Electronic Design®

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

JULY 2017 electronicdesign.com

AUTO FOCUSED

Technological advances in today's vehicles show no signs of stopping p12

$10.00

Powered by Penton

"You guys are big enough now that if you don't stock a part, it's a safe bet that I shouldn't design around it. You guys are great. Keep up the good work!"
Manufacturer, Digi-Key Customer since 2005
What Our
Customers
Are Saying...
"Service and support are excellent, had to speak with reps to change orders, they are always helpful and professional."
Aerospace Engineer, New Digi-Key Customer in 2016
"The simplicity of your website and the ease to use it outclasses and blows every other distributors' out of the water. Keep up the good work!"
Geophysics Engineer, Digi-Key Customer since 1997
World's Largest Selection
of Electronic Components Available for Immediate Shipment!
SAME-DAY SHIPPING
on orders entered by 8:00 pm CENTRAL
NEXT-DAY DELIVERY
1.800.344.4539
DIGIKEY.COM
Digi-Key ELECTRONICS
6 MILLION PARTS ONLINE | 650+ INDUSTRY-LEADING SUPPLIERS | 100% AUTHORIZED DISTRIBUTOR
Digi-Key is an authorized distributor for all supplier partners. New products added daily. Digi-Key and Digi-Key Electronics are registered trademarks of Digi-Key Electronics in the U.S. and other countries. © 2017 Digi-Key Electronics, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA
ecia MEMBER
ecsn member
CEDA

July 2017 VOLUME 65, ISSUE 7

IN THIS ISSUE

12

16

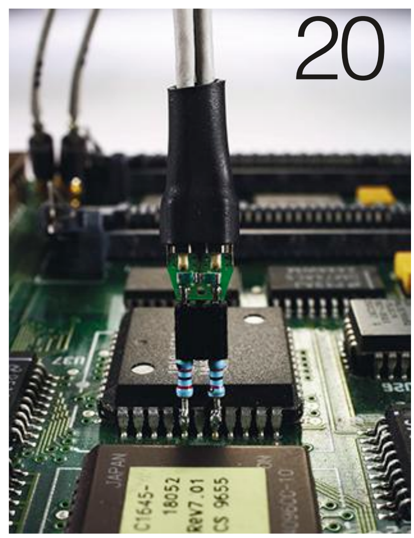

EDITORIAL MISSION:
To provide the most current, accurate, and in-depth technical coverage of the key emerging technologies that engineers need to design tomorrow's products today.

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

Build Your Own Drone on Intel's Platform

Intel's presence at this year's 2017 Xponential show was significant, including CEO Brian Krzanich as the keynote speaker. The company has been active in buying up and partnering with a number of vendors in this space, and tying them in with other major investments such as augmented and virtual reality as well as artificial intelligence.Intel has highlighted its "Shooting Stars" multi-drone technology at events like Lady Gaga's Super Bowl halftime show, where 300 drones put on an impressive light show.

http://www.electronicdesign.com/automotive/build-your-own-drone-intel-s-platform

Meet the Winners of the 2017 Formula Hybrid Competition

The Formula Hybrid Competition is an interdisciplinary design and engineering challenge that targets undergraduate and graduate university students. Focusing more on the engineering side than on racing, the teams must collaboratively design and build a formula-style electric or plug-in hybrid race car and subsequently compete in a series of events. The competition is part of the Society of Automotive Engineers (SAE) Collegiate Design Series.

http://www.electronicdesign.com/power/meet-winners-formula-hybrid-competition-2017

Forging a New Distribution Model for Innovators

By providing essential supplies and support, distributors of technology components and services are the engines that keep business moving for electronics innovators. However, in the past five years, these distributors have not only seen their customers' needs change, but also experienced an expansion of their market to include new types of buyers. These two factors prompted a shift in buying behavior, which has ricocheted up and down the electronics-component supply chain.

http://www.electronicdesign.com/industrial-automation/forging-new-distribution-model-innovators-hobbyists-manufacturers

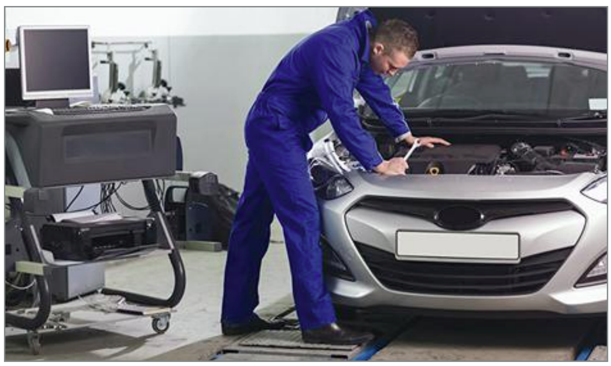

Will Right to Repair Affect Your Design?

Remember when devices were repairable? Many still are, but an increasing number are not. These days, in fact, devices are designed never to be repaired. This is especially true for many Internet of Things (IoT) devices, like bridge pressure sensors designed to be installed when a bridge is constructed, or sensors designed to run in the field for more than 20 years.

http://www.electronicdesign.com/embedded-revolution/will-right-repair-affect-your-design

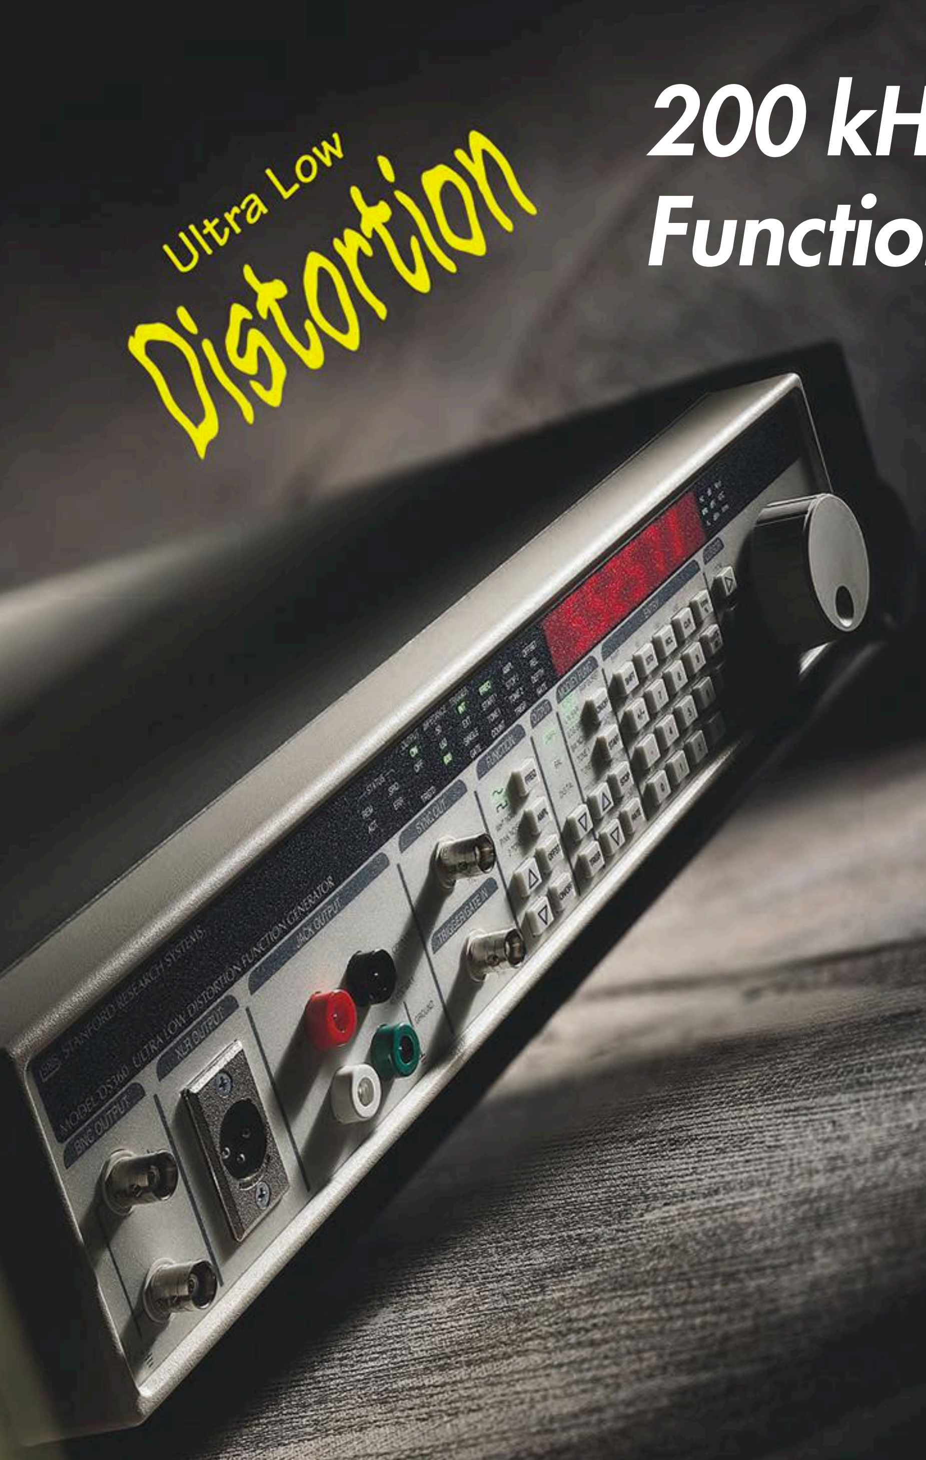

200 kHz DDS Function Generator

0.001 % THD

(–100 dB)

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

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

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

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

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

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

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

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

NANCY FRIEDRICH | Director of Editorial Content
nancy.friedrich@penton.com

Our Schools Need You

Technology will increasingly change the way we work, which means it's crucial that we revisit how we're teaching and training our kids to build the next generation of workers.

Charter schools aside, most U.S. public schools—primary through high school—still approach education in much the same way that they did for the last generation, and the generation before that, etc. Sure, there's the new Common Core math approach in use today, which is widely mocked and denigrated. Another change is that cursive and typing classes have mostly gone extinct, although typing could still be useful (unless you assume everything will eventually operate via voice control).

Largely, however, children are taught the same basics at the same time—even if they're taught in a slightly different way. As I help my own children with their homework and look at their familiar curriculums with a sense of nostalgia, I wonder: Is this the approach we should still be using to teach them?

As a career technology editor, I certainly am not an expert in education. Yet I'm well aware of today's engineering developments and how they continue to impact how we live, work, and create. We exist in a moment—it's hard to tell precisely how long it will last—right before major medical, manufacturing, transportation, lifestyle, and other marvels drastically change the world (think of examples like driverless cars and robotic surgery).

In the face of such dramatic change, however, much of what our children learn in kindergarten through 8th grade is still the same rudimentary building blocks: Starting off with learning colors, numbers, and letters and moving on to counting, reading and writing, multiplication and division, earth science, history, reading, etc. I'm not saying we should necessarily be teaching toddlers to code, but maybe engineering- and technology-related concepts should make an earlier and more regular appearance in the curriculum?

Many schools do try to add new programs. The problem is that public schools in particular are not set up to have much fluidity around their curriculums. Budgets are fixed in most areas, and many teachers and administrators, while excellent at what they do, aren't well-versed in today's fast-moving technology. This is where you come in—well, not just you, but any engineering or science-minded person who could offer advice on a basic science lab, educate a STEAM teacher on things like building a basic circuit, or offer some other contribution (calling all makers and hobbyists!).

For instance, many elementary schools are starting with a community-assisted effort like a student makerspace. And they're often looking for donations—everything from Legos to shoeboxes and batteries, old toys and games, etc. So if you have anything to donate, please reach out to your local schools and see if they're interested. But start with the elementary schools first. High schools often have some good stuff already in place; the elementary schools are the ones that need a boost. Do some summer cleaning and reach out. And if you have any knowledge you can share, see if there's a way for you to do that. Your employer may even want to sponsor a program by providing kits or some tools as a donation.

To better keep pace with technology, our schools need support, aid, and donations—of things and knowledge, not necessarily funds—to help children grow up with technology as a part of the standard curriculum, just as it has become a standard part of life. If you've been thinking maybe you should "give back" or have extra time on your hands, see what resources you can gather or solicit and get involved. Even if you just end up passing along directions on a simple starter kit of some kind with some materials, it will be appreciated. We'd love to spotlight STEAM programs that are already making a difference, so if you know of any, send me a note at nancy.friedrich@penton.com. ed

News

ACHRONIX'S SUCCESS
Reflects Growth For FPGA Accelerators

Achronix Semiconductors recently revealed that its revenue had grown sevenfold over the last year and will surpass $100 million this year. And there are lessons in those figures.

The recent gains reflect the growing market for FPGAs and other chips that act like accelerators for tricky tasks like classifying images to autonomous driving.

As improvements in traditional CPUs have slowed, Nvidia has been leading the push into accelerators for servers and self-driving cars. Nvidia's chief Jensen Huang has been an evangelist for its graphics chips, which can be used to train and execute machine learning code. Nvidia's stock price tripled to $150 in the last year.

Achronix's recent success shows that FPGAs remain in the conversation. Though difficult to reprogram precisely, the chips can do the same parallel processing as graphics chips. Nvidia's chips are still the gold standard for training machine learning models, but cloud giants like Microsoft and Baidu are using FPGAs to find patterns in large amounts of data.

The competition for that business includes Xilinx and Intel, which last year acquired Altera for $16.7 billion. The world's largest maker of server chips plans to package Altera's FPGAs with CPUs in future platforms and later squeeze them onto the same silicon. Xilinx has been releasing software stacks to help cut down the complexity of FPGA design.

Achronix reads from a similar playbook. "[We] are entering a new high growth era where our customized core FPGA technology can accelerate a broad range of complex compute tasks in machine learning, artificial intelligence, software defined networks and 5G base stations," said Robert Blake, Achronix's chief, in a statement.

Its products include Speedster 22i FPGAs, which contain large amounts of SerDes and I/O to interface with a central processor using PCIe interconnects. The 22nm chips come with 86 megabits of embedded memory and support up to a million look-up tables, also known as LUTs, which are central to an FPGA's processing power.

Achronix's revenue largely comes from its Speedster products, which first went into production in 2015 and will be refreshed in 2018. But the fabless supplier has increasingly split its business into new directions: Last year, it started licensing the intellectual property behind Speedster.

The blueprints, sold under the brand Speedcore, let engineers integrate 16nm FPGAs into custom system-on-chips, also known as SoCs. These embedded FPGAs are faster and more power efficient than standalone chips because they share the same silicon as computer processors and other circuitry.

With Speedcore chip designers can change the number of LUTs as well as embedded memory and DSP blocks to suit specific applications. Before the end of the year, Achronix will allow customers to work with the next generation of Speedcore based on TSMC's 7nm production process.

The moneymaking opportunity in embedded FPGAs is not secret. Flex Logix, a newcomer run by the former chief executive of Rambus, has raised almost $14 million to sell its eFPGAs for networking hardware and data centers, and has recently won contracts with DARPA and Harvard. Last year, QuickLogic began licensing eFPGAs to be used in wearables and sensors.

Achronix's victory lap came after years of promise. The company had been running off $132 million in venture funding since its 2004 founding, defying investor apathy for the chip industry. But it finally hit profitability this year and plans to expand its 75-person staff by a third over the next six months. ■

AFTER SKIPPING A STEP, GlobalFoundries Prepares For 7nm Production

LAST YEAR, GlobalFoundries said that it had skipped the 10nm chips that other contract manufacturers consider the next major step after current 14nm technology.

Instead, the company bet that 7nm chips will provide more meaningful improvements and boost its competitiveness with rivals like TSMC and Samsung. It recently started giving customers the option to design chips for applications like smartphones, servers, and networking to be manufactured in its 7nm production process.

GlobalFoundries said that the new technology provides around 40% more processing power and twice the area scaling than previous 14nm chips. The company plans the first products using 7LP production to appear in the first half of 2018, while volume production will begin in the second half at the Fab 8 plant in Saratoga County, N.Y.

GlobalFoundries has "multiple product tapeouts planned in 2018," said Greg Bartlett, senior vice president of the CMOS business unit, in a statement. Other foundries only started selling 10nm chips in the last few months.

The rejection of 10nm chips underlines both the harsh realities of the chip industry and the growing competition for foundry business. Fewer companies than ever can afford to spend billions of dollars on the backbreaking research and factories required to make the smallest and fastest chips.

Those that can muster the capital include Samsung, which recently spun out its foundry business into a new unit, and Intel, which is increasingly willing to take manufacturing orders. TSMC can also afford to develop new technology nodes: its manufacturing contracts with Apple and others last year came out to $29.5 billion, up from $24.1 billion in 2014.

These foundries have stuck largely to the script. In March, Samsung said that it had already shipped 70,000 wafers of 10nm technology, while TSMC said that it plans to rapidly increase 10nm output later this year. Intel waited longer to reveal 10nm technology but claimed that it is equivalent to rival 7nm chips.

In March, Intel divulged that its 10nm chips will contain 100.8 million transistors when shipped later this year. The spacing between the transistor fins measures only 34nm, down from the 42nm in previous 14nm chips. The shorter distances mean that more transistors can be packed into a given area.

But advances like Samsung's and Intel's could be short-lived. TSMC has announced plans to complete several 7nm tapeouts this year, while Intel is building a pilot plant for testing 7nm chips. Last month, Samsung said that it intends to start producing 8nm technologies this year, 7nm next year, and 4nm chips in 2020.

"By biting the bullet and skipping 10nm, GlobalFoundries opened up the technical bandwidth to attack 7nm head-on. Others have been dividing their resources and going for half- or even quarter-nodes," said Dan Hutcheson, chief executive of VLSI Research, in a statement.

The foundry has also announced a new 7nm platform for application-specific integrated circuits, also known as ASICs. The chips are intended for applications in data centers, wireless networking, and machine learning. The ASIC platform will enter full production in 2019.

For now, GlobalFoundries will create 7nm chips using optical lithography to tattoo circuits onto layers of silicon. But that will shift to extreme ultraviolet lithography – which takes fewer steps and can etch smaller circuits – for volume production. The company plans to install two EUV tools this year.

GlobalFoundries started making 14nm chips last year out of its Fab 8 plant, whose capacity it plans to expand 20% by early next year. The company now has more than 20 customers for the technology including AMD, which is using the production process for its Polaris graphics chips.

There has been much hand-wringing over 5nm chips, which will require creative new ways to prevent current leakage from transistors. But along with IBM and Samsung, GlobalFoundries recently landed a coup, proving that it was feasible to create 5nm chips using silicon nanosheets surrounded by gates to prevent leakage.

A "clean house" at GlobalFoundries, where the company manufactures its semiconductor chips.

Some industry analysts have been pleasantly surprised by the recent announcements. "I don't think many would have called that GlobalFoundries would be on a bit of a roll right now," Patrick Moorhead, principal analyst for Moor Insights and Strategy, wrote in a recent column. ■

RESEARCHERS TEST AN OPTICAL CHIP For Deep Learning

THE COMPUTER CHIP learned to recognize four basic vowel sounds, and it guessed correctly three out of every four times when tested. Other processors are more accurate—correct around 90% of the time—but few are as unique.

Researchers at the Massachusetts Institute of Technology revealed that new chip, called a nanophotonic processor, which uses light instead of electricity to solve the unforgiving math at the heart of machine learning. Using microscopic lenses instead of transistors, it could vastly improve algorithms that learn to make decisions in smartphones to sensors.

The ten researchers, writing in the journal *Nature Photonics*, said that the experimental chip could be carefully tuned to control light and solve so-called matrix multiplications much faster and more efficiently than traditional computers and graphics chips. The idea is that light moves much faster than the electrons cascading through transistors.

These calculations are vital to deep learning algorithms, which mimic how the brain learns from an accumulation of examples to improve voice recognition, image classification, or autonomous driving software. But these multiplication tables not only take the most time but also consume the most power.

The optical chip contains multiple waveguides that shoot beams of light at each other simultaneously, creating interference patterns that correspond to mathematical results. The proof-of-concept performs matrix calculations with a thousandth of the power used by traditional chips, the researchers said.

"The chip, once you tune it, can carry out matrix multiplication with, in principle, zero energy, instantly," said Marin Soljacic, an electrical engineering professor at MIT, in a statement. The accuracy leaves something to be desired for making out vowels, but the nanophotonic chip is still a work-in-progress.

Years from now, it could be useful "whenever you need to do a lot of computation but you don't have a lot of power or time," said Nicholas Harris, an author of the paper, in a statement. It could also improve signal processing because light, which is analog, would not have to be converted into a digital signal.

The nanophotonic chip is still not a complete system, and Soljacic said that further advances would only be possible with more time and investment. "We've demonstrated the crucial building blocks but not yet the full system," he said in a statement.

Other platforms have targeted matrix multiplications, which help strengthen links between the virtual neurons in deep learning software. The Volta graphics chip released by Nvidia last month, for instance, contain specialized "tensor cores" that solve matrix multiplication for training and inferencing faster and more efficiently than other chips. ■

ANALOG DEVICES Shuffles Executives

THREE MONTHS AFTER closing its $14.8 billion acquisition of Linear Technology, Analog Devices has been quiet about internal changes. But in mid-June, the company shuffled its executive ranks to accommodate its new analog and power businesses.

The Massachusetts-based supplier also added three new business units, which should fit its strategy to sell electronics for smarter cars and factory equipment. Analog Devices "has long been known as a living company founded on continuity and thoughtful evolution," said Vincent Roche, its chief executive, in a statement.

Steve Pietkiewicz is the only holdover from Linear Technology to be promoted to Analog's executive ranks. Pietkiewicz, the former vice president for power management products, seems to have an almost identical role as Analog's senior vice president of power products. He held senior engineering and executive roles for 30 years at Linear Technology.

It was not immediately clear if changes further down the corporate ladder had occurred but not been announced. This year, Analog Devices' shareholders voted Robert Swanson, the founder, former chief executive, and former chairman of Linear Technology, to its board of directors.

Swanson founded Linear Technology in 1981 along with Robert Widlar and Robert Dobkin, its former vice president of engineering and chief technology officer. He helped build it into one of Silicon Valley's most profitable chip makers and something of an apprentice program for analog engineers.

The latest changes suggest that Dobkin, 74, will have a reduced role at Analog Devices or none at all. After the announcement of Swanson's board seat in January, a spokeswoman said that Analog Devices was "still determining the roles of Linear Technology's executives upon completion of our acquisition."

Analog Devices is one of the largest makers of analog chips that convert the surrounding world into electronic signals, allowing industrial robots to sense temperature changes or networking equipment to process radio waves. But the Linear Technology deal expanded it into chips that regulate power within automobiles and other devices and control signals in assembly lines.

The deal is Analog Devices' biggest bet yet that these chips are a winning combination for the growing Internet of Things. It is still too early to know if the acquisition will succeed, but the company last month reported that it reaped $1.1 billion in second quarter revenue, up from $778 million. ■

WILLIAM WONG | Technical Editor

The Changing Nature of Automotive Technology

From advanced sensors to artificial intelligence, vehicles of all types are quickly becoming home to the latest electronics technology.

The transportation space has seen a burst of technology—not in one particular area, but rather across the board from improvements in electrical power systems to extremely sophisticated telematics to self-driving cars. Cars today have more electronics that ever before. Much more is coming, though, as features such as advanced driver-assistance systems (ADAS) become standard features instead of expensive options.

These changes are being made possible by improvements in sensors, processors and memory, software, and even human interfaces that need to be integrated in real time *(Fig. 1)*. Here are some of the latest technologies and how their relationship with other technologies makes them even more important in automotive environments.

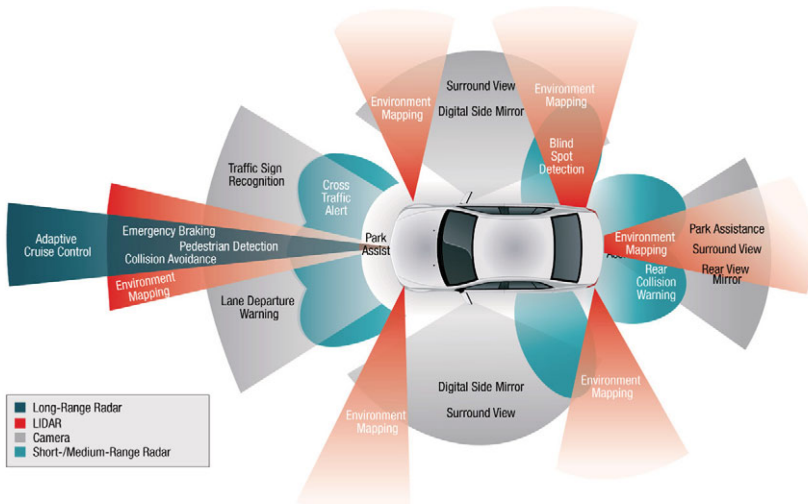

1. Multiple, overlapping sensors are needed to provide information for systems to build situational awareness in order to implement safety-critical ADAS support.

SENSOR ADVANCES

Smartphones have turned tiny digital cameras into commodity items in a way that other applications—digital cameras, for instance—could not. Automotive applications continue to benefit from the availability of cameras that can stream 4K video. High-definition cameras are being used for obstacle and object recognition for forward-looking ADAS applications in conjunction with artificial-intelligence (AI) machine-learning (ML) software. Here the higher resolution is important, and it's useful for backup cameras, too.

Multiple cameras are also being used to provide a birds-eye view around the car. Renesas' R-Car development kit knits together video streams from four cameras into a 360-degree view *(Fig. 2)*. This is very useful when parking or navigating in tight quarters. More advanced ADAS systems highlight areas of potential oncoming collisions.

2. Renesas' R-Car SoC is able to generate a 360-degree, birds-eye view around a vehicle by knitting together video streams from four cameras.

Two other range sensors that have shown significant improvements lately are LiDAR and phased-array radar. The general technology is not new, but major advances in miniaturization and cost reductions will affect when and where these systems are being utilized.

For example, Innoviz *(Fig. 3)*, LeddarTech, Quanergy, and Velodyne are just a few companies delivering 3D, solid-state LiDAR systems. These systems, which are applicable in other areas like robotics *(see "Bumping into Cobots")*, are getting so small that multiple units will be hidden around a car.

3. Innoviz is just one of many vendors delivering 3D LiDAR technology. The InnovizOne has a 200-m range with better than 2-cm depth accuracy. It maintains a 100- by 25-degree field of view with 0.1- by 0.1-degree spatial resolution. The device delivers 25 frames/s with a 3D resolution rate of over 6 Mpixels/s.

Phased-array radar overcomes many of the limitations of LiDAR, allowing it to operate in rain and snow that can otherwise fool optical systems. Radar can be used to complement LiDAR and image systems. A number of companies are working to deliver technology in this area. For example, Texas Instruments' (TI) single-chip millimeter-wave sensor, mmWave, handles 76- to 81-GHz sensor arrays for sensor and ADAS applications *(see "Low-Cost Single Chip Drives Radar Array")*.

All of these technologies have applications in other areas from manufacturing to security and even 3D scanning and printing.

SOFTWARE ADVANCES

AI and ML are garnering the limelight these days because they bring efficient image recognition to ADAS that's critical for safe self-driving or augmented driving experiences. The underlying technology is based on deep neural networks (DNNs) and convolutional neural networks (CNNs) *(see "What's the Difference Between Machine Learning Techniques?")*.

Neural networks will not replace conventional software applications, even in automotive environments, but they solve hard problems. Combined with new hardware, they can also do it in real time, which is needed in safety-critical applications such as self-driving cars. Multicore processors help in this case, but GPUs work even better *(Fig. 4)*. Custom hardware bests them *(see "CPUs, GPUs, and Now AI Chips")* all, and even specialized digital-signal processors (DSPs) can handle machine-learning chores *(see "DSP Takes on Deep Neural Networks")*.

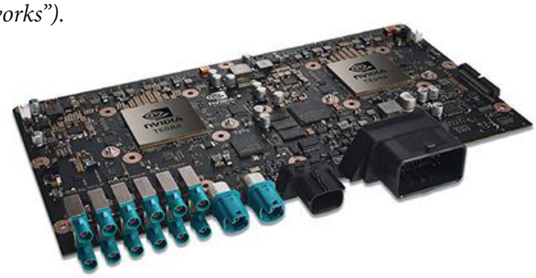

4. The Drive PX2 from Nvidia is just the latest of a series of multicore CPU/GPU solutions targeted at automotive applications.

The parallel-processing nature of these solutions plays well to the multicore and transistor count growth in designs, even as upper-level clock frequencies have peaked. The more tailored solutions also have lower power requirements compared to more conventional processor solutions.

The in-vehicle infotainment (IVI) system advance is changing what drivers and passengers are able to visualize, as well as how they can link their smart devices and cloud-based applications to their car. Cellular-based Wi-Fi hot spots in a car are available from all vehicle manufacturers. The plethora of options requires a more robust and open approach. On that front, the GENIVI Alliance *(see "Automotive Technology Platform Developed for Linux-Based Systems")* fosters open standards that are operating-system agnostic.

The The Linux Foundation's Automotive Grade Linux (AGL) is one example of an IVI system that has received wide vendor support. AGL will be used in Toyota's 2018 Camry *(Fig. 5)* as well as future Toyota vehicles *(see "Toyota Including Automotive Grade Linux Platform in 2018 Camry")*.

5. Toyota's 2018 Camry will be running Automotive Grade Linux (AGL) for its in-vehicle infotainment (IVI) system.

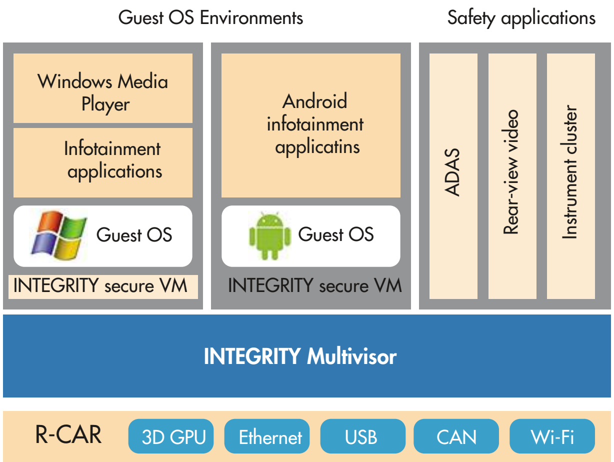

6. Green Hills Software's Multivisor provides the virtual-machine isolation of a Type 1 hypervisor, which is becoming commonplace in automotive environments that host safety and critical subsystems on the same hardware as IVI subsystems.

The number of applications and tasks running on automotive systems can be staggering when one considers the amount of information being produced from the large collection of sensors, to the data processed and generated by AI systems, to streaming video moving over in-vehicle networks. Managing data distribution in safety-critical areas can benefit from standards like the Object Management Group's (OMG) Data Distribution Service (DDS) that can provide secure, real-time, publish-subscribe managed data exchange throughout the system *(see "Should DDS be the Base Communication Framework for Self-Driving Cars?")*. This approach scales better than many point-to-point solutions typically found in designs that require fewer connections between applications.

The hypervisors is another tool that has been around for a long time, but not in vehicle-control settings *(see "What's the Difference between Separation Kernel Hypervisor and Microkernel?")*. However, that's changing, and in a big way, due to the number and complexity of multicore solutions as well as the need to mix safety- and security-critical components with IVI and non-critical systems. Hypervisors targeting the automotive space are available from a number of vendors like Blackberry QNX Hypervisor, Wind River VxWorks, Green Hills Software INTEGRITY Multivisor *(Fig. 6)*, and Mentor Graphics Embedded Hypervisor.

Hypervisors allow partitioning of virtual machines (VMs) such that safety and security certifications can be similarly divided. This means that non-critical components don't also require the same level of certification that takes times and is very costly. Likewise, the safety- and security-related components that normally have limited third-party additions over time are becoming more common in IVI.

Type 1 hypervisors like Blackberry's QNX Hypervisor 2.0 are designed to be small with a low memory and performance overhead, but features can be critical to performance and security *(Fig. 7)*. QNX provides priority-based virtual CPU (vCPU) with a configurable scheduling policy. The hypervisor is based on the QNX SDP 7.0 RTOS, and thus provides fine-grain management and security. The QNX Neutrino RTOS is a likely candidate for VM that will need safety certification. The QNX OS for Safety is certified for ISO 26262 at ASIL D and IEC 61508 SIL3, and has been used in systems certified to EN 50128 at SIL 4.

(Continued on page 40)

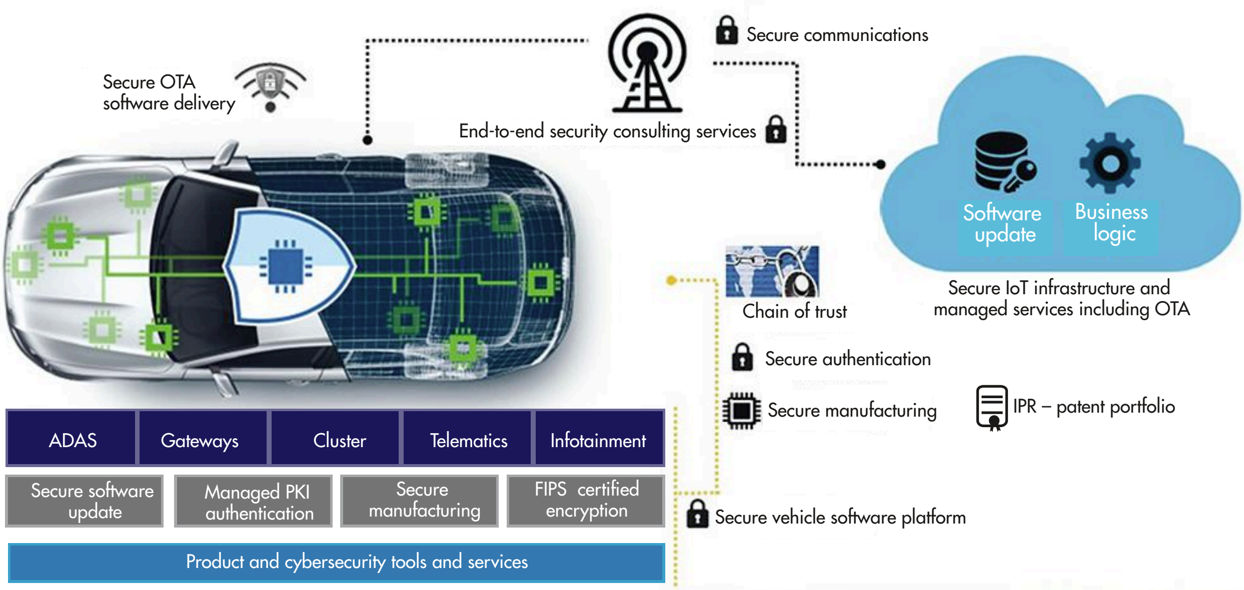

7. Security permeates all aspects of the automotive environment, from manufacturing through secure over-the-air updates. Coordinating and supporting this infrastructure can be challenging, although companies like Blackberry are providing a complete solution.

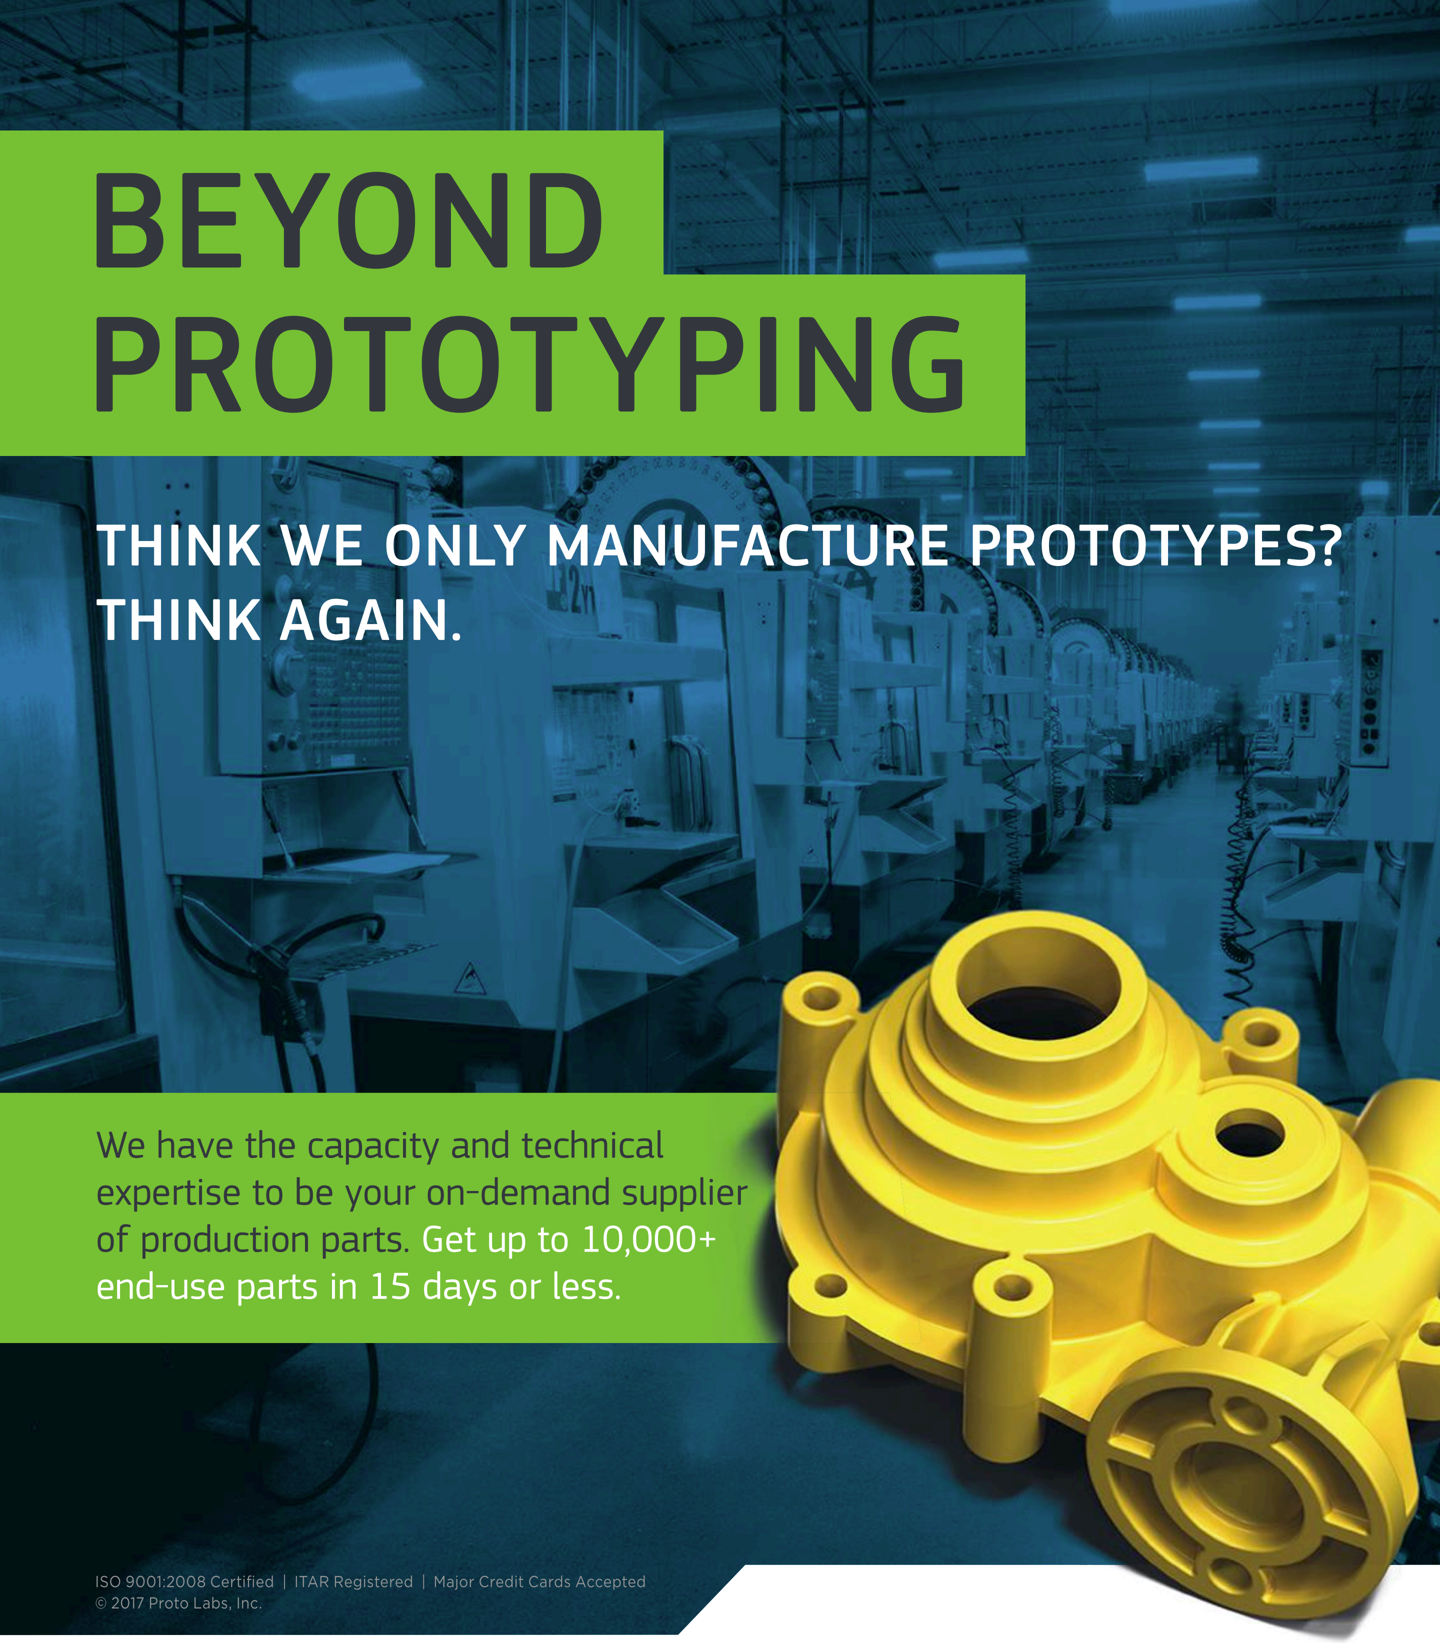

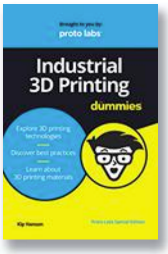

proto labs
Industrial 3D Printing
dummies

MARIA GUERRA | Analog/Power Editor

The Power of IoT Devices

Multiple function integration in power management integrated circuits is maximizing the battery life of Internet of Things devices.

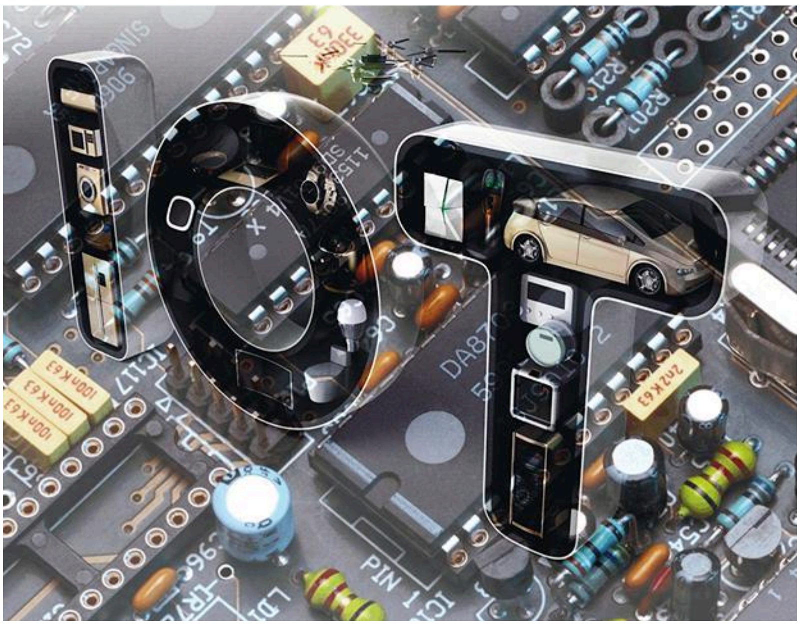

According to IC Insights total semiconductor sales for Internet of Things (IoT) systems are expected to reach $31.1 billion in 2020, with the IoT semiconductor market for wearable systems expected to show a CAGR of 17.1%. This proliferation of devices is also creating future growth in the power management IC (PMIC) market.

Designers of IoT solutions are relying on power management solutions to efficiently handle the power needed to energize a wide range of IoT devices, as maintenance and battery replacement are not cost-effective approaches. There are many power management solutions currently in the market, and depending of their characteristics, one power management solution can work better for a specific application depending on power conversion and power control options.

PMICs can now reduce power consumption of the batteries prolonging the power in IoT devices while also minimizing PCB size. PMICs are now achieving lower quiescent currents that helps to increase battery life. Maxim Integrated's new MAX20310, for instance, is a power management solution that operates with battery voltages down to 0.7V for use with single-cell zinc air, silver oxide, and alkaline batteries. This solution includes a dual-output, programmable, micro-IQ high-efficiency switching converter with a 0.5µA quiescent current. It focuses on wearable medical devices and fitness applications such as non-rechargeable medical patches, environmental and equipment monitoring, and discrete sensors for the Industrial Internet of Things (IIoT).

Some other PMICs come with integrated cold-start functionality that help to ensure the operation of the device even with low sources of ambient energy, like the ones use when harvesting energy from ambient sources (e.g., vibration, light, and RF). For example, a Belgian company called e-peas developed its AEM10940 PMIC *(see figure)* specifically to achieve more efficient energy harvesting for embedded electronics. Notable applications include very-low-power IoT wireless sensors, portable technology, domestic automation, and industrial control.

The AEM10940 harvests the available input power from 1 µW to 50 mW. It integrates an ultra-low-power boost converter to charge a storage element, such as a rechargeable battery, a capacitor, or a supercapacitor. It contains a boost converter that operates with input voltages in a range of 100 mV to 2.5 V. The cold-start functionality can start operating with an empty storage element at an input voltage as low as 380 mV and an input power of just 11 µW. As soon as a small amount of energy becomes available at the source pin, the Cold Start function raises V_{BOOST} up to a voltage of 2 V.

As energy harvesting applications gain traction, the number of semiconductor companies offering evaluation kits or

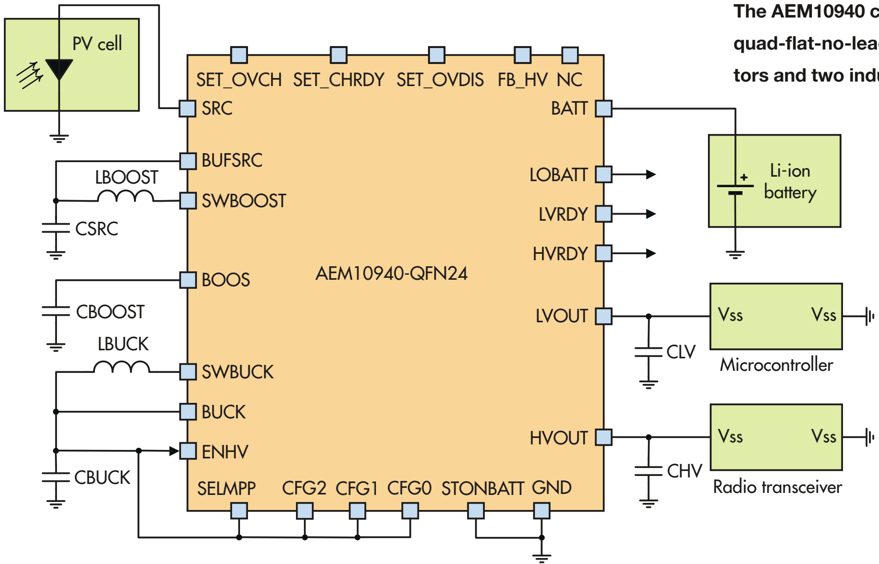

The AEM10940 comes in a 24-pin packaging option using quad-flat-no-leads packages (QFN). Five identical capacitors and two inductors are required. *(Courtesy of e-peas)*

demo boards to explore the performance of PMICs functions built for energy harvesting applications continues to grow. For example, engineers can evaluate Cypress' S6AE102A and S6AE103A PMICs, which can use solar energy to charge a super capacitor, along with a sensor board to analyze magnetic door sensor or ambient light sensor operation.

The CYALKIT-E04 kit can be used with Bluetooth Low Energy (BLE) wireless connectivity solutions from Cypress for IoT applications. The PMICs of this kit enable ultra-low power operation with quiescent currents of only 280nA and startup power of only 1.2µW.

IoT designers are now able to find more power management integrated circuit options for their designs in a wide range of applications spanning the automotive, consumer, and industrial markets. Future PMICs will aim to operate at lower voltages to provide a wide range of better reliable solutions to the harvesting energy market, which is rapidly growing as an acceptance of renewable energy resources. ed

DANIEL BOGDANOFF | Keysight Technologies Inc.

How to Pick the Right OSCILLOSCOPE PROBE

Many probes populate the T&M market—deciding on the best option requires a good understanding of bandwidth, attenuation ratio, and loading specs.

Literally hundreds of different oscilloscope probes are available *(Fig. 1)*, so how do you choose the right one? There's no single answer because all designs are different. To help in the decision-making process, here are some different probe characteristics you'll want to consider.

BANDWIDTH

A probe's bandwidth is typically its banner specification. It essentially describes how high of a frequency the probe is able to pass on to the oscilloscope. Before investing in a new probe, you should consider the speed of your signal's frequencies. As a rule of thumb, your probes should be at least three to five times faster than the fastest signal you want to see.

ATTENUATION RATIO

Probes have different (sometimes switchable) attenuation ratios that change how the signals are fed into your oscilloscope. For example, a 10:1 probe connected to a 1-V signal will pass 100 mV to the scope's input. The scope will read the probe's attenuation ratio and display the correct signal values on the screen so that you're not having to do the calculations in your head.

Having a higher attenuation ratio (100:1, 1000:1) will allow you to look at higher voltages, but it will also make the scope's internal amplifier noise more pronounced. The higher the attenuation ratio, the more scope noise you'll see. For example, a 10:1 probe will show 10X the noise. This might lead you to believe that you should always use a 1:1 probe. Not true! Lower attenuation probes typically have much higher loading on your system, which brings us to the next probe characteristic to consider.

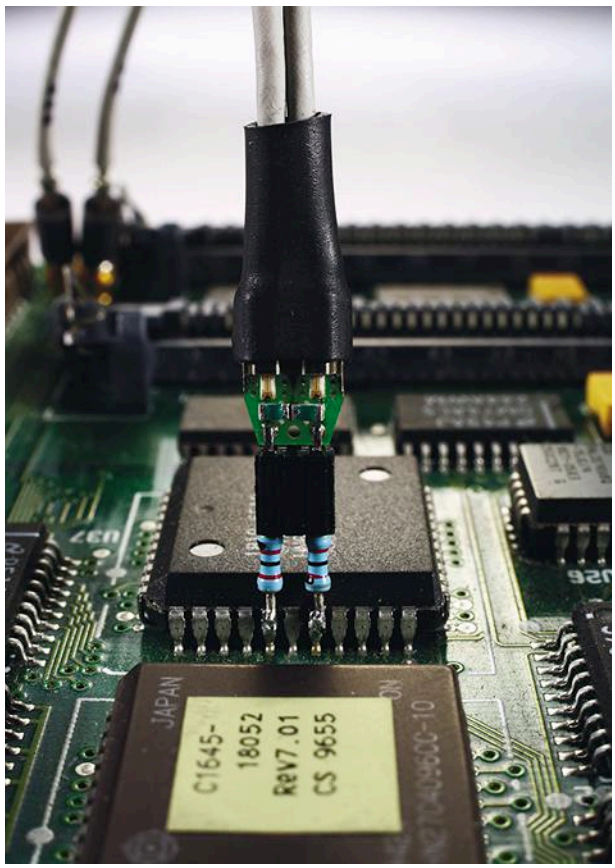

1. A probe at work on a chip.

PROBE LOADING

Connecting anything to your device under test (DUT) will change the inherent electrical characteristics of your system. Despite our best R&D efforts, probes are still subject to the laws of physics. (If we get around this pesky barrier, we'll be sure to let you know!) Therefore, any time you are probing your system, it will affect your device. Unnecessarily loading your system can lead to inaccurate measurements and even change the shape of the waveform on your oscilloscope screen!

Probing loads aren't just resistive, either; they have inherent capacitive and inductive effects. Capacitive effects are generally the most difficult to work with and are extremely responsive to frequency. Too much capacitive loading will slow down your system's bandwidth, and, as a result, effectively slow down your signal's rise and fall times. If rise/fall times slow down too much, you will start seeing bit errors in your system.

As a general rule, the higher the attenuation ratio of the probe, the lower the capacitive loading. So a 1:1 probe may not be the best choice. Oftentimes, 1:1 probes have capacitive loading somewhere around 100 pF, while a 10:1 probe is typically in the range of 10 pF. So even though a 1:1 probe helps reduce your oscilloscope amplifier noise, there's a tradeoff.

Higher-end scopes often use digital signal processing (DSP) to help compensate for probe loading, but it's simply impossible to eliminate probe loading altogether. Therefore, make sure the probes you pick have impedance values that will fit into your design parameters.

Another way to help reduce probe loading is to employ an active probe.

ACTIVE VS. PASSIVE PROBES

Many engineers will only consider an active probe if they need to look at high bandwidths. However, the probe loading from active probes is significantly lower than that of passive probes. By designing a high-frequency amplifier into the probe head, the probe cable and oscilloscope front end are isolated from your DUT. This can bring the probe loading down to 1 pF or better.

What does this mean for your measurements, though? Let's look at a system with a 1.1-ns rise time. Conventional wisdom would tell you that you only need around 300 MHz of bandwidth. Use the formula:

Required Bandwidth = 0.34/rise_time

So we set up our system and measure our edge with an active probe, and see a 1.1-ns rise time. We then disconnect the active probe and connect the free-with-the-scope 500-MHz passive probe. Subsequently, we measure the rise time and see it has changed to 1.5 ns. What happened? The passive probe's capacitance is adding more load to our system, and it's having trouble driving the extra load. The extra load from the probe is distorting the device's signals! The difference in measured edge speeds can be clearly seen in *Fig. 2*.

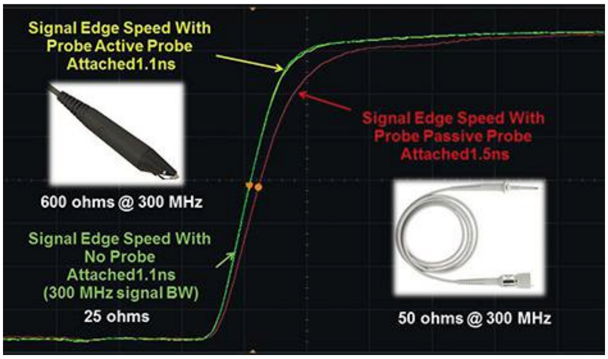

2. Shown is the difference in measured edge speed using an active versus a passive probe.

So, should I ever use the passive probes that come with my oscilloscope? Yes! Passive probes are great for *qualitative* measurements (i.e., checking clock frequencies, browsing for bugs, etc.). But, active probes excel in *quantitative* measurements like output ripple or rise times. While active probes cost more than passive probes, they can make a big difference in measurement accuracy.

If you want to know how much your probes are loading your system, you can double-probe your DUT. First, connect one probe to a known step signal on your DUT (or as close as you can get). Observe and/or save the waveform on the screen. Then, leaving the first probe connected, connect a second probe to the exact same spot and observe the change in your edge speed.

PROBE-COMPENSATION RANGE

Probe compensation makes it possible to adjust the impedance of your probes to fine-tune your acquisition system. The input impedance of an oscilloscope can vary from model to model, and even from channel to channel, due to differing component values.

Passive probes typically come with an adjustable input impedance that allows you to impedance match your probe and scope. This ensures that your probe will perform as expected for its rated bandwidth. The probe's impedance can usually be adjusted on your probe using a small (included) screwdriver. *Figs. 3, 4,* and *5* give examples of what an under-, over-, and properly compensated probe looks like:

3. This screenshot indicates an overcompensated probe; the signal should be a square wave.

4. Here's what a properly compensated probe should look like.

5. This screenshot indicates an undercompensated probe; the signal should be a square wave in this case, too.

All of these screenshots are using an identical probe and signal; only the probe compensation has changed.

Your oscilloscope will have a "typical" input capacitance value, and you want to make sure that your probe's compensation range can match. If your probe can't fully compensate for your scope, you will get subpar measurements, especially when measuring higher frequencies.

CONNECTING TO YOUR DEVICE

Another consideration when evaluating probes is how you will physically connect to your system. For low-frequency measurements, the grabber-hat on your passive probe might be good enough. But, as system bandwidth increases, you'll need to look at other options.

Can you use a ZIF tip, an SMA probe head, a solder-in probe head, or my favorite: the N2848A magnetic QuickTip head? Such questions are beyond the scope of this article, but essential to engineers needing to look at fast signals. Make sure you are designing boards that can be probed, and selecting probes that will provide a good connection to your board. The length of your ground lead is also an important part of connecting to your device, because it can act as an antenna.

SIGNAL-DEPENDENT PROBING

Lots of different styles of probes are available, depending on what signals you're probing. The type of signal can significantly narrow down your probe choices. Are your signals differential? Use a differential probe and eliminate common-mode noise. Measuring current? Use a current probe instead of a sense resistor and a voltage probe. Measuring extremely high voltages? Don't risk your life and use a high-voltage probe with its extra-high attenuation ratio.

OTHER CONSIDERATIONS

It's worth noting that your entire acquisition system's bandwidth is important, not just that of your probe. Your oscilloscope's bandwidth and probe's bandwidth both affect total system bandwidth according to the following equation:

$$\text{System bandwidth} = \frac{1}{\sqrt{\dfrac{1}{\text{Scope bandwidth}^2} + \dfrac{1}{\text{Probe bandwidth}^2}}}$$

Say you had a 200-MHz scope and used a 200-MHz probe. Your system bandwidth would be roughly 141 MHz.

In addition, probes typically come with a probe accessories kit. Before ordering, it's important to know what is or isn't included.

CLOSING THOUGHTS

Many types of probes are available for nearly every signal you'd want to measure.

To avoid confusion, stick to the basics: bandwidth, attenuation ratio, and loading. Fully understanding these specs will help guide you through the many options and help you pick the right probe. ed

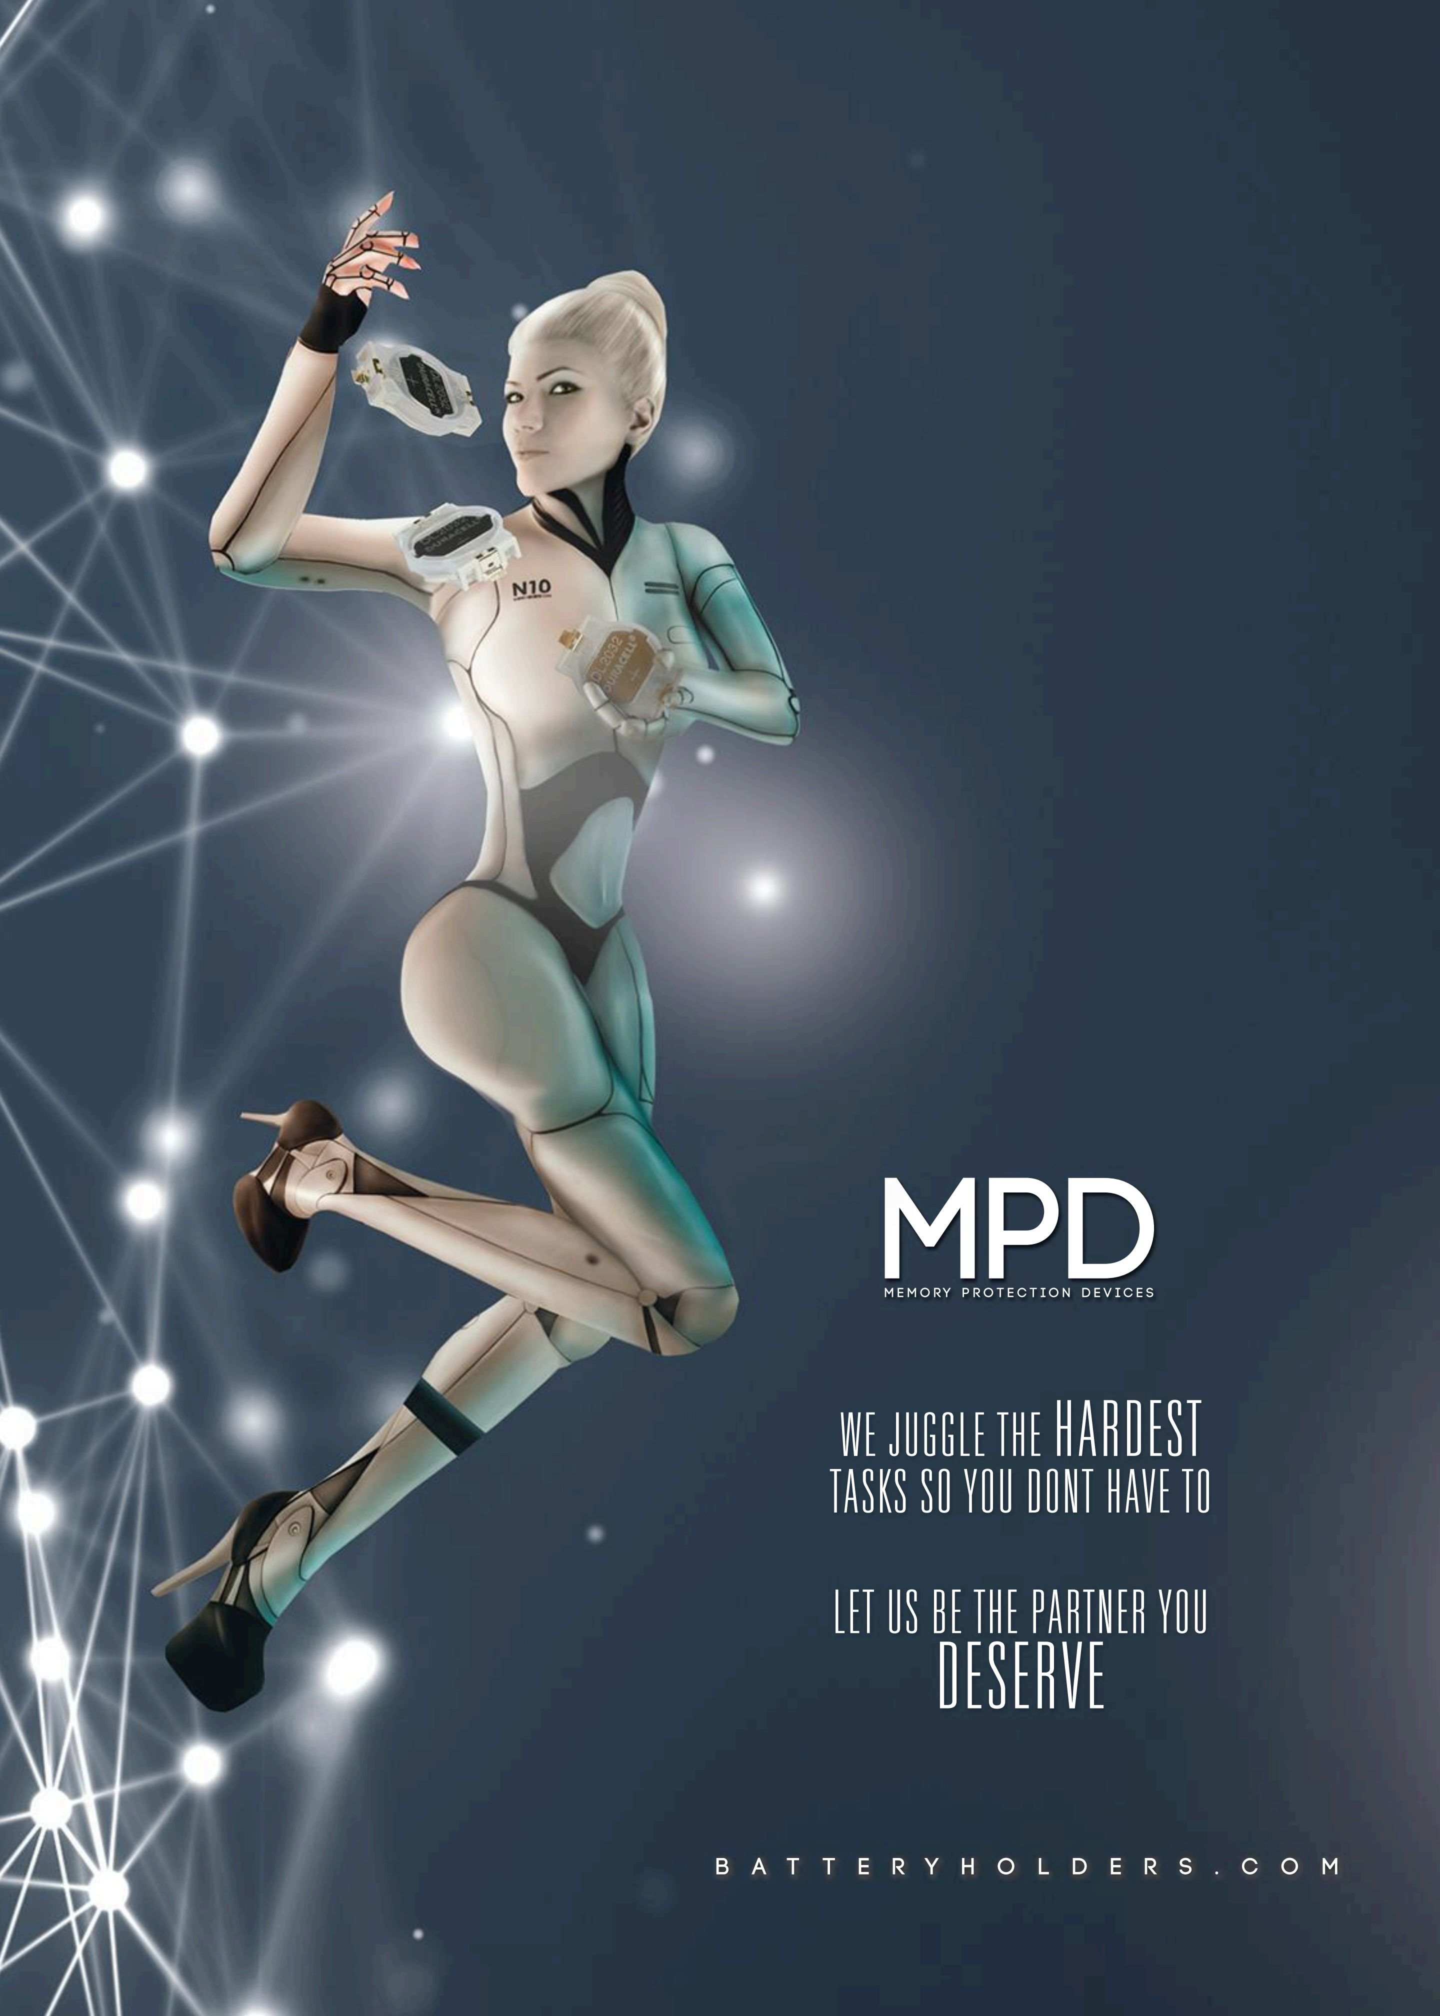
N10
MPD
MEMORY PROTECTION DEVICES
WE JUGGLE THE HARDEST
TASKS SO YOU DONT HAVE TO
LET US BE THE PARTNER YOU
DESERVE
BATTERYHOLDERS.COM

DAVE HUGHES | CEO, HCC Embedded

RAPIDLY ADVANCING FLASH is a Double-Edged Sword for Embedded IoT

Though more-enriched flash brings a number of benefits to an IoT application, it also often creates challenges when trying to meet efficiency, longevity, and reliability requirements.

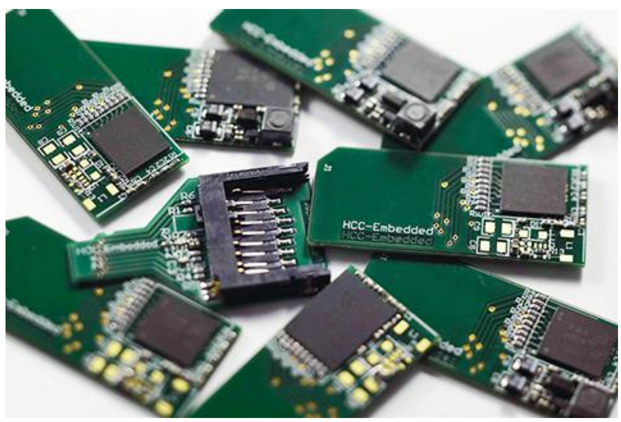

1. Packaged eMMC flash can be a fundamental building block in a reliable or deterministic storage system.

At first glance, using an external flash device with a microcontroller (MCU) appears to be a simple exercise in matching the two devices to achieve functionality. But designing an embedded product is a complex process of finding the best balance between competing solutions to meet the design requirements.

Introducing flash to the design, as most Internet of Things (IoT) devices do, adds another set of variables to arbitrate—and can make it more difficult to meet the application's efficiency, longevity and reliability requirements. This article discusses the most common flash technologies incorporated in embedded systems and the challenges that emerge in order to reliably use them.

TYPES OF FLASH FOR EMBEDDED APPLICATIONS

Flash can be either removable or non-removable. Removable media include USB pen-drives, compact flash, SD/SDHC/XDHX/MMC cards, and other similar packaged technologies. Non-removable media include NAND, NOR, and eMMC. *Figure 1* shows eMMC-packaged NAND flash.

REMOVABLE MEDIA

Practically all removable media is made from a packaged array of one or more NAND devices managed by an FTL embedded in an ASIC, FPGA or similar controller. Its main characteristics—high density, integrated flash translation layer (FTL), fast, and accessible—can be moved to a host system.

Engineers must take seriously the potential risks of using removable media in an embedded system. Most removable media are simply not capable of handling data storage reliably and they are typically not specified for use in embedded systems. The few that are tend to be more expensive. Even SD cards marked as 'Industrial' typically only have an enhanced operating temperature range. Removable media can be a benefit for consumer goods but more problematic for industrial applications like data loggers. Removable media is generally unreliable, can lead to uncertainty of data and suffers from unreliable contact and vibration issues.

Removable media can only be read by a compatible host system. This often forces the choice of file system to the eponymous file-allocation-table (FAT) system with all its known limitations. Most removable media have no chance to achieve fail-safety or establish controlled data commit points that can guarantee the state of a card. This is a real problem for reliable, fail-safe operation and presents a risk to data and to file-system integrity. Most SD cards are optimized for high speed at the expense of correctly defined behavior in the event of unexpected reset. Also, there's no method for differentiating between critical and non-critical operations.

NON-REMOVABLE MEDIA

NAND, NOR, and eMMC are amongst the most commonly used types of embedded flash memory. Compared with NOR, NAND flash presents some complex system-design issues.

NAND FLASH

NAND flash is characterized by low cost per bit, combined with fast erase and write times. The address bus is multiplexed on to the data bus to reduce electronic complexity. It's now *the* core storage technology used in a vast array of storage devices from SSDs to SD cards to eMMCs. When it comes to meeting high-density and performance requirements, NAND is the choice.

NAND flash is arranged as a set of blocks, each of which is divided into physical pages. Every page contains an area for data and a spare area that's used for NAND management data. It's increasingly common for NAND flash to consist of multiple planes that can be addressed independently and programmed or erased in parallel.

Three primary features of NAND flash contribute to the design challenge:

- An error correction code (ECC) is required to correct data that hasn't been stored correctly.
- NAND flash is delivered with a small number of bad blocks, and during the life of the device, additional blocks may become bad.
- If the device will undergo heavy usage, a system of wear leveling is required to extend its lifetime.

Many other smaller features need to be considered on a device-specific basis. NAND flash technology is experiencing rapid and dramatic changes. Originally single-level-cell (SLC) flash was used, but has since progressed to multi-level cell (MLC) and now to triple-level cell (TLC) and 3D NAND. About 15 years ago, a single NAND device could store 32 MB—nowadays it can store 128 GB or more. (Moore's law cannot keep up with NAND flash development!) This progression has led to higher densities and cheaper cost per bit, but has also changed the physical character of the flash.

SLC flash used to be able to survive 100K erase/write cycles with just 1-bit ECC correction, while the latest NAND can require 40-bit ECC correction per 528 bytes just to guarantee 3K erase/write cycles. Other technical challenges must be confronted, too, such as page-pairing issues and read disturb. Consequently, software to manage NAND flash has to progress with the new design features to meet reliability requirements.

If your design requires a NAND with more than 1-bit ECC correction, it will need to use a microcontroller with support for ECC calculations that match the NAND flash.

Serial NAND flash: One useful variant of NAND flash for embedded designs is serial NAND flash. It's not available in the same huge densities as others, but carries two primary advantages: a simpler electronic interface using SPI, and (in most cases) the ECC calculation is integrated, making it possible to integrate with a wider range of MCUs. *Figure 2* shows a traditional NAND flash controller built into an MCU.

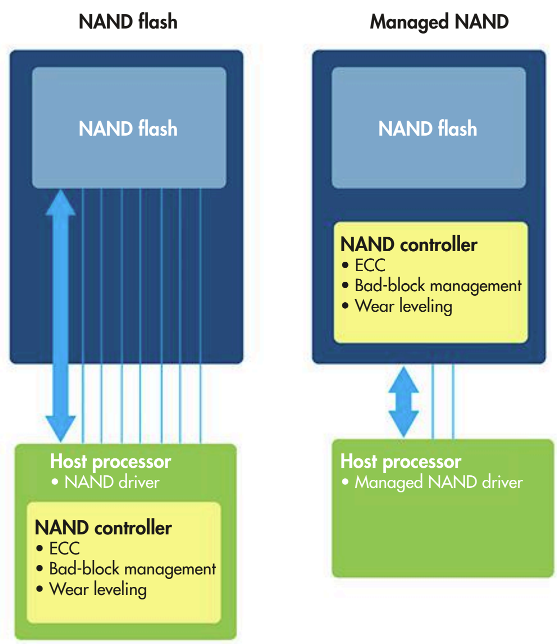

2. When comparing raw NAND versus packaged NAND, several important should be considered, such as ECC requirements and the need for complex management services like wear leveling and bad-block management.

NAND USE CONSIDERATIONS

Wear leveling: Flash cells have a limited life and can only be erased and programmed a certain number of times before becoming unreliable—in effect, they wear out. Wear-leveling algorithms are used to increase the chip's lifetime by moving the data between physical blocks to ensure some cells aren't overused in comparison with others. These algorithms can be fine-tuned to match performance requirements.

ECC: The worst-case rate at which wear occurs is defined by the flash manufacturer. ECCs are used to ensure the data is always consistent if used within the chip specification. The strength (number of bits) of the required ECC is defined by the worst-case bit failure rate.

Bad-block management: Flash memory contains blocks that may be error-prone or unusable when the device is new. During operation, data in good blocks can later be corrupted by charge leakage or disturbance from writes in adjacent parts of the chip. Software is used to provide management of bad blocks and maps unusable areas to ensure that data are not corrupted.

Read disturb: Read disturb errors occur when a read operation on one page causes one or more bits to change in other pages of the same block. Executing many read operations on pages within a block, with no intervening erase of that block, increases the risk that bit errors will develop.

While NAND flash is the leading flash-storage technology used for large quantities of data, there are several important things to consider when using it. The error characteristics of each NAND device type are different, such as the number of bad blocks, bad block markers, and ECC requirements. This means ECC requirement has to be matched to the microcontroller's NAND controller interface. NAND flash also needs complex management services such as wear leveling and bad-block management.

FLASH TRANSLATION LAYER (FTL) FOR NAND

Because of the requirement to re-map bad blocks, it's not possible to treat a NAND device as a linear array of memory in most applications. In applications where flash is erased and rewritten many times, it's necessary to introduce wear leveling to increase the life of the device. Both issues require a logical translation of the physical blocks.

These problems are resolved by using a software-based flash translation layer (FTL). From an application perspective, the FTL presents an array of logical sectors that hides the underlying complexity of the physical device, allowing a file system to use it like any other media driver.

An FTL should also address the problem of wear leveling. There are two common approaches, static and dynamic wear leveling. In dynamic wear leveling, the system chooses the least-used block for the next write operation. This method is relatively weak, since it doesn't account for data blocks whose contents remain unchanged.

With static wear leveling, under-used blocks are swapped in to distribute the wear across the whole device. Static wear leveling needs to be carefully tuned to ensure that block swapping doesn't contribute unnecessarily to device wear. This is typically achieved using thresholds to determine sensible limits for swapping blocks. The table breaks down the static and dynamic wear-leveling methods.

NOR FLASH

NOR flash, which pre-dates NAND, was used in a range of applications. NOR still has a very important share of the embedded-design market because of its relative simplicity and reliability. There's no ECC requirement to guarantee 100K erase/write cycles, no bad blocks at manufacture, and a much simpler architecture. The price to pay for the increased reliability is much longer erase and write times, though. However, like all flash, using it in terms of being reliable is complex and requires a detailed understanding of the technology.

NOR flash comes in two main forms: parallel NOR flash, where the address and data buses are directly connected to the MCU; and serial NOR flash, which only uses an SPI interface (or similar) to communicate. Typically, serial NOR flash has densities <1 MB up to 128 MB. Parallel NOR flash can be much larger.

If NOR flash is to be heavily used, then it still requires a management layer (often an FTL) to provide wear leveling and ensure that the device is used evenly. Having such a layer will also manage any developed bad blocks.

NOR flash can be more expensive than alternatives, and it has relatively slower erase and write times. There's no requirement for ECC, but to use it effectively, logical to physical mapping is required to be able to provide wear leveling,

eMMC

eMMC addresses many of these issues, thanks to several features. For instance, a simple block-mode interface hides all intricacies of the underlying flash from the user while the level of service is given by the manufacturer at the block level. Different write modes allow the user to optimize performance depending on whether the data being written is critical or not. Typically, file system meta-data is critical and file data isn't so critical, but that also depends on the application and the level of fail-safety provided by the file system.

In addition, a bidirectional parameter exchange allows the card to optimize its operations based on the properties of the host system. Finally, eMMC reduces vibration and unplanned removal issues because it's offered as a solid-state chip integrated on the target board.

eMMC devices can provide a fundamental building block in a reliable or deterministic storage system with clear advantages over other similar technologies. But as with any concept of determinism or reliability, the whole system has to be consid-

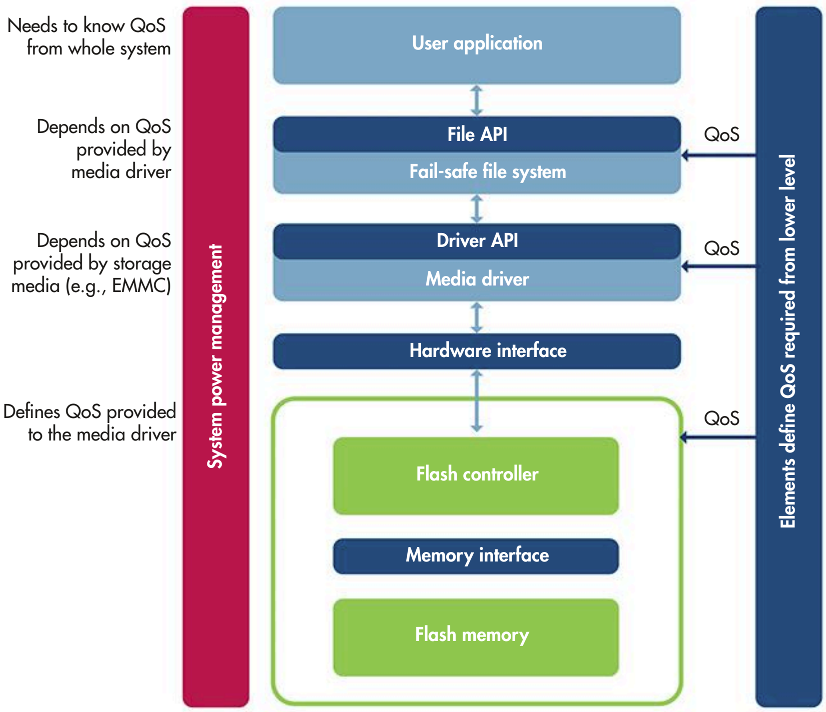

3. Fail-safety system requirements: No file system can claim to provide fail-safety on its own—a system-wide design is required.

ered and the various components validated both individually and in the context of the complete system.

CHALLENGE OF FILE SYSTEMS WHEN USING FLASH

It can be a serious quality issue for an embedded system if the file system, or its contents, become corrupt. To establish tests and verification at the design and implementation stage, the designer must address fundamental challenges, including how to handle file operations and directory structures, how to deal with the integrity of data during power loss or unexpected reset, and how to verify correct operation of the flash.

Traditionally, the handling of file and directory operations is delegated to an embedded file system. An application gets significant benefits from using a file system—the abstraction of the storage media to a set of data files.

However, a file system alone can't guarantee the integrity of data and the file system itself. Whatever method of ensuring fail-safety is used, the system remains dependent on the storage media and must define a required level of service from that media.

A system designed to ensure reliability must include a clear understanding of the critical exceptions—such as unexpected reset or power loss—and how each part of the system will meet the requirements of the components using it upon such an occurrence. Simply using a file system that claims fail-safe operation or journaling has no chance of guaranteeing reliable operation without defining these things.

POWER LOSS, UNEXPECTED RESET, AND SYSTEM FAILURE

In an embedded application, if the data (or meta-data) managed by the file system becomes corrupt, the result could be catastrophic. For example, a file system stored on NAND flash may require reformatting, which will mean loss of all data. For this reason, it's mandatory to use a fail-safe system in which the data is valuable, and furthermore, the application must be reliable.

A truly fail-safe file system will guarantee that all meta-data it manages on flash will always be consistent, and that any write to flash by an application (typically a file system) will be completed atomically. The atomicity of a write means that either the write operation is completed or the disk is left in the same state as it was prior to initiation of the write. This means the file system can guarantee the integrity of the data passed to it by the file system.

To ensure a fully fail-safe embedded system, each layer from the application level to the physical driver must specify what it requires from the adjacent layer. This is the only way to guarantee true fail-safety. For example, a generic FAT file system would require that multiple writes to different areas of the media be completed atomically. This is logically impossible to guarantee on a system where unexpected resets may occur. To build a reliable application, a fail-safe file system can be used together with another fail-safe file system.

Required characteristics of fail-safe system *(Fig. 3)* include:

- After a system reset, the file system will always be in a consistent state.
- Any file that was opened for writing at the time of an unexpected reset will be returned to its pre-open state, unless a flush or close operation on that file was successfully completed. This means that the application developer is entirely in control of when the new state of a file is set, independent of any other activity in the file system.
- Fail-safety of any file system *(Fig. 4)* can only be guaranteed if the low-level driver guarantees a defined quality of service (QoS).

pico® Technology

THE PICOSCOPE® 4444

SEE THE DIFFERENCE

NEW

A NEW STANDARD IN DIFFERENTIAL MEASUREMENT

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

WHAT CAN YOU DO WITH THE PICOSCOPE 4444?

Switch mode power supplies are notoriously difficult to troubleshoot and characterize using grounded-input oscilloscopes, as much of the circuitry is floating or electrically isolated and often at mains voltage levels.

The PicoScope 4444 differential inputs give you the confidence to probe across circuits and components without concern about shorting floating voltages to ground.

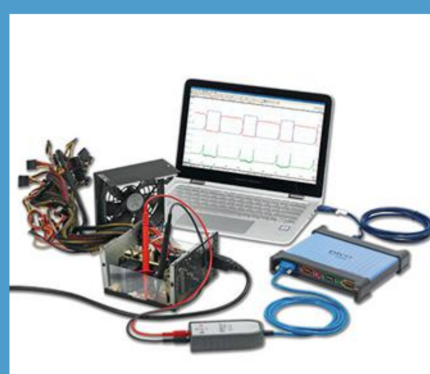

For more information please visit www.picotech.com/US109

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

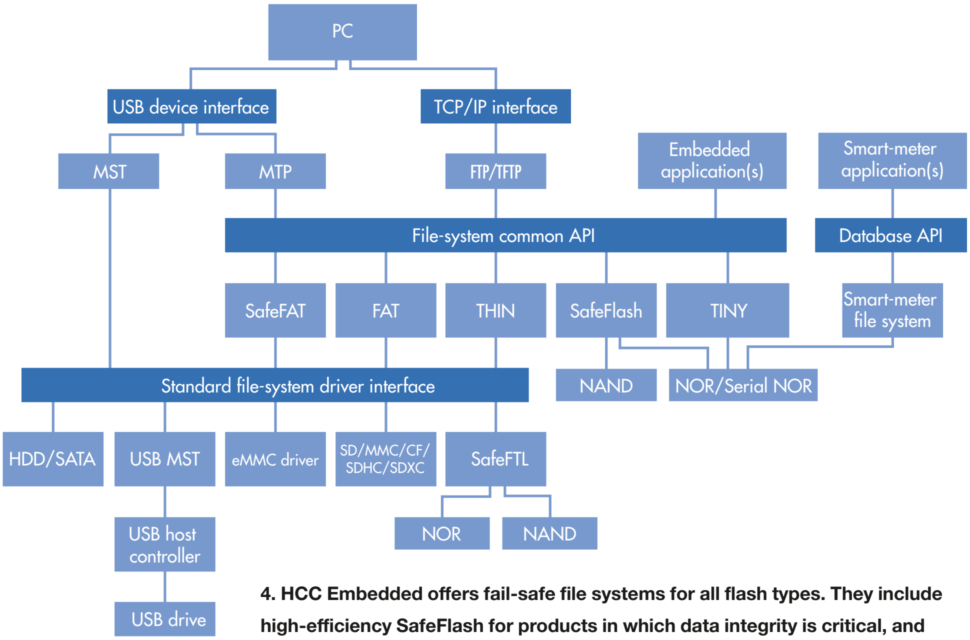

4. HCC Embedded offers fail-safe file systems for all flash types. They include high-efficiency SafeFlash for products in which data integrity is critical, and full-featured TINY for resource-constrained applications, so that designers can choose the file system that best suits their needs.

For the file system, this is defined as:

- Any write operation must complete successfully or return an error. Otherwise, the file system must be restarted.
- All writes to the media must be executed in the sequence in which they're provided to the driver.
- An erase operation must complete successfully or return an error. Otherwise, the file system must be restarted.
- In practice, this means that the hardware has to provide some level of voltage protection. This will ensure that the system can take appropriate action if the voltage provided to the flash media is falling toward the specified programming voltage.

SYSTEM PERFORMANCE AND EFFICIENCY

To achieve high performance using a complex flash device such as multi-plane NAND or an array of NAND chips, it's required that an FTL can read, write, and erase multiple devices in parallel. This could theoretically be achieved simplistically at a driver level. However, it would result in inefficient utilization of flash and, therefore, would be unlikely to be truly fail-safe. It's better to address this problem at the FTL level, and it can be done with commercially available products.

When assessing performance, it's common to measure the response by performing contiguous-read, contiguous-write, random-read, and random-write operations. In general, achieving high performance with contiguous read and random read is relatively straightforward. Contiguous-write and, more critically, random-write operations are more difficult to perform efficiently. This is the case for all flash-based memory devices. But by implementing sophisticated algorithms and system tuning, it's possible to achieve significant improvements in performance.

POWER-MANAGEMENT CONSIDERATIONS

Designers should use careful power management to ensure the flash device isn't erased or written when power levels are outside of the manufacturer's specified limits. If this isn't the case, then problems may occur with the integrity of blocks and data. A well-designed FTL will specify the requirements of the low-level driver, including power requirements, to ensure fail-safe operation. Designers should request these specifications from the software supplier. For example, when the voltage supply falls to the specified minimum, then brownout detection should notify the system so that this condition can be managed.

At first glance, using an external flash device with an MCU appears to be a simple exercise in matching the two devices to achieve functionality. However, developing an application with efficiency, longevity, and reliability requires a comprehensive system-level approach. It's hoped the issues highlighted can go some way toward helping designers address these fundamental technical issues. ed

DAVE HUGHES is the CEO and founder of HCC Embedded, experts in software for securely storing and communicating embedded data. Dave is a "hands-on" embedded specialist, who still actively contributes to the strategy and direction of HCC's core technologies. His extensive experience has made him one of the industry's leading authorities on fail-safe embedded systems, flash memory, and process-driven software methodologies. He is a graduate of the University of Sussex in England.

REFERENCES:
"Flash Reliability," HCC Embedded, https://www.hcc-embedded.com/embedded-systems-misra/flash-reliability
"Embedded NAND Flash and NOR Flash File Systems," HCC Embedded, https://www.hcc-embedded.com/embedded-systems-software-products/file-system/nand-nor-flash-file-systems.

Transient Immunity Testers from AVTECH

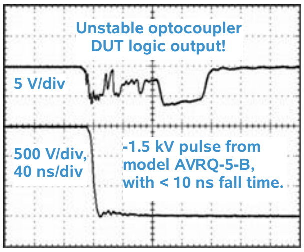

The Avtech AVRQ series of high-voltage, high-speed pulse generators is ideal for testing the common-mode transient immunity (CMTI) of next-generation optocouplers, isolated gate drivers, and other semiconductors.

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

AVTECH

Pricing, manuals, datasheets and test results at
http://www.avtechpulse.com/semiconductor
More information: info@avtechpulse.com

30, 50 and 100 Volt Lab Pulsers

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

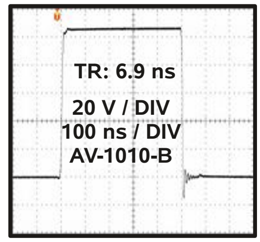

Model AV-1015-B:	**50 Volts, 10 MHz**
Model AV-1010-B:	**100 Volts, 1 MHz, 25 ns to 10 ms, 10 ns rise time**
Model AV-1011B1-B:	**100 Volts, 2 ns rise time**
Model AV-1011B3-B:	**30 Volts, 0.5 ns rise time**

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

AVTECH

Pricing, manuals, datasheets and test results at
http://www.avtechpulse.com/general
More information: info@avtechpulse.com

Nanosecond Laser Diode Drivers With Butterfly Diode Sockets

To Pulse Driver

To TEC Controller

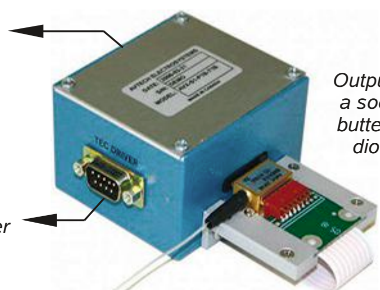

Output module with a socket-mounted butterfly-packaged diode installed.

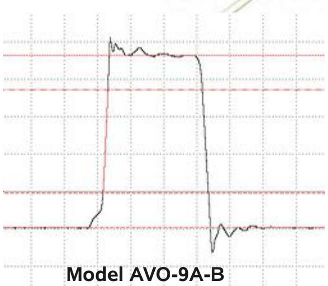

Each of the 19 models in the Avtech AVO-9 series of pulsed laser diode drivers includes a replaceable output module with an ultra-high-speed socket suitable for use with sub-nanosecond rise time pulses. Models with maximum currents of 0.1A to 10A are available with pulse widths from 400 ps to 1 us. GPIB,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

AVTECH

Pricing, manuals, datasheets and test results at
http://www.avtechpulse.com/laser
More information: info@avtechpulse.com

2 to 200 Amp Current Pulsers

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

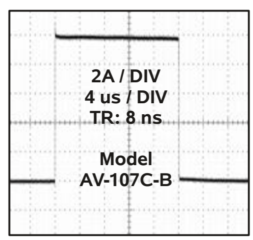

Series	I, V	PW	TR
AV-107	2 - 20 A, 60 V	0.2 - 200 us	10 - 30 ns
AV-106	5 - 100 A, 100 V	0.5 us - 1 ms	50 ns - 1 us
AV-108	12.5 - 200 A, 100V	2 us - 1 ms	5 - 15 us
AV-109	10 - 100 A, 5 V	10 us - 1 s	10 us
AV-156	2 - 30 A, 30 V	1 us -100 ms	0.2 - 50 us

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

AVTECH

Pricing, manuals, datasheets and test results at
http://www.avtechpulse.com/current
More information: info@avtechpulse.com

QUENTIN OCHEM | Lead of Business Development, AdaCore
www.adacore.com

Rust and SPARK: Software Reliability for Everyone

Programming languages often defer reliability and security issues to tools and processes. Two initiatives—SPARK and Rust—state that language is key to reaching these objectives.

Over the past few years, there's been a growing realization that reliability could be achieved more cost-effectively with programming languages designed for that purpose. A good example of that is the Rust language, developed by Mozilla to improve some of its developments. Its mission statement is expressed today in its documentation:

"Rust is a systems programming language focused on three goals: safety, speed, and concurrency. […], making it a useful language for a number of use cases other languages aren't good at: embedding in other languages, programs with specific space and time requirements, and writing low-level code, like device drivers and operating systems. It improves on current languages targeting this space by having a number of compile-time safety checks that produce no runtime overhead, while eliminating all data races. […]"

One can find some troubling similarities with the Ironman requirements from the DoD that led to Ada:

"The language shall provide generality only to the extent necessary to satisfy the needs of embedded computer applications. Such applications require real time control, self-diagnostics, input-output to nonstandard peripheral devices, parallel processing, numeric computation, and file processing. […] The language should aid the design and development of reliable programs. The language shall be designed to avoid error prone features and to maximize automatic detection of programming errors. […] The language design should aid the production of efficient object programs. […]"

As it turns out, while Rust was being developed, another effort was undertaken by a different community, also trying to raise the bar in terms of safety. The SPARK language is an evolution of the Ada language. Let's look at what each has to say.

WHAT BRINGS THEM TOGETHER

Both languages are imperative languages, compiled directly into object code, and both manage memory directly. Each provides abstractions for the usual programming paradigms (procedural and object-oriented). In addition, they both offer advanced concurrency models, providing assurance of properties such as absence of race conditions.

Both languages implement static and dynamic checks, including strict-type safety. The following doesn't compile in SPARK:

```
A : Integer := 5;
B : Float := A;
```

But will with a conversion:

```
B : Float := Float (A);
```

Similarly, in Rust:

```
let a: i32 = 5;
let b: f32 = a;
```

won't work, but can be fixed with a conversion:

```
let b: f32 = a as f32;
```

Arrays are treated as first-class citizens in both languages. In particular, equality between arrays is an equality test of their values and arrays are provided with high-level initialization (aggregates) and syntax to refer to subsets of their elements (slices).

DESIGN NOTES

Lower Power Op Amp: Utility Sine Wave

Design Note 564

Catherine Chang, Philip Karantzalis and Aaron Schultz

Introduction

Our op amp family has expanded with industry-leading speed versus supply current. The LTC®6258/LTC6259/LTC6260 family (single, dual, quad) provides 1.3MHz at a super low 20µA supply current, with 400µV maximum offset voltage and rail-to-rail input and output. In combination with a 1.8V to 5.25V supply, this op amp enables applications requiring uncompromised performance with low power and low voltage at reasonable cost.

Utility Sine Wave

One does not expect to generate a sine wave with –100dBc distortion using a 5V low power op amp. All the same, a bandpass filter using the LTC6258 can combine with an easy-to-use low power oscillator to create a sine wave at low cost, low voltage and extremely low dissipation.

Active Filter

The bandpass filter of Figure 1 is AC coupled to an input. As a result, the LTC6258 input does not place a burden on the previous stage to develop a particular absolute common mode voltage. A simple resistor divider with RA1 and RA2 provides biasing for the LTC6258 bandpass filter. Pegging the op amp inputs to a fixed voltage helps to reduce distortion that might arise with moving common mode.

This filter is centered at 10kHz. The exact resistance and capacitance values can be tweaked upward or downward, depending on whether lowest resistor noise or lowest total supply current is most important. This implementation was optimized for low dissipation by reducing current in the feedback loop. The capacitors C2 and C3 were initially 4.7nF or higher, with lower resistor values. In the end, 1nF with higher resistors optimized for lower dissipation.

Besides power dissipation, a secondary but no less important aspect of feedback impedance is loading of the op amp rail-to-rail output stage. Heavier loading, such as between 1K and 10K impedance, significantly lowers open loop gain, which in turn affects the accuracy of the bandpass filter. The data sheet suggests A_{VOL} reduces by a factor of 5 from

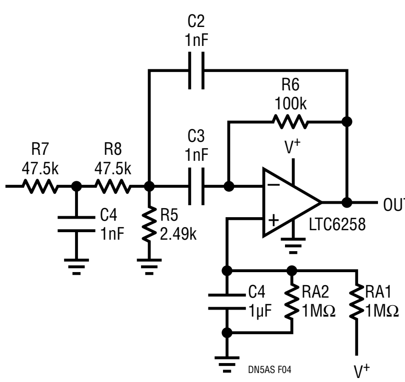

Figure 1. 10kHz Bandpass Filter

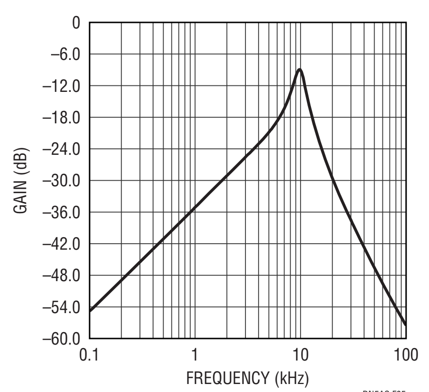

Figure 2. Bandpass Filter Gain/Phase vs Frequency

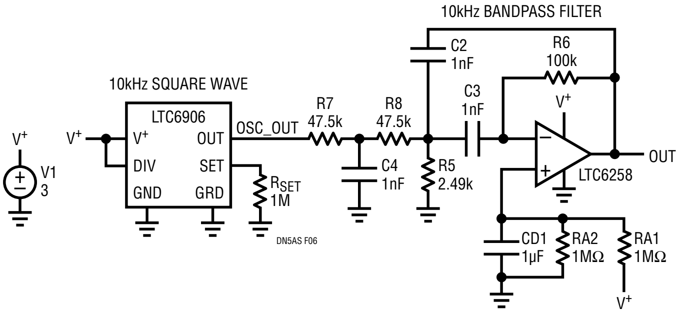

Figure 3. 10kHz Oscillator Circuit Using LTC6906 TimerBlox® Input

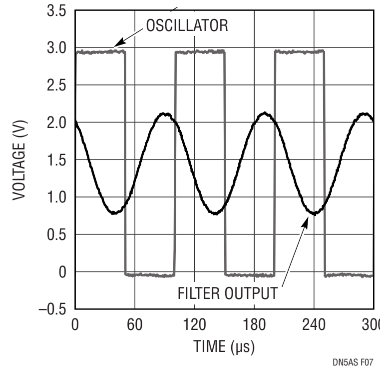

Figure 4. Voltage Waveforms Oscillator and Filter Output

100kΩ to 10kΩ. Lower C2 and C3 might be feasible, but then R6 becomes even larger, introducing more noise at the output.

The target Q of this bandpass filter is moderate – approximately 3. A moderate Q, rather than a high Q, allows for use of 5% capacitors. Higher Q will demand more accurate capacitors, and very likely higher open loop gain at 10kHz than is available with the feedback impedance load. Naturally, moderate Q results in less attenuation of harmonics than a higher Q.

Adding The Oscillator

A low power sine wave generator can be derived by driving a square wave into the bandpass filter. A complete schematic is shown in Figure 3. The LTC6906 micropower resistor-set oscillator easily configures as a 10kHz square wave, and can drive the relatively benign loading seen in the bandpass filter input resistors. Supply current of the LTC6906 at 10kHz is 32.4µA.

Figure 4 shows the LTC6906 output and bandpass filter output. HD2 of the sine wave is –46.1dBc, and HD3 –32.6dBc. The output was $1.34V_{P-P}$ to $1.44V_{P-P}$ with exact level varying slightly due to finite op amp open loop gain at 10kHz. Total current consumption is below 55µA on a 3V rail.

Other Enhancements

Figure 5 shows optional enhancements. A low power reference takes advantage of the ability of the LTC6906 and LTC6258 to operate on a very low supply. The reference provides 2.5V from a battery input. The fixed 2.5V supply stabilizes the output voltage swing in the presence of varying input voltage. In addition, even lower filter capacitor values with higher resistances reduce LTC6258 loading further, lowering dissipation and improving filter accuracy.

Conclusion

The LTC6258/LTC6259/LTC6260 family (single, dual, quad) provides 1.3MHz gain bandwidth at a low 20µA supply current, with 400µV maximum offset voltage and rail-to-rail input and output. In combination with 1.8V to 5.25V supply, this op amp enables applications requiring excellent performance with low power and low voltage at low cost.

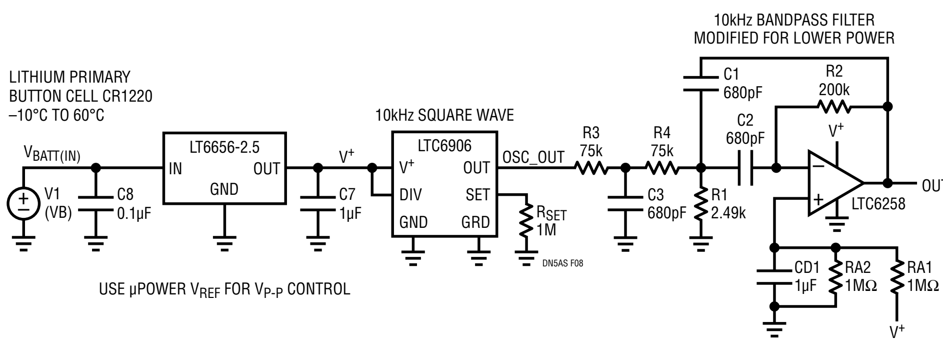

Figure 5. Oscillator and Filter with a Regulated Supply

Data Sheet Download

www.linear.com/LTC6258

For applications help,
call (408) 432-1900, Ext. 3409

DN564 LT/TP 0717 71K • PRINTED IN THE USA

Each language enforces strict safety-related rules while allowing the developer to relax them by using explicit "unsafe" operations. In both cases, this helps to make these operations visible (and thus prime suspects in case of problems) while still allowing the right level of flexibility.

RUST AND SPARK MEMORY MODELS

Rust focuses on memory integrity, providing a model that allows the use of dynamic memory without risk of memory corruption. SPARK focuses on analyzability and provability of program properties and doesn't permit direct use of dynamic memory or even pointers. Instead, it provides support for containers and objects of dynamic size and leaves other pointer requirements to the unsafe portions of the code. Both languages forbid memory aliasing in their safe subsets, and neither provides a garbage-collection capability.

Here's a Rust example on a container:

```
let mut v : Vec<i32> = Vec::new();
v.push(40);
v[0] = v[0] + 1;
```

The above will lead to the value 41 in v[0]. Access to an out of bounds element will lead to a clear error, a panic.

The core idea of Rust's memory model is that a piece of data is always owned by a unique variable, or binding. This protects against manipulations of an object from two references which can create many subtle bugs. This also eliminates the requirement of manual deallocation. When the unique binding goes out of scope, the data is freed automatically.

To fully appreciate the consequences of this, let's look at a short example:

```
let mut v : Vec<i32> = Vec::new();
v.push(40);

let mut v2 : &Vec<i32> = &mut v;
v[0] = v[0] + 1;
```

v2 is a reference (or pointer) to v. Because of the assignment, v2 borrows the value of v and owns it for the rest of the scope. As a result, it's no longer possible to directly modify v. This is checked statically, and the compiler will refuse to compile the above code. Borrowing can be done on a smaller scope as well:

```
let mut v : Vec<i32> = Vec::new();
v.push(40);
{
    let mut v2 : &Vec<i32> = &mut v;
    v2[0] = v2[0] + 1;
}

v[0] = v[0] + 1;
```

When the above code works correctly, v2 releases the ownership of v at the end of the scope. This is particularly useful in parameter passing.

With SPARK, large chunks of memory are managed through containers. The key idea is that containers not only provide safe access to a pool of objects from a memory perspective, but also allow ways of reasoning about them and statically verifying certain run-time properties. For example:

```
    V : Vector(10);
begin
    Append (V, 10);
    Append (V, 20);
    Replace_Element (V, 3, 300);
```

The example creates a vector, then adds two elements to it. Subsequently, we attempt to replace the element at index 3, which doesn't exist. This code would raise an exception at run-time. The SPARK prover is able to verify that the condition "the index of the element must exist before a replacement" cannot be verified and will actually statically issue an error.

SPARK unleashes its power when working on larger pieces of code, making sure that certain categories of errors don't happen, such as buffer overflow or division by zero.

A WORD ON FUNCTIONAL SAFETY

The goal of Rust's memory model is to provide a mechanism to reduce or avoid memory corruption. This makes a lot of sense for languages used extensively on desktop applications and servers, where memory corruption is the number one offender in terms of software vulnerability.

SPARK originated from the Ada language and has the roots of its history in high-integrity embedded software. Although memory corruption also plays a role there, the objective of safety is to verify that a given implementation is correct with regard to a given specification.

Looking at the above example, we can consider that "the index exists before calling Replace_Element" is part of the specification of Replace_Element and should be respected by any caller. SPARK allows one to specify custom requirements on user subprograms in the forms of preconditions and postconditions. For example, one could write:

```
procedure Fold_Lasts (V : in out Vectors.Vector)
with Pre  => Length (V) >= 2,
     Post => Length (V) = Length (V)'Old - 1;
```

This states that the procedure Fold_Lasts expects a vector of at least two elements and will update the vector so that the length of the new value is one less than the length of the old value. The implementation of the procedure will therefore verify that, given the precondition, the postcondition holds. For example:

```
procedure Fold_Lasts (V : in out Vectors.Vector) is
begin
     Replace_Element
     (V,
      Length (V) - 1,
      Element (V, Length (V) - 1) + Element (V, Length (V)));
     Delete_Last(V);
end Fold_Lasts;
```

This subprogram accesses the last two elements of the vector (which is possible and correct, as we established there are at least two elements) and removes the last one (hence ensuring the length of the returned object is equal to the length of the old object minus one).

This can then be used in an actual piece of code:

```
     V : Vector(10);
begin
     Append (V, 10);
     Append (V, 20);
     Fold_Lasts(V);
     Fold_Lasts(V);
```

In the above example, the second Fold_Last will issue a proof error, as it should: We're adding two elements to the vector and folding one leaves only one, which is inconsistent and can be detected statically.

These functional properties, or contracts, can be used by the prover but they may also generate assertions in the final executables to test their validity at run-time. This is particularly useful in cases of contracts too complex to be proven or, while developing contracts, to allow for debugging of their behavior at run-time.

OBJECT ORIENTATION

Both Rust and SPARK provide support for object orientation. The SPARK object-orientation model is relatively close to the usual paradigms implemented with C++ or Java. The interesting aspect of the SPARK OO model is that it supports verification of substitutability between a class and its subclasses.

In the following example, a root class handles an RGB image and is derived into a specialization that handles grayscale images:

```
type Pixel is array (Integer range <>) of Float;
type Image is tagged null record;

procedure Set (This : in out Image; X, Y : Integer; P : Pixel)
     with Pre'Class => P'Length = 3;

type Image_Gray is new Image with null record;

override
procedure Set
 (This : in out Image_Gray; X, Y : Integer; P : Pixel)
with Pre'Class => P'Length = 1;
```

In the Set method, the "Image" parameter requires a pixel to be an array of three float values, while "Image_Gray" requires a pixel to be an array of one. On a dispatching call from "Image," the user doesn't know which type is going to be the actual type and doesn't know if the length of a pixel is 1 or 3.

SPARK will verify consistency of behaviors declared with methods (expressed in the form of contracts), flag any error, and detect the inconsistency of the Set method. Generally speaking, all inputs accepted by the parent class must be accepted by the child.

Rust doesn't provide any means of specifying behavior or verifying class consistency. However, it offers an OO model that's much more modern through the notion of "traits." In Rust, declaring a type Image like the above can't be done directly in a type, but rather must be done in a trait:

```
trait Image_Trait {
     fn set(&self, x: i32, y:i32, p : Vec<f32>);
}
```

There's no notion of deriving a type in Rust, so creating an architecture comparable to the above requires two different types. Such code would look like:

www.rocelec.com

The Semiconductor Lifecycle Solution™

Long term support – Solved.

- 100% Authorized by over 70 semiconductor manufacturers
- Over 15 billion units in stock
- Licensed manufacturing solutions – 70,000+ device types

100% Authorized, licensed, certified and guaranteed

Rochester Electronics®

```
struct Image { }
struct Image_Gray { }

impl Image_Trait for Image {
     fn set(&self, x: i32, y:i32, p : Vec<f32>) { }
}


impl Image_Trait for Image_Gray {
     fn set(&self, x: i32, y:i32, p : Vec<f32>) { }
}
```

As you can see from the example, a structure isn't inherently derived from a trait, as it would be in the case of a Java/Ada interface derivation. A given trait is implemented for a given type in a separate block. This provides a nice solution to the known problem of the tyranny of the dominant decomposition that forces a hierarchy of types to be decomposed on one axis only. This also allows third-party trait implementation.

WHAT SETS THEM APART

It would take a long time to go over each feature one by one and highlight each axis that can be used to compare these two languages. We talked a lot about what brings them together and how they look at different angles of similar problems and come up with different solutions. There are, however, interesting elements that also make them very different.

The most visible one relates to specification. Rust doesn't provide a clear separation between a specification file and an implementation file. As with C or Java, this can be done through language features (you can decide to have traits in a "specification-only" file), but this is left to the discretion of the developer. SPARK enforces a clear distinction between these two notions and provides language-level verification of consistency between specification and implementation.

Another aspect is the tradeoff between writability and readability. In Rust, when typing can be statically inferred, it becomes optional. So I can write:

```
let x  = 0;
```

The compiler will know that x is an i32. This becomes particularly useful when the type is actually a complex generic instance, reducing the amount of code to be written. On the other hand, the downside is that it may be more difficult to figure out what the type actually is when reading the code.

SPARK takes the opposite stance: it limits to the absolute minimum what's to be inferred by the compiler, forcing the developer to specify every type precisely at the cost of sometimes pedantic constraints.

Both languages provide strong static typing features. However, in Rust, native types are just machine types, such as i32, f64, etc. In Ada, a type is a semantic entity associated with constraints that may or may not be defined, such as the range of values or memory size. Moreover, a type is associated with semantics that make it clear that even if it's implementation is the same as another type, it should be considered a different entity. For example:

```
type Percent is new Float range 0.0 .. 100.0;
type Miles is new Float range -10_000.0 .. 10_000.0;

V1 : Percent := 95.2;
V2 : Miles := V1;
```

The compiler knows that V1's assignment into V2 doesn't make sense because they're different types, even though they may both end up being implemented as 32-bit-machine floating point. Moreover, extra constraints—here, a range of values—can be verified either by static methods or dynamically via testing.

Rust provides an extremely powerful and structured macro language that allows one to effectively create expansions of pieces of code based on patterns. There's nothing similar in SPARK. Rust also provides a so-called "matching" control structure, which is much more powerful than the SPARK equivalent, closer to C/Java switch statements.

Differences exist in the error recovery model: SPARK has limited support for exceptions (complete support if you use the unsafe code mode), while Rust uses a much more constraining notion of "panic" when something is unexpected. Rust supports lambdas and closures, which don't really have an equivalent in SPARK.

And the list goes on. Determining which set of features is the best for a given project is a tradeoff that depends on many different things external to the language but are specific to that project, team, market, etc. But it's useful to have some understanding of these differences in order to make an educated choice.

A WORD ON ECOSYSTEMS

The choice of using a given language is also tied to its ecosystem. In this regard, Rust and SPARK are both very similar and very different. They're similar in the sense that they each rely on an open-source community, though at different levels.

The Rust compiler is LLVM-based, supported by the Mozilla foundation and a large community of hobbyists. The SPARK compiler is GCC-based, and while an open-source community does exist around it, it's more restricted. The compiler element of the tool is mostly maintained by a software vendor called AdaCore (of which I happen to be part of) and the formal proof engines underneath are maintained and developed in collaboration with various universities.

Both technologies are available for a large set of targets, native and embedded. They're two open-source strategies, but shaped differently.

It would be wrong, however, to consider this picture to be static. With Rust being deployed on a growing number of projects, there's no reason why commercial entities would not start to provide official support on it. There's no reason why professional tools wouldn't appear either.

On the SPARK side, these tools exist from various vendors, inherited from the Ada history. AdaCore has a history of scarce interactions with the communities. That facade is also breaking down, with lots of development now done on GitHub and several other community-friendly initiatives being started.

As to industrial usage, the first users of Rust are coming from the IT world, while SPARK users come from high-integrity embedded A&D applications. Both technologies target embedded application such as medical devices, industrial automation, automotive, and IoT.

The question of availability of developers must also be considered. There won't be many people who include Rust or SPARK in their résumé when getting out of university this year. But the language shouldn't be an obstacle for a decent embedded software engineer and the internet is filled with tutorials and training resources for both.

CONCLUSION

Rust provides a unique approach to software reliability and safety and is being adopted by a growing number of projects. The language addresses reliability concerns in a much better way than C or C++, without the overhead of VM-based languages such as C# or Java. It's a young language, and some features or tools may still be missing, but there's no reason why that shouldn't be overcome in the future.

SPARK provides a unique approach to software reliability and safety and is being adopted by a growing number of projects. The language leverages its Ada foundation to target a large set of use cases in the embedded domain. SPARK is also a relatively young language, at least in its current form, and the language is rapidly evolving.

The two languages can coexist. SPARK remains very well-suited for safety-critical embedded applications, while Rust looks like a good fit for the IT domain. Generic embedded applications may lean on one side or the other, depending on various factors. Both languages bring interesting ideas to the table—and both suffer from shortcomings. Perhaps there's room for cross-fertilization.

Above all, it's very nice to see new languages considering safety and reliability at their core. The market can only be enriched by a larger offer of technologies, which will undoubtedly push better practices in industrial settings. As a matter of fact, I'm not sure if I care so much whether you're using SPARK, Rust, or a super-constrained MISRA-C—as long as my car, my money, and my blood pressure are handled with safe and robust pieces of software! ed

---

QUENTIN OCHEM has a software engineering background, with a special focus on development and verification tools for safety- and mission-critical systems. He has over 10 years of experience in the Ada programming language. He has conducted customer training on topics including the Ada language, AdaCore tools, and the DO-178B and DO-178C software certification standards. He is currently leading the technical account management and business development activities at AdaCore, in connection with projects from the avionics, railroad, space, and defense industries, both in Europe and North America.

Flip-Flop Design Provides Frame Sync for Received Satellite Telemetry

By NITHIN KUMAR, CHIRAG R, and SAMBASIVA RAO | PES Institute of Technology, Bangalore, India

SATELLITE-TELEMETRY DATA is digitized, multiplexed, and formatted into frames at a 1-kb/s data rate typically, and transmitted on a downlink frequency using data formats such as shown in the *table*.[1-3]

TELEMETRY SPECIFICATIONS	
Frame length	1024 bits
Word length	8 bits/word
Number of words in the frame	3 words (24 bits) for frame sync 125 words for data
Bit rate	1024 bits/s
Frame sync code (used as example)	101011001100101000011111 ACCA1F (Hex, 24 bits)

The telemetry data is transmitted using a multilevel-modulation technique, usually via phase-shift keying/phase modulation (PSK/PM) to overcome the effects of phase noise. After the signal is received, demodulated, and decoded at the ground station, frame synchronization is needed to extract the telemetry parameters and obtain the data. This can be done with a simple circuit built almost entirely with flip-flops. *Fig. 1* shows the overall block schematic as represented in MathWorks' Simulink.

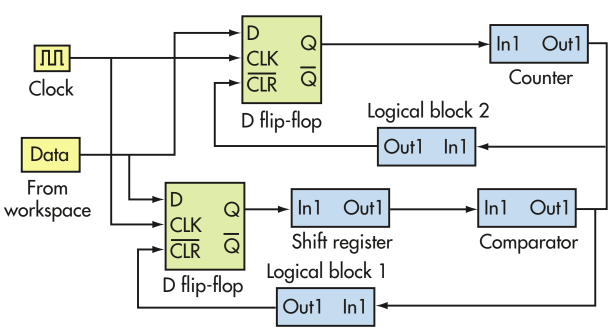

1. The overall block schematic using a Simulink-based model (from MathWorks) of the decommutation system shows the relative simplicity of the approach.

Each frame consists of a frame-synchronization sequence of 24 bits (assumed to be hex "ACCA1F" in the example here) followed by 125 words of 8-bit length representing the data. The main circuit contains two D flip-flops, D1 and D2, which act as switches to send and stop the data flow into following circuit. The output of the D1 is fed to a 24-bit serial-in, parallel-out (SIPO) shift register that's synchronized with the generated clock.

At each clock transition, the data clocks into the flip-flop's bank (shift register). After 24 clock pulses, the first 24 received bits are taken inside the shift register and the parallel outputs are fed to a comparator connected to the output of the shift register. The comparator is a single-bit implementation where each output of the shift register is connected to 24 comparators and compared with the 24 bits of the frame-sync code ACCA1F.

The data keeps entering the shift register until the frame sync is matched. Once matched, the comparator output goes high, which changes the state of D1 to stop the flow of data into the shift registers. At the same time, D2 (which was inactive during this period) is enabled by the comparator output. The data that comes after the 24 bits is routed through D2 to a 10-bit counter that's configured as a 1000-bit counter with Simulink model *(Fig. 2)*, and which counts the clock pulses from the instant data starts flowing out of D2.

Once the 1000-count is reached, the counter resets the operation to initial condition by disabling D2 and enabling D1, which sends the subsequent data bits into the shift registers to detect frame sync of the next frame. The cycle repeats as the data keeps coming. The output of D2 is stripped out for every eight bits and displayed.

The shift-register block contains the cascade of the 24 D-type flip-flops. The data is input to the first flip-flop and the output of the flip-flop is the input to the second flip-flop, and so on. This cascade connection implements the SIPO shift register. The comparator is a combinational circuit of logic gates that checks for the condition of A = B, where A and B are single-bit inputs. The 24 constant bits of the frame sync are the reference inputs (A) to the 24 comparators, while the 24 shift-register outputs are second (B) inputs.

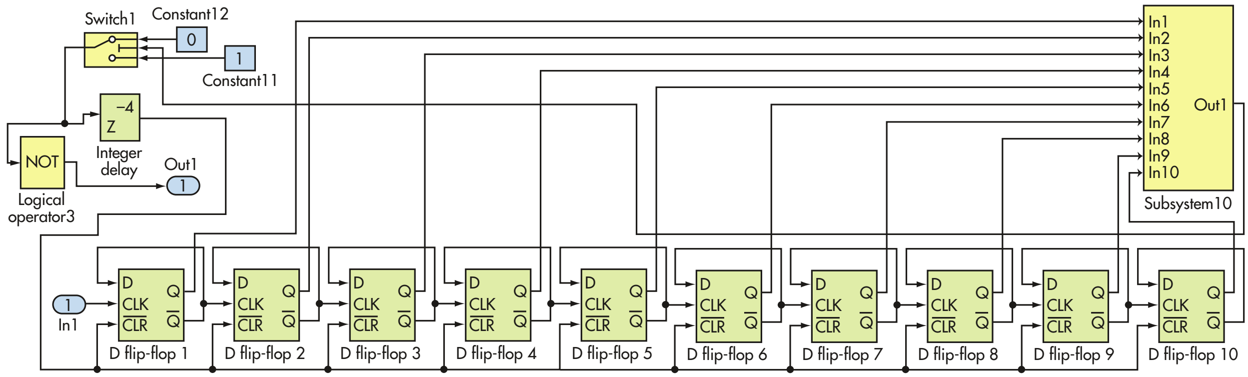

2. The Simulink model of the frequency counter shows the 10 D-type flip-flops that implement the counting function.

The system is implemented on the Xilinx XUPV5-LX110T FPGA Board[4] using the System Generator Software tool, also from Xilinx, with the system-generator circuit shown in *Fig. 3*. The data format contains bits corresponding to ACCA1F in the first 24 bits. Thus, the output of D2 for the first 24 pulses is zero, and then the 1000 bits (125 words) of data are obtained from the 25th clock pulse onwards, until the comparator output goes high.

Figure 4 shows the received data (upper trace) and the data at the output of D2 (lower trace). The output is zero until it matches the frame sync, then the next 1000 bits are matched with the input data stream. These data bits are processed to separate the words, and each word data is decoded and displayed in engineering units for the user.

ACKNOWLEDGEMENT: The authors would like to thank Dr. K.N.B. Murthy, Vice Chancellor, Dr. T.S. Chander, Chairperson, Electronics and Communication Dept., and Dr. V.K. Agrawal, Director, CORI, PES University for the encouragement which they provided.

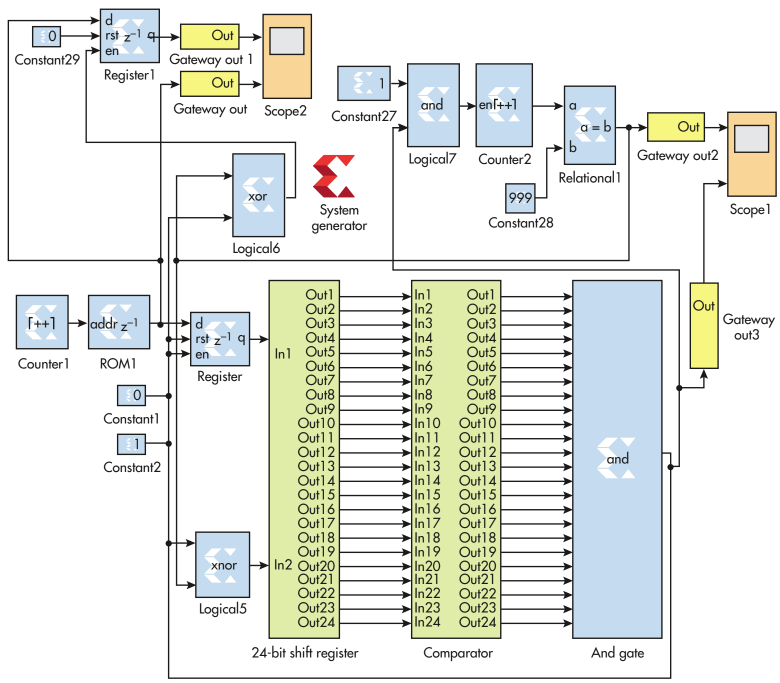

3. The System Generator Model uses the Xilinx XUPV5-LX110T FPGA Board and associated System Generator Software tool.

4. A comparison of input and output data streams shows matching of the frame sync, which then enables data decoding.

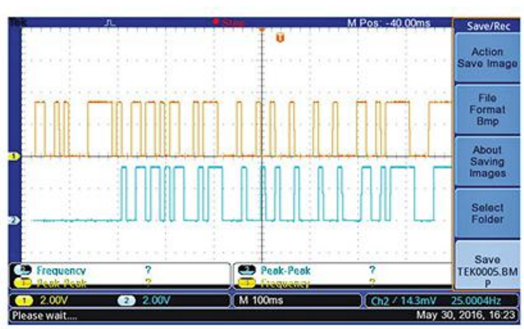

REFERENCES

1. Consultative Committee for Space Data Systems (CCSDS) Green Book, CCSDS 100.0-g-1, December 1987.
2. D.V. Raju, R.K. Rajangam, P.S. Rajyalakshmi, "The Telemetry System," *Proceedings of Indian Academy of Science, Volume C1, No.2*, September 1978, pages 184-198.
3. S.N. Prasad and S. Pal, "Telemetry, Tracking and Command Systems of Spacecraft,"*IETE Technical Review,* Volume 20, Issue 6, 2003.
4. Xilinx User Guide, ML505/ML506/ML507 Evaluation Platform, UG347 (v3.1.2), May 15, 2011.

NITHIN KUMAR was a student with an under-graduation degree of Bachelors in Telecommunication; **CHIRAG R** was a student with an undergraduate degree of Bachelors in Telecommunication; and **DR. V SAMBASIVA RAO** is a professor at Department of Electronics and Communication with Satellite Technologies as his major area of research. All work at PES Institute of Technology.

(Continued from page 14)

COMMUNICATION ADVANCES

Even the wirings running around the cars are changing what they are conducting. CAN and LIN are still mainstays for control but CAN-FD, FlexRay, MOST and Ethernet are in the mix as well. Some of these are part of the IVI system but ADAS sensor connections are being added to the mix as well.

RJ0-45 jacks and CAT5 and CAT6 are unlikely to be found under the hood but the OPEN Alliance SIG's (One-Pair EtherNet) 100BASE-T1 and 1000BASE-T1 standards are trading power and simpler wiring for distance *(see "Automotive Ethernet Was the Hidden Trend at CES 2016")*. The IEEE 802.3 standard incorporated 100BASE-T1 in IEEE 802.3bw-2015 Clause 96. This allows conventional micros with Ethernet connections to handle this wiring.

Even Time-Sensitive Networking (TSN) is being used in automotive applications, at least in the IVI side of things for audio and multimedia synchronization *(see "Time-Sensitive Networking for Real-Time Applications")*. TSN actually started out as Audio Video Bridging (AVB) support for automotive applications.

Not all automotive communication will be wired. Wireless communication for Wi-Fi access point support and cellular links to the outside world are already here, but much more will be coming: Vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) are part of the V2x mix designed to improve ADAS as well as provide self-driving cars with more information about their environment *(see "Car Technology Drives CES 2017")*.

The V2X discussion is ongoing, but test deployments are in progress at select cities around the world. This testing is often part of a smart-city approach that offers information ranging from cross traffic detection and smart stoplights to finding parking spaces or providing vehicle tracking.

USER-INTERFACE ADVANCES

Advanced IVI and ADAS work better with improved displays. Many technologies that have proven popular elsewhere, like OLEDs, are creeping into concept cars, providing features from improved viewing to curved displays.

Heads-up displays (HUDs) are becoming more common and more complex, delivering more information in a larger space. Even portable and aftermarket HUDs are now available.

TI applied its DLP technology to HUD use *(Fig. 8)*. Its DLP3000-Q1 digital micromirror device (DMD) and chipset are automotive-qualified. The platform, which provides a 12-degree field of view, is light source agnostic, and doesn't require a polarized source. The latter means the system will not conflict with glasses that have polarized lenses.

We've been talking about all of the high tech going into a car. Occasionally, a low-tech approach combined with some high-tech support can provide users with an interesting solution, such as the Mpow Universal HUD *(Fig. 9)*. It uses your smartphone to display information, using the same approach as other HUD implementations. An app can display speed based on sensors in the smartphone. This isn't the same as a HUD built into the car, but it's hard to beat the $30 price.

This overview should convince you that cars are going to be one of the most sophisticated devices around. Slapping a steering wheel on a frame that has four wheels and an engine is by no means easy, but that isn't the extent of what will be rolling off the assembly line. Not to mention, some of those cars will likely be missing that steering wheel in the future. ed

8. Texas Instruments' DLP3000-Q1 digital micromirror device (DMD) and chipset are automotive-qualified and provide a 12-degree field of view.

9. The Mpow Universal HUD uses a smartphone to provide a drive with a heads-up display.

1,200 V SiC Schottkys Lower Switching Losses

THE GEN2 SERIES of 1,200 V SiC Schottky diodes are the first in a series of products from Littelfuse to be based on a technology platform created through partnership with Monolith Semiconductor. Additional silicon carbide products based on the technology platform, including 1,200 V SiC MOSFETs, are in the pipeline.

Available in ratings of 1,200 V at currents from 5 A to 10 A, the SiC diodes help reduce switching losses and increase efficiency and robustness in power electronics. They can accommodate large surge currents without thermal runaway, and operate at higher junction temperatures than standard silicon diodes. Stored capacitive charge and forward voltage drop are improved, and a maximum junction temperature of 175°C provides for a larger design margin and relaxed thermal management requirements. GEN2 Series SiC Schottky diodes are available now in either TO-220-2L (quantity 1,000) or TO-252-2L (quantity 2,500) packages.

LITTELFUSE, *www.littelfuse.com*

IPM Integrates PFC and 3-Phase Inverter

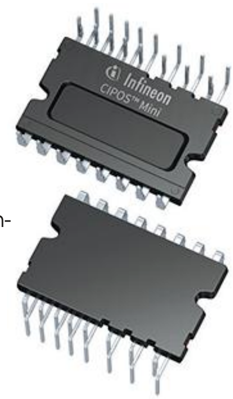

INFINEON TECHNOLOGIES' NEW CIPOS Mini intelligent power module combines a single switch boost PFC stage and a 3-phase inverter in one package. Integrating PFC into the inverter module helps reduce system size. The PFC stage makes use of an external driver circuit and a Trenchstop IGBT combined with an optimized SOI gate driver.

The module is designed to control induction motors and permanent magnet synchronous motors with single-phase PFC in variable speed drives typically found in applications like air conditioning and low-power motor drives of up to 2 kW power capability.

The IPM will be offered in inverter current ratings of 4, 6, 10 and 15 Amp, with a choice between PFC switching frequency of 20 or 40 kHz, and in either DCB or fullpack versions both featuring integrated over-current protection, under-voltage lockout and temperature monitor.

INFINEON TECHNOLOGIES, *www.infineon.com*

DIN Rail Mount Power Supplies 94.5% Efficiencies

THE TIB FAMILY of DIN Rail mount power supplies from TRACO POWER feature efficiencies up to 94.5%, enabling a slimmer package design with a DIN Rail clip that can be mounted on either the rear (typical) or side (flat panel) of the unit. Reduced heat dissipation enables a -20°C to +60°C full load operating temperature range and up to 70°C with only 20% de-rating.

The supplies offer 150% boost power for a minimum of 4 seconds to support stepper motors, solenoids or actuators with true power-back immunity. EMC immunity is compliant for applications in industrial environments (EN 61000-6-2) with emissions compliant with the standard for residential and commercial environments (EN 61000-6-3). The outputs are radio-interference-suppressed to impede radiation at long output lines, reducing common mode current to within limits of telecommunication ports. The units operate with a high power factor up to 99% with active PFC.

The TIB Family is in stock and available now in four power ranges of 80, 120, 240 or 480 Watts with 24 Vdc outputs (adjustable from 23.5 to 28 Vdc).

TRACO POWER, *www.tracopower.com*

RZ/G1E-Based SMARC Module Accelerates Development

MISTRAL SOLUTIONS IS now offering a high-performance SMARC module that leverages the benefit of Renesas' RZ/G1E microprocessor. The smart and compact module consists of the powerful MPU, which incorporates dual ARM Cortex-A7 CPU cores operating at up to 1.0 GHz, and supporting circuits, including memories and power supplies. The module also supports dual cameras, dual display ports, dual Ethernet and most of the interfaces of SMARC 1.0 standard.

The SMARC Dev Kit provides display, video, audio, wireless, USB, connectivity and other interfaces. Based on requirements, application-specific interfaces can also be brought on to a custom carrier board and used along with the SMARC module.

Designed in a 82 x 50 mm package with 1.2 mm PCB thickness, the product is compatible with any design that follows SMARC architecture and can be plugged in to provide enhanced performance and functionality. Suited for applications such as industrial HMI, intelligent cameras, IoT gateway, embedded vision, industrial controllers and V2X infrastructure, the RZ/G1E-based SMARC Module and the SMARC Dev Kit are available now.

MISTRAL SOLUTIONS, *www.mistralsolutions.com*

Compact Motion Controller Raises Resolution

PHYSIK INSTRUMENTE IS now offering a higher performance version of its Mercury stepper motion controller. The C-663.12 Mercury controller is designed to deliver 2,000 more steps per revolution than typical motion controller models. The compact controller is designed for 2-phase stepper motors, in open-loop or closed-loop operation commanded by USB or RS-232. Up to 16 units can be combined via daisy chain to operate multi-axis motion systems. A 48-V wide-range-input power supply and all cables required for operation are included with the product. Programmable digital and analog I/O lines, and input lines for limit and reference point switches controlled via TTL signals, are also integrated.

The motion controller is equipped with a data recorder for high speed tracing, ID chip compatibility for quick start-up and on-the-fly parameter changes, and exchange of system components without recalibration. The C-663.12 Mercury stepper motion controller is delivered with extensive software packages, including drivers for LabVIEW, as well as dynamic libraries for Windows and Linux.

PHYSIK INSTRUMENTE, *www.pi-usa.us*

Mini Linear Smart Stage Travels Further

NEW SCALE TECHNOLOGIES' latest Linear Smart Stage, M3-LS-3.4-15, is a direct-drive precision piezoelectric micro stage with embedded controller designed for integration into compact OEM instruments. The new stage is a higher-force, longer-travel addition to the company's M3-LS line of all-in-one micro stages with built-in controllers.

Able to be assembled into two-axis and three-axis systems, the stage has a 100-gram capacity for vertical loads and 200-gram capacity for horizontal loads. It has a 15 mm travel range and 500 nm position resolution for precise, repeatable positioning of optics, probes, sensors, etc. All electronics are completely integrated in the 32 x 32 x 11 mm stage housing. An embedded controller eases integration of the smart stage into OEM instruments.

The DK-M3-LS-3.4-15 developer's kit includes M3-LS-3.4-15 Linear Smart Stage, USB adapter / breakout board, power supply, software and accessories. It will be available in Q3 for $2,500.

NEW SCALE TECHNOLOGIE, *www.newscaletech.com*

XBee Embedded Modem Provides 4G to Devices

DIGI INTERNATIONAL'S NEW XBee Cellular LTE Cat 1 embedded modem brings together the power and flexibility of the Digi XBee ecosystem with the latest 4G cellular technology. The embedded modem enables the integration of cutting edge 4G cellular technology into devices and applications without dealing with FCC and carrier end-device certifications. A bundled data plan will be included with every development kit including 6 months of free data, with the product fully pre-provisioned and ready to communicate over the cellular network right out of the box.

The product features the authentic Digi XBee footprint and compact form factor, XBee Transparent and API modes, 3.3 V power with LP options for battery powered applications, OTA firmware updates, and built-in security with Digi TrustFence security framework.

With the full suite of standard Digi XBee API frames and AT commands, seamless transition from existing XBee products to the XBee Cellular modem is achieved with only minor software adjustments.

DIGI INTERNATIONAL, *www.digi.com*

Programmable Controller Aids USB-C Charger Design

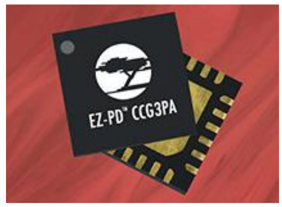

CYPRESS SEMICONDUCTOR IS tnow sampling the EZ-PD CCG3PA USB-C controller for streamlining the design of power adapters, chargers and banks. Supporting the PD 3.0 standard with PPS and Qualcomm's Quick Charge 4.0 protocol, the one-chip solution's high level of integration includes: error amplifiers for constant voltage, constant current and PPS applications; 30 V regulator allowing direct operation from VBUS; VBUS short protection on CC pins; gate drivers for high-voltage power FETs; low side current sense amplifier; dedicated hardware for legacy charger detection protocols; system level ESD protection; and an ARM Cortex-M0 and 64 KB Flash with read-while-write function for firmware upgradeability. The EZ-PD CCG3PA controller with Power Delivery is now sampling in 24-pin QFN and 16-pin SOIC, and will be in production in the third quarter of 2017. The CY4532 CCG3PA evaluation kit is available now for $149. CCG3PA solution reference designs for a 45-Watt notebook PC power adapter, 27 W mobile phone charger, 60 W car charger, and power bank are also available.

CYPRESS SEMICONDUCTOR, *www.cypress.com*

AC-DC Power Suppy Uses Transphorm GaN FETs

DEVELOPED BY BEL POWER SOLUTIONS, the ac to dc front-end TET3000-12-069RA power supply uses a GaN-based bridgeless totem-pole PFC topology to achieve greater than 96% efficiency. The 3 kW ac-dc PFC dc/dc power supply converts standard ac mains power or high voltage dc bus voltages into a main output of 12 Vdc for powering IBA in high performance and enterprise reliability servers, routers and network switching subsystems. It provides 31.7 W/in3 volume power density from a 2.72" x 1.59" x 21.85" device that meets 1U end-system design requirements.

The isolated power supply uses Transphorm's TPH3205WSB 49 mΩ GaN FETs, which offer lower gate charge, faster switching speeds and smaller reverse recovery charge than typical silicon devices. Compared to a standard interleaved boost converter, GaN designs implementing the bridgeless totem-pole PFC topology also achieve lower component count, lower EMI filter size, removal of bridge rectifier and increased on-resistance. The TET3000-12-069RA power supply provides a 244 A output current, 100 - 277 Vac input and 0°C - 50°C operating temperature range.

BEL POWER SOLUTIONS, *https://belfuse.com/power-solutions*

Superscalar Processors Improve RISC and DSP

THE DESIGNWARE ARC HS4x and HS4xD family from Synopsys, the highest performance processors in their ARC HS line, implement a dual-issue superscalar architecture that delivers up to 6,000 DMIPS per core and supports more than 150 DSP-optimized instructions. Available in single-, dual- and quad-core configurations, the configurable processors are designed to meet the requirements of embedded applications like SSDs, wireless baseband, home networking, advanced HMI, and industrial and automotive control..

SYNOPSYS, *www.synopsys.com*

1.25 mm High-Reliability Connectors Boost Pin Counts

HARWIN IS DOUBLING the range of their high-reliability Gecko Screw-Lok 1.25 mm pitch connector series with four new pin count sizes (20-, 26-, 34- and 50-contact). All sizes are available in all existing connector designs including vertical PCB throughboard and SMT in male and female, horizontal PCB throughboard in male, cable housings in male and female, and metal backshells for male and female cable housings.

All products feature ruggedized 'mate before lock' stainless steel screw fixings with inbuilt corrosion resistance against salt spray and humidity. The patented single piece 4-finger beryllium copper female contacts can withstand up to 1,000 repeated mating cycles without any loss of performance.

The micro signal contacts are rated to 2.8 A individually, and to 2.0 A for all contacts simultaneously. Provided in packages up to 45% smaller than conventional Micro-D connectors, Gecko-SL connectors can operate at temperatures from -65°C to +150°C while enduring 20 g vibration and Z axis 100 g 6 m/s shock, with no discontinuity greater than 1 µs.

HARWIN, *www.harwin.com*

Advertisers Index

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

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

Revisiting the Planet of the Apes

"The War for the Planet of the Apes" opens July 14th. Technology Editor Bill Wong takes another look at the film's behind-the-scenes marvels.

It's tough being an ape—especially with Woody Harrelson's obsessed Colonel McCullough looking to wipe you out. Andy Serkis, as the genetically enhanced ape Caesar, will need a bit of help to turn the tables. "The War for the Planet of the Apes" is the third in the series. Each has been more ambitious from a technical standpoint than the previous version.

This time around, visual effects producer Ryan Stafford had his work cut out for him and his crew *(see Go Behind the Scenes with "War for the Planet of the Apes")*. There are more chimpanzees, gorillas, and apes in this movie and the conditions are more complex. Rain and snow have their own challenges, but when the fur is flying it also gets wet *(Fig. 1)*. This required significantly improved software to handle the large number of entities within a frame, such as the battle scenes, as well as the requirement to provide realistic renditions of the non-human characters.

The actors playing the non-human characters often wore cameras on small booms *(Fig. 2)* so their facial expressions could be properly rendered. The dots on the actor's face were added daily in the same positions. This approach simplifies generation of the special effects but it can be challenging especially when acting in close quarters. Motion capture also necessitates repeatedly shooting a scene with and without various actors to properly record the background for use when rendering the entire scene.

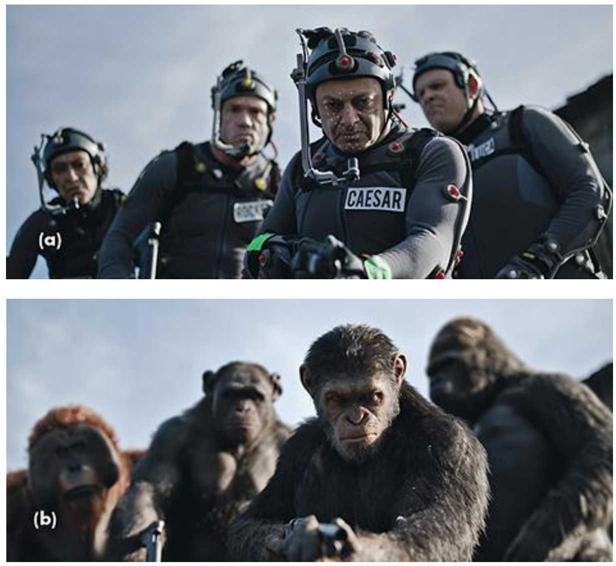

2. The booms that each actor is wearing have a small camera (a) so the facial details can be properly rendered (b).

1. Simulating fur is hard. Wet fur is harder, but Andy Serkis' Caesar (a) looks as natural as can be once the final image is rendered (b).

Computer animation software does a lot of the heavy lifting when it comes to turning actual video into the final product, but there's still a significant amount of artistry involved because of the level of accuracy that is being delivered in the final film. For example, many shots were done with snow where a trail and footprints need to be properly rendered. The prosthetics used by the actors make much difference paths through the snow. Software does not always handle this adjustment well so artists need to clean up the final version. It is a major logistics challenge.

You will need to check out the "War for the Planet of the Apes" when it opens on July 14th to find out what director Matt Reeves has in store for viewers, as well as how well it will look. Maybe on the second viewing you can look for the visual details. I'm definitely looking forward to it. ed