AUTOMATION 2022

VOLUME 6

IIoT & Industry 4.0

- Powering Remote Wireless Devices
- AI Inference at the Rugged Edge
- Digital Transformation and TSN
- Open Process Automation
- Next-Gen Biotech Manufacturing

AUTOMATION.COM | ISA A subsidiary of the International Society of Automation

Introduction

Introduction

As companies strive to achieve the benefits of Industry 4.0, they rely on development of smart manufacturing technologies. Many of the advances are occurring around data—how we can capture it, share it, secure it, and make it maximally useful.

In this edition of AUTOMATION 2022, you'll gain insights into ensuring continuous data gathering through IIoT devices supported by maintenance-free batteries that can power wireless devices for decades. You'll learn how time-sensitive networking (TSN) is enabling data gathering and data transfer to support the digital transformation of industry. You'll see how AI inference at the edge processes data at speeds high enough to create autonomous systems, and how open process automation can bring the best of information technology (IT) and operational technology (OT) together. These and other advances in data analytics and operations come together in a case study on Genentech's newest biotechnology manufacturing facility: a Level 5 Adaptive Plant that incorporates advanced manufacturing and environmental sustainability. This issue showcases the importance of collaboration—between IT and OT, and industry and technology providers, as well as the data at the heart of smart systems.

Renee Bassett
Chief Editor
Automation.com

About AUTOMATION 2022

The AUTOMATION 2022 ebook series covers Industry 4.0, smart manufacturing, IIoT, cybersecurity, connectivity, machine and process control and more for industrial automation, process control and instrumentation professionals. To subscribe to ebooks and newsletters, visit: www.automation.com/newslettersubscription.

AUTOMATION 2022 is published six times per year (February, April, June, August, October, December) by Automation.com, a subsidiary of ISA—the International Society of Automation. To advertise, visit: www.automation.com/en-us/advertise.

AUTOMATION.COM
groups/68581
automationdotcom
automation_com

International Society of Automation
Setting the Standard for Automation™
company/internationalsocietyofautomation
InternationalSocietyOfAutomation
ISA_Interchange

Renee Bassett, Chief Editor
rbassett@automation.com

Chris Nelson, Advertising Sales Rep
chris@automation.com

Richard T. Simpson, Advertising Sales Rep
rsimpson@automation.com

Gina DiFrancesco, Advertising Sales Rep
GDiFrancesco@automation.com

Table of Contents

AUTOMATION 2022 VOL 6
IIOT & INDUSTRY 4.0

IIoT devices run longer on Tadiran batteries.

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

Remote wireless devices connected to the Industrial Internet of Things (IIoT) run on Tadiran bobbin-type $LiSOCl_2$ batteries.

Our batteries offer a winning combination: a patented hybrid layer capacitor (HLC) that delivers the high pulses required for two-way wireless communications; the widest temperature range of all; and the lowest self-discharge rate (0.7% per year), enabling our cells to last up to 4 times longer than the competition.

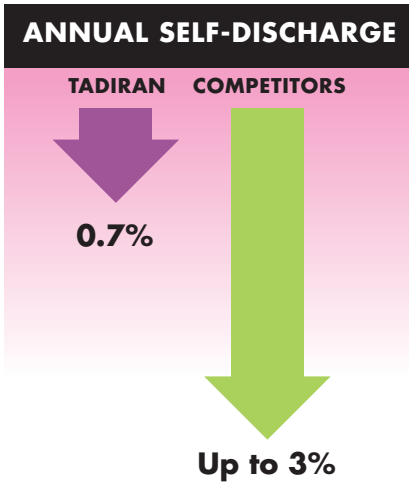

Looking to have your remote wireless device complete a 40-year marathon? Then team up with Tadiran batteries that last a lifetime.

TADIRAN BATTERIES

* Tadiran $LiSOCL_2$ batteries feature the lowest annual self-discharge rate of any competitive battery, less than 1% per year, enabling these batteries to operate over 40 years depending on device operating usage. However, this is not an expressed or implied warranty, as each application differs in terms of annual energy consumption and/or operating environment.

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

www.tadiranbat.com

Powering Remote Wireless Devices to Run a 40-year Marathon

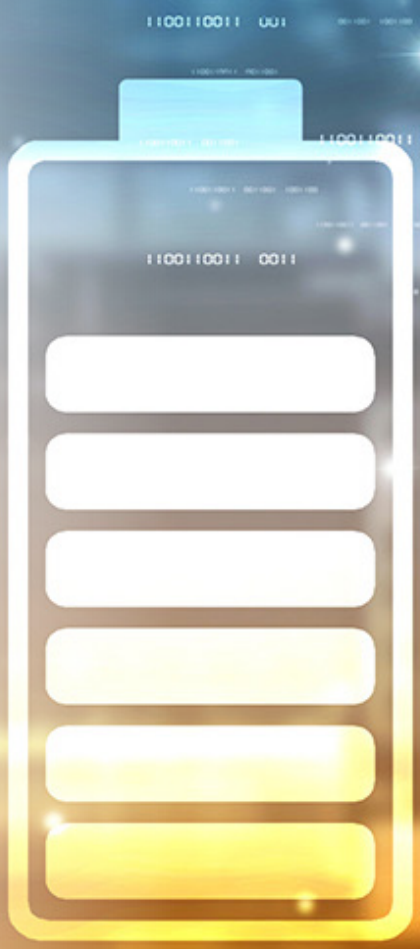

Low-power remote wireless devices can operate for up to 40 years on a single battery while powering two-way communications.

By Sol Jacobs, Tadiran Batteries

The Industrial Internet of Things (IIoT) is exploding, resulting in the dramatic growth of remote wireless devices being powered by primary (non-rechargeable) lithium batteries. Often, these devices need to operate maintenance-free for decades without having to replace the batteries.

To achieve such extended battery life, energy losses must be minimized. This is especially true for low-power devices that draw average current measurable in micro-amps while operating mainly in an energy-saving “standby” state, fully awakening only to sample or transmit data.

Battery-powered wireless devices deployed in extreme environments and hard-to-access locations are generally ill-suited for consumer grade alkaline cells that are extremely short-lived, primarily due to a high self-discharge rate of up to 60 percent per year. Lithium chemistries can last much longer.

Lithium features a high intrinsic negative potential that exceeds all other metals, functioning within an operating current voltage (OCV) ranging from 2.7 V to 3.6 V. Lithium chemistries are also non-aqueous, able to endure extreme temperatures with less risk of freezing. Chemistries such as iron disulfate ($LiFeS_2$) and lithium manganese dioxide ($LiMNO_2$) can deliver medium to high rates of energy discharge. However, there is a trade-off, as these chemistries feature annual self-discharge rates that are far lower than alkaline but far greater than lithium thionyl chloride ($LiSOCl_2$) (Table 1).

Primary Cell	$LiSOCL_2$ Bobbin-type with Hybrid Layer Capacitor	$LiSOCL_2$ Bobbin-type	Li Metal Oxide Modified for high capacity	Li Metal Oxide Modified for high power	$LiFeS_2$ Lithium Iron Disulfate (AA-size)	$LiMnO_2$ Lithium Manganese Oxide
Energy Density (Wh/Kg)	700	730	370	185	335	330
Power	Very High	Low	Very High	Very High	High	Moderate
Voltage	3.6 to 3.9 V	3.6 V	4.1 V	4.1 V	1.5 V	3.0 V
Pulse Amplitude	Excellent	Small	High	Very High	Moderate	Moderate
Passivation	None	High	Very Low	None	Fair	Moderate
Performance at Elevated Temp.	Excellent	Fair	Excellent	Excellent	Moderate	Fair
Performance at Low Temp.	Excellent	Fair	Moderate	Excellent	Moderate	Poor
Operating life	Excellent	Excellent	Excellent	Excellent	Moderate	Fair
Self-Discharge Rate	Very Low	Very Low	Very Low	Very Low	Moderate	High
Operating Temp.	-55°C to 85°C, can be extended to 105°C for a short time	-80°C to 125°C	-45°C to 85°C	-45°C to 85°C	-20°C to 60°C	0°C to 60°C

Table 1. Bobbin-type $LiSOCl_2$ batteries are preferred for use in remote wireless applications. These cells deliver higher capacity and energy density, up to a 40-year operating life, and the widest possible temperature range, which is ideal for hard-to-access locations and extreme environments.

Not all batteries can handle a 40-year marathon

Among all commercially available lithium chemistries, bobbin-type $LiSOCl_2$ batteries are preferred for long-term deployments in remote locations and extreme environments due to their ruggedness and ultra-long-life potential. Bobbin-type $LiSOCl_2$ cells deliver the highest capacity and highest energy density of any lithium chemistry, along with an extremely low annual self-discharge rate (less than 1 percent per year), enabling certain low-power devices to operate for up to 40 years on their original battery. Bobbin-type $LiSOCl_2$ batteries also offer the widest possible temperature range (-80°C to 125°C), along with a glass-to-metal hermetic seal that resists battery leakage. Common applications for $LiSOCl_2$ cells include AMR/AMI metering, machine-to-machine (M2M), supervisory control and data acquisition (SCADA), tank-level monitoring, asset tracking, and environmental sensors, to name a few (Figure 1).

Figure 1. Researchers from Cardiff University studying water channels beneath glaciers use the Cryoegg, which monitors temperature, pressure, and electrical connectivity by transmitting data underwater via radio waves. Bobbin-type $LiSOCl_2$ cells were specified for their high capacity, high energy density, extended temperature range, and high pulse capabilities. *Courtesy: Cardiff University.*

$LiSOCl_2$ batteries can also be manufactured with a spiral wound construction that permits a higher rate of energy flow but also results in a higher annual self-discharge rate that limits their potential operating life.

Harnessing the passivation effect

All batteries experience some amount of annual self-discharge as chemical reactions occur even when the battery is disconnected or not in use. Self-discharge can be significantly reduced by harnessing the passivation effect, which is unique to $LiSOCl_2$ chemistry.

Passivation occurs when a thin film of lithium chloride (LiCl) forms on the surface of the lithium anode to impede the chemical reactions that result in battery self-discharge. Whenever a load is placed on the cell, the passivation layer causes high initial resistance, resulting in a temporary drop in cell voltage until the discharge reaction slowly dissolves the passivation layer—a process that repeats each time a load is applied.

The ability of a bobbin-type $LiSOCl_2$ cell to harness the passivation effect can vary significantly based on the quality of the raw materials and the method by which the battery is manufactured.

Passivation levels can be affected by several variables, including the current capacity of the cell, the length of storage, storage temperature, discharge temperature, and prior discharge conditions, as partially discharging a cell and then removing the load can increase the amount of passivation relative to when the cell was new. While passivation can serve to reduce a battery's self-discharge rate, too much of it can cause energy flow to be blocked.

The ability of a bobbin-type $LiSOCl_2$ cell to harness the passivation effect can vary significantly based on the quality of the raw materials and the method by which the battery is manufactured. For example, lower quality bobbin-type $LiSOCl_2$ batteries can lose up to 3 percent of their original capacity annually due to self-discharge, thus exhausting 30 percent of their initial capacity every 10 years, making 40-year

battery life impossible. By contrast, the highest quality bobbin-type $LiSOCl_2$ batteries can feature a self-discharge rate as low as 0.7 percent per year, retaining 93 percent of their original capacity after 10 years, thus enabling up to 40-year operating life.

The analogy to marathon running

Distance. Distance is equivalent to the battery/device operating life. The farther a runner can travel equates to the more years a device can potentially operate (Figure 2).

Incline. An incline is equivalent to the rate of battery self-discharge. The larger the incline, the higher the rate of self-discharge (Figure 3). When athletes run up a steep incline, they expend greater amounts of energy, which shortens the maximum duration of the run. Similarly, higher rates of battery self-discharge expend greater amounts of energy to reduce battery operating life.

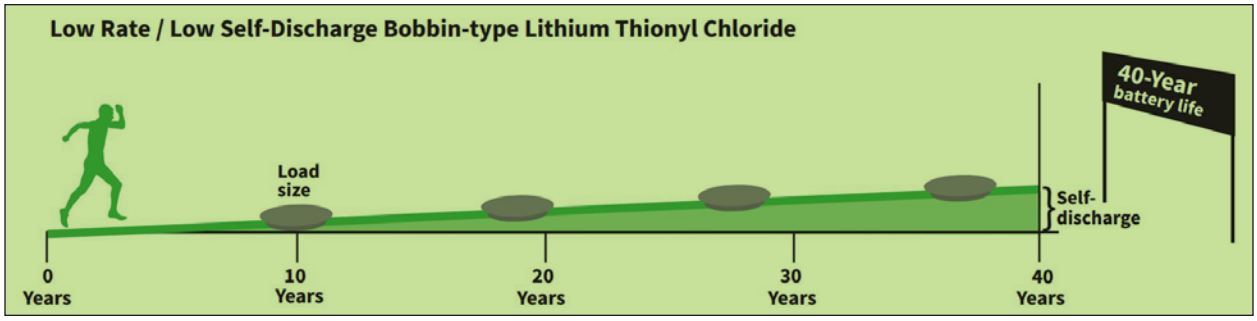

Figure 2. The farther a runner can travel equates to the more years a device can potentially operate.

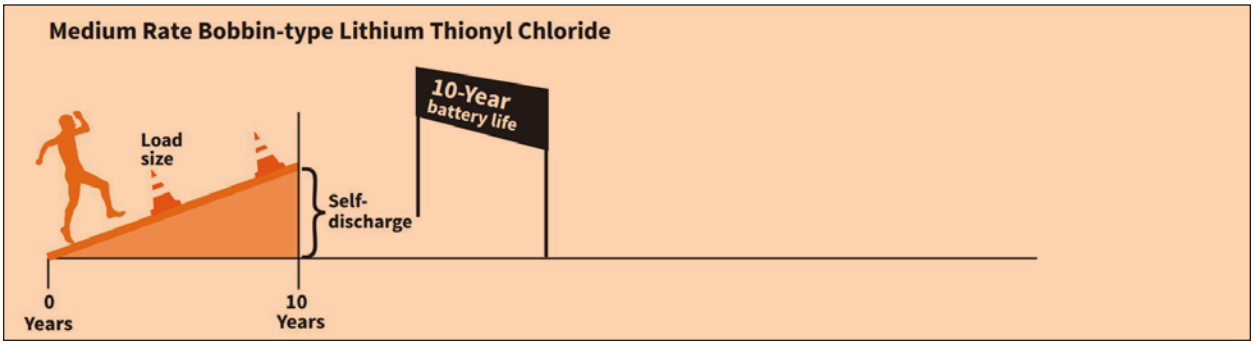

Figure 3. When athletes run up a steep incline, they expend greater amounts of energy, which shortens the maximum duration of the run.

Hurdles. High pulses of energy are similar to hurdles: the higher the hurdle, (obstacle) the higher the pulse being drawn by the battery (Figure 4).

Pole vault. Certain applications draw high rates of energy as well as high pulses (Figure 5). A prime example is a surgical power tool, which draws average current measurable in amps. As a result, these devices may be better suited to be powered by a lithium metal oxide battery.

Consumer devices with cell discharge rates measurable in the milli-amp to amp range, such as powering a flashlight or a consumer toy, are typically powered by an alkaline, LiFeS, $LiMNO_2$, or rechargeable Li-ion cell. By contrast, low-power remote wireless devices that draw average current measurable in micro-amps to conserve energy are typically powered by a bobbin-type $LiSOCl_2$ battery featuring an exceptionally low self-discharge rate.

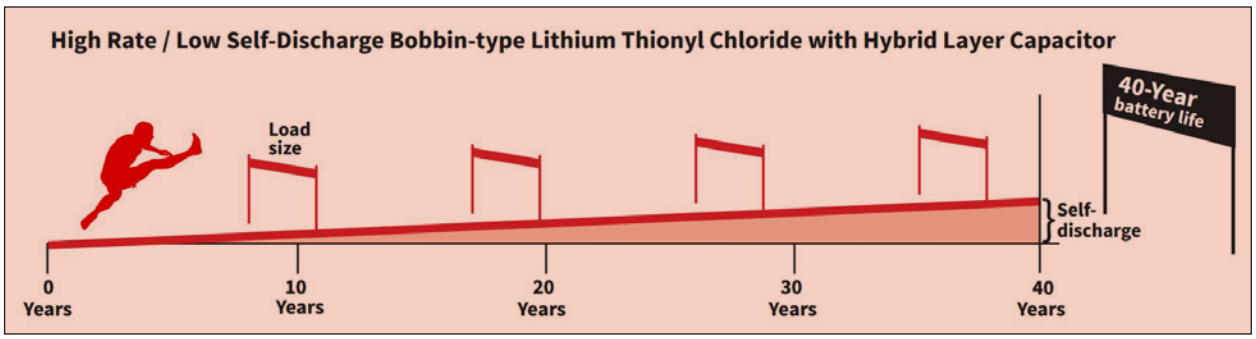

Figure 4. The higher the hurdle (obstacle), the higher the pulse being drawn by the battery.

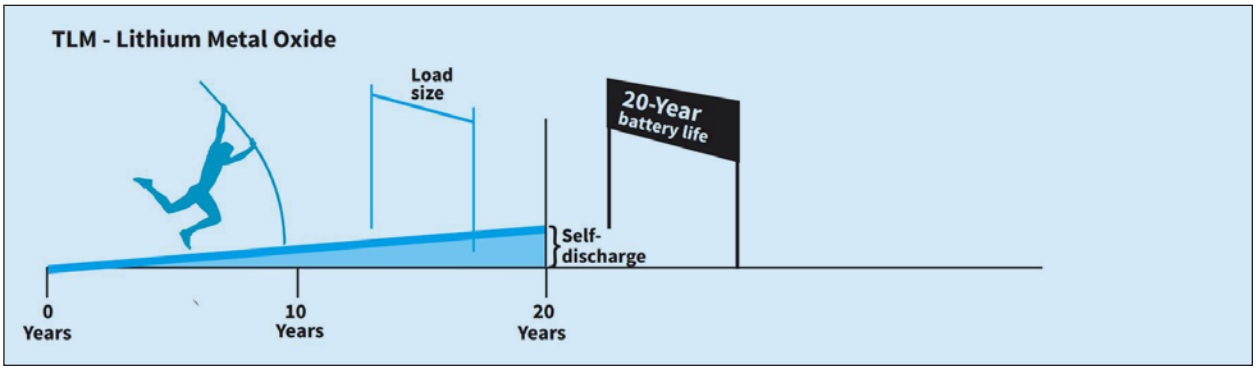

Figure 5. Certain applications draw high rates of energy as well as high pulses.

Standard bobbin-type $LiSOCl_2$ cells cannot deliver high pulses due to their low-rate design but can be modified with the addition of a patented hybrid layer capacitor (HLC) (Figure 6). With this hybrid approach, the standard bobbin-type $LiSOCl_2$ cell delivers low-level background current to power the device during stand-by mode while the HLC delivers the high pulses required to power two-way wireless communications (steeple jumping). The HLC also features a unique end-of-life voltage plateau that can be interpreted to deliver low-battery status alerts.

Figure 6. Bobbin-type $liSOCl_2$ batteries can be combined with a patented hybrid layer capacitor (HLC) to offer up to 40-year operating life while providing high pulses to power two-way wireless communications.

Many consumer devices use supercapacitors to deliver high pulses electrostatically rather than chemically. However, supercapacitors are rarely used in industrial applications due to performance limitations such as short-duration power, linear discharge qualities that prevent use of all the available energy, low capacity, low energy density, and high annual self-discharge rates (up to 60 percent per year). Supercapacitors linked in series also require the use of cell-balancing circuits that add expense, increase bulkiness, and consume additional energy to further accelerate their self-discharge rate.

Drawbacks of simulated tests

Long-term battery performance cannot be easily duplicated using short-term testing procedures. As a result, specialized test methods must be used to predict expected battery operating life. These testing methods include:

- **Long-term laboratory testing.** The ideal way to monitor long-term battery self-discharge is by continually testing random samples under various operating and environmental conditions. Over time, the accumulated data points can be used to accurately predict expected battery life based on cell size, temperature, load size, etc.

- **Accelerated testing.** The Arrhenius equation uses a two-fold increase of reaction rate for every 10°C rise in temperature to simulate long-term battery operation. Arrhenius tests are run at 72°C, which is equivalent to about 32 times the theoretical lifetime of a battery stored at 22°C. Unfortunately, the Arrhenius equation tends to yield inaccurate results with short-term tests.

- **Calorimeter testing.** An extremely accurate tool for predicting battery life is to measure the amount of actual heat energy being lost using a high-quality microcalorimeter. This device can detect small amounts of energy being dissipated as low as 0.1 W. This heat energy can be generated several ways: entropy change, often referred to as reversible heat; cell over-protection, often referred to as irreversible heat; chemical reactions that can affect self-discharge reactions and cell capacity; and side reactions that do not affect cell capacity. Calorimeter testing can be especially useful for measuring battery capacity losses resulting from long-term storage or device operation (including self-discharge), which can be measured using thermodynamic equations and cell voltage considerations. To ensure accuracy, prior to undergoing calorimetric testing, the battery needs to be stabilized for a minimum of one year as the self-discharge rate during the first year tends to be much higher than in subsequent years.

- **Lithium titration.** Lithium titration can also be used to measure available cell capacity. The battery is cut open, and then titration is used to dissolve the remaining lithium to determine its volume. High rates of self-discharge will accelerate the reduction of lithium found in the cell. Lithium titration techniques can be useful for measuring the effects of extreme temperatures or prolonged high pulses.

- **Field results.** Perhaps the best form of validation is to test a random sample of cells being deployed in the field. For this reason, Tadiran customers are asked to provide random samples of cells being used in the field to measure the real-life impact of long-term exposure to extreme environmental conditions. For example,

AMR/AMI batteries deployed by Aclara (formerly Hexagram) in the mid-1980s were tested after more than 28 years in the field and were found to contain plenty of unused capacity after nearly three decades.

Another useful indicator is to calculate the number of failures in time (FITs) among a very large sample of batteries. Some batteries consistently achieve FITs ranging from 5 to 20 batteries per billion, which is extremely low compared to the industry average.

Final thoughts

Every application has unique power requirements. As a result, you must determine whether your device needs to be powered by a sprinter (a cell with high discharge potential), a medium distance runner (a cell capable of a moderate to high discharge rate with fairly low self-discharge), or a long-distance runner (a cell featuring an exceptionally low self-discharge rate while also being able to deliver high pulses).

ABOUT THE AUTHOR

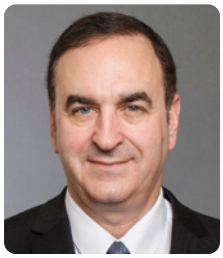

Sol Jacobs is vice president and general manager of Tadiran Batteries. He has more than 30 years of experience in powering remote devices. His educational background includes a BS in engineering and an MBA.

Interested in Edge AI?

Download Premio's M.2 Whitepaper

AI Inference At The Rugged Edge

What's Inside?

Enterprises today are seeking modern, purpose-built hardware strategies to relieve the edge AI bottlenecks traditional compute solutions face. This paper takes a deeper look into how M.2 performance accelerators look to drive Industry 4.0 edge AI applications.

Download Whitepaper

Unlock The Full Potential Of Edge AI with M.2 Acceleration Modules

Hitting an AI Wall

Compute has come a long way and AI still needs more.

Get To Know the M.2 Interface

A compact, versatile next generation option.

Throughput Matters

Understanding benchmarks for real-world AI applications.

800.977.3646 sales@premioinc.com www.premioinc.com

premio

AI Inference at the Rugged Edge: Meeting Performance with M.2 Accelerators

Data at the rugged edge enables capabilities, accuracy, and speed beyond human performance.

By Dustin Seetoo, Premio Inc.

Balancing power, performance, thermals, and footprint is the next hurdle in data-driven applications. As the number of Internet of Things (IoT) and industrial IoT (IIoT) devices continues to increase, so does the volume and velocity of data they generate. This trend, combined with the realities of continually increasing numbers and types of connected devices, creates a wealth of new opportunities for purpose-built computing solutions. And one that also demands a different approach to hardware designs that enable optimized performance.

Data drives business innovation, and most important, the ability for cognitive machine intelligence. On the factory floor, powering smart kiosks or advanced telematics, fueling surveillance and passenger services in infrastructure facilities like airports and train stations—data is everywhere and adds value when it can be revealed, captured, analyzed, and applied in real time. Yet for many of these applications performing in rigorous industrial environments, running small automated or artificial intelligence (AI) tasks from a data center is too inefficient to add true value. In this traditional centralized compute

structure, power and costs are too high. This is due to excessive, but necessary, use of compute, storage, and bandwidth resources. Performance trade-offs deepen the sacrifice, with factors such as high latency and insufficient data privacy.

For system designers, this means that yesterday's performance acceleration strategies may no longer fit the bill. While CPU/GPU-based designs have helped manage the slowing of Moore's Law, these processor architectures are now having difficulty keeping in step with the real-time data requirements inherent to automation and inference applications. This is particularly true in more rigorous non-data center scenarios. It's not just the data-intensive nature of automation that is causing change—it's where it is being implemented (Figure 1). As applications move out of the data center and into the world, more industrial and nontraditional computing settings are seeking greater competitive value from data in real time. This can be defined as the rugged edge, and it is here that performance acceleration requires a new path.

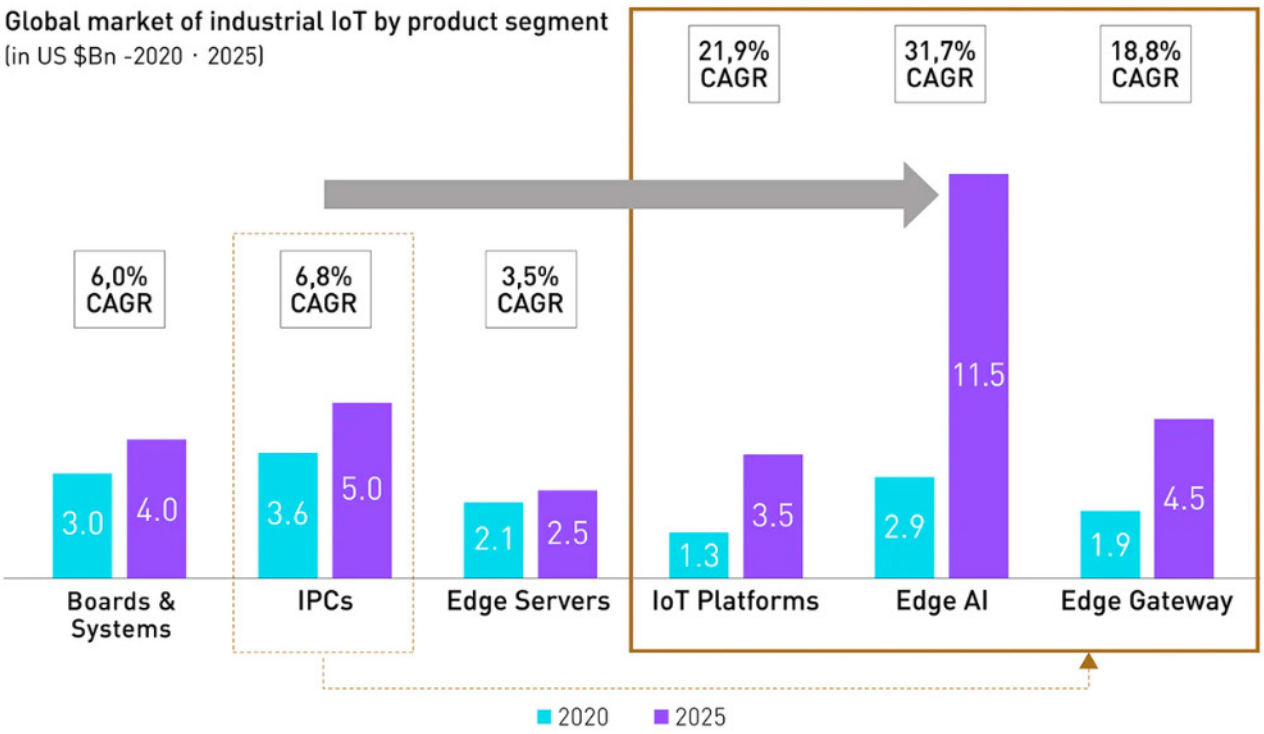

Figure 1. Gartner predicts that by 2025, 75% of enterprise data will be processed at the edge.

Combined with the mounting challenge to meet price-performance-power demands, it's more critical than ever that performance acceleration consider compute, storage, and connectivity. All these factors are necessary to effectively consolidate workload close to the point of data generation, even in rugged settings where environmental challenges are detrimental to system performance.

Using M.2 accelerators, DSAs operate deep neural networks from 15 to 30 times faster (with 30 to 80 times better energy efficiency) than a counterpart network relying on CPU and GPU technologies.

Edge computing hardware is being deployed to cope with this increasing amount of data and to alleviate the ensuing burdens placed on the cloud and in data centers. Data-intensive workloads like AI inference and deep learning models are moving beyond the data center and into factories and other industrial computing environments. In turn, designers and developers are recognizing a current shift for performance acceleration closer to data sources such as IoT sensors. The trend is pushing edge hardware further, meeting the need to interact with AI workloads on-demand. Based on data growth coupled with the complexities of edge computing environments, today's AI computing framework is moving from general CPU and GPU options to more specialized accelerators such as smaller and more power-efficient acceleration modules in an M.2 standard.

This is where M.2 form-factor accelerators come into play for eliminating performance barriers in data-intensive applications. A powerful design option, M.2 accelerators offer system architects domain-specific value to match the exact requirements of AI workloads. In contrast to a comparable system using CPU/GPU technologies, an M.2-based system can manage inference models significantly faster and far more efficiently. These increases are driving innovative system design perfect for the rugged edge where

more systems are deployed in challenging, nontraditional scenarios, and where purpose-built systems offer immense opportunity. Here, there is a clear differentiation between a general-purpose embedded computer and one that's designed to handle inferencing algorithms by tapping into more modern acceleration options like M.2 domain-specific acceleration modules.

Domain-specific architectures handle only a few tasks, but they do so extremely well. This is an important shift for system developers and integrators, as the ultimate goal is to improve the cost versus performance ratio when comparing all accelerators, including CPUs, GPUs, and now M.2 options.

Using M.2 accelerators, DSAs operate deep neural networks from 15 to 30 times faster (with 30 to 80 times better energy efficiency) than a counterpart network relying on CPU and GPU technologies. While general-purpose GPUs are commonly used to provide enormous processing power for advanced AI algorithms, they are not optimized for edge deployments in remote and unstable environments. Drawbacks of size, power consumption, and heat management generate additional operating costs on top of the upfront cost of the GPU itself. Specialized accelerators such as TPUs from Google and M.2 acceleration modules are new solutions that are compact, power efficient, and come purpose-built for driving machine learning algorithms at the edge with incredible performance.

Diving into M.2 and domain-specific architectures

Accelerators deliver the considerable data processing required, filling the gaps caused by the deceleration of Moore's Law, which for decades was a driving force in the electronics industry. This long-established principle asserts that the number of transistors on a chip will double every 18 to 24 months. When it comes to AI, however, industry experts are quick to point to signs of wear in Moore's Law. Silicon evolution in and of itself cannot support AI algorithms and the processing performance they require. To balance performance, cost, and energy demands, a new approach must feature domain-specific architectures (DSAs) that are far more specialized. Customized to

execute a meticulously defined workload, DSAs provide a fundamental tenet for ensuring performance that facilitates deep learning training and deep learning inference.

The M.2 interface: a compact, versatile next generation option

Intel developed the M.2 (next generation form factor [NGFF]) interface for flexibility and powerful performance. M.2 supports multiple signal interfaces such as PCI Express (PCIe 3.0 and 4.0), Serial ATA (SATA 3.0), and USB 3.0. A range of bus interfaces enable M.2 expansion slots to be highly versatile for different storage protocols, performance accelerators, wireless connectivity, and input/output (I/O) expansion modules. For example, M.2 expansion slots can be used to add wired and wireless capabilities or a range of M.2 SSDs with different sizes and specifications.

Besides connectivity and storage expansion modules, performance accelerators (Figures 2-5) have quickly adopted the M.2 form factor to benefit from its compact and powerful interface. These performance accelerators include memory accelerators, AI accelerators, deep learning accelerators, inferencing accelerators, and more. These new specialized processors dedicated to AI workloads provide an improved power-to-performance ratio. This is demonstrated by domain-specific workloads handled by M.2 accelerators versus heterogeneous compute SoCs such as CPU and GPU resources.

At right are a few of the top M.2 performance accelerators that are available today.

Figure 2. M.2 Intel Optane Memory: Intel's speed-boosting cache storage in an M.2 format. Developed to accelerate cache for another drive to enable high-speed computing
Source: Intel

Figure 3. M.2 VPU: Intel's Movidius VPU (vision processing unit), developed to enhance machine learning and inferencing for edge computer vision that requires robust and compact technologies.
Source: AAEON

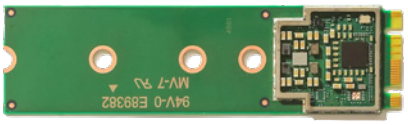

Figure 4. M.2 TPU: Tensor Processing Unit, developed by Google to accelerate training on large and complex neural network models. A powerful and energy-efficient AI accelerator in a compact M.2 form factor.
Source: Coral.AI

Figure 5. M.2 Hailo-8: AI Acceleration module—a best in class inference processor packaged in a module for AI applications; offers 26 tera operations per second and compatibility with NGFF M.2 form factor, M, B+M, and A+E keys.

Throughput matters: understanding benchmarks for real-world AI applications

Even the metrics by which industry experts measure compute performance are changing to accommodate the data-rich nature of AI applications. TOPS, or tera (trillions) operations per second, is a measure of the maximum possible throughput rather than a measure of actual throughput. TOPS identifies the number of hardware-implemented computation elements times their clock speed.

While important, it is essentially a measure of what is possible if all the stars align in a given application; that is, steady data input, clean and consistent power sources, no memory limitations, and perfect synchronization between hardware and AI software algorithms. As a theoretical measurement, TOPS also do not offer any consideration for other tasks the hardware may need to perform. Engineers focused on silicon implementation may find specific value in TOPS data, but software and hardware systems engineers may find that it does not clearly indicate true, available performance for their real-world application.

Throughput, not TOPS however, is the more precise, real-world measurement. Throughput references the amount of data that can be processed in a defined time period, for example frames per second (FPS) in vision processing terms or the number of inferences in deep learning edge application. Inferences or FPS per watt, as related to a specific neural network task or application, is not only a more precise way of evaluating and comparing hardware, but also a more clearly understood, real-world metric. On this landscape, both ResNet-50 and YOLOv3 have emerged as leading options for AI performance evaluation, as well as use as a backbone in the development of new neural models.

ResNet-50, a pre-trained deep learning model

To solve computer vision challenges, machine learning developers working with convolutional neural networks (CNNs) add more stacking layers to their models. Ultimately, a higher number of levels increases saturation such that it creates a negative impact on performance of

both the testing and training data models. ResNet-50 tackles this deterioration problem. Using residual blocks, or "skip connections" that simplify learning functions for each layer, ResNet-50 improves deep neural network efficiency while reducing errors.

YOLOv3, a real-time object detection algorithm

As a CNN, YOLOv3 (you only look once, Version 3) identifies specific objects in real time, for example in videos, live feeds, or images. YOLO's classifier-based system interacts with input images as structured arrays of data—its goal is to recognize patterns between them and sort them into defined classes with similar characteristics. Only objects with similar characteristics are sorted; others are ignored unless system programming instructs attention. The algorithm allows the model to view the entire image during testing, ensuring its predictions are informed by a more global image context. A live traffic feed provides a good example—here, YOLO can detect various types of vehicles, examining high scoring regions and identifying similarities with certain predefined classes.

Looking ahead: unlocking AI with real-time data performance in more environments

Data is key to today's business innovation, and moreover, the ability to deliver cognitive machine intelligence. Data is all around us, and it is most valuable when it can be harnessed in real time. Myriad industries are eager to make the most of data to create new services and enhance business decisions, but in many rigorous industrial environments, processing small automated or AI tasks at the data center level is just too inefficient to provide true value. Power consumption and costs are too high in this legacy centralized compute structure because of excessive, albeit necessary, use of compute, bandwidth, and storage resources. Further, high latency means that performance takes a hit, and insufficient data privacy creates another headache.

Data growth, combined with edge computing environment complexities, is driving the AI computing framework away from general CPU/GPU options and toward specialized accelerators based on domain-specific architectures that use the M.2 standard—options that are smaller and more power efficient. It's a strategy to address a data challenge that is complex, real, and not going away. Application designers and developers must recognize an urgent need for performance acceleration that resides closer to data sources and is purpose-built for the task at hand—particularly as edge computing hardware is deployed to cope with data processing and alleviate related burdens in data centers and in the cloud.

There is a clear differentiation between a general-purpose embedded computer and one that's designed to handle inferencing algorithms. M.2 is proving to be a powerful design option for system architects, offering domain-specific value that meets the precise needs of AI workloads to eliminate performance roadblocks. For system developers, the opportunity for purpose-built systems is immense, with smarter data handling poised to advance AI and inference applications even more broadly across global infrastructure industries.

To see real-world benchmarks of various deep learning models that use M.2 performance accelerators modules in purpose-built industrial computing solutions, visit Premio Inc.

ABOUT THE AUTHOR

Dustin Seetoo is the director of product marketing at Premio Inc., www.premioinc.com. For more than 30 years, Premio has been a global solutions provider specializing in the design and manufacturing of computing technology from the edge to the cloud for enterprises with complex, highly specialized requirements.

One Chip, One Driver, All Protocols

Future-Proof your solutions! Our next-gen netX 90 chips have everything you need to put your devices on the Industrial Internet of Things. All networks; All software stacks; All security features; All clouds. And all packaged in a tiny 10X10 mm footprint.

Learn more from Hilscher, call 1.630.505.5301,
email: info@hilscher.us or visit www.hilscher.com

hilscher
COMPETENCE IN COMMUNICATION

How TSN Is Accelerating Digital Transformation for Industrial Automation

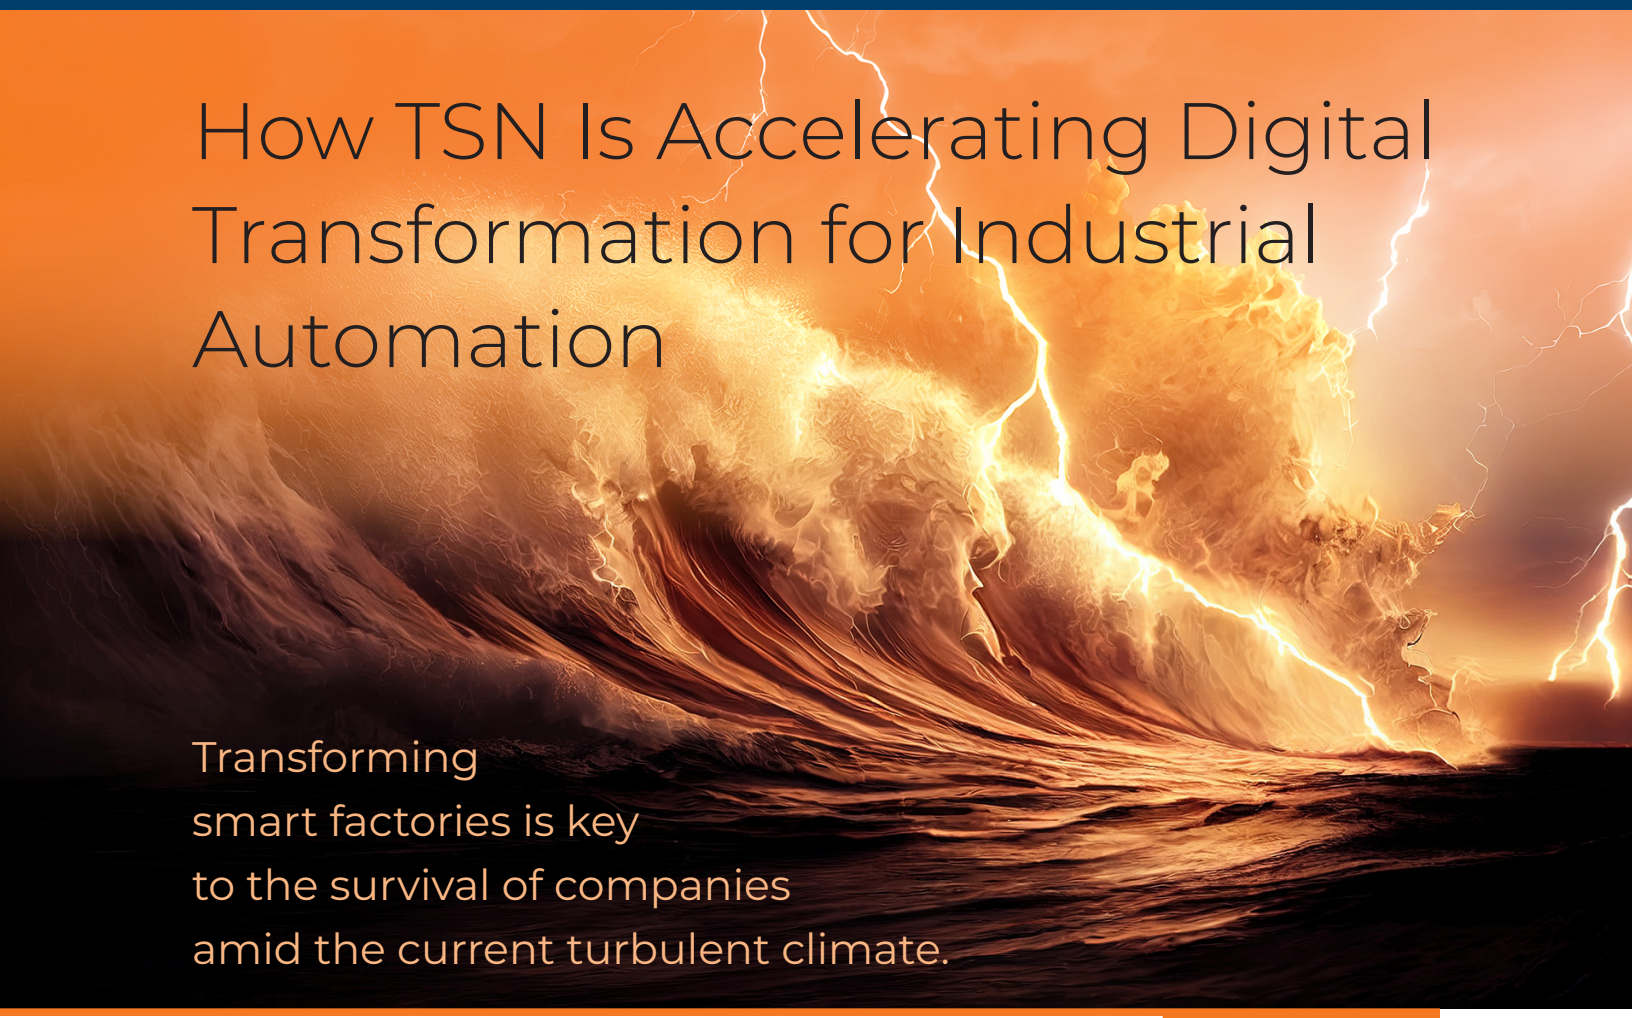

Transforming smart factories is key to the survival of companies amid the current turbulent climate.

By Jack Lin, Moxa

Most companies take advantage of economies of scale by producing a larger number of goods to gain proportionate saving in costs. However, it is not always the case that the more goods that are produced, the less each unit costs.

Currently, most companies have reached the optimum point of economies of scale, and the cost for additional units is increasing. As mass production is no longer the only answer, we find ourselves asking, what's next? The combination of markets continuing to change rapidly, and global supply chain challenges has forced companies to find new ways to adapt to become more competitive. How to implement Industry 4.0 and Industrial Internet of Things (IIoT) applications and transform smart factories using digital transformation becomes key to the survival of companies amid the current turbulent climate.

Adaptive manufacturing brings success now and in the future

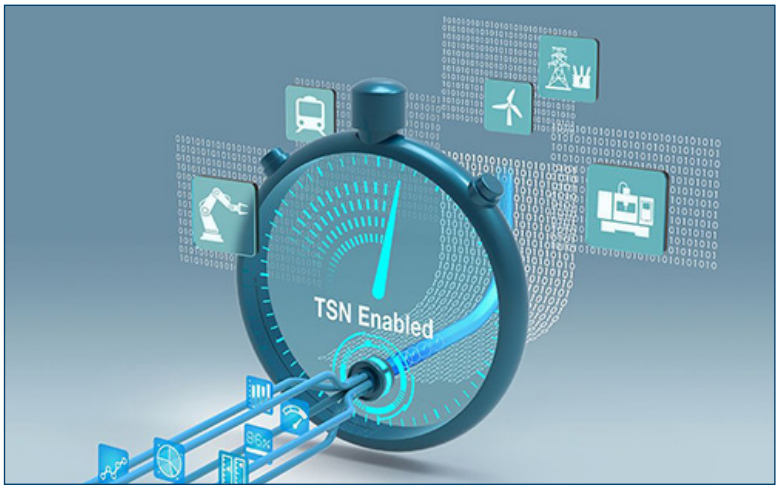

To become agile and responsive to market demands, it is important to consider adaptive manufacturing (Figure 1). The concept of adaptive manufacturing lies in a company's ability to rapidly adjust production operations to meet the new, customized demands—an increasingly important capability to stay competitive, and to cope with supply chain disruptions caused by COVID-19. Adaptive manufacturing, which includes close co-competition within ecosystems and strategies for mass customization, is the foundation for Industry 4.0 and IIoT, as it uses a production system that focuses on data to ensure everything is connected and communicated, and the right decisions can be made quickly.

Concept of Adaptive Manufacturing	
Production Aspect	Benefits and Results
Assest Management	Greater control over assets to ensure continious operation and lower the TCO.
Service-as-a-Service	Enhance connectivity to provide more efficient and innovative services.
Integration of Control Systems	Integrate control processes and automation systems on the production line to enhance production efficiency and operational excellence.

Figure 1. Adaptive manufacturing, the foundation for Industry 4.0 and IIoT.

Time-sensitive networking bridges the gap between traditional and adaptive manufacturing

It is never easy to choose the right approach to retrofit existing facilities, especially when it involves enlarging them. Time-sensitive networking (TSN) is a game changer for digital transformation in the automation industry and is the foundation of unified infrastructure. TSN offers time synchronization, low-latency communications, ultra-reliability, and better resource management. This facilitates the merging of information technology (IT) and operational technology (OT) networks to connect all end devices and centralize data collection on a standard ethernet network (Figure 2). Digital transformation can be achieved by real-time information made possible by TSN and to enable factories to optimize their production strategies for fast-changing markets and emerging challenges.

The potential of adaptive manufacturing is huge. First, consider asset management. It saves costs from the very beginning when a unified infrastructure is built. TSN can connect multiple applications and diversified systems on one unified network, which makes it easier to operate and maintain, as less equipment and fewer cables are needed at the field level.

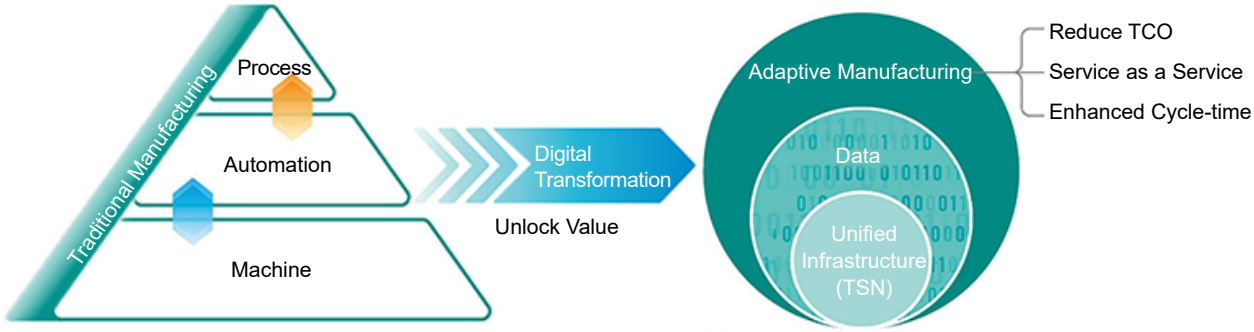

Figure 2. Merging of information technology (IT) and operational technology (OT) networks to connect all end devices and centralize data collection on a standard ethernet network.

Meanwhile, one unified network converges both IT and OT networks, which means modern IT technologies, such as artificial intelligence and machine learning, can be incorporated to assist operational efficiency and give owners more control over their assets. Second, companies can evolve from selling products to selling services, offering more customization and reducing the minimum order quantity (MOQ) without losing profitability and giving owners more possibilities. Third, integrating control systems substantially reduces the cycle time. A few years ago, it might have required weeks to shift the entire production process to producing a new product, but by integrating different control systems, it will take only a few days, making adaptive production not only feasible but also profitable.

Time-sensitive networking is a game changer for digital transformation in the automation industry and is the foundation of unified infrastructure.

Three real-world cases that tap into benefits of one unified network

Moxa has helped end users around the world tap into the benefits of adaptive manufacturing using TSN to help accelerate digital transformation for industrial automation.

One unified network to realize adaptive manufacturing

A product manufacturer planned to connect its various systems—including production, assembly lines, and logistics systems—on one unified network to realize shorter production cycle times and lower the total cost of ownership. From a production viewpoint, TSN accomplishes integration of control with reduced cycle times, reduces the total cost of ownership through a simplified topology with fewer assets to manage, and realizes service-as-a-service through one unified network.

Enabling multiple applications on a unified TSN network to reduce the production cycle time

A global leading manufacturer of industrial machinery is leveraging TSN technology to aggregate multiple applications in CNC machinery. To achieve its scalable, accelerated sensing, and advanced machine

control systems, the deterministic laser and motion control need to work harmoniously. With TSN, the devices can be efficiently integrated, and the production cycle time can be successfully reduced.

AI-driven operational efficiency for a hydropower plant to reduce the TCO

The operating company of a hydropower plant was determined to bring all its isolated networks together and implement an artificial intelligence (AI) system for its control network by embracing the TSN standard, which was a perfect fit for this kind of use case, as performing services on a converged network was much easier than on disparate networks. The hydropower plant improved efficiency and the ability to quickly adjust the total power output to the grid as needed, giving rise to a new hydropower plant with lower costs, easier maintenance, higher efficiency, and improved adaptability.

Final thoughts

TSN is an advanced technology that helps industry leaders accelerate their digital transformation in industrial automation. In this article, we briefly introduced adaptive manufacturing and the core technology of TSN that helps achieve it. Building on this point, the three real-world cases demonstrate how companies can revolutionize their industrial operations while tapping into the benefits of TSN.

ABOUT THE AUTHOR

Jack Lin is a product manager and project lead of Global TSN Initiative Project at Moxa. He has worked at Moxa since 2012 and is currently a product manager in the industrial networking infrastructure team and also the project lead for Moxa's global TSN initiative project. Lin is responsible for scouting, evaluating, and introducing advanced communication technologies that are suitable for the next generation of industrial automation projects. Furthermore, he has established and manages strategic partnerships, alliances, and ecosystems with industry-leading companies to advocate and ensure the targeted technologies are a success.

Premium performance, without premium prices

The CX52x0 Embedded PC series for PLC, motion control and IoT

www.beckhoff.com/cx52x0

With the CX52x0 Embedded PC series, Beckhoff offers a cost-effective hardware platform for universal use in automation and IoT applications. The two fanless, DIN rail-mountable versions offer users the high computing and graphics performance of the new Intel Atom® multi-core generation while greatly reducing heat dissipation. The basic configuration includes a direct I/O interface for Bus Terminals or EtherCAT Terminals, built-in IoT and cloud capabilities, two 1,000 Mbit/s Ethernet interfaces, a DVI-D interface, four USB 3.0 ports and a multi-option interface that can be equipped to accommodate a wide range of fieldbuses.

CX5230:
Intel Atom® x5-E3930,
1.3 GHz, 2 cores

CX5240:
Intel Atom® x5-E3940,
1.6 GHz, 4 cores

intel. partner
Titanium
IoT Solutions

New Automation Technology **BECKHOFF**

Next-Generation Biotechnology Manufacturing: Genentech Small-Batch Biologics

By Bill Lydon, Automation.com

A Level 5 Adaptive Plant incorporates advanced manufacturing and environmental sustainability.

I had a discussion with Genentech's Benedicte Lebreton, VP, head of technical development for Biologics US, about the new Genentech Clinical Supply Center for small-batch biologics. This flexible facility implements manufacturing concepts, innovations, and technologies that are models for future expansion. The Clinical Supply Center has put into practice much of what has been described by the BioPhorum Digital Plant Maturity Model as highest level next-generation manufacturing—a Level 5 Adaptive Plant. Single-use technologies are fundamental to achieving the facility's goals of flexible manufacturing and environmentally sustainable production.

Genentech, a member of the Roche Group, is the founder of the biotechnology industry and pioneer of the processes and technologies used to commercially manufacture complex biologic medicines. The Clinical Supply Center is part of the Roche Global Manufacturing Network and one of three manufacturing sites in California. Globally, Roche and Genentech produce medicines at 11 sites around the world. In 2021, Genentech delivered more than 55 million vials for patients.

Clinical Supply Center facility

The facility serves as a new model for the fast, efficient production of therapeutics, including personalized and rare disease medicines, for smaller patient populations. It is designed for efficiency and sustainability, and has achieved globally recognized LEED Gold Certification for healthy, efficient, carbon- and cost-saving green buildings.

The 78,520-sq. ft. facility supports digitalization from product definition to production. Its communications connect people and equipment with more than 19,000 feet of network cabling and apparatus and more than 1,000 digital connections, approximately one every 75 square feet.

Genentech, of the Roche Group, is the founder of the biotechnology industry and pioneer of the processes and technologies used to commercially manufacture complex biologic medicines.

System architecture

The digital manufacturing automation architecture (Figures 1 and 2) integrates all systems, including:

- Laboratory instrument management systems (LIMS)
- Drug substance discrepancy management system (DMS)
- Electronic document management system (EDMS)
- Environment monitoring system (EMS)
- Calibration management system (CMS)
- Product lifecycle management (PLM)
- Distributed control system (DCS)
- Building automation system (BAS)

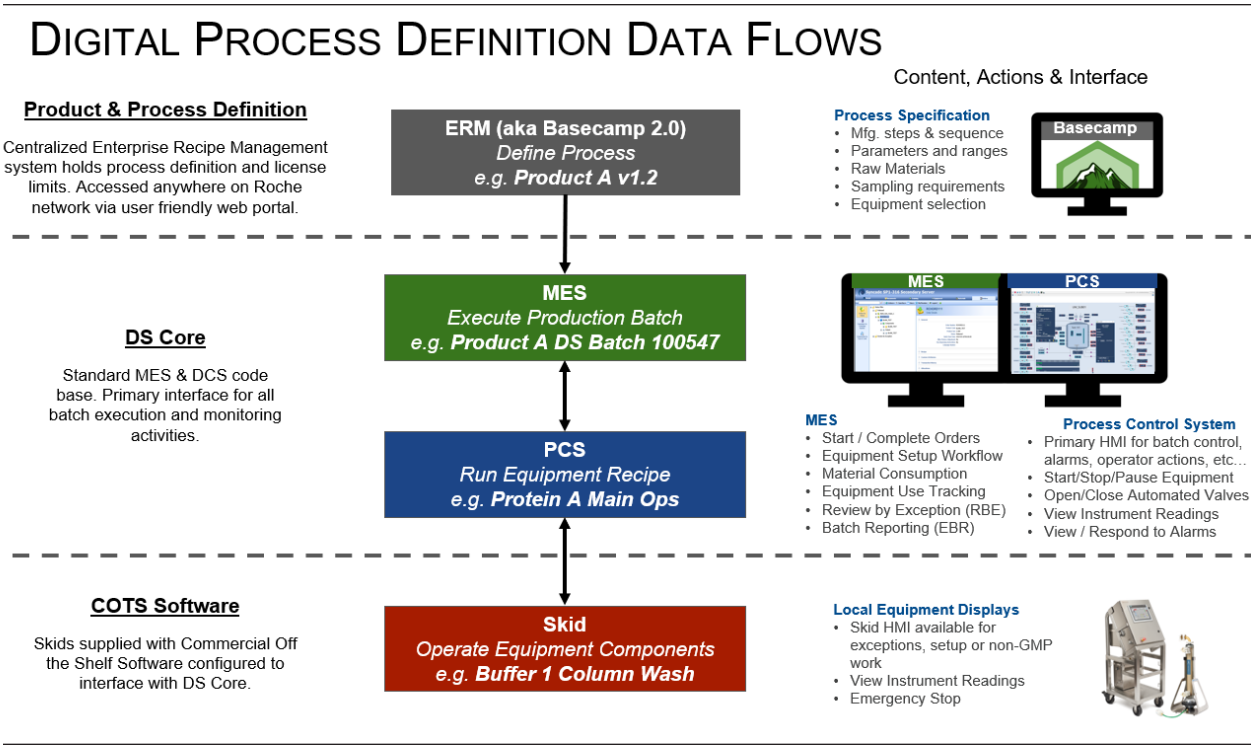

Figure 1. Digital process definition data flows.

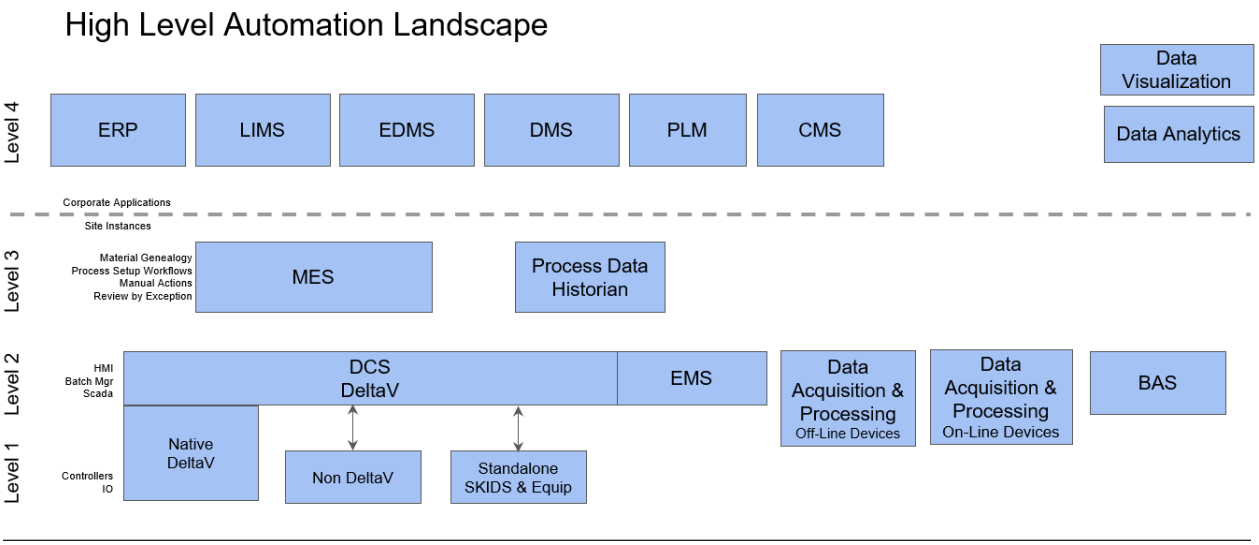

Figure 2. System configuration.

Lebreton described the project as a success, with no unexpected operational issues. She noted the project has achieved the design requirements:

- Rapid and seamless technology transfer from research to clinical manufacturing by using the PLM system
- Standard plug and play utility and data connections for single-use technology equipment
- Integrated paperless digital manufacturing
- Quick production reconfiguration changes enabled by facility, equipment design, and layout
- Reuse of workflows across unit operations and molecules enabled by parameter-driven workflows
- True ballroom configuration for the process enabled by closed systems

A major benefit has been creating a reliable design template for building and operating flexible facilities based on the Genentech team's learning experience.

Sustainability and efficiency

The facility (Figure 3) furthers Genentech's long-standing commitment with impressive sustainability and energy efficiency results proven in operation:

- 25% reduction in overall energy use
- 100% of electricity from renewable sources
- 28% less water usage compared to other facilities

A major contributor to energy conservation has been using less water and steam with single-use technologies that do not require clean in place (CIP).

Single-use technologies have decreased the time between scientific discovery and clinical manufacturing, enabling faster delivery of new medicines to clinical trial patients, while using fewer resources.

"Rather than disposable, we would say renewable or reusable. Nothing from the facility goes to landfill," says Lebreton, noting that they have a collaboration with the local waste management company to bring their solid waste to a nearby facility to incinerate for energy. "The waste-to-energy process takes our waste and combusts it to generate steam for electricity generation. This process, combined with our reductions in water and energy, translate to a 14% lower carbon footprint." There are also recycling programs for gloves, cardboard, and other items, a wash-and-reuse program for safety glasses, and plans for "many upcoming pilots for soft plastics to reduce our solid waste stream," for example.

"We are very proud of another collaboration with a local hard plastic recycling company where we have established a true circular economy," adds Lebreton. "Our hard plastics are sent to this company and then are converted to items like pipette tips and centrifuge tubes we repurchase for use across our organization."

One Genentech sustainably partner is Polycarbin. The resource management company is dedicated to transforming the single-use scientific supply chain into the sustainable, circular economy of tomorrow by repurposing waste plastic.

Figure 3. The most advanced manufacturing technologies and an agile modular design speed on-demand delivery of the investigational medicine pipeline to patients in clinical trials.

As a member of the Roche Group, Genentech is committed to supporting industry initiatives and organizations such as NAMUR and BioPhorum Operations Group (BPOG), and participates in many of the groups that help align the industry with opportunities to improve in spaces that are not proprietary in nature. Genentech and Roche have influenced many of the white papers and output from BPOG and have incorporated best practices from the groups where applicable.

As a side note, Genentech has been an active contributor to the Pharmaceutical Automation Roundtable (PAR) that I have been involved in and describe in such articles as 2019 PAR Insights from the Annual Pharmaceutical Automation Roundtable and Survey.

Thoughts and observations

I asked Lebreton what was on her wish list. She said she would like more ideal sensors for single use. This is a new area, and I have spoken with several companies working on and refining designs for single-use sensors. We both agreed that, considering the short time that the single-use field has existed, progress has been good.

It is interesting to note that the consensus in the 2019 Pharmaceutical Automation Roundtable was that the technology wasn't quite ready to achieve the BioPhorum Digital Plant Maturity Model Level 5 Adaptive Plant. This Genentech project illustrates tremendous progress since that time toward achieving flexible small batch production.

Related Articles

- Leveraging Solid Disruptive Innovations Increases Competitiveness
- Digital Manufacturing Facts and Myths
- 2019 PAR Insights from the Annual Pharmaceutical Automation Roundtable and Survey
- BioPhorum MTP Testing Plugfest
- Biomanufacturers' MES Manifesto Specifies Requirements
- Other pharmaceutical related news on Automation.com

ABOUT THE AUTHOR

Bill Lydon brings more than 10 years of writing and editing expertise to Automation.com, plus more than 25 years of experience designing and applying technology in the automation and controls industry. Lydon started his career as a designer of computer-based machine tool controls; in other positions, he applied programmable logic controllers (PLCs) and process control technology. Working at a large company, Lydon served a two-year stint as part of a five-person task group, that designed a new generation building automation system including controllers, networking, and supervisory and control software. He also designed software for chiller and boiler plant optimization. Bill was product manager for a multimillion-dollar controls and automation product line and later cofounder and president of an industrial control software company.

Take Charge

To take charge and be sure your operations and system are in sync at optimal capacity, you need to be atop autonomy. And to get there, Yokogawa delivers resilient solutions for you, a process using our smart manufacturing and IA2IA (Industrial Automation to Industrial Autonomy), deploying OpreX as our true enabler to achieve total optimization throughout the supply chain. Integrating discrete systems in society, we move together with you toward the system of systems in which everything is intricately connected and goals are achieved beyond those of a single system. Yokogawa. Atop autonomy for the planet.

OpreX™

yokogawa.com/ia2ia/

YOKOGAWA

Co-innovating tomorrow™

Understanding Open Process Automation

OPAF members think OT systems should adopt IT tools and techniques; vendors and integrators are on board.

By Jack Smith, Automation.com

Cutting-edge automation and technology topics ranging from NASA's Artemis missions to clean energy to industrial cybersecurity were highlights of the 2022 Automation and Leadership Conference put on by International Society of Automation in November in Galveston, Texas. And one of the most important was the IIoT and Smart Manufacturing session on "Open Process Automation: Moving from Concept to Reality." Global industrial and chemical companies are collaborating with process automation vendors and system integrators to bring a change in automation architectures through the adoption of standards.

The Open Process Automation Forum (OPAF) sponsored a panel discussion at the conference revealing how its members are driving change away from proprietary architectures and toward an open architecture that allows the use of new state-of-the-art equipment and functions as well as the continued use of proprietary software applications.

Unlocking innovation and successfully implementing digital transformation requires full access to data from all operational levels, they said, from the highest levels of business systems down to the operational edge. That requires a system architecture that prioritizes interoperability, modularity, standards conformity, compliance with security standards, scalability, and portability.

"The vision of OPAF is to change from a close proprietary-controlled system architecture to an open, interoperable, and secure process," said Aneil Ali, director of The Open Group, which manages OPAF.

"Instead of the traditional hierarchical process control architecture, OPAF is moving toward a much flatter architecture, which places most of the control functionality where it is most needed, instead of taking the control functionality up several layers and then bringing it back down to the control element," Ali said. "We must ensure the data can flow freely from where it is generated to where it's needed, which enables us to connect both the business systems as well as the process-controlled manufacturing systems."

Coalition of companies

Other participants on the panel represented the coalition of companies at work within OPAF:

- Mohan Kalyanaraman, technology acquisition advisor, strategic planning, ExxonMobil
- David DeBari, process control engineer, ExxonMobil
- Sharul Rashid, group head, Petronas
- Tom Clary, director, Schneider Electric
- Bob Hagenau, CEO, CPLANE.ai
- Jacqueline Allen, director, business development and applied intelligence, Wood

OPAF is tackling market education and adoption of open process automation from three different angles, said Ali. “First, we’re developing the technical standard. Second is developing the certification program so end users can have confidence in the products that enter the market as ‘open’: They’re certified to pass the standard, and they’ll work according to what it says on the box,” he explained.

The third aspect is developing the business ecosystem—the end users, suppliers, and system integrators who contribute to the success of the standard through publication and adoption of guides, whitepapers, and other tools.

A coalition of vendor companies needed to come together “because, by definition, an open system can’t be delivered by a single vendor,” said Bob Hagenau, CEO of CPLANE.ai. “We saw the need to bring together vendors that had components into a single group called the Coalition for Open Process Automation, or COPA, to start to build these systems and show that they’re viable.” COPA members also wanted to create commercial products that could be purchased.

Built on standards

OPAF, formally launched in November of 2016, develops and maintains the Open Process Automation Standards (O-PAS, figure 1). To ensure interoperability, the parts of O-PAS follow accepted reference standards including ISA 95, ISA 99, ISA 18, OPC UA, and others. The O-PAS Standard, Version 2.1 *Preliminary* was published in May 2021 and is available now. The O-PAS Standard, Version 2.1 *Final* will be available in 1Q 2023.

A “Standard of standards”		
O-PAS part	**Subject matter**	**Referenced standards**
Part 1	Technical architecture	IEC 62264 (ISA 95)
Part 2	Security	IEC 62443 (ISA 99)
Part 3	Profiles	n.a.
Part 4	Connectivity framework	IEC 62541 (OPC UA)
Part 5	System management	DMTF (Redfish)
Part 6	Information and exchange models	IEC 62714 (AutomationML IEC 62682 (ISA 18) IEC 61131 IEC 61499
Part 7	Physical platform	“whitespace”

Figure 1. OPAF is focused on developing a multivendor standards-based open, secure, and interoperable process control architecture.

The final version expands the standard based on feedback from the industry, and suppliers can begin building products for certification to it. Latest version of The O-PAS Standard is always available at: https://publications.opengroup.org/standards/opa.

According to Mohan Kalyanaraman, technology acquisition advisor and strategic planning for ExxonMobil, version one of the standard dealt with interoperability and communications. Version two dealt with configuration portability and control functionality. “For the next set of features, we are introducing our application portability, orchestration, and physical platform. That will be the next set of standards,” he said.

David DeBari, process control engineer for ExxonMobil, commented on the end user experience. “A lot of the technologies we talk about are standards based; they have a lot of appeal. You don’t have to know the proprietary methods if everybody understands the standards.”

From an end user perspective, OPAF members have a range of opinions, DeBari explained. “We have the spectrum. I am out here waving the flag for OPA, but I do have peers that some days wish we would leave it alone because they like it the way it is. Change management is that way,” he said

> **“There needed to be a coalition** of companies to come together because by definition, an open system can’t be delivered by a single vendor.”
> *—Bob Hagenau, CPLANE.ai*

ExxonMobil has “been doing prototypes and test beds, and we’re moving to a field trial where the technologies are available,” said DeBari. “We believe we should be adopting some of these tools and techniques that our information technology (IT) friends have been using for years.”

“From a commercial and a business aspect, we must test it to make sure it’s ready, make sure it’s reliable and safe to use on our types of processes. Those challenges are where the evolution is, and the recent momentum in the industry [shows it is] really taking off,” DeBari added. “We’re seeing product from vendors. We’re seeing support from system integrators and vendors. And we’re getting buy-in from many other end users who are waking up to the fact that standards-based architecture is the way of the future.”

Test beds and technology partners

Sharul Rashid, group lead at Petronas, added, "Open process automation is the only solution that has worked for us so far. The value created from it is very obvious. So the next step was to convert the non-believer within my company, my management, because they control the budget. If the budget is not made available to me, I cannot proceed with my test bed. But we were able to convince them of the security behind it. We have to move forward as part of the transformation. We are now moving to the test bed and intend to deploy it by first quarter of 2023."

Acronyms and references

OPAF: The Open Process Automation Forum (OPAF) is a consensus-based group of end users, suppliers, system integrators, standards organizations, and academia. It addresses both technical and business issues for process automation. OPAF is a part of The Open Group, a global consortium that enables the achievement of business objectives through technology standards with more than 800 participating organizations including users, systems and solutions suppliers, tool vendors, integrators, academics, and consultants across multiple industries.

O-PAS: The Open Process Automation Forum (OPAF) develops and maintains Open Process Automation Standards (O-PAS).

COPA: Coalition for Open Process Automation is a partnership program to encourage collaboration among operational technology (OT) and information technology (IT) vendors in the creation of commercial control systems based on the Open Process Automation Standard (O-PAS).

At a tipping point?

When Automation.com's Bill Lydon reported on OPAF activities earlier this year, he said, "The computer industry has advanced innovation, performance, ease-of-use, and application portability dramatically with open systems, while industrial systems have remained proprietary at their core. The middle ground major industrial automation vendors have created is gated and managed or controlled by an ecosystem of partners, but this does not unleash broad innovation. [OPAF's work] is another tipping point for industry with implications [as profound as] the introduction of Programmable Logic Controllers (PLCs) and Distributed Control Systems (DCSs)."

OPAF members and their followers believe industry may be at a new tipping point where open process automation finally becomes an innovation—and performance-improving, reality. "Tipping points" said Lydon, "spread like wildfire, changing industry." The development and adoption of open process automation may be about to generate the spark that turns into that wildfire.

"Schneider Electric is all open," said Tom Clary of Schneider Electric. "We're very actively involved in the forum to help define and develop the standard, and we've put OPAF product on our roadmap already."

Schneider Electric has an OPAF-compliant distributed control node (DCN) that it co-developed with Intel. "We announced our partnership with Intel in June 2022. This is where we are in our development. It's on the roadmap, and we will continue to push and continue to stay active in the forum to make sure that this marketplace and this ecosystem becomes viable, and adoption becomes widespread," said Clary.

Jacqueline Allen from Wood represented the system integration perspective. "We announced at the ARC forum earlier this year that we are a center of excellence for open process automation (OPA)." Being a part of that will allow Wood to develop many of the things a system integrator could offer, with training being one of them.

"We are interested in helping develop more function block libraries," said Allen. "For the last couple of decades, we've established ourselves as a system-independent integrator, and with that comes a lot of clients who are looking to make assessments. We've built up a consultancy, a practice around selection of vendor platforms. As our customers are starting to figure out what system they want next, they have already engaged with us on conversations around OPA."

Related Resources

- Open Process Automation Forum Modernizing Manufacturing
- COPA: Experts Discuss Open Process Automation Issues
- The Roles of DCS and SCADA in Digital Transformation
- Open Standards, Global Initiatives Help Modernize Operations
- Evolving Industrial Edge Computing Architecture

ABOUT THE AUTHOR

Jack Smith (jsmith@automation.com) is a contributing editor for Automation.com and ISA's *InTech* magazine. He spent more than 20 years working in industry—from electrical power generation to instrumentation and control, to automation, and from electronic communications to computers—and has been a trade journalist for 22 years.