Electronic Design®

THE AUTHORITY ON EMERGING TECHNOLOGIES FOR DESIGN SOLUTIONS

MAY/JUNE 2020
electronicdesign.com

BIG SOLUTIONS for Machine Learning p8

$10.00

IN THIS ISSUE

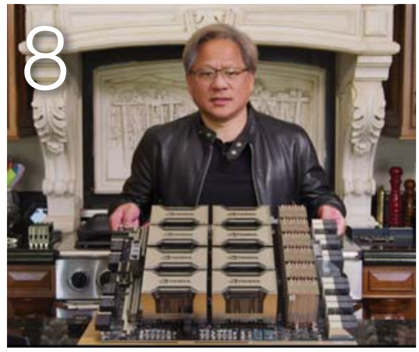

FEATURES

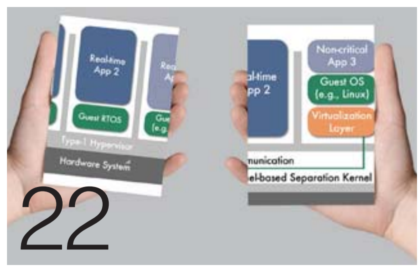

COLUMNS & DEPARTMENTS

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

ELECTRONIC DESIGN (ISSN 0013-4872) is published bi-monthly in Jan/Feb, Mar/Apr, May/June, July/Aug, Sept/Oct and Nov/Dec by Endeavor Business Media, 9800 Metcalf Ave., Overland Park, KS 66212-2216. Paid rates for a one-year subscription are as follows: $60 U.S., $90 Canada, $120 International. Periodicals postage paid at Kansas City, MO, and additional mailing offices. Editorial and advertising addresses: ELECTRONIC DESIGN, 605 Third Avenue, New York, NY 10158. Telephone (212) 204-4200. Printed in U.S.A. Title registered in U.S. Patent Office. For subscriber services or to order single copies, write to Electronic Design, PO Box 3257, Northbrook, IL 60065-3257. POSTMASTER: Send change of address to Electronic Design, PO Box 3257, Northbrook, IL 60065-3257. Canadian GST #R126431964.

Altech® ALL-IN 1 BACK-UP

Available power is automatically distributed between load and battery.

up to 35A

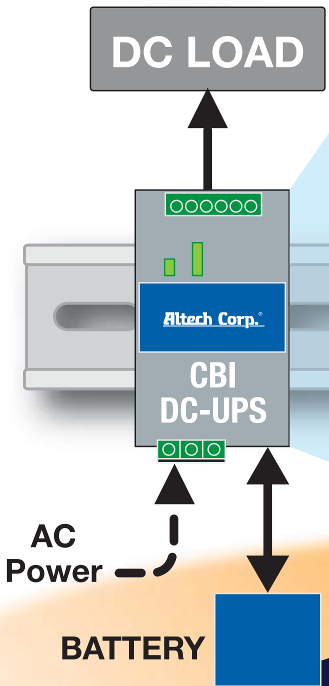

The CBI Series does all this:

1 Power Supply Unit
2 Battery Charger
3 Battery Care Module
4 Backup Module

up to 10A

up to 20A

FEATURES

- Three charging modes
- Compact, rugged metal case
- Available in 12, 24 and 48V DC
- Suitable for most common battery types
- Adjustable charging current
- Easy battery diagnosis and fault identification either by LED or external devices
- Status contacts
- High efficiency up to 91% through switching technology
- Output protection features such as short circuit, overload, deep battery discharge etc.
- DIN rail mounting
- Small size
- 3 year warranty

RoHS COMPLIANT · C UL US LISTED · Pb LEAD FREE · CE · ENERGY SAVER

Others offer a 2-3 components solution ...

... only Altech offers the CBI All-In-One solution.

Altech Corp.®

Your Source for Automation & Control Components

info@altechcorp.com • 908.806.9400

AltechCorp.com/POWER

ElectronicDesign. **MAY/JUNE 2020**

EDITORIAL

SENIOR CONTENT DIRECTOR: **BILL WONG** bwong@endeavorb2b.com
SENIOR EDITOR: **DAVID MALINIAK** dmaliniak@endeavorb2b.com
SENIOR STAFF WRITER: **JAMES MORRA** jmorra@endeavorb2b.com
CONTRIBUTING EDITOR: **LOUIS E. FRENZEL** lfrenzel@endeavorb2b.com
ASSOCIATE EDITOR/COMMUNITY MANAGER: **ROGER ENGELKE** rengelke@endeavorb2b.com

ART DEPARTMENT

GROUP DESIGN DIRECTOR: **ANTHONY VITOLO** tvitolo@endeavorb2b.com
ART DIRECTOR: **JOCELYN HARTZOG** jhartzog@endeavorb2b.com

PRODUCTION

GROUP PRODUCTION MANAGER: **GREG ARAUJO** garaujo@endeavorb2b.com
PRODUCTION MANAGER: **DEANNA O'BYRNE** dobyrne@endeavorb2b.com

AUDIENCE MARKETING

USER MARKETING MANAGER: **DEBBIE BRADY** dmbrady@endeavorb2b.com
FREE SUBSCRIPTION / STATUS OF SUBSCRIPTION / ADDRESS CHANGE / MISSING BACK ISSUES
OMEDA **T** | 847.513.6022 **TOLL FREE** | 866.505.7173 **FAX** | 847.291.4816 electronicdesign@omeda.com

SALES & MARKETING

REGIONAL SALES REPRESENTATIVES:
AZ, NM, TX: **GREGORY MONTGOMERY** **T** | 480.254.5540 gmontgomery@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: **JAMIE ALLEN** **T** | 415.608.1959 **F** | 913.514.3667 jallen@endeavorb2b.com
CT, DE, MA, MD, ME, NH, NJ, NY, PA, RI, VT, EASTERN CANADA:
ELIZABETH ELDRIDGE **T** | 917-789-3012 eeldridge@endeavorb2b.com
AK, NORTHERN CA, NV, OR, WA, WESTERN CANADA: **STUART BOWEN** **T** | 425.681.4395 sbowen@endeavorb2b.com

INTERNATIONAL SALES:
GERMANY, AUSTRIA, SWITZERLAND: **CHRISTIAN HOELSCHER** **T** | 011.49.89.95002778
christian.hoelscher@husonmedia.com
BELGIUM, NETHERLANDS, LUXEMBURG, UNITED KINGDOM, SCANDINAVIA, FRANCE, SPAIN, PORTUGAL:
LARA PHELPS lara.phelps@husonmedia.co.uk
ITALY: **DIEGO CASIRAGHI** diego@casiraghi-adv.com
PAN-ASIA: **HELEN LAI** **T** | 886 2-2727 7799 helen@twoway-com.com
PAN-ASIA: **CHARLES LIU** **T** | 886 2-2727 7799 liu@twoway-com.com

ENDEAVOR REPRINTS: **WRIGHT'S MEDIA** **T** | 877.652.5295
LIST RENTALS/ SMARTREACH CLIENT SERVICES MANAGER: **MARY RALICKI** **T** | 212.204.4284 mralicki@endeavorb2b.com

DIGITAL

SENIOR DIGITAL INNOVATION & STRATEGY DIRECTOR: **RYAN MALEC** rmalec@endeavorb2b.com

DESIGN ENGINEERING & SOURCING GROUP

VICE PRESIDENT, DESIGN & ENGINEERING: **TRACY SMITH** **T** | 913.967.1324 **F** | 913.514.6881 tsmith@endeavorb2b.com
VICE PRESIDENT OF MARKETING SOLUTIONS: **JACQUIE NIEMIEC** jniemiec@endeavorb2b.com

ENDEAVOR BUSINESS MEDIA, LLC

331 54th Ave N., Nashville, TN 37209 USA | www.endeavorbusinessmedia.com
CEO: **CHRIS FERRELL**
CRO/CMO: **JUNE GRIFFIN**
CFO: **WILLIAM NURTHEN**
COO: **PATRICK RAINS**
CTO: **ERIC KAMMERZELT**
CHIEF ADMINISTRATIVE AND LEGAL OFFICER: **TRACY KANE**
EVP KEY ACCOUNTS: **SCOTT BIEDA**
EVP KEY ACCOUNTS: **LINDA REINHARD**
EVP SPECIAL PROJECTS: **KRISTINE RUSSELL**
EVP GROUP PUBLISHER – DESIGN & ENGINEERING, ENERGY, BUILDINGS & CONSTRUCTION: **REGGIE LAWRENCE**

Electronic Design | Machine Design | Microwaves & RF | Hydraulics & Pneumatics | Source ESB | Source Today | Evaluation Engineering

Editorial

WILLIAM WONG | Senior Content Director
bwong@endeavorb2b.com

Don't Do It:
Ignore Best Practices

"The Code is more what you call guidelines, than actual rules."
—Hector Barbossa, *Pirates of the Caribbean: The Curse of the Black Pearl*

Best practices are generally accepted and prescribed policies and procedures often set by management, an authority, or governing body that are designed to improve quality, prevent problems, and streamline processes. They include, for instance, accredited management standards such as ISO 9000 and ISO 14001. And standards like MISRA-C are best practices for programming.

At this point, most of us are following best practices to avoid spreading and catching COVID-19. This involves things like social distancing, washing hands or using hand sanitizer regularly, wearing a face mask, avoiding coughing or sneezing toward someone, and cleaning and disinfecting surfaces.

The social isolation has led to the rise of video conferencing and telecommuting. These were often discounted, discouraged, or outright banned by many, but they have become the norm whether we like it or not. Though best practices for these exist, people often don't know them, or they ignore them.

For example, Zoom is a video-conferencing service that has become invaluable. In *"Rescuing Our Science Fair from COVID-19" (www.electronicdesign.com)*, I highlight how the Mercer Science and Engineering Fair was able to happen when Rider University shut down due to the virus. We used Zoom to do the judging and student interviews. Some were with students that had returned home to overseas destinations.

Zoom has been criticized because of a lack of security when users weren't following best practices like using password protection for meetings. Part of the problem is that there are many best practices. It's not really a good practice to default to password-less meetings.

Unfortunately, ignoring best practices isn't something unique to one set of practices, groups following those practices, or those defining them. Hector Barbossa highlights how many people regard best practices as just guidelines. Ignoring them is just something that's done.

Numerous best practices have been formalized and following them is often a requirement. Program coding standards are one example where the rules need to be followed and compliance can be enforced by software. Ada and SPARK programming language contract support takes this to the extreme, where rules are part of the code rather than just recommendations.

The thing is that best practices are typically specified to simplify a job, and changing circumstances often means modifying best practices or adopting new ones. Security and backup are two examples that crop up when it comes to telecommuting and remote access. Is your laptop, used at home for work, being backed up locally or to the cloud? Is the connection secured by a VPN? Is data being moved offsite that should not be? Are there controls implemented to manage this process, and are they automatic or manual? What are the best practices for all of these situations?

Finding best practices for a particular situation is an easy task. The internet can help, but given the plethora of recommendations and the quality of those recommendations, it's possible to wind up with a set of *best* practices that aren't the best for a particular situation. In fact, they may wind up being the worst practices for that situation.

These days one needs to be flexible. Adjusting the current set of best practices should be done carefully—however, sticking rigidly to some may not be the best practice.

There are potential benefits to reevaluating currently employed practices that have been—or need to be—changed based on the current climate. For example, we plan on having students post more information about their science fair projects online. It wasn't as burdensome as many had assumed and brought about the added benefit of the information being available to everyone. Essentially, part of the fair is permanent rather than transient.

So don't ignore best practices, but do try to adapt and find the best set for your situation.

WILLIAM G. WONG | Senior Content Director

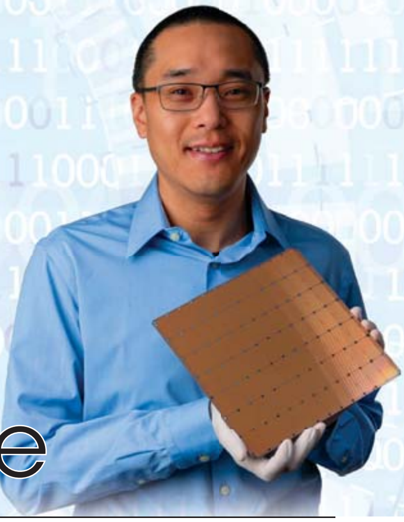

Accelerating Machine Learning Means New Hardware

Machine-learning algorithms can run on microcontrollers, but for complex applications, one really needs hardware acceleration.

Machine learning (ML) is only one aspect of artificial intelligence (AI). ML also has many parts to it, but those having the biggest impact now are based around neural networks (NNs). Even drilling down this much doesn't narrow the field a lot due to the multitude of variations and implementations. Some work well for certain types of applications like image recognition, while others can handle natural language processing or even modification and creation of artwork.

There are deep neural networks (DNNs), convolutional neural networks (CNNs), and spiking neural networks (SNNs). Some are similar while others use significantly different approaches and training techniques. All tend to require more significant amounts of computing power than conventional algorithms, but the results make neural networks very useful.

Though ML applications run on lowly microcontrollers, the scope of those applications is actually limited by the hardware. Turning to hardware tuned or designed for NNs allows designers to implement significantly more ambitious applications like self-driving cars. These depend heavily on the ability of the system to employ NNs for image recognition, sensor integration, and a host of other chores.

Hardware acceleration is the only way to deliver high-performance ML solutions. A microcontroller without ML hardware may be able to run an ML application for checking the motor it's controlling to optimize performance or implement advanced diagnostics, but it falls short when trying to analyze video in real time.

Likewise, processing larger images at a faster rate is just one ML chore that places heavy demands on a system. A plethora of solutions are being developed and delivered that provide orders of magnitude more performance to address training and deployment. In general, deployment needs are less than systems doing training but there are no absolutes when it comes to ML.

This year's Linely Spring Processor Conference was almost exclusively about AI and ML. Most of the presentations addressed high-performance hardware solutions. While many will land in the data center, a host of others will wind up on "the edge" as embedded systems.

WAFER-SCALE INTEGRATION TARGETS MACHINE LEARNING

Creating new architectures are making ML platforms faster; still, there's an insatiable need for more ML computing power. On the plus side, it's ripe for parallel computing and cloud-based solutions can network many chips to handle very large or very many ML models.

One way to make each node more powerful is to put more into the compute package. This is what Cerebras Systems' Waferscale Engine (WSE) does with identical chips, but it doesn't break up the die *(Fig. 1)*. Instead, the connections between chips remain, making the

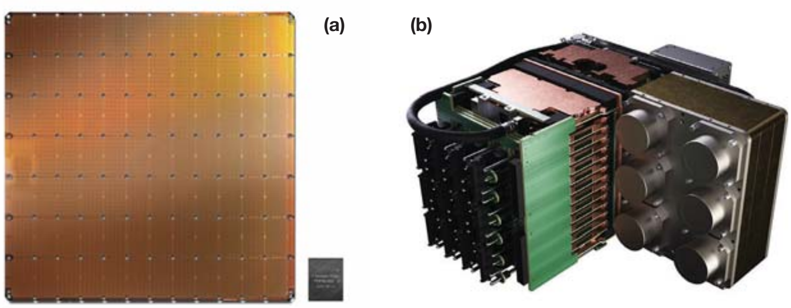

1. Shown is Cerebras' Wafer Scale Engine (WSE) machine-learning solution (a). It's designed to be used as is, not broken up into individual chips. Cerebras' WSE needs a water-cooled system to keep it running without a meltdown (b). *(Source: Cerebras Systems)*

46,225-mm^2 silicon solution the largest complete computing device with 1.2 trillion transistors that implement 400,000 AI optimized cores. The die has 18 GB of memory with 9 petabytes per second (PB/s) of memory bandwidth. The fabric bandwidth is 100 petabits per second (Pb/s). The chip is implemented by TSMC using its 16-nm process.

Each chip is power-efficient; however, packing this much computing power in a small package leads to lots of heat. Multiple die are put into one water-cooled system. Multiple systems can fit into a standard rack with Ethernet connections, allowing very large systems to be constructed. The interconnect and computational support have been optimized to handle sparse neural networks that are common for most applications.

SPIKING NEURAL NETWORK HARDWARE

Spiking neural networks (SNNs) have different characteristics than DNNs. One advantage with SNNs is that the support for learning is on par for deployment, whereas DNNs require lots of data and computational capabilities for training compared to deployment. SNNs can also handle incremental training. Furthermore, SNNs require less computational overhead because they only process neurons when triggered.

BrainChip's AKD1000 Neural Network SoC (NSoC) can handle both DNNs and SNNs. The architecture supports up to 80 neural processing units (NPUs)—the AKD1000 has 20 NPUs *(Fig. 2)*. A conversion complex implements a spike event converter and a data-spike event encoder that can handle multivariable digital data as well as preprocessed sensor data. The SNN support only processes non-zero events.

The AKD1000 benefits from sparsity in both activations and weights. It supports quantizing weights and activations of 1, 2, or 4 bits, leading to a small memory footprint. NPUs communicate events over a mesh network, so model processing doesn't require external CPU support.

Tenstorrent also targets SNN applications with its Tensix cores. The cores have five single-issue RISC cores and a 4-TOPS compute engine. A packet processing engine provides decoding/encoding and compression/decompression support along with data-transfer management.

As with most SNN platforms, Tensix cores can be used on the edge or in the data center. They provide fine-grained conditional execution that makes the system more efficient in processing SNN ML models. The system is designed to scale since it doesn't use shared memory. It also doesn't require coherency between nodes, enabling a grid of cores to be efficient connected via its network.

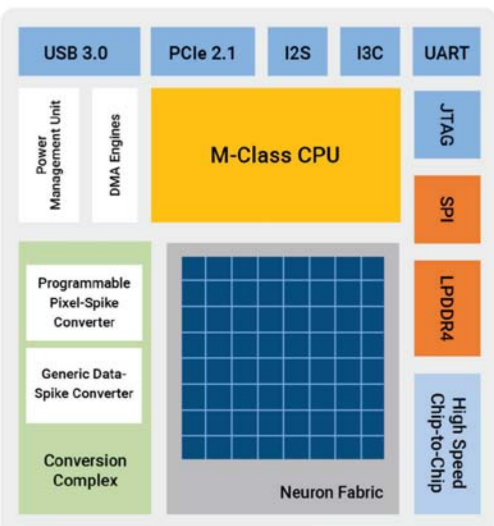

2. The AKD1000 neuron fabric developed by BrainChip supports spiking neural networks. A Cortex-M4 manages system resources.

GrAI Matter Labs also targets this event-driven ML approach with its NeuronFlow technology. The GrAI One consists of 196 neuron cores with 1024 neurons/core, which adds up to 200,704 neurons. A proprietary network-on-chip provides the interconnect. No external DRAM is needed. The SDK includes TensorFlow support.

CNNs, DNNs, AND MORE

Convolutional neural networks are very useful for certain kinds of applications like image classification. SiMa[ai] optimized its chip for CNN workloads. By the way, sima means "edge" in Sanskrit. The chip is also ISO 26262 ASIL-B compliant, allowing it to be used in places where other chips aren't suitable, such as automotive applications.

Flex Logix is known for its embedded FPGA technology. The company brought this expertise to the table with its nnMAX design and the InferX X1 coprocessor. The nnMAX array cluster is designed to optimize memory use for weights by implementing Winograd acceleration that handles input and output translation on-the-fly. As a result, the system can remain active, while other solutions are busy moving weights in and out of external memory. The chip supports INT8, INT16, and BFLOAT16. Multiple models can be processed in parallel.

Groq's Tensor Streaming Processor Chip (TSP) delivers 1 petaoperations per second, running at 1.25 GHz using INT8 values. The chip architecture also enables the system to provide this high level of performance by splitting data and code flow. The 20 horizontal dataflow superlanes are managed by the vertical SIMD instruction flow. Identical east/west sections let the data flow in both directions. There are 20 superlanes with 16 SIMD units each.

PROCESSORS AND DSPs

Special ML processors are the order of the day for many new startups, but extending existing architectures garners significant performance benefits while keeping the programming model consistent. This allows for easy integration with the rest of an application.

Cadence's Tensilica HiFi DSP is now supported by its HiFi Neural Network library in addition to the Nature DSP library that handles vector math like FFT/FIR and IIR computations. The 8-/16-/32-bit SIMD and Vector FPU (VFPU) support provides efficient support for neural networks while enabling a custom design DSP to include customer-specific enhancements.

CEVA's SensPro sensor hub DSP combines the CEVA-BX2 scalar DSP with a NeuPro AI processor and a CEVA-XM6 vision processor. The wide SIMD processor architecture is configurable to handle 1024 8-×-8 MACs, 256 16-×-16 MACs, or dedicated 8-×-2 binary-neural-network (BNN) support. It can also handle 64 single-precision and 128 half-precision floating-point MACs. This translated to 3 TOPS for the 8-×-8 network's inferencing, 20 TOPS for BNN inferencing, and 400 GFLOPS for floating-point arithmetic.

The DesignWare ARC HS processor solution developed by Synopsys takes the tack of having lots of processors to address the ML support. This isn't much different than most solutions, but it's more along the lines of conventional RISC cores and interconnects that are typically more useful for other applications.

AMD isn't the only x86 chip producer. Via Technologies has its own x86 IP and its Centaur Technology is making use of that. The x86 platform is integrated with an AI Ncore coprocesessor tied together by a ring. The Ncore utilizes a very wide SIMD architecture organized into vertical slices to provide a scalable configuration, making future designs more powerful. The chip can deliver 20 TOPS at 2.5 GHz.

I've previously covered the Arm Cortex-M55 and Ethos-U55 combination. The Cortex-M55 has an enhanced instruction set that adds a vector pipeline and data path to support the new SIMD instructions. The DSP support includes features like zero overhead loops, circular buffers, and bit reverse addressing.

Still, as with other architectures, a dedicated AI accelerator is being added to the solution—namely, the Ethos-U55 micro network processor unit (microNPU). It supports 8- or 16-bit activations in the models, but internally, weights will always be 8 bits. The microNPU is designed to run autonomously.

While the V in RISC-V doesn't stand for vectors, SiFive's latest RISC-V designs do have vector support that's ideal for neural-network computational support. What makes this support interesting is that the vector support can be dynamically configured. Vector instructions work with any vector size using vector-length and vector-type registers. The compiler vectorization support takes this into account. The VI2, VI7, and VI8 platforms target every application space through the data center.

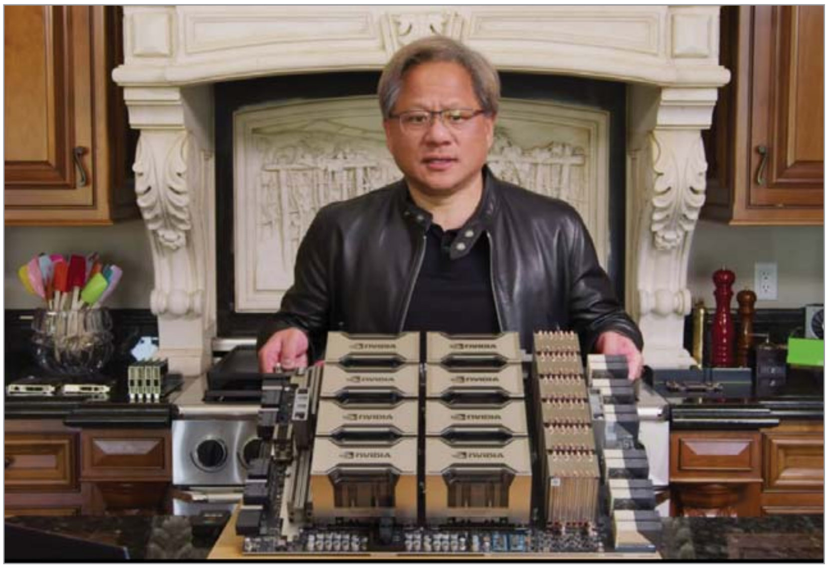

3. Nvidia's Jensen Huang just pulled this motherboard out of the oven. It has eight A100-based GPGPUs specifically designed for AI acceleration.

EXTENDING FPGAs AND GPGPUs

Xilinx's Versal adaptive compute acceleration platform (ACAP) is more than just an FPGA. The FPGA fabric is at the center, providing low-level customization. However, there are hard cores and an interconnect network surrounding it. The hard cores range from Arm Cortex CPUs for application and real-time chores along with AI and DSP support.

I left Nvidia to the end, as the company announced its A100 platform at the recent, virtual GPU Technology Conference *(Fig. 3)*. This GPGPU incorporates a host of ML enhancements, including sparsity acceleration and multi-instance GPU (MIG) support. The latter provides hardware-based partitioning of the GPU resources that allow more secure and more efficient operation. Large-scale implementations take advantage of the third-generation NVLink and NVSwitch technology that tie multiple devices together.

The plethora of machine-learning options includes more than just the platforms outlined here. They reflect not only the many ways that ML can be accelerated, but also the variety of approaches that are available to developers. Simply choosing a platform can be a major task even when one understands what kind of models will be used, and their possible performance requirements. Likewise, the variants can offer performance differences that are orders of magnitude apart. System and application design has never been more exciting or more complicated. If only we had a machine-learning system to help with that.

WILLIAM G. WONG | Senior Content Director

How and Why Standardization Will Benefit the E-Textile Industry

In e-textiles, developers spend too much time finding out which component works with which. What the industry needs is a collaborative effort in the electronics and textile industries to work toward common standards.

E-textile is a booming market in the electronics sector—and in the textile sector. Being at the crossroads of two industries has problems of its own. Christian Dalsgaard, founder and former CTO of Ohmatex in Denmark, learned that the hard way while developing state-of-the-art wearables for the European Space Agency, among others. He argues for standards in e-textile. Senior Technology Editor Bill Wong chats with Christian about this growing need.

WHICH PROBLEMS ARE SHAPING E-TEXTILE NOWADAYS?

There are two clusters of problems—the first cluster has to do with the central challenges of e-textile. The most important one is unifying something hard with something soft, something rigid with something flexible. Also, you need to make a device as small as possible, while maintaining its performance. And then there's the challenge to make cabling, connectors and electronics sealing that will survive hostile environments—sweat, washing machines, tumble dryers, to name the most obvious.

The second cluster of problems is about the fact that two supply chains are involved in e-textile. One is concentrated around the traditional textile industry, where there's a group of companies involved in the production of conductive yarn, sensing fabrics, and confectioning garments. The other is concentrated around the production of electronics and mechanical components, involving robotics and automated processes.

Textile and electronics are coming together in e-textile, but they come from very different backgrounds. Some of the elements include price points that are very different and the manufacturing process for cloth is still labor-intensive.

Moreover, the turnaround time in fashion is short. There needs to be a new collection every three months, whereas the turnaround time in electronics is on average two years. Of course, there are exceptions—certain headphones from 2007 are still hot and being sold.

People working in electronics are generally nerds rooted in STEM, while in textile they hate mathematics and science. There's also a gender factor. In the textile industry, there's a predominance of women, while in electronics you'll find mostly men.

The gender aspect influences the industry on many levels, from the people working in it to the solutions consumers are interested in. Women choose clothes for being comfortable, or for expression, while expression is much less of an issue for men, who rather look for functionality in their clothes—to protect themselves from heat or cold. These factors play a minor role in electronics.

In short, each industry has its own standards and accepted solutions, embedded in very diverse industrial history and cultures.

WHICH PROBLEMS IN E-TEXTILE ARE SUFFERING MOST FROM THE LACK OF STANDARDIZATION?

Imagine a world of mobile phones without USB or Bluetooth standards. A supplier would need to develop headsets, hands-free car kits, loudspeakers, and many other accessories—each device individually designed for a specific brand. This is the situation e-textile is in, which we need to get out of as fast as possible.

First, let's discuss connectors. They need to be soft and easy to integrate, but most of them are clumsy, rigid, and hard, especially in the military. Connectors in the consumer market, such as USB and USB-C, are much more consumer-friendly because they're small and widely used. However, they're fragile when applied to cloth, and tend to gather dirt and washing powder when being washed.

We struggled with the connector in the PLACE-it project—a project Ohmatex was involved in with, among others, Philips. For this project, we developed a connector for light-emitting textiles used in phototherapy medical devices for treating babies with jaundice, or people with skin disease. To properly make the connector, we had to develop it from scratch. It was a beautifully thin and flexible interconnect that we unfortunately couldn't afford to upscale because there was no standard we could apply to the thin design.

Another problem is wiring for both power and data transfer, which needs to be flexible and must be sewn or glued into seams. There are no bus standards, textile cables, or washable interconnects that can be used straight out of the box. You need to combine a mix of wires, textile yarns, and elastic weaving methods with custom, specific connectors. And you need to solve the strain relief to prevent the wires from breaking quickly. Finally, you need shielded cables to isolate radiation and radio interference.

Making wearable processing units are a nightmare, too. A skiing jacket with heating elements, for example, needs a controlling unit with an on/off switch, an LED panel, and controls for power level, temperature levels etc. The same goes for ECG shirts, where the interface controls are very similar. If you were building a computer, you would buy, for instance, a standard network board, a graphic card from Asus or G-force—components that are plug-and-play.

In e-textile, you sit together with an electronic engineer who has, by that time, already been working on it for months—ordering resistors, capacitors, ICs; working at a breadboard, etc. And then you're only just getting started! Compared to making computers, making a garment in e-textile is like making the garment AND having to invent zippers, reflectors, yarn, a sewing machine and a washing machine—and don't forget the washing powder, too!

Developers are struggling not only with the absence of standards, but also with the non-existence of industrially proven solutions. You can find breathable and stretchable conductive fabrics that are used for many bio-sensing applications, like ECG, skin temperature measurement, and motion capturing.

But such fabrics are hard to fixate to a foil-based electronics board, which is industry standard in electronics. There's no glue specifically developed for this purpose, so as a developer you need to conduct a lot of experiments yourself with pressure, heat, etc., before you can find the right combination that works for the wear and tear endured by that e-textile.

Speaking of wear and tear, washing is a good example of where the two supply chains need to join forces and be willing to venture out in each other's field of expertise. Take, for instance, soft displaying—they have to endure washing when built into a sleeve. Here responsibility lies with both parties, but some people in electronics have been known to elegantly drop the responsibility like a hot potato, saying "washing is not our task."

Yet I have also witnessed the textile side saying the same thing: "Oh, that's electronics, we know nothing about that." Yes, but in the case of washing there is, for obvious reasons, far more expertise on the textile side than on the electronics side! You can see how this slows down the development process. More expertise in stressing electronics, and standards for it, would benefit the whole e-textile industry significantly.

SO, WE CAN AGREE ON THE NEED FOR STANDARDIZATION IN E-TEXTILE. WHAT'S NECESSARY TO MOVE FORWARD?

What e-textile needs is a collaborative effort, in both the electronics industry and textile industry, to work toward common solutions for interconnects,

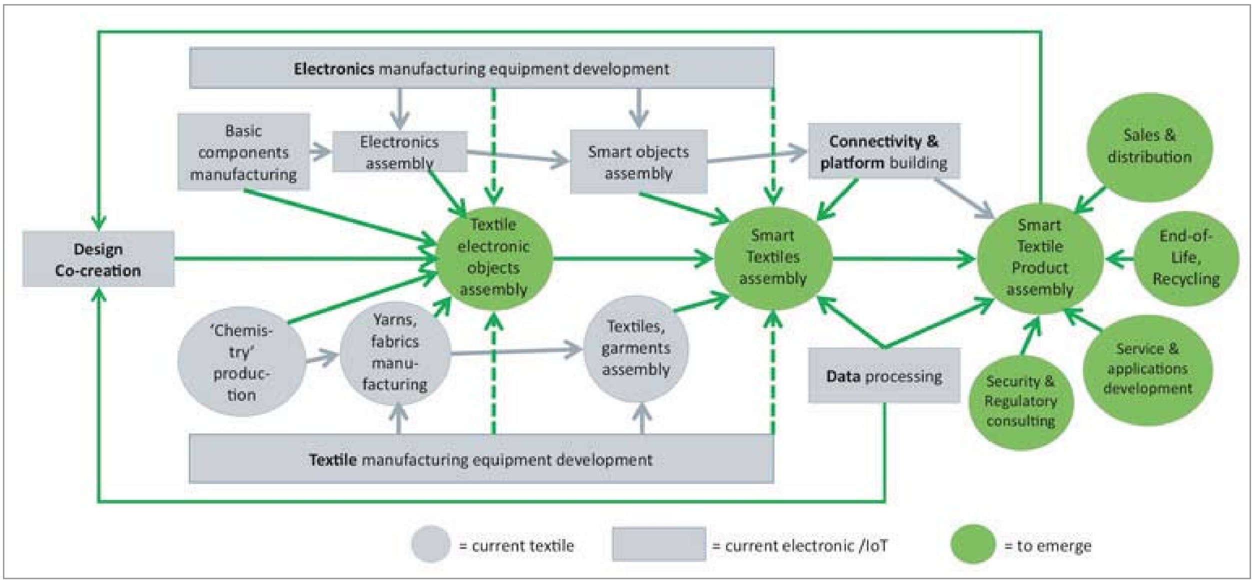

1. This novel industrial value chain in e-textile shows the complexity of the industry. *(Courtesy of Centexbel, DITF and the SmartX project)* data and power buses, wearable processing units, soft displaying, and energy harvesting. An e-textile developer should in principle only be concerned about sensor integration, the industrial design of body-worn devices, and embedded software for the applications.

The big question is, of course, how to bring about such common solutions. In the 1980s, Bill Gates formed a partnership with IBM to bundle Microsoft's operating system with IBM computers. That was a tremendous breakthrough for the whole software industry, because now you didn't need to manage program execution, files access, display interface, and devices like the keyboard and mouse.

A similar revolution is needed for e-textiles, but it's not going happen in a startup, like Microsoft was at the time, because the cost is too high for a single company. Also, many skills like industrial design, electronics design, textile engineering and embedded software are necessary to make it happen. It is more a merger of existing technologies rather than a completely new approach, like the Microsoft/IBM combination was. Such a merger will not take place overnight. It will be an incremental process rather than a paradigm shift, and a process that's hard to predict, given the number of players involved.

With Bluetooth, inventor Ericsson organized a Special Interest Group with four other companies initially, who pooled their knowledge and experience and added to the Bluetooth technology through the years. This helped develop and refine Bluetooth technology, and it got even better when more companies came on board and helped establish the standard.

On the hardware side, USB is an example of a standard solution that gained permanence in the market thanks to collaboration between producers of computers and producers of devices. Growth of USB was also greatly helped by the release of Windows 98 and the adoption of USB by Apple, which dropped all other connections in favor of a USB port with the introduction of the first iMac.

WHAT CAN E-TEXTILE LEARN FROM THESE EXAMPLES?

Trends and movements toward standardization and/or merging are already happening in textile and electronics. In 2004, when I started out in e-textile, we had to develop our own conductive yarns. This is no longer necessary because many yarn manufacturers have developed a variety of very good products with excellent conductive properties and strength. The IPC D-72 E-Textiles Materials Subcommittee is now able to finalize a standard IPC-8921, covering requirements for conductive fibers and conductive yarns, including standardized key characteristics, durability testing, and industry test methods.

Such standards are the result of a retrospective effort. What's needed to speed it up is a proactive development in both the electronics and textile industries to define standards for a more comprehensive microsystems architecture.

SmartX-Europe has taken a stab at doing this *(Fig. 1)*, but as you can see, it's still a pretty complicated framework in which the two industries have yet to really merge. Also, End-of-Life needs more elaboration, since sustainability has become a major factor in the whole process of e-textile.

If we stick to this model, I think that the e-textile industry as a whole should concentrate on the two left-hand green "bubbles": Textile Electronics Objects Assembly and Smart Textiles Assembly.

The OSI seven-layer model for communication stacks became the reference model for not only protocols, but also physical components like Ethernet connectors, routers, and switches. Having such a framework could be a great step for smart textiles. I could envision a model consisting of four layers *(Fig. 2)*.

One layer covers standards for yarns and fabrics and the transition between the flexible textile substrate and hard connections point. It's basically to adapt existing flexprint standards for foil to also include textiles. The next layer is the connection layer, where we desperately need a flat magnetic USB-T connector that's washable and that could be attached to textiles like any other snap fasteners from underwear to outdoor jackets.

On the IoT layer, I want standard electronics modules that can be sliced or snapped to a connector, providing a suitable programming platform with a full Bluetooth BLE/Wi-Fi stack, wireless charging, and a standard I/O bus for sensors based on USB, so that we have an interoperability between the connection layer and the IoT layer. On the care layer, standards for washing and tumble-drying wearables and other textile integrated technology can be defined.

I would also argue for choosing possible areas or ranges of components for standardization carefully. Some are more mature for standardization than others. Soft displaying, as mentioned above, and power are fields where lots of exciting new things are happening. Many companies are working on autonomous harvesting, meaning that the wearer is simultaneously consuming and producing power—power is harvested from the wearer's movements.

Autonomous harvesting is at a very dynamic stage, where standardization would delay further development and optimization. But in other areas where development is leveling now, like housing or interconnects, standardization would greatly help the whole of the e-textile industry.

PRACTICALLY SPEAKING, HOW WOULD YOU ORGANIZE THIS PROCESS?

I would suggest setting up an alliance of important players in the whole value chain, meaning both textile and electronics, pledging to work out a Smart Textile Standard for main elements that are leveling now in e-textile. I would start with connectors. I think this seems to be the most promising approach with the least amount of effort. We can learn from that process, then take on another element and work our way forward.

To do this, we need to find financial support before June 2020 so that the alliance can start. Then we could have a preliminary standard ready by the end of 2021. Together with skilled and knowledgeable people from the smart textile industry, we are right now working on establishing a Centre of Excellence in London, which could be up and running by June to support the alliance.

The vision is to make electronics soft and textiles smart by globally aligning and standardizing the smart textile industry on the four layers described in *Figure 2*. The mission would be to develop standards and lead successful cross-industry alliances.

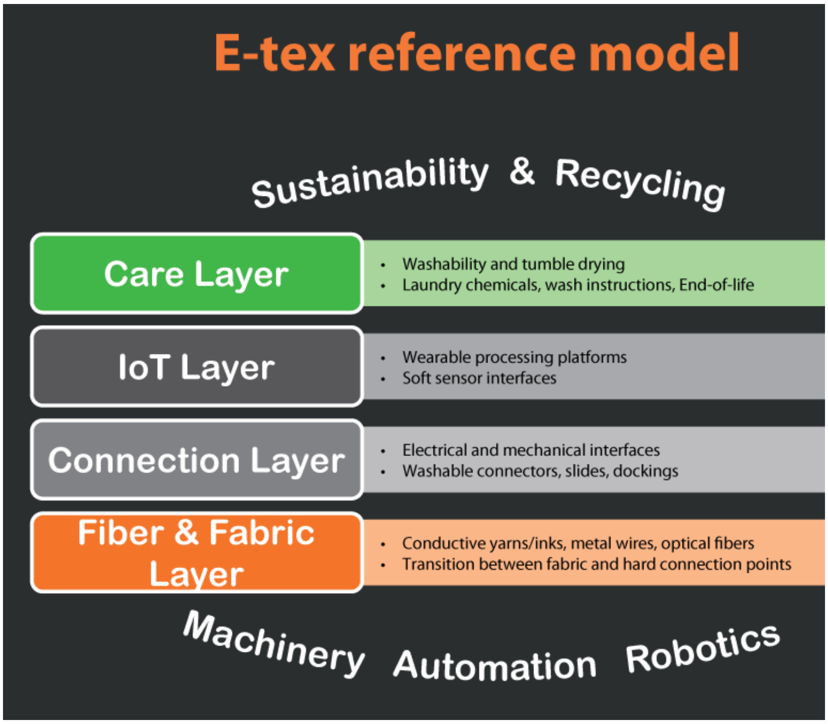

2. This e-tex reference model echoes the OSI seven-layer model. *(Christian Dalsgaard)*

HOW WOULD YOU CONVINCE THE MAIN PLAYERS IN THE E-TEXTILE INDUSTRY TO JOIN YOU?

E-textile is on the brink of entering new markets in medical, gaming, and sport if we can provide standard components and reach a price point where we can compete with handheld devices.

Take the Google Jacquard by Levi's, where an interactive cuff communicates with your smartphone through an app. That way you can access vital apps and functions just by tapping the cuff and hear information and music through your earphones. Right now, Levi's starting price of 198 U.S. dollars for the classic Trucker Jacket with Google Jacquard technology is more than 100 dollars higher than a traditional Trucker Jacket. For the technology to become more competitive, this price needs to go down.

But an even bigger impending breakthrough is a thin emulating glove, which is basically a textile hand-shaped sensor and accelerometer. It registers movements and translates them into 3D images. The same principle is already being used in the film industry, only with painted dots on an actor's head, body, and limbs. Such a glove technique renders work with laser setups in studios and labs obsolete. An athlete, actor, or patient now only needs to wear emulating garments. In addition to the film industry, those within the sports, medical and gaming industries can use this technique to enhance user experience, but at much lower costs than required by present-day techniques.

The crucial issue here is standardization of the required components. And for this we need an attitude of finding common interests between two totally different industries. Importantly, people should work together and share an interest in finding solutions. Another key element is that we need to work with open standards.

In return, the partner companies in the alliance will benefit from being first-movers and gain close technical insight into the solutions discovered by the alliance. Think of Bluetooth, its logo visible on all sorts of devices. Whoever finds a new solution will have the support of all partners in the supply chain because they are the clients. They rely on the new technology that's been developed. This creates a critical mass in a short timespan, which is crucial for creating a new market or penetrating existing ones. We're just one step away from making this happen.

ADHV4702-1
24 V to 220 V Precision Operational Amplifier

Yihang Yang, Applications Engineer

Introduction

The ADHV4702-1 is a high voltage (220 V), unity-gain stable precision operational amplifier. The next generation of proprietary semiconductor processes and innovative architecture from Analog Devices enable this precision operational amplifier to operate from symmetrical dual supplies of ±110 V, asymmetrical dual supplies, or a single supply of 220 V.

The Industry's First 220 V Precision Operational Amplifier

The ADHV4702-1 has a 170 dB typical open-loop gain (A_{OL}) and a 160 dB typical common-mode rejection ratio (CMRR). The ADHV4702-1 also has a 2 μV/°C maximum input offset voltage (V_{OS}) drift and an 8 nV/$\sqrt{Hz}$ input voltage noise. The exceptional dc precision of the ADHV4702-1 is complemented by excellent dynamic performance with a small signal bandwidth of 10 MHz and a slew rate of 74 V/μs. The ADHV4702-1 has an output current of 20 mA typical. In addition, its unique features such as adjustable supply current, slew boosting circuitry, and flexible exposed pad bias voltage make this part the ideal high voltage solution for a wide range of applications.

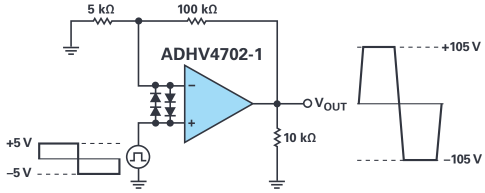

Figure 1. ADHV4702-1 output swing circuit and capability.

Applications

ADHV4702-1 is the first amplifier in the market that offers high voltage and precision performance in a small form factor. The ADHV4702-1 solves difficult design challenges and can be used in many different applications such as automated test equipment, life sciences, LIDAR, and healthcare. In automated test equipment applications, the device can be used for high-side current measurement and high voltage precision supply generation. For life sciences, the ADHV4702-1 can be used to provide precision high voltage control for mass spectrometry systems. In LIDAR applications, it can be used to accurately control the APD bias voltage. In healthcare applications, the product can be used to tightly control the bias point of silicon photomultipliers.

12-Lead, 7 mm × 7 mm LFCSP Compliant with IEC 61010-1 Spacing

The ADHV4702-1 is available in a 12-lead, 7 mm × 7 mm lead frame chip scale package (LFCSP) with an exposed pad compliant to International Electrotechnical Commission IEC 61010-1 creepage and clearance standards. This package significantly reduces the solution size and simplifies system architectures by eliminating supporting components like dc-to-dc converters and floating supplies.

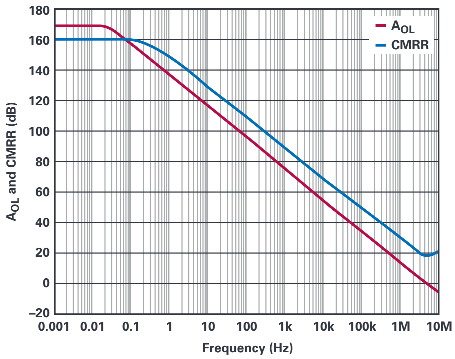

Figure 2. ADHV4702-1 precision performance.

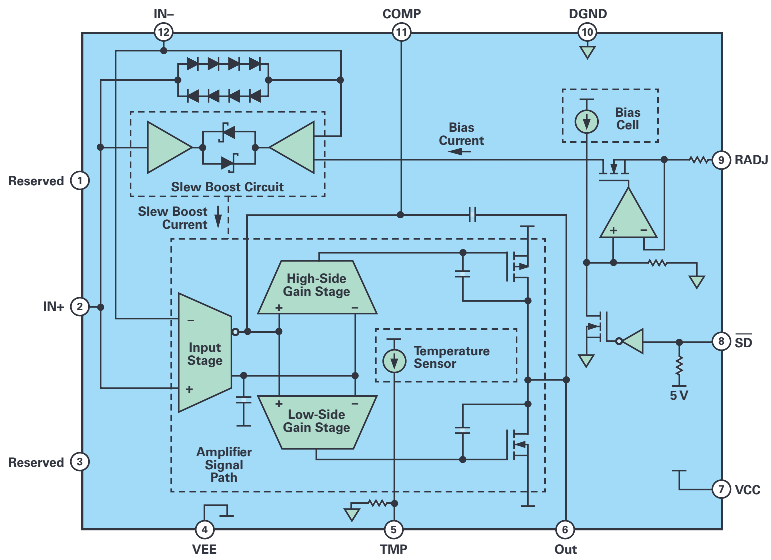

Figure 3. ADHV4702-1 functional block diagram.

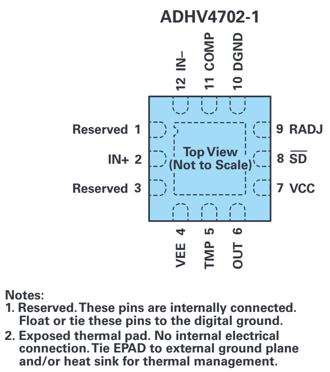

Figure 4. ADHV4702-1 pin configuration.

VAL CHRYSOSTOMAU | Video Display System Product Marketing Manager, Curtiss-Wright Defense Solutions
www.curtisswrightds.com

Cutting-Edge Video Solutions Essential for Enhanced Situational Awareness

Situational awareness is a critical capability in the battlefield. Tools and technologies, including real-time sensors, video displays, mission computers, and video-distribution systems, are evolving to bring new advanced solutions to the warfighter.

1. Raw data must be transformed into valuable information before it can provide the knowledge and understanding essential for establishing situational awareness.

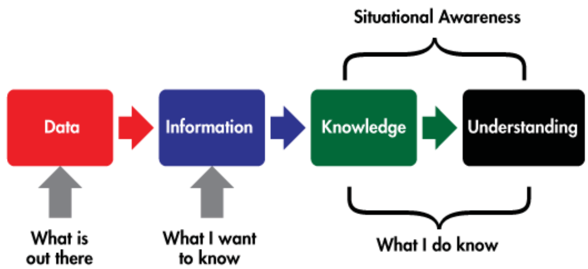

2. Today's cutting-edge video solutions, leveraging the power of visual representation and information, provide the most effective way to deliver actionable intelligence to the warfighter.

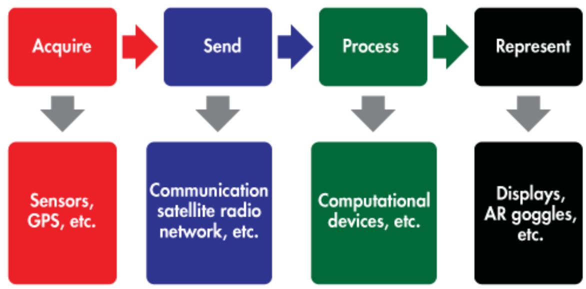

While several definitions exist for situational awareness and situational understanding, they all share the ideas expressed in the Reconnaissance, Surveillance, and Target Acquisition (RSTA) squadron definitions provided by Major Brad C. Dostal, a military analyst at the U.S. Center for Army Lessons Learned (CALL) *(https://www.globalsecurity.org/military/library/report/call/call_01-18_ch6.htm)*. Situational awareness is "the ability to maintain a constant, clear mental picture of relevant information and the tactical situation, including friendly and threat situations as well as terrain."

According to Dostal, situational understanding allows leaders to avoid surprises, make rapid decisions, and choose when and where to conduct engagements, as well as achieve decisive outcomes. It's "the product of applying analysis and judgment to the unit's situational awareness to determine the relationships of the factors present and form logical conclusions concerning threats to the force or mission accomplishment, opportunities for mission accomplishment, and gaps in information."

Raw data becomes valuable information that leads to knowledge and understanding *(Fig. 1)*.

Situational awareness is also a key component in the Observe, Orient, Decide, Act (OODA) loop developed in the mid-20th century by U.S. Air Force Colonel John Boyd, a military strategist. The OODA loop is a four-step approach to decision-making that focuses on filtering available information, putting it in context, and quickly making the most appropriate decision while also understanding that changes can be made as more data becomes available.

SITUATIONAL AWARENESS TOOLS AND TECHNOLOGIES ARE ALWAYS EVOLVING

The tools and technologies that provide situational awareness are continuously evolving. In the earliest days of battle, it was a huge advantage to have a map. Later, sextants became an important situational awareness tool. And on it goes through the invention of still cameras, video cameras, global positioning systems (GPS), and other innovations.

As technologies continue to advance, increasingly sophisticated onboard sensors and mission systems have become the eyes and ears for everyone on ground and air platforms. And cutting-edge video solutions have become the best way to get all of this information to warfighters in a fast and effective way *(Fig. 2)*.

A look at the different levels of video systems available for ground vehicles illustrates the effect that increasingly advanced technologies have on situational awareness:

- The most basic video systems allow warfighters who are in the vehicle or operating it remotely to view the images from one vehicle-mounted camera at a time, providing a rudimentary level of visibility.
- Multi-display and picture-in-picture solutions enable warfighters to simultaneously view images from multiple vehicle-mounted cameras so that they can consider their surroundings on all sides of the vehicle at all times. With potentially more than a dozen camera views to choose from, warfighters can easily access the optimal combination of views for the task or maneuver they're executing.
- 360-degree video systems give warfighters an even higher level of situational awareness. These systems blend accurate, fully stitched images from all vehicle-mounted cameras into a seamless, panoramic image that most closely resembles what the human eye sees.

The next step is to move this crucial insight and visibility right into warfighters' helmets or visors. In the next five years or so, we can expect to see AR solutions that put the video feeds, maps, GPS coordinates, speed indicators, and other information warfighters need right in their field of view. They will be able to use intuitive gestures to control the information they see at each phase of task execution, similar to movie depictions.

REQUIREMENTS AND CONSIDERATIONS FOR ENHANCED SITUATIONAL AWARENESS SOLUTIONS

The goal of every enhanced situational awareness solution is to get the right information to the right person at the right time in a form that can be rapidly assimilated and used. The only way to achieve this goal is to deploy end-to-end video solutions where all system components are seamlessly integrated.

Any other approach introduces significant risks, including interoperability issues, increased latency, increased complexity, and more challenging maintenance and upgrades. These shortcomings will make it far too difficult to acquire, distribute, and intuitively present the extremely high volumes of data that will be available in the battlespace of the future.

The end-to-end architecture for enhanced situational awareness varies depending on the specific application and platform. However, most solutions require sensors in addition to video, mission computer, and a video distribution system (VDS).

As technologies evolve, architectures will also evolve. As a result, each system component must be:

- Flexible
- Adaptable
- Expandable
- Interoperable
- Easy to deploy, use, and upgrade
- Low SWaP

Extremely low video-system latency is also essential for providing warfighters with complete confidence that what they're seeing is the reality at the time. If information is delayed and warfighters are uncertain about their situation, they're far more likely to hesitate when responding to threats, collide with obstructions or humans, or unknowingly enter dangerous territory.

To reduce latency end to end, each video system component must function in as close to real time as possible:

- A United Kingdom Ministry of Defence study found that military vehicle drivers can safely drive a vehicle through a visual display when the overall video-system latency is 40 ms or lower.
- A study looking into the effects of video latency on general situational awareness found that warfighters remain adequately aware of their surroundings when overall video-system latency is 160 ms or lower.

REAL-TIME SENSORS

The sensors that capture information range from high-definition, thermal, and infrared cameras to GPS, radar, speed indicators, and other platform systems. While the data from video cameras is typically compressed before it's sent to the mission computer for processing, data from other systems may be sent in the raw form.

MISSION COMPUTER

With increasingly sophisticated adversarial threats and advances in video technologies, many enhanced situational awareness solutions now require computer-processing capabilities.

When computer processing is added, video systems can manipulate video feeds to enable video blending and video stitching. They can layer regular camera feeds with infrared and thermal-imaging feeds. And, they can display mapping and telemetry data and image metadata along with video streams.

When these advances first came along, it seemed logical to add computer-processing capabilities to the back of the mission display. Since then, technology has continued to evolve, and all-in-one video-display and computing solutions no longer have an obvious advantage over simple, standalone displays. In many cases, a simple video display connected to a separate computing component is the better choice because it:

- Reduces the cost of computer processing compared to smart displays
- Simplifies upgrades to computer processing capabilities
- Increases deployment flexibility and improves thermal management in SWaP-constrained spaces

VIDEO DISTRIBUTION SYSTEM

The video distribution system, or VDS, is the central "brain" of the video system. These video gateways, switches, or multiplexers manipulate the image

Trust Coilcraft CPS for the battle-proven magnetics you need when your mission is on the line

Our critical-grade RF and Power inductors are built to withstand the most punishing conditions imaginable, including extreme temperatures, thermal shock, G-forces, EMI and vibration.

- Tin-lead (Sn-Pb) terminations for the best possible board adhesion. RoHS terminations also available.
- Extended temperature ranges (-55°C to +200°C)
- Vibration testing to 80 G / shock testing to 1000 G

We also offer comprehensive product testing and validation services in accordance with MIL-STD-981, as well as custom screening to your electrical and physical specifications.

Learn more about how our battle-tested components will keep your mission on target. Call or visit us online today!

Coilcraft CPS
CRITICAL PRODUCTS & SERVICES

800.981.0363 847.639.6400 www.coilcraft-cps.com and information feeds before outputting the results to the mission displays.

To increase situational awareness, many video-distribution solutions provide a range of image configuration options. The most sophisticated solutions can simultaneously deliver multiple video streams to multiple displays. They also allow warfighters to quickly and easily manipulate views with the touch of a button. Key viewing features to look for include:

- Real-time views with zoom capabilities
- Simultaneous camera views using picture-in-picture technology and window overlays
- Video streaming and blending capabilities

The VDS must also provide high-performance processing and extensive I/O to support connections to multiple sensors and displays in an extremely SWaP-friendly form factor. And it must require minimal wiring. For example, it's mandatory that all video streams sent to a single display travel over a single connection between the distribution technology and the display.

Finally, the VDS must be able to synchronize the video streams from all video cameras and other sensors; add diagnostic data, text, and graphics to put the visuals in context for warfighters; and then deliver it to mission displays with extremely low latency.

MISSION DISPLAYS

Mission displays must be ruggedized, lightweight, high-resolution, and extremely responsive. There have been numerous advances in display technologies over the last decade, and each one plays an important role in helping warfighters quickly and easily assimilate mission-critical information:

- High-brightness LED backlights make displays easy to read in low-light conditions.
- Fully bonded displays are highly reliable, reduce reflections, and provide superior clarity and contrast compared to legacy non-bonded displays.
- Projected capacitive (PCAP) touchscreens support intuitive, multitouch gestures, don't require a stylus, and can be used while wearing gloves. They're also lighter weight and more durable than resistive touchscreens.
- LCD technology provides more viewing angles, allowing warfighters to easily see the information on the screen from the side as well as the front.
- Anti-reflection and anti-glare coatings provide better visibility in bright light conditions.

END-TO-END VIDEO SOLUTION FOR ENHANCED SITUATIONAL AWARENESS

Curtiss-Wright's end-to-end video solutions help improve situational awareness today and pave the way to even more enhanced situational aware-

3. Curtiss-Wright's ground vehicle display units (GVDUs), designed specifically for ground vehicle system applications, feature multipoint PCAP touchscreen technology for ease of use.

4. The RVG-MS1 Multi-Sensor Rugged Video Gateway is a multi-input solution for processing sensor inputs for distribution to a vehicle display. Its graphical user interface (GUI) enables the operator to select the operational view that best benefits the mission.

ness solutions of tomorrow. To maximize flexibility and simplify evolution, the company takes a building-block approach to video-system design, providing ruggedized mission displays, video distribution systems and video recorders, as well as complete video management systems.

Examples of situational awareness system building blocks include:

- Ground vehicle display units (GVDUs) *(Fig. 3)* that feature ruggedized PCAP touchscreen displays designed for the unique requirements of ground vehicles, where there's less light, more sand and water in the air, and greater need to operate the display while wearing gloves.
- The RVG-MS1 Multi-Sensor Rugged Video Gateway *(Fig. 4)*, which provides 25 inputs and 20 outputs in a unit that weighs only 3.25 kg (7.17 lbs.) and requires only 80 W of maximum power. This SWaP-optimized video gateway supports single, dual, triple, and quad views in a variety of layouts.
- Advanced video display units (AVDUs) and single video display units (SVDUs) that are ideal for large- and medium-sized video solutions on airborne platforms.
- The VRDV7000, a lightweight, small-form-factor, dual-channel HD video recorder that allows warfighters to easily review recently captured video footage to verify situations.

If critical computer processing power needs to be added to these video solutions, small-form-factor, modular mission computers like the ruggedized Parvus DuraCOR 8041 are ideal.

These ruggedized video solutions are designed and built to withstand extreme temperatures, shock, vibration, sand, water, and other challenging environmental conditions to ensure long-term, reliable operation on any terrain and in any weather conditions. They meet key industry standards, including:

- MIL-STD-461 for radiated emissions and electromagnetic compatibility
- MIL-STE-1275E for power and electrostatic discharge
- MIL-STD-810 for environmental engineering design and testing
- DO160-G for environmental and EMC/EMI

In addition, these solutions address the key technical challenges associated with reducing latency in individual solution components and across the end-to-end solution. ed

REFERENCES

"Projected Capacitive (PCAP) Touch Screens In Defense and Aerospace" *https://www.curtisswrightds.com/infocenter/white-papers/projected-capacitive-pcap-touch-screens-defense-aerospace.html*

"Reducing Latency in Ground Vehicle Video Systems" *https://www.curtisswrightds.com/infocenter/white-papers/reducing-latency-in-ground-vehicle-video-systems.html*

"Why Smart Video Displays May Not Be the Cleverest Choice" *https://www.curtisswrightds.com/infocenter/white-papers/why-smart-video-displays-may-not-be-the-cleverest-choice.html*

RICHARD JAENICKE | Director of Marketing, Safety- and Security-Critical Products, Green Hills Software

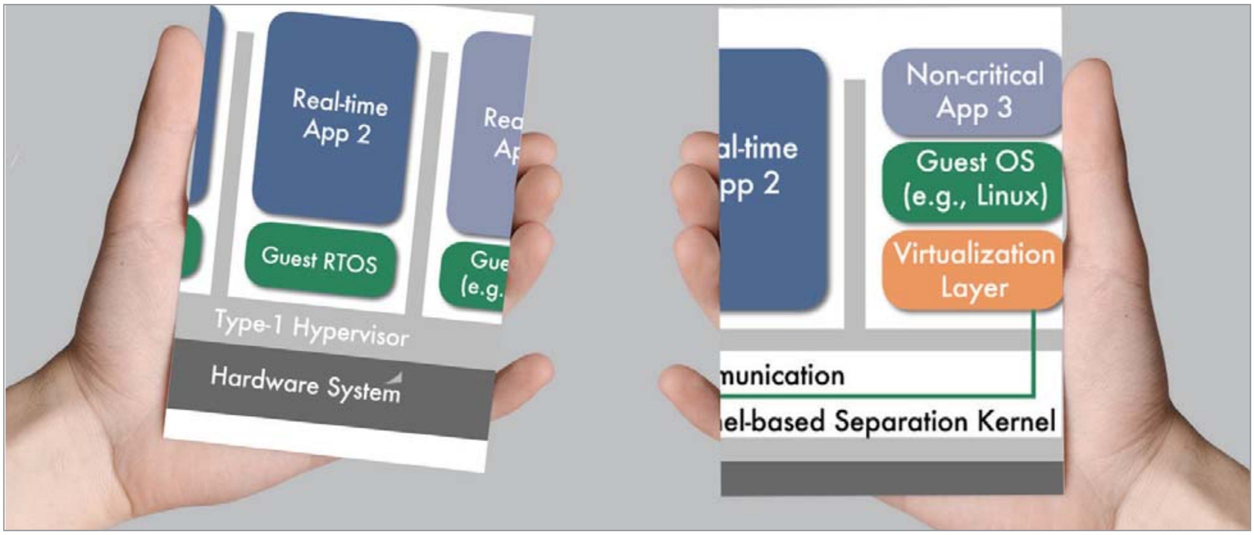

What's the Difference Between an Embedded Hypervisor and Separation Microkernel with Virtualization?

Both hypervisors and separation microkernels with a virtualization layer support multiple guest OSes, but one focuses more on virtualization features while the other targets security and real-time performance.

Hypervisors are used widely in enterprise servers, and the number of offerings in the embedded space continues to ramp up. Hypervisors are used for virtualization and provide some level of isolation, but they're not the only option. Microkernels originated the embedded world, and a separation microkernel is specially designed for isolation and security. The same virtual-machine technology used in hypervisors can be added to a microkernel to provide a virtualization solution when needed.

Both solutions provide the ability to run multiple operating systems (OSes) in a virtualized environment, including mixing OS types. Over time, the two technologies have grown closer together, but some significant differences still exist in terms of latency, determinism, and security.

A hypervisor, also known as a virtual-machine monitor, is software designed to create and run virtual machines (VMs), each of which abstracts the hardware platform and runs a guest OS. The hypervisor is responsible for isolating each VM so that actions by one VM can't compromise another *(Fig. 1)*. In keeping with their enterprise server origins, hypervisors and VMs are heavyweight constructs that typically consume a great deal of resources in terms of the code base, memory footprint, and execution latency.

Hypervisors can be categorized as Type 1 or Type 2. Type 1 hypervisors run natively on the host hardware (i.e., "bare metal"), whereas Type 2 hypervisors run on top of a host OS. Type 1 hypervisors generally provide higher performance by eliminating one layer of software. That assumes, however, that all or most applications require virtualization. If only a small percentage of applications require virtualization, the overall system performance can be higher with a Type 2 hypervisor.

For example, if most applications run on a real-time OS (RTOS) and only one or two other applications need to run on Linux or Microsoft Windows, then the real-time applications can run on the host OS and aren't burdened by going through a hypervisor. Only the applications that require virtualization need to pay the virtualization performance penalty.

A microkernel is an OS where the only essential services are implemented in kernel mode, and all other services are implemented mainly in user space, including device drivers, file systems, networking stacks, and virtualization. This leads to increased modularity, flexibility, and robustness, as well as a smaller, trusted compute base (TCB).

One specific type of microkernel is a separation kernel, which allocates all exported resources under its control into partitions, and those partitions are isolated except for explicitly allowed information flows. Separation kernels that are designed for the highest security meet the Separation Kernel Protection Profile (SKPP) defined by the U.S. National Security Agency (NSA), which was created for the most hostile threat environments.

A separation microkernel that includes a virtualization layer has some similarities to a Type 2 hypervisor in that the virtualization layer runs on top of a host OS and can be applied selectively to only the applications that require virtualization. However, it differs in that the isolation function is provided by the separation microkernel (in this case, the host OS), and isolation is enforced even between different instances of the virtualization layer *(Fig. 2)*.

VIRTUALIZATION FEATURES

Both hypervisors and microkernels with a virtualization layer make use of hardware acceleration for full virtualization. On Intel processors, for example, that includes Intel VT-x, EPT, and Intel VT-d. Intel VT-x provides hardware instructions for entering and exiting a virtual execution mode where the guest OS sees itself as running with full privilege, while the host OS remains protected. Extended Page Tables (EPT) provide virtualization of the page tables that hold the mappings from physical to virtual memory. Intel VT-d provides a hardware assist for remapping direct-memory-access (DMA) transfers and device-generated interrupts. Even with those hardware-acceleration technologies and similar ones for other processors, a fair amount of virtualization is still done in software.

Because of their enterprise heritage, hypervisors sometimes provide a higher limit on the number of VMs per host and broader support of virtualization management features. Those advanced management capabilities might include dynamic resource allocation, failover, and live migration. Such capabilities can vary greatly by implementation.

LATENCY AND DETERMINISM

Before common microprocessors supported hardware acceleration of some virtualization functionality, full virtualization solutions were notoriously low performance. Although modern microprocessors provide a range of virtualization support, a fair amount of I/O virtualization still needs to be handled in software, such as device emulation, bus emulation, and interrupt emulation and routing.

In general, virtualization has lower latency and higher performance with a Type 1 hypervisor because it doesn't have to also go through a host OS. With a Type 1 hypervisor, however, every application must pay that virtualization penalty. That includes real-time and safety-critical applications.

In contrast, a microkernel with a virtualization layer can have real-time applications running directly on the host OS while the non-real-time applications run on top of a virtualization layer. The virtualization performance penalty is larger for this approach, but it's only paid by the non-real-time applications. In that manner, all hard real-time and safety-critical applications get the low latency and determinism inherent in the host RTOS.

Some hypervisors attempt to further reduce latency by giving an application the ability to run without a guest OS in "bare-metal" mode, but that's a misnomer. Even if the application doesn't run on a guest OS, it's still running on top of the hypervisor. If the goal is to run a real-time, safety-critical, or security-critical application with lower latency, running without an OS isn't the way to get there. Almost all of those types of target applications require tasking services, semaphores, or message passing—none of which are available in bare-metal mode.

SECURITY

Security can be the area of most significant difference between a hypervisor and a separation microkernel with a virtualization layer. Although security can be the primary reason for considering a hypervisor, it's a myth that hypervisors are inherently secure just because they utilize hardware features to enforce virtual address spaces and virtual I/O to isolate VMs.

It would be more accurate to say that the primary security protection provided by most hypervisors is simply whatever the underlying hardware supports through the MMU and IOMMU. First, that hardware security only covers some of what's needed to isolate VMs, and second, that hardware security only helps isolate VMs and does nothing to make the hypervisor itself more secure. Because hypervisors run below the guest operating system, a compromised hypervisor isn't detectable by the VM. Such exploits even have a catchy name: hyperjacking.

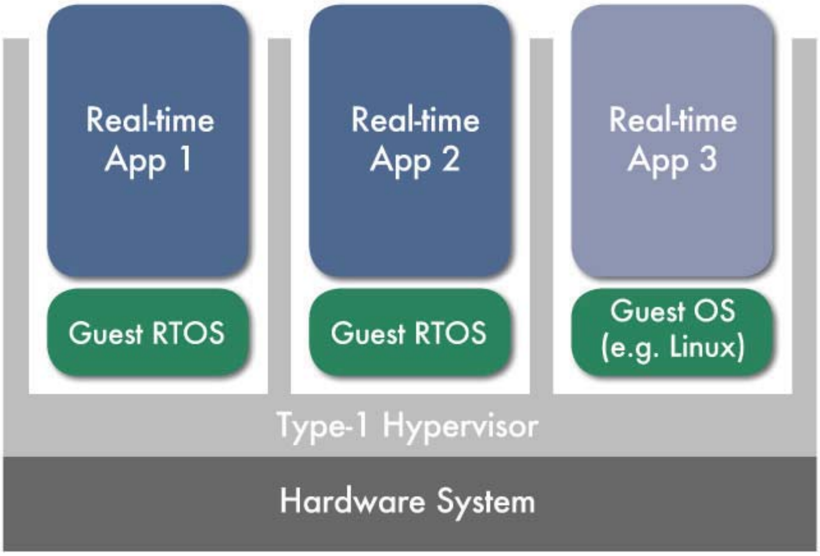

A Type 1 hypervisor runs each application in a VM with its guest OS while not requiring a host OS.

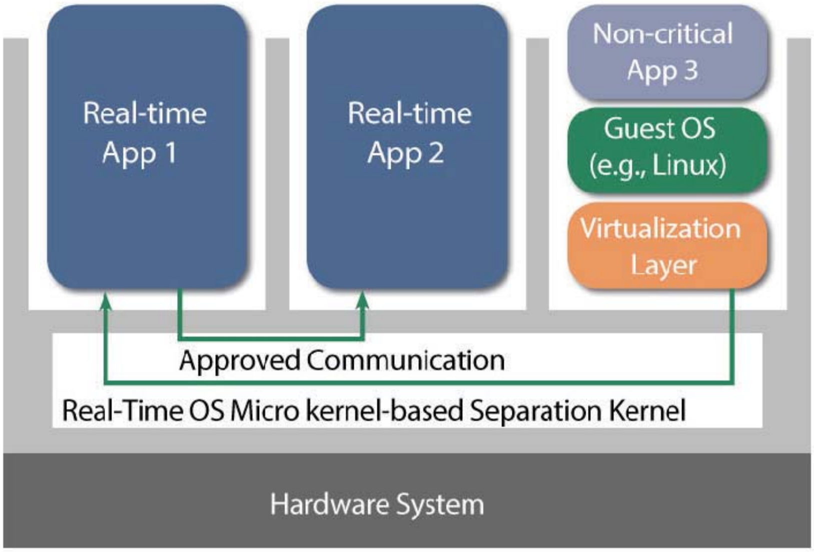

A separation microkernel runs each application in a separate partition, and only those that need a different OS run inside a VM with a guest OS on top of a virtualization layer.

The likelihood of vulnerabilities generally is proportional to the size of the attack surface. Hypervisors have the virtualization software running in kernel mode, making it part of the trusted compute base (TCB), and that virtualization code can be huge. Almost every call from the guest OS to the kernel needs to be trapped, examined, and determined if the guest OS is permitted such access.

For a virtualization solution to be efficient, it needs to virtualize sequences of instructions instead of single instructions. Such look-ahead capability is just one example of increasing the already large code base of a hypervisor in pursuit of minimizing the virtualization performance penalty.

By definition, a microkernel is built with as little as possible running in kernel mode, so the virtualization layer isn't part of the TCB. That doesn't mean that the VMs are unprotected. Rather, the separation microkernel provides the isolation for every partition, whether virtualized or not. Even if a vulnerability occurs in the virtualization layer, such a vulnerability can't spread even to another instance of the virtualization layer for another VM. In that way, a separation microkernel prevents hyperjacking.

Microkernels using separation kernel technology can have the highest levels of security and isolation. The proof of that security level is certification to the SKPP published by the NSA or similar security standards such as Common Criteria EAL6. Some hypervisors include some separation kernel principles to improve security, but no commercial hypervisor has been certified to the SKPP or Common Criteria EAL6.

An example of a separation microkernel with an optional virtualization layer is the INTEGRITY-178 tuMP RTOS from Green Hills Software. INTEGRITY-178 is a safety- and security-critical OS that has been certified by the National Information Assurance Partnership (NIAP) to meet the SKPP, High Robustness, as well as Common Criteria EAL6+. The INTEGRITY-178 tuMP RTOS extends that pedigree to multicore processors.

Its optional virtualization layer runs within partitions isolated by the separation kernel and allows one or more guest OSes to simultaneously run alongside safety- and security-critical functions on the same microprocessor. The result is a system that optimizes for safety, security, and real-time performance while enabling applications that require additional general-purpose or legacy OS support. ed

REFERENCES

1. John Rushby, "The Design and Verification of Secure Systems," Eighth ACM Symposium on Operating System Principles, pp. 12-21, Asilomar, CA, December 1981.
2. Information Assurance Directorate, National Security Agency, "U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness," Version 1.03, June 2007.

PHILIPPE VEREECKEN | Scientific Director, Electrochemical Storage, imec

Solid-State Battery Tech for Electric Cars: Key to Greater Autonomy

With solid-state batteries, electric vehicles should be able to achieve an autonomy (driving range) matching—and eventually surpassing—that of cars with an internal combustion engine.

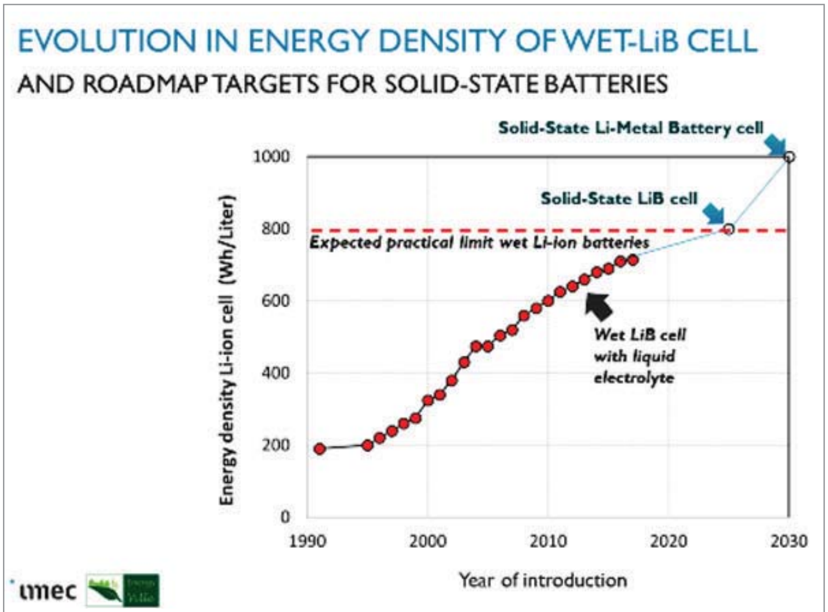

1. The energy density of the Li-ion battery (LiB) cell has more than tripled since its market introduction by Sony in 1991. Continuous improvements in LiB components with $LiCoO_2$-graphite chemistry resulted in an average increase of 25 Wh/l per year from 1995 to 2010. Introduction of new active cathode materials, such as the NiCoAl-based and NiMnCo-based lithium-metal oxides (NCA and NMC), and the gradual addition of silicon to the graphite anode, have maintained the energy density increase ever since. However, it's expected that with the materials known today, we will reach a practical limit for wet LiBs of around 800 Wh/l. Solid-state battery technology will be needed to break through this barrier and achieve an energy density of 1000 Wh/l—and more.

The great thing about solid electrolytes is that they allow for the integration of different active materials (such as a lithium metal anode) and cell architectures (such as bipolar arrangements), making for a higher energy density at cell and/or battery-pack level. In parallel with these all-solid-state batteries, developments such as "smart" battery cells with sensors will increase the amount of usable energy in the battery pack even further.

In this article, Professor Philippe Vereecken (imec, KU Leuven, EnergyVille) highlights the importance of these new technologies for future electric vehicles (EVs).

TODAY'S Li-ION BATTERY CELLS ARE RUNNING OUT OF STEAM

For electric cars to really take off, they ideally should match—or even exceed—the driving range of vehicles with an internal combustion engine. The key to achieving this lies in the battery itself, more specifically in the use of solid-state battery-cell technology.

At the heart of today's EVs are lithium-ion battery (LiB) cells that contain liquid electrolytes. The best-in-class of these "wet" LiB cells have an energy density of around 700 watt-hours per liter (Wh/l), accommodating a maximum driving range of about 500 km. Yet, their energy density is expected to stagnate at around 800 Wh/l due to the characteristics of their active materials *(Fig. 1)*.

A higher energy density can be expected from solid-state batteries, which contain a solid electrolyte instead

of a liquid one. In combination with new battery-pack and battery-module developments, the autonomy of electric cars could significantly be extended.

1000-Wh/l BATTERY CELLS ENABLE MORE AUTONOMY

Today, the maximum driving range of an electric vehicle is determined by the amount of energy that's contained in the individual LiB cells of the car's battery pack. Cells are connected in parallel and series arrangements to provide the high currents and high voltage needed to power the electric engine.

To achieve a driving range of 700 km, cells are required with an energy density as high as 1000 Wh/l (or 500 Wh/kg). Since today's LiB cells can "only" deliver 700 Wh/l (or 230 Wh/kg), a significant boost in energy density is needed.

Battery-cell roadmaps foresee that cells of 1000 Wh/l should become available in 2030—in the form of solid-state lithium-metal batteries. We will come back to the potential of those solid-state batteries, but let's first discuss an alternative approach to increase the amount of energy that can be squeezed out of an electric car's battery pack—an approach leveraging so-called "smart" battery cells.

"SMART" BATTERY CELLS: A COMPLEMENTARY APPROACH

Enlarging the energy density at cell level is one prerequisite to increasing the driving range of the electric car. But, next to the many individual cells in large battery packs, a car battery module also contains electronics and sensors to manage that battery's usage. For example, to ensure a long (enough) lifetime, the battery management system (or BMS) will typically only use part of the cells' energy to avoid damage to the cell chemistry.

In other words, the energy that can actually be used by an electric car for driving might only be 60-80% of the overall battery's capacity (depending on the type of electric car).

Improving the battery's energy management can be another factor to enhancing the electric vehicles' driving range. This could be realized by using "smart" battery cells, with micro-sensors built into the cell to better monitor its state of charge and state of health.

At EnergyVille—an association of the Flemish research institutes KU Leuven, VITO, imec, and UHasselt in the field of sustainable energy and intelligent energy systems—research into smart battery cells with multi-array sensors and integrated electronics for communication with the BMS is currently under way.

FIRST-GENERATION SOLID-STATE BATTERIES: CELLS WITH SOLID-STATE ELECTROLYTE

The first generation(s) of solid-state LiB cells will have solid electrolytes instead of liquid electrolytes. However, at cell level, they will not yet feature a higher energy density than LiBs containing a liquid electrolyte, because they will initially contain similar types of active electrode materials.

Why, then, develop solid-state batteries at all, you may ask? Well, for one, solid-state batteries allow for different, more compact arrangements in the battery pack. They can be built in bipolar arrangements providing higher voltages at cell level. This, in turn, simplifies connecting the cells and creates extra space in the battery pack to include more cells.

Solid-state LiBs will also be more intrinsically safe. Thus, less safety monitoring electronics are needed in the periphery of the battery module. Finally, solid-state batteries can have a larger voltage window than wet LiBs, which means that the risk of damaging the cell during charge and discharge is reduced and a larger part of the cell energy becomes available for use.

Combining all of those factors, the amount of energy that's available in the battery pack will be higher for solid-state LiBs—even if the energy density of the first generation(s) of solid-state cells will be the same, or perhaps slightly lower, than the energy density of wet LiB cells.

The first electric cars equipped with these first-generation solid-state batteries are expected to launch somewhere in the middle of this decade.

IN PURSUIT OF THE HOLY GRAIL: SOLID-STATE LITHIUM-METAL BATTERIES

As described above, simply replacing a liquid electrolyte in a LiB with a solid electrolyte doesn't make for an increase in energy density. On the contrary, ceramic solid electrolytes in powder form are likely to take up more space—and weigh more—than their liquid equivalents in current LiBs. Hence, solid electrolytes would somewhat lower the energy density of a cell with the same active electrode materials.

The benefit lies in the fact that some solid electrolytes can provide a larger electrochemical window. In other words, they remain stable also at very high voltages (where current liquid electrolytes do not). This means that higher-voltage cathode materials can be used than those employed in 3.6- to 3.8-V LiBs today, and a higher cell voltage leads to an increase in cell energy, provided the lithium-ion capacity remains the same.

For example, $LiMn_{1.5}Ni_{0.5}O_2$, or LMNO, has an electrode potential of 4.7 V and can't be used with liquid electrolytes that react at these high voltages. Solid electrolytes such as lithium lanthanum zirconium oxide (LLZO), however, are stable up to 5 V and are thus compatible with these high-voltage cathodes. As such, solid-state LiBs can break the 800-Wh/l barrier.

The real game-changer is the use of metallic lithium as an anode. Lithium metal has the highest energy density as active anode material and provides the highest cell voltage from the anode contribution in the cell. To reach the much desired 1000-Wh/l energy density—and to even go beyond—all-solid-state lithium-metal battery (LMB) cells are therefore targeted.

So far, the use of lithium metal didn't work because charging of the battery leads to the formation of lithium-metal needles (or dendrites). These short the battery internally, resulting in thermal runaway and explosions. It was Dr. Akira Yoshino *(Fig. 2)*, one of last year's Nobel prize winners in chemistry, who came up with graphite as a safe solution, enabling the commercialization of the current wet LiBs in the early 1990s.

2. Dr. Akira Yoshino was awarded the Nobel Prize in chemistry in 2019. *(Source: Wikipedia)*

Inserting lithium-ions in between the layers of graphite was a safe compromise, but it came at the expense of energy density and cell voltage. Several solid-state electrolytes are stable against metallic lithium; thus, the use of thin lithium as an anode becomes a possibility. In fact, lithium-metal batteries with solid polymer electrolytes are already commercially available *(Fig. 3 on page 30)*.

However, these batteries only operate at temperatures of 70°C and therefore aren't suitable for an electric family car. The reason for this is lithium-ion conductivity is too low for these solid polymer electrolytes. As a result, for many years, solid-state battery research has focused on finding solid electrolytes with a high-enough ionic conductivity.

EXPLORING THE OPTIONS FOR A MATCHING SOLID ELECTROLYTE

In the last few years, several good

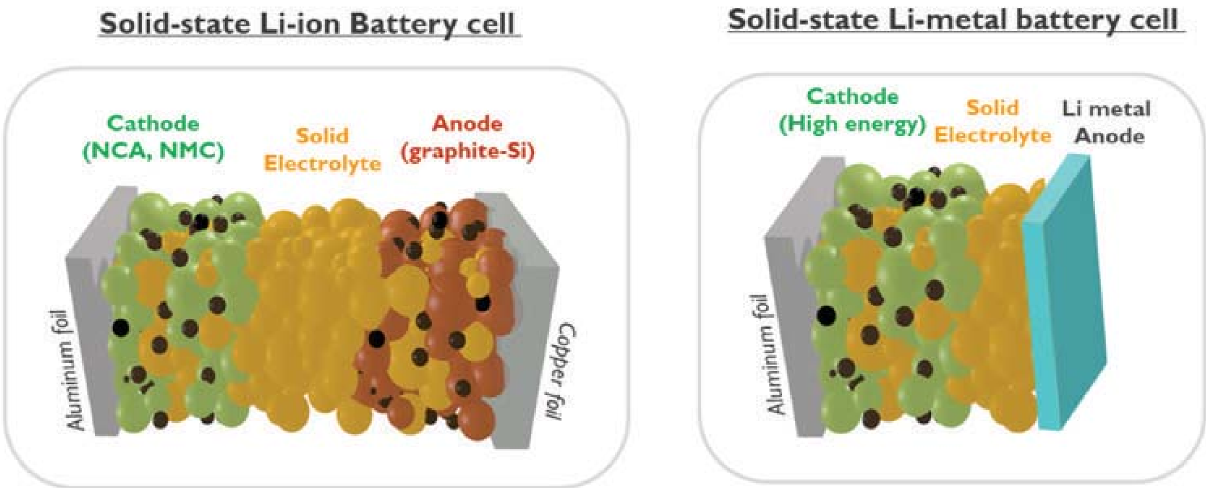

3. Shown is a schematic of a solid-state lithium-ion battery with graphite-silicon electrode and a solid-state lithium-metal battery with thin lithium-metal anode. *(Source: Xubin Chen, Philippe Vereecken, Fanny Bardé)*

options have emerged with lithium-ion conductivities matching, or even surpassing, the ion conductivity of liquid electrolytes that are currently used for lithium-ion batteries. Toyota is the current champion with a sulfidic inorganic solid electrolyte with Li-ion conductivity about 3X that of liquid electrolytes. Another contender is LLZO, a garnet-type oxidic inorganic electrolyte that, even though it has somewhat lower lithium-ion conductivity than liquid electrolytes, is of interest because of its large electrochemical window. On the downside, these inorganic electrolytes are extremely sensitive to moisture, which makes the assembly of the solid-state cells difficult (and potentially costly).

In contrast, imec has developed a nanocomposite electrolyte that's fabricated from liquid and turned solid once it's inside the cell *(Fig. 4)*. This makes the material compatible with current lithium-ion cell-fabrication processes, and thus potentially low cost.

TAKING THE NEXT STEP: OVERCOMING CELL-ASSEMBLY CHALLENGES

Now that several options are available, solid-state battery research has shifted its focus toward the assembly of the cells and the integration of all components into a functional cell. One of the issues is the reactivity of the materials forming resistive or blocking layers between the functional components, severely inhibiting the cell's operation. Special thin-film coating (or so-called "artificial interphase coatings") are being developed to prevent these reactions from happening.

At EnergyVille, upscale and cost-effective methods are developed for the deposition of these sub-nanometer thin coatings inside the thick battery electrodes. These nanoscale coatings will be instrumental not only for the solid-state lithium-metal batteries under development, but also for the extendibility of LiBs with a liquid electrolyte. Indeed, much of the solid electrolyte know-how can be applied to wet LiBs as well.

Who knows, the solution may lie in a hybrid approach after all.

4. Imec's solid-state electrolyte is a new material for use in making next-generation LiBs for electric vehicles (artist impression). *(Source: imec)*

Advertisers Index

FREE WEBINAR

May 21, 2020 | 2:00pm EDT

Registration URL: https://www.electronicdesign.com/resources/webcasts/webinar/21128288/heat-treatment-through-multiphysics-simulations

Heat Treatment Through Multiphysics Simulations

In this webinar, you will learn about solving challenges in heat treating through multiphysics simulations.

Heat treatment engineers experience numerous challenges to obtain the correct temperature history of their parts. Constructing a new facility, or modifying an existing facility, introduces significant challenges to engineers attempting to maintain a specified temperature history for numerous parts. Thermal gradients within the furnace increase scrap rate and reduce profitability.

Multiphysics simulation provides engineers with the information needed to solve these difficult challenges. These simulations help engineers visualize thermal gradients within the furnace and within the parts, as well as phase changes. In addition, simulations provide engineers with an understanding of how the part temperatures change in time.

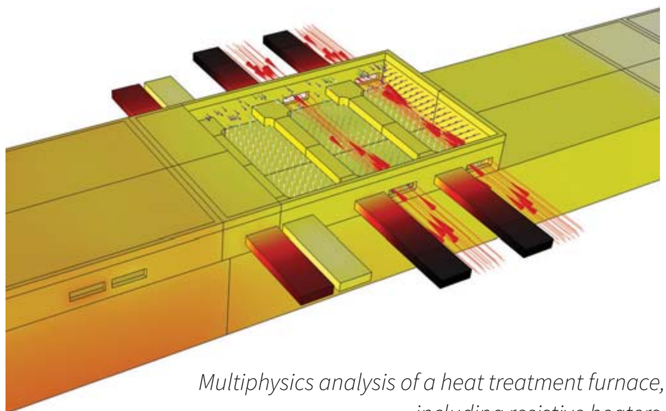

Multiphysics analysis of a heat treatment furnace, including resistive heaters.

In this webinar, AltaSim Technologies will share their experience in helping engineers increase productivity of their heat treatment facilities using multiphysics simulations. A demonstration using the COMSOL Multiphysics® software will show a simulation of the temperature within parts in a furnace. This simulation includes the metallurgical phases that develop. You can ask questions at the end of the webinar during the Q&A session.

SPEAKER: Kyle Koppenhoefer, PhD, Principal, AltaSim Technologies

Kyle Koppenhoefer has been one of the principals and leaders at AltaSim Technologies for 15 years. He works with customers to identify how computational analysis can be used to further develop their products and manufacturing processes. Prior to cofounding AltaSim, Kyle worked for the Department of Defense and the Edison Welding Institute. He holds a PhD in civil engineering from the University of Illinois.

SPEAKER: Mats Danielsson, PhD, Technical Product Manager, COMSOL

Mats Danielsson joined COMSOL in 2017 as a developer specializing in structural mechanics and mechanics of materials. He earned his PhD in mechanical engineering from MIT in 2003. Before joining COMSOL, he worked in the automotive industry and applied materials research.

Hosted with ElectronicDesign

COMSOL