Electronic Design®

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

APRIL 2017 electronicdesign.com

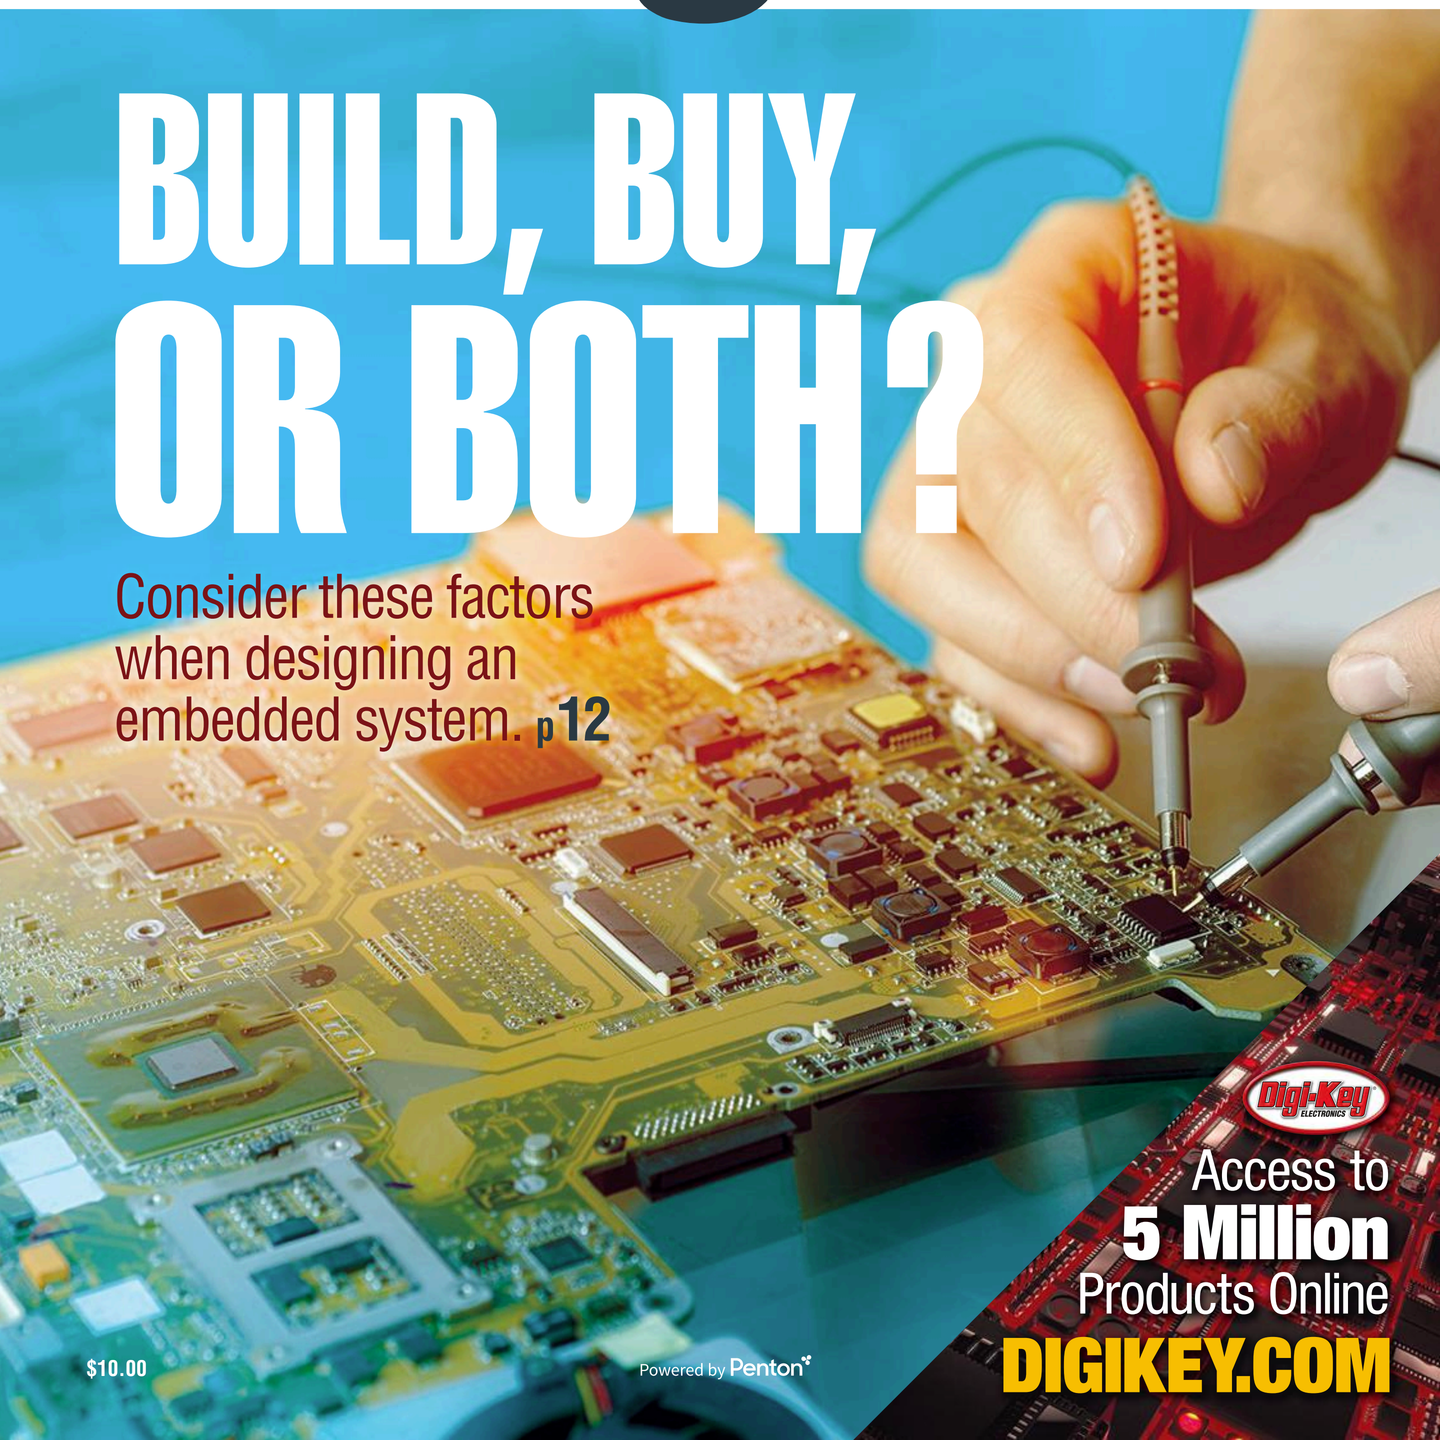

BUILD, BUY, OR BOTH?

Consider these factors when designing an embedded system. p12

$10.00

Powered by Penton


SAME-DAY SHIPPING
on orders entered by 8:00 pm CENTRAL
NEXT-DAY DELIVERY
Digi-Key ELECTRONICS
THE WORLD'S LARGEST SELECTION OF ELECTRONIC COMPONENTS
AVAILABLE FOR IMMEDIATE SHIPMENT®
NOW more than 5 MILLION PRODUCTS from over 650 SUPPLIERS
DIGIKEY.COM
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. © 2017 Digi-Key Electronics, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA
ecia MEMBER
ecsn member
CEDA

IN THIS ISSUE

FEATURES

12 **Embedded Design: Build, Buy, or Both?**
It has never been easier to build an embedded system, whether one makes or buys the hardware. Of course, there are many ways to do this.

16 **The Advantages—and Challenges—of Energy Harvesting**
In an interview with Maurizio Di Paolo Emilio, the author and telecommunications engineer talks about energy harvesting technology and what to expect in this space.

18 **Voltage-Regulator ICs for Wearables: How High Is Your I_Q?**
Steadily evolving switching regulators are a good option for designers tasked with creating smaller and more efficient solutions to achieve low quiescent current in applications that must be energy-efficient.

20 **Major Trends in Wireless Test**
The 5G, IoT, Wi-Fi, and automotive markets are driving wireless instrument development.

28 **Securing IoT Medical Devices—Are We There Yet?**
To ensure secure communications in a given design, developers must consider integrating key security- and safety-related features that help to harden a medical device against any malicious activity.

36 **What's the Difference between Separation Kernel Hypervisor and Microkernel?**
As the embedded world looks to security solutions to protect connected critical computing functions from external threats, two software platforms have emerged that, at first glance, appear to offer similar functionality.

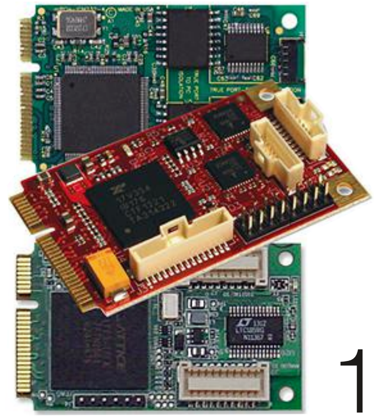

12

20

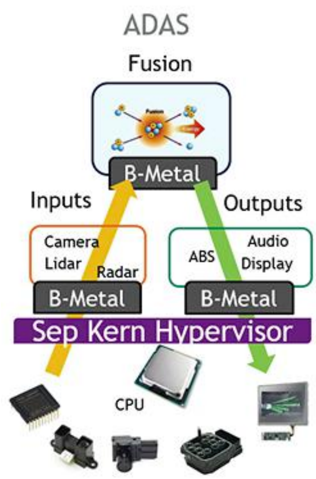

36

IDEAS FOR DESIGN

40 Interactive Graphical Interface Enhances Micro-Programmed Sequential Machine Design

42 "Cap-Drop" Approach Implements Offline Power Supply for Light Loads

COLUMNS & DEPARTMENTS

4 **ON ELECTRONICDESIGN.COM**

7 **EDITORIAL**
Security Needs to be Simple

8 **NEWS & ANALYSIS**

48 **LAB BENCH**
Will ISA Survive?

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

ELECTRONIC DESIGN (ISSN 0013-4872) is published monthly by Penton Media Inc., 9800 Metcalf Ave., Overland Park, KS 66212-2216. Paid rates for a one-year subscription are as follows: $120 U.S., $180 Canada, $240 International. Periodicals postage paid at Kansas City, MO, and additional mailing offices. Editorial and advertising addresses: ELECTRONIC DESIGN, 1166 Avenue of the Americas, New York, NY 10036. Telephone (212) 204-4200. Printed in U.S.A. Title registered in U.S. Patent Office. For subscriber services or to order single copies, write to Electronic Design, PO Box 2100, Skokie, IL 60076. POSTMASTER: Send change of address to Electronic Design, PO Box 2100, Skokie, IL 60076. Canadian Post Publications Mail agreement No. 40612608. Canada return address: IMEX Global Solutions, P.O. Box 25542, London, ON N6C 6B2.

ON ELECTRONICDESIGN.COM

Factory Automation Banks on Wireless and AI

As more factories become automated, it's clear that communications between machines, robots, and the computers that direct them is critical. The more information that's communicated, the greater the productivity. Integration of wireless sensor networks and artificial intelligence (AI) helps meet those demands.

http://electronicdesign.com/industrial/factory-automation-banks-wireless-and-ai-technology-succeed

Is RISC-V Ready for the Spotlight?

The introduction of SiFive's Freedom E310 RISC-V chip in 2016 and a high profile at Embedded World are helping raise the awareness of RISC-V, an open processor Instruction Set Architecture (ISA). Despite the buzz, some engineers discuss express concern that RISC-V isn't quite ready for primetime for their applications.

http://electronicdesign.com/embedded/risc-v-architecture-path-mainstream-likely-be-bit-slog

Where's the Fun in Assisted Driving?

Advances are happening in tandem for networked cars and automated driving, promising a world in which every car can "see" and respond to collision and other safety warnings, and drivers simply become passengers.

http://electronicdesign.com/automotive/where-s-fun-assisted-driving

7 Energy Storage Disruptors to Watch

The energy storage industry is growing quickly as battery technologies evolve and costs continue to fall, spurring innovations that change how we use, store, buy, and sell energy.

http://electronicdesign.com/analog/7-energy-storage-disruptors-watch

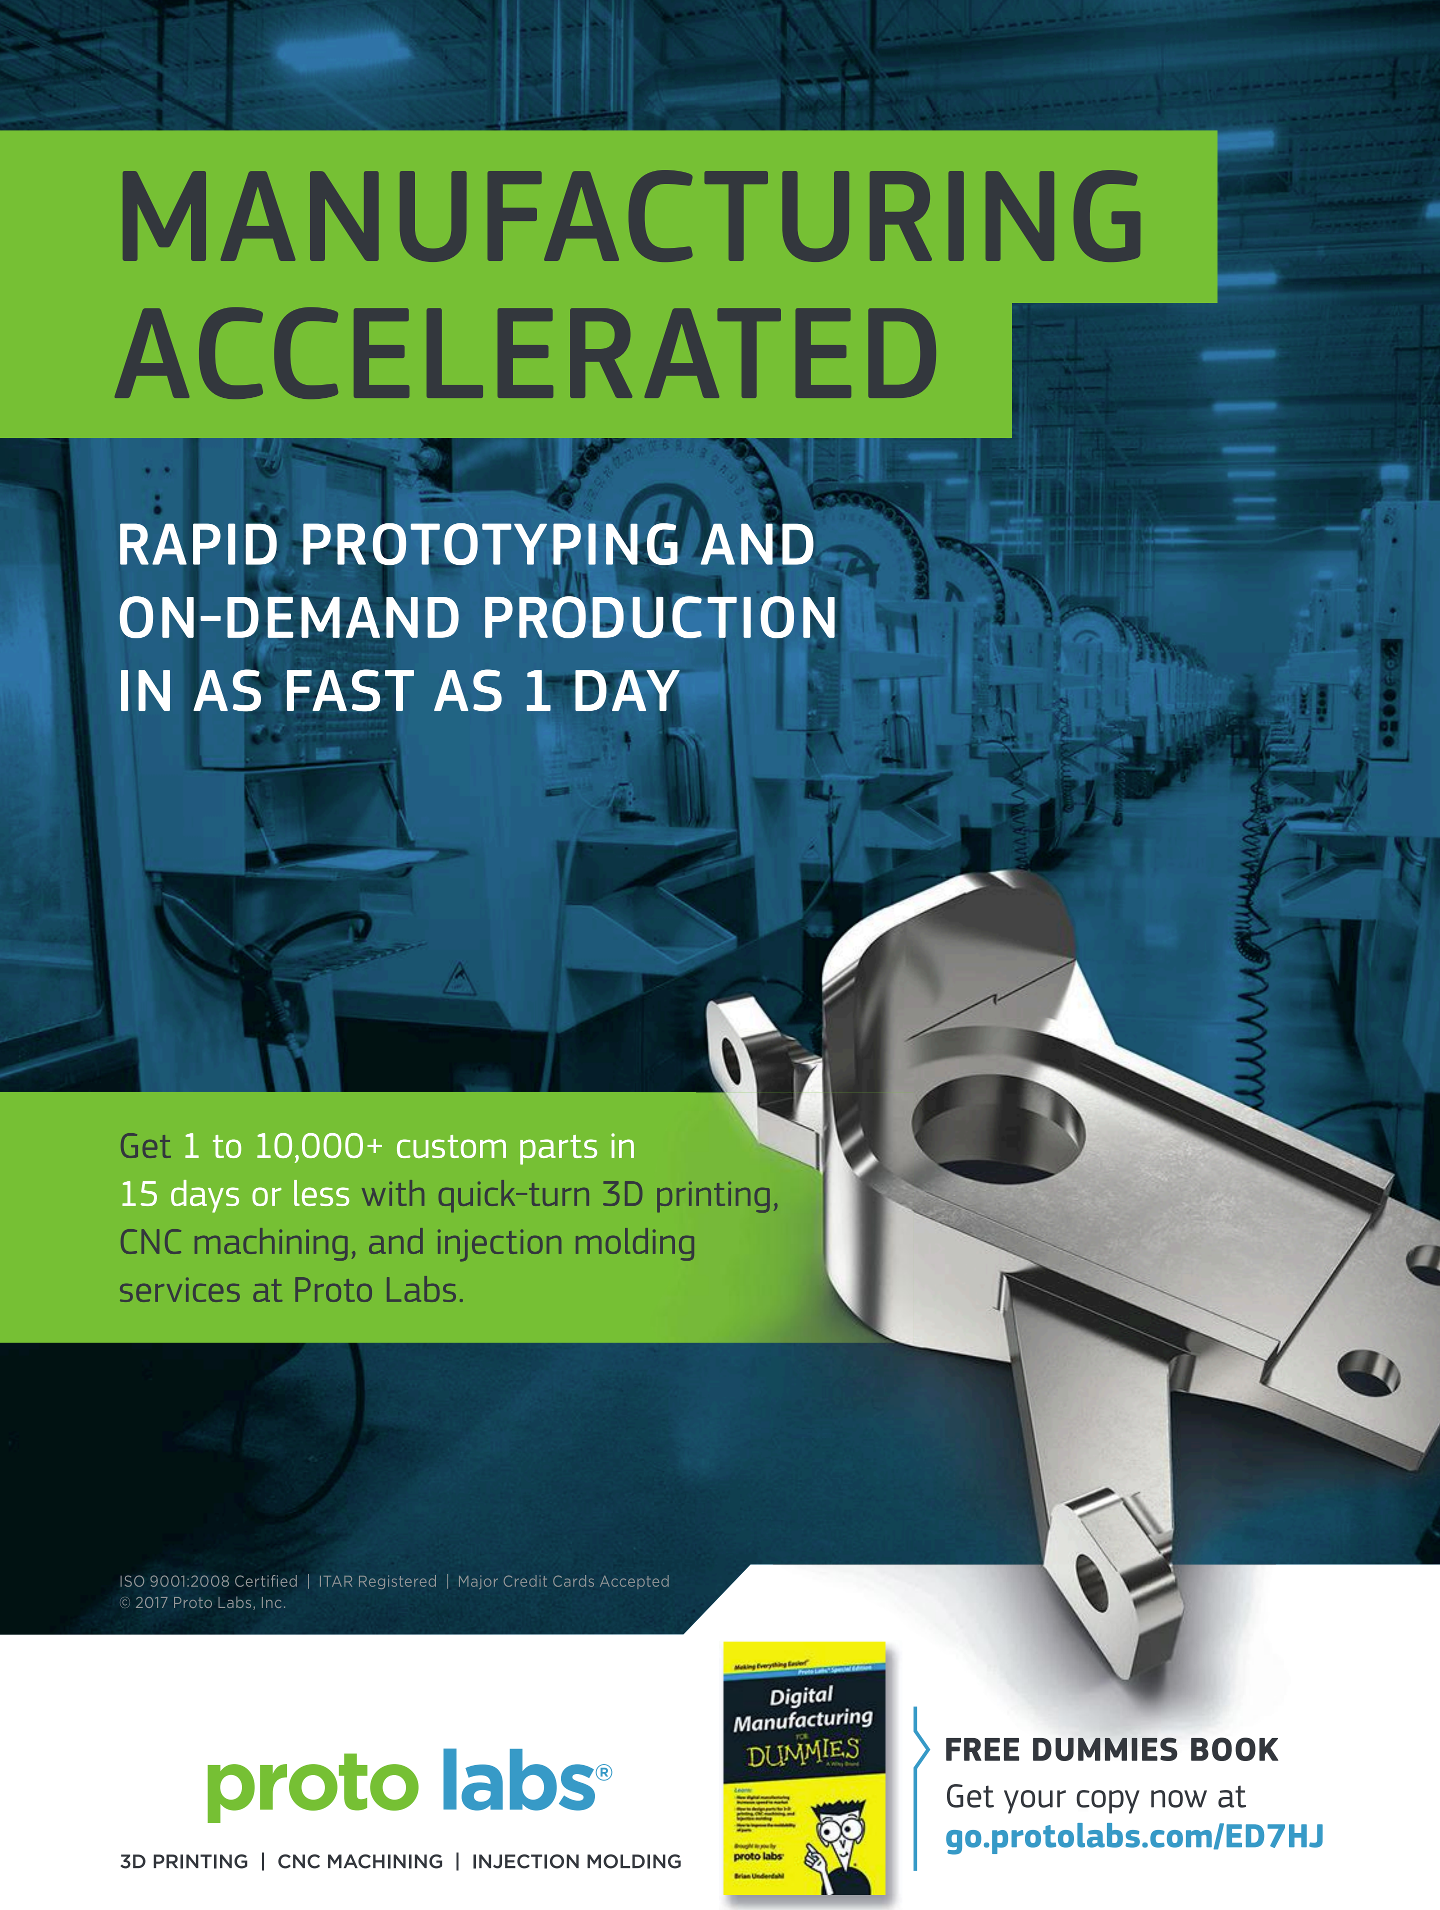
MANUFACTURING ACCELERATED
RAPID PROTOTYPING AND ON-DEMAND PRODUCTION IN AS FAST AS 1 DAY
Get 1 to 10,000+ custom parts in 15 days or less with quick-turn 3D printing, CNC machining, and injection molding services at Proto Labs.
ISO 9001:2008 Certified | ITAR Registered | Major Credit Cards Accepted
© 2017 Proto Labs, Inc.
proto labs®
3D PRINTING | CNC MACHINING | INJECTION MOLDING
Digital Manufacturing FOR DUMMIES
FREE DUMMIES BOOK
Get your copy now at
go.protolabs.com/ED7HJ

ElectronicDesign. **APRIL 2017**

EDITORIAL

CONTENT DIRECTOR: **NANCY K. FRIEDRICH** nancy.friedrich@penton.com
CONTENT PRODUCTION DIRECTOR: **MICHAEL BROWNE** michael.browne@penton.com
CONTENT PRODUCTION SPECIALIST: **ROGER ENGELKE** roger.engelke@penton.com
PRODUCTION EDITOR: **JEREMY COHEN** jeremy.cohen@penton.com
EMBEDDED/SYSTEMS/SOFTWARE: **WILLIAM WONG** bill.wong@penton.com
ANALOG/POWER: **MARIA GUERRA** maria.guerra@penton.com
CONTENT OPTIMIZATION SPECIALIST: **WES SHOCKLEY** wes.shockley@penton.com
ASSOCIATE CONTENT PRODUCER: **LEAH SCULLY** leah.scully@penton.com
ASSOCIATE CONTENT PRODUCER: **JAMES MORRA** james.morra@penton.com
CONTRIBUTING EDITOR: **LOUIS E. FRENZEL** lou.frenzel@penton.com

ART DEPARTMENT

GROUP DESIGN DIRECTOR: **ANTHONY VITOLO** tony.vitolo@penton.com
SENIOR ARTIST: **JIM MILLER** jim.miller@penton.com
CONTENT DESIGN SPECIALIST: **JOCELYN HARTZOG** jocelyn.hartzog@penton.com
CONTENT & DESIGN PRODUCTION MANAGER: **JULIE JANTZER-WARD** julie.jantzer-ward@penton.com

PRODUCTION

GROUP PRODUCTION MANAGER: **CAREY SWEETEN** carey.sweeten@penton.com
PRODUCTION MANAGER: **FRAN VAUGHN** fran.vaughn@penton.com

AUDIENCE MARKETING

USER MARKETING DIRECTOR: **BRENDA ROODE** brenda.roode@penton.com
USER MARKETING MANAGER: **DEBBIE BRADY** debbie.brady@penton.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@penton.com

REGIONAL SALES REPRESENTATIVES:
AZ, NM, TX: **GREGORY MONTGOMERY** **T** | 480.254.5540 gregory.montgomery@penton.com
AL, AK, AR, CA, CO, FL, GA, HI, IA, ID, IL, IN, KS, KY, LA, MI, MN, MO, MS, MT, NC, ND, NE, NV, OH, OK, OR, SC, SD, TN, UT, VA, WA, WI, WV, WY, W/CANADA: **JAMIE ALLEN** **T** | 415.608.1959 **F** | 913.514.3667 jamie.allen@penton.com
CT, DE, MA, MD, ME, NH, NJ, NY, PA, RI, VT, EASTERN CANADA:
SHANNON ALO-MENDOSA **T** | 978.501.7303 shannon.alo-mendosa@penton.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

PLEASE SEND INSERTION ORDERS TO: orders@penton.com
PENTON REPRINTS: **WRIGHT'S MEDIA** **T** | 877.652.5295 penton@wrightsmedia.com
LIST RENTALS/
SMARTREACH CLIENT SERVICES MANAGER: **JAMES ADDISON** **T** | 212.204.4318 james.addison@penton.com

ONLINE

PRODUCT DEVELOPMENT DIRECTOR **RYAN MALEC** ryan.malec@penton.com

DESIGN ENGINEERING & SOURCING GROUP

EXECUTIVE DIRECTOR OF CONTENT AND USER ENGAGEMENT: **NANCY K. FRIEDRICH**
GROUP DIRECTOR OF OPERATIONS: **CHRISTINA CAVANO**
GROUP DIRECTOR OF MARKETING: **JANE COOPER**

PENTON

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

Penton

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

PENTON IS AN INFORMA BUSINESS

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

SECURITY Needs to be Simple

Implementing and using security tools in the embedded space can be challenging, but Green Hills INTEGRITY Security Services' new appliance makes some security chores easier.

One of the prominent discussion points concerning the Internet of Things (IoT) involves security, and public key infrastructures (PKIs) are typically the backbone for this security. The asymmetrical public key system *(see "Crypto Essentials" on electronicdesign.com)* is used for everything from TLS network protocol stacks to blockchains *(see "What's the Difference Between Blockchains, Cryptocurrency, Audit Trails, and Databases?" on electronicdesign.com)* to secure network updates.

PKI services can run on most operating systems like Linux and Microsoft Windows Server, but only to manage keys. Chores such as signing code for secure network updates usually aren't integrated into the PKI server. Setting up this and more key-related options can be challenging. There are many vendors and projects that provide developers with customized security services—like Mender's open-source, over-the-air update client and server software that targets specific applications, such as Linux device updates—but they often focus on a particular security aspect like secure updates.

Green Hills INTEGRITY Security Services' Device Lifecycle Management (DLM) TRUST appliance addresses a number of security chores. It comes configured with its own secured hardware and software, including secure boot and encrypted storage, which is often overlooked in roll-your-own systems. It provides a web-based interface to create, manage, and deploy certificates, digitally signed certificates, and code for secure updates. There is a drag-and-drop interface where items can be signed, and it supports features like requiring multiple signatures for actions such as deploying a signed, secure update.

DLM TRUST can generate trust assets compatible with many of the secure microprocessors and controllers from vendors such as Texas Instruments, Infineon, ST Micro, NXP/Freescale, and Renesas. The key issuance service can securely distribute individual and group keys to devices, and it supports x.509 and custom certificates for authentication with protocols including like SSL/TLS, SSH, and HTTPS.

The chain of trust does not start with an IoT client. It starts with servers like DLM TRUST that manage keys and signed content, and they need to be properly secured.

IoT has many attack surfaces, and IoT security needs to address them all. Platforms like DLM TRUST are only part of the solution, but it pays to implement them well. After all, compromising them can compromise a very large collection of devices, including smartphones, tablets, and PCs. ed

News

INTEL BUYS MOBILEYE: A Chip Maker Lends Its Vision to Cars

Over the last two years, Intel seemed to have bought up enough vital parts of self-driving cars to rival Mobileye, the supplier of vision processors, software, and digital maps. And yet the company announced last month that it had acquired Mobileye for $15.3 billion.

From left, Intel's chief executive Brian Krzanich, Mobileye's chief executive Zi Aviram, and Mobileye's chief technology officer Amnon Shashua, at the Consumer Electronics Show in January. *(Image courtesy of Intel)*

In previous deals, Intel had acquired vision software from Itseez and hardware from Movidius for wiring image recognition into drones and security cameras. The $16.7 billion deal it signed with Altera two years ago gave it programmable chips that would later be used in an autonomous driving development platform.

Last year, the chip maker acquired the deep-learning expertise of Naveen Rao, Arjun Bansal, and Amir Khosrowshahi, the founders of Nervana Systems, which could be applied to the data center. Intel also acquired a 15% of the digital mapping firm Here, which is owned by BMW, Daimler, and Audi.

Intel had all the relevant parts, but it apparently lacked the time to combine them in products for the fast-moving market for autonomous cars. Intel is trying to provide not only vastly powerful computers to traditional automakers like Ford and Delphi, but also the software that could give them an alternative to tech companies like Google. And that is where Mobileye comes in, according to Intel.

Mobileye has sold vision processors and related software to over 25 major automakers for processing images from front-facing cameras, enabling safety features like blind spot warning and lane change assist. It has also expanded into digital road maps and software for fusing sensor data inside its chips.

Mobileye has also planned more sophisticated chips for fully autonomous driving on highways that it will sell to BMW, General Motors, and Volkswagen in 2018. The EyeQ5 chip will fuse data from up to 20 sensors, as well as provide a development kit for writing custom neural networks. Those programs will handle deluge of data flowing into and out of self-driving cars.

"Put just one million autonomous vehicles on the road and you have the data equivalent of half the world's population," said Brian Krzanich, Intel's chief executive, in a letter to employees. Mobileye's technology provides "the intelligent set of eyes that will allow a vehicle to see and define the world around it," he added.

The two companies already have a relationship. Earlier this year, they announced plans to build 40 autonomous cars in the United States and Europe as part of trials with BMW. They also partnered with Delphi to develop a chips package capable of handling 20 trillion operations per second.

In view of Intel's other recent deals, the move also seemed to suggest that Intel was falling behind rivals. Nvidia is raking in money for its graphics chips, which it claims are vital to running artificial intelligence programs. Qualcomm is set to become the world's biggest automotive chip with its $47 billion deal for NXP Semiconductors.

As part of the deal, Intel's automated driving group will merge with Mobileye. The company will be based in Israel and led by Amnon Shashua, Mobileye's co-founder and chief technology officer. Ziv Aviram, its other founder and C.E.O., said in a letter to employees that he would also remain at the firm.

Aviram said that the deal would give Mobileye the resources to continue growing. "By pooling together our infrastructure and resources, we can enhance and accelerate our combined know-how in the areas of mapping, virtual driving, simulators, development tool chains, hardware, data centers and high-performance computing platforms," he said in a statement.

"The saying 'What's under the hood' will increasingly refer to computing, not horsepower," said Krzanich. The deal has been approved by both Intel and Mobileye boards, but it is still subject to regulatory and shareholder approval. The deal is expected to close by the end of this year. ■

MICROSOFT EMBRACES ARM-Based Server Chips

MICROSOFT'S LONG-RUNNING PARTNERSHIP with Intel has been growing wrinkles. The company built an emulator into its Windows 10 operating system so it could run applications built for the x86 architecture pioneered by Intel. And in March, it complicated Intel's future in its growing cloud business.

Microsoft, which is based in Redmond, Wash., said that it would run several internal cloud services with chips based on ARM designs, which are also used in the vast majority of smartphones and inside gadgets like thermostats and sensors embedded in manufacturing plants.

The move amounts to one of the most significant endorsements of ARM servers, showing that the architecture has the performance and standardized code base to potentially encroach on Intel's dominant position in the market for server chips, analysts said. The other large cloud computing companies, Google and Amazon, plan to use Intel's next-generation Skylake chips.

Microsoft reworked its Window Server operating system to run over the ARM servers, which will be used internally for applications in search, storage, databases, and machine learning. The announcement came at the Open Compute Platform Summit in Santa Clara, Calif.

"We feel ARM servers represent a real opportunity and some Microsoft cloud services already have future deployment plans on ARM servers," wrote Leendert van Doorn, a distinguished engineer at Azure, Microsoft's cloud computing unit, in a blog post.

Microsoft also signed design partnerships with Intel, Advanced Micro Devices, and Nvidia, underlining how many different types of chips will ultimately live together in data centers.

Qualcomm and Cavium both devised Windows Server hardware based on ARM technology. Qualcomm built a server using its Centriq 2400 server chip with 10 nanometer technology and 48 computing cores. Qualcomm is giving out a small number of samples for testing, but the chip will start selling later this year.

It has been a bumpy road to legitimacy for ARM server chips. Since first appearing five years ago, sales of the chips from companies like AMD, Applied Micro, and Broadcom have fallen flat. The hardware has struggled to match the standard code base and performance offered by Intel processors, analysts said.

Paul Teich, an analyst with Tirias Research, said that Microsoft's announcement shows that the latest batch of ARM server chips were able to meet those conditions. The move, he added, also shows how the industry is craving more competition–an idea echoed by ARM itself.

"Faced with different requirements for workload optimization, data center architects have been calling for increased choice and flexibility in server architectures," ARM said in a statement. "There is more opportunity now to increase dynamic competition in the server market."

ARM chip makers are still facing an uphill battle to loosen Intel's grip on the server market. The research firm IDC estimates that the x86 instruction set had been used in over 99% of all servers sold in the first three quarters of 2016. Intel reaped $17.2 billion in revenue from its data center business last year.

One stumbling block is that software written for x86 chips must be reworked for ARM hardware. But Microsoft's van Doorn said that the ability to carefully match hardware and software would pay off. He said it is "economically feasible to optimize the hardware to the workload instead of the other way around." ■

PREMIUM PROCESSING FOR H1-B VISAS Suspended for Six Months

THE UNITED STATES CITIZENSHIP and Immigration Services said that it would suspend premium processing for H1-B visa applications, which allow thousands of foreign engineers and computer scientists to work in the United States for companies that sponsor them.

It usually takes several months for an H1-B application to be processed and entered into the random lottery that allots the work permits. But for a $1,295 filing fee, USCIS will decide on the application within 15 days. Starting April 3, the premium processing option will be thrown out for six months.

The USCIS said that it stopped offering the service so that it could get through applications that had piled up in recent months. The move could create headaches for foreign workers applying for H1-Bs the first time, as well as visa holders changing jobs or seeking an extension. Though a small move, it could also hint at coming changes to the H1-B program.

This year, the U.S. distributed around 85,000 of the visas out of almost a quarter of a million applications. The H1-B program is widely used to bring top engineering talent to large U.S. technology firms, but it has drawn criticism for cutting off jobs for American engineers. The biggest H1-B recipients are foreign outsourcing firms.

The program has also been targeted by the Trump administration, which also recently issued an order to stop issuing new visas to people from countries including Iran and Syria. In January, a leaked draft of an executive order indicated that the White House intends to review the H1-B application process to ensure that only "the best and brightest" were approved.

President Trump has said that he favors overhauling the H1-B visa program, but has been short on specifics. Lawmakers and groups like the IEEE-USA have proposed getting rid of the H1-B lottery, giving priority to firms that pay higher wages. That approach could leave more entry-level jobs to American engineers, they say.

The White House missed its deadline to issue a notice on changing the lottery process in March. Whether that deadline could be circumvented with an executive order is unclear. But some experts believe that the suspension is meant to buy time to figure out which H1-B applications will get priority.

The USCIS has suspended the premium processing option in recent years. In 2015, the agency shut down the program to prepare for a flood of applications that would give employment eligibility to spouses dependent on H1-B visa holders. That suspension only lasted less than two months.

The citizenship and immigration agency says that the recent suspension will help it "process long-pending petitions, which we have currently been unable to process due to the high volume of incoming petitions and the significant surge in premium processing requests over the past few years." ■

Advertisement

LEGAL NOTICE

IF YOU BOUGHT AN ELECTROLYTIC OR FILM CAPACITOR DIRECTLY FROM CERTAIN DISTRIBUTORS SINCE 2002 YOU COULD GET MONEY FROM SETTLEMENTS TOTALING APPROXIMATELY $15 MILLION

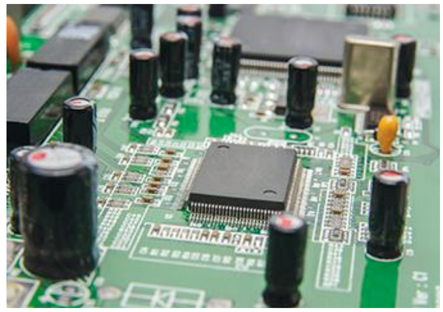

NEC TOKIN Corp. and NEC TOKIN America, Inc. ("NEC TOKIN"); Okaya Electric Industries Co., Ltd. ("OEI"); and Nitsuko Electronics Corporation ("Nitsuko") (collectively "Settling Defendants") have agreed to Settlements resolving claims that they allegedly fixed the prices of Capacitors. This may have caused individuals and businesses to pay more for Capacitors. Capacitors are electronic components that store electric charges between one or more pairs of conductors separated by an insulator.

AM I INCLUDED?

You may be included if, from January 1, 2002, through July 15, 2016, you purchased one or more Capacitors from a distributor (or from an entity other than a Defendant) that a Defendant or alleged co-conspirator manufactured. "Indirect," as that term is used below, means that you bought the product from someone other than the manufacturer, for example, from a distributor. A more detailed notice, including the exact Class definitions and exceptions to Class membership, is available at www.capacitorsindirectcase.com.

WHAT DO THE SETTLEMENTS PROVIDE?

The Settlements provide for the combined payment of $14,950,000 in cash to the Classes. The Settling Defendants have also agreed to cooperate in the pursuit of claims against other Defendants.

HOW CAN I GET A PAYMENT?

Money will not be distributed to the Classes at this time. The lawyers for the Classes will pursue the lawsuit against the other Defendants to see if any future settlements or judgments can be obtained in the case and then be distributed together to reduce expenses. The lawyers anticipate that when the settlement proceeds are disbursed to the Classes, it will be done on a *pro rata* basis based on the value of your Capacitor purchases.

If you want to receive notice about the claims process or future settlements, you should register at www.capacitorsindirectcase.com.

WHAT ARE MY RIGHTS?

Even if you do nothing, you will be bound by the Court's decisions concerning these Settlements. If you want to keep your right to sue one or more of the Settling Defendants regarding Capacitor purchases, you must exclude yourself in writing from the Classes by May 31, 2017. If you stay in the Classes, you may object in writing to the Settlements by May 31, 2017. The Settlement Agreements, along with details on how to exclude yourself or object, are available at www.capacitorsindirectcase.com. The U.S. District Court for the Northern District of California will hold a hearing on July 6, 2017, at 10:00 a.m., at 450 Golden Gate Avenue, 19th Floor, Courtroom 11, San Francisco, CA 94102 to consider whether to approve the Settlements. Class Counsel may also request at the hearing, or at a later date, attorneys' fees of up to 25% of the Settlement Funds, plus reimbursement of costs and expenses, for investigating the facts, litigating the case, negotiating the Settlements, providing notice to the Classes, and/or claims administration. The total amount of these costs shall be no more than $2,558,454.00. You or your own lawyer may appear and speak at the hearing at your own expense, but you don't have to. The hearing may be moved to a different date or time without additional notice, so it is a good idea to check the above-noted website for additional information. Please do not contact the Court about this case.

If the case against the other Defendants is not dismissed or settled, Class Counsel will have to prove their claims against the other Defendants at trial. Dates for the trial have not yet been set. The Court has appointed the law firm of Cotchett, Pitre & McCarthy, LLP to represent Indirect Purchaser Class members.

FOR MORE INFORMATION: 1-866-217-4245 | WWW.CAPACITORSINDIRECTCASE.COM

WILLIAM WONG | Embedded/Systems/Software Editor

Embedded Design: Build, Buy, or Both?

How you choose to build an embedded system depends on a lot of factors, starting with whether you make or buy the hardware.

It has never been easier to build an embedded system, whether one makes or buys the hardware. Of course, there are many ways to do this.

The choice depends upon a large number of factors from in-house expertise to time-to-market requirements. Factor in size, mobility, and other variables like quantity and the approach may become clear. For example, delivering millions of inexpensive, wearable devices may warrant a custom chip design whereas a couple hundred expensive boxes might be done using COTS (commercial-off-the-shelf) or even standard systems.

Here we will focus on custom approaches starting with the lowest level of customization including systems built from motherboards, modules, and custom boards, and trends like peripheral interfaces on tiny Mini-PCIe boards *(Fig. 1)*. Developers can also turn to third parties for all these options so developers do not need expertise in all areas. There is a price to farm out parts of a design, but there are also costs for doing it in-house.

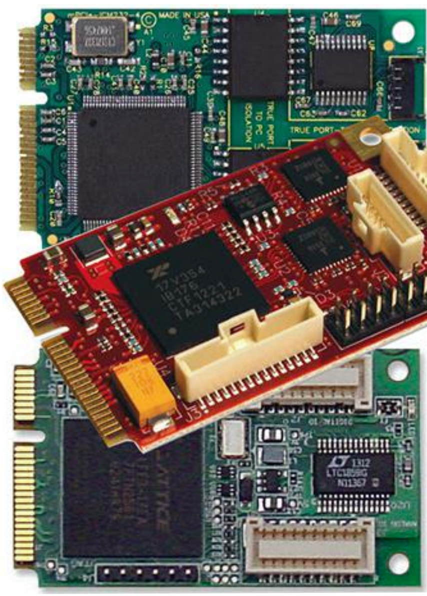

1. Mini-PCIe modules are not just for adding wireless support to laptops. Mini-PCIe sockets can populated with modules like (a) ACCESIO's mPCIe-ICM232-4 quad serial port with 1500 V isolation, (b) VersaLogic's Serial + GPIO, and (c) Diamond System's DS-MPE-DAQ0804 Analog I/O Module.

BUILDING WITH BOARDS

Building a board-level system typically starts with a motherboard that usually provides one or more expansion options. There are also backplane-based systems like CompactPCI and VME and their high-speed serial counterparts CompactPCI Serial and VPX *(see "What's the Difference Between VME and VPX?" on electronicdesign.com)*. These provide expansion using 3U and 6U boards. There are other board-level standards such as Advanced TCA and Micro TCA, as well as test and measurement standards that add features like timing.

The PXI System Alliance's PXI (PCI eXtensions for Instrumentation) and PXIe are variants of CompactPCI and CompactPCI Express.

Conduction and liquid cooling are sometimes options for more rugged environments in addition to air cooling. Boards may also have headers for mezzanine cards using standards like PMC, XMC, and FMC.

PC motherboards like ATX, Micro-ATX, and Mini-ITX provide a similar expansion approach but the add-on card mountings tend to be less rugged and they are typically air-cooled.

Smaller form factor systems are available as well, although the expansion options tend to be more limited. Form factors like Pico-ITX have become popular because of their small size and availability of enclosures that are commonly available for the larger form factors such as the PC motherboards.

Intel's Next Unit of Computing (NUC) is another form factor that is popular in many embedded applications such as digital signage *(see "What are the Differences Between Consumer and Industrial NUCs?" on electronicdesign.com)*. NUCs tend to be very compact and this can limit expansion options. NUCs like Logic Supply's ML100G NUC *(Fig. 2)* target rugged industrial applications.

Stackables are an alternative to the motherboard and backplane approach to building systems. The quintessential stackable form factor is PC/104, of which there are many varia-

2. Logic Supply's ML100G is a rugged NUC system designed for industrial applications.

tions. Typically, a system has a single board computer (SBC) somewhere in the stack, usually at one end, with additional peripheral cards stacked along with it. Connectors on each board connect to the board above and below, if they exist.

The PC/104 Consortium's PC/104 standard is based on the parallel ISA (industry standard architecture) bus. It is easy to design an ISA interface device, but it is slowly disappearing because of a number of factors including its slower speed and lack of support with new chips. PC/104 Plus added the parallel PCI bus to the mix.

Of course, PCI Express (PCIe) is now the dominant interface and it is a point-to-point connection instead of the bus interface of ISA and PCI. This requires that stacking systems employ a staggered connection as signals move from one board to another *(Fig. 3)*. For example, an SBC may provide four x1 PCIe lanes on a connector. An expansion board uses the first lane and shifts the other three before passing the signals onto the next board in the stack. The fourth connection is no longer in play and the stack can handle up to four boards that use a PCIe lane.

PC/104 variations that employ PCIe are defined. The OneBank *(see "Plenty of Platform Choices for Embedded Designers" on electronicdesign.com)* has four x1 PCIe lanes, two USB 2.0, and power signals on a single connector. USB is also a point-to-point connection and the staggered connection approach works just as well with USB as it does with PCIe. The Micro/Sys StackableUSB takes the same approach, but only uses USB on a very small form factor board that is a quarter the size of a PC/104 board.

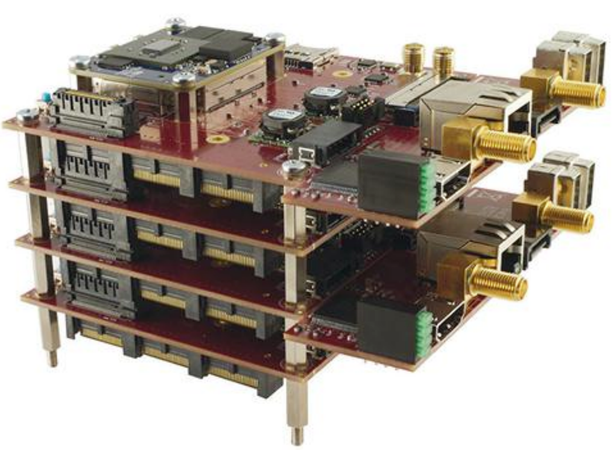

3. Sundance Multiprocessor Technology's stack is an example of the PC/104 form factor that utilizes the OneBank connectors.

Another approach to expansion used on everything from motherboards to SBCs are Mini-PCIe and M.2 sockets. These can be used for storage and peripherals, although typically M.2 is being used for flash memory with either SATA or PCIe interfaces *(see "What's the Difference Between M.2 Modules?" on electronicdesign.com)*. Mini-PCIe sockets can support a number of interfaces including PCIe, USB and SATA. The latter is used to implement the mSATA interface for storage.

The Mini-PCIe sockets were initially used for communication options on laptops but have since been used to support all sorts of peripherals from ADCs and DACs to motor controls. Even wired network interfaces can be found on these tiny cards *(see "PCI Express Mini Card Tackles Compact Embedded Expansion" on electronicdesign.com)*. Many SBCs are being designed with multiple Mini-PCIe sockets. The challenge tends to be the connectors since the boards are tiny. Luckily, tiny connectors are available as well.

BUILDING WITH MODULES

Often SBCs and other standard board-based solutions are not the answer as custom connections or peripherals are involved, but developers do not want the hassle of designing a system from scratch. In this case, a module approach may be more appropriate.

There are dozens of standard form factor modules like COM Express, SMARC, and QSeven *(Fig. 4)* and even more proprietary ones. They also vary significantly in size, performance, and cost so there is likely one suitable for most applications.

COM Express is a computer-on-module (COM) standard from PICMG. The modules support PCI Express Gen 3, USB

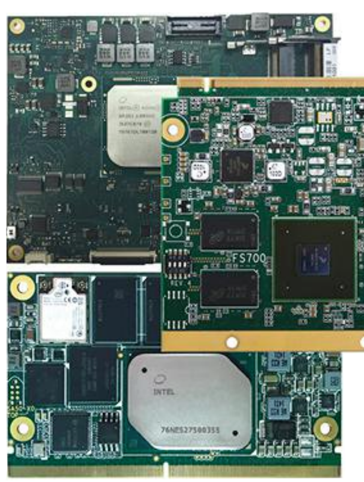

4. COM modules like (a) ADLINK's Express-BD7 COM Express board supports 10 Gbit Ethernet, (b) Congatec's conga-SA5 SMARC that supports the Intel Atom, and (c) DFI Technologies' QSeven module that is available for Intel Atom and NXP i.MX6 processors.

3.0, SATA, LVDS, CAN, HDMI, and Ethernet. The latest supports 10 Gbit Ethernet and 32 PCIe lanes *(see "COM Express Modules Handle 10 G Ethernet" on electronicdesign.com)*. The smallest Mini COM Express module is only 84 mm by 55 mm.

SMARC and QSeven are two COM from the Standardization Group for Embedded Technologies (SGET). SMARC is a 82 mm by 50 mm board designed for ultra-low power COM (ULP-COM). QSeven is 40 mm by 70 mm and plugs into an MXM connector.

Of course, there are many alternatives that are smaller like the Gumstix series of 17 mm by 58 mm ARM-based modules. Digi International has a number of modules that target the Internet of Things (IoT) in addition to some RJ-45 jacks that contain their own network server. Linear Technology's SmartMesh modules implement a mesh network that is used in industrial IoT (IIoT) settings. There are even modules that target wearables.

And the list goes on and on. There is even a Raspberry Pi Compute Module 3 *(Fig. 5)* that can be used to deploy Raspberry Pi applications, as the chip is only available in very large quantities.

Modules can deliver very high performance like NVidia's Jetson TX1 and new TX2 module that delivers 2 TFLOPs of performance and targets artificial intelligence and machine learning with support for deep neural networks (DNN). The platform is ideal for vision applications in industrial settings.

Modules with wireless support also have advantage over custom designs because they will already have the necessary approvals for wireless operation.

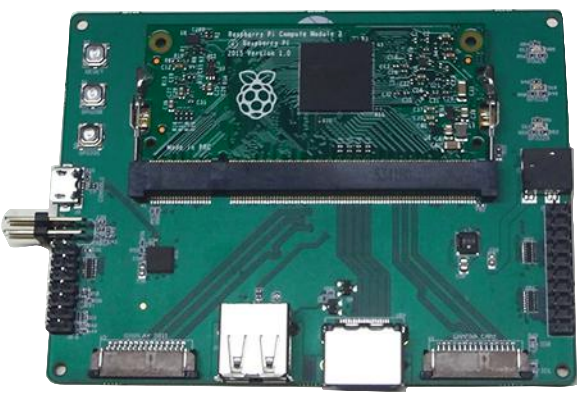

5. This Raspberry Pi Compute Module 3 is mounted in Gumstix Pi Compute Dev Board created using its Geppetto web-based design-to-order (D2O) system.

CUSTOM BOARDS

Modules do not operate as standalone devices like SBCs so typically a module is plugged into a carrier board. Most modules have development carrier boards available but many have carrier boards targeted at deployment. For example, Connect Tech sells a line of carried boards for the Jetson TX1 and TX2 *(see "Carrying a Jetson TX1" on electronicdesign.com)*. These have connectors for interfaces like USB and Ethernet not found on the module. Pixevia's Core X1 *(Fig. 6)* is a carrier board for the Jetson that targets drones and vision applications. It includes a 9-axis sensor system including 3D accelerometer, 3D gyro, and 3D compass. It also has a barometer, CSI camera ports plus CAN support, and is designed to support GPS and long range wireless.

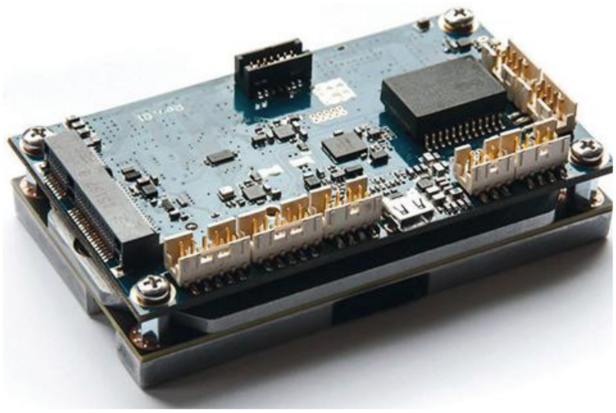

6. Pixevia's Core X1 supports the Jetson TX1 and TX2. It targets drone and video applications with on-board sensors and camera interfaces.

The alternative to off-the-shelf solutions is to design your own carrier board. This is usually easier than designing a board that also incorporates the processor, especially for higher-performance systems. This can be done with PCB design tools. One alternative for a select group of modules is Gumstix's design-to-order (D2O) Geppetto service (see "Best of 2015: Create Custom Capes Fast and Easy" on electronicdesign.com). It supports Gumstix's modules as well as Intel's Joule and Edison, Toradex's Colibri, and more. The web-based design tool uses a drag-and-drop interface that hides the routing details and requires all connections to be made before the design is complete. It has a $1,999 setup fee and a per-board fee based on the design with delivery of finished boards in days.

There are engineering firms that will provide more advanced design and manufacturing services. Many developers may also do board design in-house and even manufacturing. This requires more design and production expertise. The advantage is that almost any component can be included and the boards can meet more exacting standards in performance and construction in addition to providing a lower per board cost.

If board production is all that is required then there are a number of fast-turnaround PCB firms that can deliver boards in a day. Just upload the design files and place an order. Many, like Advanced Circuits, Sunstone Circuits, Imagineering, and ExpressPCB provide free CAD PCB tools that integrate with the websites. These firms often work with or provide complete board construction.

In-house PCB construction using conventional techniques tends to be expensive and hazardous for the inexperienced, but there are alternatives for prototyping and small production runs that may be useful to developers. These include additive and subtractive devices for creating PCBs.

Nano Dimension's DragonFly 2020 *(Fig. 7)* uses an additive 3D inkjet printer technology create PCBs *(see "Q&A: A Behind-the-Scenes Look at the DragonFly 2020 3D Printer" on electronicdesign.com)*. It can print multilayer boards and the ink and process used to create the boards is designed to replicate the electrical characteristics of a board created using conventional techniques.

LPKF Laser and Electronics has a number of subtractive devices that are essentially milling machines for creating PCBs *(see "Build Your Own PC Board: A Designer's Dream" on electronicdesign.com)*. The system starts with a PCB and then grinds off unwanted copper. A two-layer board requires a flip while multilayer boards are done a layer at a time.

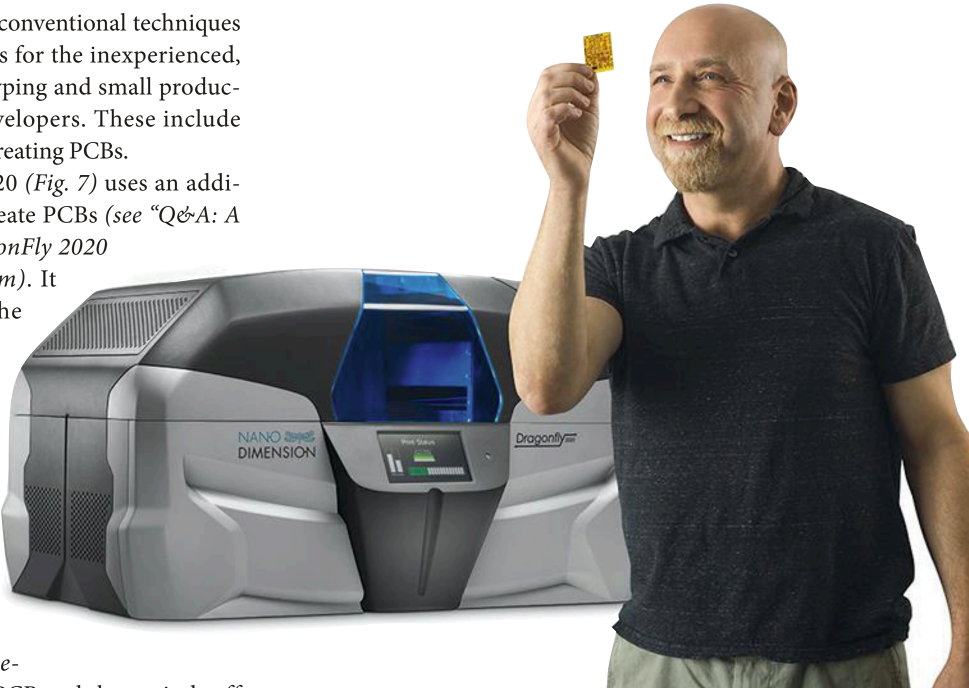

7. Nano Dimension's DragonFly 2020 is a 3D printer that generates multilayer PCBs.

CUSTOM CHIPS

Custom boards will often meet a designer's needs, but sometimes a custom chip may be required. The switch to chip foundries has opened the custom chip design options to a much wider audience, but the approach can still be daunting especially as the size of the chip grows making this alternative viable only for those with deep pockets.

On the other hand, chip designers have never had better tools and intellectual property (IP) options. Chip design is still not for the uninitiated, but there are engineering firms that will take a design and turn it into a chip. Still, it has many advantages such as making cloning more difficult.

An alternative to a fully custom chip are configurable chips. The simplest case is a programmable logic device while more complex designs require FPGAs. Designers do not have to drop to the gate level, but even logic designs for FPGAs can be quite complex. Luckily these tools have progressed significantly and are employing techniques like using IP blocks connected to standard interfaces allowing drag-and-drop designs to turning program code like C and C++ to be turned into designs automatically. There are tools that will convert graphical designs like those done in National Instrument's LabView into FPGA code.

FPGAs are also available with hard-core logic like microprocessors and microcontrollers with system-on-chip (SoC) functionality. These are tied into the FPGA fabric providing a mix of logic and software support that is applicable to a wide range of options. It is also possible to use soft-core processors like the ARM Cortex-M1 and RISC-V *(see "RISC-V (Five) Is Alive!" on electronicdesign.com)*.

FPGAs can be used for prototyping with resulting custom chips being less expensive and more power-efficient, but in many cases FPGAs have proven to be better than custom chips for a variety of reasons including long term support.

One option for microcontroller solutions that falls between standard chips and FPGAs is Cypress Semiconductor's PSoC. The latest PSoC 6 is a dual-core, asymmetric design with a Cortex-M4 and Cortex-M0+ *(see "New PSoC 6 Sports Dual, Asymmetric 32-bit Cores with Enhanced Security" on electronicdesign.com)*. The PSoC has configurable digital and analog peripherals that have a functionality similar to an FPGA but with restrictions that simplify their operation and configuration. Most developers work with predefined peripherals allowing the chip to appear as a typical microcontroller with an arbitrary selection of interfaces. More advanced designers can create custom peripherals from the digital and analog blocks.

It has never been easier to build an embedded system whether for prototyping or production but the choices are vast. Application form and function will limit the available options, but rarely will an application be limited to one approach or solution, although one alternative may be much better than the others. ed

MARIA GUERRA | Technical Editor

The Advantages—and Challenges—of ENERGY HARVESTING

In a recent interview with Maurizio Di Paolo Emilio, the author and telecommunications engineer talks about energy harvesting technology and what to expect in this space.

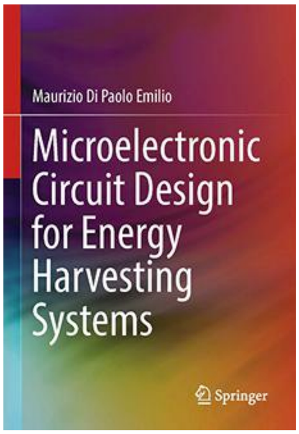

Maurizio Di Paolo Emilio is the author of Microelectronic Circuit Design for Energy Harvesting Systems, *a book covering the design of microelectronic circuits for energy harvesting, broadband energy conversion, and new methods and technologies for energy conversion. Di Paolo Emilio recently talked with* Electronic Design *about energy-harvesting technology, its advantages, and its challenges, as well as current trends and future predictions for this space.*

How important is the role of energy-harvesting in IoT?

In recent years, much attention has been placed on the Internet of Things (IoT), or a whole range of commercial and industrial devices (IIoT) interconnected by wireless (and wired) protocol. Analysts estimate an increasing number of IoT devices—about 2 billion over the next five years. The advent of these devices poses a serious powering problem: the batteries that have to be purchased, maintained, and disposed of.

The energy harvesting technique is a simple solution to recharge low-power devices easily and economically, at the same time, by using clean energy. A fundamental requirement for IoT is the power management: Mobile devices obviously require batteries, but the possibility to replace them completely or restrict the replacement/recharge is a considerable factor of importance, given the advent of further devices connected in the near future.

Energy harvesting technologies can help. They use electric power generation elements such as solar cells, RF sensors, and piezoelectric and thermoelectric elements for converting light and vibration. Furthermore, when using energy harvesting, there is a point to be considered—a balance between generation and energy consumption. This is because the device does not work if power generation is less than the required power.

What's the advantage of microcontrollers in energy-harvesting systems?

Integrated power management circuits designed for energy harvesting, as well as low-power MCUs, will help the growth of the Internet of things. Energy harvesting is a valid (but not a simple or easy) option for the IoT recharge design.

Although the generation of characteristics of the energy production elements is improving from year to year, and microelectronics continues the step of an ultra-low power management, it is difficult to continuously provide enough power for a device on an ongoing basis, and the need of a collection technique in a first phase could help.

The progress in microcontrollers with low consumption decidedly opens the target for products and applications in which a collection of light, kinetic energy, or heat can be used to power an intelligent product without an external power supply or battery replacement. The microcontrollers have a fundamental role in all electronic devices. This improved fuel consumption and power management methods are a design factor to allow easy adoption of the energy harvesting.

What's your take on powering microsystems? Why are they important? What are the most common wrong assumptions made by energy-harvesting designers?

The power conditioning circuits play an essential role in an energy harvesting system through various parameters (such as the input impedance), at the same time carrying out processing functions (such as power control and filtering). Advanced techniques actively influence the behavior of harvesting devices (such as piezo pre-biasing). The power limit of a system to be used was considerably reduced with conditioning circuits that operate at lower levels of power, by reducing the losses to increase the maximum efficiency of the harvesting system.

Power supplies are often intermittent and the excitation parameters may change over time. The purpose of the conditioning circuit is to avoid an oversized design, with a storage system for providing a correspondence between the temporal profiles of the power demand from the load source.

The challenge is always to optimize the energy and the associated conditioning circuits to cope with a system where the correspondence of the power profiles and operation dynamics hours is in some way optimized.

What are the major limitations of energy-harvesting systems nowadays?

One of the many challenges for the designers is to assess the energy factor, estimating the energy requirement and designing its power configuration. The fundamental requirement of IoT is power management. In many application scenarios of low-power devices, it is difficult to obtain a continuous active mode of energy, resulting not only in electronics but also the environmental situations (low light). Then a method of accumulation, even within the tiny size—smaller than the general cases—is still necessary to ensure the proper functioning of a device. The recent trend is to the replace the rechargeable batteries with super capacitors characterized by charge-discharge "cycles unlimited" (> 100,000).

What are some of the energy-harvesting applications available in the market?

Energy harvesting makes use of ambient energy to power small electronic devices such as wireless sensors, microcontrollers, and displays. Typical examples of these environmental sources are sunlight and any artificial source such as vibration or heat from engines or the human body. The energy transducers such as solar cells, thermogenerators, and piezoelectrics convert this energy into electrical energy. A first field of application is the automation with self-powered switches. Further applications are the monitoring systems for large industrial plants or structural monitoring of huge buildings. Another promising market is the consumer area with purses and clothing, which show the energy transducers integrated in the form of solar cells or TEG or RF transmitters to recharge consumer products such as mobile phones or audio players.

The application of energy harvesting techniques in the railway sector is a very promising field. They conducted studies and experiments on the potential of energy recovery devices, with the aim to provide information on the electric power actually generated by the use of harvester placed on board railway wagons.

One of the emerging market segments covered under the IoT is the wearable electronics category. Regardless of the application, most of these devices require a battery as the main power source. The rapidly expanding of wearable devices is set to revolutionize the technology and current processes in all areas by creating new market opportunities and new business models.

What are the most promising energy-harvesting technologies? How about RF wireless power technology?

Solar radiation has the advantage to be a green energy without the production of polluting waste. The limits that may be encountered are discontinuities caused by the alternation of day and night and weather conditions. The other is the low intensity, which implies the need to have large areas of energy storage.

Many quantities of energy harvesting devices are currently sold in the building automation industry and in the consumer market. In the consumer market solar cells represent the main choice, but more solutions in terms of vibration and RF are entering the market. The development of IoT promises to be exciting and have a great economic impact due to the development of semiconductor devices, and to the advancement of wireless technology, by allowing devices that are smaller and more efficient at the same time.

The RF solutions represent a harvesting technique to permanently remove the batteries from smartphones and ensure the auto-recharge by RF sources themselves. The wireless charging technology has demonstrated perfectly how it is possible, but this time you have to consider the possibility of gaining energy from the environment. With the advent of IoT, we are bombarded by many secondary RF sources, and then we will have more chances to get enough electrical current from the surrounding environment. The power supply is definitely the future of energy harvesting!

MARIA GUERRA | Technical Editor

Voltage-Regulator ICs for Wearables: How High Is Your I_Q?

Steadily evolving switching regulators are a good option for designers tasked with creating smaller and more efficient solutions to achieve low quiescent current in applications that must be energy-efficient.

According to research firm MarketsandMarkets, the wearable technology market is expected to grow from USD $15.74 billion in 2015 to $51.60 billion by 2022, which translates into a CAGR of 15.51% between 2016 and 2022. Wearable devices contain, among other components, voltage regulators of the switching or low-dropout (LDO) variety. No doubt, then, that the voltage-regulator market is bracing for growth as well.

It's quite common to find switching-regulator ICs in power-supply designs for wearable applications. Designers find them more flexible because they can achieve peak efficiency in the 90th percentile. Although switching regulators generate more noise than linear regulators, their efficiency is far superior as well.

Some existing power solutions combine LDOs and switching regulators, while others just use switching regulators. For example, Intersil's ISL9120 *(Fig. 1)* operates in either buck or boost mode with a pulse-frequency-modulator (PFM) controller that automatically alternates between buck, boost, and automatic bypass modes to maintain a steady output voltage. The PFM controller helps attain up to 98% efficiency at higher load and greater than 86% at lower load conditions. In doing so, it achieves less power drain with less heat, which in turn extends battery life.

During stay-alive operation, the ISL9120 goes into forced bypass mode, which reduces power consumption to a quiescent current of less than 0.5 µA. To illustrate a scenario where forced-bypass-mode operation would come into play, think of powering an LDO that may be in standby mode with near-zero output current. When the buck-boost is in bypass mode, it saves the regulator 41 µA of quiescent-current consumption.

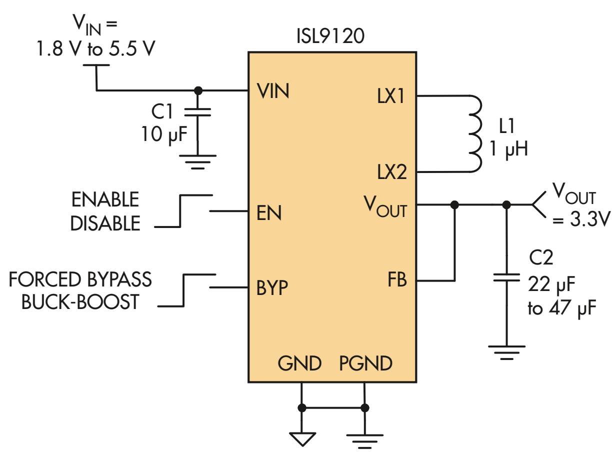

1. Only a single inductor and very few external components are needed for the ISL9120 buck-boost regulator. *(Courtesy of Intersil)*

KNOW YOUR I_Q

The forced-bypass mode is just one design option that may be used when designing wearable electronics. Among other aspects that need to be considered is the current consumed by the device in different power modes, like standby mode and shutdown.

The quiescent current (I_Q) is very important, because wearable devices spend a lot of time in standby mode. I_Q is the current consumed when a circuit is not driving a load. Therefore, the lower the value of the I_Q, the more you can extend the life of a battery.

Designers have been trying to achieve better energy efficiency in wearable devices by lowering I_Q values. Several solutions are available in the marketplace, with I_Q values con-

(Continued on page 46)

pico®
Technology

THE PICOSCOPE 4444

SEE THE DIFFERENCE

NEW

A NEW STANDARD IN DIFFERENTIAL MEASUREMENT

- 20 MHz bandwidth, 12 - 14 bit resolution
- 4 fully differential inputs
- 1000 V CAT III probes
- Low voltage probes and current clamps

With four true differential inputs, 12- to 14-bit resolution and wide differential and common-mode voltage ranges, the PicoScope 4444 and its accessories offer accurate and detailed measurement for a multitude of applications, from low-amplitude biomedical and electronic uses to 1000 V CAT III design and test.

For more information please visit www.picotech.com/US105

LOU FRENZEL | Contributing Editor

Major Trends in Wireless Test

The 5G, IoT, Wi-Fi, and automotive markets are driving wireless instrument development.

You've got to applaud the test-equipment manufacturers for their pioneering work in developing new wireless technologies. Were it not for all of their advanced work in defining new standards, test procedures, and test gear, product R&D wouldn't be possible. If you can't test and measure it, you can't design it. Thankfully, the test companies are doing a great job in meeting the needs of wireless designers. Here is a look at the current trends in this area, as well as some of the latest equipment available.

HOT NEW TRENDS INFLUENCING TEST

As you probably know, the key trends motivating wireless test development are 5G cellular technology, the Internet of Things (IoT), the ever-evolving Wi-Fi, and the growing automotive market. 5G development is well under way, but the standards have not been finalized. While field trials are ongoing, there are more years still to go before we see the finished standard and initial deployments.

The primary challenges of 5G center around millimeter-wave frequencies. In the U.S., the 27-, 37-, and 39-GHz bands have been assigned. New waveforms are being defined, creating a need for a flexible signal source—not to mention, the usual vector signal analyzer. What's required is a way to test beamforming antennas and higher degrees of MIMO in different configurations.

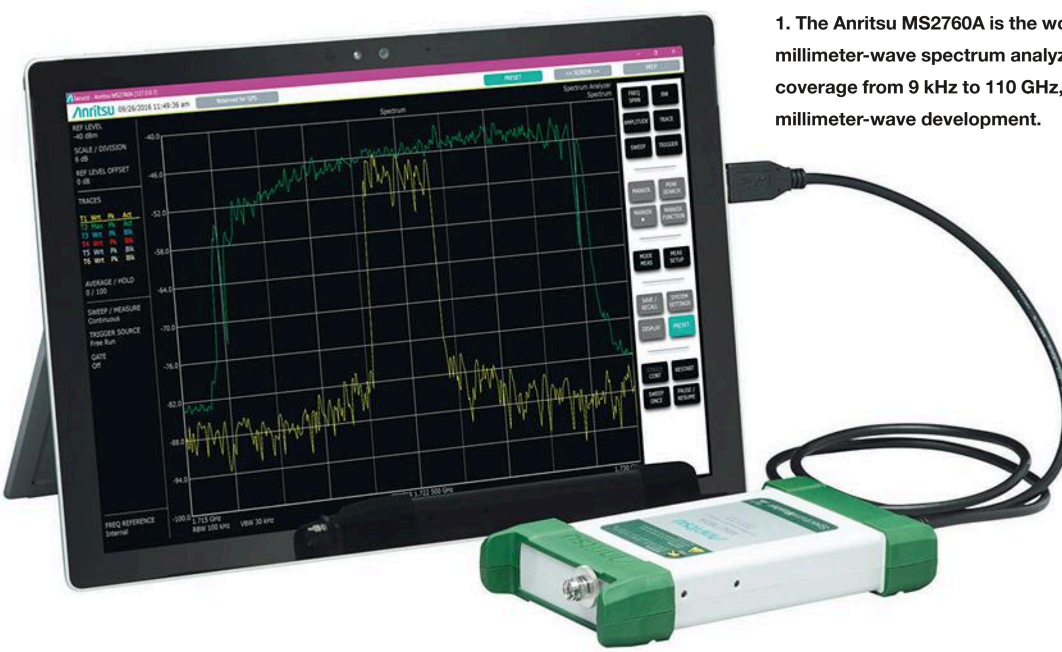

1. The Anritsu MS2760A is the world's first ultra-portable millimeter-wave spectrum analyzer. With its continuous coverage from 9 kHz to 110 GHz, it is a good choice for millimeter-wave development.

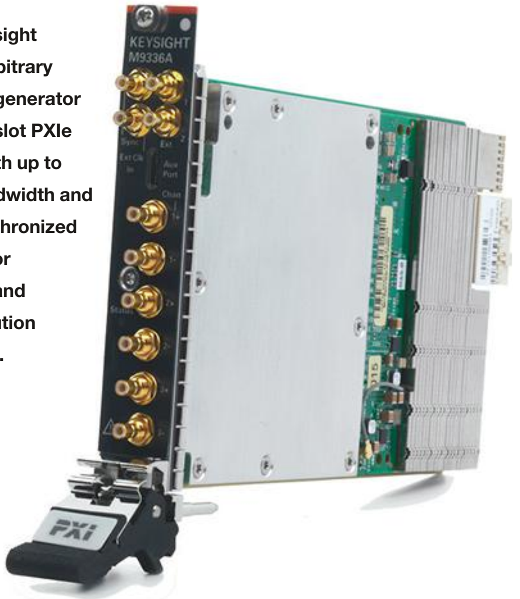

2. The Keysight M9336A arbitrary waveform generator is a single slot PXIe IQ AWG with up to 1-GHz bandwidth and highly synchronized channels for wideband and high-resolution waveforms.

"When most people think of 5G, they think of the new cellular standard that will replace 4G like it replaced 3G," says Adam Smith, director of product marketing at LitePoint. "And while the emerging 5G new radio technology (NR) won't be in carrier trials until 2018, other 5G wireless LAN, low-powered WAN (LPWAN), and broadband wireless standards will be available in 2017—and will drive the wireless testing agenda for most manufacturers."

As for IoT, many products in both the consumer and industrial sectors have already been developed and deployed. With IoT a work-in-progress, many more products are on the way. A wide mix of wireless technologies is being adopted. These include Bluetooth, IEEE 802.15.4, ZigBee, Z-Wave, Thread, and others. Most use the 2.4-GHz band, so readily available test gear can be used. Security testing, however, involves special requirements.

An emerging need is test gear for the new Low Power Wide Area Network (LPWAN) standards. These include LoRa, Sigfox, and the LTE-M standards for implementing cellular IoT for M2M applications. The NB-IoT standard is particularly popular.

According to Tektronix, the IoT is a primary driver in microwave and RF design today, and by extension a key driver in test-and-measurement instrumentation. With the IoT comes exponentially increasing numbers of devices from thousands of manufacturers, along with a proliferation in wireless standards and proprietary protocols (currently more than 20). There is a growing need for engineers to have RF design experience and the time-to-market windows keep shrinking. Systems are being designed and built for a broad range of products, so test instrumentation versatility is vital.

Another major trend that Tektronix sees is the "Interference of Things." An increasingly crowded spectrum is putting new focus on spectrum management for installation, ongoing maintenance, and signal characterization tasks. In addition to a need for portable spectrum analyzers to track interferers in the field, spectrum managers also need the ability to record and analyze hours of wide bandwidth data.

Traditionally, RF experts have shared expensive high-end equipment. The reason for this is that few organizations can afford to buy more than one or two high-end instruments, and low-end signal analyzers lacked the necessary capabilities.

As for Wi-Fi, it continues to evolve. The 5-GHz 802.11ac standard is being rolled out worldwide. New products using the 60-GHz 802.11ad standard are showing up, creating a need for some really high-frequency test gear in the millimeter-wave band. The 802.11ay standard is coming, and this higher-speed upgrade to 802.11ad will require additional millimeter-wave testing.

Next in line is the new, faster 802.11ax standard. It uses high levels of MIMO to achieve super-high data rates. Smith indicated that 802.11ax represents a fundamental change in the way that Wi-Fi will operate. "Today, Wi-Fi often doesn't have the user capacity required to deliver good wireless performance," he says. "The major goal of 802.11ax is to service 10 times more simultaneous users without sacrificing the

3. The LitePoint IQxel-MW is the first system-level test solution for the next generation of wireless challenges driven by 802.11ax. Ideal for both R&D and high-volume production, IQxel-MW delivers high-performance calibration and verification for wireless devices.

throughput of a connection. With 802.11ax, Wi-Fi is borrowing from cellular technology by adopting orthogonal frequency-division multiple-access (OFDMA) modulation techniques used in LTE to get more simultaneous users in the same spectrum."

Smith goes on to say that, unlike cellular—which operates in a managed, licensed spectrum, and with a well-defined connection interchange between the mobile device and the base station—Wi-Fi operates in the unlicensed spectrum, with connections that are established on an ad hoc, on-demand basis. A successful launch of 802.11ax technology will require network infrastructure companies to manage the chaos of the unlicensed spectrum.

Verifying that these devices play by the rules of 802.11ax requires test equipment that is "packet-aware" and supports near real-time power and timing control. Additionally, 802.11ax threatens to drive significantly longer test times compared to the previous generation of 802.11ac.

Automotive applications are beginning to ramp up with the increased deployment of Advanced Driver Assisted Systems (ADAS) and self-driving car technology (with its millimeter-wave radar). The forthcoming Dedicated Short Range Communications (DSRC) technology will add vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) radios in the 5-GHz band. These will begin to appear in many models beginning next year. 5G will probably play a role in vehicle-to-everything (V2X) communications systems, too.

What these trends translate into are the following guidelines, as summarized by Lawrence Wilson, a product manager for Rohde & Schwarz:

- The ability to generate and analyze wideband signals.
- The generation and analysis of digital signals at the higher microwave and millimeter-wave frequencies.
- Keeping the test cost low enough to meet the demands of each segment of wireless.

"The challenge for the wireless engineer is to get test equipment that can provide these signals quickly, accurately, and with repeatability," says Wilson. "Engineers just do not have the time to build, manage, and maintain complex multiple box test systems. They want off-the-shelf equipment that can meet their needs. And the test solutions need to be affordable."

"Today's market requires test solutions that address legacy technologies, while also supporting emerging IoT and 5G verification requirements," says Anritsu's national sales manager, Larry Davis. "Controlling the cost of test is another key attribute. Instruments built upon flexible platforms that support current demands—and can be easily upgraded to meet emerging standards and technologies—address these conditions, so companies have confidence that their equipment can expand as their test needs evolve."

4. National Instruments' 5840 vector signal transceiver (VST) combines a 6.5-GHz RF vector signal generator, 6.5 GHz vector signal analyzer, user-programmable Virtex-7 690T FPGA, and high-speed serial interface into a single two-slot PXIe module.

THE LATEST TEST EQUIPMENT MEETS THE CHALLENGES

Test instrument makers are offering multiple solutions to address these trends. Here is an overview of what's available today:

Anritsu is offering its new Bluetooth Test Sets that address the recent Bluetooth 5 low-power standard for IoT applications. Engineers developing chipsets and modules integrating Bluetooth 5 now have a comprehensive solution that can test products during the design stage. In addition, the test set can be used by manufacturers to verify the performance of products implementing Bluetooth 5, including smartphones, connectivity modules, wearables, smart-home systems, and connected-car systems.

Anritsu's next-generation Bluetooth platforms will also perform interoperability tests on systems-on-a-chip (SoCs). One software solution adds low-energy measurement and packet-generation capability to verify that devices are operating at the new 2-Mb/s rate.

An interesting recent addition to Anritsu's line is the MS2760A, which is claimed as the first ultra-portable millimeter-wave spectrum analyzer (*Fig. 1*). With its continuous coverage from 9 kHz to 110 GHz, it is ideal for the growing 5G network development market, as well as other fast-growing millimeter-wave applications like automotive radar, 802.11ad/WiGig, E-band, satellite communications, and electronic warfare (EW). The MS2760A family meets the need to test

CONNECTIVITY FOR THE HOME

TE Connectivity (TE) is creating the connector and sensor solutions that make connected homes a reality. From smart thermostats to connected appliances, TE's unrivaled portfolio and global manufacturing capabilities can help you turn the wildest ideas into the smartest ways of life.

Innovation for the connected home starts here. te.com/connectedhome

EVERY CONNECTION COUNTS

TE connectivity

at higher frequencies while maintaining performance and affordability.

The MS2760A conducts essential measurements such as spectrum analysis, channel power, adjacent channel power, spurious emissions, and occupied bandwidth. High accuracy is achieved because the MS2760A allows measurements to be taken directly at the DUT, unlike benchtop models requiring expensive cables and mixers that can add loss and complexity.

Lower cost-of-test in high-volume manufacturing applications is also achievable using the MS2760A, as it is considerably less expensive than benchtop alternatives. Furthermore, its ultraportable size adds flexibility to production test of larger products, where test must be taken at multiple points.

Keysight Technologies is also addressing the 5G trend, as well as aerospace and defense needs, with 10 new PXIe instruments. Included are digitizers, oscilloscopes, and arbitrary waveform generators (AWGs). The digitizers and AWGs include a programmable FPGA to customize the tests. The scopes offer 200-MHz, 500-MHz, and 1-GHz bandwidths. *Figure 2* shows the M9336A IQ AWG with 16-bit resolution and up to 1 GHz of modulation bandwidth.

Keysight also recently announced its latest software-centric design and test solutions for LTE-A, 5G, IoT, and connected-car technologies at Mobile World Congress in Barcelona. Examples include:

- An NB-IoT testing solution to help designers accelerate the deployment of cellular long-range IoT technology.
- A 5G wideband real-time beamforming reference solution that allows engineers to quickly and accurately test analog, digital, and hybrid beamforming systems. This includes transmit and receive massive MIMO with beamforming technology and algorithms.
- A Virtual Drive Testing (VDT) Toolset (Anite), a unique lab-based solution that helps the automotive industry cost-effectively verify wireless connectivity in the connected car. The toolset replicates drive test conditions by using field data to create test cases that are replayed in a repeatable and controllable laboratory environment.

LitePoint's solution to wireless testing is its IQxel-MW product (*Fig. 3*). Designed for both R&D and high-volume production test, this 16-port instrument supports a huge range of wireless technologies. The focus is on Wi-Fi with its support for:

- Wi-Fi, 802.11ax (including multi-user MIMO and MIMO configurations up to 8×8, as well as OFDMA)
- Wi-Fi, 802.11a/af/ah/b/g/j/n/p
- Wi-Fi, 802.11ac Wave 1 and Wave 2 (including 80- and 160-MHz bandwidth combinations)

But that's not all. The IQxel-MW can also test the following wireless technologies:

- Bluetooth, Classic/EDR (1-4.x), Low Energy (4.0, 4.1, 4.2, 5.0)
- ZigBee, Z-Wave, Wi-SUN, IEEE 802.15.4
- DECT
- TD-LTE
- Navigation: GPS, GLONASS, COMPASS

National Instruments (NI) is also addressing the trends. NI, the University of Bristol, Lund University, and one of the world's leading providers of communications services, British

5. With its 27.5- to 75-GHz frequency range, the Rohde & Schwarz NRPM over-the-air (OTA) power-measurement solution provides a way to test millimeter-wave wireless devices such as 5G and 802.11ad/ay Wi-Fi.

Telecommunications (BT), recently announced a collaboration on massive-MIMO trials for next-generation, highly efficient 5G wireless connectivity. Massive MIMO is a crucial component of future networks, and this partnership is helping 5G become more of a reality through indoor and outdoor testing, which is key to working in real-world environments.

MIMO is already used in Wi-Fi and 4G LTE cell-phone systems, but as network capacity demands increase, there is an obvious need for more spectrum efficiency. NI and the two universities first began working together here in 2016 and set world records in 5G wireless spectral efficiency using massive MIMO. NI's LabVIEW Communications System Design Suite, MIMO Application Framework, and SDR hardware were used for this partnership.

The NI MIMO Prototyping System is a 128-antenna, real-time, massive-MIMO testbed. Using this testbed, NI and the universities set two world records in spectral efficiency. They achieved over 79 b/s/Hz of spectral efficiency over a 20-MHz bandwidth, fully bidirectional, real-time, over-the-air link at 3.5 GHz with 12 simultaneous users. Later the team extended the system to achieve over 145 b/s/Hz of spectral efficiency, increasing the number of users to 22 sharing the same time-frequency resource.

One core product that addresses wireless test trends is NI's second-generation PXI vector signal transceiver (VST). This combo unit offers 1 GHz of instantaneous RF bandwidth for signal generation and analysis. The PXIe-5840 VST (*Fig. 4*) combines a 6.5-GHz RF vector signal generator, 6.5-GHz vector signal analyzer, user-programmable Virtex-7 690T FPGA, and high-speed serial interface into a single two-slot PXIe module.

The FPGA is programmable in LabVIEW software, a feature that allows engineers to modify and customize the instrument's firmware. The NI 5840 VST provides an upgrade to the first-generation VST in nearly every aspect, supporting a wider frequency range with a wider instantaneous bandwidth of 1 GHz.

Rohde & Schwarz is offering some interesting products for millimeter-wave testing. An example is the R&S NRPM over-the-air (OTA) power measurement solution (*Fig. 5*). With its frequency range of 27.5 to 75 GHz, it targets 5G as well as Wi-Fi 802.11ad and 802.11ay. The antenna modules include a diode detector. The antennas use beamforming to control the direction of radiation of the transmit antenna, thereby maximizing the power level at the receiver.

6. The R&S SMW200A may be the only vector signal generator available that can create signals 2 GHz wide at frequencies up to 40 GHz to address the key 5G bands of 28 GHz and 37/39 GHz

The OTA solution is designed to calibrate the transmit antenna output power and test the beamforming function. The power measurement range is 30 pW (–75 dBm) to 3 µW (–25 dBm). This instrument can implement both 2D and 3D beamforming tests and is scalable to customer needs.

Another product meeting the needs of current trends is the R&S SMW200A vector signal generator (*Fig. 6*). The generator comes in models with frequency ranges up to 3, 6, 12.75, 20, 31.8, and 40 GHz. It has a 2-GHz I/Q modulation bandwidth. Its extensive features support most common digital standards, such as LTE, 802.11, and even 5G air interface candidates.

Tektronix has also introduced some new products driven by the trends. An example is Tek's line of USB-based spectrum and signal analyzers. USB-based RF test instruments possess fundamental advantages over their traditional desktop equivalents. In addition to signal-capture circuitry, desktop spectrum analyzers must also include a full personal computer, adding considerable cost and complexity to the instrument.

In contrast, the USB-based spectrum analyzer eliminates this cost from the instrument itself and simply takes advantage of the processing power sitting on every engineer's lab bench. USB-based instruments are controlled from PC-based software, offering full real-time capture and advanced analysis capabilities in a compact package. In terms of price, organizations can now afford to buy 10 to 20 USB-based spectrum analyzers for the same price as one $50,000 or $100,000 instrument and get similar analysis capabilities.

7. Shown are Tektronix's USB-based spectrum analyzers (l-r): the RSA500, RSA306B, and RSA600.

8. The Tektronix RSA7100A is a USB-based desktop spectrum analyzer that offers 800-MHz real-time capture bandwidth and up to two hours of streaming storage.

Tektronix's new RSA500 and RSA600 series of analyzers (*Fig. 7*) offer frequency coverage from 9 kHz to 7.5 GHz with 40-MHz acquisition bandwidth, a measurement dynamic range from –161 dBm/Hz displayed average noise level, and up to +30-dBm maximum input. Both work with a standard laptop via a USB 3.0 connection. The RSA500 series units target field RF measurements and are battery-powered. The RSA600 series is designed for wireless lab applications. Both instruments are based on the design of Tek's successful RSA306 real-time spectrum analyzer. The instruments use an updated version of the SignalVu-PC software with DPX signal processing capability.

To address the needs of spectrum managers who need to capture and analyze wide bandwidth data, and designers of radar and EW communications systems and components who need to perform modulation and pulse analyses at wider bandwidth, Tektronix recently introduced its USB-based RSA7100A wideband signal analyzer (*Fig. 8*). It offers 800-MHz real-time capture bandwidth and up to two hours of streaming storage to internal RAID storage. It, too, uses the new DataVu-PC software. ed

FELIX BAUM | Senior Product Manager, Embedded Systems Division, Mentor Graphics
www.mentor.com

Securing IoT Medical Devices—Are We There Yet?

To ensure secure communications in a given design, developers must consider integrating key security- and safety-related features that help to harden a medical device against any malicious activity.

As the internet of medical devices grows exponentially, much of the attention has been given to the safety aspect, with less time devoted to protecting the private information transmitted and stored on these devices. However, that imbalance is shifting, as protecting patient data has become a critical concern. In fact, a KPMG 2015 healthcare and cybersecurity survey found that more than 80% of health plans and healthcare providers acknowledged that patient data had been compromised—even after making significant cyber-related investments.

This article discusses key security-related procedures developers need to adopt during the design process to harden the device, secure communications, and prevent malicious exploits. By implementing just a few of these precautions, software developers can design state-of-the-art connected medical devices that protect patient data and are safe for all to use *(Fig. 1)*.

In these days of Internet of Things (IoT) proliferation, medical-device manufacturers are working hard to simultaneously address three trends:

1. The integration of multicore SoCs into their designs.
2. Adoption of Linux as an operating system.
3. Meeting safety and security requirements brought on by the IoT.

INTEGRATION OF MULTICORE SoCs

Many devices on the market today include both single-core and multicore processor architectures. Single-core designs ruled the day when embedded devices were purpose-built.

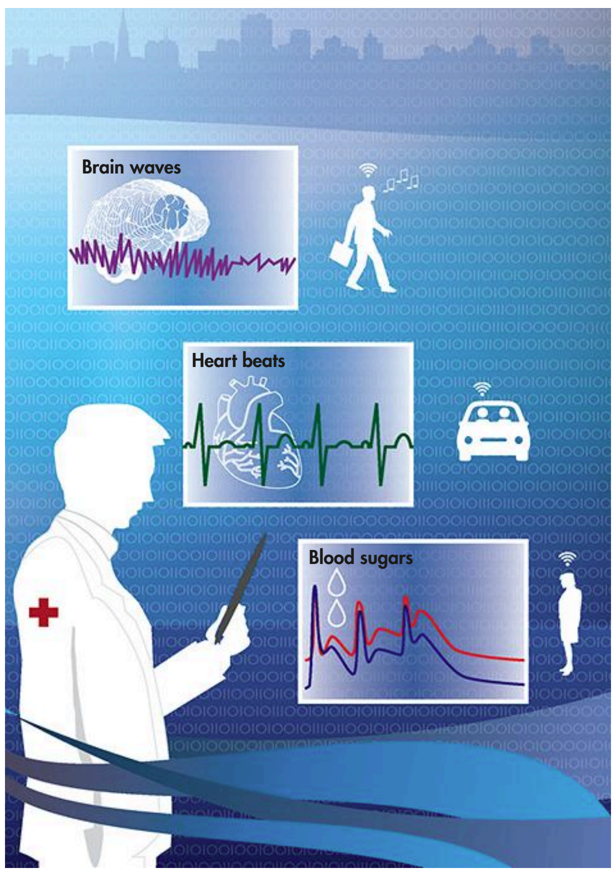

1. IoT medical devices allow medical practitioners to remotely monitor a patient's vital signs throughout the day.


WITHOUT US...
HE WOULDN'T BE
GOING...
AND GOING...
AND GOING...

With the advent of graphical user interfaces (GUIs), graphics processing units (GPUs), faster and more powerful silicon, and the need for a system to process multiple tasks simultaneously, hardware designers have made the transition to multicore designs.

When designing a medical IoT device, care must be taken to select the right silicon. With tight development cycles compressing the limited available time dedicated to security features, processors with built-in security IP blocks establish a solid foundation. Today's system-on-chip (SoC) processors comprise the hardware features to authenticate software prior to execution, encode data-at-rest, sign data to maintain integrity, and partition the device to prevent malware from entering the system. Processor security features to look for include secure boot, boot fuses, crypto engines, and device partitioning.

MORE LINUX-BASED DEVICES

With more platforms based on multicore SoCs, software developers are now able to follow the trend of the general embedded market by adopting Linux as the primary operating system. Linux is already used in a wide range of medical devices, from vital sign monitors and hospital bedside monitoring systems to complex imaging equipment.

A few reasons why Linux is gaining in popularity:

- The innovation and maturity of Linux has made it a practical choice in medical-device development.
- Free distributions are available, and they can be modified and redistributed under the GNU General Public License (GPL) and other licenses.
- Linux has been widely adopted by thousands of developers, making it easier to find developers who frequently use the operating system and intimately know it.
- Linux has a large ecosystem of semiconductor, board, and software providers who use proven toolchains and application programming interfaces (APIs).
- Linux is feature-rich, including tools, connectivity, security, and graphics — important for device UIs that require clarity and readability.

Medical devices are certified as a system by the U.S. FDA. The challenge is to design a system that can leverage Linux while meeting safety requirements.

With Linux being adopted more widely, the approach starting to take root is Software of Unknown Provenance (SOUP), or off-the-shelf (OTS) software. As such, in most of the devices, Linux is running alongside an operating system that has been certified to the IEC 62304 standard tasked with safety-related operations.

Security through connectivity should be everyone's concern. All of the trends are bringing safety and security concerns to the forefront, and in some instances, slowing down the internet-of-medical-device adoption.

Safety-certified operating systems isolated from Linux help avoid situations that might result from software faults that could cascade across the system and result in security vulnerabilities in unrelated sections of code. For example, Nucleus SafetyCert RTOS from Mentor Graphics makes it possible for developers to shorten the path to regulatory certification with a certified solution that includes all necessary documentation to develop mission-critical applications.

IoT CONNECTIVITY

Connectivity and IoT requirements are making these multicore and multi-OS systems, in which one of the OSs is Linux, more easily accessible to outside actors. Security through connectivity should be everyone's concern. All of the trends are bringing safety and security concerns to the forefront, and in some instances, slowing down the internet-of-medical-device adoption.

But there are ways to turn these challenges around.

WHAT'S YOUR SURFACE AREA?

When addressing how to secure a medical IoT connected device, it's important to look at the "surface area" vulnerable to attacks. The area of attack varies from device to device, but generally, the more sophisticated the device, the greater the area of attack. It's also important to understand that most of the threats today target data not for the sake of data, but for the ability to manipulate the data.

One example of manipulating data would be an attack on an algorithm that affects the operation of the very system it depends on for operation, such as a banking application at an ATM terminal or a point-of-sale device at a retail check-out. Within the medical field, an example might be corruption of the sensor information that's used to determine the patient's vital information. When it comes to protecting data, developers need to be aware of three critical stages: data at rest, data in use, and data in transit.

Let's look at how developers can secure data in these three stages.

DATA-AT-REST PROTECTION REQUIREMENTS

These requirements address the state of a device from being powered down to fully operational in the "on" mode. Among the issues to consider in this stage include:

Secure storage

Where is the bootable image stored? Is it on a media that could be encrypted? Are there anti-tampering methods supported by hardware to inform the device? And, if it's being tampered with, is there a way to prevent it from booting into a vulnerable state? Have the executables been encrypted, or could anyone who is able to gain access remove EEPROM, dump the memory, or attempt to reverse-engineer the application?

Root of trust

Is the SoC used in the hardware platform capable enough to provide root of trust? If not, does the hardware platform include a secure element that would offer similar capabilities in SoC fashion? Even if the answer is no to both of these questions, there are security vendors who offer purely software-based root-of-trust solutions.

Chain of trust

With the root of trust present *(Fig. 2)*, the boot process establishes a chain of trust:

- Has the hardware validated and authenticated the bootloader?
- Has the bootloader validated and authenticated the operating system or other systems?
- Have the operating systems validated and authenticated application code?
- Has the validation and authentication been done by the hardware or software, and is this different at each stage?

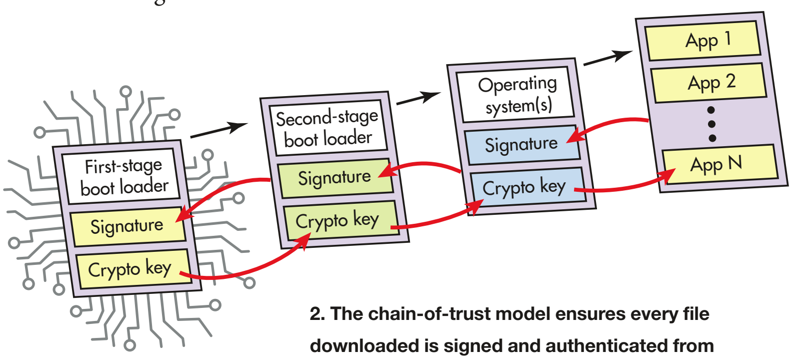

2. The chain-of-trust model ensures every file downloaded is signed and authenticated from the bootloader up to the application level.

By utilizing the technologies listed above, you can be assured that your device will be running authentic and unaltered application code—with an extremely low risk of the device being hijacked.

DATA-IN-USE PROTECTION REQUIREMENTS

These requirements are applicable to a device that's operating normally, with the data being generated and processed. A few technologies to utilize at this stage might include:

Hardware-enforced isolation

Some software developers physically isolate applications by placing multiple SoCs side by side. Not everyone can afford to do this, so for the rest of us, the next best thing is to use ARM TrustZone. ARM TrustZone technology implemented in a SoC can be leveraged to address the network, application, and data aspects of the layered security model.

ARM TrustZone architecture provides a solution that's able to carve out or partition a hardware subset of the full SoC. It does this by defining processors, peripherals, memory addresses, and even areas of L2 cache to run as "secure"

DESIGN NOTES

High Voltage, High Efficiency Positive to Negative Converter

Design Note 561

Victor Khasiev

Introduction

Positive to negative conversion (inverting output) is widely used in LCD devices, OLED displays, audio amplifiers, industrial machinery, semiconductor manufacturing process control equipment, measurement tools, testing systems, LED drivers and battery chargers. Many of these applications require high power levels and extended input voltage ranges, two inherent features of the LTC3896 positive to negative controller. This controller is especially well suited for automotive applications due to its ability to handle very high input and output voltages (total voltage stress can reach 150V), its ability to drive standard level MOSFETs and its low operating quiescent current 40μA (10μA in shutdown).

Wide 7V to 72V Input to –12V at 5A Output

Figure 1 shows a positive to negative converter based on the LTC3896. This solution delivers –12V at 5A from a 7V to 72V input voltage range. In the automotive market, the LTC3896's ability to handle high input voltages eliminates the need for bulky and costly voltage suppressors, while the low minimum input voltage keeps sensitive systems operational even during cold crank conditions.

The power train of the converter comprises MOSFETs Q1, Q2 and inductor L1. The output filter is based on ceramic capacitors C_{OX}. The $EXTV_{CC}$ pin of U1 is connected to GND, producing a 12V potential at this pin relative to V_{OUT}^-. If a power good signal is required, an external voltage source referenced to GND should be employed. The LTC3896's control and interface signals, including RUN, PGOOD and PLLIN

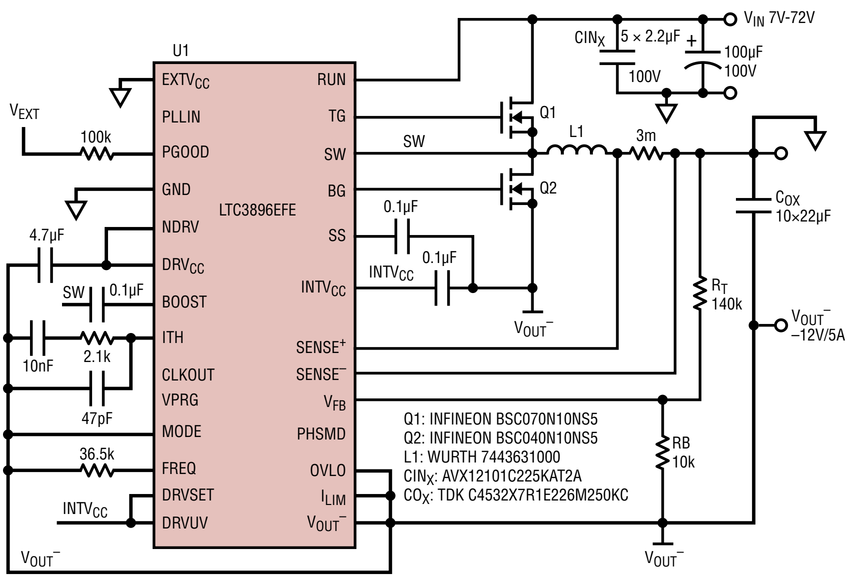

Figure 1. The –12V Output Converter (U1) Provides 5A to the Load in the Input Voltage Range from 7V to 72V. Note Control Signals RUN, PGOOD and PLLIN are Referenced to GND

are referenced to the system GND to eliminate the need for level shifters in processor controlled systems.

The guidelines for calculating voltage and current stress on the components surrounding the LTC3896 are detailed in the data sheet [1]. For a basic evaluation, duty cycle (D), average inductor current (I_L) and MOSFET voltage stress (V_{DS}) can be calculated by the following:

$$D = \frac{|V_O|}{V_{IN} + |V_O|}$$

$$I_L = \frac{I_O}{(1-D)}$$

$$V_{DS} = |V_O| + V_{IN}$$

Demonstration Circuit DC2447A [2] illustrates the versatility of the LTC3896. Designers can test the numerous functions of this controller, including synchronization to an external clock, the ability to employ an external linear regulator to reduce thermal stress on the IC at high output voltages and easy solutions for generating –5V or –3.3V outputs.

Figure 2 shows efficiency at various input voltages. Figure 3 shows a thermal image of the converter in action.

Conclusion

The LTC3896 is a highly integrated controller, specially designed for positive to negative conversion. Solutions based on this controller are highly efficient, with extremely low quiescent current—important for battery-operated systems. It is also highly versatile, featuring programmable frequency, a wide 150V input voltage range and output voltages to –60V. It simplifies the design of automotive and industrial supplies with control signals referenced to host ground.

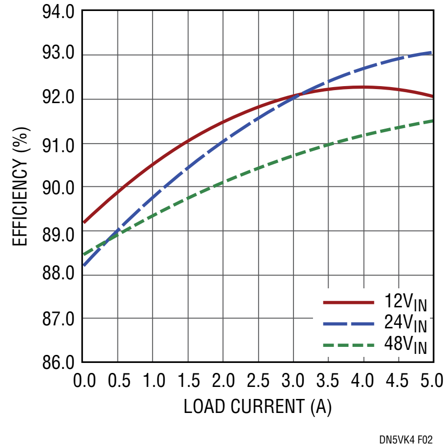

Figure 2. Efficiency Curves for the Circuit in Figure 1. Output Voltage is –12V and Maximum Load Current is 5A

References

1. LTC3896: 150V Low IQ, Synchronous Inverting DC/DC Controller.

http://cds.linear.com/docs/en/datasheet/3896f.pdf

2. DC2447A: LTC3896EFE Demo Board|Sync Inverting Controller, 7V ≤ VIN ≤ 72V; VOUT = –12V at 5A.

http://www.linear.com/solutions/7378

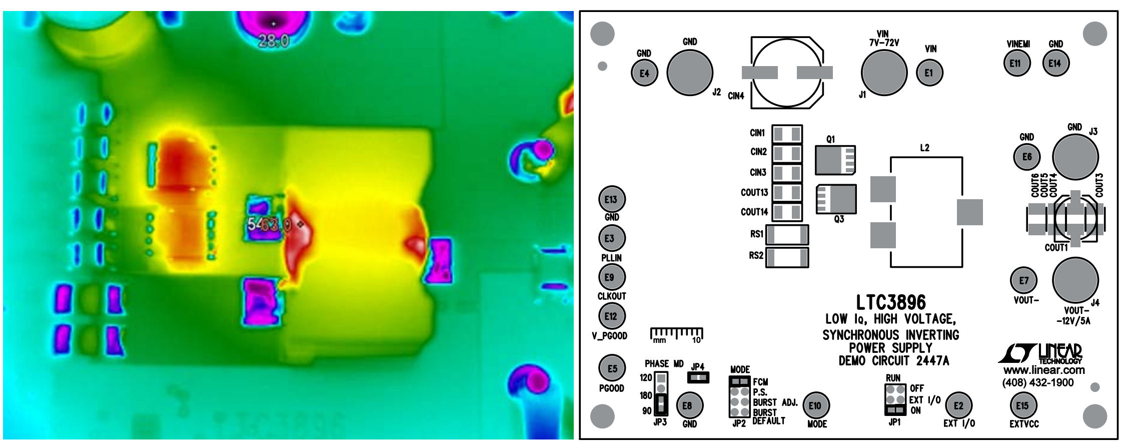

Figure 3. Thermal Image of DC2447A is Taken at 36V Input Voltage, –12V at 5A Output. On the Right Side, an Assembly Drawing of the Demo Board

Data Sheet Download

www.linear.com/LTC3896

For applications help,
call (408) 432-1900, Ext. 3161

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

DN561 LT/AP 0417 71K • PRINTED IN THE USA

or "non-secure" hardware. A SoC that employs TrustZone technology has the ability to dynamically expose the full SoC to secure software, or else expose a subset of that SoC to normal software.

The normal world (non-secure world) created and enforced by TrustZone is typically a defined hardware subset of the SoC. TrustZone ensures that a non-secure processor can access only non-secure resources and receive only non-secure interrupts.

For example, a normal world hardware subset might include the UART, Ethernet, and USB interface, but exclude controller-area-network (CAN) access. The CAN might instead be dedicated to the secure world, where a separate RTOS or application runs for the sole purpose of managing CAN traffic, independent of the normal world software stack. In this example, one can run Linux with an internet connection and be certain that even if someone hacks and brings it down, the communication via CAN with the off-the-chip/board unit that might be attached to the patient will not be jeopardized.

Unlike the hardware subset that runs normal world software, software running within the secure world has complete access to all of the SoC hardware. Thus, from the perspective of the secure software's execution, the system looks and acts nearly identical to what would be seen on a processor that doesn't have TrustZone. This means that secure software has access to all resources associated with both the secure and normal worlds. As such, updating software in the normal world via code executing in the secure world is very easy.

Software-enforced separation

If the hardware separation isn't an option, the next best thing is to use the software to isolate and protect applications. In the past, particularly on single-core SoCs, some designs utilized the concept of a separation kernel. These days, with multicore SoCs supporting virtualization extensions in the silicon, more designs are utilizing embedded hypervisors.

Hypervisors allow multiple instances of the same or different operating systems to execute on the same SoC as a virtual machine. Each virtual machine can be isolated, and through use of a system memory management unit (MMU), it's possible to virtualize other bus masters. This separation can be used to protect and secure resources and assets in one virtual machine from other virtual machines.

To fit the real-time and performance demands of medical devices, embedded virtualization places certain requirements on a hypervisor design. First, any hypervisor will not do—it has to be Type-1 (bare metal) hypervisor supporting hardware virtualization extensions with a minimal footprint and available in source code. These requirements seem obvious to some: The smaller the code, the less performance degradation one would expect when moving from native to virtualized execution.

In addition, the better support for virtualization in the SoC, assuming that there's support for the hypervisor, the less code will be required to support the same features in the software, and the faster things will run (aka hardware offload). Finally, most embedded designs are custom. No two designs are alike, even if they use the same underlying SoC, board, and guest OSs. The ability to tinker with source files to optimize the hypervisor performance and behavior to better fit a particular SoC/board/design is crucial to your success.

User-space isolation

Many operating systems today offer some type of MMU-enforced isolation of the application code running in the RAM; Linux has user space, and Nucleus RTOS has process model *(Fig. 3)*. The idea here is that while the kernel of the operating system runs at a more privileged level, such as EL1 or EL2, applications run in the EL0 and use various memory-isolation techniques to protect code and data on a per application basis.

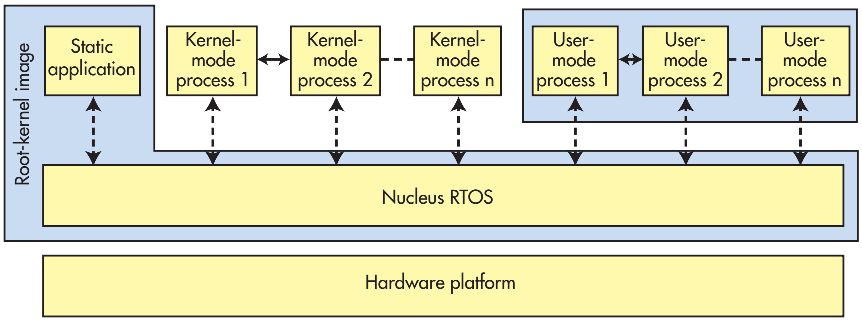

3. Nucleus Process Model is a lightweight approach to space partitioning that creates protected memory regions.

Information or data obfuscation

While most developers are careful about hiding and encrypting the password, care must also be taken to obfuscate variables and text strings stored in the memory or storage. This will make it more difficult for bad actors to modify variables and text strings, or even try to re-engineer the device operation.

In addition to the aforementioned isolation techniques, realize that these devices are connected to many instances to the outside world and one must be mindful to secure all connections. As such, the least a developer could do is enable and configure a firewall, and utilize various publicly available tools to perform network stress testing and network penetration analysis.

DATA-IN-TRANSIT PROTECTION REQUIREMENTS

Data-in-transit relates to data entering or leaving a device while the device is "on." A good design should always address the following two areas.

- Prior to sending the data, will the device utilize any mutual attestation? Various tricks and techniques could be deployed to authenticate the receiver of the information prior to sending it.
- How is the data protected in the event the device is hijacked? Encrypting the data while at rest and during transit provides some level of protection. For efficient cryptography operations, SoCs with crypto engines should be considered. The crypto engine is a

self-contained module designed to offload the encryption/decryption work from the core processor. Because the crypto engine is a self-contained IP block, hackers can find it difficult to derive techniques to gain access to the encryption process. The implementation of crypto engines can have a huge impact on an application's ability to secure data quickly and efficiently.

THE REST OF THE STORY

Now that we've outlined several of the technologies that medical-device manufacturers should taken into consideration, we can't skip mentioning the process itself. Recently, while working on a highly secure piece of software, the author of this article followed the widely accepted Security Development Lifecycle (SDL).

Some might find the SDL process a bit heavy and involved. However, developers should take the proper care in selecting the correct tools to ensure the traceability of requirements, code, and test cases.

It's also important to use appropriate tools for the verification stage of the development process. The verification and validation could be as simple as using static-code-analysis tools for compliance to MISRA (safety) or Cert C (security). With enhanced connectivity options, it's also important to perform network stress testing and employ third-party testing and certification, such as Achilles certification offered by Wurldtech, a Mentor Graphics partner.

Once the development is completed and devices are shipped, one has to determine the remediation strategy. This might be the deciding factor in using a commercial Linux offering like Mentor Embedded Linux. Mentor Embedded Linux customers will have a Mentor security team proactively monitoring the US-Cert website to identify Common Vulnerabilities and Exposures (CVE) that affect their products. If a vulnerability is discovered, Mentor Graphics will notify customers with patches to the affected products according to the product support policy. It's then up to the device manufacturer to apply the patch to the code base used in the product, and if needed, roll out the software update.

SUMMARY

Security for medical IoT devices is a complex subject. Designing an embedded system from the ground up requires many of the proven security capabilities used for years by other electronic system designs—secure boot, code authentication, and chain of trust, to name a few. These are fundamental security capabilities that every connected device needs to include in some shape or form.

Further, as these devices blend seamlessly into our daily lives, it's incumbent on software developers to design each new device with security as a paramount concern. Through the use of ARM's TrustZone technology, together with a Type-1 hypervisor, developers can provide a strong, robust, and secure base for system-on-chip designs that meet the demands of our ever-expanding IoT world.

FELIX BAUM is a senior product manager with Mentor Graphics' Embedded Systems Division, overseeing embedded virtualization/multicore/multi-OS technologies. He currently leads product marketing and management efforts for real-time operating-system capabilities. Prior to his current position, Baum worked in business development with strategic alliance partners around the globe. He has also consulted customers on the development of highly optimized devices for a broad range of industries. Baum started his career at NASA's JPL, designing flight software for various spacecraft and satellites. He holds a Master's in computer science from the CSUN and an M.B.A. from the University of California at Los Angeles.

WILL KEEGAN | Lynx.com

What's the Difference Between Separation Kernel Hypervisor and Microkernel?

As the embedded world looks to security solutions to protect connected critical computing functions from external threats, two software platforms have emerged that, at first glance, appear to offer similar functionality. However, under the hood, the approaches and the technologies are quite different.

The Separation Kernel Hypervisor and Microkernel technologies have emerged as the leading contenders in hosting next-generation embedded safety and security critical compute platforms. Both technologies share a great deal in common, stemming from least-privileged design principles, and aim to provide a more robust application runtime environment than traditional monolithic kernel-based OSs. The technologies are similar enough that in the commercial world, the terms are regularly interchanged based on audience or industry requirements, and hence become very confusing for consumers. Despite the similarities in using minimalistic approaches to control CPUs, the kernels are only useful when vendors construct application development platforms on top of them. Once in the hand of the developer, the delivered products can have wild differences at the CPU control level, system assurance properties, development models, and application behavior.

This article embarks on providing an introductory comparison between the two technologies by providing more context of on the origin, expound on their intent to host applications, and highlight some fundamental differentiating properties of the Separation Kernel Hypervisor from Microkernel based platforms.

ORIGIN AND INTENT

The Separation Kernel Hypervisor and Microkernel concepts have existed for over 30 years with definitions widely available through online and formal publications. The Separation Kernel Hypervisor is an extension of the Separation Kernel originally defined by Dr. John Rushby in his seminal 1981 paper entitled "Design and Verification of Secure Systems" and essentially describes "the task of a separation kernel is to create an environment which is indistinguishable from that provided by a physically distributed system: it must appear as if each regime is a separate, isolated machine and that information can only flow from one machine to another along known external communication lines." The Separation Kernel Hypervisor incorporates the premise and principles of Rushby's Separation Kernel, but imposes an explicit implementation constraint to leverage native CPU virtualization capabilities to autonomously host applications and further enforce separation.

The Microkernel concept comes from many generations of research and development from all over the world, striving to construct efficient CPU control models with the least amount of privileged code necessary to implement an Operating System (*Fig. 1*).

Microkernels aim to provide a safer and more secure run-

1. Here is a diagram of the Microkernel architecture.

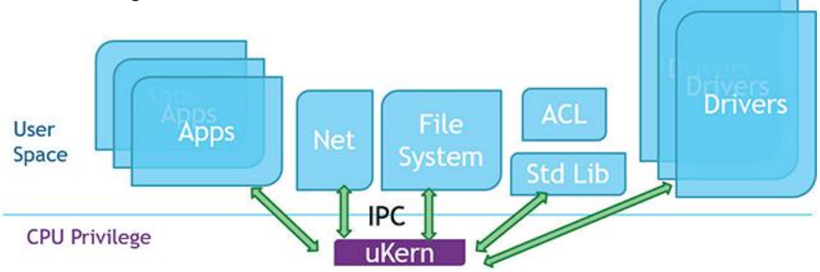

time environment over the popular monolithic kernel based Operating Systems *(Fig. 2)*. Unlike the monolithic kernel approach where all device drivers, I/O, and administrative services run in the same privileged address space to reduce the penalty of context switch time; the Microkernel requires all OS service components to be broken out into separate address spaces at the cost of extra context switching but gaining better integrity protection among internal service components. From an application platform perspective, a Microkernel-based OS and monolithic kernel-based OS look very similar, they are both Operating Systems governed by administrative and configuration policy.

The Separation Kernel Hypervisor, however, aims to support a very different thing—the construction of an n-way fully independent virtually distributed runtime architecture. This architecture demands that there is absolutely no existence of a central controlling Operating System; instead there are a number of operating systems, each one fully independent from the other, and none of which can fully control the physical host.

The Separation Kernel Hypervisor offers a different runtime architecture referred to as a Distributed Heterogeneous Architecture. This independent distributed runtime architecture serves as the main differentiator between the two kernel technologies and has significant advantages over operating systems noted in the following sections.

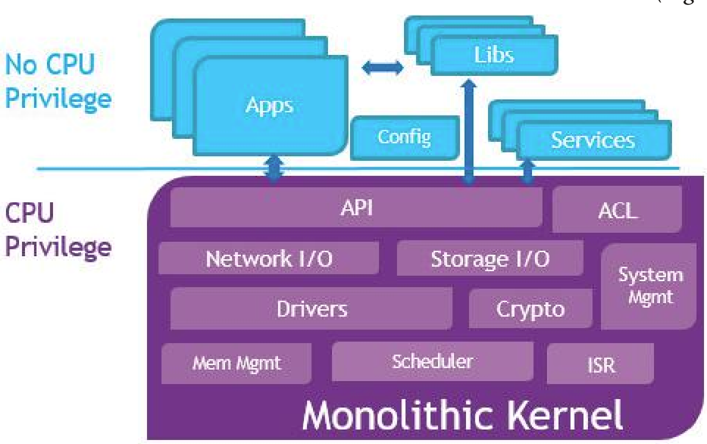

2. Above is the Monolithic Kernel architecture

DIFFERENTIATING PROPERTIES OF THE SEPARATION KERNEL HYPERVISOR

There are many comparable properties that one can examine between a Separation Kernel Hypervisor and Microkernel, like performance, deterministic behavior, trusted codebase, etc. However, these comparisons are only useful if the kernels have the same CPU controlling features and similar application runtime profiles. Imagine comparing the trusted code base of a Separation Kernel Hypervisor capable of scheduling tasks, isolating task memory, and isolating I/O between tasks and physical interfaces, versus a Microkernel that is only capable of scheduling tasks. Obviously, the Microkernel with the lesser of CPU control capabilities will look favorable in a source code line count comparison but it wouldn't be a fair comparison.

Microkernels have existed in the market much longer than Separation Kernel Hypervisors and have taken many forms from, each one different from the other influenced by the user community and the creative direction of its authors. This introductory comparison does not aim to single out any particular Microkernel implementation. Instead it aims to call out key Separation Kernel Hypervisor properties that have stood out as strong differentiators from Microkernel-based OSs over the years to educate consumers on available capabilities, benefits, properties that they may not have previously been exposed to. Some of the properties noted in these sections are not necessarily be exclusive to Separation Kernel Hypervisors and may over time be incorporated into Microkernel-based platforms.

SCOPE OF KERNEL FUNCTIONALITY

The Separation Kernel Hypervisor is a pure CPU Control Plane, not an application runtime framework like an OS *(Fig. 3)*. Its focus is to map raw resources into guest spaces that can be locally controlled by the guest. It is therefore, up to the guest environment to provide a suitable application runtime environment, obviating scope of functionality for the separation kernel hypervisor to provide OS service interfaces as seen in microkernels.

PROVABLE SYSTEM STATE

A driving motivation of Rushby's Separation Kernel invention was to provide a runtime architecture amenable to proving system safety and security properties. The Separation Kernel Hypervisor offers a system configuration model that clearly constrains the runtime architecture of all execution contexts to an explicit tamper-proof set of physical resources. By doing so, evaluators with a discerning eye for safety

3. The Separation Kernel Hypervisor is a pure CPU Control Plane, not an application runtime framework like an OS.

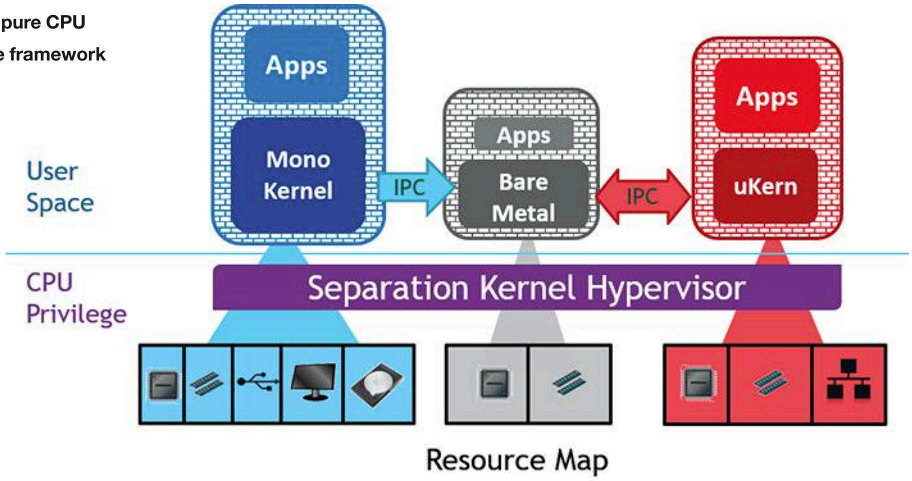

and security properties can easily argue about the system assurance properties and performance characteristics without depending on a highly complex black-box piece of software. By simply mapping every logical actor in the distributed system to its physical resource, evaluators can clearly calculate performance metrics or model assurance states by inheriting the immutable physical properties of the hosted logic in question, rather than inheriting assurance properties enforced by a software kernel that has not proven itself enforceable. On modern SoCs with many CPU cores, hosting many applications, a single OS can become a common choke point of vast complexity that can be riddled with security side channels and incomputable deterministic states.

AUTONOMOUS RUNTIME DOMAINS

The autonomy of the application runtime environments is a crucial aspect of a Separation Kernel Hypervisor hosted platform. The autonomy of a guest runtime greatly improves the ability to prove the system states of a software defined system. With autonomous runtimes, software components can be factored out of scope of the proof exercise. If applications were not truly autonomous and were instead dependent on a central resource manager and set of system services, then the scope of the proof exercise must consider the all relationships the central controlling software has with all actors in the system.

Consider a practical system vehicle controller design where two applications need to be separated, e.g., internet-facing application and CAN bus- facing application. On a Separation Kernel Hypervisor, these two applications can run with zero logical links two each other, removing all potential software based exploitations from adversaries to bypass the separation. Even on a least privilege OS, processes are at the mercy of both the integrity of the microkernel and OS system policies, making it difficult to prove that the applications are truly separate. As seen in the infamous Jeep hack (Greenberg, 2015), the attackers used central microkernel-based Operating System services to reverse engineer the system configuration policy to find an exploit that allowed an internet-facing application to control the CAN bus. On a Separation Kernel Hypervisor, the system could have been configured such that the CAN bus interface was never mapped into the internet domain, making CAN bus control a physical impossibility from the perspective of the internet domain.

ISOLATION ENFORCEMENT

Without resource isolation, autonomous runtimes cannot be achieved. The Separation Kernel Hypervisor relies explicitly on recent advancements of CPU memory and I/O controller elements to support hardware assisted CPU virtualization. These advancements under CPU virtualization allow the Separation Kernel Hypervisor to assign portions of hardware resource to guest environment to locally manage without the assistance of the Separation Kernel Hypervisor. These hardware assignments include the ability of guests to independently control their own CPU interface, system memory, and peripheral controllers. On a Separation Kernel Hypervisor, the assignment of resources to guest environments is actually enforced by immutable polices set in the CPU controllers. Having hardware enforce separation provides a tremendous advantage when considering the strength of function, performance efficiency, and strength of evidence in making system assurance property evaluations. This model is starkly different than the service-oriented architecture seen in OS designs where applications would rely on the kernel to manage memory, CPU execution contexts, and broker connectivity to I/O. Many of these OS models rely solely on software based policies that could be potentially tampered with by a software actor that gains CPU or OS administrative privilege.

(Continued on page 45)

ideas for design

Interactive Graphical Interface Enhances Micro-Programmed Sequential Machine Design

By ROBERTO TOVAR, ROMÁN OSORIO, and TANIA TOVAR | Universidad Nacional Autónoma de México

DIFFERENT DESIGN APPROACHES, techniques, and architectures are widely used for synchronous sequential machines. In particular, the micro-programmed machine approach offers important features such as speed, flexibility, a fixed hardware architecture, and a reduced instruction set (Count Conditional "CC"; Count Unconditional "CU"; Branch Conditional "BC"; Branch Unconditional "BU"; Count / Branch Conditional "CBC"). This has the objective of adapting the machine-design solution to the fixed hardware. Its main disadvantages are that it's tedious and there's a high probability of making mistakes in the process of manual instructions assembly.

This article presents a Visual Basic (VB6) application that uses a graphical interface to allow for design process automation of a "four-bit slice" micro-programmed machine, with up to six input and four output variables and 16 possible states E0-E15 *(Fig. 1)*. The solution generated with this software tool is a HEX file to be stored in the memory of the machine (IC4-IC5). The architecture of the machine can be simulated with some of the free tools offered by different hardware manufacturers and synthesized with medium-scale-integration (MSI) ICs or by using a complex programmable logic device (CPLD).

Like most Windows applications, the one shown here has a series of friendly menus that allows for the user file management, document editing, and micro-instructions assembly. To start the machine-design process with this tool, it's necessary to have a state diagram or an ASM chart *(Fig. 2)*, which represents the logical operation of the machine. The ASM chart shows the names and codes assigned to each state; input variables and micro-instructions to move from one state to another; and the immediate type outputs generated in the states.

This information will be entered into the graphical interface in the next possible order: inputs, outputs, names for states, and micro-instructions. To enter an input variable, place the mouse on one of the inputs of the multiplexers: Count Load (MuxCntLd) IC1 or Count Enable (MuxCntCe) IC2, click the left mouse button, and a gray text box will be displayed to edit the variable name; then click again the left mouse button to record changes. The same procedure is used to enter an output variable *(Fig. 3)*.

To view: Enter or edit the name of each one of the states used and their corresponding micro-instruction; place the mouse over any of the E0-E15 labels located on IC4 and IC5; and then click the left mouse button to display two controls located to the right of the selected label. The first con-

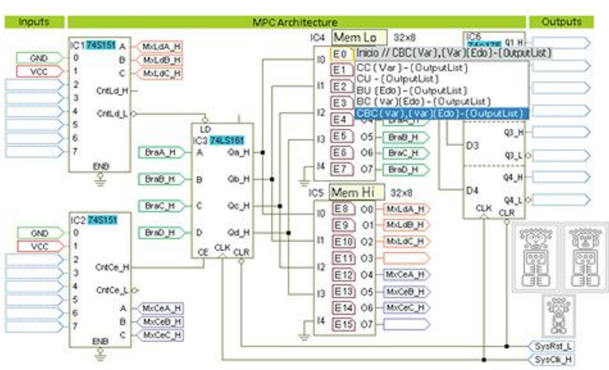

1. The graphical interface of this Visual Basic (VB6) application enables automation of the design process of a "micro-programmed" machine.

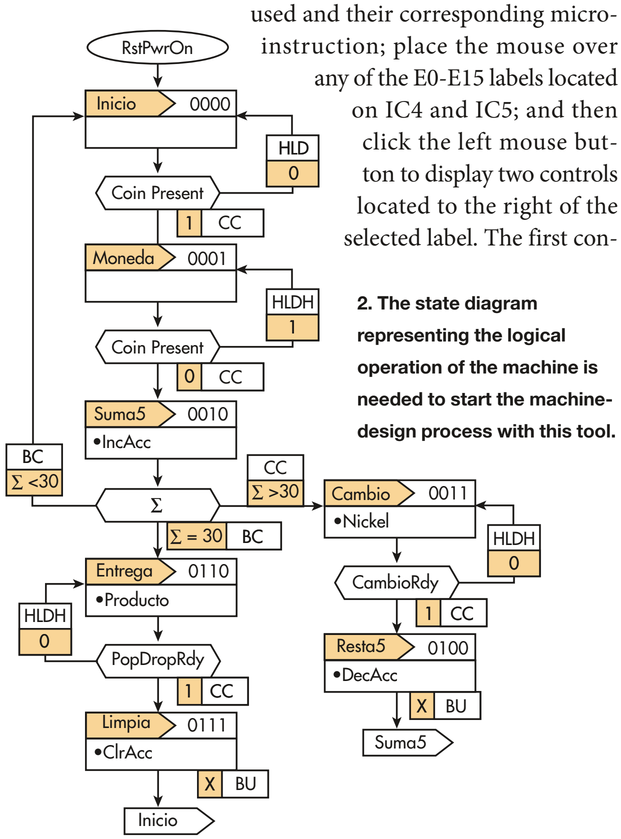

2. The state diagram representing the logical operation of the machine is needed to start the machine-design process with this tool.

trol displays the name assigned to the state, while in the second is the corresponding micro-instruction.

To change the displayed information: Place the mouse over one of the controls; make the necessary changes; and once again click the left mouse button over the control to record changes. When entering or changing the micro-instruction, click the left mouse button on the micro-instruction control and one control box with the possible micro-instructions displayed on the bottom of the control. Select the micro-instructions by clicking the left mouse button and the selected text is placed in the control of Micro-instruction.

In the background screen (which usually is hidden), there's a window titled "Assembler" that contains two text boxes—Code box and Hex box. All information presented in the graphical interface is stored in the Code box: input variables, output variables, names assigned to each of the states, and micro-instructions to move from one state to another. The Hex text box is used to display the HEX file as the result of the machine design generated when compiling the information contained in Code text box through the Assembly menu, as show in Fig. 3.

This tool has been used for educational purposes in Digital Design lectures offered at the School of Engineering of the Universidad Nacional Autónoma de México. However, it can be applied in other fields of knowledge with minor changes.

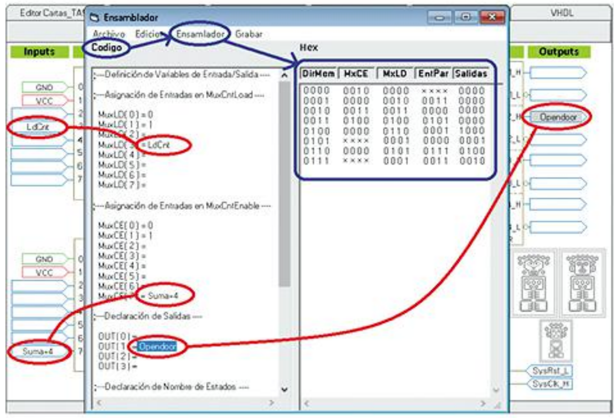

3. This screen is used to enter an input or output variable, beginning with placement of the mouse on one of the corresponding multiplexer ports.

To obtain an executable file of this application with limited functions, send an email request to rtovar_2000@yahoo.com to download it from a Dropbox application folder.

ROBERTO TOVAR and **ROMÁN OSORIO** are at the Institute of Applied Mathematics and Systems, Universidad Nacional Autónoma de México; **TANIA TOVAR** is a postgraduate student in the School of Engineering at the same institution.

"Cap-Drop" Approach Implements Offline Power Supply for Light Loads

By AKSHAY MEHTA | Texas Instruments

LOW-POWER APPLICATIONS such as e-meters (electricity or energy meters) often require a simple ac line-powered supply where a 3.3-V rail powers a microcontroller and charges a lithium-ion battery to 4.2 V. You can implement this with a mains-frequency power transformer or with a more complicated ac-dc (offline) power supply. Both approaches have well-known disadvantages related to a combination of weight, size, and complexity. Two simpler options are a full-wave capacitor-drop circuit with the bridge rectifier *(Fig. 1)* and similar half-wave circuit.

The front end of these circuits is known as a “capacitive-dropper” or “cap-drop” topology. The idea behind both the full-wave and half-wave implementations of this circuit is that line capacitor C1 acts as a lossless resistance, and the reactance of the capacitor will set the maximum current that can be provided to the input of the dc-dc regulator.

The Zener diodes clamp the input voltage to the dc-dc converter under no-load conditions, thus converting the mains voltage to an intermediate dc rail (V_{DC}). The input voltage to the dc-dc converter ($V_{DC} = V_{IN}$) is set to a relatively high value so that the current required from the cap-drop will remain low. You can then down-convert the intermediate unregulated dc rail to the regulated dc voltage required by the load by using a buck regulator with a wide V_{IN} range.

A high step-down ratio allows for a lower input current to the regulator, and is possible with devices such as the LMR14006, LMR16006, and LM46000 buck regulators. This higher ratio enables you to use a smaller value of C1, resulting in lower apparent power drawn from the mains. Applications such as smart-grid e-meters can benefit from this feature because of strict regulations on the maximum apparent power consumption; the typical maximum is limited to 8 VA.

Figure 2 shows the implementation with a half-wave cap-drop circuit. Since the half-wave circuit will disregard the negative cycle of the line voltage, this will result in a lower current delivered to the input of the wide V_{IN} buck topology than in the full-wave circuit. Thus, for applications like battery charging, where a faster charge would require a relatively higher

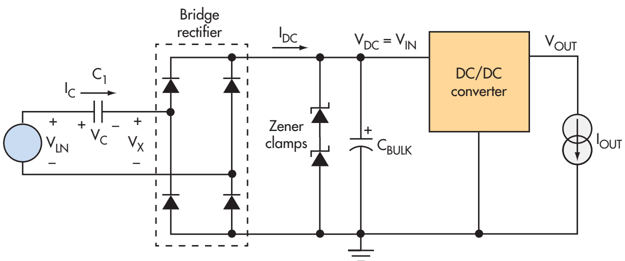

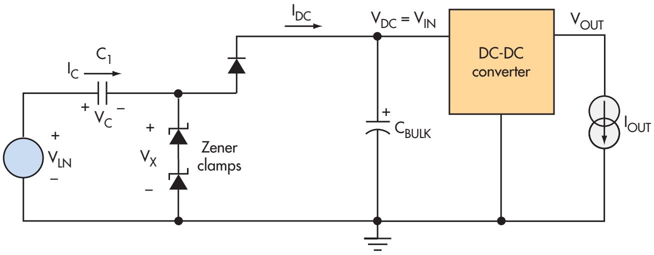

1. No transformer is needed in the full-wave capacitor-drop circuit using a bridge rectifier. The reactance of line capacitor C1 determines the maximum current.

2. The half-wave circuit delivers lower current to the input of the wide V_{IN} buck topology than the full-wave circuit can supply.

load current from the dc-dc regulator, the full-wave circuit is preferable.

The biggest advantage of these circuits is their size. In recent years, smart-grid e-meters have been shrinking in size, severely constraining available printed-circuit-board (PCB) space. To try and install a more conventional ac-dc circuit would not only increase the PCB area, but would also be quite complicated. This increase in board area would also directly relate to the cost. Cap-drop circuits are much more cost-effective, as the C1 capacitor is the only component required to be rated for ac voltages.

While these circuits are easy to configure, you should take utmost care to create a bench prototype and add an appropriate filtering and protection circuit to avoid potentially fatal injuries. The user must ensure that the intended application for this power supply, including its load, is completely isolated from any contact with grounded entities, including people, animals, and test equipment.

AKSHAY MEHTA is a systems and applications engineer with Texas Instruments' SIMPLE SWITCHER product group, where he is responsible for product definition and development, bench validation, and customer support. Akshay received a master's degree in electrical engineering from the University of Texas at Arlington.

(Continued from page 39)

DISTRIBUTED HETEROGENEOUS RUNTIME ARCHITECTURE

The distributed heterogeneous runtime architecture of the Separation Kernel Hypervisor is the most distinguishing characteristic from the centralized homogenous architecture commonly seen in in a Microkernel-based OS.

Building upon the resource isolation, and autonomous runtime domain properties previously described, the Separation Kernel Hypervisor adds in the ability to host a variety independent Operating Systems through CPU virtualization, including Microkernel based OSs. Under this architecture, each independent OS locally manages their own virtual machine contained with a single CPU context, thus implementing Rushby's vison to "create an environment which is indistinguishable from that provided by a physically distributed system."

In the advent of next generation multi-core SoCs, the Separation Kernel Hypervisor can provide the ability to map in all aspects of theses SoC and independently control portions with operating systems and development languages that best suites the need for the target environment. Developers are no longer constrained by a single operating system runtime with limited capabilities, and are now freed to mix and match runtime environments *(Fig. 4)*.

It is important to note that both Separation Kernel Hypervisor and Microkernel run-time models feature a least-privilege design, but the Separation Kernel Hypervisor model opens the door in providing greater evidence of assurance properties and offers tremendous flexibility for end users to construct highly advanced systems where software can take the most advantage out the underlying resources with minimal development efforts.

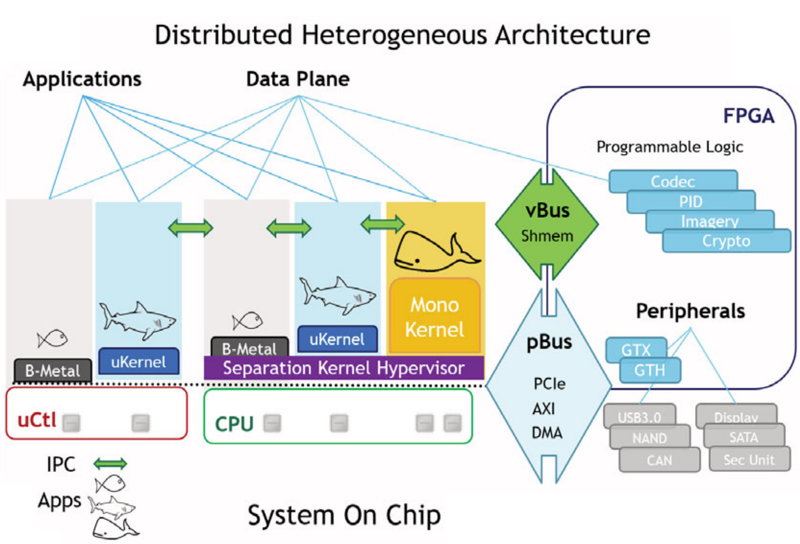

4. Developers are no longer constrained by a single operating system runtime with limited capabilities, and are now freed to mix and match runtime environments.

BARE-METAL RUNTIME

One of the most interesting runtime profiles the Separation Kernel Hypervisor offers for the embedded space is the 'bare-metal' runtime to accommodate a vast set of applications commonly seen in microcontroller designs that do not require the support of the services of an OS. Consider a sensor fusion application that simply requires polling on memory mapped I/O interfaces (*Fig. 5*). With the bare-metal context, these applications can run in a completely raw CPU and flat memory context with zero interference from any other application or kernel services.

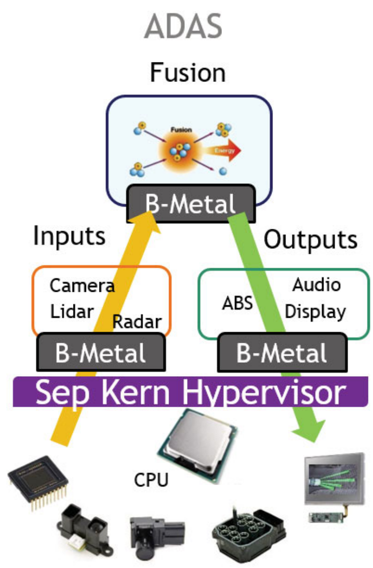

5. With the bare-metal context, applications can run in a completely raw CPU and flat memory context with zero interference from any other application or kernel services.

SUMMARY

Comparing the two classes of technologies is indeed a complicated task. In many respects the Separation Kernel Hypervisor and Microkernels are very similar. Both have minimal code bases, and both can support least privilege architectures. However, looking in more closely, the two technologies aim to be different things, and their least-privilege runtime architectures can be constructed in different ways.

Where Microkernels aimed to provide a safer runtime environment over monolithic kernel based OSs, the Separation Kernel Hypervisor aims to be something different – to not be an operating system. This entails having no centralized system control, and providing a variety of application development and portability options, and offering a runtime architecture that can support strong assurance claims through the utilization of modern CPU virtualization capabilities.

(Continued from page 18)

tinuing to drop as designs evolve. Nowadays, it's more common to find a switcher regulator for low-power applications that reaches nanoampere I_Q values using less external components. The result is a smaller solution with high efficiency that minimizes heat dissipation during active operation.

At Embedded World 2017 in Nuremberg, Germany, Maxim Integrated announced its latest boost regulator. The MAX17222 *(Fig.2)* features 300-nA I_Q at its output and 95% peak efficiency to minimize heat dissipation. The switch regulator also provides a shutdown mode that does not allow any quiescent current from the load after shutoff. With this shutdown mode, it's possible to use alternate sources to regulate the output. The regulator draws just 0.5 nA in shutdown mode.

With switching regulators, Maxim's designers are able to achieve lower I_Qs using less external components. The MAX17222 utilizes a fixed on-time, current-limited, pulse-frequency-modulation control scheme that allows for both continuous conduction mode (CCM) and discontinuous conduction mode (DCM).

Today's designers are increasingly merging pulse-width modulation with a power-saving PFM mode when operating under low loads. Doing so helps to maximize the battery life of wearables. They also are developing smaller and more flexible designs that can further extend battery life in wearables.

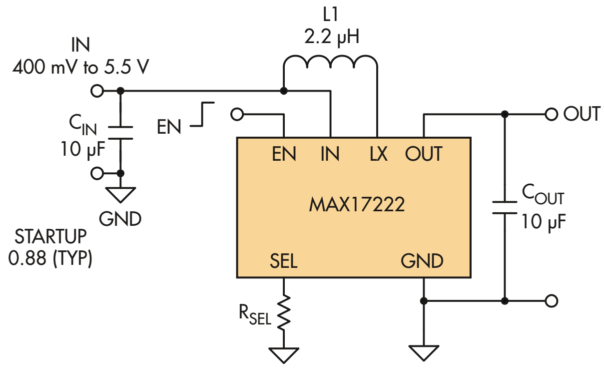

2.The MAX17222 boost regulator requires a single configuration resistor and small output filter. *(Courtesy of Maxim)*

Clearly, many aspects need to be considered when designing a low-power system. Many solutions currently circulate the market, and new options are quickly emerging thanks to evolving design approaches. Depending on the needs of the project, one solution might be better than another. No matter what solution is chosen, however, knowing and controlling quiescent currents will always help to extend a battery's lifetime. ed

Ad	Page
Beta-Layout USA	46
Carlo Gavazzi Automation Components	27
Coilcraft	1
Dean Technology	IBC
Digi-Key	FC, IFC
Edgewood Media Solutions	11
Hammond Mfg. Co. Inc.	39
Ironwood Electronics Inc	46
IXYS	38
IXYS Colorado	30
Kensington Electronics	32
Linear Technology Corporation	32A-B, BC
LPKF Laser & Electronics	42
Memory Protective Devices	29
Micro Crystal AG	41

Ad	Page
PCBcart (General Circuits, Ltd.)	25
Pico Electronics Inc	9
Pico Technology	19
Proto Labs, Inc.	5
Radicom Research	6
Stanford Research Systems	31, 33, 35
Tag Connect	32
TDK-Lambda Americas Inc	43
TE Connectivity	23
Zilog	34
Zurich Instruments	2

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.

ED 5/17 ISSUE PREVIEW

Ad Close: 4/4/17
Materials Due: 4/11/17

TECHNOLOGY:
Automotive Power

INDUSTRY TRENDS
5G
*IMS 6/4-9

PRODUCT TRENDS
Software

*Embedded World, *APEC, *IMS, *EDS * electronica
http://electronicdesign.com

BILL WONG | Embedded/Systems/Software Editor
bill.wong@penton.com

Will ISA Survive?

Is it time for Industry Standard Architecture (ISA) and Peripheral Component Interconnect (PCI) to go away?

How long does old tech need to survive? It might be half an hour in the consumer space, but on the embedded side it is measured in multiple decades, and in the form of interfaces like serial ports and RS-232 *(see "What's The Difference Between The RS-232 And RS-485 Serial Interfaces?" on electronicdesign.com)*.

The Industry Standard Architecture (ISA) Peripheral Component Interconnect (PCI), and Versa Module Europa (VME) are the parallel buses that once dominated the embedded and PC space. But these days PCI's successor, PCI Express (PCIe), runs faster and uses fewer circuit-board traces because of its serial nature. Even Ethernet has joined PCIe on some backplanes.

Form factors like PC/104 have been utilizing multiple interfaces like Diamond System's Aries PC/104-Plus *(Fig. 1)* with an Intel E3800 Atom that supports PCI Express. Compare this to the latest Diamond System's Venus *(Fig. 2)* that has a PCI, two Mini-PCIes, and OneBank Plus sockets. Two of the three interfaces use PCIe. FPGAs or bridge chips are needed to handle interfaces like PCI and ISA when the processor only provides PCIe. Chip availability has become an issue for these parallel platforms, even as they continue to thrive because of the long-term nature of embedded systems.

The challenge is that even availability of compatible support chips is waning despite their prices rising. One VME support chip has been discontinued, which forced many vendors to buy a supply to support their customers for another decade.

The other problem with something like ISA is that it is a 5 V system, so buffers need to provide level shifting. These days it is not uncommon to have this on the processor and peripheral boards that use the latest processor and interface chips. Likewise, peripherals like flash memory typically do not run at 5 V without their own level shifters. The latter take up board space and use more power in an era when both are at a premium.

So, where does your design fit into this discussion? ed

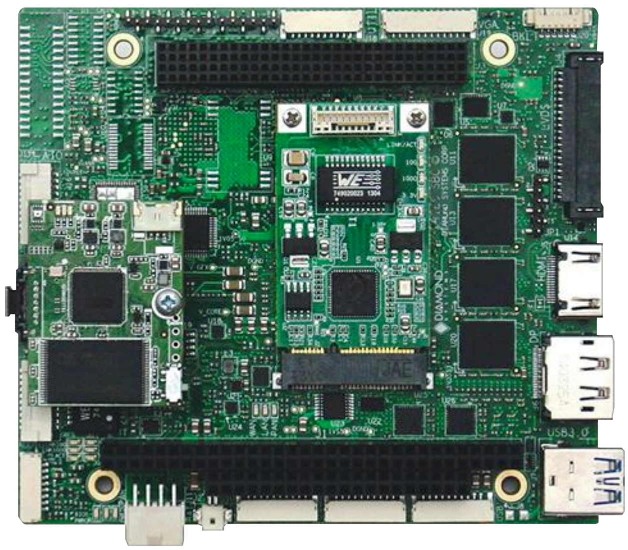

1. Diamond System's Plus uses an Intel E3800 Atom that supports PCI Express, but it has ISA and PCI parallel buses requiring bridge interfaces.

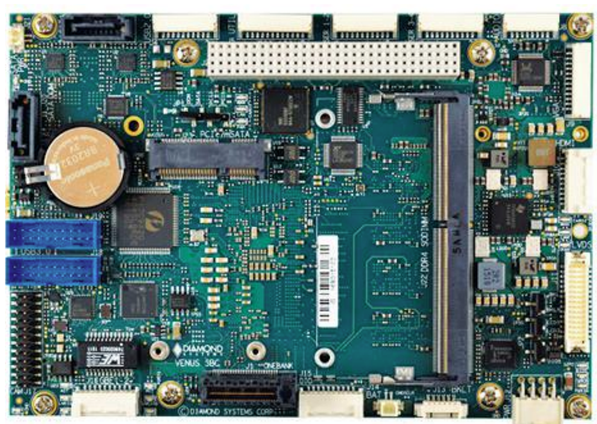

2. The Venus also has PCI, but the two Mini-PCIe and OneBank Plus sockets can take advantage of the processor's native PCIe support.

200mA “Refulator”

10ppm/°C Max Drift • ±0.05% Accuracy • 1.5ppm$_{P-P}$ Noise

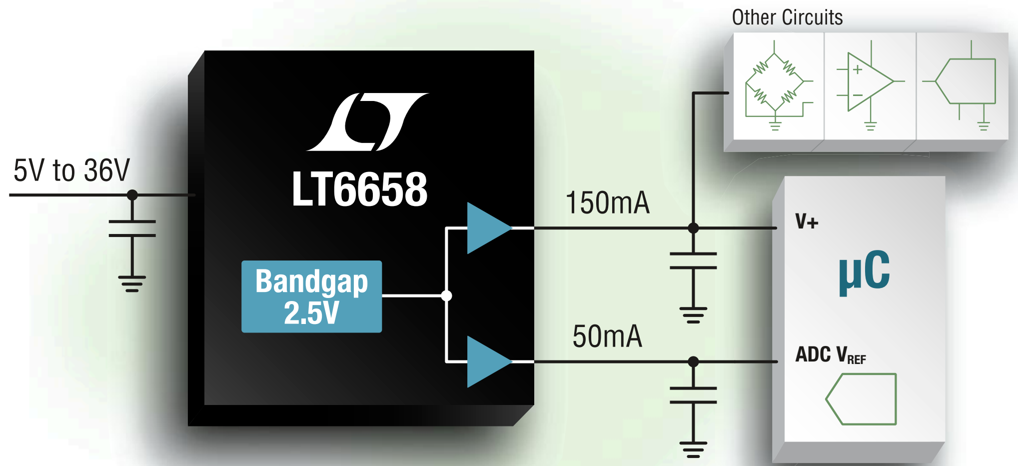

The LT6658 expands the capability of precision voltage references by including dual output buffers capable of sinking 20mA, and driving 150mA and 50mA, respectively. Six options are available with different internal reference voltages, and each output voltage can be adjusted using external resistors. With two outputs, the LT6658 can be operated as a reference and regulator, as a reference plus a virtual ground, as a dual tracking ratiometric supply, or combined as a single 200mA voltage reference.

Features

- Dual Output Tracking Reference:
  - Output 1: 150mA Source/20mA Sink
  - Output 2: 50mA Source/20mA Sink
- Low Drift:
  - A-Grade: 10ppm/°C Max
  - B-Grade: 20ppm/°C Max
- High Accuracy:
  - A-Grade: ±0.05% Max
  - B-Grade: ±0.1% Max
- Low Noise: 1.5ppm$_{P-P}$ (0.1Hz to 10Hz)
- Load Regulation: 0.1ppm/mA
- Voltage Options: 1.2V, 1.8V, 2.5V, 3V, 3.3V & 5V

Load Regulation

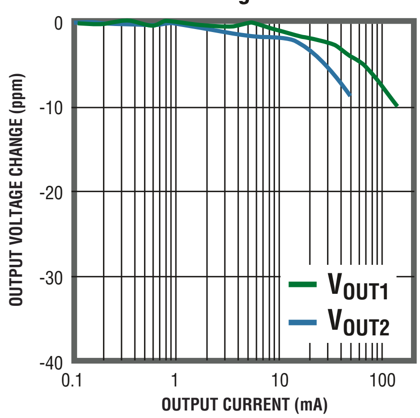

Info & Online Store

www.linear.com/product/LT6658

1-800-4-LINEAR

video.linear.com/6035

LINEAR TECHNOLOGY | NOW PART OF ANALOG DEVICES

Find your local sales office: www.linear.com/contact