Altech Corp.®

PCB Terminal Blocks

for Automation & Control Applications

EXPANDED - *now with over* **20,000** *parts*

Selection & Styles

visit altechcorp.com/PCBs

Spring Clamp & Push-In

FIXED PUSH-IN

- 2.50 / 2.54 / 3.50 / 3.81 / 5.00 5.08mm pin spacing
- 2A/ 6A/ 10A current rating
- 150V/ 300V
- Vertical/ horizontal
- Lever version

PUSH-IN PLUGS

- 3.50/ 3.81 / 5.00 / 5.08mm pin spacing
- 10A current rating
- 300V
- Screw flange

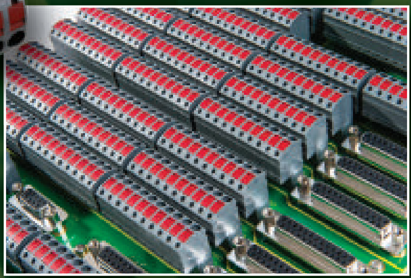

PCB POWER TERMINALS PUSH-IN

- 5.00 / 7.50 / 11.50mm pin spacing
- 20A / 35A current rating
- 300V
- Vertical/ horizontal

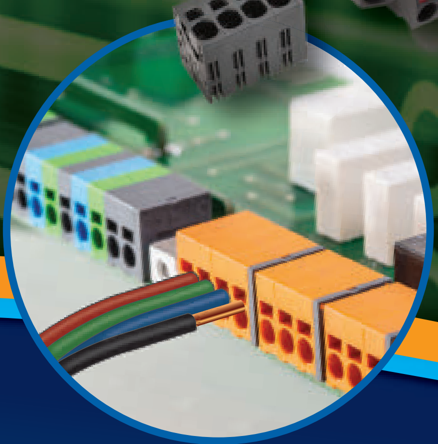

TENSION SPRING TERMINALS

- 2.50/ 2.54 / 5.00 / 5.08 / 7.50 / 10.00mm pin spacing
- 4A / 6A / 15A / 16A current rating
- 150V / 300V
- Lever version
- Double level

Altech Corp.® 40 years 1984 - 2024

Automation & Control Components

AltechCorp.com
908.806.9400

Electronic Design®

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

MARCH/APRIL 2024
electronicdesign.com

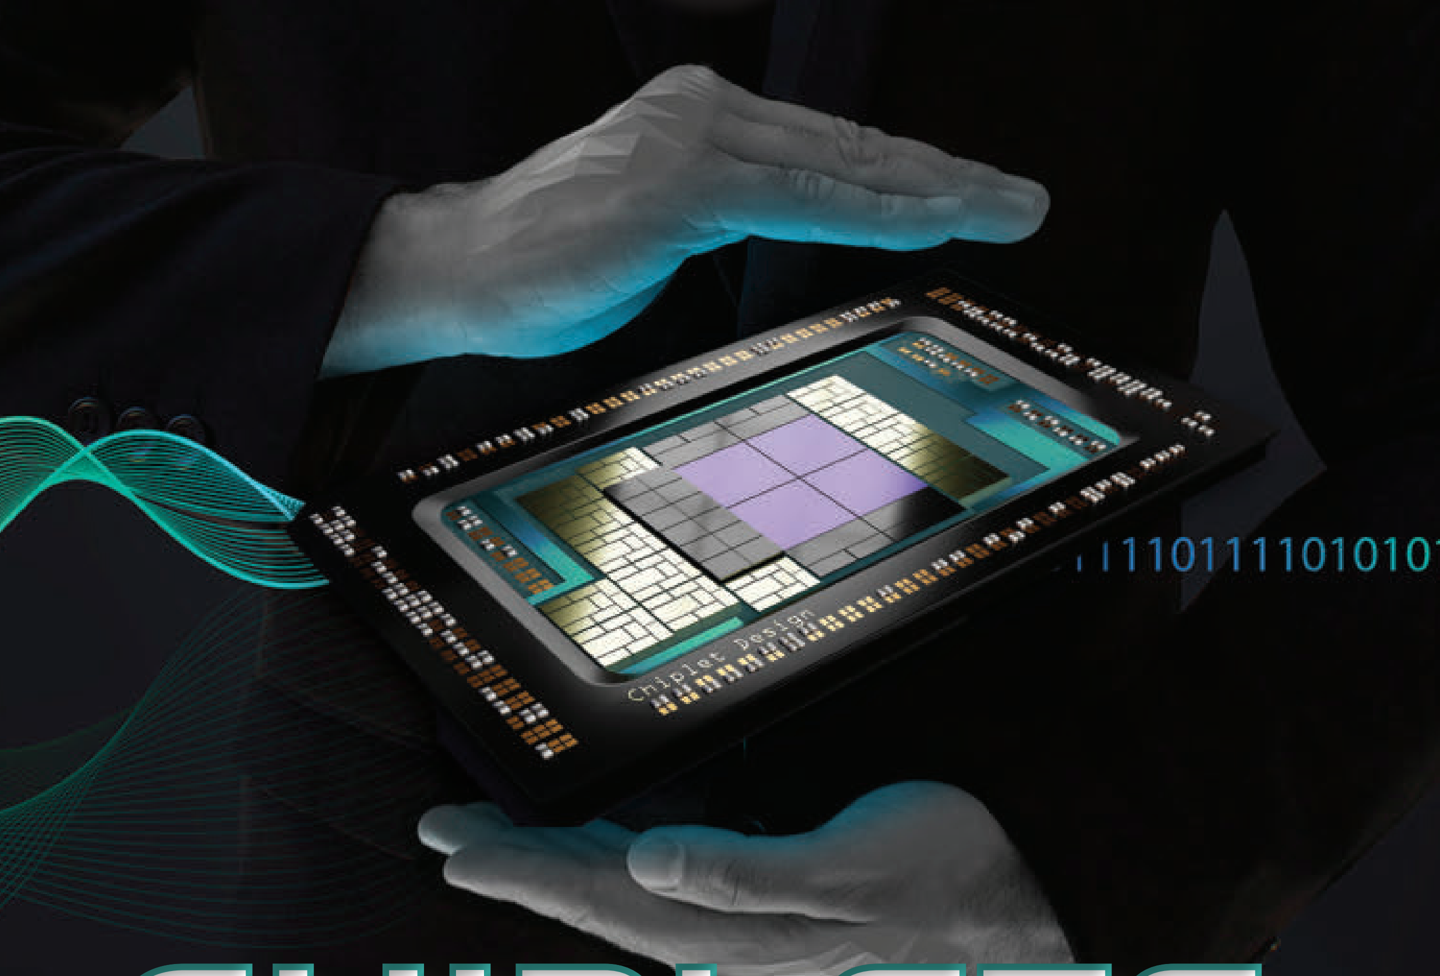

CHIPLETS:

The Agent of Change in Packaging

THERE'S A KEYSTONE IN EVERY GREAT INVENTION.

REFRIGERATION & THEIR ELECTRONICS

Featuring Metal Key-Pad Dome Switches

Cold storage has had a profound impact on the development of human culture and the modern world. Cold storage technology has enabled people to migrate into hotter, dryer climates without concern for food spoiling in the heat. The first evidence of cold storage was recorded in Syria more than 3,700 years ago. Archeologists and historians have also verified that the ancient Chinese and Persians also harvested ice.

The first commercial ice-making machine was invented in 1854. In 1913, refrigerators for home use were invented and in 1923 Frigidaire introduced the first self-contained unit.

Refrigerator and freezer technologies have progressed far beyond just cold storage. Today's smart fridge and freezer units feature water & ice dispensers, user control panels & displays, Wi-Fi connectivity, voice recognition & control, and more.

You'll find a wide range of Keystone products in modern refrigerator and freezer systems such as our featured Metal Key-Pad Dome Switches as well as • Battery Clips, Contacts & Holders • SMT & THM Test Points • LED Spacers & Lens Caps • Pins, Plugs, Jacks & Sockets

IT'S WHAT'S ON THE INSIDE THAT COUNTS

KEYSTONE®
ELECTRONICS CORP.

Designers & Manufacturers of Electronic Components and Hardware

View our Dynamic Catalog M70 at www.keyelco.com
(516) 328-7500 • (800) 221-5510

For More Details
See our Blog

IN THIS ISSUE

FEATURES

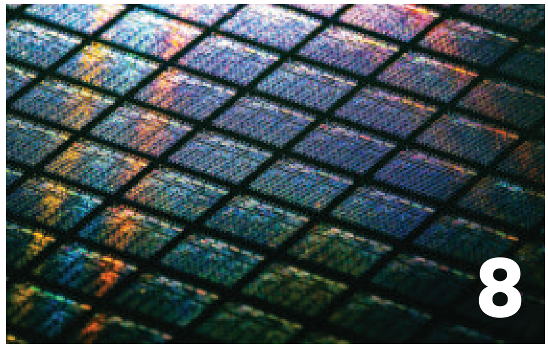

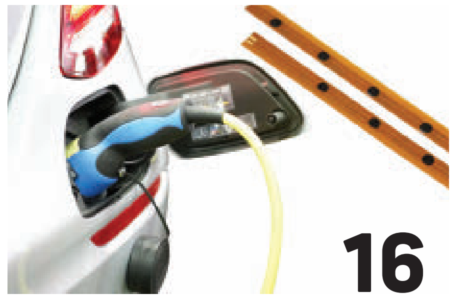

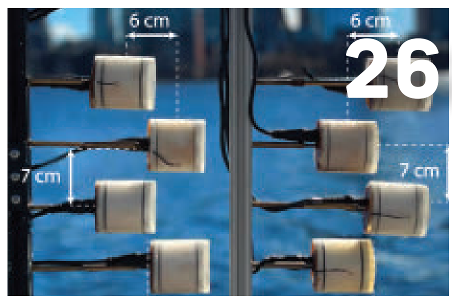

COLUMNS & DEPARTMENTS

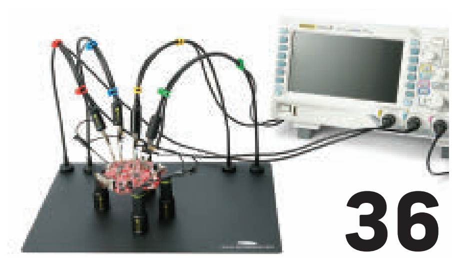

Cover images: Devrimb_1169974480 | iStock / Getty Images Plus
Peter Hansen_1417034864 | iStock / Getty Images Plus

Editorial

BILL WONG | Senior Content Director

Packaging Chiplets
for Performance and Profit

When it comes to chiplets, it's all about packaging technology.

MULTI-DIE SYSTEMS UTILIZING technologies such as chiplets and the Universal Chiplet Interconnect Express (UCIe) are essential to deliver high-performance solutions and compact applications that require a mix of technologies like memory, compute, and analog. Advanced packaging fosters die-to-die, die-to-chip, and chip-to-chip solutions often referred to as chiplet technology.

Chiplets are used regularly in high-performance solutions where the cost of using new technology is offset by cutting-edge performance. NVIDIA's latest Blackwell GPU and AMD's Ultrascale FPGAs are examples of chiplet-based solutions that incorporate features like high bandwidth memory version 3 (HBMv3) to deliver the speed and capacity that would not be possible with a single monolithic chip.

I hosted the Chiplets in 2029: How We Get There webinar at this year's Chiplet Summit, where we talked about chiplet-related technologies such as UCIe and interposers. We're just at the start of the chiplet revolution, but it will be getting more affordable and gaining wider support. It can potentially reduce the cost of developing a new chip while offering a way for companies to deliver their intellectual property (IP) in the form of a chiplet.

Though the comparison of integrated circuits (ICs) and printed circuit boards (PCBs) to chiplet technology is easy to understand, the comparatively low pin count of ICs pales in comparison with a die-to-die connection that might include thousands of connections. There's also a variety of techniques that enable connecting from the Intel/Altera Embedded Multi-die Interconnect Bridge (EMIB) to vertical stacking with direct bonding.

The idea of a chiplet store is still that—an idea—but it may play out within the next decade. Designing a chiplet may be easier than a very large SoC, though not by much. While EDA remains an expensive game, it's one where utilizing the latest chiplet technology can have a big payoff. ed

Inside Electronics Podcast:

The Importance of **Connectivity in Electronic Systems**

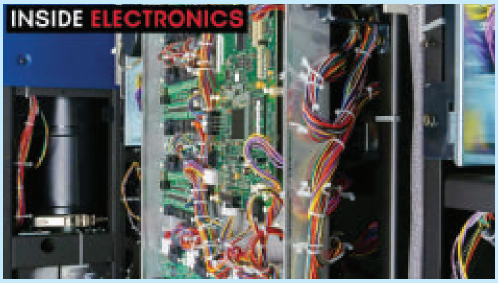

Host Alix Paultre discusses the latest in connectivity solutions, from cables to connectors, with Cinch's Scott Miller, Director of Product Management.

Click below to check out more content in the TechXchange: Inside Electronics Podcast.

CLICK HERE

AD INDEX

FOLLOW US ON FACEBOOK AND TWITTER

(ElectronicDesign) (@ElectronicDesgn)

ElectronicDesign

EDITORIAL

Senior Content Director: **Bill Wong**, *bwong@endeavorb2b.com*
Executive Editor: **David Maliniak,** *dmaliniak@endeavorb2b.com*
Managing Editor: **Roger Engelke,** *rengelke@endeavorb2b.com*
Editor-at-Large: **Alix Paultre,** *apaultre@endeavorb2b.com*
Senior Editor: **James Morra,** *jmorra@endeavorb2b.com*
Technology Editor: **Cabe Atwell,** *catwell@endeavorb2b.com*

CONTRIBUTING EDITORS

Lee Goldberg, Bill Schweber, Murray Slovick, Steve Taranovich, Alan Earls

ART DEPARTMENT

Group Design Director: **Anthony Vitolo,** *tvitolo@endeavorb2b.com*
Art Director: **Jocelyn Hartzog,** *jhartzog@endeavorb2b.com*

PRODUCTION

Production Manager: **Brenda Wiley,** *bwiley@endeavorb2b.com*
Ad Services Manager: **Deanna O'Byrne,** *dobyrne@endeavorb2b.com*

AUDIENCE MARKETING

User Marketing Manager: **Debbie Bouley** *dbouley@endeavorb2b.com*
Article Reprints: *reprints@endeavorb2b.com*

SUBSCRIPTION SERVICES

OMEDA, 847-559-7598 or 877-382-9187
electronicdesign@omeda.com

LIST RENTAL

List Rentals/Smartreach Client Services Manager: **Kelli Bery**
kberry@endeavorb2b.com

SALES & MARKETING

Gregory Montgomery, *gmontgomery@endeavorb2b.com*
AZ, NM, TX

Jamie Allen, *jallen@endeavorb2b.com*
AL, AR, SOUTHERN CA, CO, FL, GA, HI, IA, ID, IL, IN, KS, KY, LA, MI, MN, MO, MS, MT, NC, ND, NE, OH, OK, SC, SD, TN, UT, VA, WI, WV, WY, CENTRAL CANADA

Elizabeth Eldridge, *eeldridge@endeavorb2b.com*
CT, DE, MA, MD, ME, NH, NJ, NY, PA, RI, VT, EASTERN CANADA

Stuart Bowen, *sbowen@endeavorb2b.com*
AK, NORTHERN CA, NV, OR, WA, WESTERN CANADA
AUSTRIA, BELGIUM, FRANCE, GERMANY, LUXEMBURG, NETHERLANDS, PORTUGAL, SCANDINAVIA, SPAIN, SWITZERLAND, UNITED KINGDOM

Diego Casiraghi, *diego@casiraghi-adv.com*
ITALY

Helen Lai, *helen@twoway-com.com*
PAN-ASIA

Charles Liu, *liu@twoway-com.com*
PAN-ASIA

DIGITAL

VP of Digital & Data Innovation: **Ryan Malec,** *rmalec@endeavorb2b.com*

ENDEAVOR BUSINESS MEDIA, LLC

30 Burton Hills Blvd, Ste. 185, Nashville, TN 37215 | 800-547-7377
www.endeavorbusinessmedia.com

CEO: **Chris Ferrell**
President: **June Griffin**
COO: **Patrick Rains**
CRO: **Paul Andrews**
Chief Digital Officer: **Jacquie Niemiec**
Chief Administrative and Legal Officer: **Tracy Kane**

DESIGN & ENGINEERING GROUP

EVP, Design & Engineering Group: **Tracy Smith**, *tsmith@endeavorb2b.com*
Group Content Director: **Michelle Kopier,** *mkopier@endeavorb2b.com*

Electronic Design, Machine Design, Microwaves & RF, Power & Motion, SourceESB, Source Today, 3DX

VIDEO ▸

GaN Delivers More Compact Power Supplies

Gallium-nitride (GaN) transistors are able to handle large currents at high switching speeds, leading to switching power supplies that are significantly smaller than their silicon-based counterparts. That's because faster switching speeds and lower on-resistance of GaN FET transistors let designers employ smaller magnetics, reducing the size of the solutions. Unfortunately, GaN transistors can be a bit fragile, requiring additional protection to provide a more robust solution. Editor Bill Wong talked with Stephen Oliver, VP Corporate Marketing & Investor Relations at Navitas Semiconductor, about where the company's latest GaN technology is being deployed.
www.electronicdesign.com/21281913

VIDEO ▸

ePaper Not Just for e-Readers Anymore

Many myths abound about ePaper technology employed in e-readers, such as that it's only slightly more efficient than LCD displays. That's far from the truth. It's very power efficient and one of the few display technologies that can retain a static display after power is removed, making for even more efficient power use. Editor Bill Wong talked with Timothy O'Malley, Assistant Vice President of U.S. Business at E Ink, about some of the new uses of ePaper and some of the latest advances.
www.electronicdesign.com/21282236

How Smart is Smart Energy?

Smart energy is the mode of using various methods and devices for energy efficiency. It focuses on powerful and sustainable renewable-energy sources that will benefit greater eco-friendliness while lowering costs.
www.electronicdesign.com/21283937

The Origins of **Electrical Engineering Education**

Electrical engineering is often revered as one of the most difficult educational curriculums available. It's a well-respected, tough, but rewarding field of work. This article reaches back to the late 1800s, when electricity was first beginning to make an impact on society.
www.electronicdesign.com/21284104

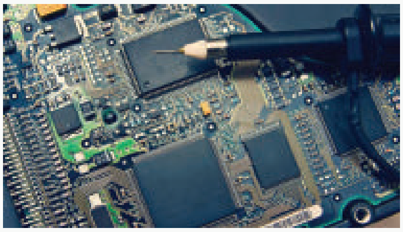

"Probing" the Future of **Power Integrity** Test & Measurement

Power integrity is all about delivering the appropriate power from the voltage regulator module to the load to maintain reliable operation under dynamic loading conditions. Thus, design margins must be maintained. But AI accelerators and other high-performance chips are creating challenges.
www.electronicdesign.com/21283129

JAMES MORRA | Senior Editor

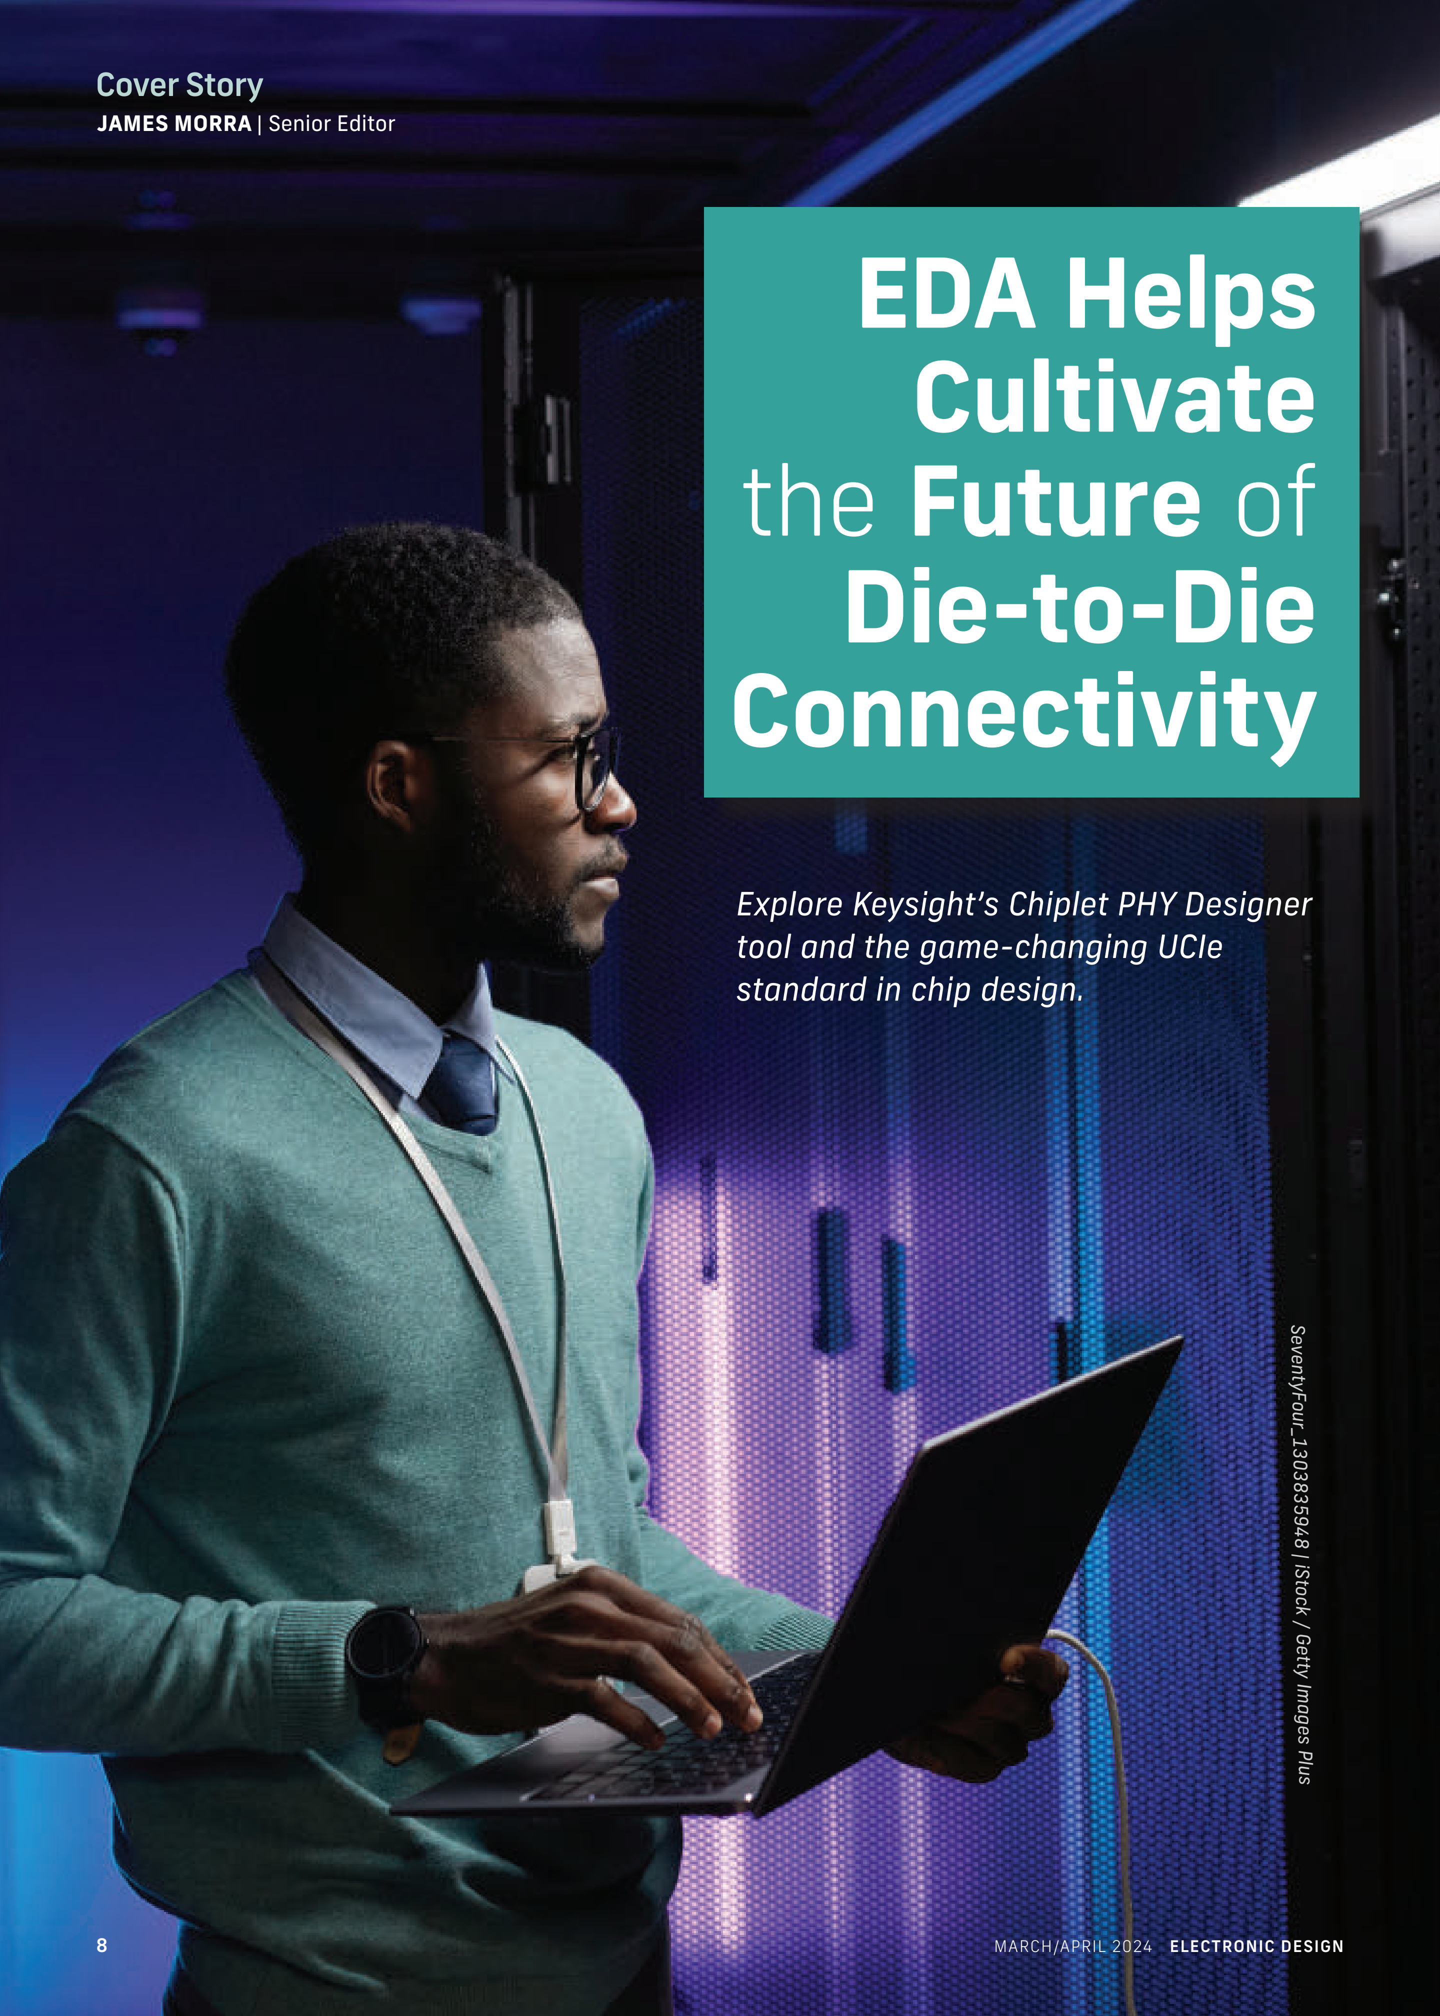

EDA Helps Cultivate the Future of Die-to-Die Connectivity

Explore Keysight's Chiplet PHY Designer tool and the game-changing UCIe standard in chip design.

SeventyFour_1303835948 | iStock / Getty Images Plus

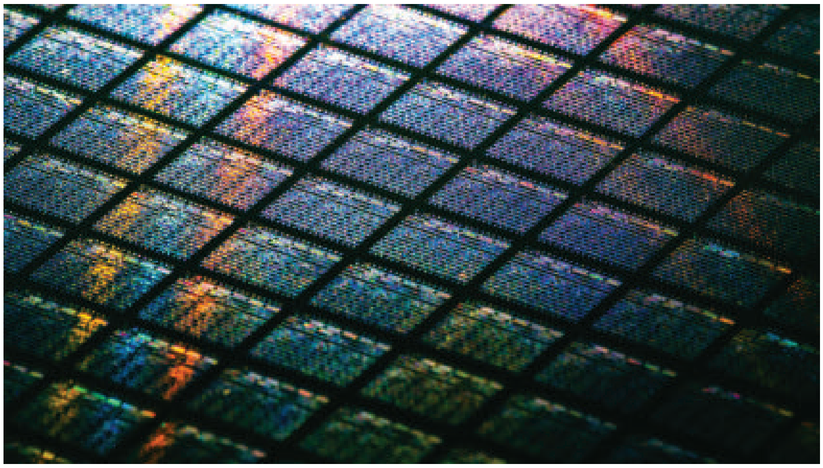

91824901 © Forance | Dreamstime.com

TODAY, MANY OF the world's most advanced processors are no longer single monolithic slabs of silicon. They're comprised of a collection of smaller silicon die commonly referred to as chiplets that use advanced 2.5D or 3D packaging to mimic a single large chip.

On the edge of every chiplet is a PHY that enables high-bandwidth, low-latency connectivity with the other devices in the package, which all interact with each other over proprietary or industry-standard protocols. As success in the semiconductor industry becomes more about what companies can cram into a package instead of a monolithic chip, the quality of these die-to-die (D2D) interconnects, typically based on high-speed SerDes, can make or break the design.

While they typically bridge very short distances, these D2D interconnects send data around the package at high speeds. This can negatively impact the bit error rate (BER) in the channel as well as the performance of the chiplets.

Keysight Technologies is trying to stay a step ahead with a new simulation tool called Chiplet PHY Designer, which I learned about first-hand at DesignCon 2024. Using it, Keysight said chip designers can quickly and accurately validate the PHY-based D2D interconnects in a package before investing in actual silicon from a foundry.

The ability to model and test D2D connectivity is key to the future of heterogeneous chip design. Doing so helps make sure the chiplets inside a 2.5D and 3D package can interact seamlessly.

"Chiplet PHY Designer accelerates validation of chiplet subsystems, from one D2D PHY through interconnect channels to another D2D PHY, much earlier in the design cycle," said Niels Faché, VP and GM of Keysight EDA.

Chiplet PHY Designer is the latest tool in Keysight's Pathwave ADS platform for high-speed digital design and simulation. It's also uniquely able to verify that chiplets meet the Universal Chiplet Interconnect Express (UCIe) standard. UCIe is one of several emerging standards in D2D interconnects. UCIe gives the chiplets an easier way to communicate, reducing friction when

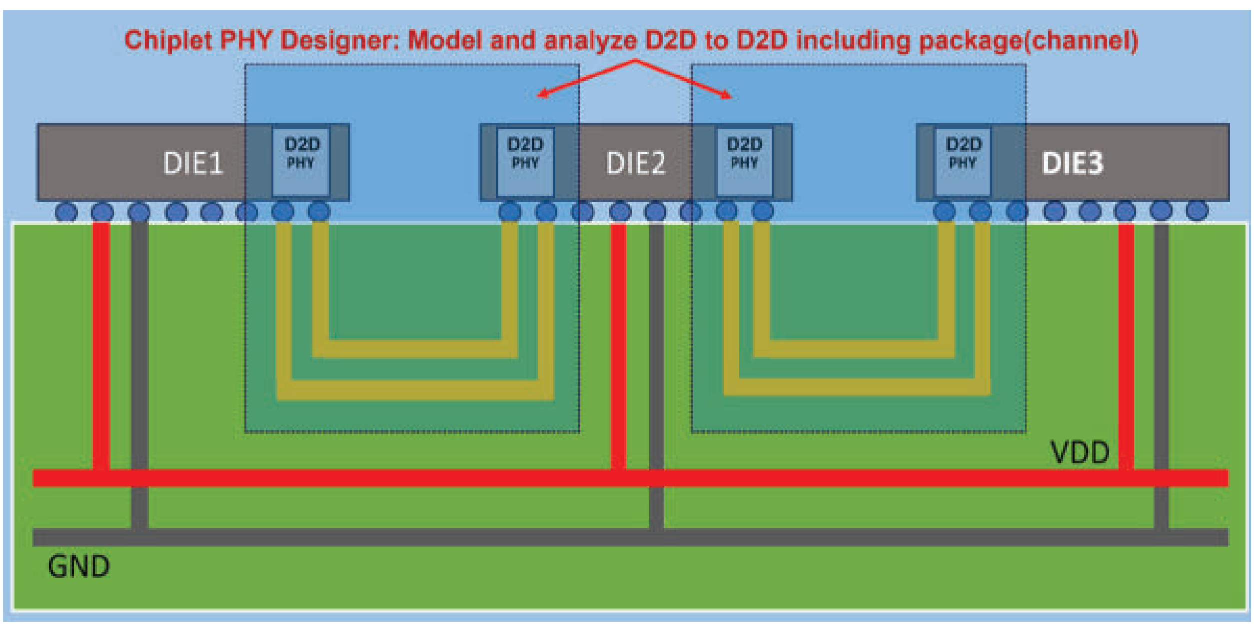

A cross-section of a chip based on 2.5D advanced packaging. Keysight is trying to apply the Chiplet PHY Designer in this arena. *Images courtesy Keysight* it comes to mixing and matching IP from third parties.

The company said Chiplet PHY Designer presents a faster and easier way to simulate the electrical performance of D2D interconnects based on UCIe and test out designs against the latest revisions of the specification.

"Chiplet PHY Designer accelerates validation of chiplet subsystems, from one D2D PHY through interconnect channels to another D2D PHY, much earlier in the design cycle."
—Niels Faché, VP and GM, Keysight EDA

What Are the Pros and Cons of Chiplets?

Flexibility is a hallmark of chiplet technology, said Stephen Slater, product manager for Keysight's high-speed digital simulation technology. The concept of chiplets allows companies to adapt to evolving markets and new technological advancements by mixing and matching chiplets, he told *Electronic Design*.

Instead of placing everything on a single slab of silicon as usual, the movers and shakers in the semiconductor industry are breaking up complex systems-on-chips (SoCs) into smaller silicon dies that have separate pieces of IP, including CPUs, GPUs, AI accelerators, memory, I/O, and various other chip functionalities. After being tested, verified, and validated, the chiplets can be mixed and matched in a package using silicon interposers or other 2.5D packaging technologies, such as TSMC's CoWoS or Intel's EMIB, or more advanced 3D stacking.

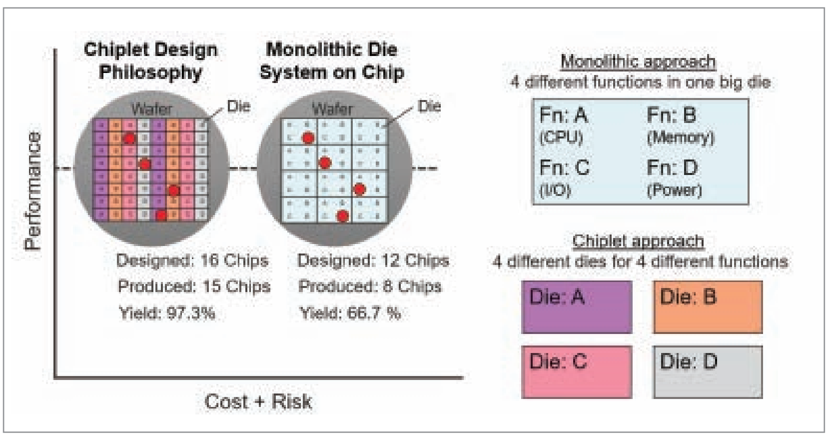

A summary of the advantages of chiplets.

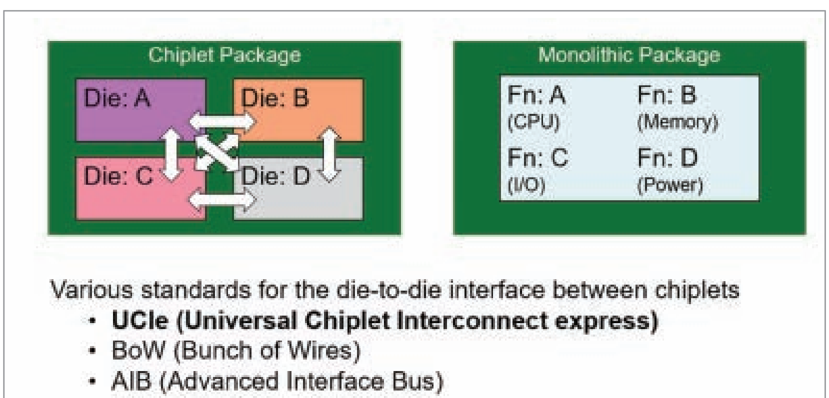

A summary of the role of die-to-die (D2D) interconnects.

In many cases, companies are pulling apart these larger processors and repackaging them because they're too expensive or impossible to build on a single silicon die due to the physical limits of Moore's Law. By spreading out the SoC's functionality over several chiplets, you can cram more transistors in the package than you can in a single processor due to the "reticle limit" of the wafer manufacturing process, said Slater.

The chiplets themselves interact with each other using high-speed, short-range SerDes PHYs and establish connections over UCIe or various other die-to-die interfaces such as the Open Compute Project's (OCP) Bunch-of-Wires (BoW) and the Common Electrical I/O (CEI) Extra-Short Reach (XSR) standard.

Dismantling SoC designs into a smaller set of chiplets makes for easier heterogeneous integration, noted Slater. Every chiplet can be built on the best process technologies for the job, helping reduce complexity and costs.

As process technology continues to evolve, specific chiplets can be upgraded while others are left alone. Slater added that the capacity of chiplets for swift customization and upgrades means faster time-to-market and lower costs.

He said chiplets are also physically smaller than the large SoCs that they're stripped out of. Consequently, a larger number can be manufactured per wafer without imperfections, increasing yields. Higher yields translate to lower production costs.

Partitioning a processor design and repackaging the building blocks in a way that mimics a monolithic chip is no easy feat. These chips resemble SoCs on a module, so they tend to be vulner-

able to signal integrity, timing, and other issues that are more often encountered by system integrators. Dissipating heat from these IC packages is also complicated, requiring thermal simulation both early in the design and during integration into products.

But the pros tend to outnumber the cons for advanced heterogeneous chip designs. Chiplets are considered key to the future of high-performance chips required for everything from artificial intelligence and machine learning (AI/ML) to augmented reality (AR).

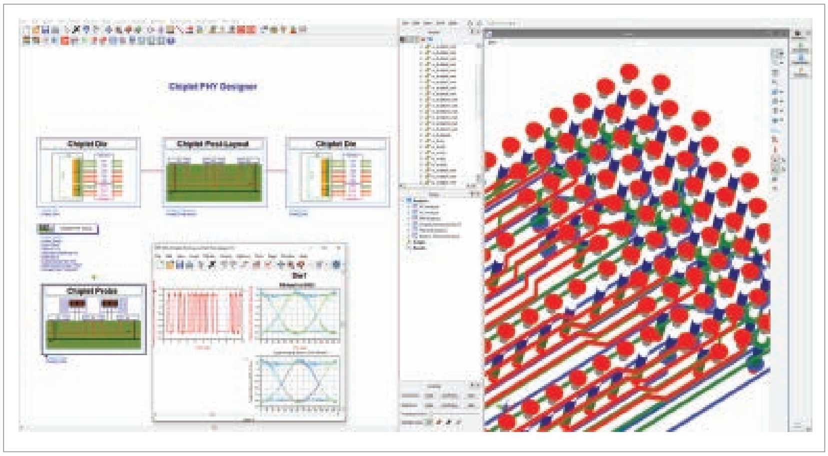

Keysight's Chiplet PHY Designer in action.

UCIe: Bridging the Connectivity Gap for Chiplets

Chiplets are unlocking a new wave of innovation in chips. The issue for the semiconductor industry is the lack of a standard plug-and-play way for chiplets to communicate once they're bound together in an IC package.

Today, companies tend to use proprietary D2D interconnects and protocols that lock them into using IP designed internally or sourced externally and then validated and tested in-house. In short, they lose the flexibility they're accustomed to with third-party IP. While they still have the freedom to utilize any foundry's packaging technologies, using third-party chiplets can be more trouble than it's worth without a standard D2D interface.

UCIe presents one way to solve these problems. It fills the gap for industry-standard D2D interconnect that allows for the mixing and matching of chiplets—no matter the company that designed it, the foundry that manufactured it, or the type of advanced packaging (2D, 2.5D, or 3D) used. First released in 2022, the UCIe spec covers the physical I/O with data paths on physical bumps grouped into lanes to the protocol stack as well as everything required for testing and validating the D2D interconnect.

The goal is to create a vibrant ecosystem for chiplets, so that companies can buy and sell each other's pre-validated chiplets and plug them into a package almost as seamlessly as components on a circuit board. Instead of spending money on a full SoC, engineers can focus on developing the specific chiplets they need and rely on commoditized technologies for other aspects of their chip designs. Call it the "chiplet economy."

UCIe is already being supported or adopted by many of the top semiconductor equipment and electronic design automation (EDA) vendors as well as leading foundries and chip designers like AMD, Intel, NVIDIA, and TSMC.

The Role of Keysight's EDA Software in Chiplet Design

With Chiplet PHY Designer, Keysight is catering to semiconductor companies with a stake in UCIe's future.

The test-and-measurement giant said it adapted its high-speed digital design and simulation technology to help deal with the single-ended signaling and forwarded clocking of UCIe-based D2D interconnects.

Chiplet PHY Designer can automatically parse the signals traveling through the UCIe-based interconnect. It's also able to automate wiring connections between the chiplets inside the package to evaluate and perhaps improve the signal integrity. Furthermore, according to Keysight, the EDA software supports a standard-driven simulation setup such as speed grade and a special probe component for more intuitive measurement setup.

Chiplet PHY Designer can also model and simulate the voltage transfer function (VTF) of the interconnect—a ratio of the input voltage of a signal as it enters the interconnect to output voltage on the other side—to make sure that the PHY meets the UCIe standard. Since it's influenced by signal loss and crosstalk, the VTF must be measured very accurately to guarantee the delivery of signals from one PHY to another, noted Slater.

Keysight said Chiplet PHY Designer helps companies accurately model and simulate D2D interconnects, and it's accurate enough to rival the precise time-domain measurements of high-speed oscilloscopes—it can assess the BER of the system to within 1e-27 or 1e-32. It can also measure eye diagram height, eye width, skew, mask margin, and BER contour. The term "eye diagram" refers to the pattern of a high-speed signal.

By meeting the latest IBIS modeling specification for electrical I/O, physical integrators can dig even deeper into the PHY electrical validation process, leveraging detailed models to do "pre-silicon" performance predictions.

High Voltage Power Inductors

Coilcraft offers a wide range of voltage-rated power inductors that assure the safe, reliable operation of your designs.

Coilcraft voltage-rated inductors support applications such as high-voltage DC-DC and AC-DC power supplies, industrial and home automation controls, appliances, and automotive, including electric vehicles (EVs).

Our XEL and XGL Families of molded power inductors offer operating voltage ratings from 60 V to 120 V, and many of our LPS and MSS Family inductors are suitable for 400-volt operation or higher.

Our new MSS1514V Series offers an 800 V rating in a 15.5 x 15.5 mm footprint, with excellent current handling and low DCR.

Learn more about operating voltage ratings for inductors and see our full line-up at **www.coilcraft.com/HighVoltage.**

Coilcraft®

WWW.COILCRAFT.COM

STELIOS DIAMANTIDIS | Distinguished Architect, Synopsys Center for Generative AI

Generative AI to Usher in Big Opportunities in 2024

Learn what's in the forecast for generative AI this year and how it will impact the electronics industry.

IT'S OFFICIAL! In less than a year, generative AI (GenAI) has managed to reach Gartner's Peak of Inflated Expectations, likely in record time. Yet, all of the talk and new applications around ChatGPT this year have demonstrated that GenAI is not only here to stay, but it's poised to truly alter the way knowledge work is done.

GenAI is, indeed, one of those transformative technologies that only comes about once in a blue moon. We're seeing examples of it in action in everything from real-time chat functions on e-commerce sites to code generation in software development and much more.

Last year, when I peered into my crystal ball to highlight predictions for AI in 2023, I saw a number of trends taking shape. I predicted that companies would simplify the integration of heterogeneous components (which is happening with solutions for multi-die systems), that we'd see growing adoption of AI design tools (at Synopsys we've tracked over 300 AI-driven commercial tapeouts, and the trend is accelerating), that generative AI would accelerate application development (it's quickly becoming the world's new application platform), and that AI would be instrumental in the pursuit of net-zero carbon emissions (this is a work in progress).

What's on tap for AI in 2024?

Let's take a closer look at how GenAI might shape the electronics industry in 2024 and the key considerations that will need to be addressed as AI becomes more prevalent.

Copilots: Coming to a Design Tool Near You

If you follow Microsoft's Satya Nadella, then you already know this is "The Era of Copilots" (if not, I recommend watching Satya's recent keynote at Ignite 2023).

The semiconductor industry is facing a convergence of challenges and opportunities. Scale and systemic complexities continue to grow, while a looming engineering talent shortage threatens to stifle innovation. At the same time, supply-chain pressures intertwined with geopolitical headwinds persist.

Still, engineering ingenuity continues to deliver breakthroughs that are extending the advantages of Moore's law, enhancing engineering productivity, and, ultimately, producing some of the most sophisticated electronic devices and systems we've ever had.

Against this backdrop, GenAI figures to play an increasingly prominent role in the new year. Considering the very rich knowledge domain that's electronic design automation (EDA), and the short supply of new talent to complement deep experience, chip design copilots can provide an efficient path for knowledge-sharing and enhance productivity of existing engineering resources.

For example, rather than having to pause to research solutions in any given design context—potentially wading through multiple documents and/or contacting an applications engineer—a copilot embedded into a tool could deliver responses to queries almost instantaneously.

As 2024 progresses and GenAI-driven copilots gain a deeper foothold across different industries, we can expect to see these copilots becoming smarter and ubiquitous—being deeply integrated into the workflows used to design and verify silicon chips. From a chip design perspective, copilot capabilities could help engineers do things in seconds that previously required hours or even days.

Data Ecosystems Will Begin to Take Shape

GenAI holds tremendous promise for large, established companies with deep design expertise. Industry leaders will deploy their own copilots to operationalize treasure troves of methodology, architecture, and other domain-specific data accumulated over decades of experience (check out NVIDIA's ChipNeMo paper for an excellent example).

At the same time, in 2024, GenAI will further expand democratization of the chip design process. This will enable new silicon pioneers to innovate and scale faster than ever before, and to focus on their core value proposition while tapping into industry-standard reference flows and optimization knowledge.

EDA companies and IP providers will play an important role in bringing together their deep expertise, flows, and IP with customers' own domain data to create powerful GenAI solutions across the entire technology stack. In 2024, we will likely see the formation of some early data ecosystems in chip design, similar to the ones discussed by OpenAI's Sam Altman at OpenAI DevDay 2023.

Such data partnerships will drive the availability of large-scale datasets that reflect chip design domains across multiple modalities of code, specs, register-transfer-level (RTL), simulations, etc. This will make it possible to train better, more efficient models for GenAI applications.

These partnerships, though, will not take flight unless new, scalable, and secure business models allow for sharing of data in a secure and economically viable manner. Here, the emergence of hosted micro-services offers a model that companies will look to explore and put into production use (for a good discussion, see NVIDIA's recent announcement at AWS re:Invent).

Chip Design Will Become Even More Software-Centric

Software product development AI, such as GitHub Copilot, has gained rapid adoption and was powering over 46% of new code as of earlier this year. Hardware systems design for semiconductors has already started to evolve along a similar software-centric design trend. GenAI has the potential to transform the chip design flow, from design authoring and RTL, through the various steps in design verification and implementation, leading toward systems automation.

In an engaging blog on "Bringing Generative AI to Semiconductor and Electronics Design," Microsoft recently outlined how pervasive, agile software development practices are likely to continue to influence the hardware design world. This trend will more than likely pick up pace in 2024, as several capabilities in the semiconductor industry are already becoming mature and challenging the current "waterfall" model for chip design.

For example, GenAI will make it possible to create ultra-fast prototyping flows. It will bring highly automated optionality to design authoring and creation, and abstract away the labor-intensive, fine-grain development of supporting collateral, such as verification coverage models, complex assertions, or constrained-random test stubs.

The AI Economic Wall Will Drive More AI Hardware Innovation

It's no secret that today's AI workloads are highly computationally intensive. Pre-training large language models (LLMs), in particular, requires access to AI supercomputing systems and silicon chips that are in very high demand.

As the number of parameters in neural networks continues to grow exponentially (per one estimate, over 200X every two years), the upcoming economic "wall" for AI is becoming very clear. Pioneers across the CPU, GPU, xPU, and system integration worlds are intensifying investments in architectures that offer energy-consumption and total-cost-of-ownership benefits to expand the era of AI.

In 2024 we will continue to see the debut of new AI processing architectures, including much-anticipated neuromorphic computing devices, optical, and even quantum computers that hold the promise of pushing the AI economic wall further out in the future. The integration of these heterogeneous compute elements will amplify the industry's move toward multi-die systems and complement the continued strong push to angstrom-scale digital CMOS devices.

While these compute devices will target the data center, the autonomous edge will continue to integrate very significant compute power into autonomous vehicles, industrial applications, and personal-computing solutions, where the innovation focus will be placed on sensor integration.

Responsible Use of AI Will Open New Paths

GenAI operationalizes data in a way that we've never seen before. In fact, it's made possible in large part due to the proliferation of data in our digital world. Several thought leaders in the AI space have warned that GenAI innovation may be outpacing our ability to make sense of its implications and respond to its consequences.

In 2024, data privacy and governance will be explored deeply. There already are many discussions in industry and government around AI's impact on sustainability, society, and business, as well as its ethical considerations. For example, the recently published NIST AI Risk Management Framework offers an important tool for AI product development. As AI becomes increasingly ingrained in our world, the emergence of regulatory frameworks to provide guidelines and protections would not be surprising.

Responsible use may sound like a restriction; however, it is instead an essential precondition. In my keynote at the Silicon Valley Leadership Group Retreat 2023, I explored the path forward for elevating the role of AI in chip design from today's "tell-ask," low-trust interactions to potential future "share-discover," or co-creator, problem-solving networks. To reach its full potential, and earn its place among human creators, AI will need to demonstrate fairness, reliability, privacy, inclusiveness, transparency, accountability—in other words, all elements of responsible use.

New Year, New Innovations

Innovation is an inherently human trait, yet AI is already demonstrating how it can augment and accelerate our abilities. In time, it may even win our trust. In 2024, we should think about our interaction with AI and how we can collaborate with it. While we humans continue to drive unique ideas, AI can act as our co-creator, helping us accomplish more than ever before. The future is closer than we think. ed

JAMES MORRA | Senior Editor

Thermal-Monitoring "Tape" Helps EV Batteries Beat the Heat

A unique peel-and-stick thermal sensor works to keep the high-voltage battery packs in EVs from overheating—and potentially from exploding into flames.

LITHIUM-ION (Li-ion) BATTERIES have their pros and cons. But their unique combination of high energy density, moderate self-discharge rate, and endurance to handle many thousands of charge-discharge cycles before reaching the end of their lifecycle has turned the battery chemistry into the gold standard of the electric-vehicle (EV) era.

As it turns out, Li-ion batteries are only getting better over time. However, despite the power-handling capabilities of these batteries, they have deficiencies that must be carefully managed.

One of the risks inherent to Li-ion batteries is their tendency to be highly combustible. If the cells comprising the high-voltage battery pack inside the EV are damaged or incorrectly charged or drained, they can overheat uncontrollably or even erupt in flames. Given that explosive potential, it's important to have accurate thermal monitoring of the battery cells at any instant in their lifecycles to ensure safe operating conditions.

To help enhance the thermal management of high-voltage EV battery packs, Littelfuse launched a new thermal-monitoring device called TTape. The thin, flexible tape is embedded with a series of tiny, printed temperature sensors that can be wrapped around several Li-ion cells at a time. By tracking the local temperature of the cells instead of monitoring the entire pack, the tape can detect overheating before it brings about a thermal runaway, offering superior safety and battery life enhancement.

144424498© Surasak Petchang | Dreamstime.com

With a response time of less than a second, the tape can alert the battery-management system (BMS) to abnormal levels of heat building up in the cells faster than existing solutions, such as thermistors, said Littelfuse.

The tape offers more flexibility in terms of its placement in the battery pack. It's possible to use a single strip of temperature-monitoring tape to monitor several cells in the pack at the same time, increasing the "spatial resolution" of overtemperature monitoring. This feature also simplifies the cell installation process. The tape is backed by adhesive so that you can peel and stick it directly on the Li-ion battery cells in the EV.

According to Littelfuse, the thermal sensing tape is suited not only for Li-ion batteries in consumer and commercial EVs, but also large-scale commercial batteries that act as energy storage for the grid.

Battery-Management Systems Keep the EV Cool

Every EV is packed with as many battery cells as possible—physically and economically—to increase the storage capacity of the packs housing them, which can account for approximately 30% to 40% of the EV's total production cost.

Since the capacity of Li-ion batteries varies over time and usage—and even more so when operating under harsh conditions out on the road—every cell must be carefully managed and regulated to perform at its best and protect the battery cells from permanent damage.

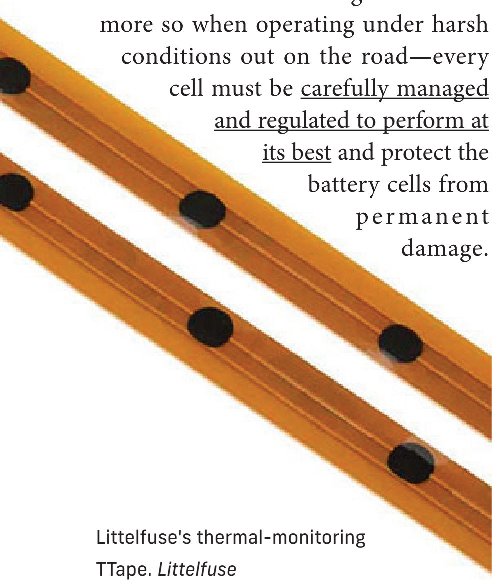

Littelfuse's thermal-monitoring TTape. *Littelfuse*

If the cells are damaged in any way, the EV battery can suffer from reduced driving range and a shorter useful lifespan, or it can raise the risks of explosion.

The cells are typically wired to a BMS that plays the pivotal role of monitoring the voltage, current, and other facets of each cell over its lifetime. It's also engineered to adapt to the constantly shifting conditions that the EV battery cells are exposed to and, in that way, to maximize the state of charge (SOC) of the cells at any point in time and over their lifetime, also called state of health (SOH).

The primary role of the BMS lies in balancing the level of charge stored in each battery cell, as overcharging or undercharging the battery cells may also cause physical stress, leading to premature cell aging. Cell balancing helps prevent cell degradation. The system's other role lies in correcting the temperature fluctuations inside and outside the cells comprising the battery pack, which in some cases can imply the onset of a thermal runaway.

Exposure to severe heat or cold can also lead to large variances in a battery's performance and lifespan that must be accounted for. Frigid temperatures, for instance, can inhibit chemical reactions in the battery cells, sapping their usable power and slowing down the charging process.

While there are many ways to keep tabs on the temperature of the battery's cells, one of the most widely used is a type of thermal sensor called the NTC thermistor. A thermistor is a type of temperature sensor that exhibits a large change in resistance proportional to a change in temperature. Thus, by measuring the resistance, you can figure out the amount of heat afflicting the system.

However, these thermal sensors may not be the best fit for the Li-ion cells in EV batteries. Specifically, thermistors lack the rapid response times and accuracy necessary for safe battery management, said Littelfuse. They also tend to take a broader view of the temperatures affecting the battery pack. Placing these bead-, disk-, or cylinder-shaped devices inside the battery pack in such a way to perform "localized" cell monitoring is a major challenge, according to the company.

Distributed Temperature Monitoring at the Cell Level

However, you can never be too careful when it comes to protecting a Li-ion battery from harm. The tape is intended to supplement the thermistors and other sensors in EVs instead of replacing them. Since the tape has a faster response time, the BMS may start cooling the EV battery cells at an earlier stage and help stop thermal runaways before they get out of control, negatively impacting the battery's longevity.

Because no calibration or temperature lookup tables are needed, the tape can be easily integrated into existing battery-management units along with NTC thermistors, delivering enhanced over-temperature detection at the cell level.

In the event it detects overheating in the battery cells, the tape wakes up the MCU at the heart of the BMS from sleep mode over a single input. The device uses a dual-wire interface to sense if any of the cells in the battery runs over a critical temperature limit. The tape measures 10 mm wide and stretches up to 1 m. The product comes in tape-and-reel or box packaging, depending on the length required.

The thin form factor of the tape is also ideal for "conformal installations," meaning it can bend around sharp corners, rounded edges, or pouch-shaped bellies of the Li-ion cells in battery packs. As industry insiders point out, EV batteries are becoming huge clusters of high-voltage energy storage. The 400/800-V batteries being plugged into the latest EVs typically comprise 200 cells connected in a series.

Since the tape is automotive grade, it's rigorously "qualified" against a long list of industry standards to ensure durability and reliability for cars. It can tolerate operating temperatures of −40 to 85°C.

Design Feature

GEIR KJOSAVIK | Product Director – PMIC, Nordic Semiconductor

My Application Needs Accurate Fuel Gauging—

WHAT OPTIONS ARE AVAILABLE?

spaxiax_1300457374 | iStock / Getty Images Plus | fujiwara_1842465761 | iStock / Getty Images Plus

Inexpensive methods to measure battery energy can result in low precision measurements, and techniques to increase precision often add complexity and cost. One alternative uses a software algorithm that provides reliable battery-capacity measurement.

USERS OF BATTERY-POWERED products demand accurate cell fuel gauging so that they can conveniently plan recharges. But there's a second, less obvious, reason for precise battery energy measurement: Inaccuracy leads to bad user experiences such as a product suddenly failing while still indicating to the user there's some runtime left.

To avoid this situation, vendors often add safety margins that force the user to charge the product sooner than needed. That's not only inconvenient, but it also creates the impression that the product has a shorter battery life when it's not the case. And while charging a consumer product like a wearable more often than required is one thing, in an industrial context, where there might be hundreds or even thousands of battery-powered devices, frequent unnecessary charging is quite another.

Until now, developers were faced with a dilemma: Should they contain costs by using an inexpensive fuel-gauge method and risk consumer disappointment, or add design complexity, PCB real estate, and expense with a more accurate fuel gauge? Now, however, there's a third option—a software-based solution that promises high precision without the added complexity and cost of dedicated hardware solutions.

The Inexpensive Way: A Voltage-Based Lookup Table

There are two established techniques for estimating how much energy is left in a rechargeable battery, and each comes with pros and cons.

Let's first consider the inexpensive approach—the voltage-based lookup table. This method is simple and low cost to implement, which is why most designers tend to use it.

It's a case of measuring the battery's terminal voltage and software and then referring to a lookup table derived from the cell's discharge curve *(Fig. 1)*. The measured voltage is mapped to a battery state-of-charge (SOC) figure and reported to the user. While this sounds good in principle, when you look closer, several inherent problems pop up.

The first challenge comes with varying battery load. Since the battery is an integral part of the circuit, it's impossible to measure the open-circuit voltage of the battery's terminals. This means that there will always be some voltage drop across the cell's internal resistance depending on how much power is delivered by the battery. In other words, the measured voltage will be lower than the actual voltage.

For an application such as a Bluetooth speaker, which will exhibit great variations in power consumption, the calculated SOC could be much lower than reality. Worse yet, the SOC error can't be estimated since there's no way of knowing the battery load when the sample was taken.

Second, SOC measurements are influenced by temperature. The internal resistance of a rechargeable lithium battery varies greatly with temperature.

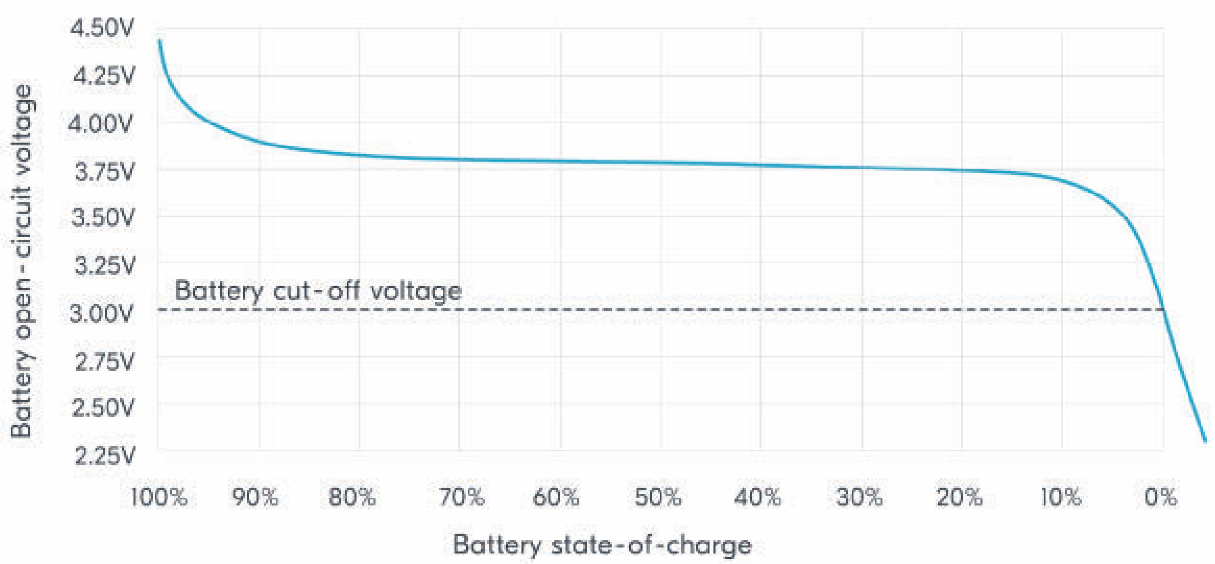

1. Shown is a typical Li-ion battery discharge curve. Note how flat the curve is between 80% and 20% state-of-charge (SOC); this makes it very difficult to ascertain true SOC. *Images courtesy Nordic Semiconductor*

That means temperature fluctuations amplify the SOC inaccuracies imposed by battery load.

The third challenge is that the middle portion of the discharge curve is almost flat. For the battery *(Fig. 1, again)*, the battery voltage drops only 80 mV or 2% from an 80% to 20% SOC. This means that even without considering the effects mentioned above, it's very difficult to determine the SOC across this part of the discharge curve. If we do try to take battery load and temperature effects into account, it becomes almost impossible.

The Accurate Way: A Dedicated Fuel-Gauge Chip

Until now, the dedicated fuel-gauge chip has been the only way to achieve precise battery SOC for embedded applications. Many of these dedicated fuel-gauge chips employ a so-called coulomb counter. It's a circuit that accurately measures the battery current and then pairs this information with battery voltage measurements to keep track of how much energy is drawn or added to the battery. SOC precision is significantly enhanced compared to the voltage-based lookup table, but there are some tradeoffs.

Until now, the dedicated fuel-gauge chip has been the only way to achieve precise battery SOC for embedded applications. Many of these dedicated fuel-gauge chips employ a so-called coulomb counter.

First, the chip adds cost to the bill of materials (BOM) and consumes PCB area; it also requires the support of several passive components. Second, the dedicated chip increases the product's power consumption. The chip must be active when current is drawn from the battery, which happens even when the system is in sleep mode.

While the measurement interval can be extended during sleep, even in the lowest power modes, a dedicated fuel-gauge chip can increase system sleep current by as much as an order of magnitude in an otherwise power-frugal Bluetooth application.

Third, to transform the battery current and voltage measurements into an accurate SOC estimate, the system needs accurate information on the battery characteristics during charging and discharging. This "battery profiling" is typically done as a service by the fuel-gauge chip vendor and requires the designers to ship some samples of the battery to the vendor and wait several weeks for the results.

The drawbacks of both these common solutions for SOC measurement can make it hard for designers to choose one solution over the other. But now there's a third option that addresses the disadvantages of both of those conventional approaches.

A New Approach: SOC Precision Without Adding Cost or Power Consumption

The company I work for, Nordic Semiconductor, recently announced the nPM1300 power-management IC (PMIC). The PMIC optimizes the power

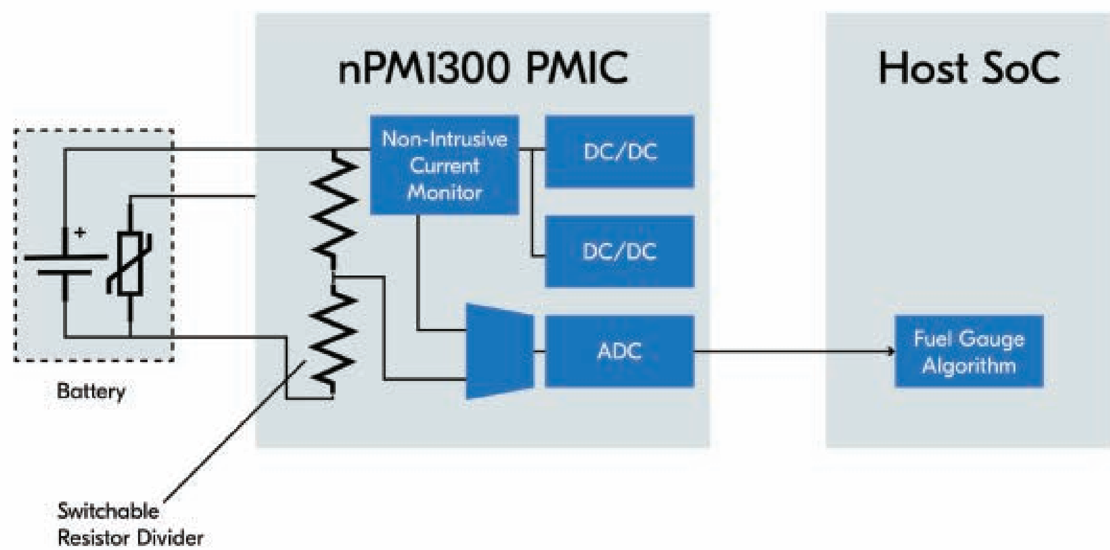

2. The nPM1300 PMIC offers hardware hooks that enable the host processor to measure vital battery parameters. The host processor then uses a free software algorithm to precisely calculate battery SOC. There's no additional hardware cost or impact on power consumption.

consumption of the company's Bluetooth LE systems-on-chips to maximize battery life.

The PMIC doesn't directly measure battery SOC. It's not equipped with a coulomb counter or other complex hardware that would add to the cost or power consumption of the product. Rather, the PMIC has hardware hooks that enable a host processor to measure three vital battery parameters: voltage, current, and temperature. Nordic Semiconductor then provides a free-of-charge fuel-gauge algorithm as a part of the software driver library for the nPM1300. The cost to the developer is a few kilobytes of flash memory to store the algorithm code.

The code does slightly add to the product's power consumption. However, while the algorithm runs once per second in active mode, it doesn't need to run during system sleep mode. Therefore, the overall increase in power consumption is minimal compared to dedicated fuel-gauge chips. When the system wakes up, the algorithm can determine how much net energy was lost or added while the system was in sleep mode *(Fig. 2).*

The Benefits of a Mathematical Battery Model

How does the Nordic's method compare to the traditional alternatives? To answer that question, we ran some benchmark tests comparing its new technique against the voltage lookup table approach, and then against a Keysight N6705B DC Power Analyzer set up to work as a calibrated coulomb-counter reference. The load pattern used throughout the test was a train of current pulses alternating between 55 and 80 mA with about 10 mA of current consumption in between those pulses.

In the upper diagram of *Figure 3,* the voltage lookup table and the nPM1300 PMIC methods are plotted against the coulomb-counter reference. In the middle diagram, the reference is subtracted, hence showing the absolute error of the voltage lookup table and the nPM1300 solution, respectively. The voltage lookup table yields errors of up to 20% with a large spread that represents the variations introduced by the series internal resistance of the battery and the variations in current draw.

In comparison, the nPM1300 PMIC and algorithm show a maximum deviation from the reference of approximately 3%, which is comparable to what's achieved by a decent dedicated fuel-gauge chip *(Fig. 3, again).*

Temperature fluctuations amplify the weaknesses of the voltage lookup table method because of the impacts on battery internal resistance already discussed *(Fig. 4).* At −5°C, the voltage lookup table exhibits between 10% and 30% error depending on current draw, whereas the error of the nPM1300 PMIC plus algorithm is at worst 5% and typically under 2%. This is again comparable to that achieved by dedicated fuel-gauge chips.

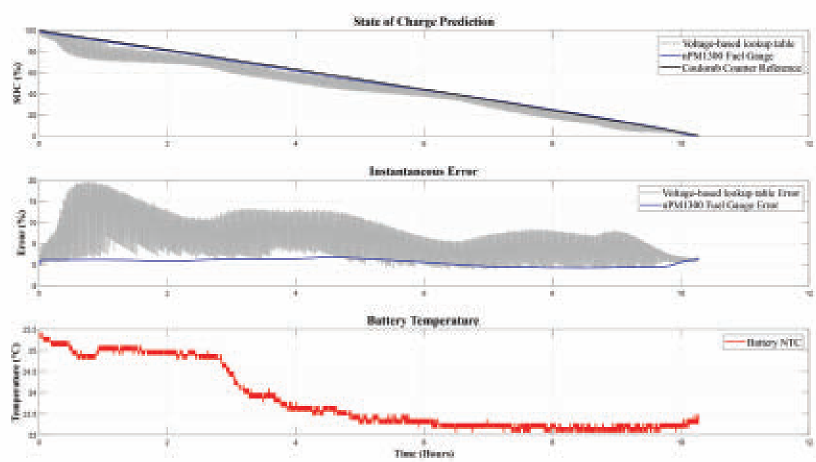

3. The top graph illustrates the error between a coulomb-counter reference, a voltage lookup table, and Nordic's nPM1300 and algorithm. The middle graph compares the instantaneous error between the reference and the alternative SOC measurements. The final graph shows the battery temperature over time.

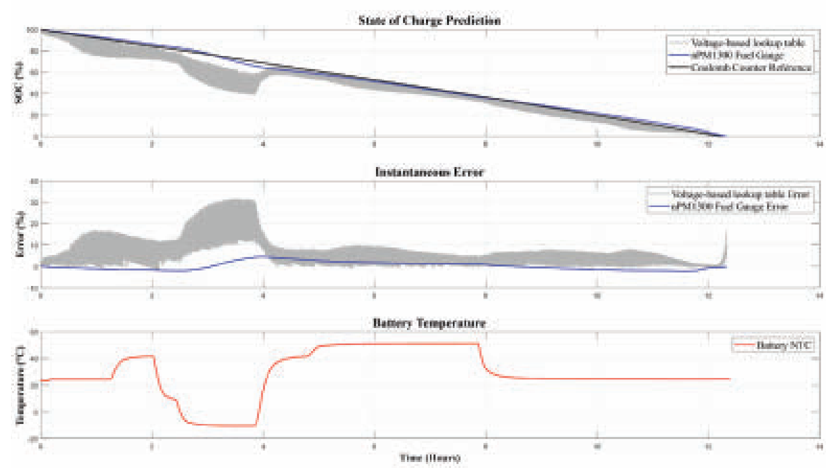

4. The top graph illustrates the error between a coulomb-counter reference, a voltage lookup table, and Nordic's nPM1300 and algorithm with wide temperature variations. The middle graph again compares the instantaneous error between the reference and the alternative SOC measurements with the final graph showing battery temperature over time.

DIY Battery Profiling

The nPM1300 PMIC and algorithm, in common with dedicated fuel-gauge chips, requires battery profiling. However, unlike other suppliers, Nordic provides designers with the ability to do this in house, eliminating the time and expense of sending the batteries to an external agency for profiling.

The company's nPM PowerUP PC app, which is part of its nRF Connect for Desktop development tools suite, is a GUI-based tool for evaluation and development of the company's PMICs. It allows for designers to evaluate and configure the products and export configuration code. Within the configuration scope, nPM Power UP also supports battery profiling, enabling a developer to build an accurate model of the battery for the fuel-gauge algorithm in-house.

The battery to be profiled is plugged into an nPM1300 evaluation kit and in turn is connected to Nordic's dedicated Fuel Gauge board that enables a constant current draw from the battery. Then it's just a case of clicking on the "Profile Battery" in the nPM PowerUP tool *(Fig. 5)*. This instruction initiates a slow discharge-rest-discharge-rest sequence customized to the battery, which is closely monitored to build the battery model.

The procedure should be repeated at three different positive temperatures (for example, +5, +20, and +40°C) to ensure accurate fuel gauging across the product's likely operational temperature range. Once complete, the fuel-gauge algorithm contains a mathematical model of the battery that will yield the typical accuracies shown in the previous section. The battery profile can then be exported from the tool into the nPM1300 PMIC's configuration file. Nordic's application note "Using the nPM1300 Fuel Gauge" details how to do this operation.

Targeting the Best Fuel-Gauge Solution

Selecting a fuel-gauge solution that satisfies end customer demand for accuracy is a difficult task since it adds cost and complexity to the product and increases power consumption. Hence, designers often go with the cheaper voltage lookup table option and then build in a safety margin, forcing consumers to charge more frequently than needed.

The SOC measurement method eliminates the drawbacks of the two most common fuel-gauge solutions and makes the choice easier for designers. ed

Selecting a fuel-gauge solution that satisfies end customer demand for accuracy is a difficult task since it adds cost and complexity to the product and increases power consumption.

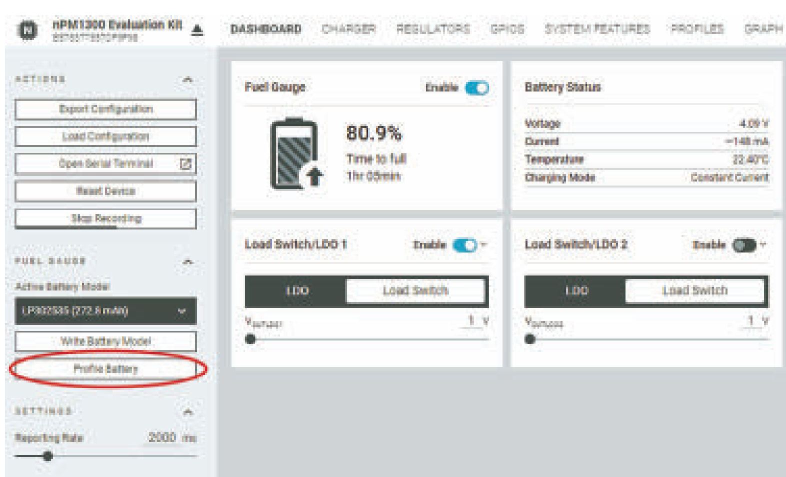

5. The nPM PowerUP tool supports battery profiling, enabling a developer to build an accurate model of the battery for the fuel-gauge algorithm in-house. The "Profile Battery" instruction initiates a slow charge/discharge cycle, which is closely monitored, to build the battery model.

New

Automotive series products
Click for more information

Trusted manufacturer by leading industry leaders
for high quality Inductors and transformers

Possessing 45 years of professional experience.

- High quality product with customization available
- Products are designed for various applications
- Automotive grade products are compliant with AEC-Q200

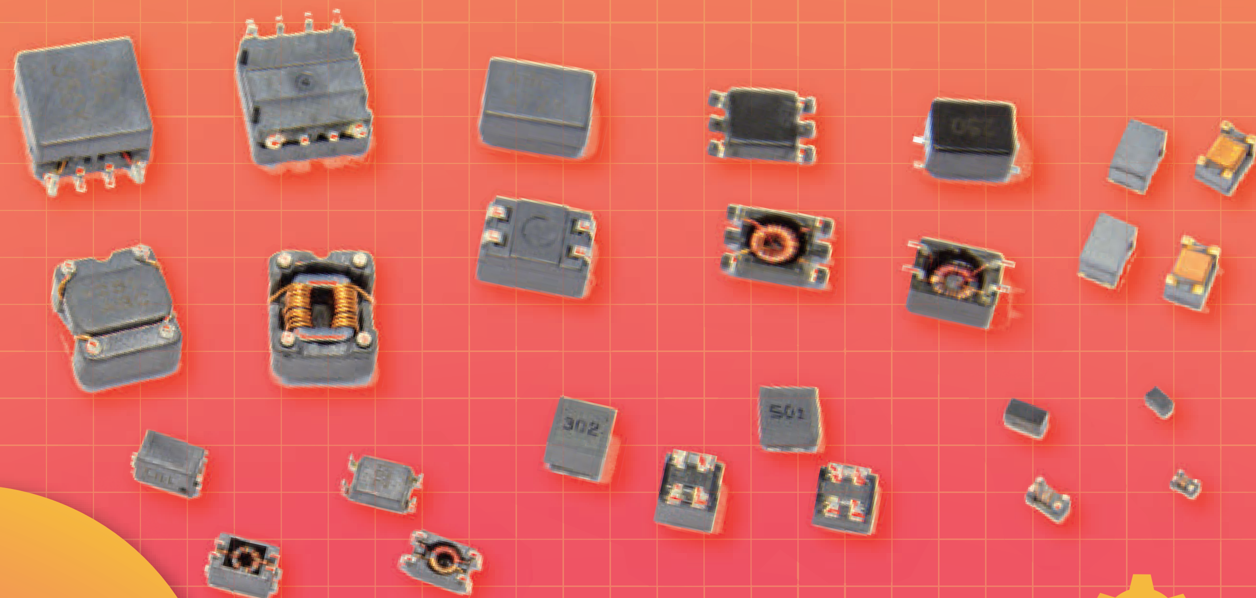

● ISO 9001 ● IECQ QC080000 ● IATF 16949

ABC-ATEC ELECTRONICS GROUP
ABC TAIWAN ELECTRONICS CORP.

Tel:+886-3-478-8188 Fax:+886-3-488-0648
Email: abcsa@atec-group.com www.atec-group.com

ABC
ATEC Electronics Group

FREDERIK DOSTAL | Analog Devices

anigoweb_843002094 | iStock / Getty Images Plus

Dostal's Designs: **Compact UPS** with **Supercap**

Continuous supply voltage via an uninterruptible power supply is crucial for a number of applications, but it can be tough to ensure at all times. This Idea for Design offers a reliable, compact solution that's easy to integrate.

IN MANY APPLICATIONS, it's important for the supply voltage to be continuously available no matter what the circumstances. This isn't always easy to ensure, though. A new concept can provide an optimal solution for an uninterruptible power supply (UPS) with an extremely compact design.

Several applications require an uninterruptible power supply. One example is the RAID system for redundant data storage, which must be protected so that no data is lost in the event of a power failure at an inconvenient time, such as during data backup activity. Systems with real-time clocks also must be supplied continuously with power. This can come from a battery or another backup solution. Other applications include telemetry applications in the automotive sector and systems for administering medications—for example, controlled insulin pumps used in the healthcare sector.

Figure 1 shows a typical industrial application for an uninterruptible power supply. Here, an industrial sensor is supplied with power. The reliability of the system mainly depends on this sensor's power supply.

A linear charge regulator IC is used to charge a supercapacitor when there's available system voltage. If the system voltage drops, the energy from the energy-storage system is raised to the required supply voltage level with a boost regulator.

This system works well, but it's difficult to implement because many different energy converters are needed. Moreover,

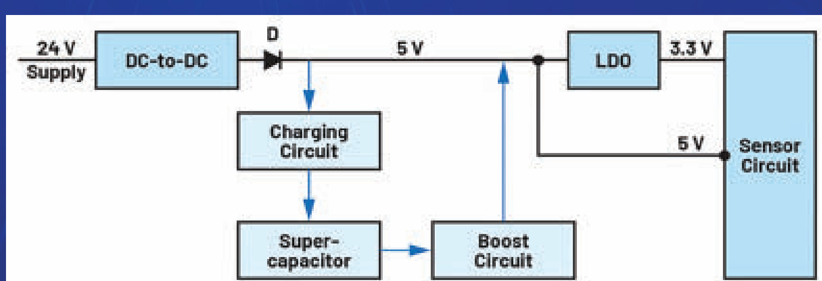

1. Shown is a typical application for an uninterruptible power supply. *Analog Devices*

(Continued on page 37)

Design Feature

BILL SCHWEBER | Contributing Editor

Backscattered, Harvesting **Underwater Acoustic Comms Link** Reaches Kilometers

Using retrodirective techniques and backscatter, low-bit-rate acoustic links are feasible for surprisingly long underwater distances.

IT'S WELL KNOWN that underwater communications is a very difficult proposition. Water absorbs a significant fraction of electromagnetic energy and thus attenuates it severely; acoustic energy is also attenuated (though not as severely), though it's dispersed and distorted.

The low signal levels and low signal-to-noise ratio (SNR) means that establishing a viable underwater link over more than a modest distance—say, around a meter—using a reasonable power source has been a challenge since the earliest days of electronics and electroacoustics.

Underwater Piezo-Acoustic Comms System

Despite these difficulties, researchers at MIT devised a scheme that supports ultra-low-power networking and links at kilometer-scale distances, albeit at modest data rates on the order of 500 bits/s. Applications include monitoring water-based sensors and buoys, aquaculture, coastal storms, and even marine-animal migration patterns.

Unlike many underwater acoustic communication systems, this piezo-acoustic design communicates information via harvesting of backscatter, which is a reflection of external acoustic signals, rather than generating its own carrier. Their design enables backscatter sensor nodes to sense and communicate at five to six orders of magnitude lower power than present-day underwater modems, even at similar data rates.

The system uses battery-free piezoelectric-based transducers, which are the standard acoustic source in underwater settings. They then adapted a 70-year-old radio arrangement called a Van Atta array, where symmetric pairs of antennas are connected so that the array reflects energy back in the direction from which it came in a variation of phased-array beamforming *(Fig. 1)*.

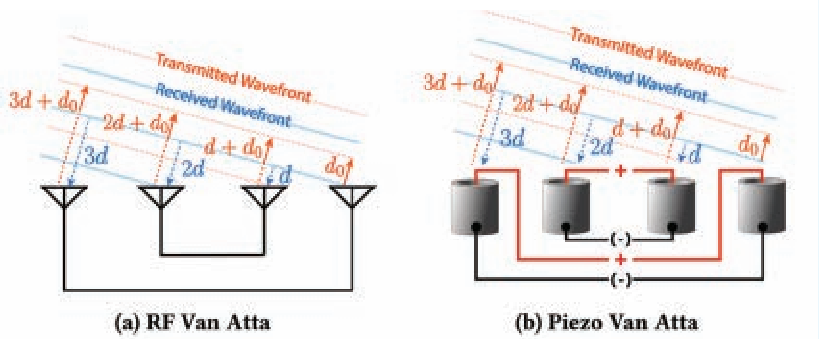

1. Van Atta structures: (a) By connecting antennas symmetrically, the received signal is retransmitted in the same direction. (b) They realized the same concept via differential connections with piezo-acoustic nodes. *Images courtesy MIT*

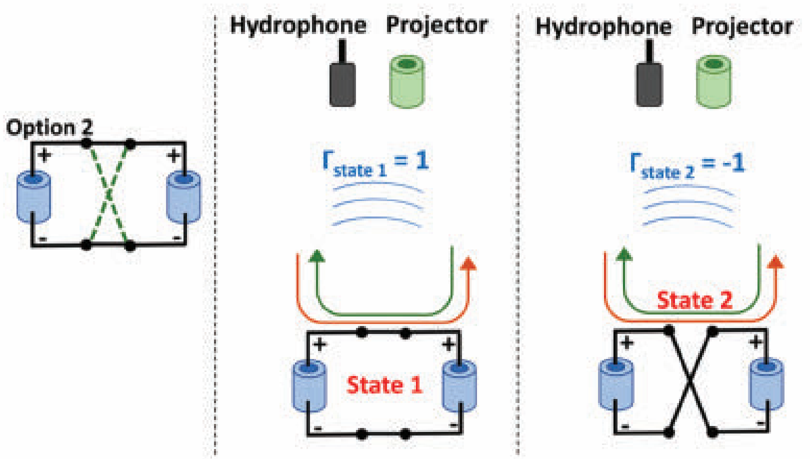

2. Cross-polarity switching: When switches connect like-polarities, the hydrophone sees a reflection coefficient $\Gamma_{state1} = 1$; when switches connect cross-polarities, the hydrophone sees a $\Gamma_{state2} = -1$. This yields the highest modulation factor of 4, resulting in a highly detectable backscattered response.

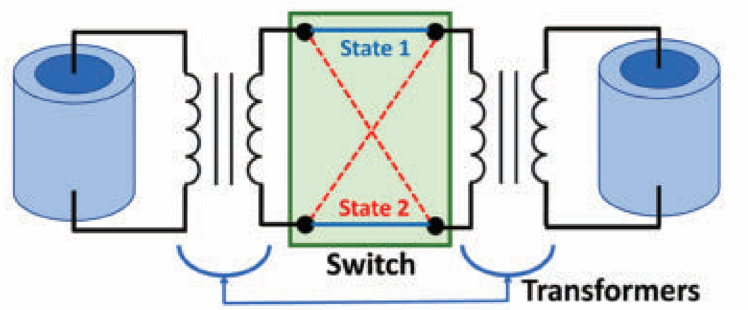

3. Van Atta unit cell: A unit cell consists of two piezoelectric transducers, two transformers, and a cross-polarity switch.

This article will focus on the transducers, but briefly look at the link arrangement itself as well.

The Van Atta Cell

An RF Van Atta structure consists of an array of antennas interconnected in symmetrical pairs. This concept was the starting point for the multi-element array, but with major differences due to the nature of the piezo transducer and interconnection. They began with a switching technique known as pass/absorb, where the switches are simultaneously controlled to alternate between these two states to achieve retroreflective backscatter.

However, applying the same idea in underwater backscatter isn't that simple. Although it's in some ways similar to an RF phased array and associated beamforming, there are significant differences. First, unlike antennas that typically have one main external feeding point (meaning they are single-ended), piezoelectric transducers are differential. Therefore, they needed a new switch architecture to alternate between different reflective/absorptive states and optimize the backscatter SNR.

Second, in addition to the need for a novel switching architecture, the piezoelectric transducers must be matched to maximize the retroreflective SNR. Unlike antennas that can be easily designed or purchased to have 50- or 75-Ω impedance, it's very difficult to design a piezoelectric transducer to have a real and consistent impedance. A piezoelectric device is capacitive by nature and even with tight mechanical tolerances, the fabrication of a sealed transducer results in significant process variations in the measured impedance.

Instead, they used a second switching technique via an arrangement which maintains that the connection between the two nodes be "ON" at all times, but toggles it between in-phase and counter-phase polarities *(Fig. 2)*.

There's also the issue of impedance matching, as the transducer is highly capacitive and has significant variability. Thus, they needed a mechanism to ensure matching for proper Van Atta operation. To address this problem, they integrated a transformer-based matching network between each node in every pair *(Fig. 3)*.

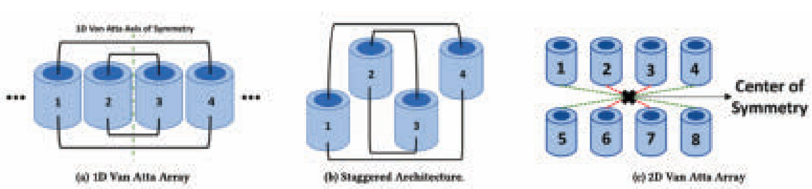

4. Multi-node Van Atta architectures: (a) 4-element Van Atta where pairs are connected around a symmetry axis. (b) The architecture creates a difference in elevation between the elements in a pair of the Van Atta. (c) A 4×2 architecture, where elements are connected around a center of symmetry.

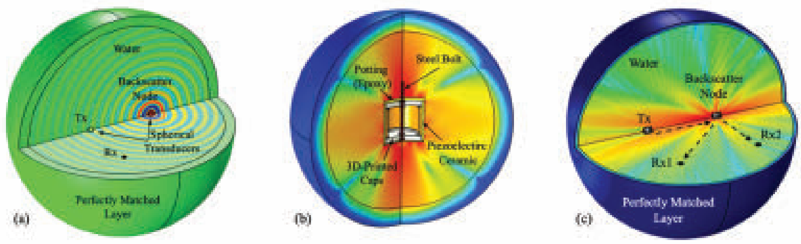

5. Numerical simulations: (a) The backscattered pressure from a spherical transducer. Positive pressure is in red and negative pressure in blue. (b) A model for characterizing cylindrical transducers. The "heatmap" shows the sound level generated by the transducer at 30 kHz (red indicates higher sound level). (c) The backscatter level from a cylindrical transducer at 40 kHz. Black arrows highlight acoustic propagation from transmitter Tx to two receiver locations, Rx1 and Rx2.

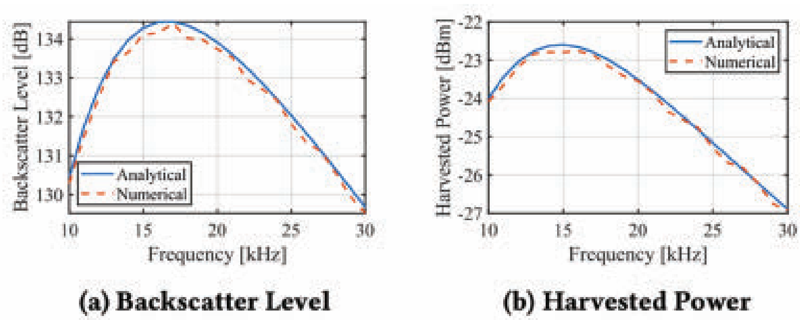

6. Numerical validation for spherical transducers shows the analytical (blue) and numerical (red) (a) backscatter levels and (b) harvested power for two spherical transducers separated by 0.5 m, as observed by a hydrophone 0.5 m away from the backscatter node.

Having a single Van Atta cell is critical, but it's only a first step. They had to go from that signal node to a one-dimensional array and then a two-dimensional array. However, doing so brings problems of occlusion and self-blocking.

To achieve retrodirectivity, the elements are connected around an axis of symmetry that's implemented choosing the inner transducers (nodes 2 and 3) as one pair, and the outer ones (nodes 1 and 4) as the other pair *(Fig. 4)*. Connecting each pair and switching them simultaneously enables retrodirectivity.

Modeling and More

The team developed detailed models of the downlink channel and signal path, as well as the uplink budget, and related this to the transducer and channel parameters. As an example of their attention to detail and fidelity, they first applied the models to spherical-shell piezoelectric transducers. These are the simplest form of underwater transducers, and the only type of underwater transducers with closed-form expressions for all of their acoustic parameters.

However, spherical transducers are omnidirectional regardless of their operating frequency, so this analysis only validated a special case of their analytical expression *(Figs. 5 and 6)*.

Therefore, they extended their analysis to potted cylindrical transducers with a directional radiation pattern that depends on frequency, but doesn't have closed-form analytical expressions for most of its acoustic parameters at practical backscatter frequencies. To overcome this dilemma, they developed a semi-analytical approach to predict the backscatter performance of practical transducers.

Their overall implementation had three major components: the piezo-acoustic backscatter nodes, fabricated in-house; a transmit-side audio amplifier, also fabricated in-house, which was connected to software radio to generate the downlink signal; and a receiver consisting of an omnidirectional hydrophone with a differential receive-voltage sensitivity of −180 dB (re 1 V/μPa).

The output of the hydrophone was fed to a four-channel oscilloscope controlled through a LabVIEW code. It recorded the experimental data related to input electrical power as well as the voltage at the backscatter node during downlink power-harvesting verification experiments.

Test and Results of the Underwater Comms System

They evaluated their design in a river (4-meter depth) and the Atlantic Ocean (30-meter depth)—most experiments were performed in the river—conducting over 1,500 trials. In the first iteration of the array, they found that if the connected nodes were too close, they would lock out each other's signals. To overcome this, they devised an improved arrangement where the nodes were staggered *(Fig. 7)*.

These evaluation environments had standard sources of underwater noise and

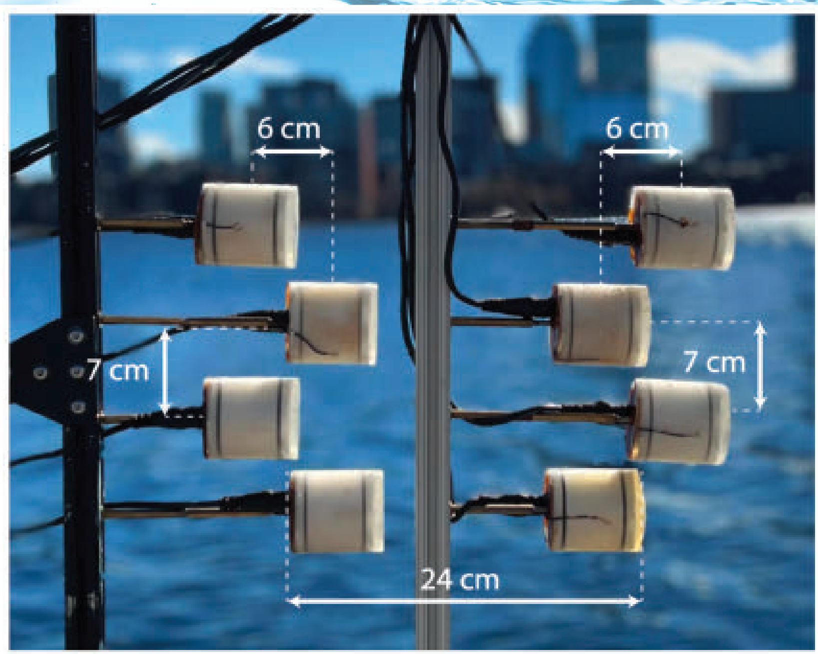

7. The final 4×2 Van Atta array has modest dimensions and a staggered arrangement to prevent self-blocking.

interference, including marine animals, motorboats, sailboats, and ships, as well as other underwater acoustic communication systems. They varied the location, depth, and orientation of the transmitter, receiver, and backscatter nodes.

Experiments were performed across different weather conditions, including wind, rain, and snow. They even built a custom rotating device using a stepper motor and belt drive to rotate a shaft to which the backscatter device was attached, in order to vary the angle of signal path.

Among the many factors they characterized were link gain (actually, loss); backscatter level from 10 to 18 kHz in steps of 100 Hz (15 kHz provided optimum results); harvested power versus frequency; SNR versus frequency; impact of directivity and spreading loss; and, of course, bit error rate (BER) *(Fig. 8)*.

They compiled a large set of test results with variations in operating parameters and scenarios, using a 4×2 staggered Van Atta array vs. range at 500 bits/s. Applying SNR versus distance for input power of 1 W (blue), 150 W (red), and 2 kW (yellow), they found that the SNR drops below the decodable threshold (6 dB) for 1 W of input power around 20 meters. Upping the power to 150 W increases the decodable range to roughly 60 meters, and it further extends beyond 100 meters with 2 kW of input electrical power *(Fig. 9)*.

As you would expect from a multidisciplinary MIT project, the associated documentation is both thorough and readable. It consists of two detailed papers, one focused primarily on the Van Atta array and transducers ("Enabling Long-Range Underwater Backscatter via Van Atta Acoustic Networks") and the other of system design ("The Underwater Backscatter Channel: Theory, Link Budget, and Experimental Validation").

Together, the two documents provide insight and discussion of the underlying problems, equation-laden physics and analysis, review of attempts at solutions and course chosen, detail of modeling of all aspects of the system, description of the final arrangement, references, and test results. ed

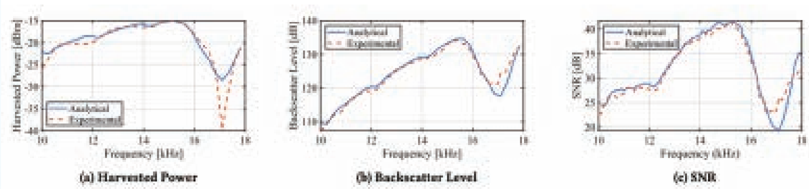

8. Experimental verification: (a) plots the harvested power vs. frequency on the downlink for the analytical model (solid blue) and experimental measurements (dotted red). (b) plots the backscatter level as a function of frequency for the analytical model (solid blue) and experimental measurements (dotted red). (c) plots the SNR against the frequency for the analytical model (solid blue) and experimental measurements (dotted red).

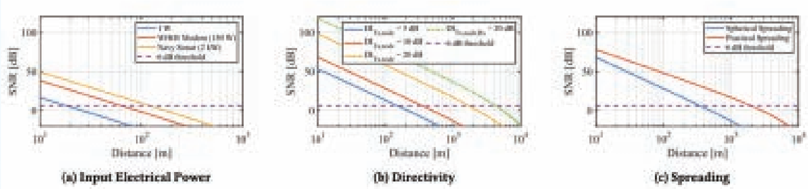

9. Theoretical uplink communication range: (a) plots SNR vs. distance for the in-house transducers while varying the input electrical power between 1 W (blue), 150 W (red), and 2 kW (yellow). (b) plots the SNR vs. distance for a transmit power of 150 W; a Tx/node directivity index of 5 dB (blue), 10 dB (red), and 20 dB (yellow); and a Tx/node/Rx directivity of 20 dB (green). (c) plots SNR vs. distance for two different spreading factors. The red curve plots the practical spreading (k = 1.5) and the blue curve plots the spherical spreading (k = 1.5). The dashed line shows the 6-dB decoding threshold.

JON PEDDIE | President, Jon Peddie Research

NVIDIA RTX 3080: Real-Time Ray Tracing on a Chip

NVIDIA's RTX 3080 ushered in the era of real-time ray tracing for consumers everywhere.

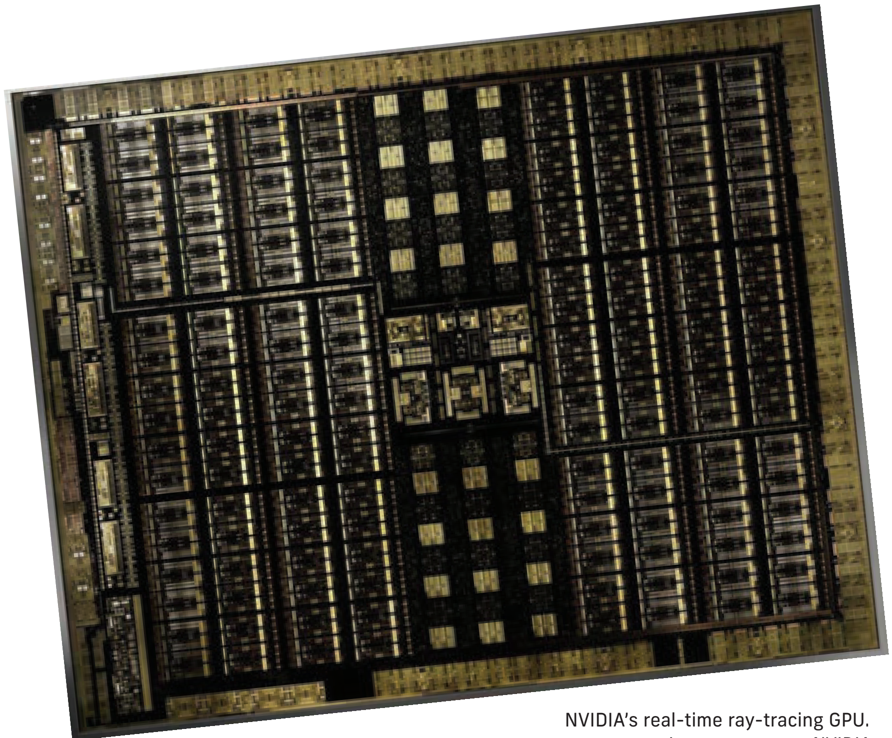

NVIDIA's real-time ray-tracing GPU. *Images courtesy NVIDIA*

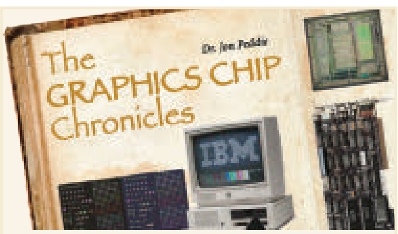

Jon Peddie's long-running "Graphics Chip Chronicles" series highlights the graphics chips that have made a major impact over the years. It presents a fun look back at how we got to where GPGPUs are the norm.

IN EARLY 2018, NVIDIA showcased hardware-accelerated real-time ray tracing using its DGX system and the DXR API. The $50,000 DGX contained four NVLink 4-Way interconnected NVIDIA Tesla V100 add-in boards (AIBs). Each AIB featured a 5,120-shader Volta GPU running at 1.46 GHz with 16 GB of HMB2 local memory and was capable of 7 TFLOPS double-precision, for a total of 40,960 shaders and 128 GB of memory.

The cluster of AI accelerators was paired with a 2.2-Hz, 20-core Intel Xeon E5-2698 processor equipped with 256-GB DDR4 that could produce 800 GFLOPS, which in total pumped out a theoretical compute capability of 28.8 TFLOPS. But given that it was necessary to run the workload on a small supercomputer, I predicted that it would take more than half a decade to bring real-time ray tracing inside a single GPU.

Upon the introduction of the "Turing" GPU architecture in late 2018, the company signaled it was ready to bring real-time ray-tracing technology to consumers everywhere. NVIDIA surprised everyone, showing off real-time ray tracing running for the first time on a single chip—the RTX Turing TU102 GPU, placed in the GeForce RTX 2080 AIB. The industry had a new acronym—RTRT—for real-time ray tracing. NVIDIA also introduced its Quadro RTX AIBs at Siggraph at the same time.

What is Real-Time Ray Tracing?

Ray tracing is a complex process of generating light rays that illuminate a visible scene combined with the bounces and reflections of those visible rays. Millions, if not billions, of rays are generated, and their luminosity and color must be calculated for every instance of interaction they experience. It's a horrendous calculation workload that's aggravated by resolution because the number of visible pixels goes up as the product of X times Y.

The RTX AIBs based on the Turing GPU featured fixed-function "RT" inference cores designed to accelerate the operations needed to simulate rays, such as bounding volume hierarchy traversal. NVIDIA also introduced DLSS, leveraging AI to improve gaming performance.

Turing GPUs incorporated tensor cores for matrix math processing. Texture-Space Shading allowed for efficient caching and sampling of shading results. The Turing architecture laid the foundation for dedicated ray tracing and tensor engines, full mesh shading capabilities, and advances in real-time rendering.

When it formally introduced the GeForce RTX 20 series AIBs a month after unveiling its Turing GPU, NVIDIA promoted DLSS as a prominent feature.

However, initially, its application was limited to a few video games, such as Battlefield V and Metro Exodus. This was because the algorithm required individual training for each game, and the outcomes were generally not as impressive as basic resolution upscaling.

In the context of geometric objects, a bounding volume hierarchy (BVH) is a tree structure utilized to organize and categorize objects. The objects themselves serve as the nodes or leaves of the tree, and they're enveloped by bounding volumes.

Those nodes are further organized into smaller sets and enclosed within larger bounding volumes. This recursive grouping process continues, with the nested nodes being enclosed within increasingly larger bounding volumes, ultimately forming a tree-like structure where a single bounding volume encompasses the entire hierarchy. BVHs find application in collision detection (inference detection) in ray tracing.

Solving the ray-tracing algorithm was step one. The other, maybe larger, problem was managing the screen resolution.

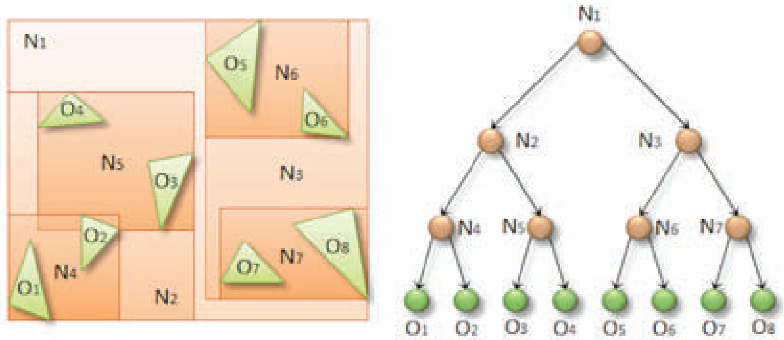

How a bounding tree works in a GPU.

Turing: The First Consumer AI Processor?

In 2017, NVIDIA was enjoying the take up of GPUs in the data center for compute acceleration. At the same time, AI training was getting into full swing, and as fate would have it, AI training was (is) an ideal candidate for parallel processing—the GPU's major strength.

However, AI needed another type of recursive processing that involved matrix math. Matrix can be done on a GPU, but it's not the most efficient engine. Other approaches tried by other firms were digital signal processors (DSPs) and dedicated matrix-math processors constructed in field-programmable gate arrays (FPGAs). DSPs were—and are—used successfully in mobile devices such as smartphones, and Qualcomm is a leader in the area. FPGAs are programmable and functionally effective but not fast.

A new, dedicated matrix-math processor was needed. Google was one of the first to offer such a device, naming it a tensor processor unit, or TPU. In mathematics, a tensor is an algebraic object that describes a multilinear relationship between sets of algebraic objects related to a vector space. The straightforward way to think about it is that a tensor processor does 3D processing in real-time.

The NVIDIA Turing GPU being employed in the RTX 2080 AIB (and Quadro RTX AIB) was equipped with custom math engines—inference engines for ray tracing and tensor engines for matrix math.

Another feature of the new GPU was Deep Learning Super-Sampling (DLSS). NVIDIA claimed the tensor cores in the GPU could be used to improve gaming performance using DLSS. DLSS leverages a deep neural network to extract multidimensional features of the rendered scene and intelligently combine details from multiple frames to construct a high-quality final image. Thus, Turing GPUs can use half the samples for rendering and AI to fill in the information to create the final image.

The result, said NVIDIA, is an image with similar quality as traditional rendering used in many of today's latest games but with higher performance. NVIDIA said DLSS was an easy integration for developers, and at the time, developers announced 25 games would have DLSS support.

Therefore, it's not a stretch to call the Turing GPU the first general-purpose consumer AI processor. On top of the thousands of SIMD processors, video/display/audio processors, raster operation processors, memory managers, and I/O processors, the Turing GPU gained two new processing blocks: the ray-tracing, or RT, and Tensor cores.

The Turing architecture introduced important enhancements to the core execution data paths. Previous shader workloads consisted of a combination of floating-point (FP) arithmetic instructions, such as floating add (FADD) or floating multiply (FMAD), along with simpler instructions like integer adds for addressing and data retrieval, floating-point compares, or min/max for result processing, and more.

In prior shader architectures, the floating-point math data path remains inactive whenever one of those non-FP-math instructions is executed. Turing addressed that limitation by incorporating a second parallel execution unit alongside each CUDA core.

Within the new architecture, the primary enablers for Turing's improvement in graphics performance were a new GPU processor (streaming multiprocessor, or SM) architecture with improved shader execution efficiency and a new memory system architecture that includes support for the latest GDDR6 memory technology. As NVIDIA explained, the Turing SM is partitioned into four processing blocks, each with 16 FP32 cores, 16 INT32 cores, two Tensor cores, one warp scheduler, and one dispatch unit.

In parallel with Turing's development, Microsoft announced both the DirectML for AI and DirectX Raytracing (DXR) APIs in early 2018.

The Turing GPU was full of superlative numbers, equipped with 18.6 billion transistors (fabricated on TSMC's 12-nm FinFET process node) that fit into a silicon area of 754 mm^2—a new high on both fronts. In terms of performance, it pumped out 14.2 TFLOPS of peak single precision (FP32) performance, and NVIDIA touted several new metrics of GPU performance: 113.8 Tensor TFLOPS and 10 Giga Rays/s. The tradeoff is that it was a heavy user of power, consuming up to 280 W.

Turing packs 4,608 shaders (or as the company calls them, CUDA cores) capable of mesh shading. Mesh shading advanced geometry processing architecture by offering a new shader model for the vertex, tessellation, and geometry shading stages of the graphics pipeline, supporting more flexible and efficient approaches for the computation of geometry.

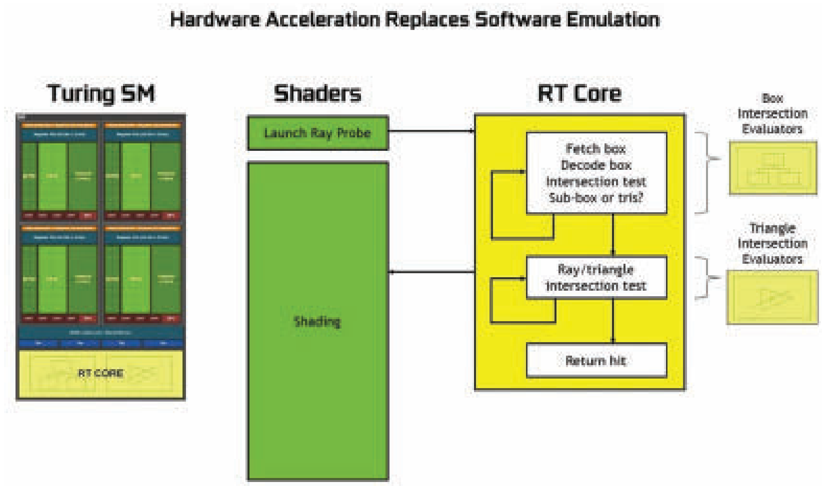

In addition to thousands of SIMD processors, the Turing GPU added several other processing units—the ray-tracing (RT) and Tensor cores.

Shown is NVIDIA's RTX 2080 AIB.

New GPU, New Graphics Capabilities

Another feature introduced with the Turing architecture was variable rate shading (VRS). VRS offers dynamic control over the frequency of shading in graphics rendering. Developers can adjust the shading rate, ranging from once per 16 pixels to eight times per pixel. VRS enables more efficient shading by reducing unnecessary work in areas of the screen, where full-resolution shading would not visibly enhance image quality.

Various categories of VRS-based algorithms have already been identified, including Content Adaptive Shading (which adjusts shading based on content level of detail), Motion Adaptive Shading (which varies shading based on the rate of content motion), and Foveated Rendering (which considers lens resolution and eye position, particularly in virtual reality).

Texture-Space Shading was yet another new development in the Turing architecture. Objects are shaded in a private coordinate space (a texture space) that's saved to memory, and pixel shaders sample from that space rather than evaluate results directly. With the ability to cache shading results in memory and reuse/resample them, developers can eliminate duplicate shading work or use different sampling approaches that improve quality.

Turing was the foundation of a new generation of GPUs with dedicated ray tracing and tensor engines as well as full mesh-shading capabilities. NVIDIA paid for being first, and it was forced to encourage and support game developers to embrace the new rendering approach. Most of them welcomed it, but game development takes time, so major titles didn't show up for a year or more.

AMD held back on implementing RT acceleration. However, when Intel entered the market with the launch of its Xe GPUs in September 2020, AMD offered it and its XLSS scaling capability. The company launched its first hardware RT accelerated GPU, the Radeon RX 6000, a month after Intel. By then, NVIDIA was on its second generation of real-time ray-tracing GPUs and preparing its third.

By mid-2023, NVIDIA had become the undisputed leader in real-time ray tracing. It was shipping DLSS 3.0, which improved the performance of the feature, allowing for 4K RTRT at respectable frame rates. ed

BILL SCHWEBER | Contributing Editor

Single-Channel, High-Voltage Output DACs Include Precision Reference

The basic digital-to-analog converter may not get a lot of attention, but a high-performance device can make a major difference in systems such as test instruments and precision analog-output modules.

IN THE BROAD-BRUSH view of the analog-to-digital converter (ADC) versus its complement, the digital-to-analog converter (DAC), there's no doubt that the DAC has an easier functional challenge.

What's the difference? The ADC must capture an external, unknown, and random signal to translate it into its digital equivalent. In contrast, the DAC starts with a predetermined, internally sourced, noise-free bit pattern and must create the corresponding analog version. As a result, the DAC's role is somewhat easier, and sometimes doesn't get the same consideration and design-in respect received by ADCs.

Nonetheless, creating a high-precision, high-resolution DAC function is a difficult undertaking, especially for higher-voltage output. Issues related to drift, stability, and linearity must be identified and overcome, as well as the need to provide a voltage reference with characteristics that are commensurate with the DAC's own performance objectives.

That's where DAC ICs such as Texas Instruments' 16-bit DAC81401 and 12-bit DAC61401 (DACx1401) excel *(see figure)*. These devices are pin-compatible, single-channel, buffered, high-voltage output DACs that have a 2.5-V internal reference with drift of just 10 ppm/°C. They don't include a lot of extraneous "bells and whistles," but instead are tightly focused on their primary role.

These devices, based on the well-known R-2R DAC topology, are specified to be monotonic and provide an exceptional nonlinearity of less than ±1 LSB (maximum), coupled with ultra-low glitch energy of under one nanovolt-second (nV-sec). Applications include semiconductor test and ATE, lab and field instrumentation, precision analog-output modules, and servo drive-control modules.

The high-voltage DACs offer user-programmable bipolar output voltages of ±20, ±10, and ±5 V, along with full-scale unipolar output voltages of 40, 10, and 5 V. A power-on-reset (POR) circuit powers up the DAC output and keeps the device in power-down mode until the outputs are enabled to avoid random power-up output voltages.

The DAC double-buffered architecture allows for data updates without disturbing or glitching the existing analog output. Data updates are performed asynchronously. In the update mode, a minimum wait time of 2.4 μs is required between updates of the output.

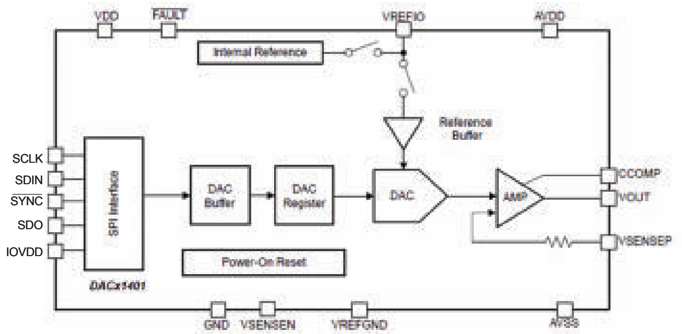

TI's 16-bit DAC81401 and 12-bit DAC61401 digital-to-analog converters are functionally focused on delivering a high-precision and stable analog output. *Texas Instruments*

Communication to the devices is done via a 50-MHz, 4-wire serial SPI-compatible interface supporting 1.7- to 5.5-V operation. If the device is used in a noisy environment, CRC polynomial-based error checking can be invoked by the user to check the integrity of SPI data communication between the device and the host processor.

The DACx1401 devices are fully characterized for operation over the temperature range of –40 to +125°C, offered in a small, 20-pin TSSOP package. The 50-page datasheet provides comprehensive specification and applications insight.

There's also a schematic and fairly detailed discussion of a higher-voltage analog-output circuit based on these DACs, which adds additional active and passive components to reach 0- to 81-V unipolar/±41-V bipolar operation. For those who are serious about this circuit, an even more-detailed discussion is presented in a nine-page application note.[1]

Potential users of even a basic-function IC such as this DAC can benefit from an evaluation unit. The DAC81401EVM is an easy-to-use platform for evaluating the functionality and performance of the DAC81401.

This EVM, with its own 29-page user manual, connects to a Windows-compatible PC and is exercised (setup, operation, evaluation) using a graphical user interface (GUI) application that can be downloaded from TI's website. ed

REFERENCE

1. "High-Voltage Gain Stage Design Circuit for DAC81401," Analog Engineer's Circuit SLAAED7, Texas Instruments.

BILL SCHWEBER | Contributing Editor

Tiny, Power-Miserly **Pushbutton Controller** Also "Saves" the Battery

This focused IC solves the debounce problem and "seals" the battery for freshness despite long shelf-storage periods.

DESPITE THE WIDESPREAD use of touchscreens and their flexible operation, there's still a widespread need for electromechanical pushbutton switches due to their ruggedness, tactile functionality, visual appearance, and standout distinctiveness. Furthermore, interfacing to them should be a "no big deal" issue in a design. However, doing so often has a disproportionate impact on circuitry complexity or the bill of materials (BOM).

Consider the situation where you need to interface with just one or two panel pushbuttons, while conserving battery energy in a unit that may sit on a shelf for months. Using a basic, general-purpose I/O (GPIO) pin or similar may be design overkill or push you into a larger I/O interface, and these solutions or others may also consume extra standby-mode quiescent current.

In such cases, a smaller, tightly focused, application-optimized device may be the right answer. One possibility is the MAX16169 nanoPower Pushbutton On/Off Controller and Battery Freshness Seal from Analog Devices, which could provide a crisp solution *(see figure)*. This IC is an extremely low-power, pushbutton on/off controller with a switch debounce and built-in latch. It accepts a noisy bounced input from a mechanical switch and produces a clean latched output along with a one-shot interrupt output in response to a switch closure exceeding the debounce period.

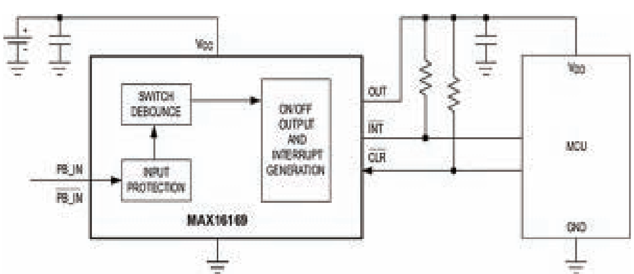

SWITCH DEBOUNCER/FRESHNESS SEAL PROVIDING POWER DIRECTLY TO LOAD

Analog Devices' MAX16169 nanoPower Pushbutton On/Off Controller and Battery Freshness Seal addresses issues associated with basic pushbutton debounce as well as minimizes battery drain before use. *Analog Devices*

Debounce circuits are nothing new, of course, and countless techniques are available ranging from crude R-C circuits to classic 555-based designs, and even specialized ICs that assess bounce timing using the system clock. While these may work acceptably in various circumstances, they all have a drawback that makes them unattractive for many of today's battery-powered products: their current requirements.

A Fresh Take on the Battery

The MAX16169 operates from a supply range of +1.3 to +5.5 V and consumes less than 40 nA of supply current to ensure minimal battery drain in low-power applications—and to enable use as a battery "freshness seal."

This seal doesn't imply that it somehow places the battery in a physical zippered bag or conformally coats it. Instead, a battery freshness seal is a feature in microprocessor supervisory circuits that disconnects a backup battery from any downstream circuitry until VCC is applied the first time. This keeps a backup battery from discharging until the first time a board is plugged in and actually used, and thus preserves the battery life.

Why is this "sealing" needed? Customers expect a quick startup experience when they unwrap and use the product for the first time. The most convenient way to provide this capability is to have the battery already installed and ready for use.

Unfortunately, in the indeterminate time lag between the battery's factory installation and the product's first turn-on—a period that can be many months—enough of its energy may have been drained even in standby mode and the battery can't deliver the needed power.

Reducing Battery Drain by 1,000X

By adding the MAX16169 pushbutton on/off controller, there's a way to enhance the overall off-the-shelf battery freshness goal. It provides a 1,000X reduction of the total shutdown-mode battery drain, with the standby current dropping from tens of milliamperes to tens of nanoamperes. The on/off controller completely disconnects the remainder of the circuit from the battery, thus creating the so-called "freshness seal" for the battery.

The MAX16169 family includes two sets of devices, one in which a more extended switch closure greater than the shutdown period de-asserts the latched output, and another in which the latched output stays asserted. The output provides a system interrupt whenever a valid pushbutton signal is detected. There are also two debounce timing options of 50 ms and two seconds.

The pushbutton input can handle up to ±60 V. ESD-protection structures are incorporated on all pins to guard against electrostatic discharges. There's extra protection in terms of static electricity to guard against ESD of ±40 kV at the switch input without damage. It operates over the –40 to +125°C temperature range and comes in a 2- × 2-mm, 6-pin micro dual-flat no-leads (μDFN) package.

WILLIAM G. WONG | Senior Content Director

Many Magnetic Arms Simplify **Hands-Free PCB Probing**

Editor Bill Wong takes a hands-on look at Sensepeek's PCBite kit.

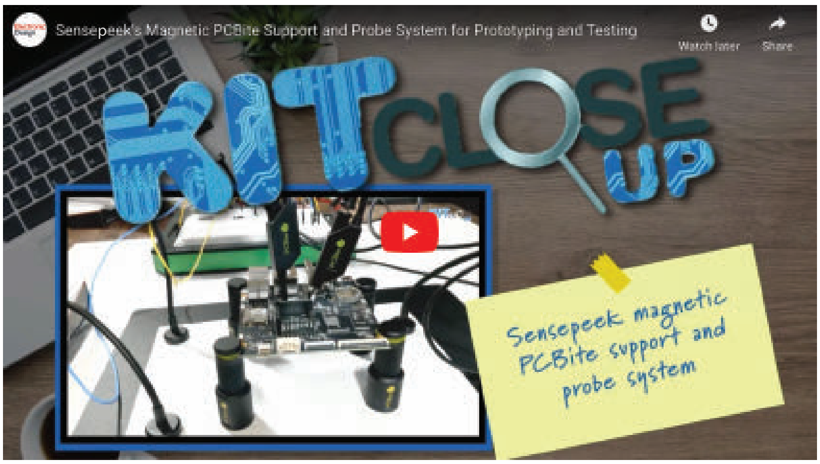

I WORK A LOT with development kits and particularly like to prototype robotic projects. I first saw Sensepeek's PCBite at the Sensors Converge conference and thought it was a great idea. Luckily, I was able to check it out myself with some of my projects *(see the video above)*.

PCBite consists of a metal base that works with spring-loaded magnetic standoffs and flexible arms with a magnetic base *(see figure)*. The latter can be oscilloscope probes or wire probes for things such as digital multimeters (DMMs). The parts are available individually or as a kit. There's even a magnifying glass with a magnetic base that can be positioned to view any area, which can be handy for checking very tiny LEDs or placement of the probes.

The standoffs have a very strong spring; just one can handle small printed circuit boards (PCBs). Large PCBs can be managed with two or more standoffs. I usually find three standoffs are more than sufficient unless I'm dealing with a very large PCB.

Support for the probe is another matter. The flexible arms keep the cabling out of the way, but they don't actually provide support. If you let go of the weighted end, it will just drop to the ground. Instead, that weight is designed to hold the probe in place. The points are spring-loaded so that you essentially touch the probe to the point on the circuit you want to track and press the weighted holder down to lock things into place.

I had a pair of oscilloscope probes and two regular probes, which turns out to be enough for most of my work. The hands-free operation becomes easier and trying to get the probe onto a very tiny PCB trace or pad was usually simple to do. Of course, attempting to get three or four of the probes into adjacent pins can be a challenge, but it's often doable with careful positioning.

PCBite doesn't work if you're trying to probe the bottom of a board. It may still pay to use clips in some instances as well, though mixing probing techniques is second nature for most developers. Whatever works is the way to go. At this point, PCBite is a new tool in my prototyping and testing toolbox. ed

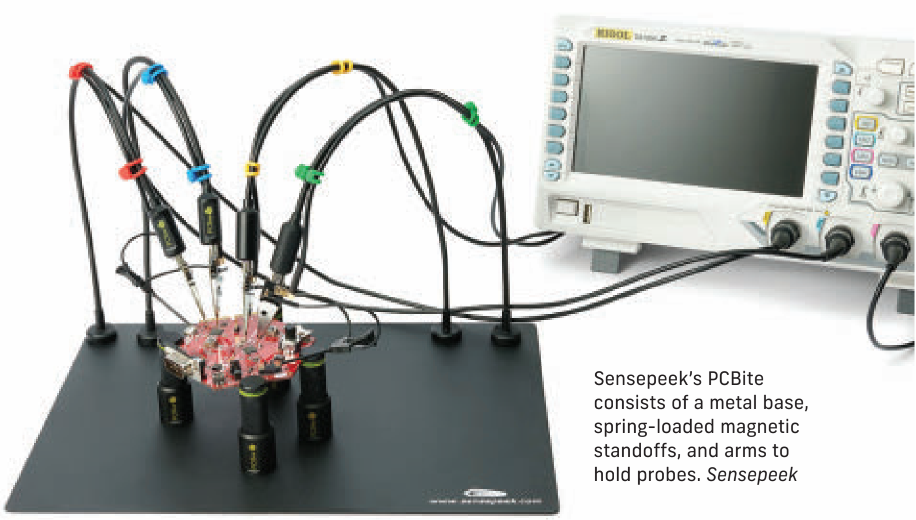

Sensepeek's PCBite consists of a metal base, spring-loaded magnetic standoffs, and arms to hold probes. *Sensepeek*

(Continued from page 24)

in many applications, it's important that no energy flows from the energy-storage system back to the power supply *(depicted in Fig. 1)*.

As shown in *Figure 1*, the supercapacitor should only power the sensor circuit and not any other electronics that may be attached to the 24-V line *(left side of Fig. 1)*. The energy-storage system is normally designed to supply the local load and not the complete system attached to the 24-V supply voltage. This makes diode D in *Figure 1* necessary.

Figure 2 shows a new concept supported by the MAX38889 from Analog Devices (ADI). It's a highly integrated backup solution called Continua for power rails up to 5 V. A single IC with a few passive external components is all that's required. The MAX38889 has an integrated half-bridge, operated alternately in highly efficient buck and boost modes.

A complete, operable circuit is illustrated in *Figure 3*. The logic and the power switches are all integrated, so just a small external chip-scale inductor and a few backup capacitors are required, apart from the supercapacitor.

The integrated high-side power switch is executed with the True Shutdown technology developed by ADI. As a result, the system voltage can be separated from the CAP voltage so that no current flows from the CAP to the system if the CAP voltage is ever higher.

While plenty of backup solutions for various voltage and current ranges exist on the market, the compact MAX38889 Continua backup solution can easily be added to the 5- or 3.3-V supply line with minimal development and implementation effort. It also offers high conversion efficiency of up to 94% in charging and discharging modes to minimize the size and cost of the energy storage.

The integrated high-side power switch is executed with the True Shutdown technology developed by ADI. As a result, the system voltage can be separated from the CAP voltage so that no current flows from the CAP to the system if the CAP voltage is ever higher.

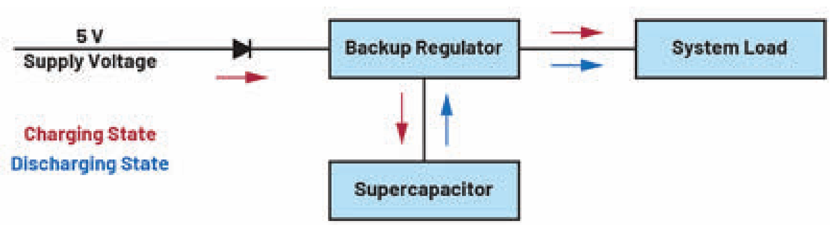

2. The Continua backup concept has numerous integrated system functions. *Analog Devices*

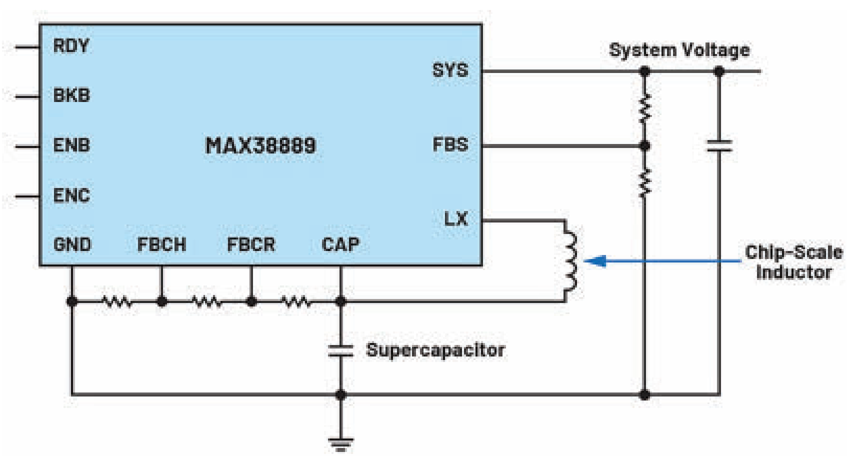

3. This implementation features a tiny Continua backup solution with the MAX38889. *Analog Devices*

REGISTER NOW FOR IMS2024

Pack Your Bags!

"Capitol"ize Your IMS Experience:

- RF to Optical • Systems Applications • The Human Dimension

IEEE MTT-S IEEE MICROWAVE THEORY & TECHNOLOGY SOCIETY


For more information,
Check out Electronicdesign.com!
ElectronicDesign
Microwaves&RF
LIBRARY
Generative
A I
APRIL 1ST