a backing store of sorts for the DRAM.

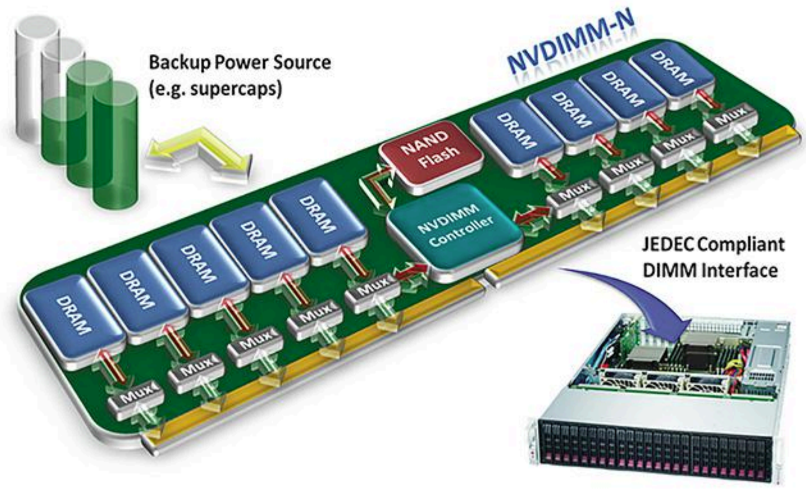

An NVDIMM-N is one type of persistent-media device.

So, when the host is busy rewriting or even refreshing the DRAM, and when the power fails, the device has enough power to persist what's in DRAM.

The devices shipping now that support this sort of persistent media are called non-volatile dual in-line memory modules, or NVDIMMs. While the main promise of persistent media is speed, it also can protect against power interruptions, allowing end users and processes to work with data even if the power is interrupted to a given device.

There are three types of NVDIMM devices: NVDIMM-N, NVDIMM-P, and NVDIMM-F. NVDIMM-N is a memory-mapped DRAM with no system access to flash. It features a low-capacity of 2 to 32 GB and very low latency of tens of nanoseconds. NVDIMM-F is memory-mapped flash while DRAM isn't system-mapped, featuring a high capacity of 100 GB to one terabyte and a low latency of tens of microseconds. Finally, NVDIMM-P is memory-mapped flash and memory-mapped DRAM with both byte and block drive-oriented access. It has the same high-capacity ability as NVDIMM-F, but carries a low latency of hundreds of nanoseconds.

PARADIGM SHIFT

Persistent Memory media is surely a sea change in our computing paradigm. A similar transition came when SSDs came along. NAND (Not-And) devices were placed into enclosures

and connected to computing devices with the same form factor as the previous spinning disk technology. Computing devices use the same interface, a serial-attached small computer system interface (SAS) or Serial AT Attachment (SATA), to talk to it. Suddenly, all of the applications that were on the hard drive work with solid state, only a lot faster.

The problem was, however, that without changes to the software, these new media devices were too fast for the underlying computer architecture. There was a latency limitation with SAS and SATA interfaces, as well as the driver stack above them, which limited the performance in that form factor and using those interfaces. A brand-new interface just for SSDs called NVMe, or NVM Express, had to be created. Now, devices can take full advantage of SSD performance.

NVMe SSDs work on a PCIe bus, so computing devices aren't restrained by lower-performance hard-disk drives. Consequently, new issues need to be resolved, like dissipating the heat from a bunch of NAND devices. Now a number of new form factors have the promise of fully taking advantage of solid-state storage.

The same will likely happen with new persistent media, but this time as well, there needs to be software changes to get full performance from the new paradigm.

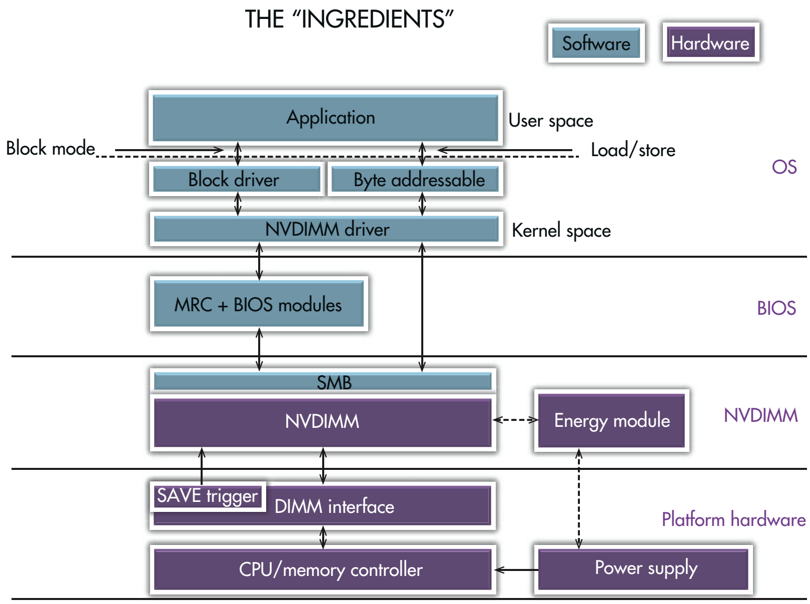

The "ingredients" of persistent media include platform hardware, OS and software.

SOLVING THE SOFTWARE

Software, from the operating system on up, needs to be rewritten to deal with new, lower latency and the end of the need to write data to slower non-volatile destinations. Many applications today assume that memory isn't persistent. Therefore, sharing that data with another machine or another instance of the application, the cluster, or somewhere, requires that it's already coded to write it out to a block device. The app will write data to memory and use it in memory, but at some point, it needs to be checkpointed and written out to some persistent media, typically an SSD or a hard drive.

For software to take advantage of persistent memory, it will have to change the operations and realize that when its reading and writing in memory, that's it. It's already persistent. So, code won't have to stage it out to another device.

The operating system itself needs to understand that memory is persistent, too, so it can pass that persistent-memory model to the applications it runs. It must add new features to the language itself to accommodate and optimize for this new model.

SNIA MODEL

An NVM programming model developed by SNIA explains the technical issues involved. For example, there's a command called flush. The reason we use flush is because some of your data might be in memory and some might be in a processor cache. So, if your data is still in the processor cache, it's not persistent, but if you flush it out to memory, then it can be persistent. Therefore, where it once had to execute a checkpoint and send it out to the SSD, now it just has to do a checkpoint, or a flush.

That gets all of the data out of the cache and into the DRAM, for one thing. And if the power fails right then, the application and its data are saved in NAND flash memory, and its data would still exist. When the application comes up again, what it thought was in memory is still there.

The bottom line is that your applications will run several orders of magnitude faster than if you're writing out to SSDs. It will initially cost more to implement this new memory architecture. That might make some business people worry. Businesses will have to pay more, but they will get more performance, and maybe the improved performance will actually make it possible to charge more, which is usually a good business choice.

Ultimately, it's about making a tradeoff. Some things won't change in an app, but others will need to be modified if they're to take advantage of this performance. Thus, you could make some interim steps that will result in better performance without having to change anything.

> The Storage Networking Industry Association (SNIA) is an organization made up of storage networking professionals. It provides a wide variety of whitepapers and educational documents, including the "NVDIMM Cookbook: A Soup-to-Nuts Primer on Using NVDIMMs to Improve Your Storage Performance."

Swapping out all of your SSDs for persistent-media devices will be very expensive, especially now, when the technology is new. Changing all of your software to take advantage of persistent media will also involve a layout of capital, so you might find a middle path and just upgrade the devices and software that offer the best return on your investment.

WHAT IS SNIA?

The Storage Networking Industry Association (SNIA) is an organization made up of storage networking professionals. It provides a wide variety of whitepapers and educational documents, including the "NVDIMM Cookbook: A Soup-to-Nuts Primer on Using NVDIMMs to Improve Your Storage Performance."

The association has a technical work group that's developing specifications, white papers, and educational material around the NVM programming model. The Solid State Storage Initiative, which is a group meant to go out and promote adoption of solid-state and persistent memory, is aiding in that effort. The Initiative has a special interest group (SIG) that's focused on NVDIMMs and persistent memory.

The same group has also done a number of BrightTALK webcasts, and presented at professional conferences. They present a substantial track every year on solid-state and persistent memory at the Storage Developer Conference every year in September.

The mission of the Persistent Memory and NVDIMM Special Interest Group (PM/NVDIMM SIG) is to educate on the types, benefits, value, and integration of persistent memory, as well as communicate the usage of the NVM programming model created to simplify integration of current and future persistent-memory technologies. It also hopes to influence and collaborate with middleware and application vendors while developing useful case studies, best practices, and vertical industry requirements around persistent memory. Furthermore, the PM/NVDIMM SIG hopes to coordinate with other industry standard groups to promote, synchronize, and communicate a common persistent-memory taxonomy.

All of that activity and energy is going into creating a future where persistent memory is the norm and persistent media is commonplace. SNIA has formed an alliance with a group called OpenFabrics Alliance (OFA), which focuses on low-latency fabric interconnects between CPUs and servers. Persistent memory gives you the ability to take an NVDIMM from a failed server and plug it into a new server so that everything is back online and available. It helps with data durability and availability—while the NVDIMM is swapped into a new server, the data is unavailable, which may mean the application is unavailable, and customers may not be paying money at that time.

Persistent memory over fabrics allows you to mirror the data that's written to local persistent memory over to remote persistent memory. That means if your server dies and can't be repaired, the application can use the data on a different server. Thus, your application is still available, the data is still available, and your customers are still paying.

That makes it so there's no impact on the customer side. A slight hiccup could occur when an app finally determines that the server really is down while it has to switch over to the secondary, but that happens pretty quickly—a lot faster than unplugging the NVDIMM and putting it in a different server.

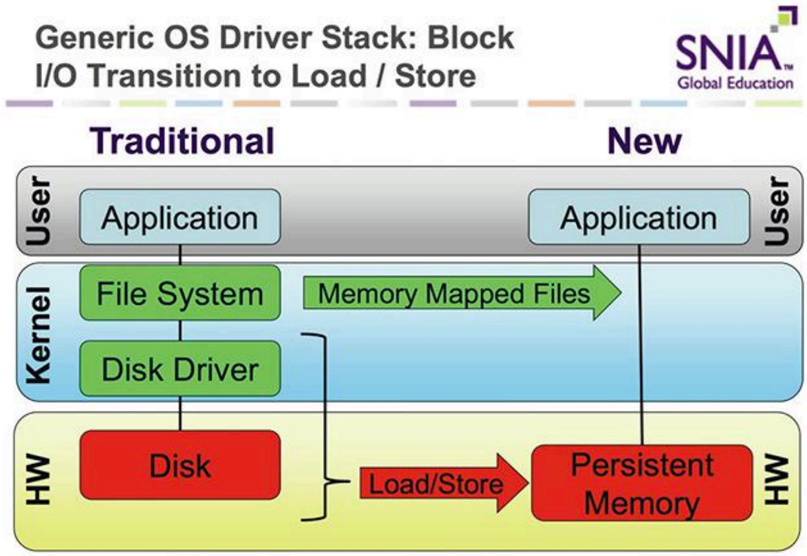

Persistent memory will change the way applications transition to a new stack.

GAZING INTO THE CRYSTAL BALL

In the distant future, it's not hard to imagine how much faster your laptop would be to get to a specific task without the need to boot up. A laptop would pick up wherever it was before, meaning whatever windows you had up and wherever your cursor was on the screen—all exactly the same as when you powered it off.

The cloud will benefit from the use of persistent memory, too. The new technology is going to make it so that any given machine can support many more clients. Businesses can have more customers deferring the cost of that one server by being able to service a greater influx of customers. All they give you is a slice, and that slice is based on how fast things can get done, and how fast they can take your data and put it on that server. If that data is already sitting in persistent memory, it's just running right now. It can start generating income right away since time isn't spent on loading your applications and data.

That will help businesses decrease overhead as well. Typically, with a normal boot, devices have to read in a little piece and then a bigger piece, start running the operating system, subsequently start running the applications, and, finally, present the prompt (or the web page). During that time, customers aren't being charged; there's no work happening. Persistent media can eliminate that downtime on every server. As a result, more applications can be run on every server, which will ultimately help defray the cost of switching over to a persistent-memory paradigm.

Ultimately, while persistent memory and associated media is just coming into play, the technology itself will have far-reaching implications for our increasingly computerized and cloud-based world. Faster, durable, and mobile memory devices mean more effective computing abilities across applications in both business and private domains.

Companies, after some costly upgrades to both hardware and software, can begin to reap the benefits with more efficient systems, far less noticeable downtime or failure, and increased ability to serve more customers across a variety of domains. As we see the tech move into personal use devices like laptops and possibly tablets and handheld devices, persistent memory is the forefront of a new, more powerful way to manage our computing needs. ed

---

TO LEARN more about the work that SNIA is doing around these topics, please visit our website at www.snia.org/pm.


WHEN YOU'RE CREATING
ENERGY EFFICIENT DEVICES THAT
SENSE, CONNECT AND CONTROL
THINK ON
ON

LOU FRENZEL | Contributing Editor

Whatever Happened to These Once Hot Communications Technologies?

Though 2G/3G cellular, WiMAX, and HomePlug are still around, they've moved to the background.

Change is inevitable. You're only as good as your last successful product or technology. If you're a follower of communications technologies, you know what I am talking about. As new and better devices and techniques are discovered and emerge, they leave the older standards to fade away. New competitive technologies quickly make the existing and older technologies obsolete. That's how the modern electronics industry works these days.

Several once promising technologies come to mind when considering this phenomenon. These are the 2G and 3G cellular standards, WiMAX, and the power-line communications (PLC) technology HomePlug. All have faded into near oblivion as they've been superseded by newer and better technology. For whatever it's worth, here's a brief look at these technologies to bring closure to your curiosity.

2G/3G CELLULAR

The second-generation cellular was digital technology that replaced the first-generation analog FM systems. Several digital standards were developed based on a time-division multiple-access (TDMA) scheme (e.g., IS-136). Most of these enjoyed a short success, but quickly the European GSM (Global System for Mobile) communications technology gradually replaced them all.

It's still around today, mainly in Europe and some developing countries. In the U.S., though, it's mostly been phased out. AT&T turned off their GSM network in January of 2017. Verizon says it will keep GSM on until the end of 2019, and T-Mobile will maintain its GSM network until the end of 2020. Sprint plans to hang on until the end of 2021. The main reason for maintaining 2G is probably the companies' responsibilities toward fulfilling the machine-to-machine (M2M) contracts.

Also on the way out are the 3G networks that were supposed to replace 2G. 3G cellular systems used spread-spectrum techniques like code division multiple access (CMDA). Two standards emerged—the CDMA of Qualcomm called cdma2000 and wideband CDMA standard of the Third Generation Partnership Project (3GPP). Verizon and Sprint built cdma2000

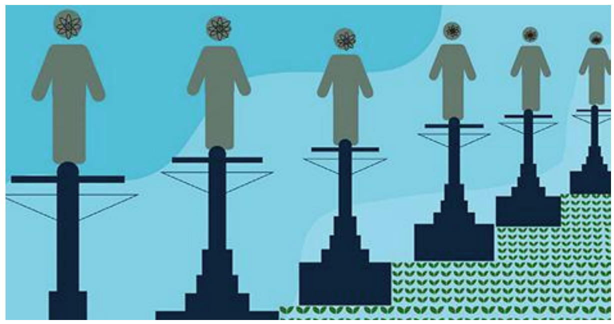

networks while AT&T and T-Mobile made UMTS W-CDMA networks. Both camps were successful, bringing higher data rates and increased network capacity.

Most of these 3G networks are still in operation, and the majority of phones have a 3G fallback provision if a local LTE network is not available. (I experienced this myself recently in the boonies of Wyoming.) The 3G technologies are also used in some M2M and Internet of Things (IoT) applications. But, invariably, these 3G networks will be phased out.

3G may last longer than desired simply because many LTE networks still don't support voice. A voice call from an LTE phone is usually handled by 3G. While some of the more advanced LTE networks have implemented voice over LTE (VoLTE), many still switch you over to 3G for a voice call. In addition, the carriers are encouraging M2M and IoT customers to use the newer Cat-M and NB-IoT modules that are compatible with their LTE networks. Carriers are no longer certifying 3G products for their networks.

The core technology in most cellular networks today is Long Term Evolution (LTE) or what we all call 4G. It uses the more spectrally efficient orthogonal frequency-division multiplexing (OFDM) to get higher data rates, increased capacity, and greater reliability. LTE Advanced, a further enhanced version, has already been implemented with more operators still to adopt.

Eventually, though, the 5G New Radio technology being developed by the 3GPP will emerge around 2020 and gradually work alongside LTE and eventually replace it. Who needs

Design Note

42V, 2A/3A Peak Synchronous Step-Down Regulators with 2.5μA Quiescent Current and Ultralow EMI Emission

Dong Wang

Introduction

The LT®8609, LT8609A, LT8609B and LT8609S are synchronous monolithic step-down regulators that feature a wide 3V to 42V input range. This device family is optimized for applications requiring low EMI, high efficiency and small solution size—suitable for demanding automotive, industrial, computing and communications applications. All regulators in this series have the same 2A continuous, 3A transient (<1 second) load current capability. Their features are summarized in Table 1.

The LT8609, LT8609A and LT8609S feature 2.5μA ultralow quiescent current, important for battery-powered systems. With integrated top and bottom N-channel MOSFETs, the regulators exhibit impressive light load efficiency. The LT8609B operates in pulse-skipping mode only, with higher quiescent current than the other devices, but offers lower ripple during light load operation.

All can pass CISPR 25 Class 5 radiated EMI regulation, the most rigorous EMI standard for automotive equipment. Furthermore, the LT8609, LT8609A and LT8609S feature spread spectrum frequency operation to reduce EMI peaks. The LT8609S displays the most impressive EMI performance in this family, based on its proprietary Silent Switcher®2 technology, described below.

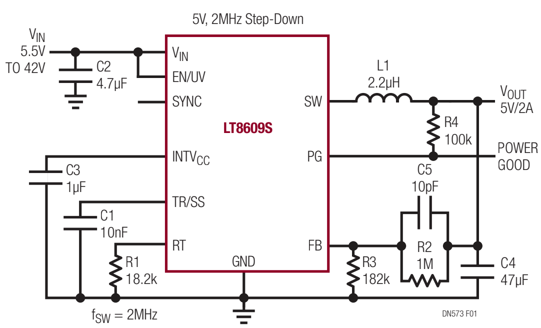

Figure 1. Ultralow EMI Emission LT8609S 12V to 5V Synchronous Step-Down Converter

5.5V to 42V Input, Low EMI, High Efficiency 5V, 2A Supply

A 5.5V–42V input to 5V/2A output power supply is shown in Figure 1. This solution features a 16-lead LT8609S regulator with a 2MHz switching frequency. Only a few components are required for the complete solution, including inductor L1 and a few passive components. Figure 2 shows that this solution can achieve 92.9% peak efficiency.

Burst Mode Operation to Improve Light Load Efficiency

During light load operation and no-load standby mode, high efficiency and low idle current are important for battery-powered applications. The LT8609, LT8609A and LT8609S feature a low 2.5μA quiescent current in Burst Mode® operation. During light load and no-load conditions, the switching frequency is gradually reduced, greatly reducing power loss while keeping the output voltage ripple relatively low. Figure 2 shows that light load efficiency remains above 85% while the power loss approaches zero at minimal loads.

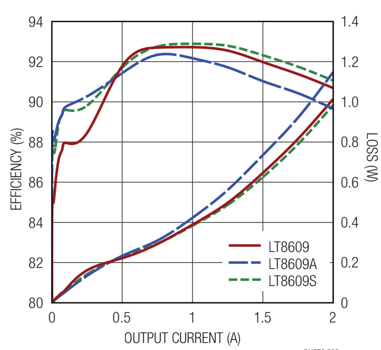

Figure 2. Efficiency vs Load Current for LT8609/09A/09S Based $12V_{IN}$ to $5V_{OUT}$ Step-Down Converter

High Switching Frequency with Ultralow EMI Emission & Improved Thermal Performance

EMI compliance is a concern in a number of environments, including automotive systems. With integrated MOSFETs, advanced process technology, and up to 2.2MHz operation, all of the devices in this family can achieve small solution size while satisfying the most stringent EMI standards. Spread spectrum frequency operation, which reduces EMI peaks, is available in all but the LT8609B. Furthermore, the LT8609S incorporates Silent Switcher 2 technology. Silent Switcher 2 devices feature integrated hot loop and warm loop caps to make EMI performance insensitive to board layout and the number of board layers. A board with fewer layers can be used to reduce manufacturing costs without sacrificing EMI and thermal performance.

Figure 2 shows that the LT8609S features the best peak and full load efficiency of the device family. Figures 3 and 4 show a CISPR 25 EMI and thermal performance comparison of the Figure 1 solution on 2- and 4-layer boards.

Conclusion

The devices in the LT8609 family are easy-to-use monolithic step-down regulators with integrated power MOSFETs and built-in compensation circuits. They are optimized for applications with wide input voltage range and low EMI noise requirements. Low, 2.5μA quiescent current and Burst Mode operation make them excellent battery-powered step-down converter solutions. 200kHz to 2.2MHz switching frequency range make them suitable for most low power to micropower applications. Integrated MOSFETs and up to 2.2MHz switching frequency capability minimize solution size. CISPR 25 scanning results show excellent radiated EMI performance, compliant with the most stringent EMI standards. Silent Switcher 2 technology in the LT8609S makes its performance immune to layout and layer changes, greatly reducing development and manufacturing costs.

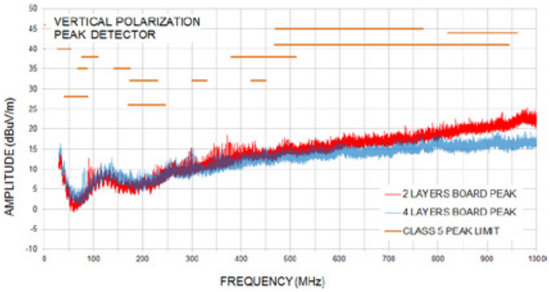

Figure 3. CISPR 25 Radiated EMI Performance Comparison Between 2- and 4-Layer Boards for the Circuit in Figure 1

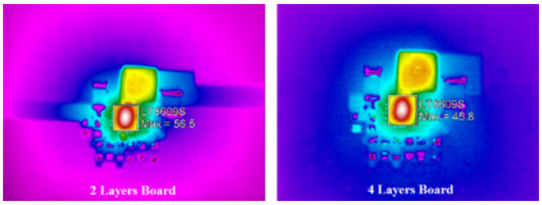

Figure 4. Thermal Performance Comparison Between 2- and 4-Layer Boards for the Circuit in Figure 1

Part	Package	Performance	Operation Mode
LT8609	MSE-10	High efficiency	Burst Mode operation Pulse-skipping mode Spread spectrum mode Sync mode
LT8609A	MSE-10	Optimized for both efficiency and EMI performance	Burst Mode operation Pulse-skipping mode Spread spectrum mode Sync mode
LT8609B	MSE-10	High efficiency	Pulse-skipping mode
LT8609S	LQFN-16	Silent Switcher 2 technology incorporated with best efficiency and EMI performance	Burst Mode operation Pulse-skipping mode Spread spectrum mode Sync mode

Data Sheet Download
www.analog.com/LT8609

THE TOP 50
Electronics Distributors

2018 Platinum Sponsor MOUSER ELECTRONICS

	Company	2017 Global Revenue
1.	Arrow Electronics, Inc.	$26.813 billion
2.	WPG Holdings LTD	$17.51 billion
3.	Avnet *	$17.44 billion
4.	Future Electronics Inc. **	$5 billion (EST)
5.	Digi-Key Electronics	$2.33 billion
6.	TTI	$2.25 billion
7.	Electrocomponents plc/Allied Electronics & Automation	$1.87 billion
8.	Smith & Associates	$1.5 billion
9.	Mouser Electronics	$1.34 billion
10.	Premier Farnell/Newark	$1.195 billion
11.	Rutronik Elektronische Bauelemente GmbH	$1.17 billion
12.	DAC/Heilind	$857.4 million
13.	Fusion Worldwide	$525 million
14.	Sager Electronics	$272.9 million
15.	Master Electronics	$245.2 million
16.	America II Electronics	$238.2 million
17.	PEI-Genesis	$215 million
18.	Advanced MP Technology	$189 million
19.	Bisco Industries Inc.	$156.95 million
20.	The Powell Electronics Group	$138.1 million
21.	Richardson Electronics, Ltd.	$136.9 million
22.	Electro Enterprises Inc.	$127.4 million
23.	RFMW Ltd.	$119.4 million
24.	NewPower Worldwide	$107.3 million
25.	Flame Enterprises	$102 million
26.	Steven Engineering, Inc.	$99.15 million
27.	Hughes Peters	$98.5 million
28.	Classic Components Corp	$90 million
29.	CoreStaff Co., Ltd.	$70 million
30.	Crestwood Technology Group	$56 million
31.	Marsh Electronics	$52.7 million
32.	Edge Electronics, Inc.	$48.4 million
33.	IBS Electronics, Inc	$47.7 million
34.	Symmetry Electronics, Corp.	$44 million
35.	NRC Electronics, Inc.	$41.6 million
36.	SMD Inc.	$39.4 million
37.	Air Electro Inc.	$29.48 million
38.	ASAP Semiconductor	$28.9 million
39.	Area51-ESG, Inc.	$24.96 million
40.	House of Batteries	$24.84 million
41.	Diverse Electronics	$24.5 million
42.	March Electronics, Inc.	$22 million
43.	Powertech Controls Co. Inc	$21 million
44.	CTrends, Inc.	$19.5 million
45.	Cumberland Electronics Strategic Supply Solutions Inc (CE3S)	$18.65 million
46.	Kensington Electronics	$18.5 million
47.	PUI (Projections Unlimited Inc.)	$18.47 million
48.	Gopher Electronics Company	$16.1 million
49.	Marine Air Supply	$14.8 million
50.	Advantage Electric Supply	$13 million

*In October 2016, Avnet closed its acquisition of Premier Farnell and the company's total includes Premier Farnell's revenue for fiscal year 2017. Source Today also lists Premier Farnell's global revenue separately
**SourceToday estimate of Future Electronics global revenue

2G/3G today when most cellular activity requires fast digital data rates? LTE will be around for many years as 5G gradually takes over. Fixed wireless access first, then handsets and mobile service, and finally IoT and M2M services. Until then, your only choice is LTE.

WIMAX

Worldwide Interoperability for Microwave Access, or WiMAX, came on the scene about the same time as LTE. For a while, the two technologies were competing for the next cellular standard. Sprint/Clearwire even had a rather extensive WiMAX cellular network, but eventually abandoned it in favor of LTE. LTE won the standards battle, big time. Both technologies use OFDM; however, LTE is primarily a frequency division duplex (FDD) technology, while WiMAX uses time division duplex (TDD).

WiMAX, an IEEE standard (802.16) is now known primarily as a last mile, fixed wireless access technology to be used for internet access where it competes with cable and DSL. Its data rates typically max out at 30 to 40 Mb/s. However, recent upgrades to the standard now provide data rates to 1 Gb/s. WiMAX is still available, but mostly outside the U.S. For more updates, go to the WiMAX Forum at *www.wimaxforum.org*.

HOMEPLUG

HomePlug is a PLC standard. It transmits data over the ac power lines within a home or building. The technology uses a form of OFDM that's superimposed on the ac mains 60-Hz voltage. Its primary use is as a home-networking technology for connecting to the internet. Simply plug the modem into an outlet and its signal will reach another modem at another outlet via the home wiring. No other wiring is needed.

The HomePlug technology is standardized as IEEE 1901. Multiple other PLC versions and standards are used in industry.

Despite the convenience of no extra wiring and data rates up to over 1 Gb/s in the latest form, HomePlug just didn't catch on. With de facto Wi-Fi so cheap, convenient, fast, and featuring regular upgrades, HomePlug didn't stand a chance. Yet HomePlug is still available. For more details, go to *www.homeplug.org*.

It's routine for multiple technologies to emerge and compete for some new application and standard. One eventually wins the battle while the others either fade away or find new but more limited applications. Sometimes it pays to hold off in committing to a technology until it emerges as the winner. On the other hand, those who bet on the winner early in the game get a head start and are usually more successful financially. It's a risk most companies take at some point. Shine up your crystal ball for the best results. ed

FREE WEBINAR

August 16, 2018 | 2:00pm ET

Registration URL: https://event.on24.com/wcc/r/1797416/CD833766CB3AB7A6BA1377ABACC4B581

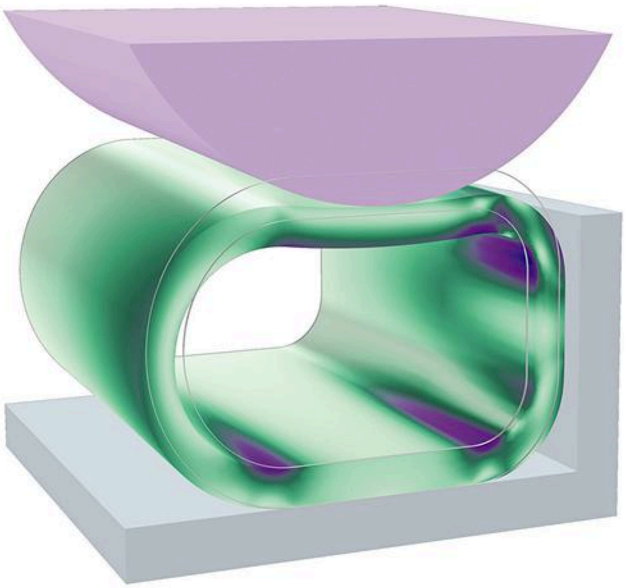

Visualization of deformation and von Mises stress in a hyperelastic car door seal.

NONLINEAR STRUCTURAL MECHANICS MODELING

Nonlinear structural analysis is concerned with structures where structural response, such as deformation or stress, is not proportional to inputs like applied loads. This behavior arises from:

- Nonlinear constitutive behavior, such as plasticity and hyperelasticity
- Large deformation, such as buckling
- Nonlinear boundary conditions, such as contact and friction

In this webinar, we will discuss the built-in and user-defined features that the COMSOL® software provides to capture nonlinear structural behavior and will highlight some best practices. The webinar includes a live demo in the COMSOL Multiphysics® software and concludes with a Q&A session.

SPEAKER:

Temesgen Kindo, *Applications Engineer, COMSOL*

Temesgen Kindo works at COMSOL as an applications engineer, specializing in structural mechanics and equation-based modeling. Before joining the Burlington, MA office, Temesgen received his PhD in civil engineering from Duke University. He is interested in nonlinear mechanics, multiphysics problems, and applied mathematics.

Hosted with
ElectronicDesign

COMSOL

BOB ZOLLO | Solution Architect for Battery Testing, Electronic Industrial Solutions Group, Keysight Technologies Inc.

Fixture vs. Charger Design Tradeoffs When Charging Li-ion Cells

Careful attention to all aspects of system design needs to be paid to optimize for capital costs, operational costs, and ease of design and construction.

When forming lithium-ion cells in manufacturing, or when testing large quantities of cells for evaluation, it's important to select your charging and discharging systems with the proper size for the task. In manufacturing, where you may be charging and discharging thousands of cells at a time, total power consumption and heat generation must be considered. Oversizing a system to give margin may seem like a good idea, but you must take into account the additional capital cost, system size, and operational expenses.

Two key factors that need proper sizing are the maximum voltage of the charging circuit and the minimum voltage of the discharging circuit. Let's look at each of these in more depth.

The cell itself will be the determining factor of the voltage that's needed to charge the cell. For example, the maximum charging voltage may be 4.2 V. This voltage, of course, is the maximum voltage that must appear on the terminals of the cell, but this will not be the maximum voltage of the charging system.

The charger must be able to source more voltage to overcome voltage drop in the wires and the contacts to the cell.

Charging	
V_drop =	I_in_wire * (R1 + R2 + R3 + R4)
V_out =	V_cell + V_drop
Power in Wire =	V_drop * I_in_wire
Power in Cell =	V_cell * I_in_wire
Power in Charger =	Power in Cell + Power in Wire

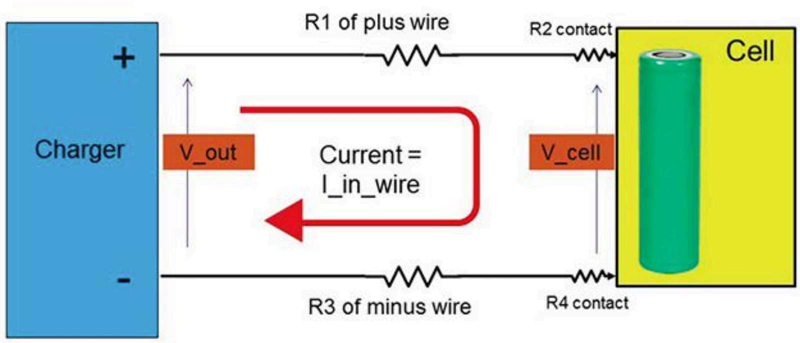

1. Voltage drops in the wires when charging a cell.

The available voltage at the terminals of the charger will need to be the sum of the maximum voltage on the cell (4.2 V in our example) plus the voltage drop in the wires, which will be the resistance of the wires times the maximum current flowing through the wires *(Fig. 1)*.

Your first inclination may be to specify a charger voltage that is 2 or 3 V higher than the cell's maximum charge voltage. This will give you a wide margin in the resistance that you can have in your wiring and contacts. By accommodating a high resistance wiring path, you can then use longer wires or thinner wires, simplifying fixture design.

However, while having extra voltage on the charger to accommodate a higher resistance path seems like a benefit, it comes at a high cost. First, consider the

extra power needed by the charger. The power in the wires, which is V_drop × I_in_wire, must be added to the power being supplied by the charger. You will also need to buy a larger charger that can source more power to supply what was lost in the wires. You will pay for this power in increased electricity consumption.

Next, you must consider that the power in the wires will be lost as heat. Normally, the wires will be bundled together, and this heat can build up. Lastly, all of the excess heat generated by the more powerful charging circuit and the loss in the wires will add a larger heat burden on the environmental systems/air conditioning, which itself must be sized larger and will cost more (increased capital), as well as require more electricity to operate (increased operational expense).

THREE POSSIBLE SCENARIOS

Let's look at three cases:

Case 1: Charging a 3-AH 18650 cylindrical cell to 4.2 V, using 18 AWG wire of 20 feet total length (includes both positive and return path) and 20 mΩ of contact resistance.

Case 2: Charging a 100-AH pouch cell to 4.2 V, using 4 AWG wire of 20 feet total length (includes both positive and return path) and 20 mΩ of contact resistance.

Case 3: Charging a 100-AH pouch cell to 4.2 V, using 0 AWG wire of 8 feet total length (includes both positive and return path) and 5 mΩ of contact resistance.

So, as you can see, simply deciding to give more voltage margin on wiring design to allow for an easier design with increased resistance of the wires has many ramifications on operational expenses and capital expenses. You must consider the tradeoffs.

WHAT ABOUT DISCHARGING?

Now, let's look at discharging the cells. In this case, while the resistance of the wires still causes a voltage drop when discharge current flows, the voltage drop reduces the voltage on the cell

THREE CHARGING SCENARIOS

Parameter	CASE 1 When charging a 3000-mAH 18650 cylindrical cell	CASE 2 When charging a 100-AH pouch cell with long wires, too thin wire, poor contacts	CASE 3 When charging a 100-AH pouch cell with short wires, thicker wires, better contacts
Max voltage on the cell (V)	4.2	4.2	4.2
Max charge current (A)	3	100	100
Wire size	18 AWG	4 AWG	0 AWG
Wire resistance per foot	0.0064 Ω/ft.	0.0002486 Ω/ft.	0.0000983 Ω/ft.
Wire length (sum of wire length in the positive path and the return path)	20 ft.	20 ft.	8 ft.
Contact resistance (Ω)	0.02	0.02	0.005
Total resistance (Ω in wire)	0.148	0.024972	0.0057864
Voltage drop in the wire (V)	0.444	2.5	0.58
Power dissipated in the wire (W)	1.332	250	58
Voltage required at the charger (V)	4.644	6.7	4.78
Power delivered to the cell (W)	12.6	420	420
Power supplied by the charger (W)	13.932	670	478
Conclusions	For low current, wire size and length have minimal effect		Switching to larger wire, making wires shorter, and lowering contact resistance can reduce lost heat to 1/5th and reduce size of charger by 30%

when it gets to the terminals of the discharger circuit. This means the power of the discharger circuit can be lower, as some of the energy coming from the cell is lost in the wire before it gets to the discharger *(Fig. 2)*.

Generally speaking, discharger circuits are electronic loads. All electronic loads have a minimum operating voltage, typically on the order of 1 to 3 V. This means that you can pull full current into the discharger (load) until you hit the minimum operating voltage. Then, as the voltage at the terminals of the discharge decreases, the available current will be linearly derated *(Fig. 3)*.

So, if you are trying to discharge your cell down to 3.0 V, and if you have 2.5 V of drop in the wires (as in Case 2), the voltage at the discharger input is only 0.5 V. This is probably below the minimum voltage of the discharger; therefore, the available current will be derated as well.

Discharging	
V_drop =	I_in_wire * (R1 + R2 + R3 + R4)
V_out =	V_cell - V_drop
Power in Wire =	V_drop * I_in_wire
Power in Cell =	V_cell * I_in_wire
Power in Discharger =	Power in Cell - Power in Wire

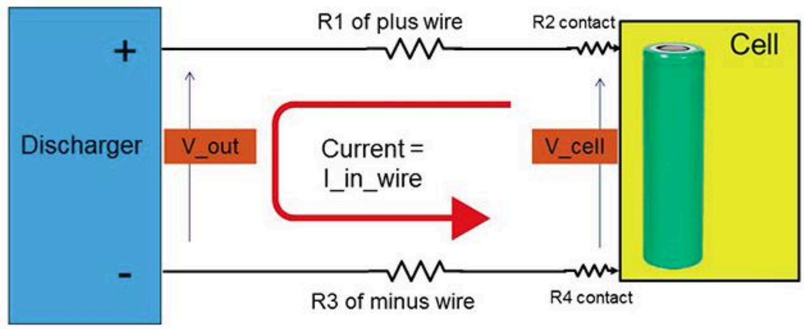

2. Voltage drops in the wires when discharging a cell, too.

You may not be able to pull enough discharge current to meet your testing requirements. It will cause you to have to specify a larger discharger in terms of maximum current, even though you won't be using that maximum current.

A larger discharger will drive up your equipment cost (capital expenses).

Like with charging, heat will exist in the wires, and they will be bundled together, causing the heat to build up.

Thus, once again, a simple decision to give more voltage margin on wiring design to allow for an easier design with increased resistance of the wires has many ramifications. You must consider the tradeoffs.

If your charging and discharging operation is on only a few cells at a time, perhaps the impact on operational costs and capital costs aren't significant. However, when testing many cells or when forming cells in manufacturing, you will have many cells that need to be simultaneously tested. Each of those cells will require its own charger, discharger, fixture, and wiring. Wasted energy, excess heat, and increased equipment cost will be multiplied many times over. ed

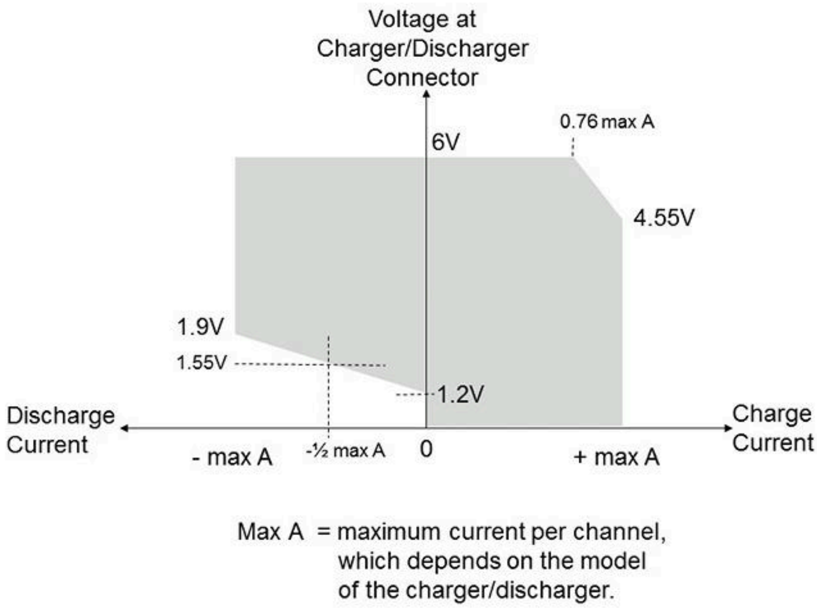

3. Shown is the operating voltage and current of a single channel of the Keysight BT2200 Cell Charger/Discharger System, used for forming cells in manufacturing. While the maximum charging voltage can go to 6 V, the drop in the wire needs to be considered. Note the minimum operating voltage during discharge and the derating from 1.9 V to 1.2 V.

ideas for design

Microcontroller Efficiently Converts Sensor Voltage to Current Loop

RICARDO JIMENEZ | Imperial Valley College, Imperial, CA
JORGE L. AGUIRRE and THOMAS F. MURRIETA | Instituto Tecnologico de Mexicali (ITM)

WHEN TAKING MEASUREMENTS control circuits use some form of flip-flop that responds to a pushbutton switch or other control input. These all have volatile memory and default to an OFF condition if power is turned off, and sometimes that's even the preferred situation. However, if your application must remember what it was doing when power failed and resume from where it left off when power is restored, you may have a problem.

The circuit in *Figure 1* will deliver an output current from 4 to 20 mA for an input voltage (V_{in}) from 0 to 5 V. PNP transistor Q1 (2N2907) is configured as a constant-current source; its emitter current and resistor value (R_e) are defined by the equation:

$$R_e = \frac{V_{cc} - V_b - V_{be}}{I_e}$$

where V_{cc} = 5 V, and V_{be} = 0.7 V.

Transistor base voltage V_b must be within the range of 3.64 V to 1.00 V to control the emitter current from 4 to 20 mA. Thus, a 165-Ω resistor is the appropriate value to generate that current loop. Since V_b will be generated within that range, a PIC microcontroller is selected to control the transistor's base voltage. A microcontroller that suits these needs is the PIC16F1783, because it has both analog-to-digital and digital-to-analog converters (ADCs and DACs).

The ADC input will be reading the sensor's input voltage on channel AN0, and with some processing, it will then control the voltage output in the DAC module. An op amp (MCP602) is configured as a driver and used to buffer the DAC's analog output. The 8-bit DAC will provide a voltage output with a resolution of 20 mV, while the 10-bit ADC has a resolution of 4.88 mV. Therefore, the ADC reading is divided by 4 to make it compatible, and the output voltage will be controlled by the sensor's input voltage.

The DAC voltage references are set to V_{Ref+} = 3.64 V and $V_{Ref}-$ = 1 V, using a pair of multiturn trimmer potentiometers. This configuration makes the DAC output vary from 3.64 V to 1.0 V. As the DAC output voltage decreases from 3.64 V to 1 V when the sensor's input voltage increases, the PIC micro delivers a voltage from 1 V to 3.64 V. Thus, the output voltage must be inverted, so that the ADC reading is stored in variable VOLT. The ADC's binary reading in this variable goes from 0 to 255; knowing this, we get the equation:

(Continued on page 44)

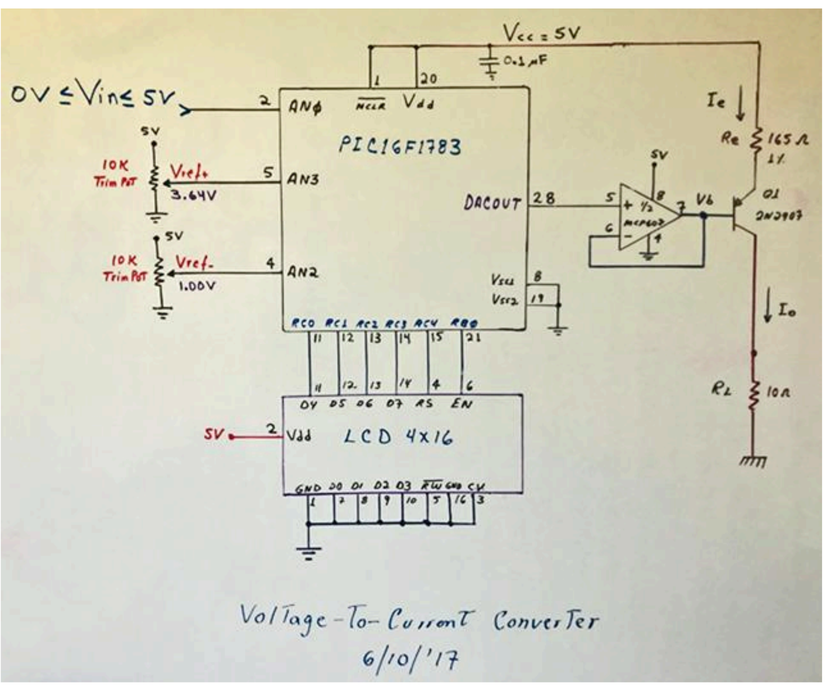

1. This circuit, based on a PIC microcontroller with a few added components, converts a sensor-based voltage into a 4- to 20-mA current loop. Shown is the circuit as originally drawn on a whiteboard.

Microcontroller's Code

Listing 1: Microcontroller's code for the PIC16F1783 to control the current source.

```
'*  Name    : V-TO-I.BAS                                                     *
'*  Authors : Ricardo Jimenez, Thomas Murrieta, and Luis Guzman  *
'*  Date    : 6/08/2017                                                      *
'*  Version : PBP3 compiler from www.melabs.com                              *
'*  Notes   : PIC 16F1783 VOLTAGE TO CURRENT CONVERTER                       *
OSCCON= %01101011          ;               4 MHz internal oscillator
OSCTUNE=0;
ASM__CONFIG_MCLRE_OFF      ;           MASTER CLEAR PIN CONNECTED TO VDD
TRISA= $FF                 ;               PORTA SET AS INPUT
TRISB= 0                   ;               PORTB SET AS OUTPUT
TRISC= 0                   ;               PORTC SET AS OUTPUT
ANSELA= $FF                ;               PORT A SET AS ANALOG
ANSELB= 0                  ;               PORT B  SET AS DIGITAL
ADCON0  = %10000001        ;             see description in article
ADCON1  = %11110000        ;
ADCON2  = %00001111        ;
DACCON0 = %10010101        ;             DAC SETTINGS WITH EXTERNAL "VREF(+)" AND "VREF(-)"

'//////////LCD CONFIGURATION////////'
DEFINE LCD_DREG PORTC              ;
DEFINE LCD_DBIT 0                  ;
DEFINE LCD_RSREG PORTC             ;
DEFINE LCD_RSBIT 4                 ;
DEFINE LCD_EREG PORTB              ;
DEFINE LCD_EBIT 0                  ;
DEFINE LCD_BITS 4                  ;
DEFINE LCD_LINES 2                 ;
DEFINE LCD_COMMANDUS 1500          ;
DEFINE LCD_DATAUS 44               ;

VOLT VAR WORD: VOLT1 VAR WORD: VOLT2 VAR WORD : V VAR WORD         ;
UNITS VAR WORD: DECIMALS VAR WORD: CENTS VAR WORD: X VAR WORD      ;

MAIN:
   GOSUB adc_read    ;
   X = 255 - VOLT    ;                          X VARIABLE inverts THE VALUE OF Vin
   DACCON1 = X       ;                          DAC output VALUE EQUALS TO X
   GOSUB ANALOG_TO_DECIMAL ;                              GOES TO method ANALOG_TO_DECIMAL
   LCDOUT $FE, $01, "Vin = ",   DEC UNITS, ".", DEC DECIMALS, DEC CENTS, "V"   ; displays input voltage Vin
   GOSUB ANALOG_TO_DECIMAL2                               ; GOES TO SUBROUTINE ANALOG_TO_DECIMAL2
   LCDOUT $FE, $C0, "VOdac = ", DEC UNITS, ".", DEC DECIMALS, DEC CENTS, "V"   ; Displays DAC's output voltage
   gosub OUTPUT_CURRENT                                   ; GOES TO SUBRUTINE OUTPUT_CURRENT
   LCDOUT $FE, $94, "Io = ",  DEC UNITS, DEC DECIMALS, ".", DEC CENTS, "mA"    ;   displays output current Io
   PAUSE 250                   ;                    250 mS pause
   GOTO MAIN                   ;                      GOES TO MAIN method

ADC_READ:                      ;                     method to  READ VOLTAGE INPUT ON AN0
   PAUSEUS 10                  ;                   10 uS pause
   ADCON0.1 = 1                ;                   TURNING ON ADC channel 1
   PAUSEUS 50                  ;                   50 uS pause to wait for analog reading
   VOLT.LOWBYTE  = ADRESL             ;            stores in variable VOLT the ADC low byte
   VOLT.HIGHBYTE = ADRESH             ;              SAVES IN VOLT VARIABLE THE ADC high byte INPUT LECTURE
   VOLT=VOLT/4                 ;                DIVIDE VARIABLE VOLT BY 4 GIVING THE MAX VALUE OF VOLT = 256
   RETURN                      ;

ANALOG_TO_DECIMAL:                                         Binary-to-decimal conversion method
   VOLT1 = Volt * 19           ;                     MULTIPLIES VARIABLE Volt BY 19
   VOLT2 = Volt * 62           ;                     MULTIPLIES VARIABLE Volt BY 62
   VOLT2 = VOLT2/100              ;                     DIVIde VARIABLE Volt BY 100
   V = VOLT1 + VOLT2              ;
   UNITS =     V DIG 3         ;                   SAVES 3rd DIGIT OF VARIABLE "V" IN  VARIABLE UNITS
   DECIMALS =  V DIG 2            ;                 SAVES 2nd DIGIT OF "V" IN THE VARIABLE DECIMALS
   CENTS =     V DIG 1          ;                   SAVES 1st DIGIT OF "V" IN VARIABLE CENTS
   RETURN
ANALOG_TO_DECIMAL2:
   X = X * 100                 ;        "X" EQUALS TO "(255-VOLT)*100" TO GIVE A result IN THOUSANDS
   X = X / 256                 ;              DIVIDE "X" BY "256" GIVING US A VARIABLE IN A RANGE OF "0" AND "100"
   X = (X * 267) + 10000          ;       MUTIPLIES "X" BY 267 AND ADDING 10000, THIS GIVE US A 5 DIGITS VARIABLE
   UNITS =     X DIG 4            ;          SAVES THE 4th DIGIT OF VARIABLE "X" IN VARIABLE UNITS
   DECIMALS =  X DIG 3                       SAVES THE 3rd DIGIT OF VARIABLE "X" IN VARIABLE DECIMALS
   CENTS =     X DIG 2            ;          SAVES THE 2nd DIGIT OF VARIABLE "X" IN VARIABLE CENTS
   RETURN                      ;

OUTPUT_CURRENT:
   X = 43000 - X   ; MULTIPLIES (VCC-VBE)*10000 TO HAVE A  5 DIGITs NUMBER AND SUBSTRACTS THE VALUE OF X
   X = X / 165                 ;       DIVIDES  VOLTAGE "X" BY 165 THAT CORRESPONDS TO THE RESISTOR VALUE
   UNITS =     X DIG 2            ;              SAVES THE 2nd DIGIT OF VARIABLE "X" INTO VARIABLE UNITS
   DECIMALS =  X DIG 1               ;            SAVES THE 1st DIGIT OF VARIABLE "X" INTO VARIABLE DECIMALS
   CENTS =     X DIG 0            ;              SAVES THE Zero DIGIT OF VARIABLE "X" INTO  VARIABLE CENTS
   return                      ;
   END                         ;
```

(Continued from page 42)

DACCON1 = 255 – VOLT

where DACCON1 is the DAC output value. Thus, the output voltage when VOLT is 255 will be 3.64 V as shown by:

$$V_{out} = [(V_{Ref}(+) - V_{Ref}(-))\left(\frac{DACCON1}{256}\right) + V_{Ref}(-)$$
$$V_{Ref}(+) = 3.64\ V$$
$$V_{Ref}(-) = 1\ V$$

From the V_{out} equation, we see that changing the DACCON1 value will change the output voltage value. After controlling the DAC's output voltage, this signal goes to the buffer op amp which drives the base of the Q1 to control the output current, using the 16 5Ω resistor to generate the current which varies from 4 mA when V_b = 1.00 V up to 20 mA when V_b = 3.64 V.

PROGRAMMING THE MICROCONTROLLER

The microcontroller's code *(see listing on page 43)* is written using the PBP3 compiler from *melabs.com*. In the PIC configuration for the DAC module, 8-bit register DACCON0 is used to configure the DACOUTPUT2 and both the V_{ref+} and V_{ref-} pins.

- The 7th bit, DACEN, turns on the DAC if it is cleared to "0," while a logic "1" enables the DAC module.
- The 6th bit is "0" because it's an unimplemented bit.
- The 5th bit, DACOE1, enables DAC output DACOUTPUT2; this bit goes to "0" to disconnect output DACOUT1.
- The 4th bit, DACOE2, enables the DACOUT2 pin; setting it to "1" enables this pin and setting a "0" value disconnects the DACOUT2 pin. In this case, we are using the DACOUT2 pin, as it is set to "1."
- For bits 3 and 2, instruction DACPS<1:0> (embedded in the DACCON0 register) determines the use of an external V_{ref+}, where setting a value of "10" enables the use of a FVR Buffer2 output. A value of "01" uses V_{ref+} while with "00" it uses the V_{dd} as an output range. In this case, we're using V_{ref+} and therefore the value is "01."
- Bit 1 goes to a "0" logic because it is unimplemented.
- The Zero bit, called DACNSS, is used to select the negative source V_{ref-} ; in this case, it is set to "1" to enable the V_{ref-} pin. Thus, the DACCON0 register configuration will be DACCON0= %10010101.

The ADC module has three configuration registers that require specific set up. Here, we want to read an analog input signal ranging from 0 to 5 V.

- For register ADCON0, bit 7 controls the ADC Result Mode and it has two possibilities, 10-bit or 12-bit results. This design

uses the 10-bit result, so the bit is set to 1.

• Bits 6 through 2 select the analog input AN0, so these bits are set 0000.

• Bit 1 is the ADC conversion status. When this bit is set, the module is processing a conversion, and when this bit is set to 0, a conversion is completed. Bit 0 is the ADC Enable bit, and it is set to 1 to enable the ADC module.

• Bit 7 of register ADCON1 is the ADC Result Format, which can be in either two's complement or sign-magnitude format. Since we are using two's complement, this bit is set to 1.

• Bits 6 through 4 are the ADC Conversion Clock Select bits, which are set to 111 because the clock which is supplied from a dedicated FRC oscillator is selected.

• Bit 3 isn't implemented and is set to 0. Bit 2 controls the V_{ref-} setup; as it is connected to V_{ss}, it is set to 0. Bits 1 and 0 control V_{ref+}, which in this case is set to 00 to connect V_{ref+} to V_{dd}.

• For Register ADCON2, bits 7 through 4 are the Auto-conversion Trigger Source Selection bits, which must be disabled; therefore, these bits are set to 0000.

• Bits 3 through 0 are the Negative Differential Input Channel Select bits, here set to 1111 for the ADC Negative reference, controlled by ADNREF.

• Finally, the setup registers are ADCON0= %10000001, ADCON1= %11110000, ADCON2= %00001111.

Figure 2 depicts the circuit assembled on a prototyping board. It includes an LCD readout to show the key variables such as input voltage, DAC output, and 20-mA loop current.

RICARDO JIMENEZ holds a Master's degree in Electronics. He is an Electronics Professor at Imperial Valley College, Imperial, Calif., and is the author of the book "The PIC Microcontroller Engineer's Notebook" (2016), ISBN: 978-06962764190. JORGE L. AGUIRRE and THOMAS F. MURRIETA are in their 4th year at Instituto Tecnologico de Mexicali, pursuing degrees in Electronics Engineering.

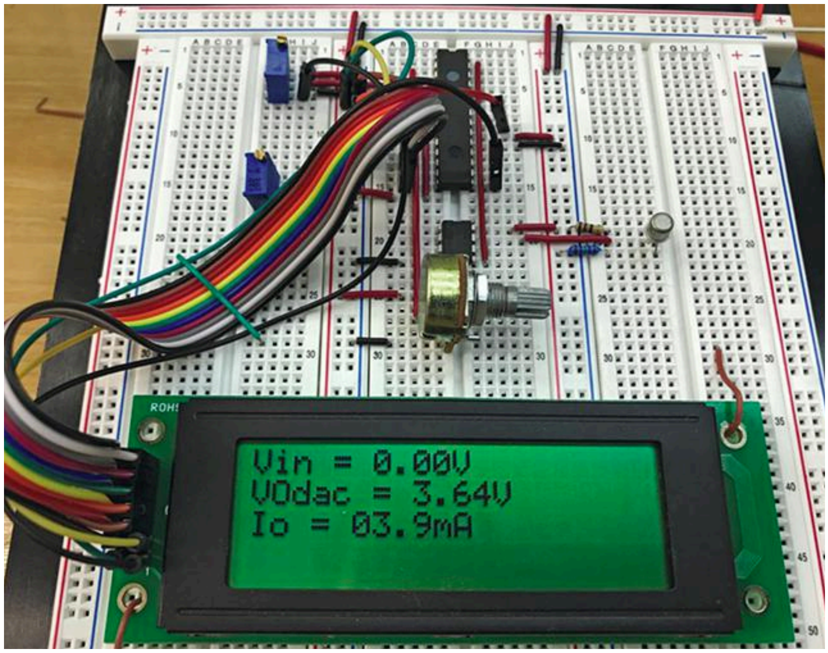

2. The circuit was built on a prototyping board with an LCD readout for convenience. The two small blue rectangles in the upper left are the trimming potentiometers used to adjust the positive and negative reference voltages.

Texas Instruments is looking to fill engineering and related positions working with semiconductor technologies. We are recruiting @ various educational and experience levels:

CITY	POSITION	APPLY FOR JOB USING THIS LINK
Dallas/Richardson, TX	Analog Design Engineer	https://tinyurl.com/y9kguyj4
	Analog Design Manager	https://tinyurl.com/yc9oyqsv
	Application Developer	https://tinyurl.com/y7qlrrwn
	Digital Design Engineer	https://tinyurl.com/ycbw4d4n
	Digital Design Manager	https://tinyurl.com/y9hg6orw
	Failure Analysis Engineer	https://tinyurl.com/y89rx9ew
	IT Security Administrator	https://tinyurl.com/ya95p4pl
	Packaging Engineer	https://tinyurl.com/y8346y85
	Process Development Engineer	https://tinyurl.com/ydejefwo
	Product Engineer	https://tinyurl.com/yaxodg6u
	Product Marketing Manager	https://tinyurl.com/y7y6c5lf
	Software Development Manager	https://tinyurl.com/ycxbprk4
	Systems Engineer	https://tinyurl.com/yburm2u9
	Systems Manager	https://tinyurl.com/ydf4o3zm
Cary, NC	Digital Design Engineer	https://tinyurl.com/ybx98ck7
Duluth, GA	Analog Design Engineer	https://tinyurl.com/y9yo85d7
Longmont, CO	Product Marketing Engineer	https://tinyurl.com/yapyqmhu
Manchester, NH	Analog Design Engineer	https://tinyurl.com/y8afrkzy
Santa Clara, CA	Analog Design Engineer	https://tinyurl.com/y7om4hwp
	Analog Design Manager	https://tinyurl.com/ybufgtak
	Digital Design Engineer	https://tinyurl.com/y8n78rrv
	Digital Design Manager	https://tinyurl.com/y7e76dz2
	Layout Engineer	https://tinyurl.com/y9wrf2lp
Phoenix, AZ	Analog Design Engineer	https://tinyurl.com/y88zvaox
Tucson, AZ	Analog Design Engineer	https://tinyurl.com/y7cfnrc3
	Test Engineer	https://tinyurl.com/yclsxcga
	Validation/Characterization Engineer	https://tinyurl.com/ycq7kok6
Warrenville, IL	Analog Design Engineer	https://tinyurl.com/ycz82u5o
	Packaging Engineer	https://tinyurl.com/yd3crgok
Warwick, RI	Systems Engineer	https://tinyurl.com/ydxwwmez

Some positions may require travel.
Please send in resumes using the individual URL code for the position listed. EOE.

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

Lab Bench Clear-Out: Compact Logic Analyzer and More

Technology Editor Bill Wong takes a hands-on look at a compact logic analyzer, a handwriting calculator app, and sewable electronics.

My Lab Bench has been piling up with stuff to check out, and I needed to clear some space for new stuff. The latest crop includes an economical, compact logic analyzer, a new version of my favorite calculator app and some sewable electronics kits. These will appeal to new developers and makers.

TINY LOGIC ANALYZER

IKA Logic sells a number of low-cost, USB-based logic analyzers and oscilloscopes. I took a look at their top end, 149€, SQ200 logic analyzer *(Fig. 1)*. This 4-channel logic analyzer has a 200-MHz sampling rate and can capture up to 4 Msamples per channel. The system also will operate in a mixed capture/generate mode using a digital pattern generator to drive one or more connections. Of course, you need to trade off inputs and outputs among the four connections.

1. IKA Electronics' SQ200 is a four-channel logic analyzer that captures up to 4 Msamples/channel at 200 MHz.

The SQ200 has support for adjustable input/output voltage and adjustable input resistance, and outputs can support open drain with an optional pull up. Input protection is ±35 V. The system can also handle two differential pair inputs.

These all work with IKA Logic's ScanaStudio *(Fig. 2)*. It runs on Microsoft Windows, Linux, and Apple's Mac OS.

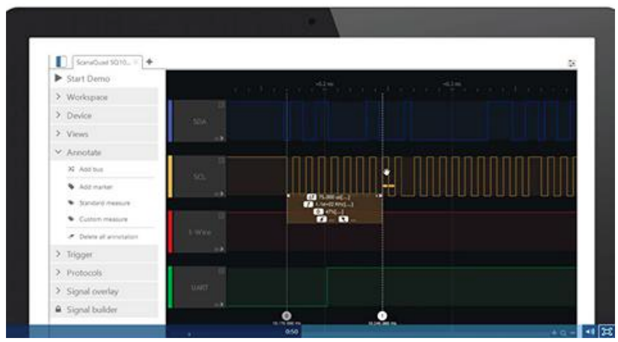

2. The SQ200 and its siblings work with the ScanaStudio.

Less-expensive versions are available with fewer inputs, smaller capture buffers, etc. The top end supports the widest range of trigger options including edge detection, level change, pulse, pattern, and serial protocol such as I^2C, SPI, RS-232, CAN, 1-Wire, and JTAG. It was easy to use ScanaStudio to trigger on patterns like I^2C address and CAN frame IDs. The protocol support is open source and available on GitHub.

Getting started with ScanaStudio and the SQ200 for basic capture and generation was relatively easy, since the number of options is significantly less than high-end logical analyzers. Still, for microcontroller projects, platforms like the SQ200 are more than sufficient and significantly less expensive.

The ScanaStudio interface was easy to use and has a good complement of time-measurement-related features, as well as protocol analysis and presentation. The Signal Overlay feature is interesting—it allows for editing of captured streams to see how changes affect timing. This can be handy for creating data streams to be sent as output streams.

In addition, the system can be extended using JavaScript, although I didn't have a chance to write any of my own code. Users can suggest features and new features are made available to ScanaStudio users.

The signal-generator script build IDE is actually used to create Javascript functions that generate signals. For example, it's possible to create a transaction for an I2C slave device using the SQ200 without having a microcontroller manage the transaction. This can be handy for testing devices or scenarios. The scripts can generate error logs and messages, which is useful for regression testing.

Overall, the system is very flexible and really limited by the capture rate. It's more than sufficient for microcontroller work, easily handling low-speed serial protocols.

IMPROVED HANDWRITING

I've taken a look at some of MyScript's handwriting technology in previous articles. Its Nebo editor uses MyScript's Interactive Ink. It works best with a tablet and pen. The Calculator app is free and runs on tablets and smartphones that work with Android and iOS.

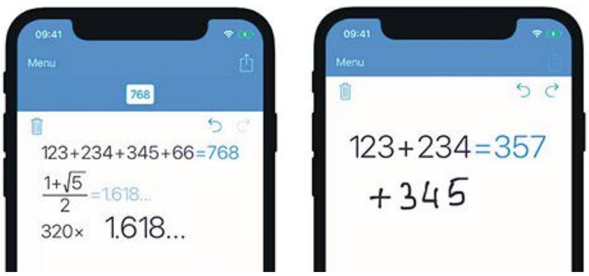

3. MyScript Calculator 2 supports limited equation-solving abilities as well as multiple expressions.

MyScript Calculator 2 *(Fig. 3)* is the latest addition. It adds a significant number of features compared to the original app, which enabled you to write arithmetic expressions that were then solved. Unfortunately, the app only lets you do one at a time. Calculator 2 makes it possible to write multiple expressions on the screen, and thus it's easier to remember results. It also allows chained modifications, such as extending an existing expression.

The new version retains a history of calculations. They can also be exported as text or as an image. The system is moving toward spreadsheet-like functionality with the ability to handle time-order independent calculations of multiple sub-equations on a page. Edit a value and the system changes the results from related equations accordingly.

Calculator 2 works well on an iPhone, but a tablet provides more workspace for more complex calculations. Though more detailed work is best with a pen, fingers work just fine in most instances.

For now, Calculator 2 is available for Apple users. Android users still have version 1.x available for free. Try out whatever version your device can handle. It's much quicker than using a calculator app.

I've taken a look at some of MyScript's handwriting technology in previous articles. Its Nebo editor uses MyScript's Interactive Ink. It works best with a tablet and pen.

SEWABLE ELECTRONICS

I didn't actually test out Sparkfun's sewable electronics kits, but I looked at them before giving them out at my presentation on wearable electronics. I had mentioned some of the other kits I gave away there, but didn't have time to write up these until now.

They included the $49 Sparkfun LilyPad ProtoSnap Plus Kit *(Fig. 4)* and the $99 LilyPad Sewable Electronics Kit. They are both based on LilyPad USB Plus microcontroller that has an ATmega32U4 microcontroller that works with the Arduino bootloader. The USB connection is for programming and debugging. A 110-mAh lithium-polymer (LiPo) battery powers the system. The more expensive kit includes additional devices.

The system is essentially a small board designed to be sewn into a garment and connected to small displays and sensors like the LilyPad Light Sensor, LilyPad Buzzer, LilyPad Button Board, LilyPad colored LEDs, and a LilyPad Slide Switch. These devices are connected to the microcontroller using uninsulated wire that can be easily threaded through fabric. The voltage and current is minute, so the potential issues tend to be things like unwanted shorts if wires are sewn too close together.

The wearable technology was developed for Sparkfun by Dr. Leah Buechley, who now runs the design firm, Rural/Digital. Dr. Buechley was also associate professor at the MIT Media Lab. There she founded and directed the High-Low Tech group. An online activity guide is designed to get developers started with wearable projects. ed

4. Sparkfun's LilyPad ProtoSnap Plus Kit is intended to be used for the creation of wearable technology.