Electronic Design®

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

MARCH/APRIL 2018
electronicdesign.com

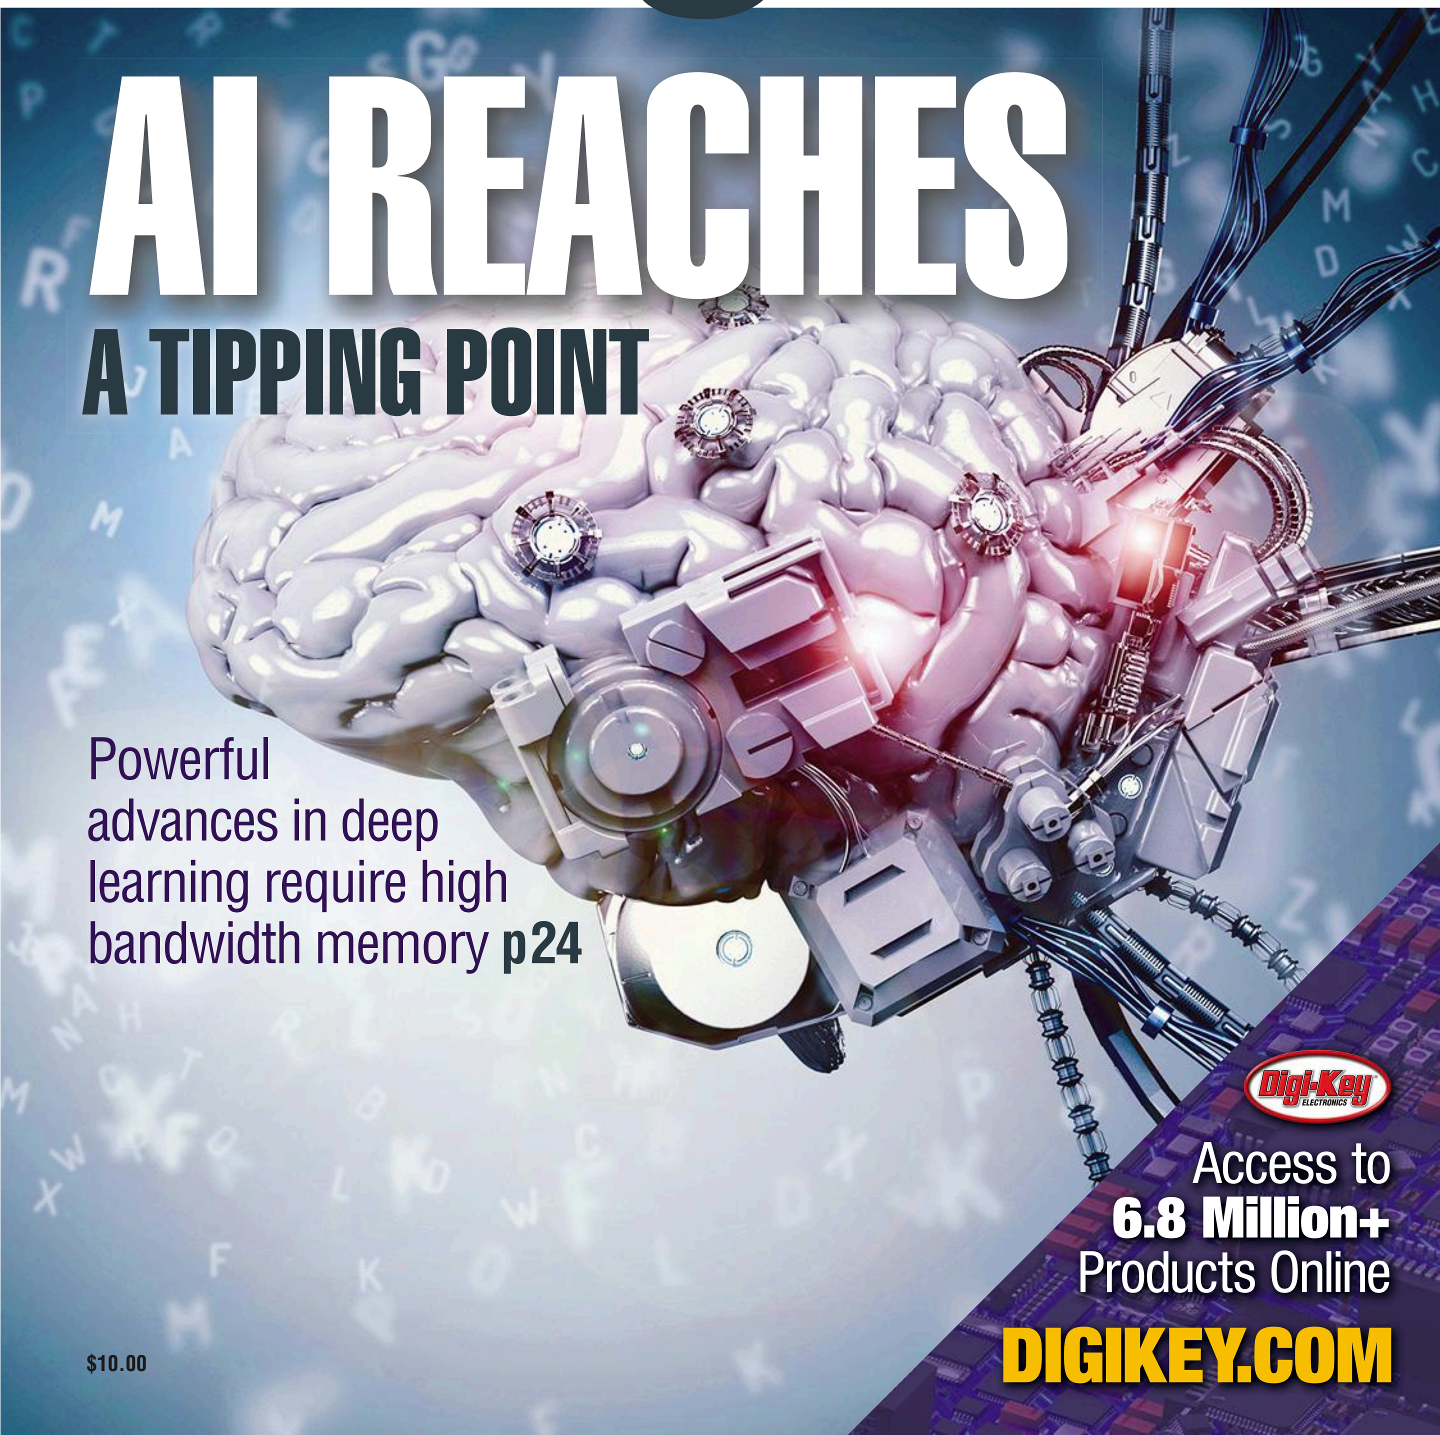

AI REACHES
A TIPPING POINT

Powerful advances in deep learning require high bandwidth memory p24

Digi-Key ELECTRONICS
Access to
6.8 Million+
Products Online
DIGIKEY.COM

$10.00


WHAT.
6.8 MILLION+ PRODUCTS ONLINE
WHEN.
99% OF ORDERS SHIPPED SAME DAY
WHERE.
WHEREVER YOU NEED IT
SAME-DAY SHIPPING
on orders entered by 8:00 pm CENTRAL
NEXT-DAY DELIVERY
1.800.344.4539
DIGIKEY.COM
Digi-Key ELECTRONICS
1,400,000+ PRODUCTS IN STOCK | 750+ INDUSTRY-LEADING SUPPLIERS | 100% AUTHORIZED DISTRIBUTOR
Digi-Key is an authorized distributor for all supplier partners. New products added daily. Digi-Key and Digi-Key Electronics are registered trademarks of Digi-Key Electronics in the U.S. and other countries. © 2018 Digi-Key Electronics, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA
ecia MEMBER
ecsn member
CEDA

Take the path of least resistance

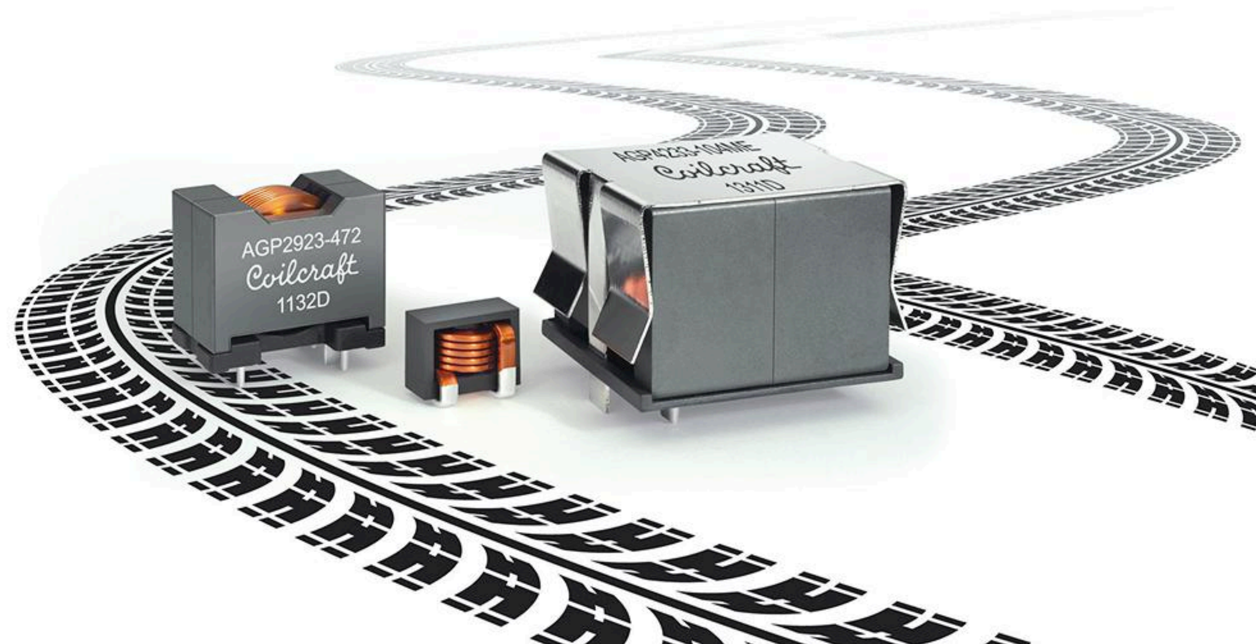

Increase the efficiency of your power electronics systems with our ultra-low DCR power inductors

Our AGP Family of high current power inductors feature flat wire construction for exceptionally low DC and AC resistance, letting your power systems run cooler and more efficiently.

They're offered with inductance values as high as 470 µH and current ratings up to 108 Amps.

AEC
Q200 Grade1

And all AGP Family inductors meet stringent AEC-Q200 Grade 1 quality standards, making them suitable for automotive and other harsh-environment applications.

Visit **www.coilcraft.com/AGP** to learn more and arrange a test drive. Free samples are always just around the corner!

Coilcraft®
WWW.COILCRAFT.COM

coilcraftdirect.com
No min order. Next day delivery.

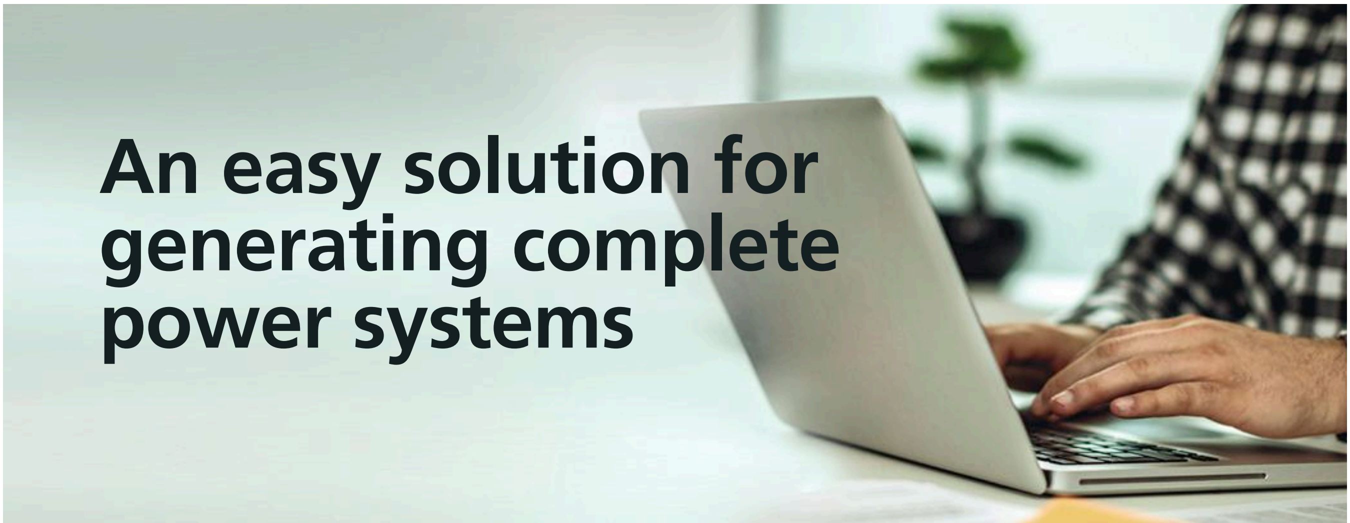

An easy solution for generating complete power systems

Just enter a few specs to design your next power system

- Design your power system in a single location
- Up to 75% faster than traditional methods
- Easy as entering your input and output as well as your basic system requirements

Advantages of using the Power System Designer

- Easy to use by both novice and expert engineers
- Meet project schedules quicker and easier
- Achieve power system performance without trial and error
- Mitigate risk of power spec changes

To find out how to design a power system in under two minutes using the Vicor Power System Design Tool, visit

vicorpower.com/power-system-designer

VICOR

IN THIS ISSUE

FEATURES

12

14

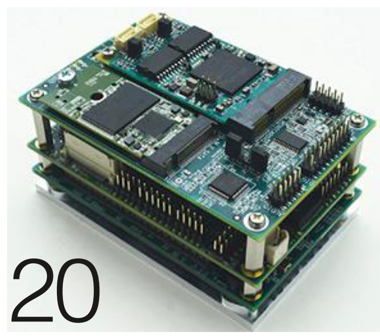

20

COLUMNS & DEPARTMENTS

24

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

ELECTRONIC DESIGN (ISSN 0013-4872) is published bi-monthly in Jan/Feb, Mar/Apr, May/June, July/Aug, Sept/Oct and Nov/Dec by Informa Media Inc., 9800 Metcalf Ave., Overland Park, KS 66212-2216. Paid rates for a one-year subscription are as follows: $60 U.S., $90 Canada, $120 International. Periodicals postage paid at Kansas City, MO, and additional mailing offices. Editorial and advertising addresses: ELECTRONIC DESIGN, 1166 Avenue of the Americas, New York, NY 10036. Telephone (212) 204-4200. Printed in U.S.A. Title registered in U.S. Patent Office. Copyright ©2018 by Informa Media Inc. All rights reserved. The contents of this publication may not be reproduced in whole or in part without the consent of the copyright owner. For subscriber services or to order single copies, write to Electronic Design, PO Box 2100, Skokie, IL 60076. POSTMASTER: Send change of address to Electronic Design, PO Box 2100, Skokie, IL 60076. Canadian Post Publications Mail agreement No. 40612608. Canada return address: IMEX Global Solutions, P.O. Box 25542, London, ON N6C 6B2.

Permission is granted to users registered with the Copyright Clearance Center Inc. (CCC) to photocopy any article, with the exception of those for which separate copyright ownership is indicated on the first page of the article, provided that a base fee of $2 per copy of the article plus $1.00 per page is paid directly to the CCC, 222 Rosewood Drive, Danvers, MA 01923 (Code No. 0013-4872/94 $2.00 + $1.00). Copying done for other than personal or internal reference use without the express permission of Informa Media, Inc. is prohibited. Requests for special permission or bulk orders should be addressed to the editor. To purchase copies on microfilm, please contact National Archive Publishing Company (NAPC) at 732-302-6500 or 800-420-NAPC (6272) x6578 for further information.

ON ELECTRONICDESIGN.COM

Full Material Declarations: Removing Barriers to Environmental Data Reporting

Suppliers and customers face challenges when it comes to full material declaration (FMD) of product content in electronics (and other industries). Thus, the IPC-1752A Standard was created, specifying an XML schema for mandatory data. The article focuses on requirements for tools that enable rapid, accurate reporting of class D FMDs.

http://www.electronicdesign.com/industrial-automation/full-material-declarations-removing-barriers-environmental-data-reporting

Advancing the Artificial Eye with a Controllable Optical Metalens

Lenses and their focusing mechanisms tend to be bulky components in electro-optical systems. To combat that, researchers used nanoparticles and electrically controlled elastomers to develop a 30-µm-thin metalens that can dynamically correct for astigmatism, focus, and image shift.

http://www.electronicdesign.com/embedded-revolution/advancing-artificial-eye-controllable-optical-metalens

Is C the Best Embedded Programming Language?

C still dominates the embedded programming space, but the Internet of Things has brought Javascript and Python into the fray. Such exposure to multiple languages can be useful, but does the language matter? Senior Tech Editor William Wong weighs in.

http://www.electronicdesign.com/embedded-revolution/c-best-embedded-programming-language

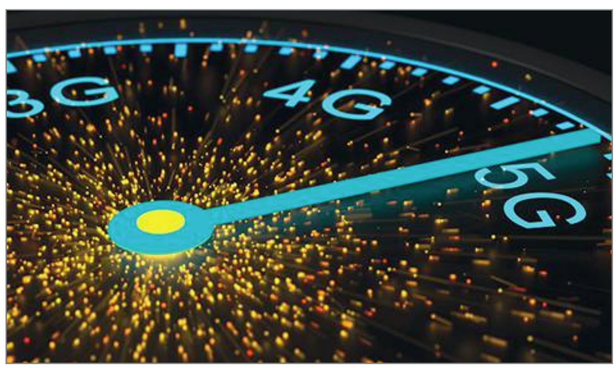

The Many Frequencies of 5G

The final standards aren't set, and we're still in the trial phase, but it's certain that speeds will be important in consumer take up of 5G. Flexible RF transceivers and software-defined-radio base stations will be key players, making it more cost-effective to switch capacity from LTE.

http://www.electronicdesign.com/embedded-revolution/many-frequencies-5g

EXTINCT NO MORE.

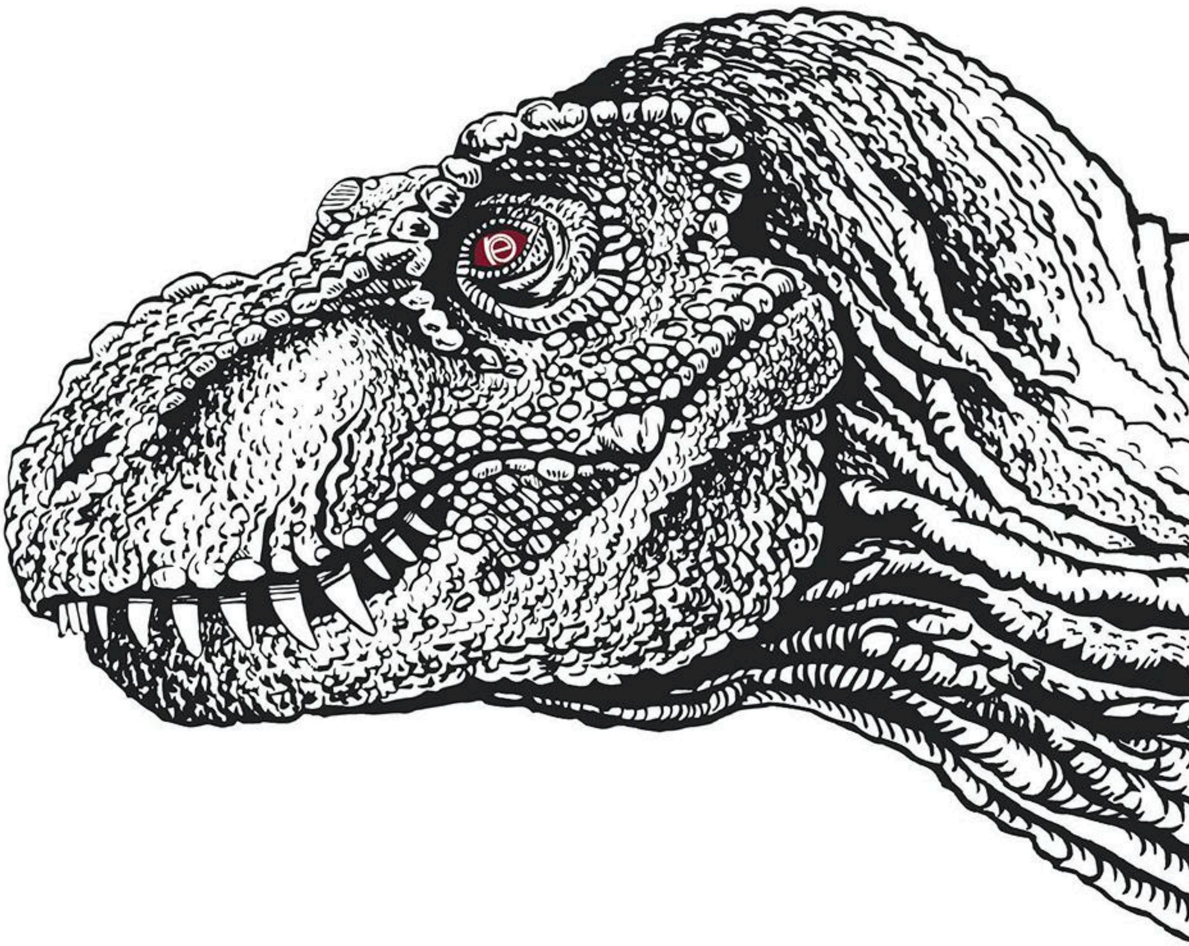

The world's largest continuous source of semiconductors.

100% AUTHORIZED AND GUARANTEED

Rochester Electronics®

www.rocelec.com

© 2018 Rochester Electronics, LLC. All Rights Reserved. 01232018

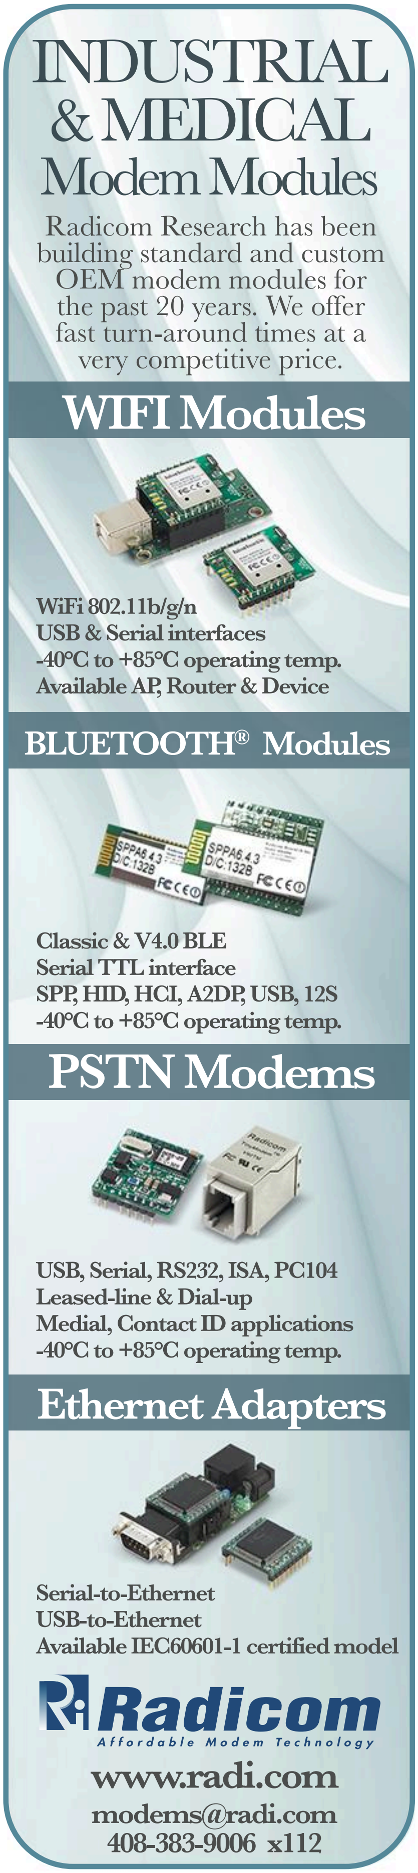

ElectronicDesign. **MARCH/APRIL 2018**

EDITORIAL

EXECUTIVE DIRECTOR, CONTENT: **KAREN FIELD** karen.field@informa.com
ASSOCIATE EDITOR/COMMUNITY MANAGER: **ROGER ENGELKE** roger.engelke@informa.com
ASSOCIATE EDITOR/COMMUNITY MANAGER: **JEREMY COHEN** jeremy.cohen@informa.com
EMBEDDED/SYSTEMS/SOFTWARE: **WILLIAM WONG** bill.wong@informa.com
ANALOG/POWER: **MARIA GUERRA** maria.guerra@informa.com
ASSOCIATE CONTENT PRODUCER: **JAMES MORRA** james.morra@informa.com
CONTRIBUTING EDITOR: **LOUIS E. FRENZEL** lou.frenzel@informa.com

ART DEPARTMENT

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

PRODUCTION

GROUP PRODUCTION MANAGER: **GREG ARAUJO** greg.araujo@informa.com
PRODUCTION MANAGER: **JULIE GILPIN** julie.gilpin@informa.com

AUDIENCE MARKETING

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

FREE SUBSCRIPTION / STATUS OF SUBSCRIPTION / ADDRESS CHANGE / MISSING BACK ISSUES
OMEDA **T** | 847.513.6022 **TOLL FREE** | 866.505.7173

SALES & MARKETING

MANAGING DIRECTOR: **TRACY SMITH** **T** | 913.967.1324 **F** | 913.514.6881 tracy.smith@informa.com

REGIONAL SALES REPRESENTATIVES:

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

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

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

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

INTERNATIONAL SALES:

GERMANY, AUSTRIA, SWITZERLAND: **CHRISTIAN HOELSCHER** **T** | 011.49.89.95002778
christian.hoelscher@husonmedia.com

BELGIUM, NETHERLANDS, LUXEMBURG, UNITED KINGDOM, SCANDINAVIA, FRANCE, SPAIN, PORTUGAL:
JAMES RHOADES-BROWN **T** | +011 44 1932 564999 **M** | +011 44 1932 564998 james.rhoadesbrown@husonmedia.com

ITALY: **DIEGO CASIRAGHI** diego@casiraghi-adv.com
PAN-ASIA: **HELEN LAI** **T** | 886 2-2727 7799 helen@twoway-com.com
PAN-ASIA: **CHARLES LIU** **T** | 886 2-2727 7799 liu@twoway-com.com

PLEASE SEND INSERTION ORDERS TO: orders@informa.com
INFORMA REPRINTS: **WRIGHT'S MEDIA** **T** | 877.652.5295
LIST RENTALS/ SMARTREACH CLIENT SERVICES MANAGER: **MARY RALICKI** **T** | 212.204.4284 mary.ralicki@informa.com

DIGITAL

GROUP DIGITAL DIRECTOR: **RYAN MALEC** ryan.malec@informa.com
CLIENT SUPPORT MANAGER: **KIM BLASKI** kim.blaski@informa.com

DESIGN ENGINEERING & SOURCING GROUP

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

INFORMA MEDIA INC.

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

informa

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

• UL Recognized Components • Easy, safe and quick way of connecting and disconnecting wires • Cost effective PCB connectivity • Sturdi-Mount features for increased PCB mounting strength • Durable Tin Plating for excellent solderability • Wide variety of configurations to suit most applications • Male tab sizes for .110"; .187"; .205"; .250" • Female receptacle sizes for .110"; .187"; .205"; .250"; .312" • Custom terminals can be manufactured to your specifications • Request Catalog M65

IT'S WHAT'S ON THE INSIDE THAT COUNTS

KEYSTONE®
ELECTRONICS CORP.

www.keyelco.com • (516) 328-7500 • (800) 221-5510

GPS disciplined **10 MHz reference ...**

... and so much more !

- GPS/GNSS disciplined 10 MHz
- TCXO, OCXO or Rb timebase
- Time tagging to GPS and UTC
- Frequency counter with 12 digits/s
- Source out: sine, square, triangle & IRIG-B
- Built-in distribution amplifiers
- Ethernet and RS-232 interfaces

The FS740 GPS disciplined10 MHz reference delivers cesium equivalent stability and phase noise at a fraction of the cost.

It's host of features includes a 12-digit/s frequency counter, a DDS synthesized source with adjustable frequency and amplitude, built-in distribution amplifiers, and event time-tagging with respect to UTC or GPS.

The optional OCXO or rubidium clock (PRS10) provide better than −130 dBc/Hz phase noise.

FS740 ... $2495 (U.S. list)

SRS Stanford Research Systems
Tel: (408) 744-9040 • www.thinkSRS.com
www.thinkSRS.com/products/FS740.htm

Editorial

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

Wishing Nancy Well

Nancy Friedrich has been the Executive Director, Content, Design Engineering & Sourcing for our group for many years and is a long-time Penton, now Informa, alum who is moving on to new projects. We wish her well, though she will be missed.

She is the latest of many notable editors with *Electronic Design*, including Joe Desposito, David Bursky, Lucinda Mattera, and Roger Allan. I worked with all of them as well as Nancy and found her to be one of the best at inspiring the group and leading the way as we transitioned to our online presence while delivering our print editions that many still prefer.

Nancy Friedrich

Nancy has a 20-year-plus track record here, including a stint as technology editor for *Wireless System Design* magazine. She was also Editor-in-Chief of *Microwaves & RF* magazine that covered the transition to LTE and now 5G. As Executive Director, Content, Design Engineering & Sourcing, she oversaw all of the publications in our group, including *Electronic Design, Machine Design, Microwaves & RF,* and *Hydraulics & Pneumatics.* Managing such a diverse group is not easy, as most of us are engineers who have taken up writing as a second career.

Nancy received her Bachelor's in English Language and Literature/Letters from Rutgers University. Nancy said, "When I started my career almost 20 years ago, I was fresh out of college with dreams of writing the next great novel or being a college professor. Taking some time to myself to decide which path I would take, I started a career in technical editing and quickly discovered a different road—one that I never would have imagined." We are glad she took that turn.

Nancy was instrumental in our move to all our new websites in the group. They now provide a more consistent presentation, especially on smartphone and tablets, which are more popular ways to consume content these days. We also started up innovation-destination.com. This is our new automotive technology website that presents the latest in self-driving cars, infotainment, and advanced driver-assistance systems (ADAS).

Electronic Design has been around a long time. It started in 1952 and is still going strong. It has changed substantially from when she started. She now joins other the editors who have retired or moved on from our publication.

Karen Field, Executive Director, Content, has picked up where Nancy left off as we at *Electronic Design* continue to bring the latest in engineering and technology to our readers. Karen's leadership is key to our **#fiercelyindependent** focus on delivering the highest-quality technical content across our publications and their many outlets, including webinars and shows such as the new *Electronic Design Connect* this fall. ed

JOIN US ONLINE!
twitter.com/ElectronicDesign
facebook.com/ElectronicDesign

News

XILINX ADAPTS TO AN ADAPTIVE Future of Computing

When he was promoted to chief executive of Xilinx, Victor Peng said that the job was a rare opportunity to capitalize on the fundamental changes taking place in computing. In his previous roles as executive vice president of product and then briefly chief operating officer of the FPGA maker, he became acutely aware of the challenges.

He had overseen the development of the last three generations of Xilinx FPGAs and has been working to thrust them further into the market for accelerated computing and artificial intelligence. After two months as C.E.O., Peng charted a similar course for the San Jose, California-based company, one focused on bringing the benefits of programmable silicon to the masses.

He announced what he called a completely new product category called the adaptive compute acceleration platform (ACAP). These heterogenous multicore chips can be programmed to accelerate a wide range of tasks, giving them similarities to traditional FPGAs. But users can program ACAPs without going into the guts of the hardware. They can use a software language like Python instead.

Peng also announced a new strategy that doubles down on data centers, where Xilinx has pushed FPGAs to accelerate applications ranging from machine learning tasks like image recognition to processing human genomes in minutes versus days. The company's FPGAs can also be used in server storage and networking.

The central component of that strategy is the ACAP and the first processor family based on the technology, codenamed Everest. The company said that the Everest line of chips can accelerate machine learning tasks with better performance-per-watt than general purpose chips like CPUs and GPUs, which are the current gold standard.

"While FPGA and Zynq SoC technologies are still core to our business, Xilinx is not just an FPGA company anymore," Peng said. With Everest, which will be first manufactured with 7-nanometer technology and released next year, Xilinx is trying to build momentum with cloud giants offering its chips like Amazon, Baidu, TenCent, and Alibaba.

Xilinx is also trying to become more than a peripheral player in machine learning chips. The benefit of programmable logic is that it can be customized for specific algorithms and altered as algorithms change. It is the opposite of an ASIC, which trades flexibility for performance. Xilinx is targeting the inference side of machine learning, not training, where Nvidia GPUs excel.

The race to artificial intelligence, and the slowing of the Moore's Law, the chip industry's economic engine for decades, has made the development of chips a high stakes game. Nvidia is spending billions into to hold onto its dominant lead in machine learning chips. Other companies, including startups as well as Google and Microsoft, are either trying to muscle into the business or wean themselves off Nvidia.

And yet no one knows what the future of computing will look like and what combination of processors will ultimately be installed in things like wireless sensors monitoring a factory floor, gateways that pool information from them, and the data centers where it is all analyzed. How important will CPU, GPU, FPGA, DSP, ASIC, or ACAP accelerators be? It's still early.

Traditionally, the FPGA has been a bulky and expensive slab of hardware suited for aerospace, industrial, medical, test and measurement, and other markets. These chips are divided into separate routing and logic blocks arranged in a checkerboard pattern, and programmers tie together the individual cells based on the workload.

Xilinx said that ACAP devices lower the bar for that type of programming. "We want to increase the number of users on our products by an order of magnitude," Peng said in a recent conference call with reporters. "There are easily a thousand times more software developers than there are FPGA developers."

Everest, and any other accelerator based on ACAP technology, will include a next generation FPGA and a software programmable, yet hardware adaptable, compute engine based on a new architecture that the company has declined to discuss in more detail. The parts will be connected to highly integrated programmable I/O blocks with a network-on-a-chip system.

(Continued on page 28)

Fast Pulse Test Systems from Avtech

Avtech offers over 500 standard models of high-speed pulsers, drivers, and amplifiers ideal for both R&D and automated factory-floor testing. Some of our standard models include:

AVR-E3-B: 500 ps rise time, 100 Volt pulser

AVRQ-5-B: Optocoupler CMTI tests, >120 kV/us

AVO-8D3-B: 500 Amp, 50 Volt pulser

AV-1010-B: General purpose 100V, 1 MHz pulser

AVO-9A-B: 200 ps t_r, 200 mA laser diode driver

AV-156F-B: 10 Amp current pulser for airbag initiator tests

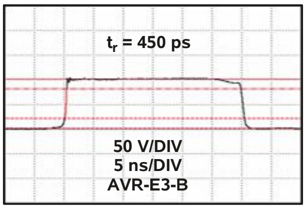

Nanosecond Electronics
Since 1975

Pricing, manuals, datasheets and test results at:
www.avtechpulse.com

AVTECH ELECTROSYSTEMS LTD.
Tel: +1-613-686-6675 Fax: +1-613-686-6679
info@avtechpulse.com

High Output - Low Risetime Pulsers from Avtech

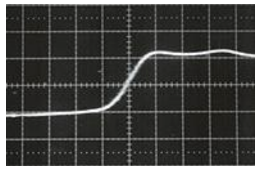

Sub-nanosecond rise time pulsers for every amplitude range!

Model AVP-3SA-C provides up to 10V with < 50 ps rise times.
← Typical waveform, 50 ps/div, 5V/div.

At the other extreme, Model AVI-V-HV2A-B provides up to 100V with < 300 ps rise times.

Ampl	t_{RISE}	Max. PRF	Model
100 V	500 ps	0.1 MHz	AVR-E3-B
100 V	300 ps	0.02 MHz	AVI-V-HV2A-B
50 V	500 ps	1 MHz	AVR-E5-B
20 V	200 ps	10 MHz	AVMR-2D-B
15 V	100 ps	25 MHz	AVM-2-C
15 V	150 ps	200 MHz	AVN-3-C
10 V	100 ps	1 MHz	AVP-AV-1-B
10 V	50 ps	1 MHz	AVP-3SA-C
5 V	40 ps	1 MHz	AVP-2SA-C

Nanosecond Electronics
Since 1975

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

AVTECH ELECTROSYSTEMS LTD.
Tel: +1-613-686-6675 Fax: +1-613-686-6679
info@avtechpulse.com

Pulsed Laser Diode Drivers from Avtech

Each of the 19 models in the Avtech **AVO-9 Series** of pulsed laser diode drivers includes a replaceable output module with an ultra-high-speed socket, suitable for use with sub-nanosecond pulses.

Models with maximum pulse currents of 0.1A to 10A are available, with pulse widths from 400 ps to 1 us.

GPIB, RS-232, and Ethernet control available.

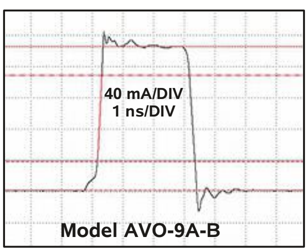

Nanosecond Electronics
Since 1975

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

AVTECH ELECTROSYSTEMS LTD.
Tel: +1-613-686-6675 Fax: +1-613-686-6679
info@avtechpulse.com

100 to 1000 Volt Lab Pulsers

Avtech offers a full line of 100, 200, 500, 700 and 1000 Volt user-friendly pulsers capable of driving impedances of 50 Ω and higher. The **AVR Series** is suitable for semiconductor and laser diode characterization, time-of-flight applications, attenuator testing, and other applications requiring 10, 20, or 50 ns rise times, pulse widths from 100 ns to 100 us, and PRFs up to 100 kHz. GPIB & RS-232 ports are standard, VXI Ethernet is optional.

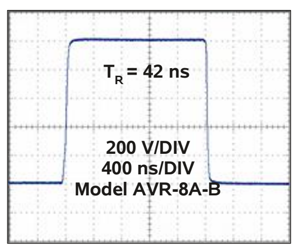

See additional test results at:
http://www.avtechpulse.com/medium

Nanosecond Electronics
Since 1975

Pricing, manuals, datasheets and test results at:
www.avtechpulse.com/medium

AVTECH ELECTROSYSTEMS LTD.
Tel: +1-613-686-6675 Fax: +1-613-686-6679
info@avtechpulse.com

WILLIAM WONG | Senior Technology Editor

Automakers: Manage Connected-Car Data Before It Manages You

Technology Editor Bill Wong talks with Dan Gittleman, CEO of Xevo, about how manufacturers must address the influx of data in the connected car through proper management techniques.

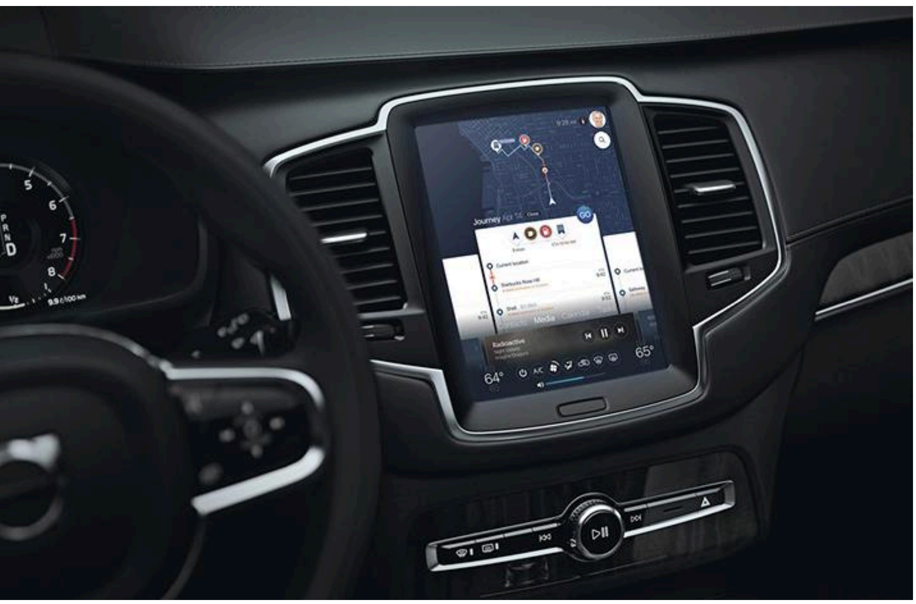

The Xevo Carware Intelligent Edge (XE) data-centric platform.

Cars now generate an enormous amount of data every day. As vehicles become more complex, the amount of information is expected to severely strain data networks and cloud infrastructure, creating a serious need for a new network and computing infrastructure solution.

Xevo's expertise is in the data-driven automotive user experience. I talked with Dan Gittleman, CEO of Xevo, who shared his take on the solutions that automakers can incorporate into their vehicles to keep everything running normally, even with the ever-increasing amount of connected-car data.

There's a lot of buzz these days around the increase in data from connected cars. How much data are we talking about?

Estimates say that the amount of data transferred between connected cars and the cloud will experience a 10,000-fold increase, likely reaching 10 exabytes per month by 2025.

Is there a way to help the network process all of this new data?

Automotive OEMs and automobile designers can ease the impact of this huge data increase by incorporating technology that combines cloud services and in-car computing power, so that even as the amount of data increases, devices and vehicle software will continue to function and communicate normally.

That sounds great—how can designers make that happen?

The first step is to implement a data-centric technology solution, like our Journeyware platform. Journeyware combines cloud-based machine learning, artificial intelligence, and powerful data-analysis tools with in-vehicle data processing tools to acquire, manage, analyze, and distribute the massive, diverse data flows generated by connected cars. The solution leverages these advanced technologies, integrating in-car computing with cloud services, to optimize networking and computing between vehicle head units, mobile devices, and the cloud.

What is distributed processing?

Distributed processing manages and categorizes car-generated data by "triaging" it. In other words, it determines which information can be processed in the car, and what needs to be addressed in the cloud. Any data that can be processed in the car stays "local"—it's not unnecessarily sent to the cloud, which helps to conserve network resources. Managing the data like this also reduces network congestion by decreasing the overall amount of data being transferred, and improves cloud platform efficiency by removing unnecessary data processing.

What happens to the data once it's been triaged?

That's where Journeyware's customizable "rules" about data collection and storage come in. These policies ensure that the highest priority data always receives the highest-priority treatment, by enabling the software to determine which data to collect, what to store in the car, and which information should be sent to the cloud. It's like collecting video data only in specific locations under certain conditions or on an intermittent basis, and storing some data permanently in the car.

For example, an engineer could set up the system so that data is only stored permanently in the car or uploaded to the cloud if certain events happen, such as cars in a certain area that slow down suddenly, or a cluster of vehicles whose windshield wipers all go on at the same time. The policies create an effective framework for managing the information, and reduce the overall amount of data that requires processing, which helps prevent the network from getting overworked.

Storage seems like it could be a mammoth task. How does your software solution address it?

The system tackles the storage challenge using a few different rules. It determines whether the information should be saved temporarily or permanently, and if it will be stored in the car or sent to the cloud. For information that goes to the cloud, it decides which data to store for quick analysis and what needs to be stored, but likely won't be accessed very often (known as "inactive" data in "cold storage"), then specifies how data will be accessed when it is needed.

Using a combination of policies and criteria to manage and sort the data as it's generated, and distributing the data according to temporary-vs.-permanent and local-vs.-cloud storage, allows automotive designers to proactively address the rapidly growing quantity of information from connected cars. This helps to ensure that devices and software continue to function correctly.

Can cell phones and mobile devices strain the network?

Yes. Consumer devices—phones and tablets, and even the apps that run on them—can affect the network. People expect the same functionality from their devices when they're in the car as they do in all other areas of their lives. A hardware and software solution that supports Bluetooth pairing and in-car Wi-Fi network hosting gives devices more than one way to connect to the network and receive information. This means resources can be distributed and conserved effectively, and that all of our phones and "toys" will work normally while on the go.

Overall, the increase in connected-car data warrants automakers' attention. Incorporating a solution that employs smart data collection and retention, with a focus on policies and criteria and in-vehicle management of consumer devices, addresses many of the potential bandwidth issues created by the influx of data. By designing their vehicles to include such a solution, OEMs can increase consumer satisfaction and brand loyalty by delivering what consumers want from their vehicles. ed

DAN GITTLEMAN is the CEO at Xevo Inc., a global leader in data-driven user experiences, with the world's largest IoT deployment in the automotive industry. He previously founded OpenPeak as well as RAID Power Services Inc. (which was rebranded in 2000 as StorageApps).

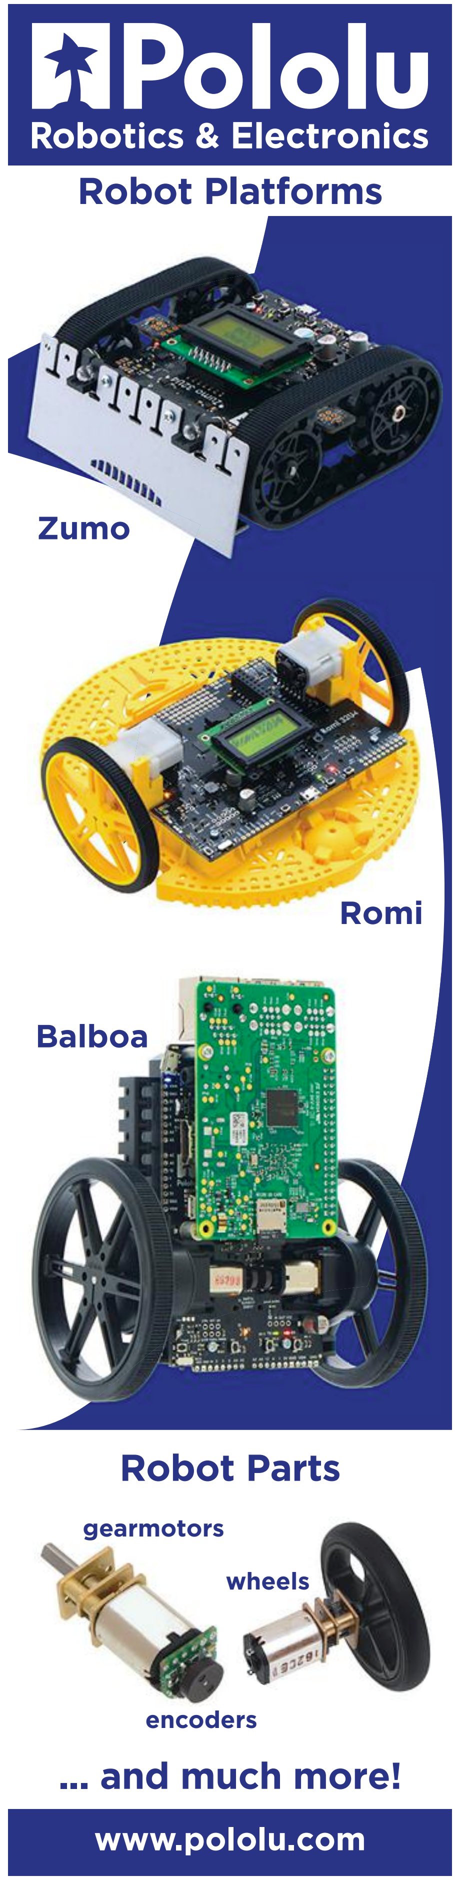

STEVEN LEE | Application Engineer, Keysight Technologies Inc.

What Can You Expect from the New Generation of POWER SUPPLIES?

Today's power supplies do more than deliver clean and stable dc power on a daily basis—they provide advanced capabilities that can save you time and money.

1. Shown are Keysight's E36300 Series triple-output dc power supplies.

Because today's designs place higher demands on the systems that power them, newer-generation bench power supplies are packed with very useful features that help engineers address those problems much earlier in the design process. This includes controlling multiple power-supply voltage sequences, measuring wide dynamic ranges of current, and varying the speed of the power-supply voltage to reflect real circuit characteristics.

Underlying all of this, bench power supplies need to be easy to use, deliver clean and stable dc power, and protect themselves and the development board (or DUT). It's not uncommon for bench power supplies to also be used in automated systems, so modern I/O connectivity might be another important consideration.

This article delves into problems that the new generation of power supplies *(Fig. 1)* can uniquely solve, and factors to consider when specifying or purchasing a power supply.

DESIGN TOPOLOGY: LINEAR VS. SWITCHER

A fundamental requirement of a power supply is an output signal with a high level of signal integrity, especially when working on a delicate design. The design topology of a power supply can tell you a lot about its signal integrity.

Generally, dc power supplies are developed with either linear or switch-mode technology, each offering different advantages. Linear power supplies tend to have low output noise, fast transient response, and high programming speed. However, they also come with a series of disadvantages, such as low

No place to replace a battery.

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

Highly remote locations call for Tadiran batteries.

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

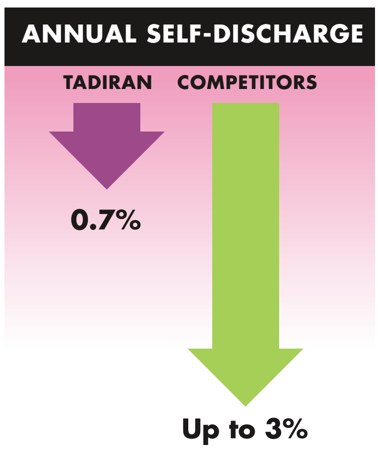

Take no chances. Take Tadiran batteries that last a lifetime.

TADIRAN BATTERIES

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

www.tadiranbat.com

* Tadiran 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.

More and more applications require low current measurements such as idle/standby current on devices while in a low-current consumption state. These measurements can only be made with a power supply that accurately measures low-current values. Newer power supplies can measure dc current in the 100-µA range. However, be sure to check the datasheet for accuracy of measurement specifications at that range.

efficiency, more cooling, higher level of low-frequency magnetic radiation, and (usually) larger-sized.

Switch-mode power supplies are typically smaller in size, provide higher efficiency, and require less cooling. Some of the disadvantages include slower transient response and higher output noise. However, design advances are enabling the newer generation of switch-mode power supplies to perform just as well as linear power supplies in those areas.

Newer bench power supplies provide the best of both topologies. To achieve higher power in a smaller package, Keysight combined a linear output stage with a phase-controlled pre-regulator. Some of the benefits of this approach are better efficiency, double the output power, and lower acoustic fan noise *(Fig. 2)*.

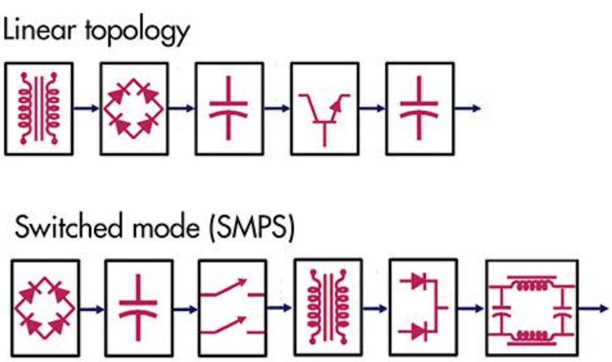

2. This simplified block diagram represents the design topology between linear and switching power supplies.

HIGH- AND LOW-CURRENT MEASUREMENTS

More and more applications require low current measurements such as idle/standby current on devices while in a low-current consumption state. These measurements can only be made with a power supply that accurately measures low-current values. Newer power supplies can measure dc current in the 100-µA range. However, be sure to check the datasheet for accuracy of measurement specifications at that range.

There may also be situations when more current is needed than can be provided by an older dc bench power supply. Newer bench supplies come with a built-in auto-parallel-mode capability that allows the engineer to combine two channels into a single, higher-current dc output channel rated as high as 25 V and 4 A. This eliminates the need to find and add another power supply into the test design. In a way, new bench power supplies with auto-parallel capability perform the work of two power supplies.

ADVANCED CAPABILITIES

Modern bench power designs fully take advantage of the latest graphical displays, microprocessors, and FPGAs. Thus, many advanced capabilities are built-into newer power supplies to save you time and money.

One example is the ability to synchronize multiple outputs on a power supply. Without this feature, a design engineer could spend hours writing code to perform this synchronization. Another example is the ability to create a dynamic user-defined output signal on a power supply to eliminate the need of a second setup requiring additional programming and instrumentation *(Fig. 3)*. With these and other advanced features, you will use less test equipment, set up tests faster and more simply, and reduce test setup errors—all of which save time and money.

3. In Data Logger view, you can log data on multiple traces. Here, the voltage of output 1, output 2, and output 3 are captured over 30 seconds.

PROTECTION AND SAFETY FEATURES

When dealing with power, protection and safety come first. When a device fails, or goes up in smoke, it may be catastrophic. It's not only important for the power supply to protect itself, but also protect the device under test (DUT) and the operator. When a DUT fault occurs, protection circuits in the power supply can limit the voltage, limit the current, and shut off all outputs with a single button push.

For example, the E36300 Series from Keysight integrates overvoltage, overcurrent, and over-temperature protection to prevent damage to the DUT. One push of a button on the front panel can shut down all power-supply outputs.

ENHANCED USER EXPERIENCE

An intuitive user interface that allows you to easily navigate, control, and track each channel is another important feature with the latest power supplies. Equally important is a large graphical display that lets you see all of the output channels at the same time. With newer power supplies, you can set up advanced capabilities like remote sense, auto-parallel/auto-series, tracking mode, voltage sequencing, voltage steps, and data logging of V/I pairs over time with one push of a button.

In addition, modern I/O control such as built-in USB and LAN is typically available in modern power supplies. Support software can simplify automated measurement setup, recording, and analysis.

It's convenient to be able to see all channel readings and settings on the same screen *(Fig. 4)*. This may not be obvious, since many of the older power supplies make the user toggle back and forth to compare the channel setting and the reading.

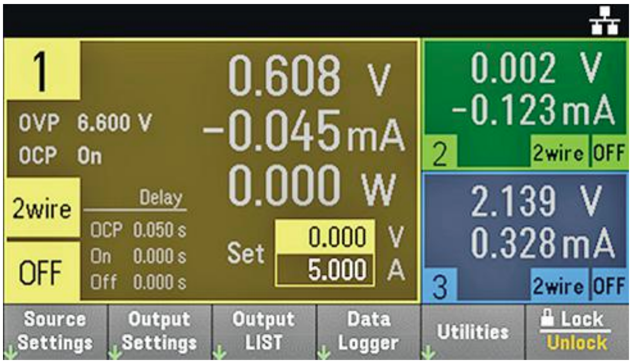

4. The user can view details of a single channel, including the measured power, OVP/OCP condition, and delays.

Solid State Relays and Contactors

CARLO GAVAZZI

Chassis mount solid state relays (SSRs) available in one, two or three pole switching designs. Single phase types up to 125 amps, two pole types up to 75 amps and three phase types up to 75 amps. Also offered in our new slim-line design (up to 90 amps) and compact fast-on type (up to 25 amps).

DIN rail or chassis mount solid state contactors and SSRs which are UL508 rated for motor loads, and feature integrated heat sinks, fans and large load terminals. Designed for switching single phase loads up to 85 amps (15 Hp) and two or three pole types for switching up to 75 amps per phase (25 Hp).

Specialty SSRs for your growing demands: System Monitoring SSR for line /load voltage and load current, fused SSRs provide more protection, and 1, 2 and 3 pole proportional controllers with several switching modes: phase angle, distributed burst (1, 4 or 16 cycle), or soft start.

Carlo Gavazzi is one of the fastest growing SSR manufacturers worldwide!
Contact us if you're interested in a free evaluation sample (qualified OEMs only).

GavazziOnline.com • 847.465.6100 • Info@CarloGavazzi.com

Unfortunately, no standards have been established for power-supply datasheet specifications. Descriptions of the specification, as well as the specifications themselves, are determined by the manufacturer and will vary from one manufacturer to another. It's also at the manufacturer's discretion whether a characteristic is even specified. Some vendors specify more than others.

Newer power supplies make this information available on the same screen.

PERFORMANCE CHARACTERISTICS

Unfortunately, no standards have been established for power-supply datasheet specifications. Descriptions of the specification, as well as the specifications themselves, are determined by the manufacturer and will vary from one manufacturer to another. It's also at the manufacturer's discretion whether a characteristic is even specified. Some vendors specify more than others. Texas Instruments' experts wrote an application note, titled How to Read Your Power Supply's Data Sheet, to help you better understand power-supply datasheets and evaluate which power supply best meets your needs.

TOTAL COST OF OWNERSHIP

The cost of an instrument includes the initial purchase price, as well as ongoing maintenance expenses plus the opportunity cost of downtime for calibration and repairs. Look for a manufacturer that will provide worldwide maintenance and support, a dedicated support team, and a solid track record for delivering products that last for 20 years or more. ed

FOR MORE information, visit www.keysight.com/find/E36300, or check out the application note "Seven Ways to Speed Up Your Testing with a Modern DC Bench Power Supply."

High Voltage High Frequency Power MOSFETs and Drivers

Optimized for high speed and high power applications

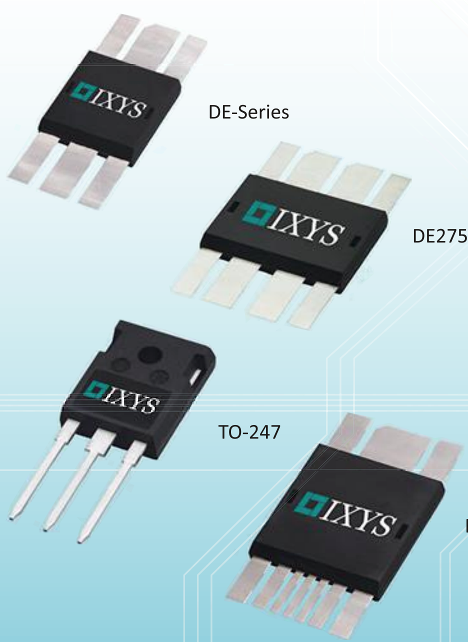

Power MOSFETs:

High Power 100V to 1200V devices
Low-Inductance DE-Series and industry-standard package styles

Optimized for:

- ISM band RF generators and amplifiers
- High voltage pulse generators and pulsed laser diode drivers
- Switched-mode and resonant-mode HF power supplies

Gate Drivers:

Class D and E HF RF and other high speed switch applications
Up to 30A peak current, ultra-fast rise times

RF Power Modules:

Integrated gate driver and MOSFET in a low-inductance package
500V and 1000V configurations
Designed for Class D, E, HF, and RF applications at up to 27 MHz

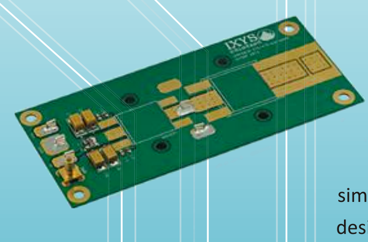

IXYS COLORADO

1609 Oakridge Drive, Suite 100
Fort Collins, Colorado 80525 USA
TEL +1-970-493-1901 sales@ixyscolorado.com

Evaluation boards to simplify your application design and development

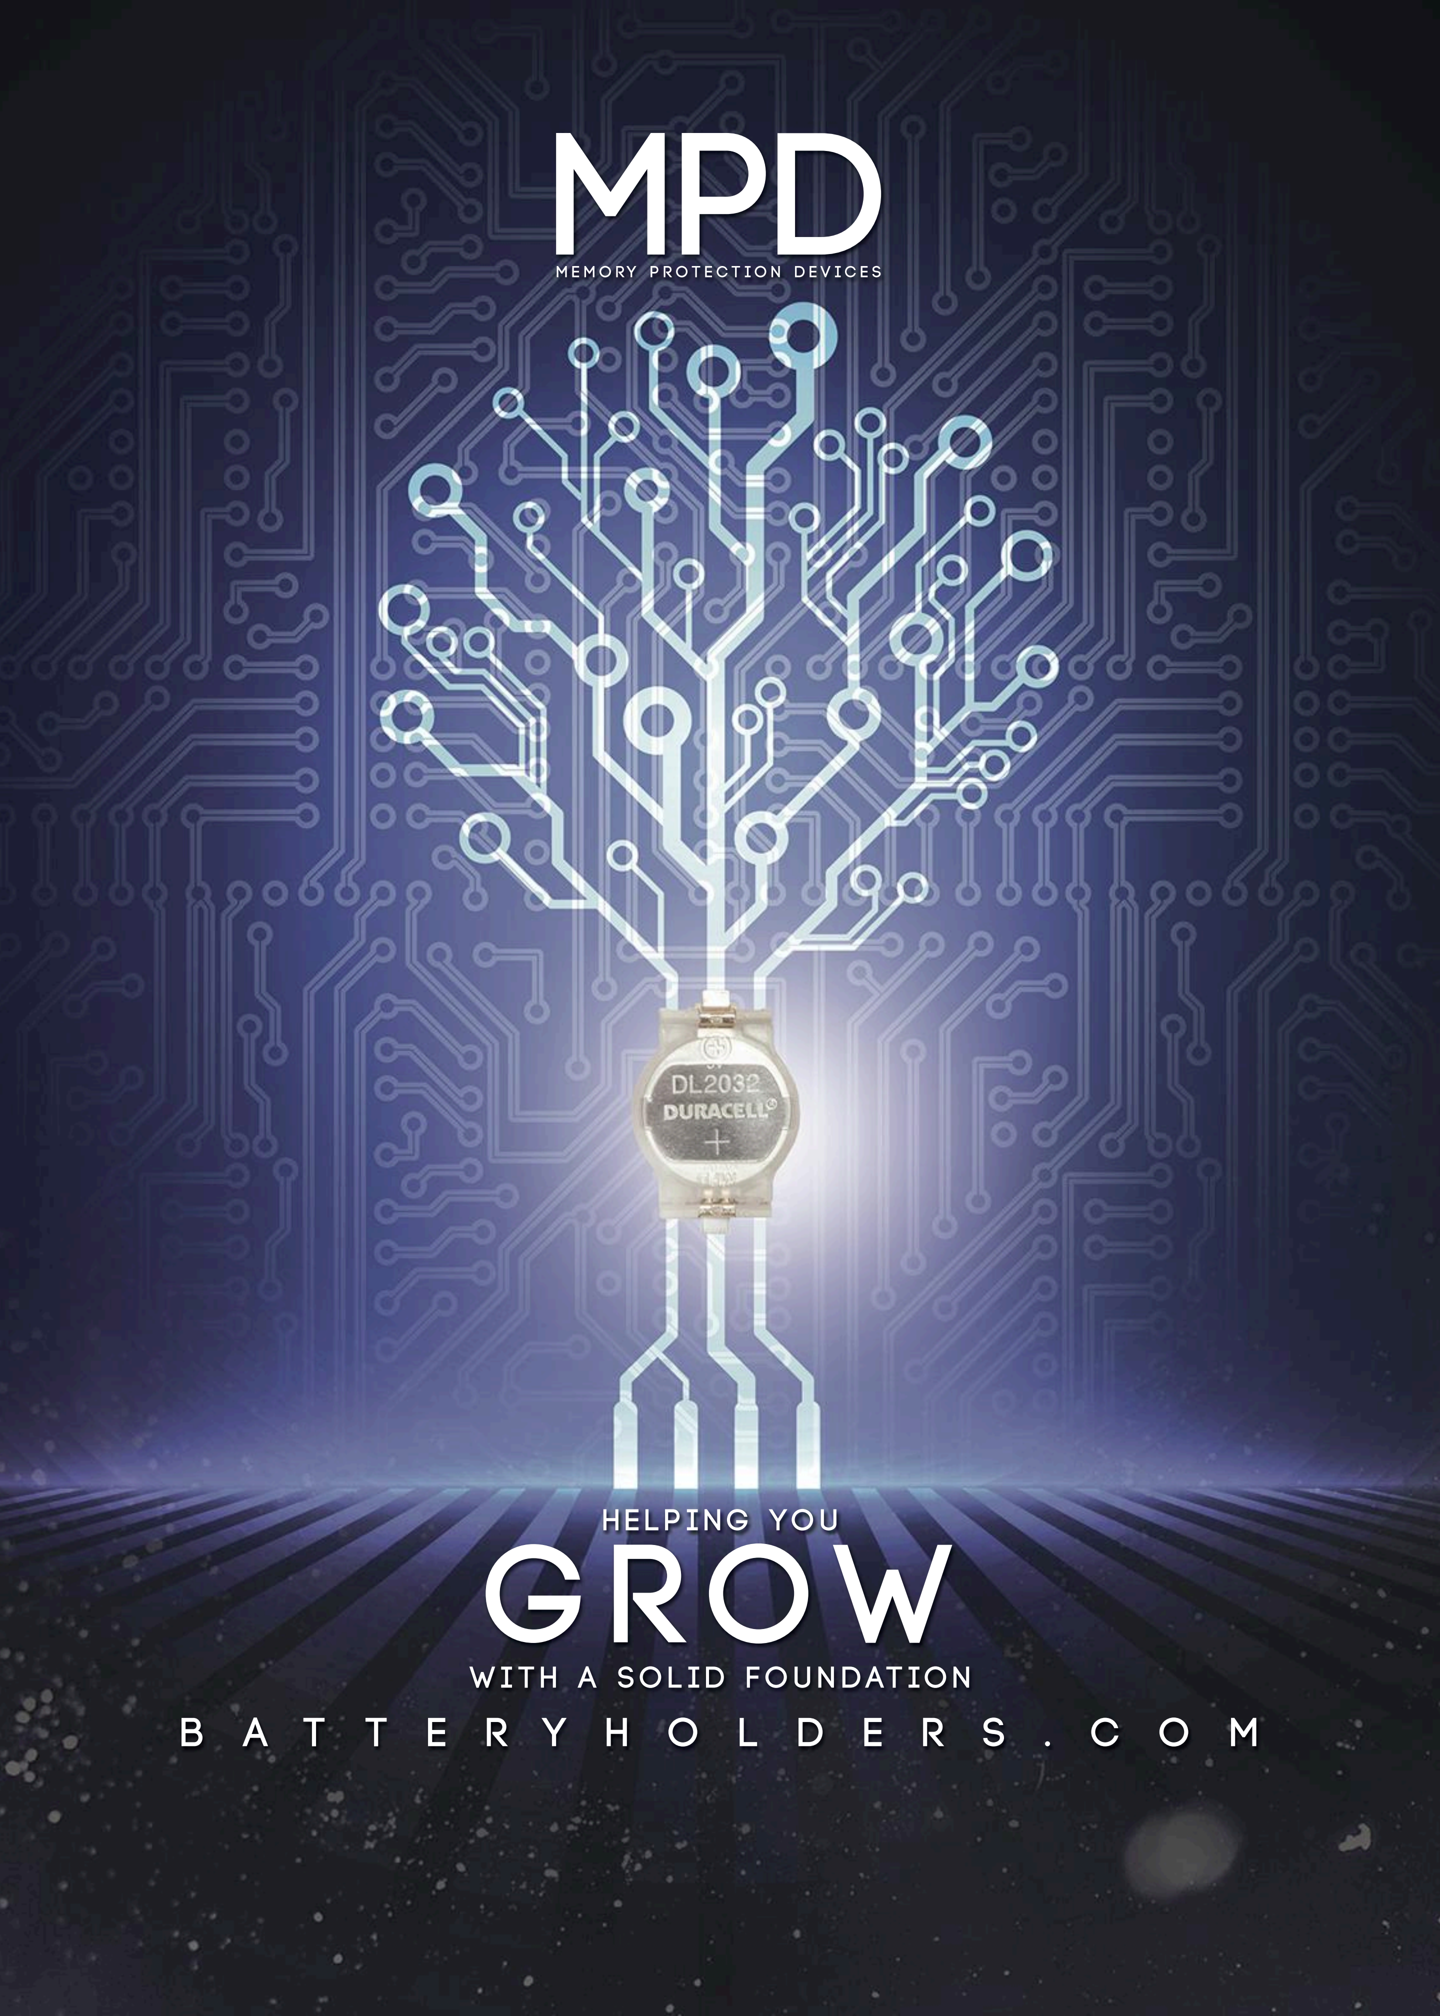
MPD
MEMORY PROTECTION DEVICES
DL2032
DURACELL
HELPING YOU
GROW
WITH A SOLID FOUNDATION
BATTERYHOLDERS.COM

WILLIAM WONG | Senior Technology Editor

M.2, NVMe Bring Performance and Capacity to Modules and Motherboards

M.2 and NVMe are starting to show up in embedded modules and motherboards, allowing them to be more compact while delivering high performance storage.

The M.2 memory and peripheral expansion form factor is very popular in PC motherboards and mobile devices like laptops where compact storage is a requirement. It is also starting to show up in embedded motherboards and modules like Diamond Systems' Zeta COM Express Type 10 SBC family *(Fig. 1)*. The Zeta 84- by 55-mm stack actually has a heat spreader on the bottom with the COM Express module in the middle. The white connector on the top carrier board provides a stackable expansion, such as Diamond Systems' expansion board with the M.2 and Mini-PCIe sockets.

The Zeta COM Express supports up to 8 GB of RAM and a 1.1 GHz, quad core, Intel Apollo Lake N4200 processor. The fanless design has a –40 to +85°C operating temperature range. The main carrier board includes a 16-channel, 16-bit ADC and a four-channel, 16-bit DAC in addition to connectors for dual gigabit Ethernet, GPIO, four serial ports, an SD card socket, and a Mini-PCIe socket.

Mini-PCIe has become a popular peripheral expansion option for embedded systems. Mini-PCIe cards are available with everything from ADCs to motor control. Of course, there are plenty of wireless options as well. The M.2 form factor can also be used for peripherals, but it is primarily being used for storage. The M.2 B Key provides support for x2 PCI Express (PCIe), SATA, USB 2.0 and 3.0, audio, UIM, HSIC, SSIC, I^2C,

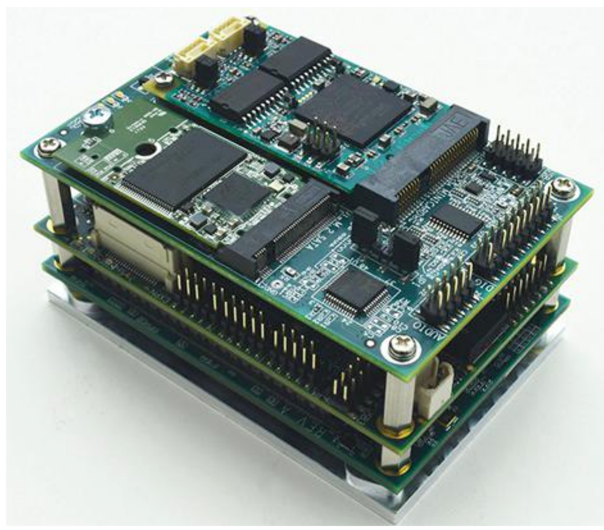

1. Diamond Systems' Zeta COM Express SBC family can work with carrier boards that include Mini-PCIe and M.2.

2. Supermicro's 5019D-FN8TP system has a motherboard with an M.2 B Key socket and Mini-PCIe socket that includes mSATA support.

and SMBus. The M.2 M Key provides x4 PCIe, SATA, and SMBus. NVMe can work with any PCIe interface.

Super Micro Computer's 5019D-FN8TP *(Fig. 2),* which is built around the X11SDV-8C-TP8F motherboard, runs an 80W, 8-core/16-thread, Intel Xeon-D D-2146NT processor with up to 512 GB of EDD LRDIMM memory. It targets applications like network appliances, digital signage, and IoT servers. It can handle a pair of 3.5-in. drives or four 2.5-in. drives.

For expansion, the motherboard has a M.2 M Key socket for storage, an M.2 B Key socket that can be used for storage or wireless network cards, and a Mini-PCIe socket with mSATA support. There is also a x8 PCIe slot. The system has four 1-Gigabit Ethernet ports, two 10GBase-T ports, and two 10G SFP+ ports. It has a dozen SATA 3 ports, too.

M.2 may eventually be the form factor of choice for peripheral expansion, but for now it tends to be limited to storage and some wireless interfaces. Mini-PCIe is the form factor where there are many options for interface cards. Most Mini-PCIe solid-state drives are mSATA devices that use a SATA interface. In

a.

b.

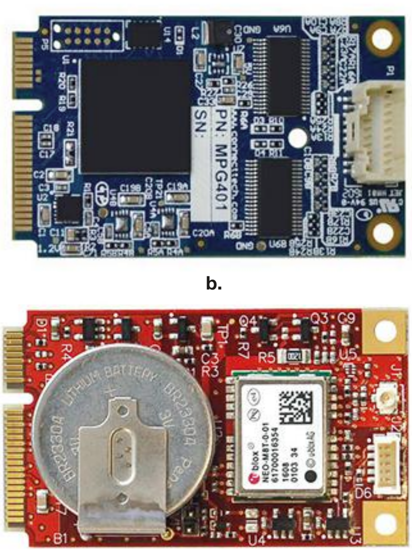

3. Some examples of rugged Mini-PCIe adapters including Connect Tech's 16-channel, 16-bit ADC (a) and VersaLogic's Advanced GPS controller (b).

One challenge for embedded developers is finding devices that meet their ruggedness requirements. M.2 devices and even many Mini-PCIe devices are designed for PCs or other applications that are not as demanding.

theory, an NVMe Mini-PCIe card could be created, but most NVMe cards are implemented as M.2 devices. The x1 or x4 PCIe interfaces are the most common.

One challenge for embedded developers is finding devices that meet their ruggedness requirements. M.2 devices and even many Mini-PCIe devices are designed for PCs or other applications that are not as demanding. Luckily, most Mini-PCIe peripheral cards are designed with embedded applications in mind *(Fig. 3)*. Connect Tech's 16-channel, 16-bit ADC operates at 500 ksamples/s. VersaLogic's Advanced GPS controller has an on-board battery to maintain satellite position data to support fast restart of the GPS chip. Both support an extended industrial temperature range of –40°C to +85°C.

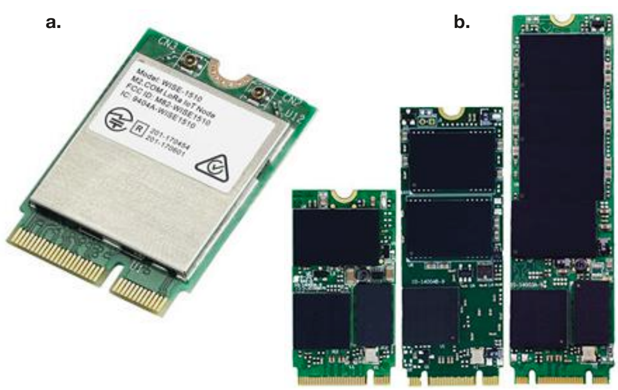

4. Examples of M.2 devices include Advantech's LoRA wireless adapter (a) and Virtium's family of rugged M.2 storage (b).

Likewise, more M.2 devices that meet embedded temperature and ruggedness are becoming available *(Fig 4)*. Advantech supports LoRAWAN, a long-range, low-speed wireless protocol using its M.2 adapter. Virtium's StorFly family of M.2 flash cards highlight rugged memory devices that are available. This includes SLC, iMLC, and MLC options as embedded applications often have different requirements for endurance and capacity. These cards are also designed for industrial temperature ranges.

From
Concept to Reality

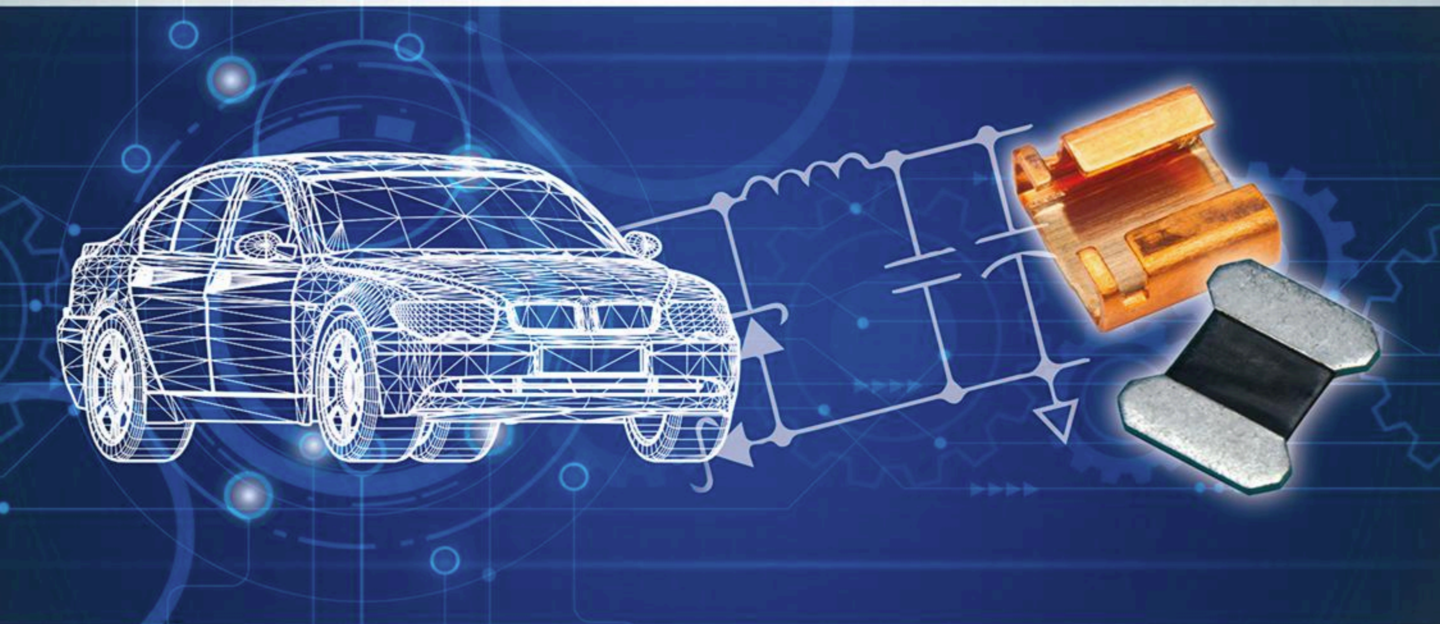

Our new passive component solutions will help you make the leap!

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

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

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

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

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

NEW! From KOA Speer

Precision Thick Film Resistors (RK73G)

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

Wide Terminal Current Sense Resistors (WK73S & WU73)

KOA®
KOA SPEER ELECTRONICS, INC.

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

TIEN SHIAH | Memory Product Marketing, Samsung Semiconductor Inc.

High Bandwidth Memory: THE GREAT AWAKENING *of* AI

We're entering a historic new era in artificial intelligence, driven by advances in deep learning, and HBM is critical to making it all happen.

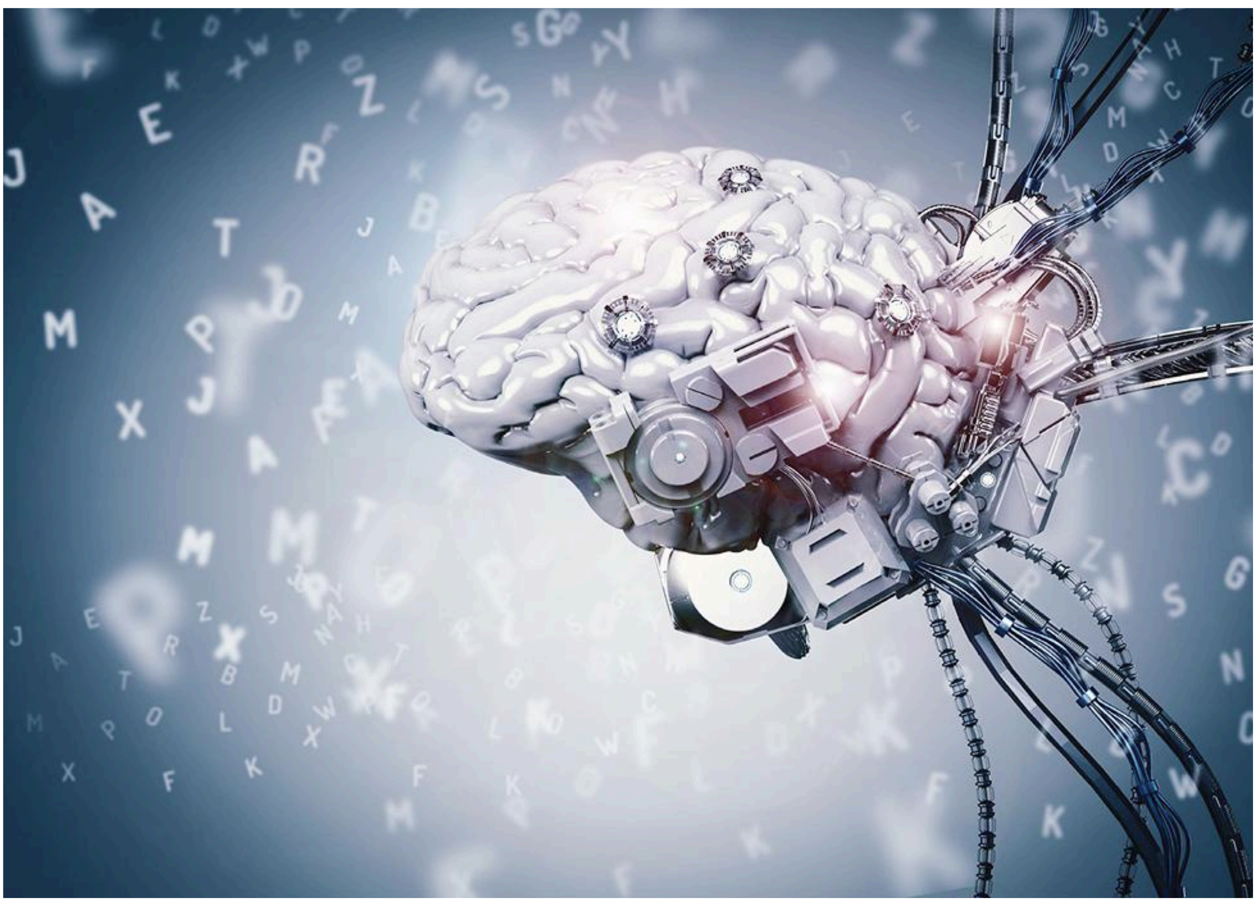

Artificial intelligence (AI) is fast becoming one of the most important areas of digital expansion in history. The CEO of Applied Materials recently stated that “the war” for AI leadership will be the “biggest battle of our lifetime.”[1] AI promises to transform almost every industry, including healthcare (diagnosis, treatments), automotive (autonomous driving), manufacturing (robot assembly), and retail (purchasing assistance).

Although the field of AI has been around since the 1950s, it was not until very recently that *computing power* and the *methods used in AI* have reached a tipping point for major disruption and rapid advancement. Both of these areas have a tremendous need for much higher memory bandwidth.

COMPUTING POWER

Those familiar with computing hardware understand that CPUs have become faster by increasing clock speeds with each product generation. However, as material physics will tell you, we’ve reached a limit as to how far we can go to improve performance by increasing clock speeds. Instead, the industry has shifted toward multicore architectures—processors operating in parallel.

Modern data-center CPUs now feature up to 32 cores. GPUs, traditionally used for graphics applications, were designed from the start with a parallel architecture; these can have over 5,000 simpler (more specialized) cores. GPUs have become major players in AI applications primarily because they’re very fast at matrix multiplication, a key element of machine learning. The shift toward more cores underscores the growing demand for data to be fed into these cores to keep them active—driving the need for much greater memory bandwidth.

METHODS USED IN AI

The fastest-advancing area within AI has been machine learning, and more specifically, deep learning. Deep learning is a subset of machine-learning methods based on learning data representations, as opposed to task-specific algorithms. In the past, AI used hand-coded “rules” to make decisions. This had very limited success.

A breakthrough came five years ago when deep-learning techniques

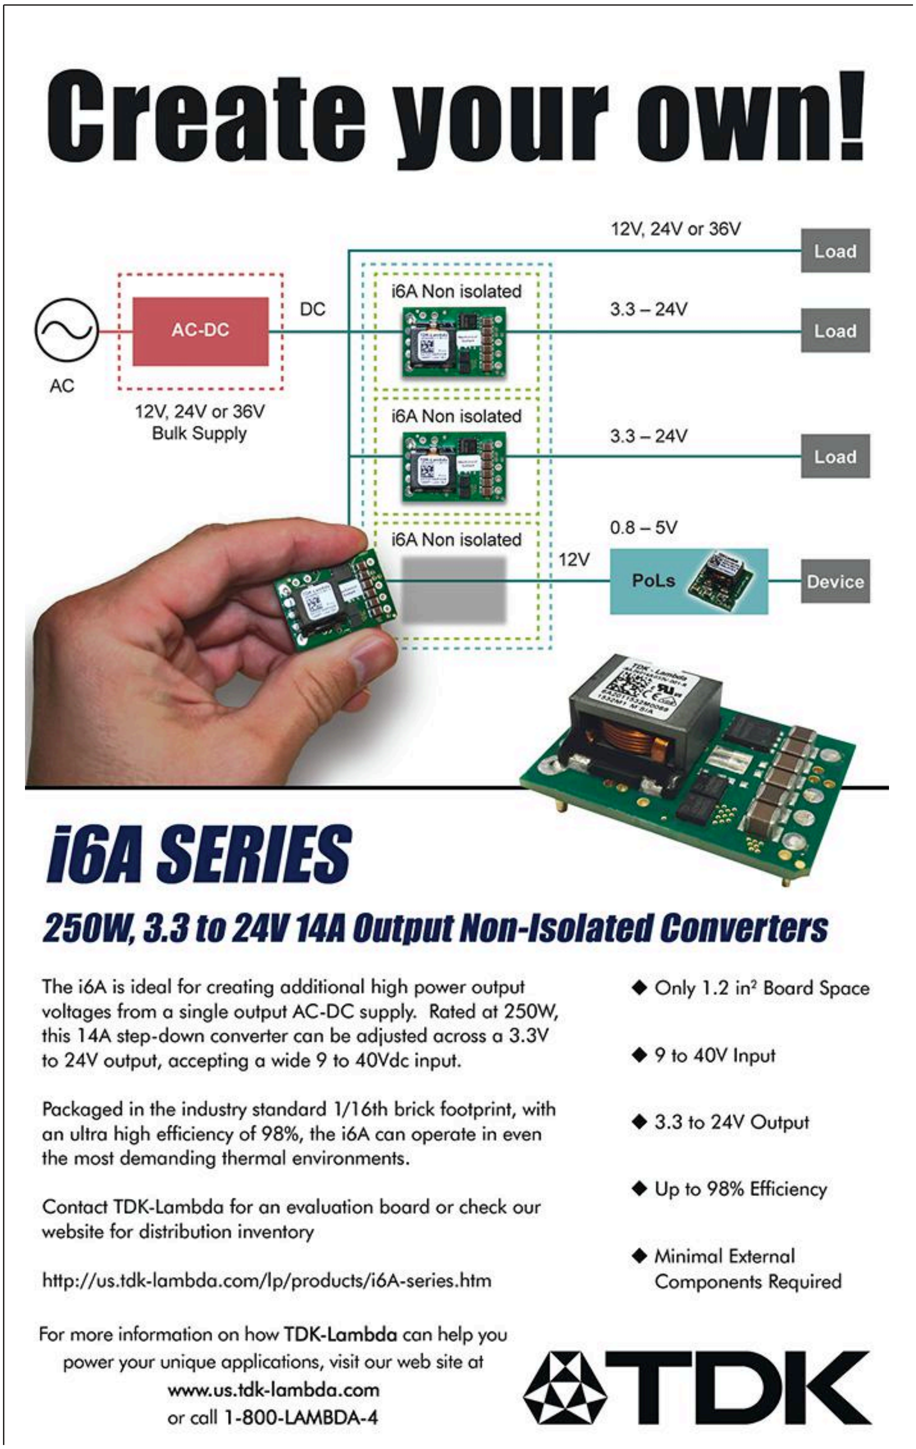

(using deep neural networks, or DNN) were implemented. Google Translate, for one, became much more accurate when Google deployed a DNN platform last year, yielding more improvements than in the previously 10 years combined. Moreover, as shown in *Figure 1,* image recognition improved dramatically when DNNs have been applied versus previous hand-coded techniques.

DNN involves training machines with vast amounts of data. Once trained, an AI system is able to make better predictions based on new input. This training process involves millions of matrix multiplications, a task that's more suited for parallel, multicore compute architectures. This is why much greater memory bandwidth is needed for AI.

THE NEED FOR HIGH BANDWIDTH MEMORY

The main memory used in today's computers is DRAM (dynamic random access memory), which has a top system memory bandwidth of 136 GB/s. GPUs, by their multicore nature, utilize a special form of DRAM called GDDR (graphics double-data-rate DRAM) that can deliver system bandwidth at up to 384 GB/s. But the data-hungry, multicore nature of the processing units needed for machine learning requires even greater memory bandwidth to feed the processing cores with data. This is where high bandwidth memory (HBM) is beginning to make a critical contribution that will grow by orders of magnitude over the next several years.

HBM is a specialized form of stacked DRAM that's integrated with processing units to increase speed while reducing latency, power, and size. These stacks are known as memory "cubes," each of which can have up to eight DRAM die.

HBM is now an industry standard via JEDEC (Joint Electron Device Engineering Council) and its current iteration, HBM2, can support 8 GB of memory per cube at 256-GB/s throughput. This will improve system performance and increase energy efficiency, enhancing the overall effectiveness of data-intensive, high-volume applications that depend on machine learning, parallel computing,and graphics rendering.

When four HBM2 cubes per processor are interconnected through use of a silicon interposer, that configuration will provide 32 GB of memory with a bandwidth of 1 TB/s. *Figure 2* provides an example of an HBM structure, as well as how it connects to a processor.

HBM is continuing to rapidly evolve to meet the needs of machine learning. The technology has already gained

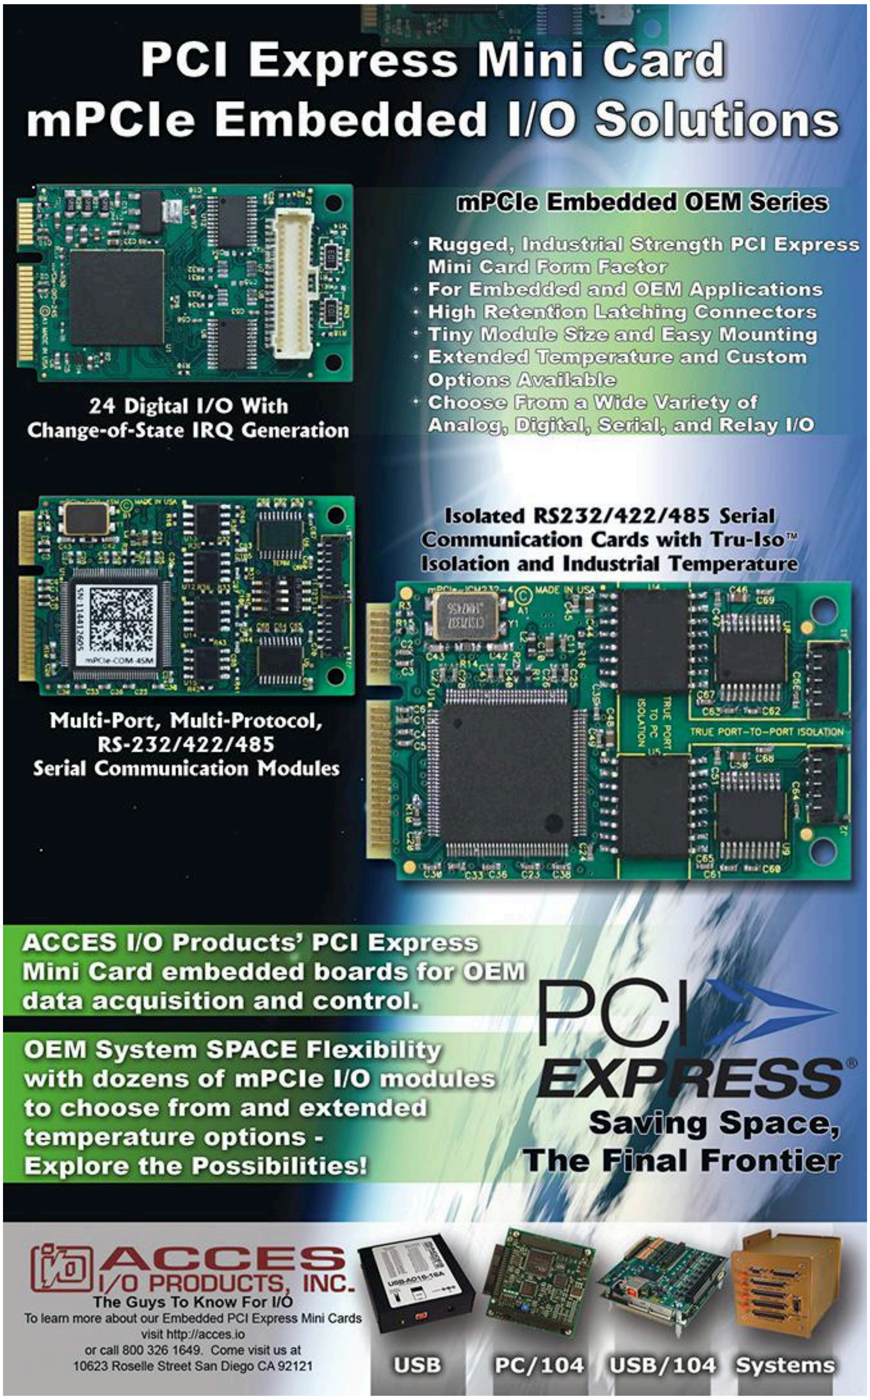

significant traction in the market, going from virtually no revenue a couple of years ago to what will likely be billions of dollars over the next few years. The next iteration, expected to become available in about three years, is now projected to enable a system bandwidth of 512 GB/s per cube. *Figure 3* compares the memory bandwidth that's anticipated with HBM technology to that of GDDR, the next fastest memory.

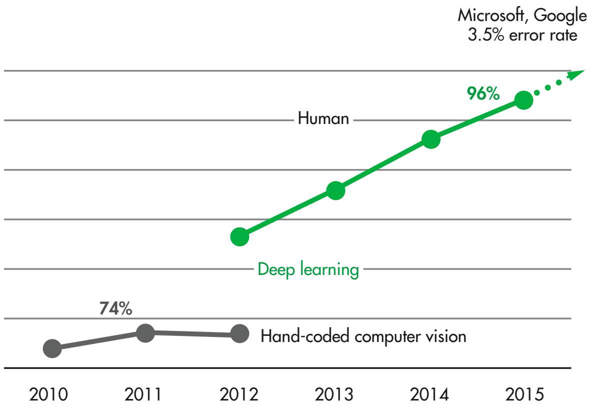

1. Deep neural networks improve the accuracy that's vital in computer-vision systems. *(Source: Nvidia, FMS 2017)*

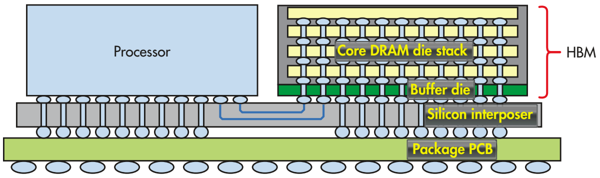

2. Shown is how a high-bandwidth-memory structure connects to processor. *(Source: Samsung)*

SMARTER, FASTER NETWORKS

Not only is HBM being used to increase the speed at which local data is fed to processors, designers are also looking at ways to leverage it to speed movement of data between computing systems.

HBM2 is being designed into high-end switching ASICs (application-specific integrated circuits) that can operate at terabit speeds, to create a significant edge in routing performance. On top of that, new breeds of data-center ASICs used in servers are leveraging HBM, suggesting routing capability much closer to the optimum AI compute node than traditional networks. Expect smarter, faster networks in the near future that will play a pivotal role in the life-enhancing expansion of machine intelligence.

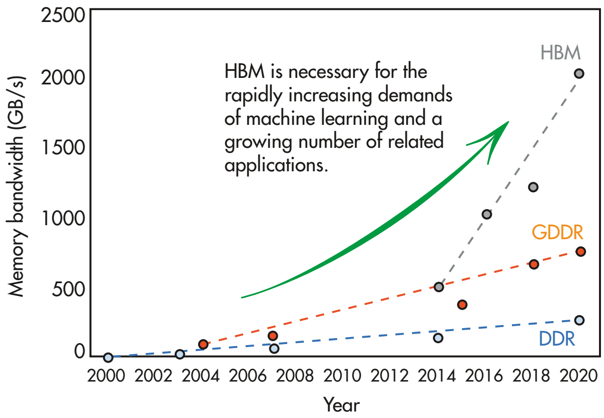

3. The graph tracks memory bandwidth progression over time, and how HBM stacks up against GDDR and DDR in the coming years.
(Source: Samsung)

DOORWAY TO DEEP LEARNING

We are in the "gold rush" era of AI. Recent developments in deep neural networks are enabling major improvements in usable machine intelligence across the globe. Computer vision has already reached a milestone now that machines are able to, on their own, learn how to recognize a cat just by being given lots of images—without needing to know the distinguishing features of a cat.

Computer vision and language represent the tip of the iceberg. Other applications leveraging machine intelligence range from high-frequency trading in the financial sector to tailoring product recommendations in eCommerce.

HBM is an essential technology to drive the AI revolution. Entire industries are about to be transformed, creating huge opportunities now for businesses that stay ahead of the curve. Make no mistake—the technology is complex. Only a couple of today's vendors are able to supply HBM2 memory in volume, as well as possess the expertise to provide future AI/HPC generations with increasing speeds and densities that will meet the growing market demand for deep learning.

However, that's sufficient to move us into the world of AI in grand fashion. When the stakes are this high, leading corporations know to obtain the memory technology that will give their AI-impacted, networked systems the most competitive edge. ed

REFERENCE

1. Interview with Jim Cramer, CNBC, 9/28/17.

(Continued from page 10)

That allows Everest and other ACAP devices to be programmed with more flexibility than FPGAs, the company said. The I/O functionality ranges from high bandwidth memory to advanced SerDes technology, which the company recently pushed to 56 gigabits per second and plans to boost further in the next few years.

The chip can be programmed with software languages including Python and OpenCL, the company said. "Simply dropping an FPGA block into an ASIC design doesn't really bring anything new and different to the game, but Xilinx designed a lot of flexibility into the rest of the ACAP design," said Paul Teich, an analyst with Tirias Research.

Reducing the programming complexity also allows the hardware to be altered faster and conform to different workloads more dynamically. "Although compiling an application to accelerate a workload can take second, minutes or hours, dynamically swapping the combined binaries takes milliseconds," a Xilinx spokesperson told *Electronic Design*.

"We are talking about data centers being able to program their servers to change workloads depending upon compute demands, like video transcoding during the day and then image recognition at night," said Patrick Moorhead, founder of technology research firm Moor Insights & Strategy.

Xilinx said that the Everest line of chips would scale to devices with 50 billion transistors. The 7nm processor provides 20 times better performance on machine learning workloads than 16nm Virtex FPGAs. The company also said that radios based on Everest would have four times the bandwidth of radios based on 16nm technology.

The company has not yet benchmarked the Everest accelerator against GPUs but claimed that it would provide up to a hundred times the performance of CPUs. The company plans to tape out Everest with Taiwan Semiconductor Manufacturing Corporation before the end of the year, with customer shipments scheduled for next year.

Xilinx has been working to tune its products for data centers. Before he was promoted to chief executive, Peng pushed the company to release software libraries that can accelerate image recognition on its hardware. The company also moved to integrate SoCs onto programmable dies and partner on the development of the Cache Coherent Interconnect for Accelerators, or CCIX.

Xilinx said that Everest would be essential to helping it push further into data centers. The ACAP architecture "is likely to give the cloud giants more of what they want from FPGA-based hardware platforms," said Teich. "Xilinx didn't design ACAP in a vacuum – early software tools are available to strategic customers." A spokesperson declined to comment on the identity of these customers.

And the programming tools for Everest should move even further up the software stack. "Fundamentally, there is a cap on the number of programmers who can understand massively parallel programming," said Teich. "It requires thinking differently. Everyone in the industry is trying to develop software tools to hide that complexity from most programmers, who just want to simulate things and look for patterns in data."

That was the same thinking behind the devleopment of TensorFlow, Caffe and other machine learning libraries. That type of software has been an equalizer in the machine learning hardware space as startups like Graphcore and Wave Computing can focus on the development of compilers rather than an expansive ecosystem on par with Nvidia's Cuda platform.

Peng said Xilinx is also working on integration with TensorFlow, Caffe and other frameworks.

Every chip maker targeting artificial intelligence is spending lavishly. Xilinx poured more than a billion dollars over four years – almost half the company's $2.15 billion in research and development expenses from 2014 to 2017 – and more than a tenth of its $9.32 billion in revenue over the same span – into roughly 1,500 employees that worked on the ACAP and Everest.

The expense underlines the lengths that companies are going to develop machine learning chips. Jensen Huang, Nvidia's chief executive officer, has said that it cost $3 billion to create the company's Volta architecture. Volta graphics chips contain custom tensor cores that accelerate the matrix multiplications used in deep learning.

Intel, struggling with the speed of the artificial intelligence market, reportedly spent $350 million in its acquisition of Nervana Systems, which has been tasked with engineering a server chip for neural network training. The company also allegedly spent about $400 million on Movidius, which built a computer vision chip architecture for edge devices.

Hundreds of millions of dollars of funding is flowing into hardware startups, too. Last year, Graphcore raised $50 million in financing led by Silicon Valley institution Sequoia Capital, while Wave Computing has solicited $117 million from investors including Dado Banatao. Cerebras Systems, which last year was reportedly valued at almost $1 billion, is funded in part by Benchmark Capital.

"The world of CPU-centric computing is over," Peng said. "What that means is that in this new era, architecture will be heterogeneous with accelerators. I say accelerators, not accelerator, because of the breadth of applications that will integrate some form of artificial intelligence is vast. There is not a single accelerator that will do all that well."

Xilinx believes that the ACAP technology inside Everest can take the place of many accelerators, switching between different workloads and algorithms dynamically. "We are in the very early stages of the emergence of artificial intelligence," Peng told reporters in the conference call. ■

WILLIAM WONG | Senior Technology Editor

AMD Ramps Up Embedded Processor Performance, Security

The EPYC Embedded and Ryzen Embedded processor families stack up well against their respective competition, the Intel Xeon D and Intel Core.

Developers looking for embedded processors tend to have more demanding requirements than those used in PCs and servers. This can include a wider operating temperature range or a longer life cycle. The latter typically extends beyond five years, with many being used for 10 or 20 years.

Along these lines, AMD announced its EPYC Embedded and Ryzen Embedded processor families to target the embedded space. These are compatible with the EPYC server chips and Ryzen processors. AMD is providing long-term support and availability for these families.

IT'S EPYC

The EPYC Embedded 3000 (*Fig. 1*) is built on the same platform as the EPYC 7000 server chips. The embedded version sports a maximum of 16 Zen cores with up to a 32-MB shared L3 cache. The EPYC Embedded 3000 has four independent memory channels and up to 64 PCI Express Gen 3 lanes. It has a primary frequency of 2.15 GHz and a boost frequency of 3.1 GHz. Systems can handle up to 1 TB of memory.

The EPYC Embedded 3000 system-on-chip (SoC) doesn't require additional support chips. It incorporates up to eight 10-Gb Ethernet ports and up to 16 SATA/NVMe ports.

The embedded family currently matches the midrange EPYC 7000 server chips that have up to 32 cores and 128 PCI Express lanes. The EPYC family also shares a security system with a secure Cortex-A5 processor that provides a secure root-of-trust. They also support Secure Memory Encryption (SME) and Secure Encrypted Virtualization (SEC).

SME works with encryption and decryption hardware that's integrated with the memory controllers. Data in DRAM is always encrypted with SME enabled. It's decrypted as it comes into the SoC and is encrypted before being written to memory. SME is a BIOS option that allows it to be used transparently with respect to the operating system and applications. The secure processor manages the encryption keys.

SEC is a bit more complicated. It also uses SME, but is employed on a virtual-machine (VM) mode. Each VM has its own encryption keys that are managed by the secure processor. This allows the system to be used with an unsecured hypervisor since the encryption is managed by the chip. The hypervisor manages memory and the VM and doesn't care that it's encrypted. Of course, some hypervisor modifications are needed to support SEC properly. These have been incorporated into the mainline Linux implementation as well as other operating systems.

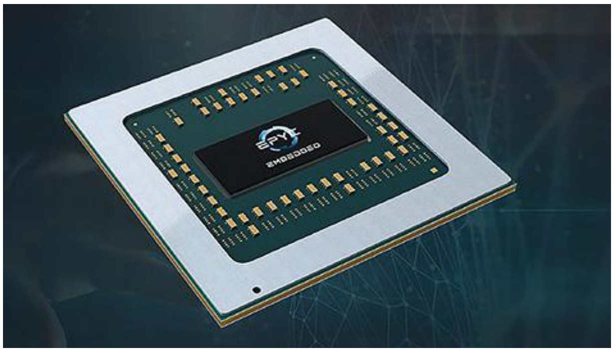

1. The EPYC Embedded brings the EPYC server family to the embedded space.

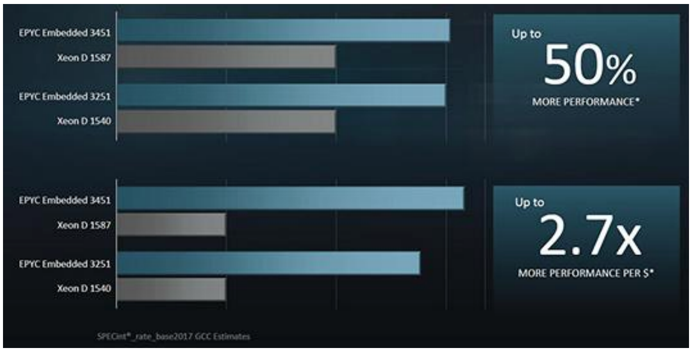

2. Price-to-performance ratio of the EPYC Embedded outshines the competition.

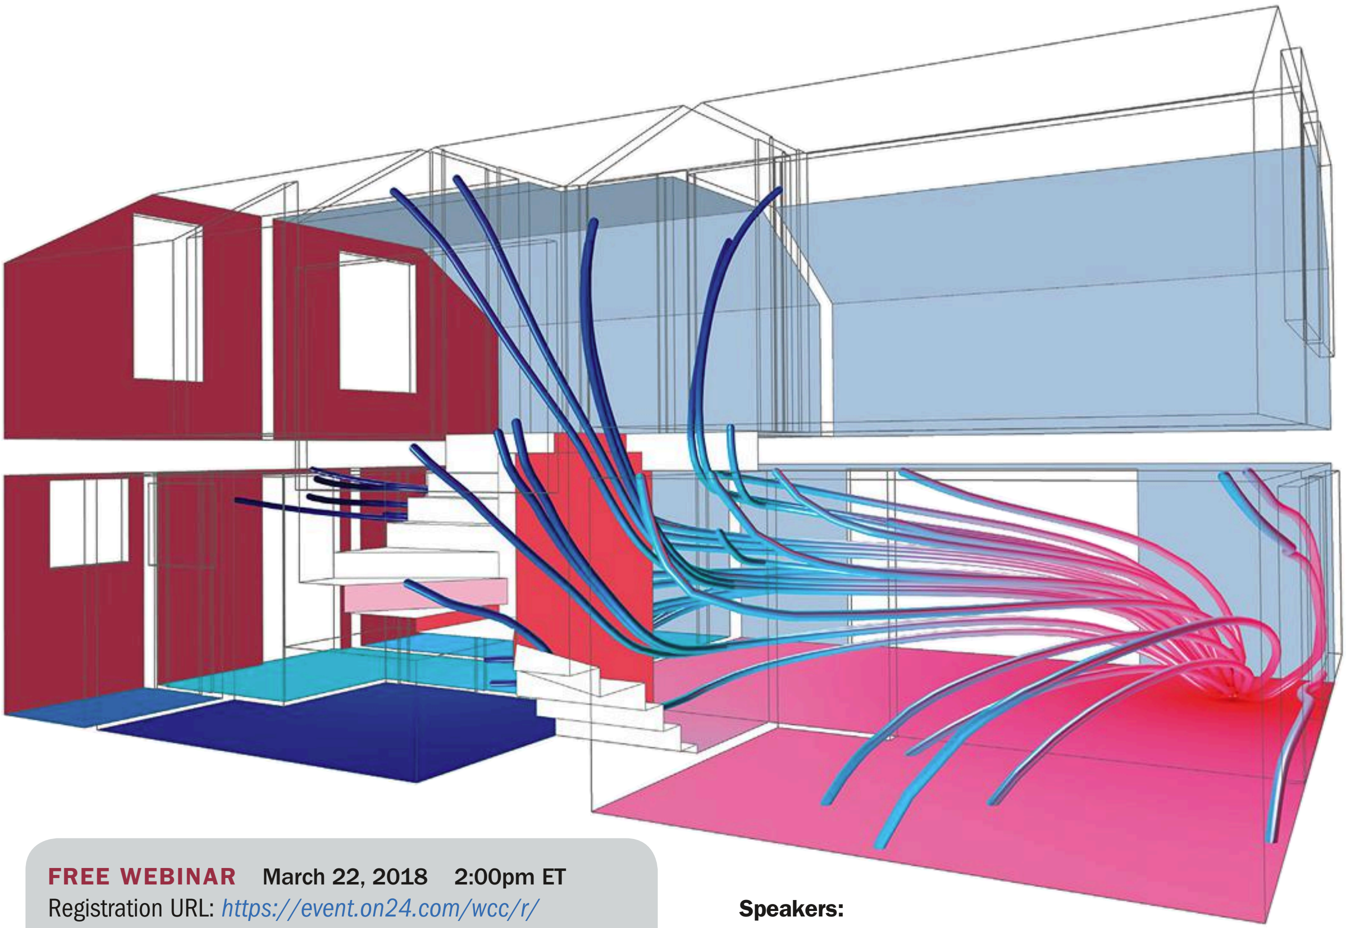

FREE WEBINAR March 22, 2018 2:00pm ET
Registration URL: *https://event.on24.com/wcc/r/1576736/2F550C2290C0B6B4176FBB04B57E2F18*

MODELING VIBRATION IN THE BUILT ENVIRONMENT

If you want to learn about simulating vibration in the built environment, then tune into this webinar with *Murray Snaith* of Xi Engineering Consultants.

Structural response to vibration is an important part of the architectural and structural design process, as the response can be critical to the overall success of a building.

In this webinar, we will discuss how the COMSOL Multiphysics® software can be used to simulate the dynamic response of buildings. We will also discuss the effect of external and environmental vibrations on a structure and how these vibrations can be modeled. During a live demonstration, you will see how a frequency response analysis can be used to predict vibration levels in a converted classroom that is used to house sensitive lab equipment.

You can ask questions at the end of the webinar during the Q&A session.

COMSOL

Speakers:

MURRAY SNAITH, *Xi Engineering Consultants*

Murray Snaith is the resident expert for buildings and structures at Xi Engineering Consultants and has worked on improving the acoustics and vibrations in buildings of all shapes and sizes. With a background in mechanical engineering, Murray is a member of the IMechE and IoA and works as one of the lead project engineers at Xi. His experience covers all aspects of noise and vibration and he regularly hosts CPD events on noise and vibration in the built environment. Murray is a team player and works closely with his clients to deliver robust engineering solutions, such as planning and design validation as well as project implementation and completion.

VALERIO MARRA, *COMSOL*

Valerio Marra is the marketing director at COMSOL in Burlington, MA. Previous roles include working in applications, technical support, and sales. He has been at COMSOL since 2008. Valerio received his PhD in fluid machines and energy systems engineering and his MSc in nuclear engineering, focusing on numerical methods for CFD.

There's an overhead for SEC but it is very small. It also applies only to an encrypted VM. The hypervisor can handle both types of VM. As with SME, security is transparent to the application or VM.

EPYC Embedded competition includes Intel's Xeon D. EPYC Embedded's performance is better, but its price/performance is the more impressive difference (*Fig. 2*). This is due to AMD's lower price. That's something Intel could match with a price drop, but for now AMD has the price/performance edge. The EPYC Embedded 3451 is priced at $880, while the Xeon D-2191 goes for $2407.

A range of the latest operating systems support EPYC Embedded, including Red Hat Linux, Canonical's Ubuntu, Mentor Graphics' Mentor Embedded Linux (MEL), Wind River's Wind River Linux Base, and the latest Yocto project.

RYZEN TO THE OCCASION

EPYC Embedded doesn't include a GPU, but one is integrated into the Ryzen Embedded (*Fig. 3*). The Ryzen Embedded V1000 APU also has Zen cores, and includes an AMD VESA GPU using a Heterogeneous System Architecture (HSA). The video support can handle up to four 4K displays. It's 4K60 encode/decode support includes H.265 and VP9. The single-chip solution handles HDMI 2.0b, DisplayPort (DP) 1.4, and eDP1.4. The chips have dual 10-Gb Ethernet, allowing them to be used for gateways. The CPU, GPU, and peripherals are tied together using AMD's Infinity Fabric.

3. Unlike the EPYC Embedded, the Ryzen Embedded incorporates a GPU.

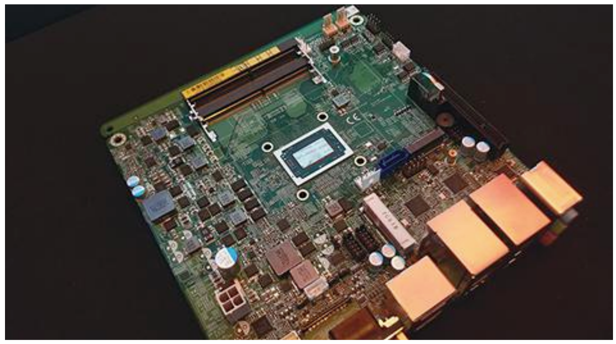

5. The iBase MI988-A is a Mini-ITX motherboard that integrates a Ryzen Embedded V1000 processor.

Versions are available with power requirements from 12 to 54 W. There are two- and four-core versions that run two threads per core. The base frequency is 1.3 GHz, with a boost speed up to 3.8 GHz. The top-end V1807B has 11 GPU compute units.

The Ryzen Embedded family competes with Intel's Core family (*Fig. 4*). AMD's processors shine when it comes to GPU performance. The family is supported by Microsoft Windows 10, Canonical Ubuntu, Mentor Graphics' MEL, Wind River's Wind River Linux Base, and the latest Yocto project.

Systems that utilize the EPYC Embedded and Ryzen Embedded chips are already finalized. In fact, a number of motherboards and modules using the Ryzen Embedded processors are being shown at this year's Embedded World. For example, the iBase MI988-A is a Mini-ITX motherboard with a Ryzen Embedded V1000 processor (*Fig. 5*). The board exposes HDMI, DisplayPort, eDP, and 24-bit dual channel LVDS displays. The board has dual Ethernet ports, a PCI Express socket, and an M.2 socket.

4. The Ryzen Embedded's main competition is Intel's Core family.

One feature that makes the Ryzen Embedded chips stand out is security. They also can handle the SME, SEV support and secure boot found in the EPYC processors.

Ryzen Embedded processors are ideal for many applications, such as medical systems and gaming systems, due to its graphics capabilities. These applications require the advanced security features found in AMD's latest chips.

advertisement

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

Design Note

170W Voltage Doubler in 23mm × 16.5mm Package

Ya Liu, Jian Li, Jeff Zhang and Brian Lin

Introduction

For high voltage input/output applications, inductorless, switched capacitor converters (charge pumps) significantly improve efficiency and reduce solution size over conventional inductor-based buck or boost topologies. By using a charge pump instead of an inductor, a "flying capacitor" is used to store and transfer the energy from input to output. The energy density of capacitors is much higher than inductors, improving power density by a factor of 10 using a charge pump. However, charge pumps have traditionally been limited to low power applications, due to the challenges presented in start-up, protection, gate drive and regulation.

The LTC®7820 overcomes these problems, allowing high power density, high efficiency (to 99%) solutions. This fixed-ratio high voltage, high power switched capacitor controller includes four N-channel MOSFET gate drivers, for use with external power MOSFETs, to produce a voltage divider, doubler, or inverter: specifically, 2:1 step-down from inputs to 72V, 1:2 step-up from input voltages to 36V, or a 1:1 inverter from inputs to 36V. Each power MOSFET is switched with 50% duty cycle at a constant preprogrammed switching frequency.

Figure 1 shows a 170W output voltage doubler circuit featuring the LTC7820. The input voltage is 12V and the output is 24V at up to 7A load, and a switching frequency of 500kHz. Sixteen 10µF ceramic capacitors (X7R, 1210 size) act as a flying capacitor to deliver output power. The approximate solution size is 23mm × 16.5mm × 5mm as shown in Figure 2, and the power density is as high as 1500W/in^3.

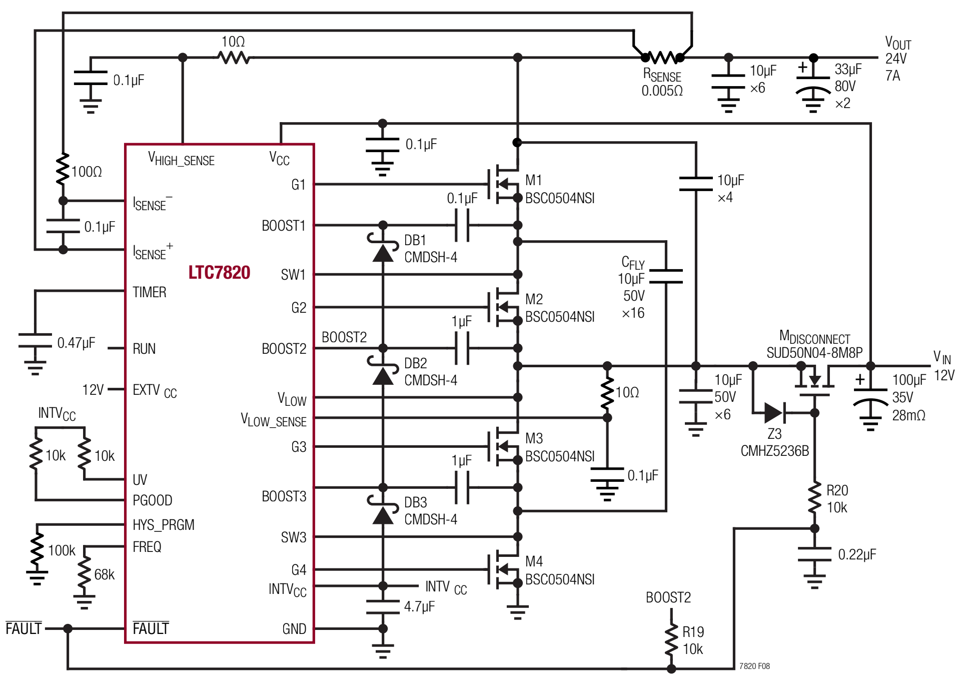

Figure 1. A High Efficiency, High Power Density 12V V_{IN} to 24V/7A Voltage Doubler Using the LTC7820

All registered trademarks and trademarks are the property of their respective owners.

03/18/571

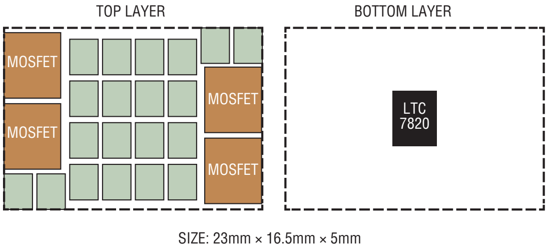

Figure 2. Estimated Solution Size

High Efficiency

All four MOSFETs are soft switched because there is no inductor used in the circuit, greatly reducing switching-related losses. Furthermore, in a switched capacitor voltage doubler, low voltage rating MOSFETs can be used, significantly reducing conduction losses. As shown in Figure 3, the converter can achieve 98.8% peak efficiency and 98% efficiency at the full load. Power loss is balanced between the four switches, spreading thermal dissipation, and simplifying heat mitigation in a smart layout. The thermograph in Figure 4 shows a hot spot temperature rise of only 35°C in free air in an ambient 23°C.

Tight Load Regulation

The LTC7820-based voltage doubler is an open-loop converter, but the LTC7820's high efficiency keeps regulation tight, as shown in Figure 3—the output voltage drops only 0.43V (1.8%) at full load.

Start-Up

In voltage doubler applications, LTC7820 can start-up without capacitor inrush charging current if the input voltage is ramping up slowly from zero. As long as the input voltage ramps up slow (in milliseconds), the output voltage can track the input voltage and the voltage difference between capacitors remains small, so there are no large inrush currents.

Slew rate control of the input can be achieved by using a disconnect FET at input or using hot swap controllers, as shown in the typical application section in the LTC7820 data sheet. In Figure 1, a disconnect FET is used at the input. Unlike voltage divider solutions, the voltage doubler must start up from zero input voltage every time, but it can start up directly with heavy loads. Figure 5 shows the start-up at 7A load.

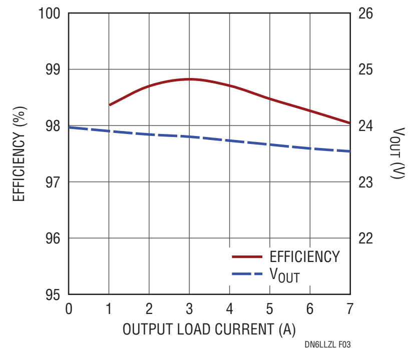

Figure 3. Efficiency and Load Regulation of 12V V_{IN} to 24V/7A Voltage Doubler at 500kHz f_{SW}

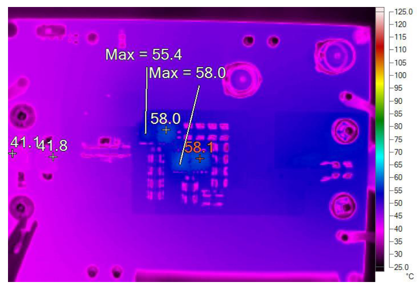

Figure 4. Thermal Test at 12V V_{IN}, 24V V_{OUT}, 7A Load, T_A = 23°C, Free Air

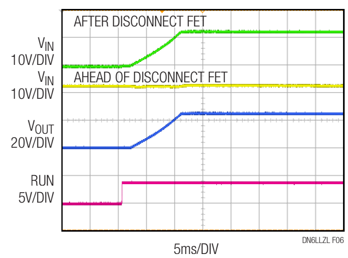

Figure 5. Start-Up Waveform at 7A Load

Conclusion

The LTC7820 is a fixed ratio switched capacitor controller that drives external MOSFETs with built-in gate drivers, achieving very high efficiency (to 99%) and high power density. Robust protection features enable an LTC7820 switched capacitor converter to fit high voltage, high power applications such as bus converters, high power distributed power systems, communications systems and industrial applications.

Data Sheet Download
www.linear.com/LTC7820

For applications help, call (408) 432-1900

DN571 LT/AP 0318 47.7K • PRINTED IN THE USA
© ANALOG DEVICES, INC. 2018

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


Discover the Power of Real-Time Spectrum Analysis
Capture
Transient & short duration signals confidently
Identify
Errors in power, frequency and time
Isolate
Signals of interest
Analyze
RF signals & the devices that control them
Spectrum Monitoring
RF Device Characterization
NEW RSA5000 Series
Real-Time Spectrum Analyzers
Combining a High Performance Swept Spectrum Analyzer with Superior Real-Time Analysis
• 3.2 and 6.5GHz models • 1ms Sweep
• 1Hz RBW • 40MHz Real-Time Bandwidth
• 7.45 µsec 100% POI • Powerful Triggering
• 7 Visualization Modes
RIGOL
Innovation or nothing
Discover the Power of Real-Time @ www.RigolNA.com/Real-Time

BONNIE BAKER | Owner, BB@ATW, LLC

What's the Difference Between DC-DC Conversion Topologies?

When modifying dc voltage from high to low or vice versa in a design, it's important to have a clear understanding of the different converter types involved in the process.

Power supplies represent *the* conversation topic for every circuit designer across the landscape of electrical designs. For instance, you can delve into sensor, Internet of Things (IoT), automation, transportation, medical, or fitness circuit solutions and in every case, one of the primary challenges is to design an appropriate power-supply network for your circuit.

Power-supply classifications come in many flavors. At the top level are the basic ac-ac, ac-dc, dc-ac, or dc-dc converters. In this article, we'll discuss the dc-dc conversion category in terms of shifting the power signal down with a low-dropout regulator (LDO) or buck topologies, and up with a boost topology.

LOW-DROPOUT REGULATOR

When designing your power-supply circuit with an LDO, you're embarking on a simple, straightforward and inexpensive way to change your higher dc voltage down to a smaller voltage.

A linear regulator, or LDO, uses a voltage-controlled current source (VCCS) or transistor to force the regulator's output to a fixed voltage. The control circuitry continuously senses the output voltage and adjusts the current source to hold the output voltage to the predetermined value. The pass transistor's current determines the maximum regulator output-current load while maintaining regulation *(Fig. 1)*.

In Fig. 1, the LDO elements are the pass transistor, error amplifier, voltage reference (V_{ref}), and the feedback/gain

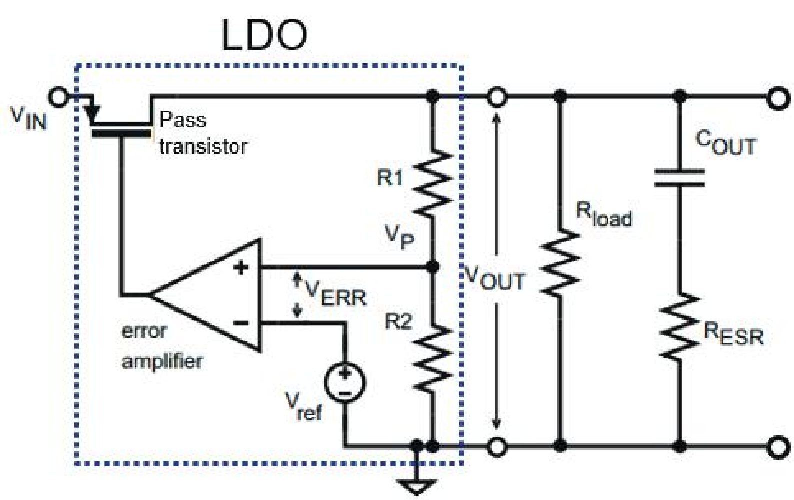

1. The basic LDO circuit design has a pass transistor to provide variable output currents and the topology, though V_{REF}, the error amplifier, and R1/R2 is pre-programmed to provide a relative stable output voltage.

pico Technology

PicoLog® 6

Data logging

Straightforward from the start

Real-time data collection and display

Visual logger and channel setup for easy configuration and viewing

Available for Windows, macOS and Linux

Virtually unlimited logging capacity to PC

Robust database format minimizes data loss and corruption

Supports all current Pico data loggers

Try out the software for free!

Download now from www.picotech.com/A71

Email: sales@picotech.com. Errors and omissions excepted. Please contact Pico Technology for the latest prices before ordering.

resistors (R1 and R2). The V_{ref} and R1+R2 establish the regulator's output dc voltage. A key element in this circuit is the pass transistor. The error-amplifier feedback loop and the voltage reference controls the pass transistor's gate and thereby influences the amount of current that goes to R_{load}. This transistor acts like an adjustable current source and responds to varying changes with the output load.

To assure loop stability, this feedback loop requires loop stability compensation. You will find that most LDOs have built-in compensation to ensure stability without added external components.

The output capacitor parasitic resistance, R_{ESR}, produces a zero into the LDO feedback system. The frequency of this zero is equal to:

$$f_{ZERO} = 1/(2 \times \pi \times R_{ESR} \times C_{OUT})$$

This zero will add a phase shift to the compensation loop. Typically, manufacturer's datasheets address this issue and provide ample guidance.

The LDO is a linear device that theoretically produces a linear power supply, but don't be fooled into using this device in your system and assume you have a good low-noise power supply. These devices will produce a relatively low-noise output voltage with an efficiency of V_{OUT}/V_{IN}; however, their power-supply-rejection capability may not rise to the task of eliminating input noise.

For example, you can have a switching power supply provide the input voltage to your LDO. If the switching power supply's frequency is high enough, that switching noise may pass right through the LDO and to the output pin.

SWITCHED-MODE POWER SUPPLY (SMPS)

The LDO is strictly a linear device that doesn't have any involvement with clock signals, which is contrary to the rest of the devices discussed in this article. In their most-basic form, the currents in the switching regulator require incremental breaks in the current source followed by a filter to reconstitute the dc signal to the output. The non-synchronous buck, synchronized buck, and boost power converters accomplish this with five devices: switch, inductor, capacitor, diode, and clock.

The classes in this technology are the switching regulator and switching controller. These devices are very similar. The configuration for both are dc-dc converters and can generally be either buck or boost.

The switching regulator's integrated circuit (IC) contains all of the necessary components except for the inductor, and a few resistors and capacitors. It's important to notice that

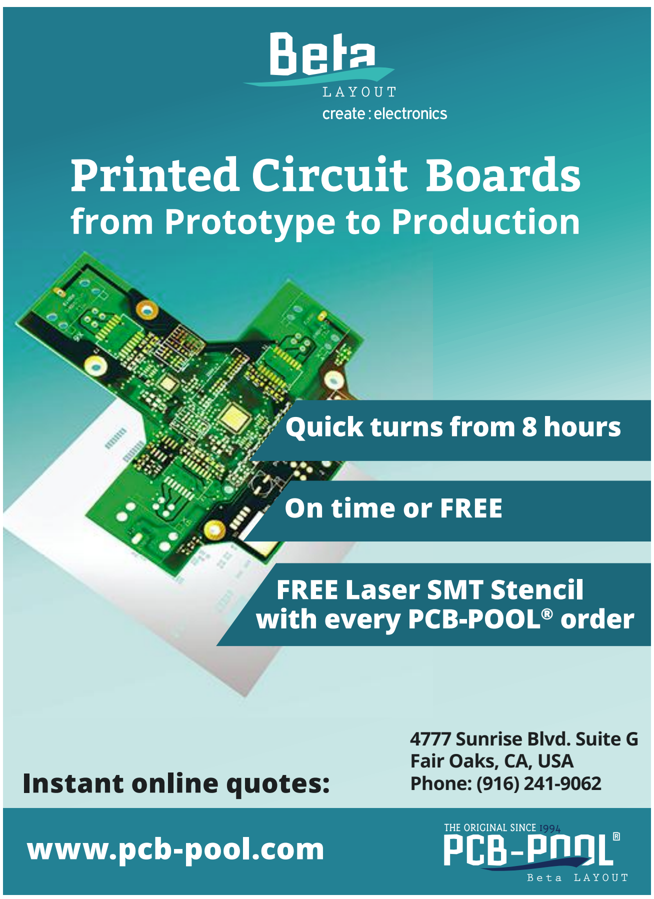

the switching regulator's IC chip includes the FET switch. As a result, switching regulators can't handle too much current. Usually the maximum output current is 1 to 2 A.

In contrast, the switching controller's IC has the same components as the switching regulator, with one exception. The switching regulator doesn't have the FET switch within the IC; it's now external. This new configuration enables higher currents than with the switching regulators. The sizing of the external FET can now match the application task for currents greater than 1 to 2 A.

NON-SYNCHRONOUS BUCK SMPS

The non-synchronous buck switching or step-down topology is the most common form among the switching technologies. The reason for this popularity is that higher voltages are easy to transport over distances, and lower voltages are most useful in electrical circuits.

A non-synchronous buck SMPS steps down an input voltage to a lower output voltage level. In its most basic form, the non-synchronous buck SMPS has an input switch, inductor, diode, and capacitor *(Fig. 2)*.

By definition, the non-synchronous buck switcher output voltage (V_{OUT}) is less than the input voltage (V_{IN}) (Equation 1). The establishment of the output voltage is a by-product of the Q1 duty cycle and the inductor (L1) / capacitor (C2) storage capability (Equation 2).

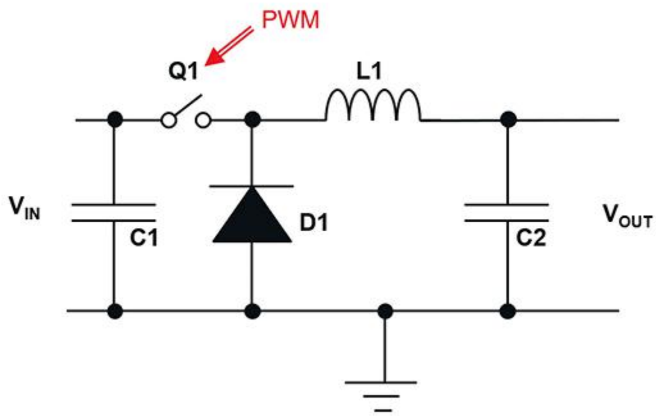

2. The ideal version of the non-synchronous buck SMPS shows the basic components that allow voltage output to be consistently lower than the input voltage.

$$V_{OUT} < V_{IN} \quad (1)$$

$$I_{L\text{-}AVE} = I_{OUT} \quad (2)$$

$$D = V_{OUT}/V_{IN} = t1/(t1 + t2) \quad (3)$$

where D is the pulse-width modulator's (PWM) duty cycle applied to the switch.

In the non-synchronous buck SMPS circuit, the driving frequency force comes from the PWM signal. The signal's duty cycle (D) is the percentage of the ratio of t1 to t1 + t2 where the inverse of the sum of these times is equal to the PWM frequency *(Fig. 3)*.

The FET switch (Q1) inversely reflects the transitions of the PWM signal with the PWM off → Q1 closed. At the same time, the diode (D1) is reverse-biased with no current conduction. The closure of Q1 conducts current to L1 and charges C1 to generate a voltage.

During the PWM t2 time, the FET (Q1) switch opens. At the same time, D1 is forward-biased and pulls current from the inductor to ground. This in turn reduces the C1 voltage.

From this algorithm, you can see that there will be a voltage ripple at the output of the non-synchronous buck SMPS. You can reduce this ripple by increasing the output capacitor (C2) value, but there are boundaries to that. It's important that the power converter remains stable.

As you compare the non-synchronous buck SMPS to the LDO, the non-synchronous buck will have higher efficiency and low thermal dissipation during operation. In addition, the non-synchronous buck can handle large output currents, especially when the FET switch is on the printed circuit board (PCB). As disadvantages, the non-synchronous buck SMPS is more complicated and will have a high part count. Furthermore, the output voltage switching noise and existing ripple can cause errors in the downstream circuit.

Larger non-synchronous buck SMPS duty cycles produce higher dc output voltages. Looking at Equation 3, the output voltage will always be equal to or lower than the input voltage.

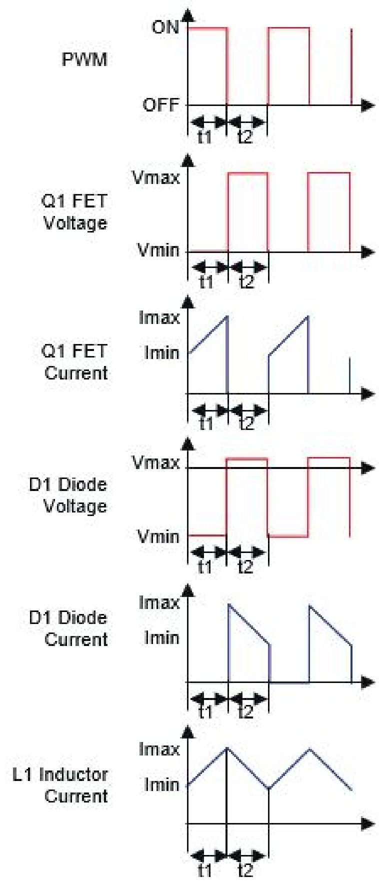

3. Non-synchronous buck timing.

The non-synchronous buck SMPS provides good inductor current filtering because of the inductor-capacitor (LC) arrangement. A continuous switching mode that prevents the inductor current from reaching zero is an ideal condition for the non-synchronous buck SMPS.

SYNCHRONOUS BUCK SMPS

A non-synchronous buck topology is an older design, which is known for its power loss across the diode. This power loss compromises the SMPS's efficiency.

The synchronous buck SMPS accomplishes an increase in efficiency because the voltage drop across the low-side FET (Q2) can be lower than the voltage-drop across the diode of the non-synchronous buck SMPS. With Q2, the current level through the inductor doesn't change, a lower voltage drop translates into less power dissipation and higher efficiency *(Fig. 4)*.


IMS
Connecting Minds. Exchanging Ideas.
The 2018 IEEE MTT-S
International Microwave Symposium
MICROWAVES | MEDICINE | MOBILITY
IMS2018
Philadelphia
Don't miss this year's exciting keynote speakers
at the Microwave Week in Philadelphia !
IMS Plenary Session Speaker (Monday, 11 June 2018):
"The Hitchhiker's Guide To the Healthcare Galaxy: The Actions That Changed the Healthcare Landscape in America From 2017-2027"
Stephen K. Klasko, MD, MBA, President and CEO, Thomas Jefferson University and Jefferson Health
IMS Closing Session Speaker (Thursday, 14 June 2018):
"Extreme Platforms for Extreme Functionality"
Nader Engheta, PhD, H. Nedwill Ramsey Professor at the University of Pennsylvania
RFIC Plenary Session Speakers (Sunday, 10 June 2018):
"Compact Silicon Integrated mmWave Circuits: From Skepticism to 5G and Beyond"
Zachary J. Lemnios, Vice President, Science, Technology & Government Programs, IBM T.J. Watson Research Center
"The Road Ahead for Autonomous Cars – What's in for RFIC"
Lars Reger, Automotive Chief Technology Officer (CTO), Business Unit Automotive, NXP Semiconductors
IMBioC Opening Session Speaker (Thursday, 14 June 2018):
"Renal Denervation for Uncontrolled Hypertension: Complexity After Symplicity"
Dr. Nicholas J. Ruggiero II, MD
IMBioC Closing Session Speaker (Friday, 15 June 2018):
"Is There a Fundamental Law of Health and Disease?"
Dr. Chung-Kang Peng, Director of the Center for Dynamical Biomarkers at Beth Israel Deaconess Medical Center / Harvard Medical School (BIDMC/HMS)
EARLY BIRD REGISTRATION ENDS
14 MAY 2018
FOR MORE INFORMATION,
VISIT WWW.IMS2018.ORG
PENNSYLVANIA CONVENTION CENTER • PHILADELPHIA, PENNSYLVANIA, USA
Exhibition Dates: 12–14 June 2018 • Symposium Dates: 10–15 June 2018
Follow us on:
#IMS2018
IEEE
MTT-S

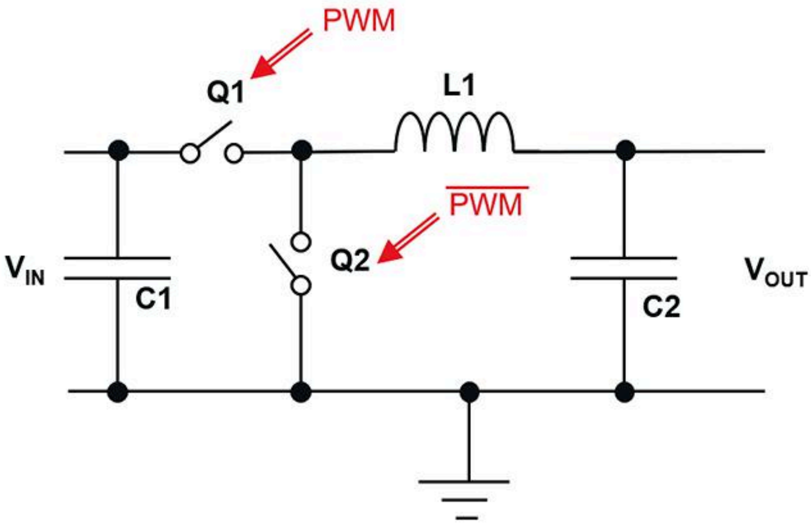

4. The ideal version of the synchronized buck SMPS uses two switches that operate with a PWM non-inverted on Q1 and inverted on Q2 signals.

The synchronized buck SMPS starts with a non-synchronous buck topology. The important distinction is that the second switch (Q2) allows the current to flow in both directions when in the on state.

With the synchronized buck SMPS circuit, the driving frequency force still comes from the PWM signal. The signal's duty cycle (D) is the percentage of the ratio of t1 to t1 + t2 where the inverse of the sum of these times is equal to the PWM frequency.

The FET switch (Q1) inversely reflects the transitions of the PWM signal with PWM off → Q1 closed. At the same time the diode (Q2) is reverse biased with no current conduction *(Fig. 5)*. The closure of Q1 conducts current to L1 and charges the C1 voltage.

During the PWM t2 time, the FET (Q1) switch opens. At the same time, Q2 pulls current from the inductor to ground. This in turn reduces the C1 voltage.

From this algorithm, you can see that a voltage ripple at the output still exists. You can reduce this ripple by increasing the output capacitor (C2) value, but there are boundaries to that as well. It's important that the power converter remains stable.

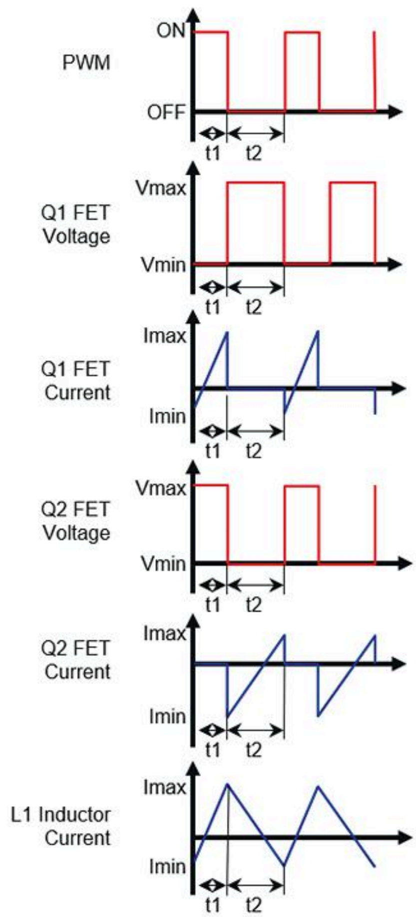

5. Synchronous buck timing.

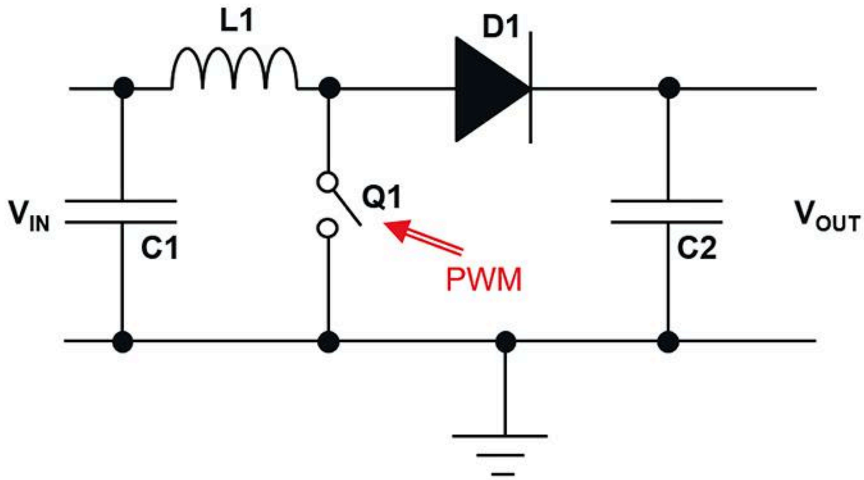

6. The ideal version of the boost SMPS shows the basic components that allow output voltage to be consistently be higher than the input voltage.

As you compare the non-synchronous buck SMPS to the synchronous buck SMPS, the synchronous buck SMPS tends to be smaller and relatively inexpensive. In the synchronous topology, the low resistance ($R_{DS\text{-}ON}$) of Q2 helps reduce losses and optimize overall conversion efficiency. Also, the synchronous buck SMPS circuit is slightly more complicated because of the risk of Q1 and Q2 simultaneously conducting. Timing is everything!

BOOST SMPS

The boost SMPS implements a step-up of the input dc voltage; the dc input source can be batteries, dc solar panels, fuel cells, or dc generators. It increases the input voltage to a higher level. A boost or step-up circuit contains the same components as the buck and synchronous buck *(Fig. 6)*.

The arrangement of components provides an output voltage that is greater than the input voltage (Equation 4).

$$V_{OUT} > V_{IN} \quad (4)$$

$$I_{L\text{-}AVE} = I_{OUT}/(1 - D) \quad (5)$$

$$D = (V_{OUT} - V_{IN})/V_{OUT} \quad (6)$$

±2000 V High Speed Pulse Generator

PVX-4000 Series All-in-One High Voltage Pulse Generators

Simply set pulse width, frequency and voltage with the intuitive front panel interface, or with your computer… *It's that easy!*

With thirty years experience making high voltage pulse generators, Directed Energy is your source for benchtop, module and custom/OEM pulsed solutions.

- 0 to ±2000 V
- Pulse widths from 500 ns to DC
- USB and RS-232
- Pulse rise and fall times ≤50 ns
- Repetition frequencies to 30 KHz

DEI SCIENTIFIC

Directed Energy, Inc. an IXYS Company
www.directedenergy.com | +1-970-493-1901

It's taking off! Can you handle 5G?

Register for the 5G Summit at IMS2018

The 5G Summit on Tuesday, 12 June 2018 at the Pennsylvania Convention Center in Philadelphia is an IEEE event that is organized by two of IEEE's largest societies – MTT-S and ComSoc. This special collaboration, for the second year running, complements MTT-S' "hardware and systems" focus with ComSoc's "networking and services" focus. The one-day Summit features talks from experts from industry, academia, and government on various aspects of 5G services and applications. It's further complemented by the 5G Pavilion at the IMS2018 exhibition where table top demonstrations and "fire-side" chats are presented at the 5G theater.

5G Summit Speakers:

"Bringing the World Closer Together"
Jin Bains
Head of Connectivity, SCL, Facebook

"AT&T Perspectives on 5G Services"
David Lu
Vice President, AT&T

Other featured presentations from Huawei, GM, Keysight, NI, Global Foundries, MACOM as well as academia will include following topics:

- Spectrum/Regulatory
- Infrastructure/Trials, Applications
- Technologies, Circuits, Systems
- Design, Test & Measurement Challenges
- Test-bed Services for 5G

Lunchtime Panel session on, "mmWave Radios in Smartphones: What they will look like in 2, 5, and 10 years"

Industry Co-Sponsors: KEYSIGHT TECHNOLOGIES, NATIONAL INSTRUMENTS

Media Sponsor: Microwave Journal

For more information visit: https://ims2018.org/5g-summit

PENNSYLVANIA CONVENTION CENTER • PHILADELPHIA, PENNSYLVANIA, USA

Exhibition Dates: 12–14 June 2018 • Symposium Dates: 10–15 June 2018

IEEE

Follow us on: #IMS2018

MTT-S   IEEE ComSoc IEEE Communications Society

You will find a power-supply conversion in nearly every circuit. This article covers two types of dc voltage conversion: high to low and low to high. When you modify the dc voltage from high to low, you will be using either an LDO or buck converter.

With the boost circuit, the diode (D1) prevents the inductor (L1) current to flow continuously to the output capacitor (C2). With this circuit, the transmission of power to the output is a two-step process. During this two-step process, the inductor serves as a temporary storage element. The Boost driving frequency force still comes from the Pulse-width-modulator (PWM) signal. The signals duty cycle (D) is the percentage of the ratio of t1 to t1 + t2 where the inverse of the sum of these times is equal to the PWM frequency.

In *Figure 7*, a closed switch (Q1) causes input current to flow into the inductor with an increase in L1's magnetic field. During this time the diode (D1) blocks the flow of current to the load. In addition, the output capacitor will hold the output voltage steady.

Then, an open switch (Q1) causes the inductor's magnetic field to fall due to a reverse voltage polarity. This causes current to go to the output and replenish the output capacitance charge and consequently increase the capacitor voltage. This two-step process creates higher peak currents and lower efficiency in higher power applications.

The boost SMPS is a dc-dc converter that steps-up a low dc input voltage to a higher dc output voltage. The boost-converter input device is an inductor, which is an energy storage device. When the switch (Q1) closes, the inductor current flows to ground. When the switch opens, the stored energy continues to generate current, which then flows through the diode to the output capacitor (C1). The inductor current charges the capacitor. One of the by-products of this switching algorithm is an output-voltage ripple.

Your Job just GOT a whole lot EASIER

NEW for 2018!

Customize your workspace with personalization features that will save you time:

- Customized Parts Library – build and organize parts into an easy-to-manage library
- Product Categories - tag only the one's that matter to you
- Part/Product Alerts – keep up-to-date on product availability
- Search "Authorized Only" Distributors – your reliable source for truly vetting authorized distributors

SourceESB has built-in engineering tools and functionality, designed just for you. If streamlining your workflow and increasing productivity are important, then you will appreciate the new features added to SourceESB in 2018.

SourceESB

SourceESB.com

The most trusted electronics parts search engine that serves the engineer from design to purchase.

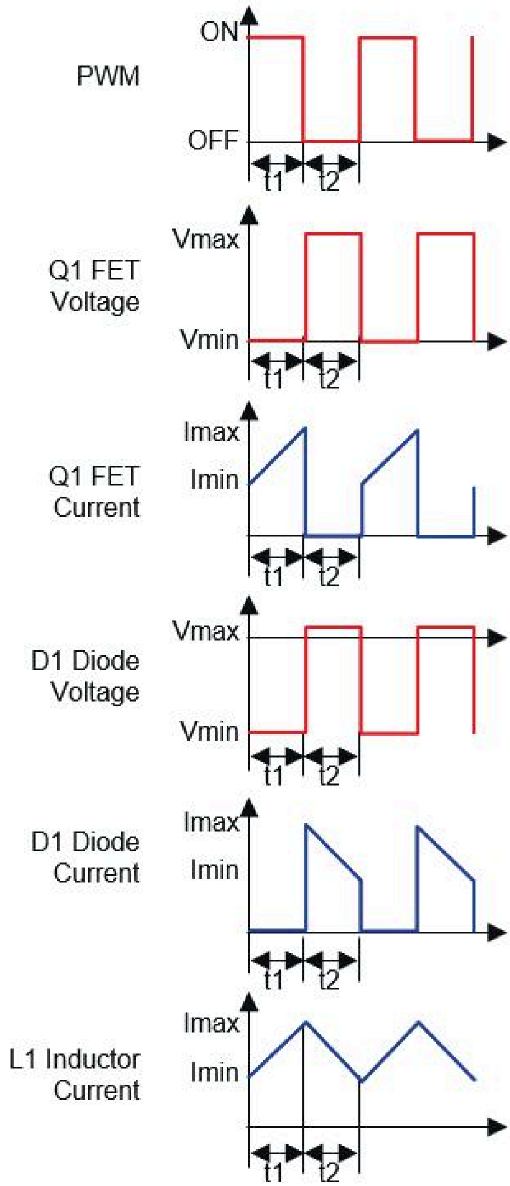

7. Boost timing.

CHARACTERISTICS TO CONSIDER

There are a few characteristics about these SMPSs to take into account as you move them into your applications. With non-synchronous and synchronous buck converters, the input has a direct connection to a switch (Q1). This causes an interruption to the input current and causes radio-frequency-interference (RFI) noise. As you can see with the capacitor C1 in Figs. 2, 4 and 6, input filtering is in order.

The boost circuit has a rectifier (Q1) blocking the output-capacitor charging current. This causes a much higher ac current in the capacitor, with potential switching noise at the output.

Of all these topologies, the synchronous buck SMPS is the most efficient. This is a nice feature, but be aware—if the step-down ratio is greater than 10:1, the PWM pulses will narrow, causing higher peak currents. The consequence of this is loss that's proportional to I^2.

This is a brief introduction to the more popular SMPS topologies. Please note that these three topologies did not illustrate the functions of the basic controls, such as soft-start, feedback compensation, shutdown, dead-time control, or slope compensation. When selecting your next product, these are issues to consider.

CONCLUSION

You will find a power-supply conversion in nearly every circuit. This article covers two types of dc voltage conversion: high to low and low to high. When you modify the dc voltage from high to low, you will be using either an LDO or buck converter. The LDO will provide a good linear output at the expense of lower efficiency. You will improve your efficiency with a buck SMPS, but acquire a ripple at the output of the device.

To change a low dc voltage to a higher dc value, the boost topology is appropriate. No linear conversion option is available for this function, but the boost SMPS effectively provides a good solution to this problem. ed

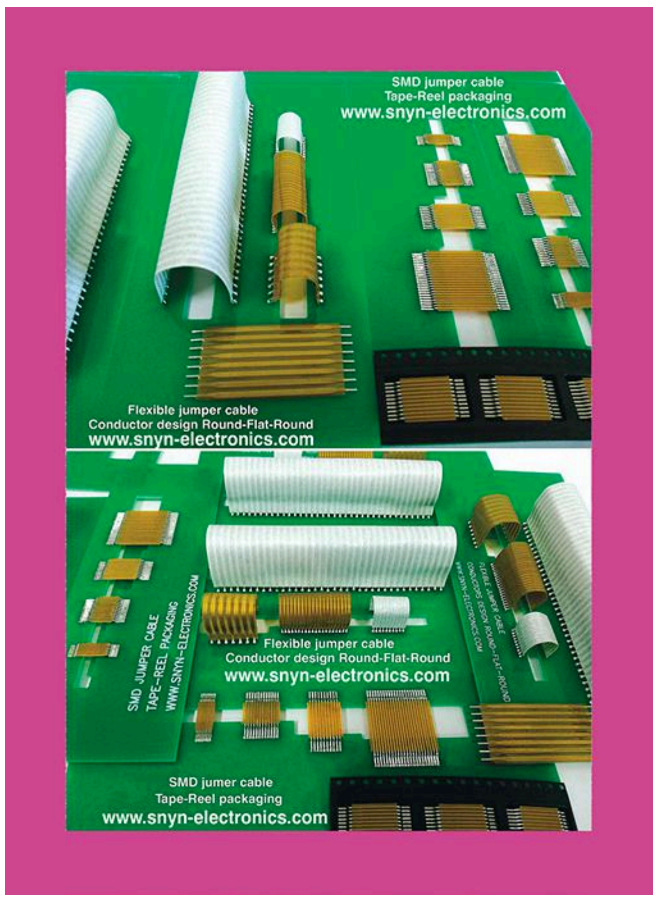

ideas for design

Simple Generator Provides Very-Low-Frequency/Distortion Sine and Square Waves

By JIM TONNE | Tonne Software

THE CIRCUIT OF *Figure 1* generates sinusoids down to very low frequencies with distortion in the region of 3% or less, yet has no feedback or gain-stabilizing components because none are needed. It uses a phase-shift oscillator with a low-pass phase-shift network configuration rather than the more-common high-pass network.

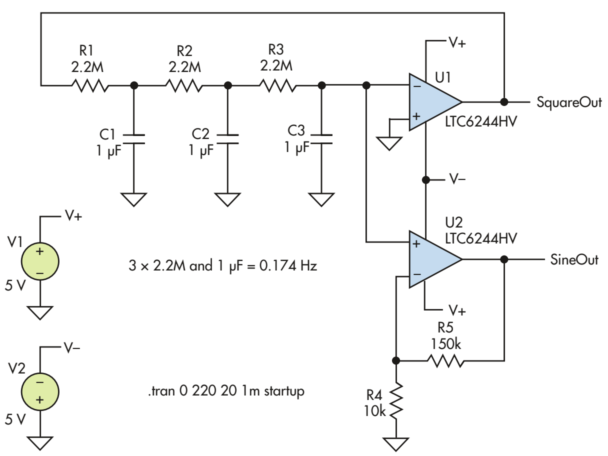

1. This very-low-frequency sine- and square-wave generator requires few components, but provides low-distortion outputs (derived from LTspice simulator).

The frequency-determining network is used as a low-pass filter to remove most of the harmonic content of the output waveform. Other phase-shift oscillators using low-pass networks have been published, but most have been more complicated (some of them far more so).

The output of op amp U1 is applied to the first section of the phase-shifting network via R1 and C1. Each stage of the network attenuates more of the harmonic content (along with some of the fundamental). The final sinusoidal waveform is fed back into U1, which is operating open-loop and thus generates square waves at its output. It's also fed into U2, which is operating as a linear amplifier to restore the low-level sinusoid up to a more-usable level at low impedance.

With three 2.2-MΩ resistors and three 1-µF capacitors, the circuit generates a sinusoid waveform at about 0.174 Hz. (Note that a three-gang variable-tuning capacitor can be used here to create an inexpensive, adjustable-frequency audio oscillator.)

The circuit has a quick startup of a few cycles regardless of frequency, and its output amplitude is very stable. U1 doesn't need to deliver output rail-to-rail, but it must clip on overload symmetrically; if it does not, then even-order harmonic components will appear in the output. If this circuit also had a true triangular-wave output, it would be called a function generator. However, this simple configuration doesn't have that output waveform.

By increasing the values of the network components, this circuit can generate sinusoids with periods of the order of a minute or more, which largely depends on the characteristics of the components. With the high resistor values shown, the op amps must be CMOS devices in order to have the necessary extremely high input impedance, and the timing capacitors must also have very low leakage. To view the circuit's waveforms at different critical points, check out the online version of this article at *www.electronicdesign.com*.

A possible negative aspect of this circuit is that it requires a second op amp as a buffer with gain to provide the sinewave output. That second op amp has the same high input-impedance requirement as the first one. Although shown as a generator of very low frequencies, the upper-frequency range of this circuit is limited only by the gain-bandwidth product characteristics of the chosen op amps. By using available wideband op amps, it will operate into the upper parts of the audio region without a serious increase in distortion.

JIM TONNE has been involved with electrical-circuit design and analysis, largely in the broadcast industry, since 1965. His specialties are filter design and speech processing, and he has been a regular contributor to the *ARRL Handbook* as well as QEX, their Forum for Communications Experimenters. His website is www.TonneSoftware.com and can be reached at Sales@TonneSoftware.com.

Electronic Design

July/August Issue Preview

Ad close: 7/6/18

Materials due: 7/13/18

TECHNOLOGY
Artificial Intelligence

PRODUCT TRENDS
FPGAs

ENGINEERING ESSENTIALS
Persistent Memory

DIFFERENCE BETWEEN
Power Supply

http://electronicdesign.com

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

CITY	POSITION	APPLY FOR JOB USING THIS LINK
Dallas/Richardson/ Sherman, TX	Analog Design Engineer	https://tinyurl.com/kgjueje
	Analog Design Manager	https://tinyurl.com/kq7bdda
	Digital Design Manager	https://tinyurl.com/l3hjj23
	Failure Analysis Engineer	https://tinyurl.com/knh5ofs
	Marketing Manager	https://tinyurl.com/kpmpljr
	Process/Equip Engineer	https://tinyurl.com/k649hms
	Software Development Manager	https://tinyurl.com/ohzev3v
	Systems Administrator	https://tinyurl.com/yam3b5kg
	Test Engineer	https://tinyurl.com/qgy7vqp
Santa Clara/ Sunnyvale, CA	Analog Design Engineer	https://tinyurl.com/ke86n9o
	Analog Design Manager	https://tinyurl.com/lbjtqab
	Digital Design Manager	https://tinyurl.com/ycvu4qoo
	Field Application Engineer Analog	https://tinyurl.com/ktu3wrr
Tucson, AZ	Analog Design Engineer	https://tinyurl.com/m6hgk9e

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

INDEX

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

Subscribe to Newsletters

ElectronicDesign®

www.electronicdesign.com/newsletters/signup

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

SIP Solutions Streamline System Design

Modules and system-in-packages (SIPs) continue to deliver functionality and performance while simplifying system design.

System-on-module (SOM) and system-in-package (SIP) solutions were out in force at this year's Embedded World. These cost a bit more than trying to put comparable logic down on a printed circuit board (PCB), but most will likely result in lower-cost solutions that are more flexible and powerful.

Take, for example, Octavo Systems' OSD335X-SM SIP design (*Fig. 1*). It packs in more than 100 components including a 1-GHz Texas Instruments Sitara AM335x ARM Cortex-A8 Processor, up to 1 GB of DDR3L memory, a TPS65217C power-management IC (PMIC), a TL5209 LDO regulator, a 4- kB EEPROM, and all necessary passives into a 21- by 21-mm, BGA. The AM335x can run Linux and includes a touchscreen display controller with a 3D PowerVR SGX graphics system; USB OTG with PHY; dual Gigabit Ethernet; an 8-channel, 12-bit SAR ADC; and a programmable real-time unit (PRU) that allows the network to support protocols like EtherCAT.

The OSD335X-SM groups signals to simplify PCB layout (*Fig. 2*). It has 16- by 16-ball grid with a 1.27-mm pitch that allows routing on a single side of a PCB. This allows the SIP to be used in a double-sided PCB, thus keeping system costs low. Octavo Systems SIPs are used in the popular BeagleBone platforms.

Another platform I saw at Embedded World was Silicon Labs' WFM200 Wi-Fi SIP (*Fig. 3*). This SIP targets the Internet of Things (IoT) and combines a secure microcontroller with a radio and built-in antenna. The 6.5- by 6.5-mm chip supports 802.11 b/g/n as well as external antennas. One advantage of the internal antenna is that the chip is pre-certified by the FCC, CE, IC, and in South Korea and Japan. This can significantly reduce system design costs and the related expertise.

The WFM200 incorporates a variety of security-related features including secure boot, secure debugging, AES and PKE hardware accelerators, and a true random number generator (TRNG). There's wireless support for packet traffic arbitration (PTA) coexistence, and it has a link budget of 115 dBm at 1DSSS. The receive power is only 48.6 mA at 1.8 V, while transmit power requirement is a mere 128 mA.

NXP's IoT-on-a-Chip (*Fig. 4*) blends its i.MX family with Wi-Fi and Bluetooth connectivity plus 512 kB of SRAM and 2 MB of SST-eMRAM into a 14- by 14-mm FBGA. The latter is via Samsung's foundry support for MRAM.

Designing with SIPs like the WFM200 and the OSD335X-SM can speed time to market. They also make it easier for designers to concentrate on their secret sauce while using the latest technology, without having to gain expertise in areas where SIP designers provide the solution. ed

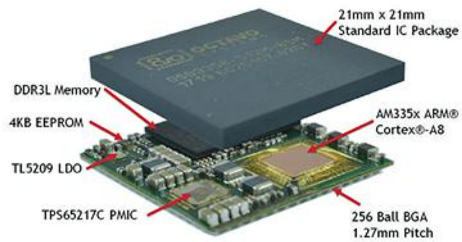

1. Octavo Systems' OSD335X-SM packs an entire system into a compact SIP.

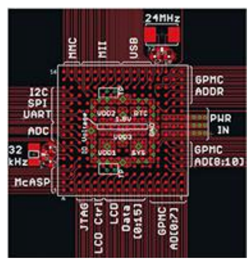

2. The SIP developed by Octavo Systems is designed to simplify PCB layout, allowing it to be used in many double-sided layouts.

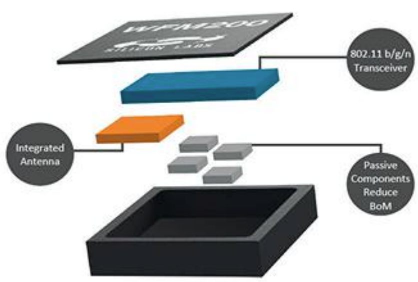

3. Silicon Labs' WFM200 Wi-Fi SIP includes an integrated antenna.

4. NXP's IoT-on-a-Chip SIP combines an i.MX SoC with Wi-Fi and Bluetooth support. The chip is shown here in the upper left of NXP's SOM board.

POWER ELECTRONICS

COMPONENTS, ASSEMBLIES, AND CUSTOM SOLUTIONS

Dean Technology offers a wide range of power electronics in all common packages, with ratings to cover most applications. As with everything we offer custom specifications, packages, and full assemblies are a specialty - we pride ourselves on helping customers get exactly what they need for their success.

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

www.deantechnology.com

HVCA CKE HVPSI

PRODUCTS BY: DEAN TECHNOLOGY

ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

CUTTING EDGE PRECISION

Your product's performance depends on precise signal chain components. We have the answers: the highest accuracy data converters, the lowest drift references, the lowest noise amplifiers and LDOs. Our industry-leading precision signal chain products ensure precise and repeatable measurement and control. Contact our signal chain experts to select a high performance solution.

LTC2387-18	Industry's Fastest No Latency SAR ADC: 18-Bit, 15Msps
LT6657	Industry's Lowest Drift Bandgap Voltage Reference: 1.5ppm/°C Drift
LTC2641	Industry's Lowest Glitch Precision DAC: 0.5nV•s Glitch
LTC2500-32	Industry's Most Accurate 32-Bit ADC: 0.5ppm INL
LT3042	Industry's Lowest Noise LDO: < 1μV Noise
LT6018	Industry's Lowest Noise Op Amp: 30nVP-P

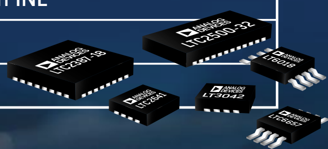

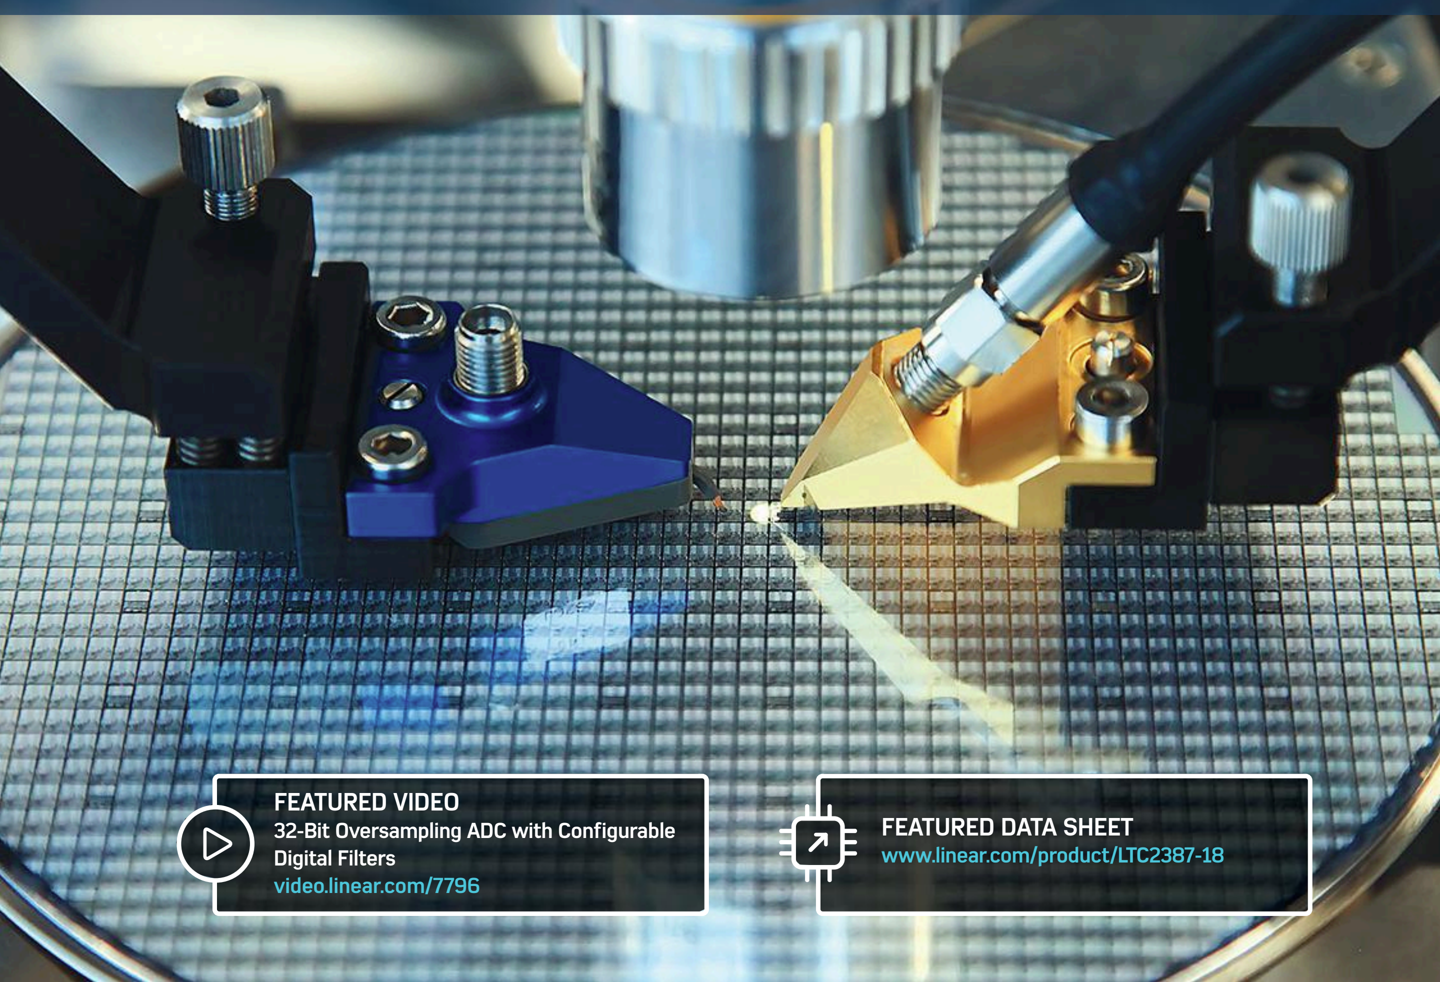

FEATURED VIDEO
32-Bit Oversampling ADC with Configurable Digital Filters
video.linear.com/7796

FEATURED DATA SHEET
www.linear.com/product/LTC2387-18