electronic design

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

September 2016 electronicdesign.com

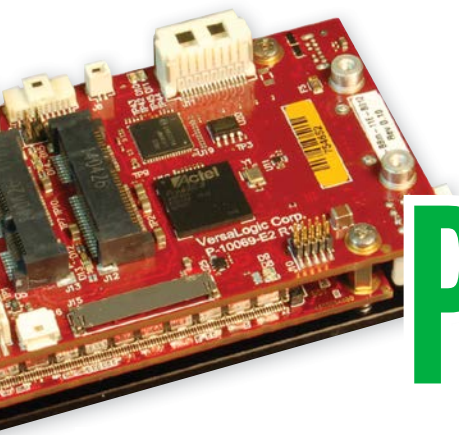

PLENTY OF PLATFORM CHOICES for Embedded Designers

p|20

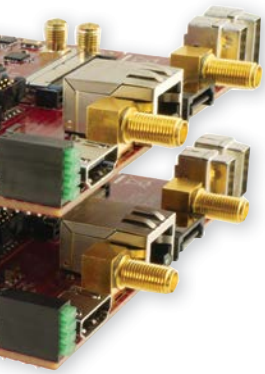

$10.00 Powered by Penton

When You Need It!

Best in Class: Consistent On-Time Delivery?*

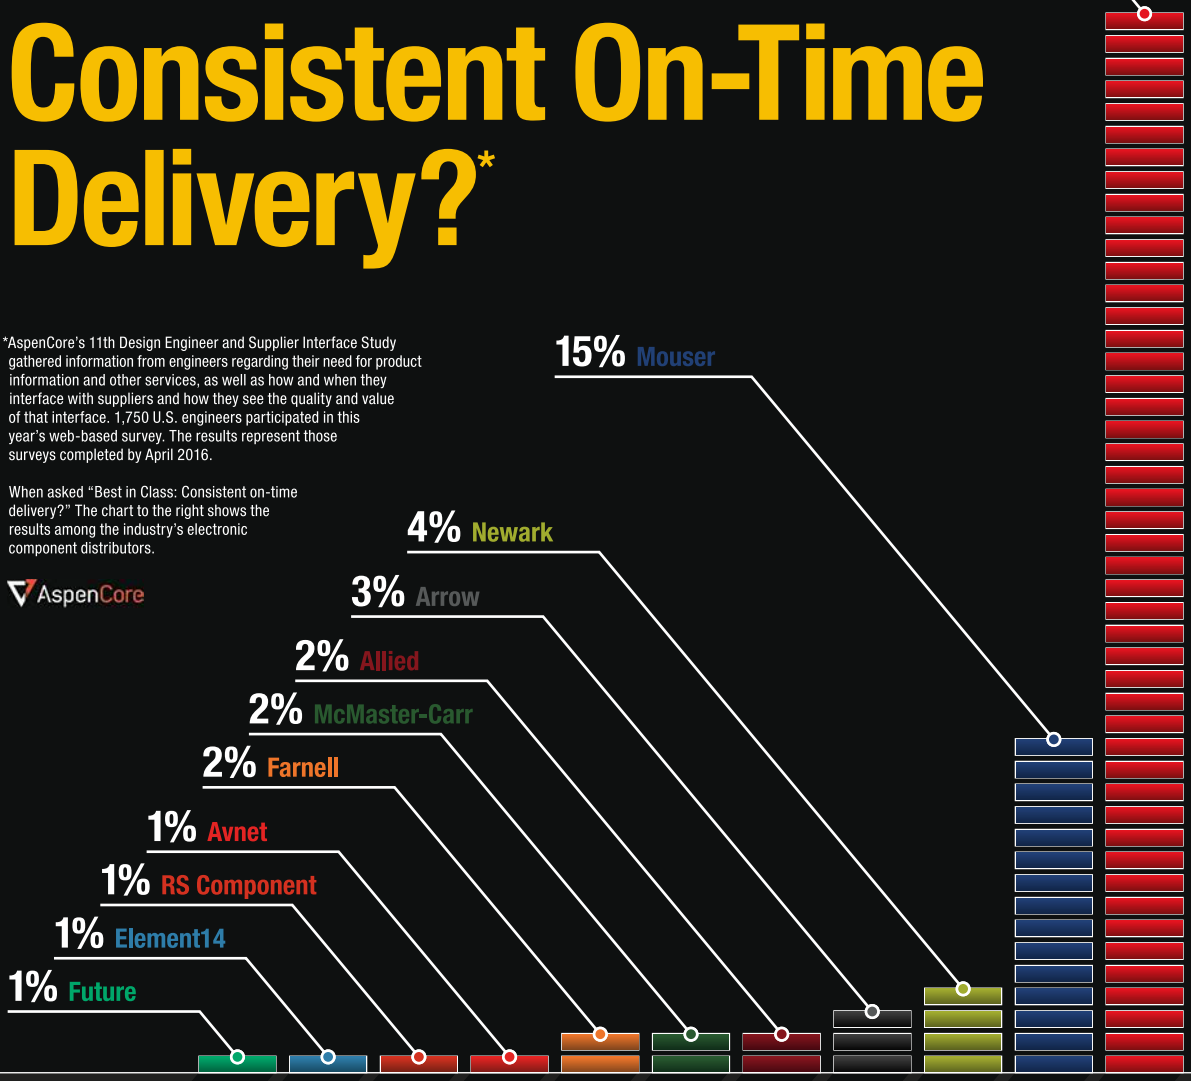

*AspenCore's 11th Design Engineer and Supplier Interface Study gathered information from engineers regarding their need for product information and other services, as well as how and when they interface with suppliers and how they see the quality and value of that interface. 1,750 U.S. engineers participated in this year's web-based survey. The results represent those surveys completed by April 2016.

When asked "Best in Class: Consistent on-time delivery?" The chart to the right shows the results among the industry's electronic component distributors.

AspenCore

SAME-DAY SHIPPING on orders entered by 8:00 pm CENTRAL
NEXT-DAY DELIVERY

Digi-Key® ELECTRONICS

1.800.344.4539

DIGIKEY.COM

4.8 MILLION PARTS ONLINE | 650+ INDUSTRY-LEADING SUPPLIERS | 100% AUTHORIZED DISTRIBUTOR

ecia MEMBER ecsn member CEDA

The Smart Choice for Small Spaces

Coilcraft is the brand engineers trust most when specifying tiny inductors for wearable technology

Boost the performance and battery life of your wearable products with these tiny RF and power inductors from Coilcraft:

- Wirewound chip inductors as small as 0201 size for wireless communications
- Shielded power inductors as thin as 0.71 mm for power management
- Coupled inductors as small as 2.2 x 1.45 mm for LED display drivers

You can get started by using our suite of web tools to quickly locate the perfect inductors for your design. Compare and analyze multiple parts based on your true operating conditions, and then order free evaluation samples with just a click.

Learn why we're the biggest name in ultra-miniature inductors. Visit us at **www.coilcraft.com.**

Coilcraft®

WWW.COILCRAFT.COM

coilcraftdirect.com
No min order. Next day delivery.

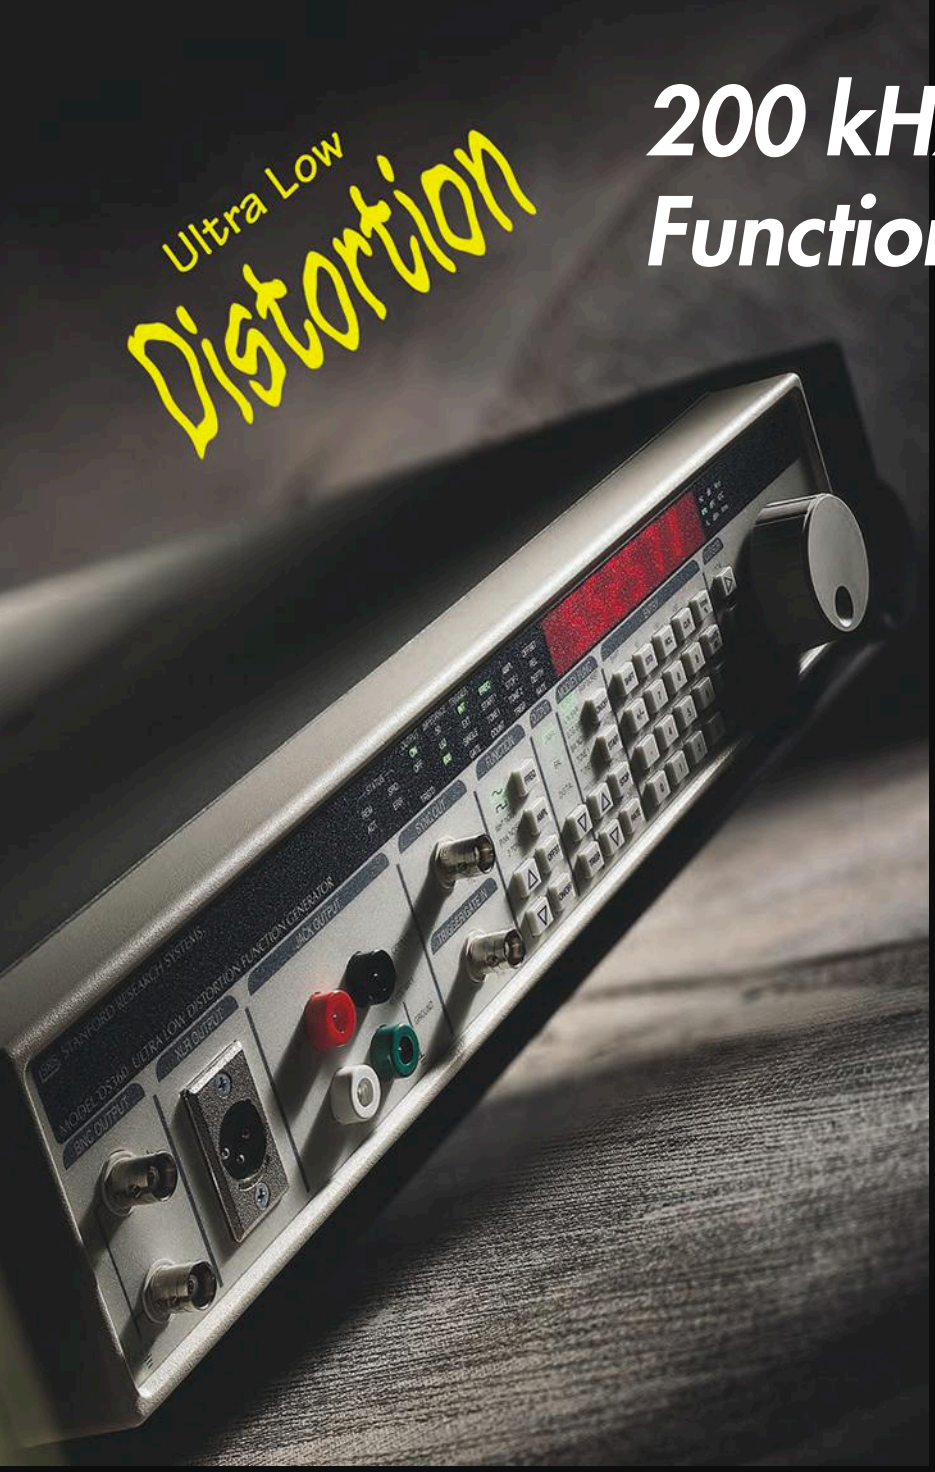

200 kHz DDS Function Generator

0.001 % THD
(–100 dB)

DS360...$2895 (U.S. list)

- 25 ppm freq. accuracy
- 1 mHz to 200 kHz
- 4-digit amplitude resolution
- Sine, square, two-tone, burst
- Pink and white noise
- Log & linear freq. sweeps
- 35 mVpp to 40 Vpp
- GPIB and RS-232 (std.)

The performance of a low-distortion signal source and the features of a digitally synthesized function generator are combined in a unique instrument, the DS360. With 0.001 % THD, 25 ppm frequency accuracy and functions that include sine waves, square waves, sweeps, and bursts, the DS360 delivers unparalleled precision and versatility. To make it complete we added two-tone capability (including SMPTE, DIM, CCIF), bandwidth limited white noise, a 10 MHz reference input, and computer interfaces.

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

In This Issue

FEATURES

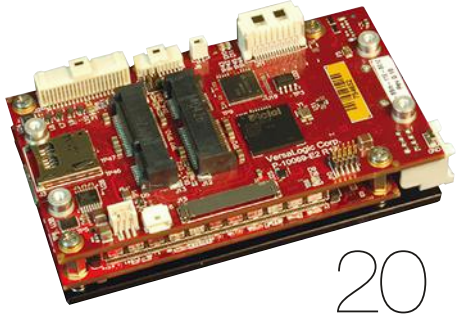

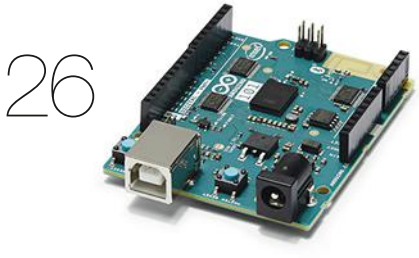

NEWS & ANALYSIS

IDEAS FOR DESIGN

COLUMNS & DEPARTMENTS

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

ELECTRONIC DESIGN (ISSN 0013-4872) is published monthly by Penton Media Inc., 9800 Metcalf Ave., Overland Park, KS 66212-2216. Paid rates for a one-year subscription are as follows: $120 U.S., $180 Canada, $240 International. Periodicals postage paid at Kansas City, MO, and additional mailing offices. Editorial and advertising addresses: ELECTRONIC DESIGN, 1166 Avenue of the Americas, New York, NY 10036. Telephone (212) 204-4200. Printed in U.S.A. Title registered in U.S. Patent Office.

on electronicdesign.com

A JEWEL OF A JOULE

http://electronicdesign.com/iot/jewel-joule

Introduced by Intel CEO Brian Krzanich at the recent Intel Developers Forum, the Intel Joule module looks to be a winner. It is designed for high-volume production and is essentially a PC on a module. Tech Editor Bill Wong takes a closer look.

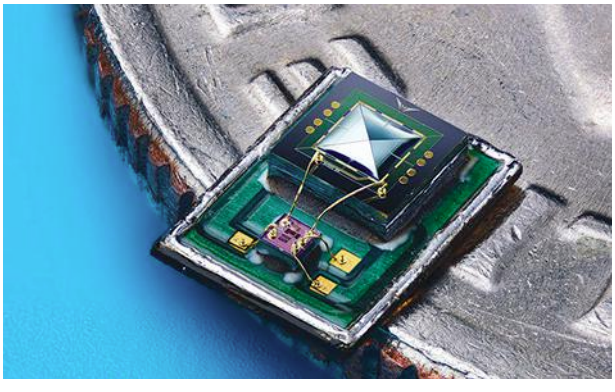

Q&A: VESPER CEO MATT CROWLEY

http://electronicdesign.com/electromechanical/qa-vesper-s-chief-executive-lends-ear-internet-things

Vesper Technologies, an acoustic sensor startup that designs microelectromechanical systems (MEMS), is betting that microphone arrays will become one of the most popular ways to interact with smartphones, household devices, and smart speakers like the Amazon Echo and Google Home.

join us online

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

INSIDE THE DRAGONFLY 2020 3D PRINTER

http://electronicdesign.com/boards/qa-behind-scenes-look-dragonfly-2020-3d-printer

When the founders of Nano Dimension decided they wanted to create the first professional 3D printer for printed electronics applications, the result was the DragonFly 2020 3D Printer that prints PCBs and combines advances in inkjet3D printing technology, advanced nanotechnology-based conductive and dielectric inks, along with sophisticated software to deliver a rapid-prototyping tool.

BEHIND THE HIGH-BAND SPECTRUM FOR 5G NETWORKS

http://electronicdesign.com/blog/us-opens-high-band-spectrum-5g-networks

The Federal Communications Commission (FCC) made history last month by unanimously voting to pass the proposal of the Spectrum Frontier Proceeding. With this move, it allocated unprecedented amounts of bandwidth for 5G wireless communications. Tech Editor Maria Guerra goes behind the proposal, outlining its history and what it means for the future.


AR
...OK GUYS...
WE GOT THIS!
MPD DEVICES CAN BE FOUND IN NUMEROUS
GODZILLA DETERRENTS AND MILITARY APPLICATIONS
BATTERYHOLDERS.COM
MPD

electronic design
Powered by Penton
SEPTEMBER 2016

WE'RE THE MANUFACTURER YOUR COMPETITION WANTS KEPT SECRET.

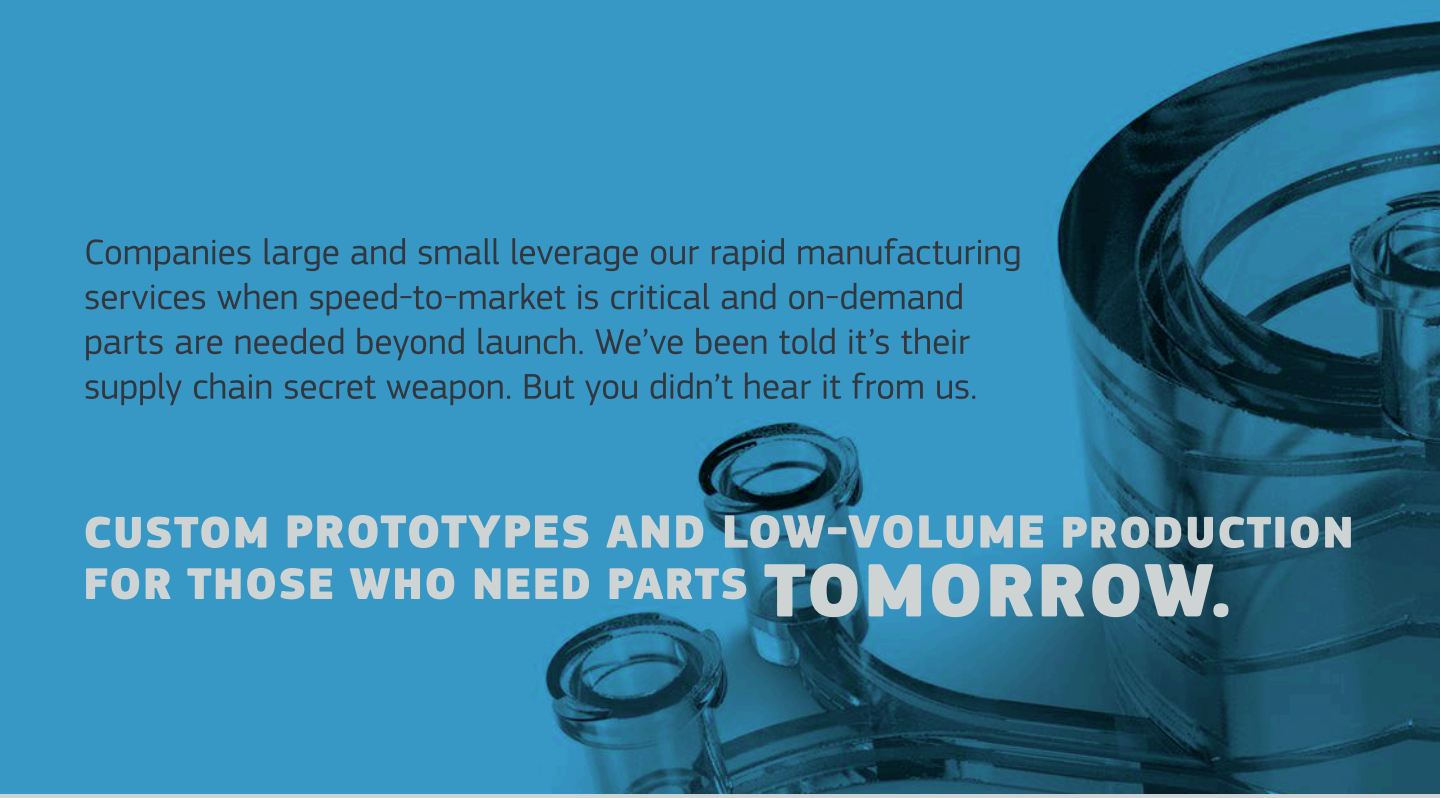

Companies large and small leverage our rapid manufacturing services when speed-to-market is critical and on-demand parts are needed beyond launch. We've been told it's their supply chain secret weapon. But you didn't hear it from us.

CUSTOM PROTOTYPES AND LOW-VOLUME PRODUCTION FOR THOSE WHO NEED PARTS TOMORROW.

proto labs®

Real Parts. Really Fast.™

3D PRINTING | CNC MACHINING | INJECTION MOLDING

ISO 9001: 2008 Certified | ITAR Registered | 2016 Proto Labs, Inc.

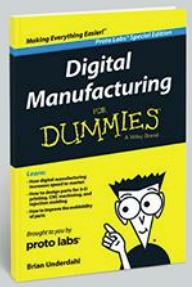

FREE BOOK

Request your Digital Manufacturing for Dummies book at **go.protolabs.com/ED6DJ.**

Looking for just the right part,
at just the right price...

Connect with over 200 distributors
on the NEW SourceESB!

Part Lists Tool

- Send multi-part RFQs
- Save your part lists to work on later
- Filter by authorized distributor

www.SourceESB.com

MARIA GUERRA | Technical Editor
maria.guerra@penton.com

London Post-Grad Team Wins Power Challenge

A post-graduate team from Imperial College London has won the inaugural GaN Systems Geoff Haynes Future Power Challenge. The winners were announced at the EPSRC Centre for Power Electronics Annual Conference 2016, which took place in Nottingham, England.

Intended to be an annual event, the GaN Systems Geoff Haynes Future Power Challenge was jointly established by GaN Systems and the EPSRC Centre. Haynes is the recently retired co-founder of GaN Systems.

The Challenge was open to all UK power electronics post-graduate students—specifically, those who submitted research papers or posters that contributed to accelerating the use of GaN transistors in future power conversion or control applications. Entries covered subjects as diverse as the design of a novel compact motor with embedded filter windings and optimized p.w.m. control strategy for an A-NPC converter.

The winning team comes from Imperial College London's Department of Electrical and Electronic Engineering. It includes Ph.D. students George Kkelis and Juan Manuel Arteaga, along with research assistants Sam Aldhaher and David Yates, and is supervised by Dr. Paul Mitcheson (*below*). The team developed two inverter prototypes, each based on a Class-EF topology using GaN Systems' GS66504B switches.

This new design maintains zero-voltage switching and delivers a constant output AC current, regardless of the load resistance value. The design allows a Class-E or Class-EF inverter to operate efficiently for any load. This was shown to significantly relax the requirement for accurate alignment of transmit and receive coils in a wireless power application.

The GaN Systems Geoff Haynes Future Power Challenge winning team, including supervisor Dr. Paul Mitcheson (second from left), with Geoff Haynes (second from right). ***(Courtesy of GaN Systems)***

"Our team had three posters competing in the GaN challenge," says Arteaga. "The judges decided to award the three posters given the fact that we had been working together in most of their experiments."

Usually posters for conferences might include pictures of the experiment, the collected data, simulations, results, analysis, and conclusions. For this particular completion, the format of the poster was completely free and there were no live demonstrations. Arteaga mentioned the posters awarded as:

1. Class-E2 Resonant Converter with Inherent Output Voltage Regulation Operating at 6.78 MHz.
2. A 13.56 MHz Load Independent Inductive Power Transfer System with 73% DC-DC Efficiency.
3. Comparison between GaN and Si Devices in 6.78 MHz and 13.56 MHz DC/AC Inverters.

"The three posters were presented by George Kkelis and me in the Centre for Power Electronics Summer School," Arteaga continues. "The posters presented the latest work developed by us, Dr. Paul Mitcheson (our supervisor), Dr. David Yates, and Dr. Samer Aldhaher, who developed the concept of the Load-Independent Class-EF Inverter—shown in his latest publication at the 2016 Wireless Power Transfer Conference."

Prof. Mark Johnson of the University of Nottingham and Prof. Barrie Mecrow of Newcastle University judged the competition at the annual summer school event, organized by the Ph.D. students of several universities that form the EPSRC power electronics center. The judges saw all the posters competing for the challenge and talked to the postgraduate students who were presenting during the poster presentation sessions.

Presenting the award, GaN Systems' president and co-founder, Girvan Patterson, underlined the importance of supporting the research initiatives between industry and academia to accelerate the adoption of disruptive technologies, as well as to inspire a new generation of engineers. ed

News

TOSHIBA AND WESTERN DIGITAL REVIEW THICKEST 3D NAND, But No Word on Cost Yet

Less than two years ago, memory chipmaker Toshiba demolished one of its major NAND manufacturing plants in Yokkaichi, Japan, and built an entirely new facility on the site. Now the facility, filled with new semiconductor equipment, is producing the latest generation of NAND built in three dimensions.

Toshiba and manufacturing partner Western Digital have begun turning out samples of the thickest version of 3D NAND to enter production. Memory chips from both companies feature 64 layers of memory cells, instead of the 48 layers that defined the previous generation of 3D NAND.

The new thickness was Western Digital and Toshiba's big announcement at the recent Flash Memory Summit, an industry conference for non-volatile memory. One of the event's themes was the industry's ongoing shift toward 3D NAND production, which has been slowed by competitive prices for planar NAND.

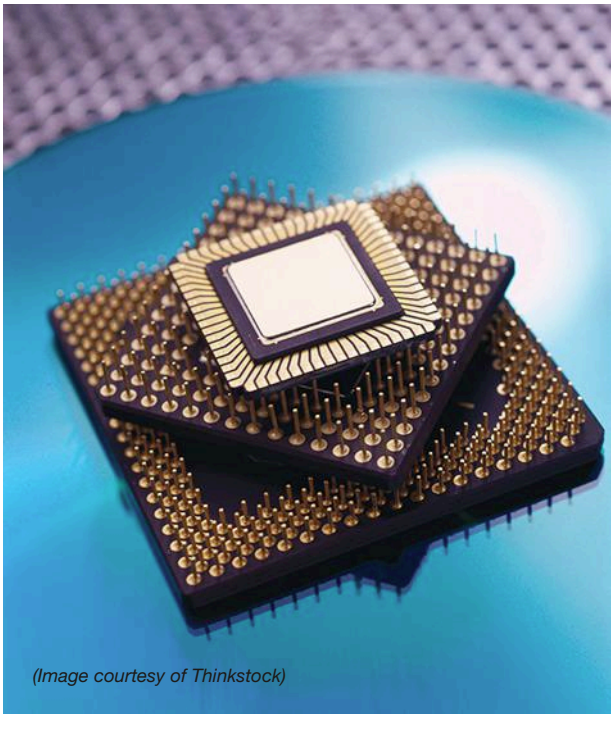

(Image courtesy of Thinkstock)

The announcement, however, wrinkled under questions about the memory's cost per bit—one of the measuring sticks for new generations of memory chips. For years, industry analysts predicted that NAND with 64 layers would match the price of planar NAND and catapult the technology into the mainstream.

But Western Digital and Toshiba have both stayed quiet on pricing and declined to share details with *Electronic Design*.

To construct 3D NAND, chipmakers pile layers of memory cells vertically on a silicon substrate, so that they act like a single device. There are several ways to connect the layers: Intel and Micron link them with the same floating gate transistors used in planar NAND, while Samsung and Toshiba slide insulators in between the layers to control the flow of electrons.

With either approach, the result is memory chips that provide significantly more capacity than planar NAND. That makes them ideal for solid-state drives inside data centers, smartphones, and personal computers, as well as applications for the Internet of Things.

Western Digital's memory is called BiCS3, since each memory cell stores up to three bits. Its first incarnation has 256 gigabit—or 32 gigabyte—capacity, but will eventually hold up to half a terabyte on a single chip, the company said. This is combined with "advances in high-aspect ratio semiconductor processing to deliver higher capacity, superior performance, and reliability," said Siva Sivaram, Western Digital's executive vice president of memory technology, in a statement.

BiCS3 was an outgrowth of an earlier collaboration between Toshiba and SanDisk, which helped build the New Fab 2 manufacturing plant in Yokkaichi. Western Digital

continued the relationship after it acquired SanDisk last year in a $19 billion deal and officially opened the new facility with Toshiba in July.

Western Digital's bet on SanDisk was a significant departure from its main business of selling hard disk drives, which are only slightly cheaper than planar NAND with the same capacity. The company fell into a memory industry increasingly strained by volatile prices, which shift with the tides of supply and demand. In recent years, the low cost of planar NAND has stifled the market for 3D NAND, causing big chipmakers like Micron to cut jobs while the market turns around.

Despite the fact that NAND prices continue to fall, it has become increasingly difficult to make smaller and more advanced planar NAND devices. Western Digital and Toshiba, along with competitors like Micron and Samsung, are betting the higher capacity of 3D NAND offers a route around that obstacle.

A spokesperson for Western Digital declined to comment on the price of BiCS3, which will enter mass production toward the end of 2016. She said that the company's long-term plan was "to maintain the cost scaling that SanDisk pioneered in the 2D NAND era," adding that its new technology will be "very cost-competitive."

Also being built in the new Yokkaichi manufacturing plant is the latest version of Toshiba's BiCS flash memory, which has been upgraded to 64 layers. Like Western Digital's version, the new architecture can hold three bits per memory cell and achieve a 256-gigabit capacity.

The latest version of BiCS—which stands for "bits cost scalable"—was built using Toshiba's new stacking process, which provides around 40% higher capacity than previous 48-layer processes. Toshiba said that the new stacking process "reduces the cost per bit, and increases the manufacturability of memory capacity per one silicon wafer."

The new BiCS flash is already sampling, with mass production out of the Yokkaichi facility scheduled for early 2017. At the Flash Memory Summit, Toshiba announced its first products built on the new memory: One example is the BG series of M.2 solid-state drives, which measure only 30mm long and have 128, 256, and 512 gigabit capacities.

Toshiba also declined to share price information with *Electronic Design*, and did not reveal the price for its new BG solid state drives.

"Planar NAND has hit the wall in terms of significant scalability, which is where 3D BiCS enables an industry cost-down path for higher-density NAND," the company said in an e-mail statement, but did not provide any timeline for when that might happen. ■

WHY DID LINEAR TECHNOLOGY SELL ITSELF to Analog Devices?

LINEAR TECHNOLOGY WAS not looking for buyers when Analog Devices offered $14.8 billion to acquire the company and its library of power management circuit designs. At the end of July, however, the companies announced that the deal had been signed, approved by both boards, and sent to regulators for review.

That the company accepted the deal was surprising to an industry familiar with its engineering culture, disdain for corporate matchmaking, and fanaticism for analog design. On social media and online forums, engineers reacted with criticism of the deal's inevitable culture clash and questions about the fate of products, customer support, and the popular LTSpice circuit simulation software.

An array of sensors inside a jet engine. Data converters, like those built by Analog Devices, turn physical signals from sensors, like movement and pressure, into digital data. Linear Technology's power management chips help to conserve the electricity required for that process. *(Image courtesy of Matt Gibson, Flickr)*

So why did Linear Technology agree to the $14.8 billion deal?

The short answer is that it's a windfall for Linear Technology shareholders. The deal represents almost 10 times the chipmaker's revenues in 2015–one of the highest multiples paid in the semiconductor industry in the last decade. Shareholders will be paid $46 and around one-fourth of an Analog Devices share for each stock.

According to Robert Swanson, who helped establish Linear Technology in 1981 and serves as the executive chairman, the price was too high to brush off. Taking the deal far outweighed the idea of keeping the business and famously high profit margins, which it makes from selling analog chips to automakers and factories.

"I think what pushed us over there was, frankly, we are a public company. They have fiduciary responsibilities, and the offer was so compelling," Swanson said in a conference call with financial analysts on Tuesday. "It just turned out to be a deal we couldn't turn down."

But in an industry where intellectual properties have been thrown around like trading cards, the terms were a little murkier than other recent chip deals. In the conference call, the two sides sketched out the long negotiations about Linear's culture that preceded the deal and how it might be preserved under new management. While the final decision might have been motivated by shareholder responsibilities, the fate of Linear's culture clearly weighed on the deal.

For years, Linear Technology has been a breeding ground not only for high-performance circuits but also for new generations of analog engineers. It has built a reputation as an engineering-driven company, carving out time every week to hear engineers pitch new designs and expecting them to interact directly with customers. It has such high retention that Linear Technology has called itself "the company that no one leaves."

"We talked extensively about what he's paying for, and he understands that clearly," Swanson said about his conversations with Vincent Roche, Analog Devices' chief executive. "Our people are our most valuable assets. I am incredibly proud of Linear's team, and I am gratified that Vince is well aware of how lucky he is to welcome them into the analog family."

Neither Swanson nor Robert Dobkin, Linear Technology's co-founder and chief technology officer, were interviewed for this article. Roche was also not interviewed.

Merging the two chipmakers is not without its technical merits. Analog Devices and Linear Technology are two of the biggest makers of circuits for translating things like sound, video, and radio signals into digital language. Analog is the industry's biggest supplier of data converters, while Linear has focused on power management. These products could be combined with wireless chips to create systems that gather data from vehicle sensors, factory equipment, and city infrastructure.

The cultural crossover was less clear. But executives on both sides said that the engineering-driven culture at both companies would smooth out the merger. "I think we have cultures that are very similar relative to innovation," said Swanson. "I think that's the reason this is a combination that can work better than any that I've seen in the past."

Linear Technology was founded 16 years after Analog Devices was established in 1965. The startup quickly turned into a bastion for analog design, in an industry shifting toward digital circuits with the rise of personal computers. Its first employees were defectors from other Silicon Valley companies—including National Semiconductor, where Swanson served as vice president for analog design—that had been leaving analog behind.

Swanson's exit created a vacuum that sucked in other analog gurus. These included Dobkin, George Erdi, and Bob Widlar, a former Fairchild Semiconductor engineer who later became known as the "god of the operational amplifier." The initial migration turned Linear Technology into a destination for analog engineers and eventually into something like an apprentice program for students of the analog craft.

Linear Technology has leaned on the expertise passed through its ranks. For most of his career with the company, Swanson rejected the idea of buying new technology through big acquisitions.

Instead, he favored developing new circuits in-house. Linear has only made one acquisition in its history—Dust Networks, a wireless sensor networking company, in 2011.

Linear Technology has around 4,865 employees worldwide, a little more than half the 9,000 workers at Analog Devices. What will happen to the combined engineering team remains unclear. Both companies have overlapping products in power management, amplifiers, interface chips, and data converters. Analog Devices said that it will operate four chip manufacturing plants after the deal.

"There isn't a great abundance of talent in this entire industry, so it's our intention to make sure that we win the hearts and the minds of everybody around this great future that we've got," said Roche.

Both companies are relatively small for the analog industry. Other chipmakers like Texas Instruments, Infineon, and STMicroelectronics have tens of thousands of employees and more diverse product lines. In contrast to its new acquisition, Analog Devices has dipped into the playbook of big corporations, buying new analog capabilities. Analog's most recent deal came in 2014 when it paid $2.45 billion for wireless chipmaker Hittite Microwave.

Both companies have been shaped by different industries, said Jonathan Liao, a semiconductor analyst at IHS Markit. Analog's products are mostly general-purpose devices, which are used in markets ranging from smartphones to cellular infrastructure, and designed on shorter cycles. Linear Technology has focused on higher-performance circuits in markets like industrial control and automobiles, where design cycles are longer and more flexible.

Roche said that any cultural rifts would be healthy for growth. "We can tolerate cultural differences," he said. "In fact, I believe that, in many ways, culture enriches by getting the best of both of these very high-performing companies. It's going to enrich the combination. But that's what we will pay very careful attention to."

Linear Technology will not exist as an independent division under Analog Devices, but the brand will be kept alive to sell power-management products. In addition, Analog Devices said that Linear's engineers will remain a strong voice in guiding new power designs.

That could give Linear Technology's culture a second life. And in the eyes of its founders, that culture is at the heart of Linear Technology's legacy.

"We have a unique culture here. We have people who like what they're doing and they'd like to continue to do it," Swanson said in a 2011 interview with *Electronic Design* magazine, in which he talked about Linear's 30-year anniversary. "And I think that our unique culture can go for another 30 years."■

SCIENTISTS MAKE LITHIUM-AIR BATTERIES That Hold Their Breath

WHEN TESLA MOTORS revealed its plan to buy solar panel maker SolarCity, the thread tying them together was battery storage. The idea, according to chief executive Elon Musk, is to combine solar panels with Tesla's stationary Powerwall batteries. On cloudy days or in the evening, the stored energy could run your dishwasher or charge your Tesla Model S sedan.

For the batteries inside its cars and Powerwall systems, Tesla is betting on the same lithium-ion chemistry used in smartphones and hover-boards. Other companies and researchers think there is a better option. Many are experimenting with new lithium-air batteries that are smaller, lighter, and more energy-efficient than their forebears.

In August, scientists from the Massachusetts Institute of Technology, the Argonne National Laboratory, and Peking University in China revealed a promising new version of lithium-air batteries. Their new device, the scientists said, can serve as a drop-in replacement for lithium-ion, while storing over five times more energy as today's batteries.

According to the scientists, the new design uses solid oxygen electrodes to overcome many of lithium-air's drawbacks. This new battery loses much less energy in the form of heat than earlier versions. The result is that it lasts longer and is more energy-efficient, making it a better option for electric cars and renewable energy storage.

"This means faster charging for cars, as heat removal from the battery pack is less of a safety concern, as well as energy efficiency benefits," said Ju Li, an MIT professor of nuclear science and engineering and author of the research.

The new device, which was described in the journal *Nature Energy*, is the latest vying to become the replacement for lithium-ion. Scientists from IBM and the startup PolyPlus are placing bets on lithium-air, while other companies are working toward zinc-air and sodium-air batteries.

The new design "has added a further promising high energy storage battery system into the mix," wrote Lawrence Hardwick, a chemist at the University of Liverpool, in a *Nature* article about the research. But he admitted that "as yet, no obvious front-runner has emerged."

Holding back lithium-air batteries has been their chemistry. Until now, lithium-air batteries inhaled outside air—driving a chemical reaction with the battery's lithium—while electric current flows out. This oxygen is released to the atmosphere during the charging cycle. The chemical reaction produces other molecules, known as lithium peroxide, that slowly clog the battery electrodes.

That raises several problems. According to Li, the solid particles formed by the reaction cause the battery to degrade faster than lithium-ion batteries, which are completely sealed from outside air. When the battery degrades, it stores less energy.

The battery is also prone to losing energy in the form of heat. Its output is more than 1.2 volts lower than the voltage needed to charge it, causing it to lose 30% of the electricity as heat. Because of that, the battery can "actually burn if you charge it too fast," said Li. Overcharging can lead to structural damage or an explosive reaction known as thermal runaway.

(Image courtesy of Thinkstock)

The new design solves these problems by closing off the battery to outside oxygen. The same electrochemical reactions take place between lithium and oxygen during charging and discharging, but they take place without ever using oxygen gas. Instead, the oxygen stays inside the battery and switches between three solid chemical compounds: Li_2O, Li_2O_2, and LiO_2. This prevents the damaging particles from forming.

The new battery was paved by an earlier study in lithium-air chemistry. Earlier this year, two of the engineers involved in the new report—Khalil Amine and Jun Lu from the Argonne National Laboratory—helped create the lithium superoxide LiO_2, while the battery was discharging. The compound can more easily split into its parts than lithium peroxide, making it more efficient and longer-lasting.

"This discovery really opens a pathway for the potential development of a new kind of battery," Larry Curtiss, an Argonne National Laboratory chemist, said about the experiment in January. "Although a lot more research is needed, the cycle life of the battery is what we were looking for."

The new battery sharply cuts voltage loss, so only 8% of the electrical energy is lost as heat. It also inherently guards against overcharging: The device can shift between different lithium compounds if it is being overcharged to stop activity that might cause damage. The scientists overcharged the battery to 100 times its capacity for 15 days without any damage. They also found that through 120 charging cycles, the battery only lost 2% of its capacity.

The lithium-air battery can also do without components to pump air inside and out of the battery. Without these auxiliary parts, it can easily be adapted to existing devices or battery packs inside cars and power grid storage.

For now, the researchers are keeping to the laboratory. The team expects to make a practical prototype of the device within about a year. ■

Accuracy first.

Keysight True*form* outperforms DDS.

Keysight True*form* waveform generators offer the industry's lowest jitter and superior signal integrity. BenchVue software enables you to easily create custom waveforms without programming. And thanks to True*form* technology, your custom signals are accurately represented with low distortion and no skipped points.

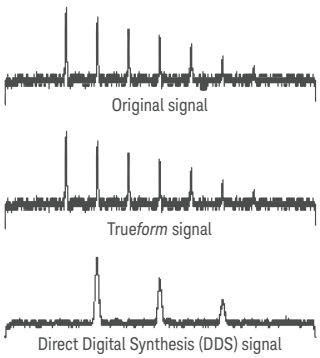

Keysight True*form* Waveform Generators

	33600A Series, 80 and 120 MHz	33500B Series, 20 and 30 MHz
Jitter	<1 ps	<40 ps
Pulse edge time	2.9 ns independent rise & fall	8.4 ns independent rise & fall
Arbitrary waveforms	64 MSa memory w/ sequencing	16 MSa memory w/sequencing
PRBS patterns	200 Mbits/s, N from 3 to 32	50 Mbits/s, N from 7 to 28
2-channel coupling	*Equal, differential, ratio and combined*	

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

Buy from an Authorized Distributor:

microlease Newark element14 TESTEQUITY

KEYSIGHT TECHNOLOGIES

Unlocking Measurement Insights

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

MARIA GUERRA | Technical Editor

Can Supercapacitors Surpass Batteries for Energy Storage?

Advances in supercapacitors are delivering better-than-ever energy-storage options. In some cases, they can compete against more-popular batteries in a range of markets.

A supercapacitor is a double-layer capacitor that has very high capacitance but low voltage limits. Supercapacitors store more energy than electrolytic capacitors and they are rated in farads (F). Supercapacitors store electrical energy at an electrode–electrolyte interface. They consist of two metal plates, which only are coated with a porous material known as activated carbon. As a result, they have a bigger area for storing much more charge.

The plates are immersed in an electrolyte made of positive and negative ions dissolved in a solvent. When voltage is applied, two separate charged layers are produced on the surface with a small separation distance. This is why supercapacitors are referred to as electric double-layer capacitors or EDLCs (*Fig. 1*).

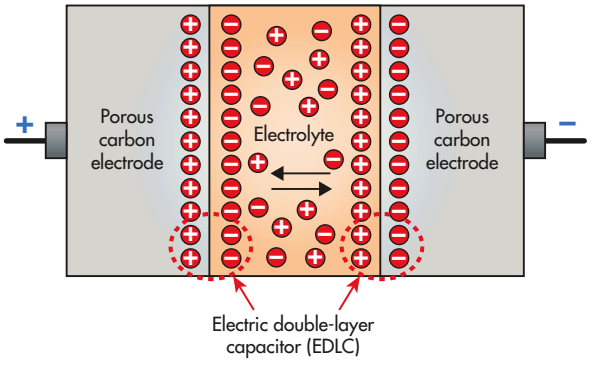

1. The separation of charge distance in a double layer is on the order of a few angstroms (0.3-0.8 nm). (Courtesy of InTech)

Supercapacitors have many advantages. For instance, they maintain a long cycle lifetime—they can be cycled hundreds of thousands times with minimal change in performance. A supercapacitor's lifetime spans 10 to 20 years, and the capacity might reduce from 100% to 80% after 10 or so years. Thanks to their low equivalent series resistance (ESR), supercapacitors provide high power density and high load currents to achieve almost instant charge in seconds. Temperature performance is also strong, delivering energy in temperatures as low as –40°C.

On the other hand, supercapacitors are offset by their low energy density. Thus, they can't be used as a continuous power source. One cell has a typical voltage of 2.7 V; if higher voltage is needed, the cells must be connected in series.

Supercapacitors are used in many power-management applications requiring many rapid charge/discharge cycles for short-term power needs. Some of these applications include:

- Voltage stabilization in start/stop systems
- Electronic door locks in the event of power failures
- Regenerative braking systems
- Distribution microgrid
- Medical devices
- Energy harvesting
- Consumer electronics
- Kitchen appliances
- Real-time clock backup
- Utility meters
- Backup power
- Grid stability
- Wind energy
- Energy efficiency and frequency regulation
- Remote power for sensors, LEDs, switches
- Memory-backup SRAM
- Burst-mode power delivery
- Forklift hybridization

2. Zap&Go allows you to charge up to three devices simultaneously, two using the twin USB 2.0 ports and the third using the Qi wireless charging disc. *(Courtesy of IEEE Spectrum)*

SUPERCAPACITOR MATERIALS

Different materials, such as various carbon materials, mixed-metal oxides, and conducting polymers, have been used for supercapacitor electrodes. Advances in carbon-based materials, namely graphene, increase the energy density to nearly the level of batteries.

Graphene is a thin layer of pure carbon in the form of a two-dimensional, atomic-scale, honeycomb lattice. Graphene is very strong and thin, with high capacitance value due to its high internal surface area. Researchers at companies developing commercial solutions for graphene supercapacitors are targeting much more efficient and eco-friendly energy-storage solutions at lower price points. Some technologies have already arrived, some are about to launch, while others are still in the making. Among the latest developments:

Skeleton Technology's graphene supercapacitors will be key players in a new transport fleet trial in the UK, turning rigid diesel trucks into hybrids through power from regenerative braking for the first time. The hybrid truck system developed by **Adgero** and Skeleton Technologies, named UltraBoost, will be installed and go through trials on UK delivery routes in coming weeks. During braking, the unit becomes a generator, recovering kinetic energy that would otherwise be lost as heat. At the heart of this technology is a bank of five high-power graphene-based supercapacitors known as **SkelMod**.

Zap&Go, a UK-based startup, is launching a new type of charger specifically for the business traveler. It uses graphene supercapacitors to charge phones in five minutes. The company had to develop a new type of power supply, because getting enough power into the charger in a few minutes wouldn't work with existing power supplies. The Zap&Go five-minute charger will be available in the market soon and can be preordered (*Fig. 2*).

Other companies are developing supercapacitors with interesting applications, too.

Eaton, a power-management company, offers supercapacitor solutions such as coin cells, large cells, small cylindrical cells, and modules. For example, its **XLR 48V Supercapacitor**

SUPERCAPACITOR STANDARDS AND STANDARDIZED TEST

WHILE TESTING CAN assure customers whether supercapacitors are reliable, manufacturers are not required to comply with a specific standard or test. A few standards intended for supercapacitors let manufacturers voluntarily test their products using different methods.

Some of the more commonly used standards and tests for supercapacitor include:

- **Restriction of Hazardous Substances (RoHS) Compliance:** This European directive specifically ensures that no hazardous materials restricted by Europe are involved during the product manufacturing process. It restricts the use of six hazardous materials (lead, mercury, cadmium, hexavalent chromium, polybrominated biphenyls, and polybrominated diphenyl ether) in the manufacturing process of electronic and electrical equipment.

- **IEC 60068-2-27 and 2-29:** These shock standards assure that a product is durable in rugged, high-demand applications, and that it can handle sudden applied forces and shocks.

- **REACH Compliance:** REACH stands for the Regulation, Evaluation, Authorization and Restriction of Chemicals. The sole purpose of REACH is to address the production and use of chemical substances, and their potential impacts on both human health and the environment.

- **UL 810A:** The standard was developed by UL especially for supercapacitor technology. The requirements of this test cover supercapacitors being used in equipment such as electronic products, uninterruptible power supplies, emergency lighting, engine starting, and power equipment.

- **IEC 62576:** This standard describes methods for testing basic electrical characteristics (capacitance, voltage maintenance, and energy efficiency) of conventional electric double-layer capacitors (EDLCs) to be used for peak power assistance in hybrid electric vehicles.

- **IEC 62391-2:** This is for fixed EDLCs for use in electronic equipment and EDLCs for power application.

As the presence of supercapacitors increases in the power electronic market, the need for a single standard keeps rising, too. It will be easier for customers to have just one standard in mind when deciding which supercapacitor to buy. Right now, for example, there are several methods for ESR testing (e.g., IEC6231, EUCAR, manufacturers' own methods).

Module (*Fig. 3*) provides energy storage for high-power, frequent-charge/discharge systems in hybrid or electric vehicles, public transportation, material handling, heavy equipment, and marine systems. It doesn't require monitoring, since the module has built-in cell balancing.

Maxwell Technologies' supercapacitors are used for regenerative-braking energy storage in the Beijing subway system. China Railway Rolling Stock Corp. (CRRC-SRI) leverages Maxwell's 48-V modules in two sets of regenerative-braking energy-storage devices for the system's No. 8 line, an urban rail network that runs north-south through China's capital.

3. The XLR modules consist of 18 individual Eaton XL60 supercapacitor cells designed to provide 48.6 V and 166 F with 5 mΩ dc resistance for incorporation into systems requiring up to 750 V. *(Courtesy of Eaton)*

Vishay offers its 220 **EDLC ENYCAP** with a rated voltage of 2.7 V. It can be used in several applications, including power backup, burst power support, storage devices for energy harvesting, micro UPS power sources, and energy recovery.

4. When the input power is lost, the bidirectional switching controller LTC3350 acts as a step-up converter to provide power from the supercapacitors to VOUT, which becomes the backup power for the application. *(Courtesy of Linear Technology)*

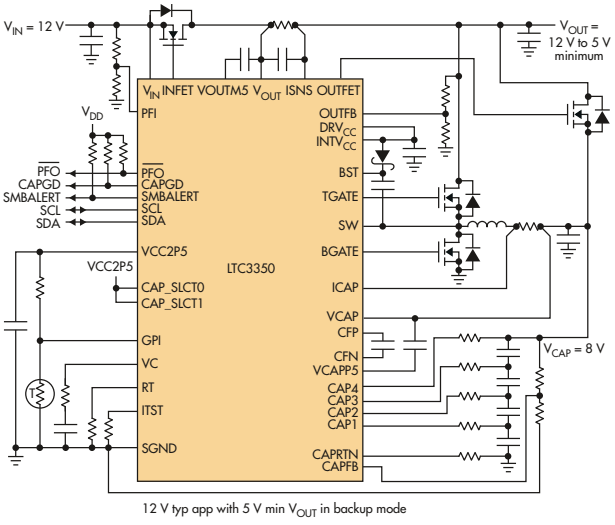

Though a single supercapacitor cell will provide 2.7 V, higher voltages can be achieved by connecting several supercapacitors in series. Just as with lithium-ion batteries, superca-

pacitors in a stack might not have the same capacitance due to manufacturing or uneven aging. Strings of more than three capacitors require voltage balancing to ensure long operational life, preventing overvoltages by keeping the voltage on each cell as low as possible to achieve the needed total stack voltage.

POWER-MANAGEMENT SUPERCAPACITORS

Engineers can choose from various devices designed specifically to manage the unique requirements of supercapacitor charge, depending on the application.

Linear Technology, for example, offers its **LTC3350** (*Fig. 4*), a backup power controller that can charge and monitor a series stack of up to four supercapacitors. The LTC3350, designed for automotive and other transportation applications, offers the following features:

- Power backup by charging a bank of up to four supercapacitors in case of a power failure. It can handle 4.5- to 35-V input voltage and over 10 A of charge/backup current.
- Balancing and overvoltage protection to the series stack of supercapacitors.
- Monitoring of system voltages, currents, and die temperature.
- Internal capacitor voltage balancers that eliminate the need for balance resistors.

Intersil's ISL78268 is a synchronous buck controller designed to support an input voltage range of 5 to 55 V. The buck controller integrates a peak current mode control for PWM generation and a constant average current loop ideal for supercapacitor charging. With an ISL78268, it's possible to charge a bank of capacitors using a constant-current/constant-voltage (CICV) methodology. In addition to supercapacitor charging, the ISL78268 is also suited for automotive power and telecom power supplies.

Maxim's MAX13256 H-bridge transformer driver is another solution for charging supercapacitors while simultaneously driving a system load. The addition of a small diode and resistor in parallel between the supercapacitor's positive lead and the output of the transformer circuit limits the current used to charge the supercapacitor during normal operation. It also allows for easy discharge when energy is required.

Advances made in super-capacitor capacity and energy density will ultimately lead to greater functionality and more overall presence of the devices throughout the energy industry. Based on all of their inherent advantages, supercapacitors should help reduce the costs to the customer by minimizing the amount of batteries needed, as well as the frequency of battery replacement. ed

WILLIAM WONG | Embedded/Systems/Software Editor

Plenty of Platform Choices for Embedded Designers

Developers have a wide range of choices when it comes to industrial single-board computers. We examine many of the options available.

Embedded developers have never had more standard board or module options than are available now. Standard form factors like PC/104 continue to grow in popularity, packing in more interfaces and computational support while vendor-specific form factors provide developers with even more options.

Expansion options have likewise expanded into a plethora of alternatives from Mini PCIe *(see "PCI Express Mini Card Tackles Compact Embedded Expansion" on electronicdesign.com)* to stackable interfaces like PC/104 OneBank and SUMIT. Perhaps the biggest change has come in non-volatile, on-board memory. Flash memory options include soldered eMMC modules to removable MicroSD slots. Rugged Mini PCIe sockets often support mSATA devices. They also have a USB and I^2C interface that are often utilized.

Compact devices like Versalogic's Osprey *(Fig. 1)* Embedded Processing Unit (EPU) are ideal for mobile applications like UAVs where performance and low power are critical requirements. Vendor-specific platforms like this allow more compact designs.

The Osprey EPU is only 55- by 95- by 27-mm but it packs a wallop, including Intel's quad-core Bay Trail processor with up to 4 Gbytes of DDR3L memory and up to 8 Gbytes of eMMC flash memory. Storage options include a MicroSD slot and a SATA port. It has high-speed connectivity including dual Gigabit Ethernet ports and USB 3.0 connections. Expansion is possible via a full-size Mini PCIe socket with mSATA support and a half-size Mini PCIe socket. Video outputs include DisplayPort++ and LVDS connections. The conduction-cooled system requires less than 7 W of power.

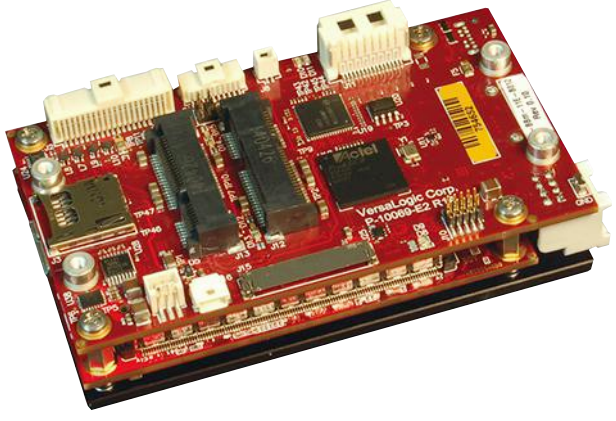

1. Versalogic's compact Osprey Embedded Processing Unit (EPU) runs Intel's quad-core Bay Trail processor and features expansion options that include a MicroSD slot, SATA port, Mini PCIe with mSATA support, and a half-length Mini PCIe socket.

Standard ARM-based parts have provided vendors with low-power alternatives to the x86 world. WinSystems' SBC35-C398DL *(Fig. 2)* runs NXP's 800 MHz i.MX 6DL SoC, which has a pair of ARM Cortex-A9 cores. It also supports ARM's TrustZone security features. The board is a popular 3.5-in. form factor with expansion options that include support for a half-size Mini PCIe card and an IO60 expansion connector. It can be powered by a Power-over-Ethernet (PoE) connection via the Gigabit Ethernet, which also has IEEE-1588 support or by a 10 to 50 VDC source.

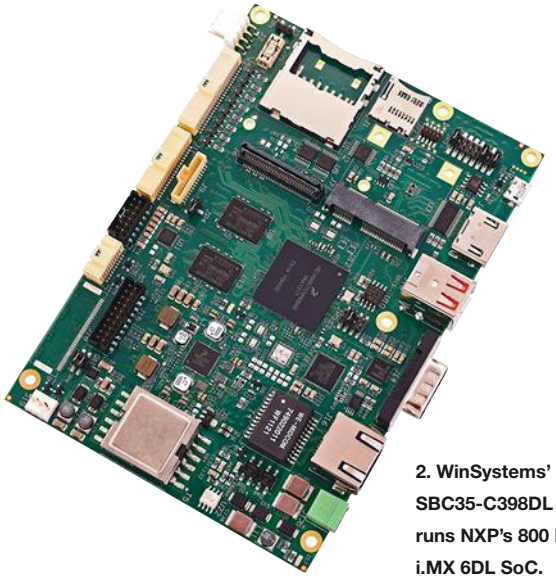

2. WinSystems' 3.5-in. SBC35-C398DL board runs NXP's 800 MHz i.MX 6DL SoC.

The SBC35-C398DL holds up to 2 Gbytes of DDR3 DRAM and a MicroSD slot. It also has a CMOS camera input, six USB 2.0 ports, 24 GPIO lines that are 30 V tolerant, HDMI and LVDS video outputs, as well as audio support. The IO60 connector exposes I2C, SPI, GPIO, and PWM signals.

Vendors have been able to pack more onto a board as system-on-chip (SoC) designs pack more into a single chip. Likewise, the use of low-power, compact FPGAs provides flexibility that vendors can take advantage of without having to expose users to the complexity of FPGA programming. The challenge these days is often how to connect these interfaces to the outside world. Some designs, like Diamond Systems' Helix *(Fig. 3)*, allocate almost half the board to connectors.

The Helix is a stackable PC/104 design with a 1 GHz, x86-based DMP Vortex86DX3 SoC with up to 2 Gbytes of memory. It has six USB 2.0 ports, 2 RS-232/422/485 and 2 RS-232 ports, a 10/100Mbps Ethernet port, a 1 Gigabit Ethernet port, and a SATA port that can be used with an on-board SATA DOM or off-board SATA device. A Mini PCIe port support mSATA. CANbus 2.0 is a custom option.

There is a 16-channel, 16-bit ADC with a 100 KHz sample rate plus four 16-bit DAC channels with voltage outputs. There are 11 programmable digital I/O lines. These peripherals are supported by Diamond's Universal Driver data acquisition programming library plus an interactive, easy-to-use graphical control panel that runs on Windows or Linux to provide to control of all data acquisition features.

3. Diamond Systems' Helix allocates almost half the board to connectors, including the PC/104 ISA and a Mini PCIe socket.

The PC/104 ISA bus is not fast compared to PCI Express (PCIe), but it is much easier to design an expansion board for ISA. Many applications can make due with ISA's speed and memory limitations typically reducing the cost of custom or standard expansion boards.

Compatible form factors like EPIC and EBX use a larger processor board, allowing more peripherals and connectors to be incorporated in the base system while providing stacked expansion.

PCIE STACKABLES BLEND EXPANSION AND COMPACTNESS

PC/104-Plus added a PCI connector to the mix with its higher performance and functionality. Having two connectors on either side of the board provides a rigid stack, but it does take up board space.

A PCI-only, PCI-104 variant is a standard configuration freeing up board space *(Fig. 4)*, but even PCI has been replaced by PCI Express in most processors and SoCs. Often a PCIe-to-PCI bridge is needed. That type of bridge is straightforward than a PCIe-to-ISA bridge because PCIe was designed based on PCI.

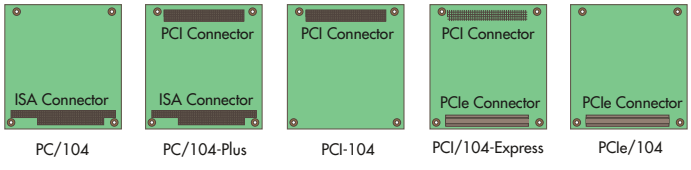

4. PC/104 has progressed through PCI Express versions including PCI/104-Express and PCIe/104 that also support OneBank boards.

The PCIe versions use a staggered pass-through versus the straight pass-through for the parallel bus PCI and ISA connectors because PCIe is a point-to-point connection. A PCIe expansion board will have multiple lanes coming in. It will route a lane for its use and shift the unused lanes so the next unused lane lines up with the one removed but on the connector on the other side of the board. Normally a single lane is needed for an expansion board. There are x1, x4, x8, and x16 versions.

There are a number of PCIe stackable standards. PCIe/104 includes a number of variations including "OneBank," which refers to the use of one connector bank rather than the three for the

full x16 PCIe lane implementation. OneBank can support PCI Express Gen 3. OneBank expansion boards can also work with PCI/104-Express processor boards. OneBank includes four x1 PCI Express lanes, two USB 2.0 connections, and an SMB/I^{2}C interface.

The board on top of the OneBank stack *(Fig. 5)* from Sundance is the EMC2-KU35. The board supports a range of Xilinx FPGA modules including the Zynq SoC, Artix and Kintex FPGAs. The board also has an FMC connector to support VITA 57.1 boards. It uses four PCI Express lanes and a 12.5 Gbit/s serial link for the FMC connector.

The PC/104 ISA bus was based on early x86 processors. The move to PCI Express platforms opens up this form factor to other processors, primarily ARM-based platforms. ConnectTech's XTreme/SBC PCIe/104 system board is a carrier board for a COM module. It currently has three ARM-based COM modules that have an NXP i.MX51, Texas Instruments OMAP, and an NVidia Tegra in addition to two Intel Atom modules.

The ConnectTech modules contain their own DRAM and flash memory. They have different peripheral complements but typically these include LVDS video, USB, Ethernet serial ports. The Atom and Tegra modules support PCI Express.

MODULES MELD CUSTOMIZATION WITH FAST TURNAROUND

Modules are the flip side to stackables. Modules can be used on standard boards like the ConnectTech Xtreme/SBC *(Fig. 6)*, but they are more often mated with custom carrier boards. In general, the carrier board design complexity is significantly less than the module, which is often a multilayer board whereas the carrier is often a two-layer board with connectors instead of dense SoCs.

Developers can often turn to outside design houses to create carrier boards that often incorporate additional peripherals along with the connectors needed to link the system to the outside world. One interesting alternative is Gumstix's Geppetto *(see "Best of 2015: Create Custom Capes Fast and Easy" on electronicdesign.com)*.

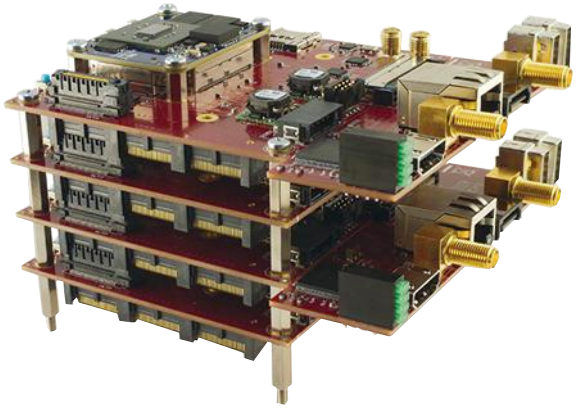

5. This stack of Sundance boards utilizes the OneBank connector. The top board hosts an FPGA module.

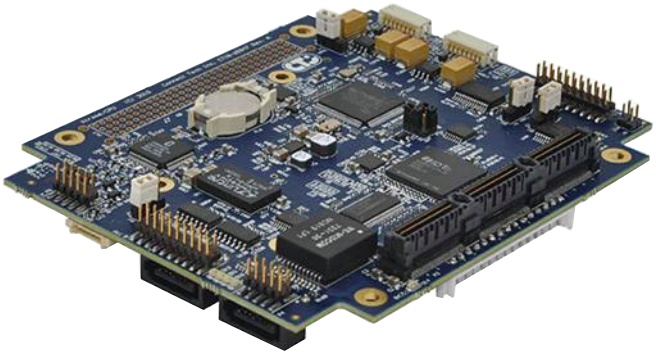

6. ConnectTech's XTreme/SBC PCIe/104 system board is a carrier board for a COM module.

7. This carrier board was designed using the web-based Geppetto from Gumstix.

Geppetto is a web-based design tool specifically for creating carrier boards *(Fig. 7)* including Gumstix' own line of modules. There is a flat setup fee of $1,999 and then a per-board charge. The result is a fully populated board. Geppetto also supports a range of third-party modules as well, such as the Colibri iMX6 SOM from Toradex. The Colibri iMX6 is available in dual- and single-core versions. They have an ARM Cortex-A9 core plus 512 Mbytes of RAM and 4 Gbytes of eMMC flash memory.

Modules can range from very tiny boards like the Revolution Robotics Warp_0x01 module *(see "What Is Inside an IoT Chip?" on electronicdesign.com)* that targets wearable applications to high-performance modules like NVidia's Jetson TX1 *(see "Module Delivers Supercomputer Performance" on electronicdesign.com)*. ed

The power to _____.

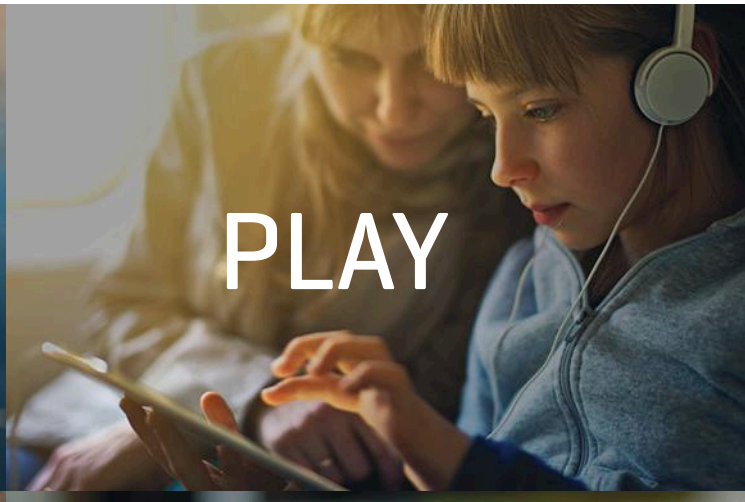

The Power of Things™

You are reinventing the way individuals play and interact with one another. You are designing the products that will preserve natural resources and biodiversity. You are behind the technology that will improve personal, home and corporate security.

While you are working on the applications that will shape the future, we are here to support you with our proven power products and collaborative design approach.

electronica

Power Hall A2, Booth 613 *www.cui.com/electronica*

CUI INC®

BILL WONG | Technical Editor

A Deeper Look at Marvell's MoChi Technology

Dr. John Shin discusses how Marvell's LEGO-inspired Modular CHIp (MoChi) technology will change the SoC industry.

This past year, Marvell Technology introduced its MOdular CHIp (MoChi) technology as the answer to the costly, complex system-on-chip (SoC) implementation. I talked with Dr. John Shin, vice president at Marvell, about how the company perceives its LEGO-inspired technology will impact the SoC industry.

What was the impetus behind the development of Marvell's Modular Chip technology?

Shin: Modern SoCs are very complicated—it often takes more than two iterations to develop an SoC for production. After decades of technology scaling, Marvell realized that the SoC mask set cost was reaching the multi-million dollar mark. At that point, we recognized that even if a company could afford the astronomically high point of entry, the cost of development would eventually take a toll on the bottom line. Also, the ROI of these SoCs in advanced technology nodes (such as FinFET) would be very low, sometimes negative. Additionally, the SoC end market is extremely dynamic with requirements changing all the time.

However, as design and process complexity increases drastically, design and debug cycles and time to production get longer, and it is enormously challenging to react. All of these factors led us to explore how to build SoCs differently.

Can you expand on how the market has changed and how it relates to MoChi?

Shin: Vertically integrated powerhouses, like Apple and Samsung, jumped into the semiconductor market for internal consumption, significantly lowering volume opportunity for the rest of the industry. You may be wondering, what's the big deal? Every industry has its own competition. But the problem with other companies doing SoC business in an already shrunken market is that the industry still competes to make more complicated SoCs just because it's possible, believing that the more we can integrate, the better the chip.

In the smartphone industry, for example, we integrate not just the CPU and GPU functionality, but also the 3G/4G modem, video/image/audio processor, display controller, GPS, and so on. Meanwhile, the smartphone market is constantly changing and, even if we develop the SoC with huge resources ahead of time in anticipation of the market, there is no certainty that the results will come to fruition.

For these reasons, we had to shift to a new way of implementing equivalent SoCs, which is more flexible, quicker to respond, and less costly. MoChi technology and architecture is poised to revolutionize the way next-generation SoCs are designed, built, and brought to market.

How does MoChi technology address these issues?

Shin: In thinking about how to solve this challenge of developing modern SoCs, we needed to accept that the optimal integration is not necessarily building one super SoC with every function required. Rather, it's a handful of simpler and modular chips (MoChis), each optimized and fabricated in the most cost-effective technology.

The Marvell MoChi solution is a radical departure from

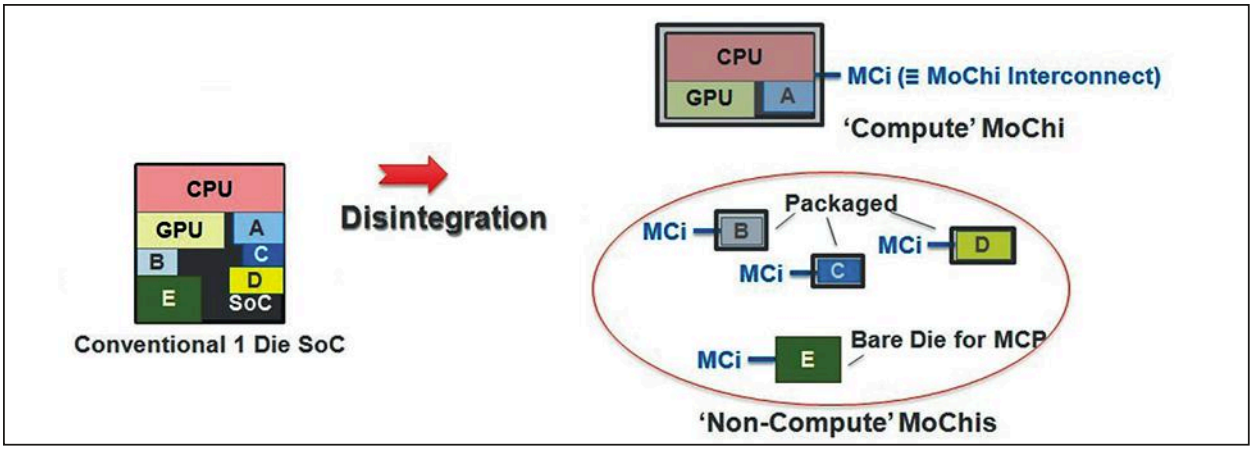

MoChis are split into "compute" and "non-compute" categories. Compute MoChis usually require high performance, thus are fabricated in most advanced technology nodes, while non-compute MoChis are typically standard fixed-speed interface functions built with older process generations. (Source: Marvell)

the current convention of integrating as much as possible onto one IC, helping solve the toll that the cost and complexity of building billions of transistors on a single device is taking on the ecosystem. Much like LEGO blocks that we all likely played with as kids, the MoChi approach focuses on splitting the SoC into cost-optimized modular chips of specific functions that can easily and seamlessly be connected with each other without compromising overall system performance.

MoChis are split into "compute" and "non-compute" categories. Compute MoChis usually require high performance, thus are fabricated in most advanced technology nodes, and non-compute MoChis are typically standard fixed-speed interface functions built with older process generations (*see figure*). With a high-bandwidth, low-latency, point-to-point MoChi Interconnect common to all of these MoChis, this approach provides the ability to connect and combine different MoChis that "talk" to each other as if they are integrated on a single piece of silicon. This leads to the creation of virtual SoCs (VSoCs) in countless configurations with great flexibility. The MoChi Interconnect is essentially a packetized and serialized extension of the widely used on-chip bus, and its area and power overhead is minimal.

Marvell has also introduced Marvell Final-Level Cache (FLC) technology? How does this fit in?

Shin: On top of the challenge that we identified in designing SoCs, we also looked at existing computer systems that have been slowly evolving. The problem is that even if SoCs were free, the cost of building a computer system is still quite high because of the DRAM cost. That said, if you look at the task manager of your PC, you'll see that most processes are idle, but take up a lot of expensive DRAM space. This gave us a clue on how to solve the problem.

Marvell's FLC innovation introduces a new way to build computer systems that redefines the main memory hierarchy by substantially reducing the amount of expensive DRAM. By offloading the storage onto less-expensive SSD memory and using a smaller amount of DRAM to cache the active processes, the FLC architecture provides unprecedented reductions in system cost, form factor, operating power, and battery size, or a significant increase in standby battery life. FLC can enable all computers to use a fraction of the energy they use today. This opens up a host of new applications and products, and provides a tremendous cost savings for higher-end devices.

How will MoChi impact the future of the SoC industry?

Shin: At an industry-level, the Marvell MoChi architecture will enable significant cost savings and flexibility as well as faster development cycles for the entire silicon market from end to end. Minor die overhead of MoChi Interconnect will be easily compensated by the reduction of upfront engineering cost. This comes as a result of the drastic reduction in the number of super SoC combinations that otherwise must be built to serve the various markets and corresponding customer requirements.

To date, Marvell has launched multiple MoChi processors ("compute") incorporating the FLC technology and MoChi South Bridge products ("non-compute") across a variety of applications, including networking, mobile computing, multimedia, printer, and storage, among others. The variety of VSoCs using these MoChis demonstrates the key advantages of MoChi, such as flexibly and quickly configuring VSoCs with optimal functions.

The days of the single-die SoC are numbered. The MoChi architecture provides a significant paradigm shift that delivers the return-on-investment required to advance the semiconductor industry.

BILL WONG | Technical Editor

What's Inside an IoT Chip?

What does an IoT SoC need? Everything from embedded security and communications to neural networks and more.

The Internet of Things (IoT) can exist with conventional microcontrollers and SoCs, but issues such as low power requirements and wireless support have pushed development of platforms designed for IoT applications. Hardware security support is now a requirement and forthcoming ARMv8-M microcontrollers will be the new benchmark for security (*see "New ARM v8 Architecture for Microcontrollers" on electronicdesign.com*). Specialized coprocessor support is also common from sensor integration and neural networks.

Many of the coprocessors found on IoT-oriented platforms are designed to do more, often targeted, processing using less power than using software alone on a conventional platform. One example is Intel's Curie based around a 32-bit Quark SoC. It incorporates a 32-bit DSP for sensor fusion support as well as a 128-neuron pattern matching accelerator. The neural network is supported by General Vision's CurieNeurons library. The module has Bluetooth Low Energy (BLE) support as well as a six-axis accelerometer and gyroscope. The Intel Curie is the compute engine for the Arduino 101 board (*Fig. 1*). The neural net support is also available in a standalone chip from General Vision (*see "Neural Net Chip Enables Recognition for Micros" on electronicdesign.com*).

1. The Arduino 101 board runs Intel's Curie chip with a 128-node neural network supported by General Vision's CurieNeurons library.

IoT platforms typically incorporate wireless support of one sort or another as well as sensor fusion support normally done by a separate, low-power processor that runs independent of the primary processor. As noted, hardware security support is becoming more important and integrated, but there are alternatives.

For example, Microchip's 3-mm by 2-mm AWS-ECC508 chip is preconfigured to provide mutual authentication with Amazon Web Services IoT (AWS IoT). There are a number of I2C-based security chips available, but this one is already registered for development work and it can be customized to handle the product production process. The chip is linked to a host processor, which runs software from the AWS Software Development Kit (SDK). The chip's keys are already set up to work with AWS.

Sometimes a single-chip solution is not always possible or preferred, even in a mobile IoT environment. Developers can create something from scratch and it often makes sense if production quantities runs into the millions, but modules are often a good alternative for prototypes in lower quantities. Revolution Robotics' Warp_0x01 module (*Fig. 2*) is an open-source hardware platform that fits the bill (see *"Tiny Modules Target IoT Applications" on electronicdesign.com*). It is based on NXP's low-power i.MX6 SoloLite with an ARM Cortex-A9 core. The module also contains 512 Mbytes of LPDDR2 memory and a 4 Gbyte eMMC flash module. It also has 802.11b/g/n WiFi and Bluetooth Low Energy 4.0 wireless support. The i.MX6 has access to a six-axis accelerometer and magnetometer sensor.

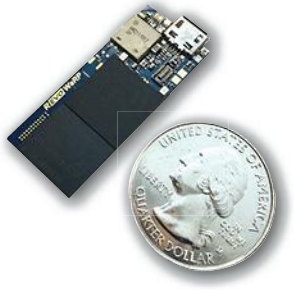

2. Revolution Robotics' Warp_0x01 module is based on NXP's i.MX6 SoloLite adding memory, 802.11b/g/n WiFi, Bluetooth Low Energy 4.0 and a six-axis sensor.

IoT chips still need a lot of software to link them to nearby IoT devices or the cloud. Still, having the right low-power hardware makes a developer's job much easier. ed

Tadiran batteries make your devices run for their lives...

PROVEN **40** YEAR **OPERATING** LIFE*

And keep them running until 2056.

The battery of the future is here today. Tadiran bobbin-type lithium thionyl chloride ($LiSOCL_2$) batteries feature an annual self-discharge rate of just 0.7% per year: so energy efficient that they allow low power consuming wireless devices to operate for up to 40 years on a single battery. No one else even comes close. Tadiran lithium batteries also feature the highest capacity, highest energy density, and widest temperature range of any lithium cell, plus a glass-to-metal hermetic seal for added ruggedness and reliability in extreme environments.

For a battery that lasts as long as your device, run with Tadiran.

TADIRAN BATTERIES

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

www.tadiranbat.com

* Tadiran LiSOCL2 batteries feature the lowest annual self-discharge rate of any competitive battery, less than 1% per year, enabling these batteries to operate over 40 years depending on device operating usage. However, this is not an expressed or implied warranty, as each application differs in terms of annual energy consumption and/or operating environment.

LOU FRENZEL | Contributing Editor

Wireless Is Headed... WHERE?

From the Internet of Things to 5G and LTE, we look at current wireless trends and what the future holds.

Wireless is everywhere. In fact, there are probably very few times during the day when you are not using wireless. Think about it—smartphones, Wi-Fi, Bluetooth connectivity, GPS, AM/FM/satellite radio, garage door openers, automatic cruise control radar and more. And the future holds even more wireless for us all. Here is an overview of the current state of wireless and what you can expect in the future.

THE NASCENT INTERNET OF THINGS

The Internet of Things (IoT) and its predecessor machine-to-machine (M2M) communications are probably the most discussed wireless applications today. While the IoT/M2M has been massively hyped, companies and other organizations are already making progress toward the predicted 50 billion connected devices by 2020. It is estimated that there are already over a billion devices adding to the already thick cloud of electromagnetic radiation that envelopes us all. Good examples are remotely controlled thermostats and home video monitoring with your smartphone. Sensor networks are also beginning to show up in manufacturing plants and factories, agriculture, smart cities, and the smart grid. Drones are IoT devices. Smart watches are IoT.

Multiple wireless technologies are being used to make all these connections. These include Wi-Fi, Bluetooth Low Energy, 802.15.4, and derivatives like ZigBee, Z-Wave, LoRa/LPWAN, and 3G/4G cellular. Even RFID and NFC are finding a place. No one technology seems to dominate. It is doubtful that any one technology will emerge as a dominant player, nor is it probably necessary. However, in some sectors such as industrial or smart cities, some say that standardization is essential. Multiple efforts are being made to identify what that standard may be. Some of the organizations currently involved in IoT standards development include IEEE, ETSI, IETF, ITU-T, ISO/IEC, oneM2M, 3GPP, and ISA. There are other cooperative efforts going on to further muddy the waters.

1a. Sequans Communications' Monarch module is a single-chip LTE CAT M1/NB1 radio designed for IoT and M2M applications. The package is only 6.5 x 8.5 mm.

Part of the problem of defining an IoT standard is the incredibly broad scope of IoT applications. Some of the more prominent sectors include home network monitoring and control, industrial wireless sensor networks, smart grid, smart cities, personal area networks, wearables, automotive, medical, and cellular M2M. It is hard to imagine one global standard that can cover such a wide range. Perhaps there will be one major standard for each major sector.

One key issue with IoT is security. It has many concerned about personal privacy, potentially stolen data, and hacking in general. Encryption, authentication, and password strategies will go a long way to provide the desired security. There are also some instances where you can implement a useful wireless network, but it doesn't need to be connected to the internet. M2M via a cellular network is an example. This can eliminate or at least minimize the security threats.

1b. The SARA-N2 module is u-blox's new NB-IoT module. Its package is 16 x 26 mm.

Another issue is regulation. While no regulations (except the FCC's net neutrality) are yet in place for IoT, most suspect that there are some regulatory organizations at the federal and state levels already structuring regulations that will control IoT and M2M. The fear is that regulation will discourage development and investment as regulations in other fields have proven. And one can only imagine the unintended consequences of regulations structured before systems have been tested and experience gained to identify regulatory needs.

One of the most recent trends in IoT/M2M is the emergence of the cellular and long-range technologies. Specifically, 4G LTE has been modified to provide lower-power, lower-cost options to other IoT variants. Release 8 of the 3GPP LTE standard identified the LTE CAT 1 a 10 Mb/s version in a 20 MHz channel. A CAT 0 version in Release 12 provides 1 Mb/s speed in a 20 Mb/s channel. These versions are still costly overkill for so many low-speed applications. Now even lower-cost, lower-speed versions are available in Release 13 of the standard. These are CAT M1 (CAT MTC) and CAT M2 (CAT NB1). M1 delivers 300-400 kb/s in a 1.4 MHz channel. M2, the so-called NB-IoT, can provide 30 to 50 kb/s in a standard 200 kHz GSM channel. Several vendors are offering modules *(Fig. 1a and b)* for these versions and the wireless carriers are anxious to offer you M2M services using these new technologies.

And that is not the end of it. The forthcoming 5G cellular systems will also serve IoT applications. With phone subscriptions saturated at this point, carriers are looking for new revenue streams. IoT/M2M applications will be implemented on new 5G systems that may come on line by 2020.

The low-power long-range movement in IoT is competitive with the cellular LTE M2M effort. One leading contender is Semtech's LoRa, which can provide connectivity over a range up to 10 miles. Other competing technologies in this sector are France's Sigfox, Wi-Fi's HaLow, and Weightless. Winners are hard to predict at this point, but Sigfox is doing a great job and expanding fast.

Finally, projections for IoT/M2M growth are positive. ABI Research forecasts the installed base of active wireless connected devices will top 47 billion by 2021. ABI also projects that more than 50% of cellular IoT nodes will adopt LTE with most using NB-IoT.

The big question remains: Is all this really necessary? Just because we *can* do this, *should* we? Not everything really needs to be monitored or controlled. Will it really lead to higher productivity and other benefits or will it just result in more technical complexity, overloaded spectrum, massive non-interoperability, and colossal amounts of unused data with no analytics, security nightmares' or other unintended consequences? The electronics industry is betting heavily on IoT, but it is still an emerging technology that has yet to fully demonstrate its hyped promises. There is considerable competition and an unusual number of company partnerships betting on IoT being the next big thing. It is here now but with lots more to come.

DREAMING ABOUT 5G

Fifth-generation (5G) cellular technology is the other major wireless technology in your future. Formally known as the International Mobile Telecommunications standard IMT-2020, this system is currently under development by the Third Generation Partnership Project (3GPP). 3GPP is expected to submit a final draft of the standard to the International Telecommunications Union (ITU) by mid-2019 with a final standard being available by the end of 2020. So it is not here yet despite the hype. However, most of the major cellular carriers are conducting 5G field trials to gain some early experience and hopefully a competitive edge when it is time to deploy.

In summary, here is what 5G is all about:

- Download speeds of 1 b/s to 10Gb/s.
- Greater subscriber capacity and density.
- Greater diversity of device connectivity.
- Small-cell architecture with short-range coverage. More indoor cells.
- Greater bandwidth with most operation in the millimeter-wave bands.
- Network architecture based on software such as network function virtualization (NFV) and software-defined networking (SDN).
- Massive and multi-user MIMO. 4 x 2 and greater.
- Adaptive beamforming.
- Short latency, 1 to 5 ms or less.
- Greater power efficiency for battery-powered devices
- New modulation schemes to improve spectral and power efficiency as well as minimize out-of-band interference.
- Fiber- and millimeter-wave backhaul.

Driving the development of 5G is the need to handle more video streaming and the ultimate growth of ultra-high-definition (UHD) 4K and 8K video formats.

Another driver is the growth of the Internet of Things. With smartphone saturation complete or near, carriers are seeking new revenue streams from IoT and M2M applications. Applications like autonomous vehicles and virtual reality are expected to require the much lower latency of 5G. Both of these are well in the future.

Perhaps the most important latest 5G development is the FCC's recent allocation of nearly 11 GHz of 5G spectrum for the United States. New licensed spectrum includes 27.5-28.35 GHz, 37-38.6 GHz, 38.6-40 GHz, and unlicensed service in 64-71 GHz. Shared-use spectrum is also available from 37-37.6 GHZ. In the future, an additional 17.7 GHz of millimeter wave may be available as the FCC determines the need and resolves potential conflicts. No assignments below 6 GHz were announced, although some carriers may adopt some of their existing low-band spectrum for 5G service.

Along with the FCC's spectrum declaration was the announcement of a $400 million research program led by the National Science Foundation (NSF). Called the Advanced Wireless Research Initiative, its goal is to help the U.S. maintain leadership in mobile technology, meaning accelerated 5G growth.

It will be years before you can sign up for 5G service or buy a 5G smartphone or IoT device. Some 5G chips are just showing up *(Fig. 2)*. Some carriers will try to offer pre-5G-like products and services to beat the competition, but don't expect full operational service before 2022 or beyond. In the meantime, learn to appreciate the excellent and evolving LTE service we have now.

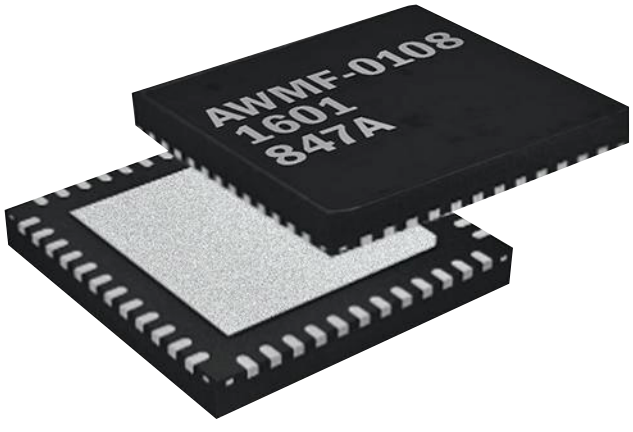

2. Anokiwave's AWMF-0108 is a single-chip millimeter-wave transceiver covering 27.5 to 30 GHz and supporting four Tx/Rx radiating elements and beam steering. The size is only 6 x 6 mm.

SETTLING FOR LTE

Long-Term Evolution is the cellphone technology we all use today. It has been around since about 2008 when the first standard (3GPP Release 8) became available. Then there was the war between the carriers over who could be the first to implement LTE or 4G. Although LTE was really 3.5G, the carriers dubbed it 4G and that designation stuck. Today, according to research firm 4G-Reports, at the beginning of 2016 there were just over 1 billion LTE subscribers mostly in China, Japan, and the United States. There are now 428 LTE networks in 155 countries and those numbers are still growing. Various upgrades continue to be made to expand subscriber capacity and boost data speeds to handle the increased video load that dominates the data capacity of most networks.

The next big upgrade is to LTE-Advanced as defined by 3GPP Release 10 and designated by the ITU as the real 4G. LTE-A added carrier aggregation (CA) and higher-level 8x8 MIMO. CA allows operators to combine up to five 20 MHz channels (contiguous or non-contiguous) into one channel as a way to boost data rate. Along with higher MIMO, the potential maximum data rate is 1 Gb/s. LTE-A is still not widely implemented; ABI Research indicates that of the smartphones shipped in 2015, 23% had LTE-A. ABI predicts that LTE-A with CA will be a part of 61% of smartphones in 2020. Now with Release 13, LTE-Advanced Pro offers additional features such as up to 32 CA, making LTE speeds encroach on 5G territory.

Another trend is the rollout of voice-over-LTE. Today, most networks still support 2G and 3G voice technology. Now, to improve voice quality and lower costs, most carriers are implementing VoLTE. Juniper Research projects that there will be 2 billion VoLTE connections by 2020.

Another development is LTE-Unlicensed and Licensed Assisted Access (LAA). Both LTE-U and LTE-LAA are techniques that cellular operators can adopt to boost data rates and capacity by using the unlicensed 5 GHz band spectrum. Offloading LTE connections to Wi-Fi lessens the carrier's burden and speeds up the data rates. The big issue is interference to Wi-Fi, especially the latest version 802.11ac. LTE-U and LAA are not yet widely implemented in phones or base stations, but progress is being made. LTE-U is ready now, but the LAA standard is still being finalized by the 3GPP. The Wi-Fi Alliance is cooperating to ensure full interoperability and co-existence.

A potential LTE application is to replace and/or enhance the traditional land mobile radio (LMR) technologies such as P25 and TETRA. The need for faster data transfers (e.g., video) as well as the usual voice by public safety organizations has the LMR industry seriously considering LTE. Once needed features are added to LTE such as push-to-talk, proximity detection, and relay capability, LTE will get further consideration. No final decision has been made, but today most major manufacturers of LMR handi-talkies and other gear are already incorporating LTE capability.

Wireless is like most electronic technologies—it changes fast. New methods and products occur almost daily. What is common today will be either obsolete or phased out within five years.

WORLDLY WI-FI

There is always something going on with Wi-Fi. Right now the latest commercial version 802.11ac is rolling out and replacing 802.11n access points (APs). This version operates in the 5 GHz unlicensed band. Using 8-stream MIMO its Wave 2 version can crank out a peak data rate of up to 7 Gb/s. Its successor standard 802.11ax is still under development with up to 8 x 8 multi-user MIMO, a unique OFDM access method, and 1024QAM that can produce a data rate to 10 Gb/s while serving more users concurrently. A successor to the 60 GHz band version of Wi-Fi 802 or 11ad (WiGig) is also under way and designated 80211ay.

Also now available are the 802.11ah and 11af standards. Both target the IoT area. The 11h version, called Halow, is just Wi-Fi in the 902-928 MHz spectrum. The 11f version uses the TV white spaces and cognitive radio techniques. Both offer a range up to a kilometer or more. With so many similar and competing technologies, it is difficult to predict their level of adoption. When commercial chips show up, you could see some IoT and LPWAN usage.

LOOKING AHEAD

Wireless is like most electronic technologies—it changes fast. New methods and products occur almost daily. What is common today will be either obsolete or phased out within five years. One way to keep track or participate in the next generation of wireless is to link up with a new non-profit organization called US Ignite. It was created to foster the next-generation internet applications that can provide some public benefit. The organization recently announced a new advanced wireless industry consortium to support the Platforms for Advanced Wireless Research (PAWR) program. The consortium is a mix of the key government organizations plus a long list of who's who in industry. For more information, go to usignite.org and advancedwireless.org.

iot emerge™

November 2-4, 2016
McCormick Place, Chicago

Learn. Disrupt. Network.

Join us in Chicago this November for real-world IoT insights, strategic and technical sessions, and live demos in four tracks: Smart Cities, Industrial IoT, IoT Engineering, and IoT Security.

Discount Codes:
Free Expo Pass: **IOTEDESEXPO**
15% OFF Conference Pass: **IOTE-DES16**

iotemerge.com

LINEAR TECHNOLOGY

DESIGN NOTES

6A Monolithic Synchronous DC/DC Step-Down Converter with Digital Power System Management

Design Note 554

Haoran Wu

Digital Power System Management (PSM) controllers usually target high current point-of-load (POL) applications. Lower current applications, with up to 6A load current, can also benefit from PSM features. The LTC®3815 is a 6A monolithic synchronous step-down converter with digital power monitoring and control features.

The LTC3815 uses a phase lockable controlled on-time, constant frequency, current mode architecture, that provides extremely fast load step transient response with minimum output capacitance. It also enables operation at the very low on-times required to regulate low output voltages at high switching frequencies for minimum solution size. The output voltage is programmable from 0.4V to 72% of the input voltage with a single external resistor or with an external voltage reference through the reference (REF) input pin. The operating supply voltage range is from 2.25V to 5.5V, making it suitable for operation from 2.5V, 3.3V or 5V input rails or lithium-ion batteries.

6A Converter at 1MHz

Figure 1 shows a typical application with an input voltage range of 2.25V to 5.5V and an output of 1.8V at 6A.

The efficiency curve for the circuit in Figure 1 is shown in Figure 2. Since the LTC3815 can achieve relatively high efficiency even at 1MHz, a small inductor can be used to minimize the total footprint and board size, making this part especially suitable for applications where board space or component height is limited.

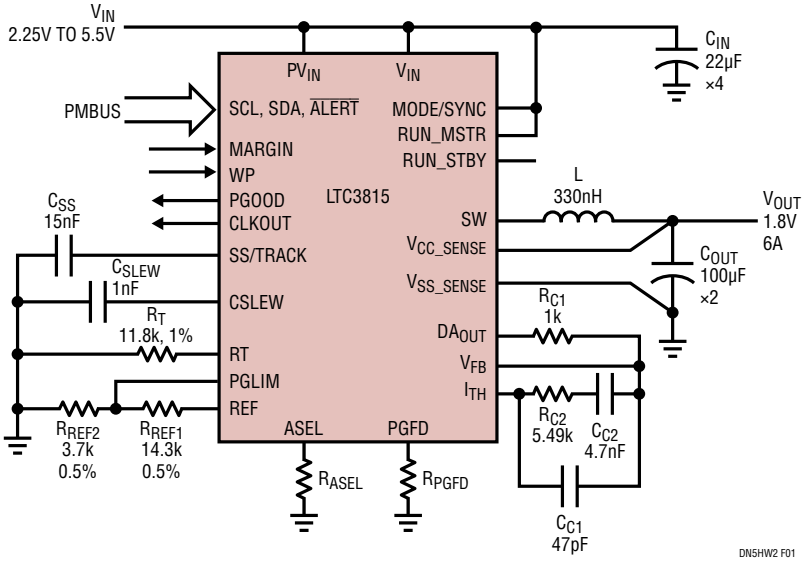

Figure 1. A 1MHz, 1.8V/6A Output High Efficiency Converter Featuring the LTC3815

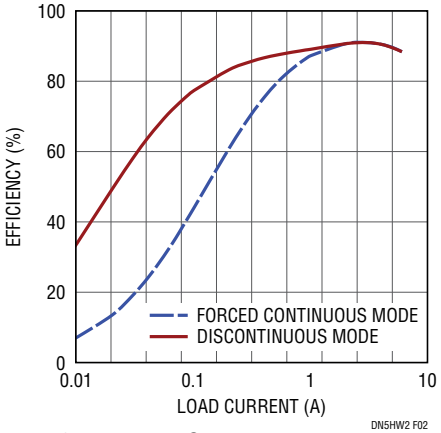

Figure 2. Efficiency Curves at V_{IN} = 5V, V_{OUT} = 1.8V, f_{SW} = 1MHz (for the Circuit in Figure 1)

Figure 3 shows the thermal footprint of the LTC3815 at full load. The IC remains cool at full load even without airflow—the temperature rise at the hot spot is only 37.4°C.

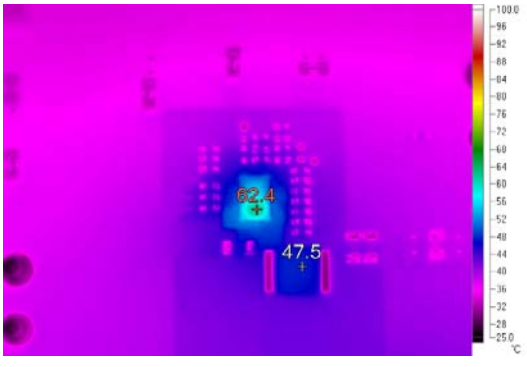

Figure 3. Thermal Picture @ V_{IN} = 5V, V_{OUT} = 1.8V/6A, FSW = 1MHz, No Airflow, TA = 25°C (for the Circuit in Figure 1)

Power System Management Features

The LTC3815 includes a number of PSM functions, including digital control of output margining (±25% with 0.1% resolution) readback of fault status and monitoring of time-averaged (~4ms) and peak input/output current, input/output voltage and temperature. The LTC3815 does not have internal NVM (nonvolatile memory), but the output voltage can be set through PMBus. PSM functions can be accessed using Linear's free LTpowerPlay® software, an easy to use PC-based GUI, shown in Figure 4.

Data Sheet Download

www.linear.com/LT3815

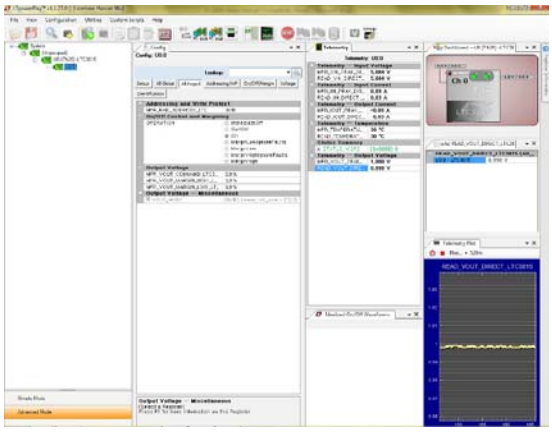

Figure 4. LTpowerPlay Main Interface

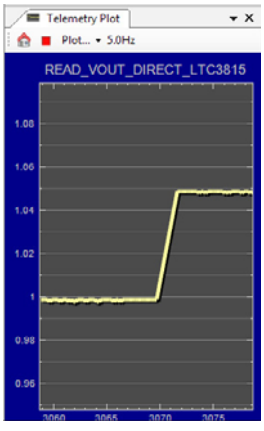

Figure 5. Telemetry Plot after Issuing a MarginHigh Command through LTpowerPlay

For example, to margin V_{OUT} to the margin high value, in LTpowerPlay, simply change OPERATION to MarginHigh, then click the "W (PC to RAM)" icon to write these register values to the LTC3815. When the write is successful, the output voltage jumps to 1.05V on the telemetry plot (Figure 5).

Conclusion

The LTC3815 is a 6A monolithic synchronous buck regulator with digital power system management. It offers a comprehensive solution for applications requiring simple, compact and high efficiency design as well as digital configuration and monitoring functions usually reserved for much higher power POL converters.

For applications help,
call (408) 432-1900, Ext. 3061

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

DN554 LT/AP 0916 71K • PRINTED IN THE USA

PAUL GOUGH | Principal Corporate Strategy, u-blox

Provisioning Reliable Wireless Connectivity for Wearables

Designers face growing demands for power-saving, more ergonomic medical and consumer wearables that maintain high-precision location capability.

As the market for trendy smartphones heads toward maturity, the next consumer "must have" appears to be a wearable device. Often designed for use with a smartphone, using its short and long-range gateway connectivity capabilities to cloud-based applications, the demand for wearables is forecast to intensify rapidly.

The term wearable is a broad classification that comprises smartwatches, Bluetooth headsets, wristbands, chest straps, sports watches, smart garments, and head-mounted displays like those used for gaming. According to a recent report from Gartner,[1] 274.6 million wearable devices will be sold in 2016, an 18.4% increase from 2015. They will generate revenue of $28.7 billion, of which $11.5 billion will be from smartwatches.

Among the 274.6 million devices expected to be sold in 2016, the largest unit sales are forecast as Bluetooth headsets (128.5 million), smartwatches (50.4 million) and wristbands (35 million) *(see table)*. While Apple set the smartwatch benchmark with the launch of its Apple Watch, increasingly strong growth is evident with fitness wearables such as sports watches, fitness bands, and vital-signs monitors used by runners, cyclists, and water-sports enthusiasts. According to Gartner, this particular category is set to maintain its average retail price over the next few years thanks to their special application-specific user interfaces, the need for environmental durability, and the continued advances of sensors and analytics.

Keen to offer a broader set of functionality that competes with those of a smartphone, such as mobile payments, wristband manufacturers are working hard to take market share away from the smartwatch sector. These manufacturers also want to develop the premium paid-for cloud-based services that analyze the data generated by the device.

Another application of the fitness-band and chest-strap sectors involves devices used by wellness programs. Initially driven by health initiatives established in the United States, the positive link between an individual's activity levels and general health continues to gain popularity with health professionals around the world. Many of these programs pay the individual for maintaining a regular exercise regime as opposed to the future high costs of providing health care resulting from lack of exercise.

FORECAST FOR WEARABLE DEVICES WORLDWIDE (MILLIONS OF UNITS)

Device	2015	2016	2017
Smartwatch	30.32	50.40	66.71
Head-mounted display	0.14	1.43	6.31
Body-worn camera	0.05	0.17	1.05
Bluetooth headset	116.32	128.50	139.23
Wristband	30.15	34.97	44.10
Smart garment	0.06	1.01	5.30
Chest strap	12.88	13.02	7.99
Sports watch	21.02	23.98	26.92
Other fitness monitor	21.07	21.11	25.08
Total	232.01	274.59	322.69

Source: Gartner (January 2016)

CONNECTIVITY OPTIONS

One key requirement for any wearable device is connectivity. Popular methods include Wi-Fi, Bluetooth, ZigBee, and cellular, each having their merits. When faced with developing, for example, a new fitness band, the engineer needs to think about how much data will need to be transferred, how frequently, and over what range it would typically need to be sent.

For nearly every application, there will be a trade off between range, data rate, and use case to be considered. Use-case questions such as "Will the fitness band communicate to a smartphone that then collects data and forwards it to the cloud?" and "Will the smartphone application perform local analysis of the data or will that be done in the cloud?" all have an impact.

Wearable devices will always be battery-powered, which ultimately influences the choice of connectivity method. Bluetooth Low Energy (BLE) is designed for low-power requirements, and is ideal for sending relatively small amounts of data. Virtually any smartphone can support this method of communication.

However, if a higher quantity of data, say a few megabytes needs to be transmitted, then the designer might best consider using Bluetooth Classic or Wi-Fi. After that, range needs to be taken into account. BLE can typically communicate over 30 meters in line of sight.

For the fitness-band example, it's assumed that the wearer will also have his or her smartphone with them, so distance is not an issue. However, some wearable applications will dictate the use of cellular communication, since they will need independence and reliance on other communication methods. An example might be the use of tracking bands used for workers in isolated locations, so called lone-workers. Other types include child and pet trackers that give locational data in near real time.

One module that provides cellular connectivity is the SARA-U2 from u-blox *(Fig. 1)*. This miniature LGA-sized package (16 × 26 × 3 mm) weighs under 3 g, suiting it for space-constrained wearable designs. It offers high speed 5.76-Mb/s (HSUPA) and 7.2-Mb/s HSDPA cellular data rates, and features low idle mode current consumption down to 0.9 mA.

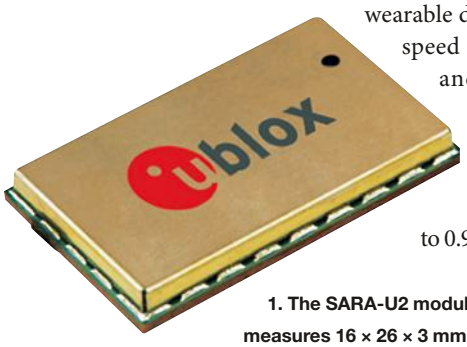

1. The SARA-U2 module developed by u-blox measures 16 × 26 × 3 mm and weighs under 3 g, suiting it for space-constrained wearable designs.

ABILITY TO LOCATE

In addition to connectivity, many wearable devices also need to track and record the wearer's location. Sports-performance monitors and cycling watches use this to overlay the wearer's heart rate to the actual latitude, longitude, and elevation.

2. An example of a GNSS module for use in wearables is the EVA-M8M. Power consumption is as low as 5.5 mA in the power-save mode, where GNSS data is updated every second.

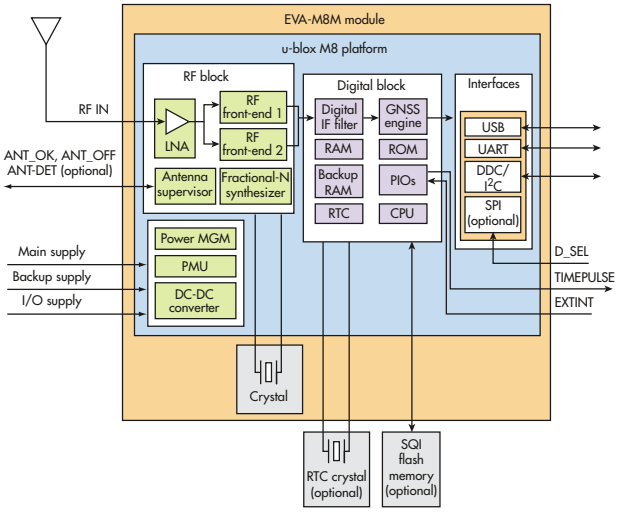

Incorporating positional capabilities into the design can be achieved in one of two ways. The most obvious is through the use of a Global Navigation Satellite System (GNSS) receiver that naturally, for any wearable application, needs to be the smallest and most power-efficient possible.

The EVA-M8M *(Fig. 2)* is an example of a GNSS module. The 43-pin, LGA-packaged device measures 7 x 7 x 1.1 mm, weighs 0.13 g, and consumes up to 25 mA in full continuous operation, but down to as low as 5.5 mA in the power-save mode, where the GNSS data is updated every second.

Like any GNSS system achieving a reliable position "fix," it relies on the antenna being able to "see" the satellite. Achieving this indoors or where the satellite signals have been reflected by large buildings, such as in dense city centers or in any area of marginal signal conditions, represents a major challenge.

Some wearable devices might need reliable indoor reception more than others, such as aforementioned lone-worker application. Should this be a design requirement, and the primary communications method is through cellular means, then a mobile-network-based positioning approach can complement the GNSS data. By maintaining a database of the positions of cellular network towers, a cellular service, such u-blox's CellLocate *(Fig. 3)* can estimate the location of the device based on previous observations from other CellLocate-enabled modules.

PRE-DEVELOPMENT CONSIDERATIONS

When trying to meet a wearable design challenge, what key steps must an engineer review before commencing development? One early step is a thorough understanding of the device requirements. Questions to ask include: "What will the wearable application monitor?"; "What sensors need to be incorporated?"; and "What are the potential use cases?"

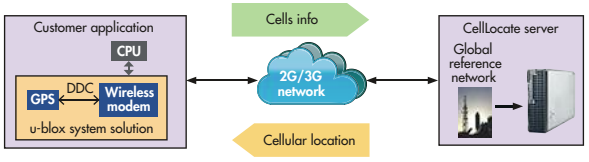

3. Mobile-network positioning technology like CellLocate can estimate the location of a device based on previous observations from other CellLocate-enabled modules.

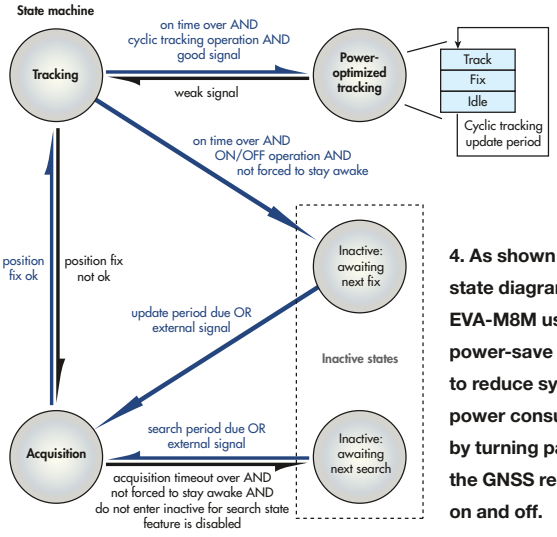

4. As shown in this state diagram, the EVA-M8M uses a power-save mode to reduce system power consumption by turning parts of the GNSS receiver on and off.

Reviewing use cases is a crucial aspect, since it will highlight key factors such as product size, available space envelope, and duty-cycle expectations. These will shape the space available for a battery and have an impact on battery capacity, which of course directly correlates to the time between charge cycles and operational duty cycle.

The various use-case scenarios will also identify the type of communications required. Does it need to use a smartphone as a gateway to a cloud application, or will it communicate directly using its own cellular data connection?

An increasingly important consideration for many Internet of Things (IoT) devices is that their firmware be updated over the air (OTA). This would eliminate the user having to download new device images to a PC from the manufacturer's site and upload to the wearable device. The specification of the host processor and the amount of memory needed to achieve OTA might need careful review should this be the case.

Most wearable devices will most likely experience the same environmental influences as the wearer. Rain, moisture, dust, and wide temperature variances all need to be taken considered for the product enclosure's design. Will it require an ingress-protection (IP) rating to satisfy the marketing specification while balancing these factors with the experience for the user?

Ending up with a wearable device that the user finds uncomfortable to wear due to its size, weight, and shape are all critical factors for future success of the product. The engineer needs to look not only at the electrical specifications of the components selected, but also their physical attributes. As noted earlier, the SARA-U2 cellular module and the EVA-M8M GNSS module together weigh just 3.13 grams, making their combination suitable for any wearable design.

GAUGING BOM VIA ANALYSIS

The marketing requirements analysis of any product will also estimate potential volumes. These will shape a number of production decisions and greatly influence the overall BOM goal.

From the connectivity standpoint, this might prompt many engineers to review whether a discrete design is better than using a module. The difference of cost against price is a hard-learned lesson. The BOM for a discrete approach might be slightly cheaper, but when factoring in the test and certification costs, the difference becomes nominal if anything. When it's not known where a device might be sold, using modules that are pre-certified to most worldwide wireless regulatory bodies can be a huge timesaver in PCB design.

Also, RF design is a specialist skill and requires equally specialist test equipment and facilities. Having to spend many weeks addressing EMI problems resulting from a poor track layout (and facing a potential PCB redesign) would negate all of the cost benefits of a discrete design.

POWER-SAVING MINDSET

Making the available power budget last as long as possible is a skill that many embedded engineers know through trial and error. When selecting connectivity and GNSS modules for a wearable design, engineers need to carefully review the module's technical documentation. They must pinpoint the methods that can both keep power consumption to a minimum while not affecting the device's responsiveness, particularly when it comes to user interaction.

Most microcontrollers and modules available today will offer several different power-saving modes; engineers need to diligently review these to find the scheme best suited to the application's needs. Typically, such modes will selectively turn parts of the module's functions on or off. For example, the u-blox EVA-M8M uses a power-save mode to reduce system power consumption by turning parts of the GNSS receiver on and off *(Fig. 4)*.

The power-saving mode is based around five different states defined as inactive (awaiting next fix and next search), acquisition, tracking, and power-optimized tracking. The power-consumption profiles differ with each state, with the acquisition consuming most power down to inactive, where most parts of the receiver are switched off. Taking full advantage of these power-saving modes in the GNSS module and similar methods available within the connectivity module will considerably enhance the wearable device's battery life and, in turn, the user experience.

REFERENCES:
1. Gartner–Forecast: Wearable Electronic Devices–Worldwide–2016, News

Prevent Oscillations in Supercapacitor Power Supplies

By RICK BOSMA, Dairy Cheq Inc.

IT'S FAIRLY EASY to add features to operating systems such as Linux when used in embedded systems. However, one of the functions needed in these systems is sufficient time to safely shut them down upon removal of system power. Not allowing for a proper shutdown can lead to data corruption to the point where the system may not reboot when power is reapplied.

One way to provide time for the system to shut down is to design-in supercapacitors (sometimes called ultracapacitors) as a backup power system. In this way, power to the processor can be maintained for up to several minutes after the mains power is cut, to allow time to write any open files to disk and properly shut down any remaining processes.

Several major vendors make ICs that control the charging and discharging of the supercapacitors. Since the voltage on these capacitors varies as they are charged and discharged, the vendors also offer dc-dc converters that convert the varying voltage to a stable voltage that's suitable for powering the rest of the system. These off-the-shelf circuits continue to provide a stable voltage until the supercapacitor voltage drops below a certain level. However, when the voltage gets to this level, commercial circuit functions are inadequate.

The voltage that's measured at the capacitor terminal is a function of the stored charge and the voltage drop due to its internal resistance. When the capacitor voltage falls below the shutoff threshold, the dc-dc converter and other loads are removed. The voltage drop that was created by the load current across the internal resistance is also removed, which increases the voltage at the capacitor terminals.

If the current draw on the system was high enough, the result of removing it can be that the capacitor voltage again becomes high enough to turn the dc-dc circuit back on. The load is reattached, the voltage drops, and the dc-dc circuit turns off again. This off/on/off cycling repeats itself and the result is continuous oscillations of the system's power network.

To prevent this power oscillation, a circuit with a controllable amount of hysteresis was needed so that it will not oscil-

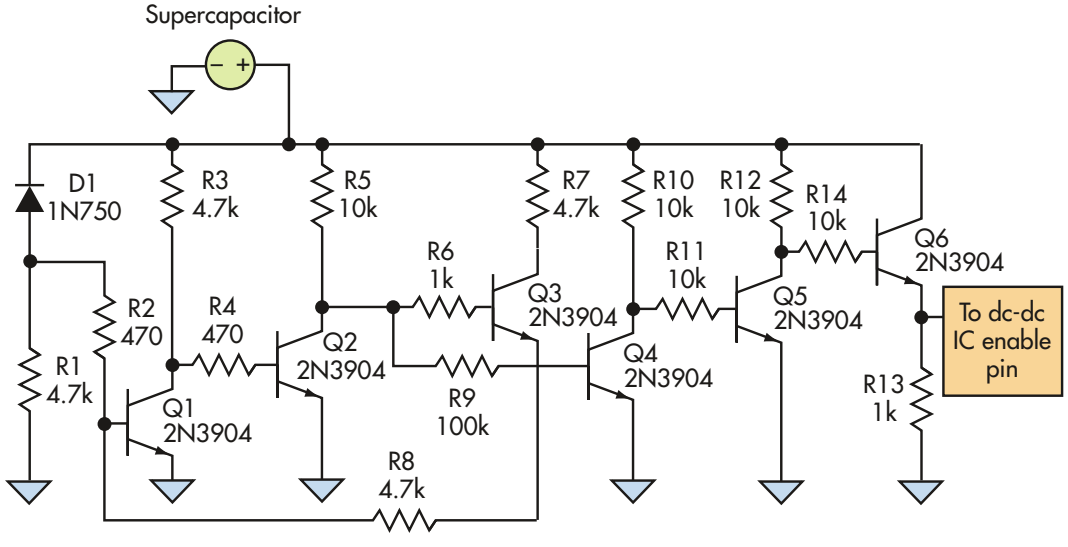

This all-analog, discrete-component circuit adds appropriate hysteresis to a charging controller for a backup-power supercapacitor. As a result, the backup rail and dc-dc power converter don't cycle through repeated turn on/off cycles due to internal characteristics of the capacitor and associated components.

late, even during worst-case load conditions. In addition, this circuit had to work from 0 to 5 V. Several op-amp circuits were considered for this application, but even the low-voltage op amps have undefined behavior when their V_{CC} rails are brought from 0 V to V_{MIN}.

The *figure* shows a circuit that can address this problem. The main circuit's power is provided by the supercapacitor voltage. As the capacitors charge from 0 V, very little current flows through Zener diode D1. As the capacitor voltage reaches approximately 4.3 V, current starts to flow though D1 and into R1. Once 150 µA of current flows and generates about 0.7 V across R1, Q1 will start to turn on. Q2 would have been turned on up to this point, but when Q1 starts to conduct through R3, it will pull the base of Q2 low and turn it off. This pulls the base of Q3 high, allowing it to start conducting current.

The current from Q3 is directed through R8. Some of this current will flow in Q1 and some will flow through R1. Therefore, rather than the voltage across R1 being dictated solely by the current through D1, it will now increase to account for the current from Q3. It's this additional current that provides the needed hysteresis and causes Q1 to turn off at a lower supply voltage compared to what it turned on beforehand.

Note that the voltage at the collector of Q2 will turn Q4 on, which turns Q5 off, and subsequently turns Q6 on. The voltage drop across R13 provides a signal to the dc-dc circuit that the voltage on the supercapacitors is high enough whereby it should turn on.

If the mains power is now removed and the system starts to draw current from the supercapacitors, the main voltage rail for this circuit will start to drop as well. However, since the voltage at R1 is due to the sum of the currents through D1 and Q3, the voltage has to fall to around 2.3 V before Q1 is turned off. This action will turn Q2 on, Q4 off, Q5 on, and finally turn Q6 off. When that happens, no current will flow through R13 and this will signal the dc-dc circuit that it should turn off as well. Once the load from the dc-dc circuit is removed, the main voltage rail will increase somewhat. However, as long as it doesn't increase by more than 2 V, the circuit will not turn on again, thus preventing power-line oscillation. Different turn-on and turn-off voltages can be realized by replacing D1 and by changing the value of R8. ed

RICK BOSMA is an electronics engineer in charge of product design at Dairy Cheq Inc. (Waterloo, Ontario, Canada). He has an Electronics Technologist Diploma from RCC College of Technology (Toronto) and a Bachelor of Engineering degree in electronics from Lakehead University (Thunder Bay). He spent seven years in the Product Engineering group at Honda Motor Company before joining Dairy Cheq in 2009.

Hybrid Switching/Linear Method Yields LED-Current Controller with Wide Dimming Ratio

By THOMAS MOSTELLER, Linear Technology Corp.

MANY APPLICATIONS FOR LED illumination require wide dimming ratios, which can be easily done using an adjustable current source (*Fig. 1*).

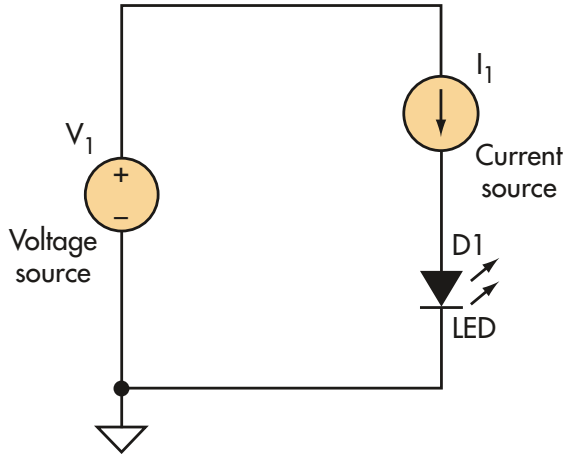

1. LEDs are current-driven, rather than voltage-driven, devices, and their output intensity is related to the current level—but not linearly.

The current source can be adjusted over a large LED current range using several different techniques. The primary problem with this technique is that the power dissipation can be quite high. Voltage source V1 has to be high enough to accommodate the largest LED voltage drop and the headroom required by the current source. LED manufacturers typically specify a maximum voltage that's higher than the average value, forcing the designer to use a higher-than-necessary input voltage. Applications that use multiple LEDs in series multiply the voltage-tolerance issue.

In addition, applications may not have access to a voltage of the optimum value to drive the LED string, forcing the designer to incorporate another supply to bring the input voltage to a usable level. This adds additional circuitry and power dissipation of its own.

This issue has been addressed by making a switching power supply have a constant-current output (rather than a constant-voltage output), and making this current adjustable to implement dimming. The concept is shown in *Figure 2* using the LTM8042.

By varying the voltage on the CTL pin from 100 mV to 1 V, the LED current can be scaled over a 10:1 range (*Fig. 3*).

This limitation in dimming ratio comes primarily from the difficulty in scaling the current-sense signal in the pulse-width-modulation (PWM) section of the switching regulator to very low levels. The dimming ratio can be extended by using the PWM signal to the LED to pulse it on and off. If this pulsing can be done at a high-enough rate—typically a few hundred hertz—the pulsing is not visible to the human eye. With this technique, the dimming ratio can be extended to 3000:1 (*Fig. 4*).

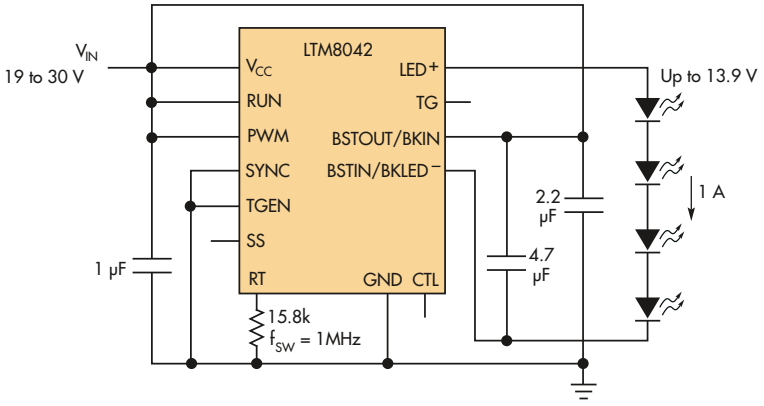

2. A power supply designed for externally adjustable constant-current output, here done with the LTM8042, meets the drive needs of LEDs.

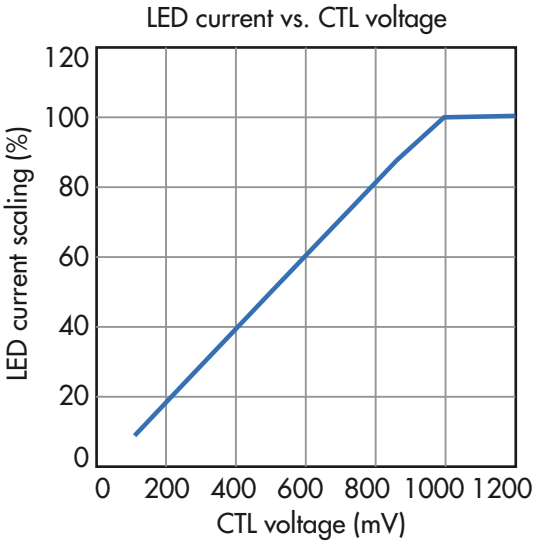

3. When the CTL pin on the LTM8042 is varied between 100 mV and 1 V, the LED current is then scaled over a 10:1 range.

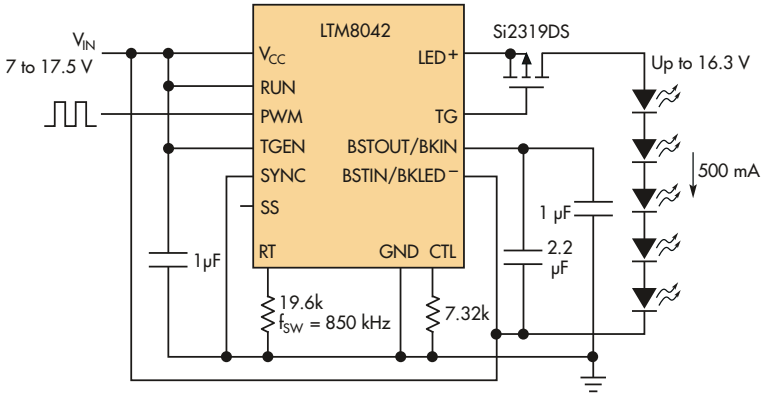

4. Using a pulse-width-modulated signal to pulse the LED on and off can extend the dimming ratio by over an order of magnitude, with some restrictions.

In this case, the LEDs are turned on and off by the clock signal applied to the PWM pin. While the pulsing isn't visible to the human eye, it can be an issue for other applications, such as when the scan rate of a digital video recorder potentially interacts with the LED flashing to produce undesirable artifacts.

The technique here allows a wide dimming ratio to be controlled by a strictly analog means. It uses a hybrid approach in which a switching regulator holds the voltage across an adjustable current source constant, and puts the LED voltage inside a feedback loop so that variations in the LED voltage don't affect the LED current. The voltage across the current source doesn't have to support the LED voltage drop and its variations, and thus can be optimized for the best choice of power dissipation versus dimming ratio.

The conceptual diagram of the hybrid approach (*Fig. 5)* leads to the implementation using the LT8614 as the voltage supply and the LT3083 as the current source (*Fig. 6)*.

With the input voltage of 12 V, the LED current is varied over a 0-to-3-A range via a 0-to-300-mV signal applied to the SET terminal of the LT3083. The IC's OUT terminal follows the SET terminal, so a constant current flows in R5 via the relation:

$$I\_R5 = V\_SET/R5$$

The current in the IN terminal of the LT3083 follows the current in the OUT terminal very precisely until it reaches a minimum load value of about 500 µA. The LT3083 needs 510-mV worst-case dropout voltage from IN to OUT, as well as 300 mV for R5 at maximum current. The LT8614 holds the voltage at the IN terminal of the LT3083 at exactly 0.97 V, which keeps the power dissipation of the LT3083 at tolerable levels, peaking at about 2.4 W at 3 A out.

The LED cathode is held at 0.97 V by the LT8614, but the anode voltage goes to whatever voltage is required due to the feedback action of the LT8614. Additional LEDs can be put in series with D1 and the current will not change until the dropout voltage of the LT8614 is reached.

The LT3083 is a convenient part for the current-source function, but (as noted previously) it does have a minimum-load requirement. In addition, there's a slight offset between the SET and OUT terminals, which is typically a few hundred microvolts but can be as high as ±6 mV

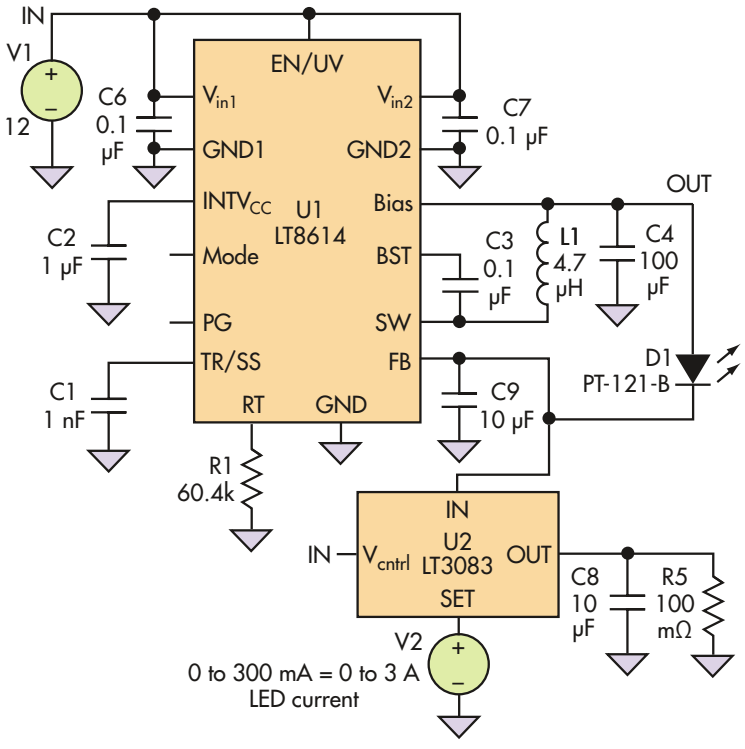

5. The hybrid approach combines a switching regulator with an adjustable current source to extend dimming range.

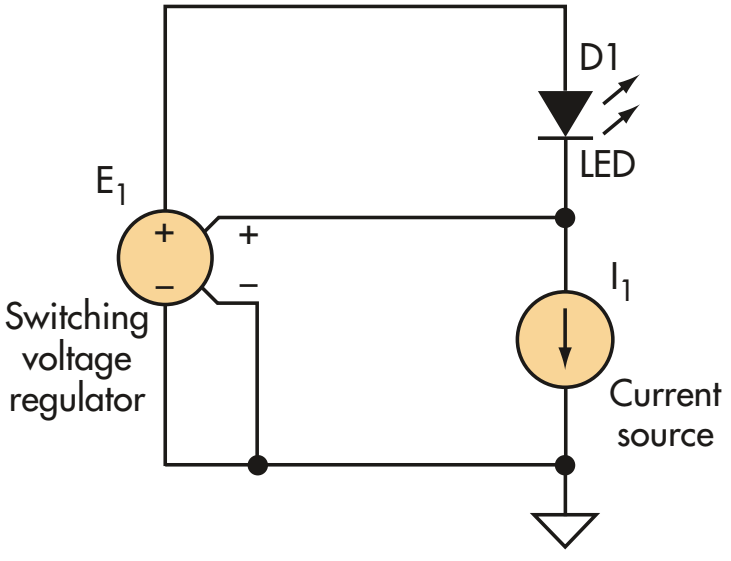

6. A basic implementation of the hybrid approach uses the LT8614 as the voltage supply and the LT3083 as the current source for a precisely controllable output.

over the temperature range. This offset places a lower limit on the dimming ratio. If a wider dimming ratio is needed, a discrete current sink using an op amp with low offset voltage can be used, such as with the LT6015 precision op amp as the control element (*Fig. 7*).

The LT6015 has a maximum voltage offset of ±250 µV over temperature, and the current source only needs about 100 mV dropout voltage. This allows the full-scale voltage to be raised to 600 mV, and permits dimming ratios in the range of 1000:1.

THOMAS MOSTELLER has been a Field Applications Engineer for Linear Technology's Middle Atlantic region since 1990. He has assisted a wide range of military, commercial, and industrial customers with designs in many fields, such as power distribution and supplies, analog signal conditioning, data conversion, and RF and communications applications. Prior to joining Linear Technology, Thomas designed medical equipment for 10 years, and holds a patent in the design of infusion pumps. He earned a BSEE from Drexel University in 1977.

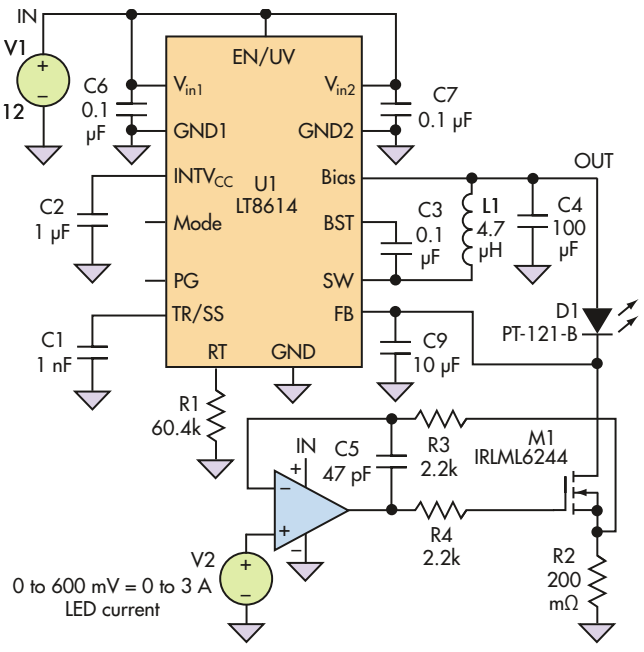

7. By adding a precision, low-offset op amp as a current sink, an even-wider dimming ratio is possible.

New Products

Measurement System Utilizes Electro-Optic Sensor

TEKTRONIX' NEW IsoVu Measurement System uses an electro-optic sensor to convert input signals to optical modulation, electrically isolating the device-under-test from a Tektronix oscilloscope. The platform incorporates four separate lasers, an optical sensor, five optical fibers, and sophisticated feedback and control techniques. For power devices involving GaN and SiC technologies, superior common mode rejection is a critical advantage that makes signals, previously buried in common mode noise, now visible. The optically isolated measurement system offers 1 Million:1 (120 dB) CMRR up to 100 MHz and 10,000:1 (80 dB) CMRR at 1 GHz.

Small differential signals (5 mV - 50 V) can be accurately measured in the presence of large common mode voltages from dc to 1 GHz, and the common mode voltage capability will not de-rate over bandwidth. IsoVu technology is available in 6 models of the TIVM Series Isolated Measurement Systems with 200 MHz, 500 MHz and 1 GHz bandwidth configurations with either 3 m or 10 m fiber optic cable lengths. The 10 m cable option offers the same performance specifications as the 3 m option, while also allowing systems to be moved away from the interference and radiated emissions of the device under test. Priced starting at $12,000, the IsoVu Measurement System is now shipping.

TEKTRONIX

www.tek.com

Enterprise Platform Addresses Rigid-Flex and High-Speed Designs

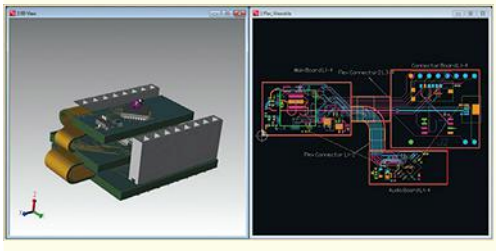

MENTOR GRAPHICS is implementing the first phase of its Xpedition PCB flow to address the increasing complexity of advanced systems designs. The platform provides advanced technologies to enable design and verification of 3D rigid-flex structures, and to automate layout of high-speed topologies with advanced constraints. The Xpedition rigid-flex technology enables a streamlined design process from initial stack-up creation through manufacturing. Complex rigid and flex PCBs can be designed in a fully supported 3D design and verification environment, resulting in a correct-by-construction methodology for optimum reliability and product quality.

Integration with high-speed analysis technology enables optimization of signal and power integrity across complex rigid-flex stack-up structures. For fabrication preparation, the PCB flow provides all flex and rigid information using the ODB++ common data format, communicating the finished board intent to the fabricator.

The new release also features Xpedition xPCB Layout, an advanced layout automation to address increasing complexity in high-speed designs and emerging guidelines for high-end computing chip-sets. The Xpedition flow for flex/rigid-flex and automated high-speed design layout is available today.

MENTOR GRAPHICS

https://www.mentor.com

Security Solution Transforms/Encrypts SSDs to eDrive

ENOVA TECHNOLOGY'S 10th generation hardware security solution for Microsoft EHDD has built-in TCG Opal 2.0 capability to transform SATA disks or SSDs to an Opal 2.0 compliant drive. When working in conjunction with another built-in IEEE1667 firmware, the product can transform any number or capacity of standard SATA disks or SSDs into an eDrive, wherein Microsoft BitLocker can manage and configure it, through the software interface of the IEEE1667 and Opal 2.0. In this application it can be configured as either a boot drive, a data drive, or a portable USB3.0/3.1 drive.

The product can be configured to perform in-line 6 Gbps full disk encryption with AES CBC/XTS 256-bit strength protecting entire drive in Data-at-Rest mode. With add-on software components, the security device can be configured to perform file level encryption, encrypting Data-in-Motion such that files stored in the Cloud storage remain MX+ hardware encrypted. Armed with built-in RSA 2048, HMAC, CMAC, AES ECB/XTS/CBC, SHA256, Hash_DRBG and TRNG, the X-Wall MX+ is able to generate signature, sign and verify each identity.

ENOVA TECHOLOGY

www.enovasol.com

Expansion Chassis Provides PCIe Slot to Thunderbolt 3

THE THUNDER3 PCIE BOX is designed by Akitio to provide an additional PCIe slot for Thunderbolt 3 computers where there is either not enough space to install another PCIe card, or on systems such as laptops where there previously was no option to add a PCIe card. The screw-less PCIe box features a 1x16 slot that supports half-length, full-height, double width cards. The 1x16 PCIe slot supplies up to 25 W of power to the card. Two Thunderbolt 3 ports provide daisy chaining of up to 6 devices, such as Akitio's Thunder3 Duo Pro dual-bay, Thunder3 Quad 4-bay or Thunder3 PCIe SSD storages, with a data transfer rate of up to 2,500 MB/s.

The four PCIe Gen3 lanes and the Thunderbolt 3 interface are capable of transfer rates up to 40 Gbps. Sony Vegas and Adobe Premier users can edit and playback multiple streams of HD and/or 4K content. The box can be daisy chained to other Thunderbolt 3 devices. The empty Thunder3 PCIe Box will be available in August 2016 with an MSRP of $299.99, and includes a 40 Gbps 0.5 m Thunderbolt 3 cable.

AKITIO

https://www.akitio.com

Multiport Minidock Reference Design Utilizes USB Type-C

TEXAS INSTRUMENTS' multiport USB Type-C and Power Delivery Minidock Reference Design (TIDA-01243) provides audio, USB data, power, and video support. The reference design presents a fully tested, verified plan for a 2-in. x 4-in. dock, and can be powered by a traditional power adapter, USB Type-C adapter, or notebook computer.

At the core is the TPS65982 USB Type-C and PD controller. All ICs within the design, including port controllers, data multiplexers, dc/dc converters, load switches, FETs, signal conditioners and circuit protection, are compliant with the current USB Type-C and PD standards. The TPS65982 enables dual-role port USB Type-C functionality, capable of delivering 60 W of power and supporting DisplayPort Alternate Mode and USB data transmission. The design also features the TPS65986 USB Type-C and PD controller, HD3SS3212 SuperSpeed multiplexer, TUSB321 dual-role port controller and HD3SS460 Alternate Mode multiplexer, which work together to provide dual-role port capability for power, data and video transfer.

The USB Type-C and Power Delivery Minidock with Video and Charging Support Reference Design (TIDA-01243) is available for download now. The USB-CTM-MINIDK-EVM is also available for $499.

TEXAS INSTRUMENTS
www.ti.com

Dual 100 G Network Interface Eliminates External Switch

THE SHARPSWITCH PCIE-9205 PCIe intelligent network interface card from Artesyn Embedded eliminates the need for an external load balancer in wireless, communications, broadcast and streaming media applications. An Intel Xeon D series processor and Intel Ethernet multi-host controller (code name Red Rock Canyon) provide dual 100G Ethernet interfaces and 100G switch, and suit the card for low power, dense computing applications.

By focusing on the minimal requirements for typical applications, and building on top of Intel DPDK, Artesyn has designed an accelerated virtual switch. By using the interface card as an Open vSwitch accelerator, the vast majority of the available processor cycles in the system can be used for payload Vms.

The SharpSwitch PCIE-9205 PCIe intelligent network interface card is optimized to work in Artesyn's MaxCore PCIe platforms which leverage off-the-shelf PCIe cards to create very low latency and low OpEx systems for applications such as vRAN, mini-CRAN, video transcode, VoLTE and security gateways.

ARTESYN EMBEDDED
https://www.artesyn.com

Multi-Vendor Calibration Services Utilize PXI, AXIe Instruments

KEYSIGHT TECHNOLOGIES is extending its One-Stop Calibration Services to include non-Keysight equipment. The services ensure ongoing accuracy and availability of test assets, as well as electrical, physical, dimensional and optical equipment from virtually any manufacturer. Keysight is also expanding its offering of high-performance PXI and AXIe instruments and reference solutions. Using benchtop, modular or a combination, the products deliver proven test configurations for critical applications such as 5G, PA/FEM, EW threat simulations and digital interconnect test. The new offerings include:

The M9019A 18-slot Gen 3 chassis, controller, I/O components and external PC option PXIe system components boosts system bandwidth by at least 2x more than a typical Gen 2 chassis.

The 4-slot PXIe vector transceiver and the PXIe high-speed SMU improve test throughput, reduce the overall footprint and are built for production test and design validation of next-generation PA/FEMs. Covering frequencies from 60 MHz to 6 GHz, with modulation and analysis bandwidth up to 160 MHz, the VXT supports FPGA-accelerated measurements and includes software supporting analog demodulation, noise figure, and cellular and wireless connectivity formats. The SMU changes voltage, stabilizes and accurately measures µA, in less than 1 ms.

A new Digital Interconnect Test reference solution features the PXI VNA with full cross-bar s-parameter calibration of 32-ports up to 26.5 GHz.

KEYSIGHT TECHNOLOGIES

www.keysight.com

Highly Integrated Hall Sensors Boost ESD Robustness

INFINEON TECHNOLOGIES' TLx496x series Hall sensors provide high ESD robustness and precise, stable magnetic switching points while consuming no more than 1.6 mA (5 V versions consume 1.4 mA). While meeting the requirements of compact designs, all of the sensors have an integrated Hall element, a voltage regulator, choppers, an oscillator and an output driver. The TLE496x-xM series is suitable for automotive applications with an operating voltage of 3.0 V to 5.5 V, which generally are not exposed to overcurrent. Target applications are all systems that need a precise Hall switch or Hall latch for an expanded temperature range of -40°C to 170°C. The TLE496x-xM is typically used in power windows and sunroofs, trunk locks, windshield wipers, seatbelts, camshafts and shift levers.

The TLI496x-xM series is functionally the same as the TLE496x-xM but is specified for a temperature range of -40°C to 125°C and is JESD47 qualified. The TLI496x-xM sensors find typical use in BLDC motors, PC fans, electric drives in building automation, and in the detection of open/close states in white goods and building security.

TLV496x-xTA/B versions have a power consumption of 1.6 mA and an ESD protection up to 4 kVH HBM. Available as latch and switch-type sensors, TLx496x series Hall sensors are specifically developed for contactless positioning.

INFINEON TECHNOLOGIES

https://www.infineon.com/

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

Lab Bench
WILLIAM WONG | Embedded/Systems/Software Editor
bill.wong@penton.com

More Storage Changes Are Coming

The Flash Memory Summit offers a host of new approaches to storage with ever-increasing capacities.

The Flash Memory Summit continues to be the place where the latest storage technology is announced. This year is no different with a host of new approaches to storage in addition to ever-increasing capacities. On the horizon is quad-level cell (QLC) flash memory that stores four bits of information in a single cell. This extends the current crop of SLC, MLC, and TLC flash memory technologies.

There were impressive multilayer solutions like 64-layer NAND from Toshiba, Western Digital, and Micron. These build on the 32- and 48-layer NAND implementations from previous years. The latest implementations can pack 4 Tbytes onto a single M.2 module. Intel's XPoint memory will be highlighted at this year's Intel Developers Forum, which is just after the Flash Memory Summit.

There were a couple of other technologies that stood out that would be of special interest to embedded developers. Everspin is delivering its 256 Mbit DDR3 chips based on its perpendicular magnetic tunnel junction (pMTJ) ST-MRAM (see *"MRAM Optimizes System Energy Consumption" on electronicdesign.com*). MRAM is faster than flash memory and does not suffer from a limited write lifetime. This is one reason that Diskfull Writes Per Day (DWPD) is a common specification for solid-state drives (SSDs).

Aupera's Aup-AXL-M128 M.2 module (*Fig. 1*) uses Everspin's new MRAM chip. The M.2 module is delivering four orders of magnitude BER reduction. It also uses 30% less power. Everspin also demonstrated an NVMe board using its chip that delivered 1.5 million write IOPS. Controllers can be simpler than their flash-memory counterparts because MRAM does not require wear leveling.

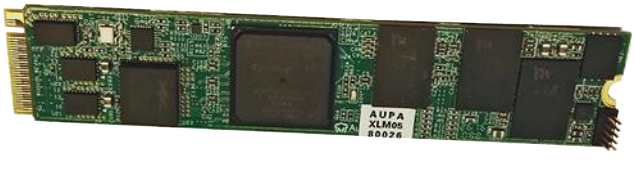

1. Aupera's M.2 module uses Everspin's 256 Mbit MRAM chips to provide storage that is faster and has a longer write lifetime than any NAND solution on the market.

Another switch from the conventional flash-storage methodologies was from Diablo Technologies, which is now shipping its 256 Gbyte Memory 1 DIMM (*Fig. 2*). This 256 Gbyte flash memory module fits in a DDR4 slot, but it is not destined as non-volatile storage when mixed with Diablo's DMX software. Instead, the combination is taking advantage of flash memory's higher density and lower cost compared to DRAM.

2. Aupera's M.2 module uses Everspin's 256 Mbit MRAM chips to provide storage that is faster and has a longer write lifetime than any NAND solution on the market.

A typical server would fill its memory slots with a combination of DDR4 DRAM and Memory 1 modules, usually with more Memory 1 modules. This allows up to 2 Tbytes of storage on the processor's memory channel, but it will not be accessed like DRAM. Instead, it acts as a second-level cache. The DMX software hooks into Linux's virtual memory support and moves data to Memory 1 as it becomes mostly read-only. This is typical for a large number of most applications like database servers. Data that is updated frequently remains in DRAM and cached data is brought back into DRAM as needed.

Storage is getting increasingly complex. Processors typically have a few cache levels. This adds another for main memory, and larger server farms typically have more for SSDs and disk drives. Even those have caching internally. On the plus side, most of this is transparent to developers and their applications.

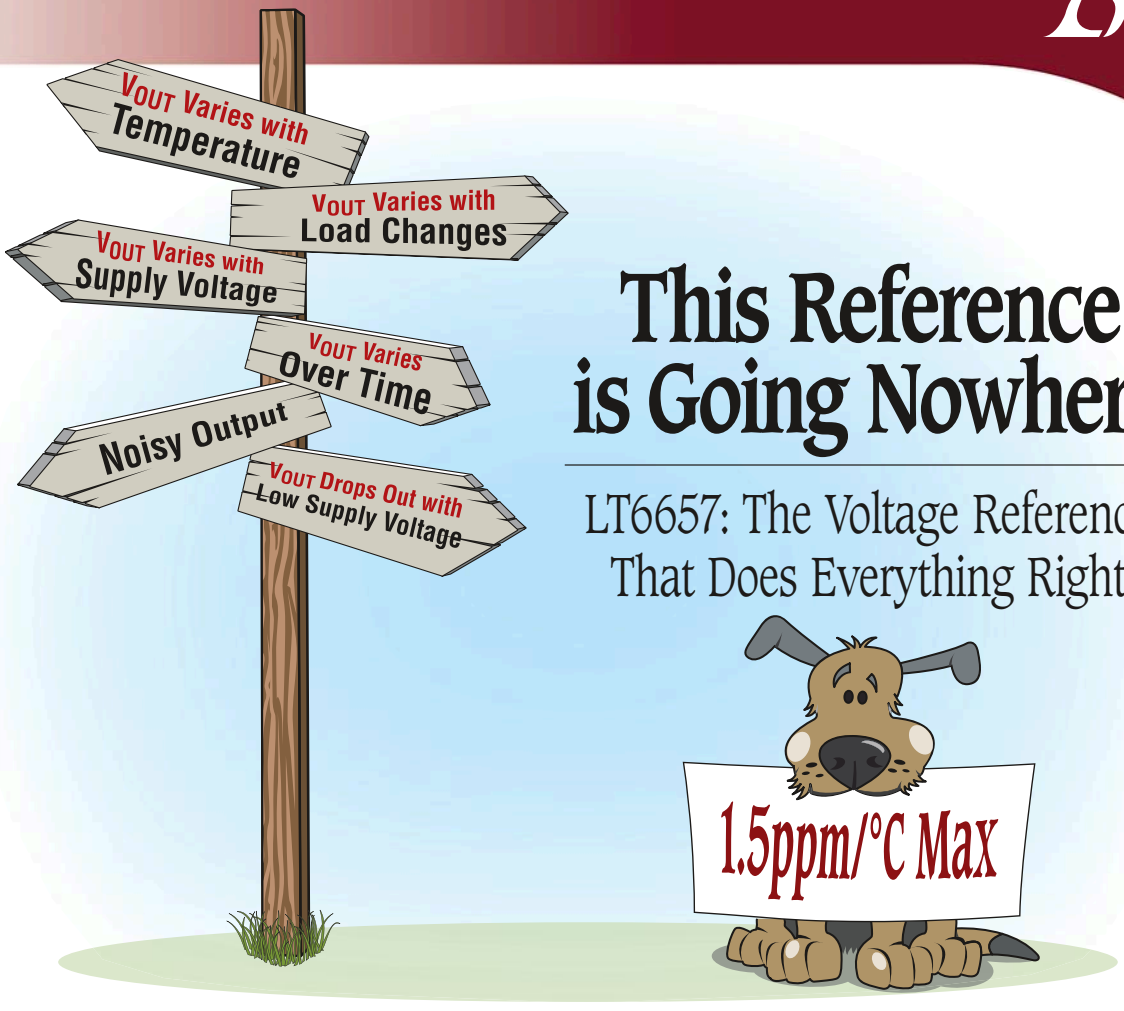

This Reference is Going Nowhere

LT6657: The Voltage Reference That Does Everything Right

The best voltage references maintain their output voltage under all operating conditions. Linear Technology's LT®6657 bandgap voltage references have less than 1.5ppm/°C of temperature drift. They also exhibit extremely low noise, and excellent long term drift, load regulation, line regulation and thermal hysteresis. All these features minimize the impact of often overlooked, but significant sources of output voltage changes. Think of the well-rounded LT6657 as a loyal companion. It never strays.

Features

- Excellent Temperature Drift:
  - A-Grade: 1.5ppm/°C Max
  - B-Grade: 3ppm/°C Max
- 0.5ppm$_{P-P}$ Noise (0.1Hz to 10Hz)
- 4ppm/mA Max Load Regulation
- 2ppm/V Max Line Regulation for V_{IN} Up to 40V
- Low Power Shutdown: <4µA Max
- Can Operate in Shunt Mode

Output Voltage Temperature Drift

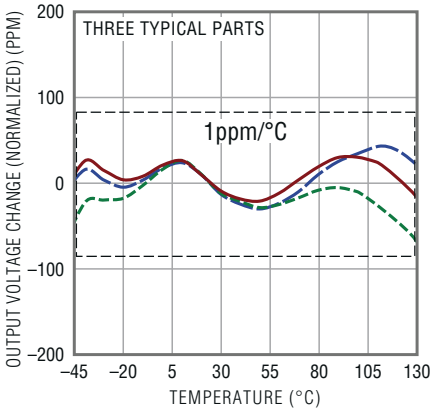

Info & Free Samples

www.linear.com/product/6657

1-800-4-LINEAR

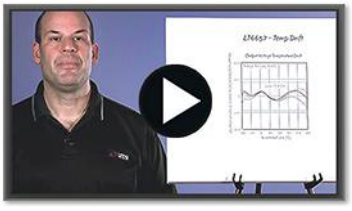

video.linear.com/6029