LYNX
SOFTWARE TECHNOLOGIES

Driving the Push from the Cloud to the Edge

In this paper we will articulate the motivation for and the characteristics of the end-to-end computing, communications, and storage infrastructure which will ideally host distributed applications for Cyber-Physical Systems found in modern Automotive, Transportation, Industrial Automation deployments, among other verticals.

We will emphasize the necessary decoupling between the infrastructure and the applications running over it, in the spirit of the Software-Defined paradigms which are progressively adopted, from networking to computing to storage. We will specifically focus on the technologies required at the critical boundary between physical systems and cyber space, between Operational technologies (OT) and Information Technologies (IT), where a challenging cultural and technological convergence needs to fully unfold. The success of the Digital Transformation rests in a significant way on this complex convergence process.

I. INTRODUCTION

One of the main objectives of the Digital Transformation is the application of human and artificial "intelligence" to the management, optimization, and control of systems that touch the physical world, to achieve improvements in Fig.1, where we also show a broad spectrum of potential applications, from Industrial to Transportation, Energy, Smart Cities, and Health Care. This intelligence is derived from human knowledge and experience as well as from physical models (e.g., Digital Twins) or artificial intelligence models (e.g., neural networks). This intelligence can be hosted within resources distributed across what is called "cyber space," can be continuously enriched through physical system sensing, and is applied to the control such physical systems through some form of actuation.

This is the domain of Cyber-Physical Systems [1]. Cyber space is supported by the vast computing, networking, storage and data organization and management resources that are progressively deployed through the evolution of Enterprises and the Internet, the Telco fixed and mobile infrastructure, Cloud and Edge Comput-ing, as well as all the resources deployed in operational the environments mentioned above. It is a complex End-to-end distributed, non-homogeneous infrastructure encompassing what are known as Information Technologies (IT) and Operational Technologies (OT).

There is a critical boundary between Cyber-space and the physical world, which is called Cyber-Physical Boundary or Operational Edge. This is where the most meaningful convergence of Information Technologies (IT) and Operational Technologies (OT) is taking place. Infrastructure and applications deployed at this boundary need to satisfy critical requirements because of the time-sensitive way digital control systems interact with the systems they control, and because systems such as those shown in Fig. 1 carry huge physical and financial risk implication in case of failure or unpredictable performance.

Thus, computing, networking, and storage at the Operational Edge need to satisfy requirements not typical in the IT world, such as determinism, time sensitivity, and functional safety in order to avoid the costly — and possibly life-threatening — implications of failures, unpredictable responses, missed alarms, and security breaches of the production process.

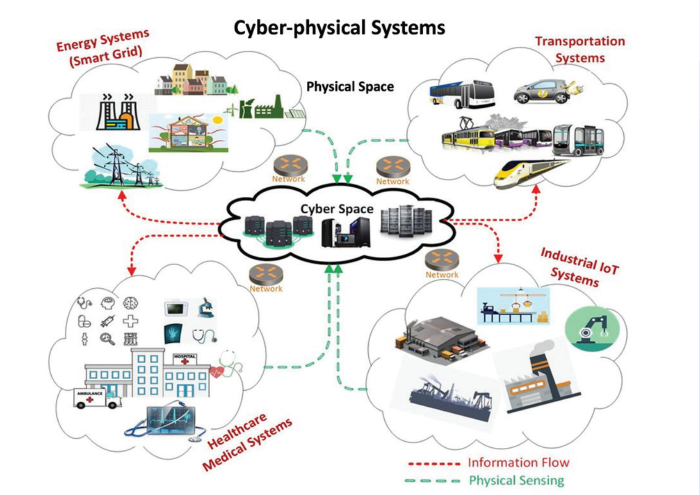

Fig. 1. Requirements at the boundary between Cyber Space and Physical Space often include time-sensitivity, determinism, physical and cybersecurity, functional safety, and reliability

Some of the expected benefits gained by this convergence of technologies at the Edge include, among others:

- the achievement of more sophisticated and intelligent behavior prediction and control of the physical systems.
- more agile and precise provisioning of such systems.
- better information interpretation from sensors associated to the physical processes for maintenance and quality control purposes.
- improved overall process efficiency through real time modeling of process and operations flows.
- real-time synthetic data creation from models that may complement physical data for AI application purposes.

II. THE SOFTWARE-DEFINED INFRASTRUCTURE FOR CYBER-PHYSICAL APPLICATIONS

Fig. 2 illustrates at the high-level the end-to-end infrastructure underpinning the complex interaction between distributed cyber intelligence and physical systems. This distributed infrastructure will be built by leveraging current pools of computing, storage, and networking resources available in Clouds, at the Service Provider Edge, in Operational Data Centers and Edges, over public and private networks, wired and wireless, all the way to Embedded Endpoints.

The infrastructure illustrated in Fig. 2 needs to offer a nonhomogeneous, distributed, orchestrated networked computing and storage fabric, hosting software applications that can be potentially deployed in a consistent way anywhere across this fabric. This requires the ubiquitous adoption of containerization and microservice models of software development and lifecycle management. This will enable the development of applications and models in the Cloud and the deployment of these applications and models across the infrastructure.

Note that as we come closer to the physical boundary, stricter timing, determinism, security, and safety requirements apply, with the resources becoming more scarce and usually more distributed, more energy and location sensitive, more exposed to physical and cyber security threats. Real-time decisions are naturally made closer the controlled Endpoints, while slower optimizations and learning naturally happen at higher levels of this infrastructure.

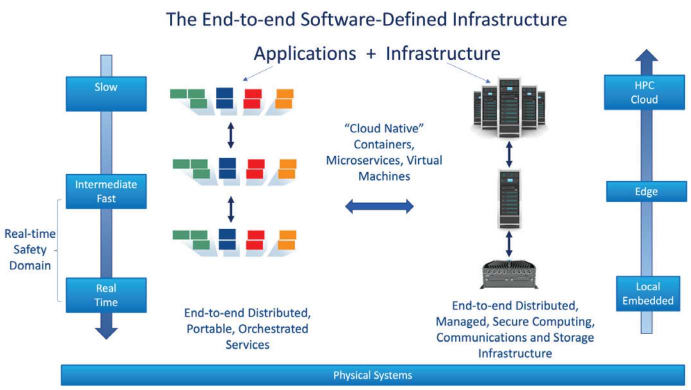
The End-to-end Software-Defined Infrastructure
Applications + Infrastructure
Slow
Intermediate Fast
Real Time
Real-time Safety Domain
"Cloud Native" Containers, Microservices, Virtual Machines
End-to-end Distributed, Portable, Orchestrated Services
End-to-end Distributed, Managed, Secure Computing, Communications and Storage Infrastructure
HPC Cloud
Edge
Local Embedded
Physical Systems

Fig. 2. The vision for an end-to-end distributed, software-defined infrastructure of computing, communications and storage resources

Software and data will need to be securely and effectively distributed across the fabric, enabling previously unthinkable cooperation and collaboration across humans and machines.

A far-reaching implication of what we are describing here is the progressive decoupling of applications and infrastructure. Applications are not bound any longer to specialized hardware components but may be moved and hosted wherever sufficient resources are available and where performance guarantees can be met. We are moving towards what is now called a Software Defined paradigm [2], [3], [4], proliferating from networking to storage to the computing infrastructure, from Clouds, Data Centers or Enterprise Networks towards the operational domain. This movement is leading to what is referred to as the Software Defined Vehicle or Software Defined Manufacturing, discussed later in this section.

Fig. 3 illustrates how intelligence residing in Embedded Endpoints, at the Edge and in Clouds, interacts with physical systems, through the creation of a hierarchy of information distribution and processing loops, characterized by different time scales. Again, learning and training happen deeper up the infrastructure, while inference models for fast decision making are active near the Endpoints. This architecture creates a powerful, adaptive and collaborative cyber space.

This virtuous cycle of data collection, analysis, and feedback can occur inside the physical system hosted by embedded computing or it can take advantage of computing capabilities distributed in the developing "continuum" involving what is known as Fog or Edge Computing [5], [6], [7], [8] resources, all the way to computing resources located in Private or Public Clouds or where High Performance Computing (HPC) is available.

Recently, much emphasis has been devoted to the deployment and impact of Intelligence (Digital Twins as well as AI) and other applications supported by Cloud computing modern intelligence closer to the boundary with the physical systems. resources. This is already bringing important benefits, because Cloud intelligence can rely on vast aggregate data as well as scalable computing and storage

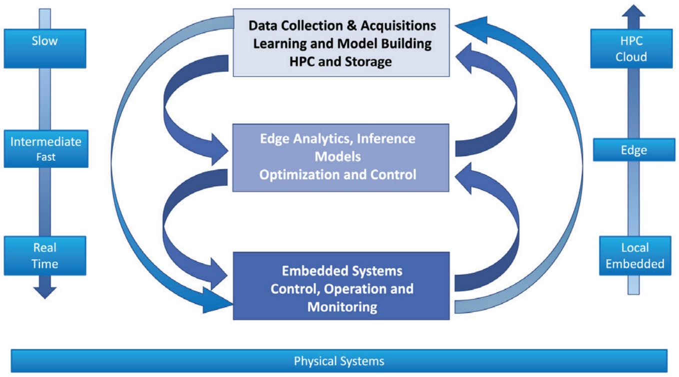

Fig. 3. Digital Twin Lifecycles require a continuous process of data extraction, analysis, model building and refinement, inference model reduction, and deployment. This process is supported by the distributed,orchestrated communications, computing, and storage resources of the underlying infrastructure

resources. With the perspective of the broad distributed objectives of Fig. 2, it is fundamental to bring more attention to the complementary deployment of modern intelligence closer to the boundary with the physical systems.

Intelligence available at the Operational Edge can respond quickly to observations and thus be more effective and enrich control, quality, and equipment reliability. This is also where data are extracted, possibly in growing amounts, and may need to be cleaned, normalized, reduced, assigned access rights, and security protected. This is where private data may need to be stored after local analysis and exposed to locally deployed intelligence.

The infrastructure perspective described here is an ambitious and will require further cultural and technological progress; however, the full realization of this vision will contribute greatly to the fulfillment of the functional and market promises of the Digital Transformation.

More concretely, let us clarify these concepts by illustrating the progress towards the distributed infrastructure discussed here in the automotive and in the industrial verticals, respectively, presented in Fig. 4 and Fig. 5. Over the past ten years, in both the automotive and the industrial domains, we have been witnessing the deployment of more powerful multicore computing in vehicles, robot, industrial floors, and industrial machines, supporting the fast evolution of more advanced sensors. Also, we have witnessed a stronger and stronger reliance on Cloud hosted applications, storage and intelligence. Finally, the demand for Edge hosted applications, storage and intelligence is also growing.

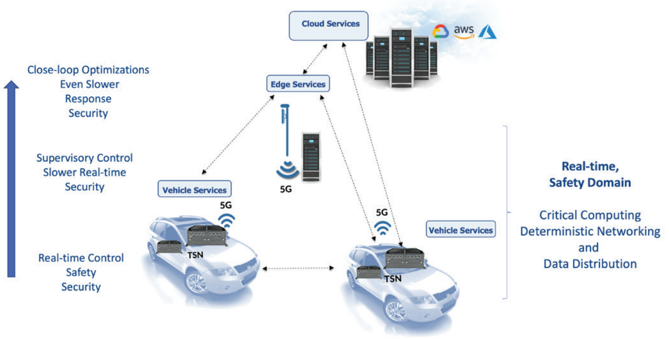

Fig. 4. End-to-end distributed computing, communications, and storage in Automotive

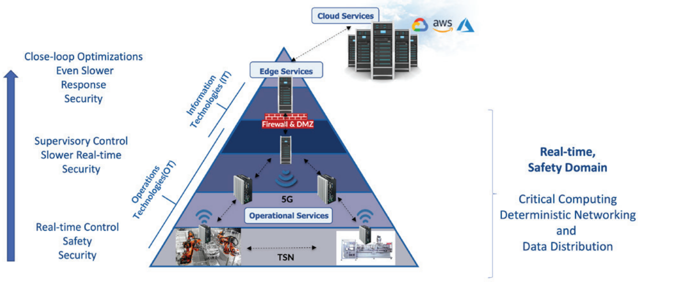

Fig. 5. End-to-end distributed computing, communications, and storage in Industrial

The full realization of the Software Defined infrastructure described here in automotive and industrial will require a deeper convergence between IT and OT technologies and practices. This convergence is dramatically accelerating in the Automotive vertical, while it is progressing more slowly, although in the same direction, in the Industrial vertical.

In the automotive vertical we are witnessing the adoption of recent advances developed in IT, such as virtualization, data distribution, analytics, digital twins [9], and software lifecycle management with remote over the air update. The progress towards a Software Defined vehicle architecture in now irreversible. A key enabler of this evolution is the progressive adoption of virtualization from the Cloud all the way to various manifestations of the Edge, whether in telco Edge deployments behind 5G towers or in cities and roadsides, supporting transportation and autonomy relevant services. The internal vehicle architecture, characterized until recently by a proliferation of small and poorly networked ECUs, is quickly consolidating its computing infrastructure around a few virtualized larger controllers, supporting applications of mixed criticality and networked via high bandwidth time sensitive Ethernet (TSN) [10].

This dramatic evolution is illustrated in Fig. 6, where we show distributed services (Cloud, Edge, Vehicle) supported by virtualized computing nodes across the distributed infrastructure. In order to satisfy the critical requirements at the cyber-physical boundary, the virtualization technologies more appropriate for the IT domain need to be complemented or replaced by more Mission Critical Virtualization, evolving from embedded, real-time technologies, which will be discussed more in detail in a later section. The implications of this new automotive infrastructure are enormous.

A slower but consistent infrastructure evolution is happening in the Industrial sector. Industrial Operations have traditionally been the domain of embedded, often mission critical electronic systems, built on microcontrollers, Programmable Logic, or Numerical Controllers other types of deterministic controllers, running real-time Operating Systems, supported by ruggedized Windows Industrial PCs,

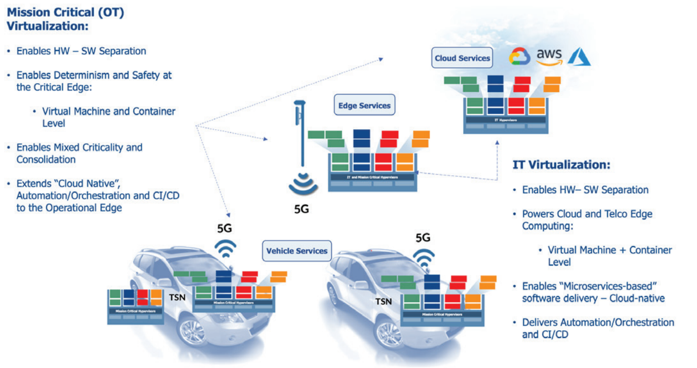
Mission Critical (OT)
Virtualization:
• Enables HW – SW Separation
• Enables Determinism and Safety at the Critical Edge:
• Virtual Machine and Container Level
• Enables Mixed Criticality and Consolidation
• Extends "Cloud Native", Automation/Orchestration and CI/CD to the Operational Edge
Cloud Services
aws
IT Hypervisors
Edge Services
IT and Mission Critical Hypervisors
5G
5G
5G
Vehicle Services
TSN
TSN
Mission Critical Hypervisors
IT Virtualization:
• Enables HW– SW Separation
• Powers Cloud and Telco Edge Computing:
• Virtual Machine + Container Level
• Enables "Microservices-based" software delivery – Cloud-native
• Delivers Automation/Orchestration and CI/CD

Fig. 6. Virtualization (IT and OT) is a key infrastructure enabler for Software Defined Automotive

often un-managed and minimally interconnected. This traditional infrastructure is not conducive to the insertion and dynamic evolution of software innovations inspired by recent IT advances in Data Analytics, AI, etc. It is static, and the software functionality is usually rigidly associated with a specific and fixed piece of hardware. Each hardware-software system is usually managed by its operator, locally. For security reasons, the Operational Edge is, in general, strictly isolated from the IT world via Firewalls and Demilitarized Zones.

The Industrial floor is currently undergoing a dramatic transition towards a new electronic architecture inheriting some of the technological progress experienced in the evolution of cloud computing, software defined networking, cloud storage and object data bases, big data, software deployment and orchestration, as well as security. The new infrastructure needs to adopt other innovations in virtualization, communications and networking, storage security, and management that specifically address the unique Operational Edge requirements.

We will progressively move towards an Edge architecture like the one depicted in Fig. 7, which is the illustration of a distributed, interconnected set of non-homogeneous virtualized computing and storage nodes, a "system of systems," a distributed computing fabric over which software applications are deployed, interworked and orchestrated. These nodes host critical (i.e., time-sensitive, predictable, reliable, safe, etc.) processes close to the interface with physical endpoints. Applications satisfy less critical requirements as they are operating further away from that interface. Data are distributed, processed, and stored at various locations across this fabric.

III. THE SOFTWARE-DEFINED INFRASTRUCTURE AT THE EDGE: MODERN EDGE COMPUTING, DATA NETWORKING, AND STORAGE

The full potential and maturity of the distributed infrastructure described in this paper requires a complex but promising technology convergence, ultimately requiring the integration of key elements, some of which are still maturing, including deterministic networking and data distribution (such as IEEE TSN and 5G), secure, safe and real-time capable virtualization, with software deployment centered on containers and microservices, time-sensitive data analytics and AI, digital twins and simulation, and distributed and QoS aware system management and orchestration. In the following sections, we will discuss achievements and work in progress in some of the principal components of the forthcoming distributed edge infrastructure, along the following dimensions:

a. Computing: Mission critical computing at the edge, supporting mixed criticality applications, with strong security and safety.

b. Data Networking: Time-sensitive, reliable networking technologies like TSN and Private 5G/6G.

c. Data Distribution: Time-sensitive, reliable Data Distribution Middleware, including OPC UA over TSN and DDS over TSN.

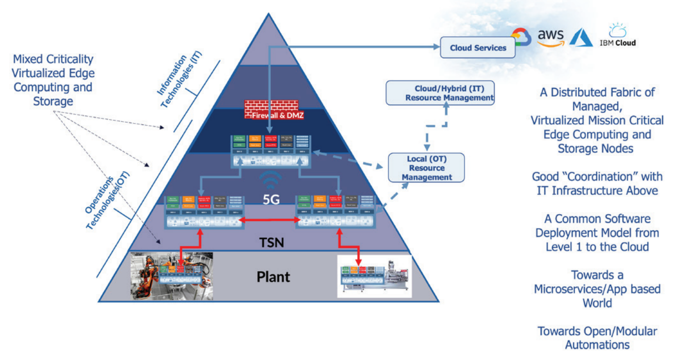
Mixed Criticality Virtualized Edge Computing and Storage
Information Technologies (IT)
Operations Technologies(OT)
Firewall & DMZ
5G
TSN
Plant
Cloud Services
aws
IBM Cloud
Cloud/Hybrid (IT) Resource Management
Local (OT) Resource Management
A Distributed Fabric of Managed, Virtualized Mission Critical Edge Computing and Storage Nodes
Good "Coordination" with IT Infrastructure Above
A Common Software Deployment Model from Level 1 to the Cloud
Towards a Microservices/App based World
Towards Open/Modular Automations

Fig. 7. Virtualization (IT and OT) and Management are key infrastructure enablers for Software Defined Industrial

d. Data Management and Storage: Data are the key motivation and driver for the Digital Transformation and requires distributed management, distributed structuring and storage in file systems, object and relational databases, and streaming databases.

e. End-to-end Management and Orchestration: A common software deployment model, from Clouds to endpoints, enabling a modern continuous integration and continuous delivery (CI/CD) software lifecycle management.

Once this innovative infrastructure gets deployed, the emphasis should shift towards the application layer, and, in particular, on the modeling description technologies that may enable an expanding community to develop, validate, and deploy distributed applications in the operational domain with greater agility and lower cost.

Existing programming languages, toolsets, and platforms have proven to be limited in dealing with the overwhelming technical challenges posed by the description of complex models for cyber-physical system behaviors that should be capable of executing in sync or even faster than the physical system being modeled. Promising technologies in this fundamental area are on the horizon and will help scaling the programmer community and reduce the cost of system modeling. Much more progress in this area is needed.

A. Edge Computing and Mission Critical Edge

Computing One of the key promises of Edge Computing is precisely to provide hosting of new applications, such as AI and digital twins, near physical systems, closely integrated with and enhancing typical monitoring, control, and supervisory applications running today's operations. The hope is also in a new edge infrastructure management and application deployment model, evolving from the lessons learned in the IT world and, particularly, in Cloud Computing (e.g., remote management, microservices, orchestration, etc.).

There is a parallel imperative to be mindful of several key “mission critical” characteristics of the manufacturing floor when introducing any such innovative Edge Computing functionality. Quality of Service (QoS), availability, determinism, security, and (ultimately) Functional Safety (FuSa) requirements must be satisfied to avoid the costly — and possibly life-threatening — implications of failures, unpredictable responses, missed alarms, and security breaches of the production process.

Mission Critical Edge Computing, based on a class of thin hypervisors such as those delivered by Lynx Software [11], Green Hills [12], Wind River [13], or from recent Open Source Fig. 7. Virtualization (IT and OT) and Management are key infrastructure enablers for Software Defined Industrial effort such as ACRN [14], Jailhouse [15], and others, fulfill these demands. These technologies deliver partitioning and monitoring of all system resources on multi-core computing nodes with optionally immutable allocation of critical resources to more critical subjects. They provide efficient and controllable inter-partition communications with real-time support and strong security separation. They can also support safety critical applications, while providing the option of dynamically allocating non-critical resources.

Some of these hypervisors are portable across a broad range of hardware platforms including x86 and ARM. They can be adopted in industrial systems in applications ranging from embedded controllers to high-performance servers, turning them into Mission Critical Edge nodes. A technology of this kind allows systems architects to subdivide systems into smaller independent partitions and stacks, as illustrated in Fig. 8. This promotes the design of more consolidated, traceable, and efficient architectures reducing the development overhead associated with integration and enabling security and safety validation.

The most appropriate architecture for a mission critical hypervisor is based on the concept of a Separation Kernel [16]. Lynx Software LynxSecure is one of the most mature implementations of such an architecture. With reference to Fig.8, the hypervisor supports multiple VMs (or “subjects”) which host different types of guest operating systems including Linux, Windows, traditional fully featured real-time operating systems (RTOSes), minimal scheduler-like RTOSes, safety certifiable OSs and applications, and true bare-metal applications. Fig. 8 illustrates some of the options available for internal and external connectivity to I/Os and devices. Efficient shared-memory-based peer-to-peer connections can be configured across VMs, while device sharing across VMs is enabled using bridged connections. These are implemented via a special Linux VM that hosts device drivers, bridging, management, and other functions.

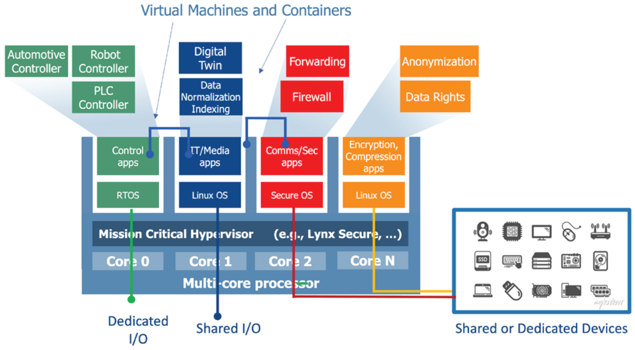

Fig. 8. Mission Critical Hypervisors, evolving from the real-time embedded application towards modern Edge Computing are a foundational element of the infrastructure, within Endpoints, at the Edge, but also at the Network Edge and in Clouds.

The hypervisor has no visibility into what is happening above the guest OSs. Specifically, while containerization can be supported inside a Linux OS, the management of containers is performed by complementary functionality, offered by partners or Cloud Service Providers, as discussed in a later section.

B. Data Communications

Networking and Data Distribution In this section, we will discuss both the evolution of networking in the Operational Edge and the all-important Data Distribution middleware technology, simplifying the development and deployment of distributed and interworking applications, such as Digital Twins.

1) Networking: Towards TSN and 5G Networking at the operational edge needs to satisfy the same type of critical requirements on

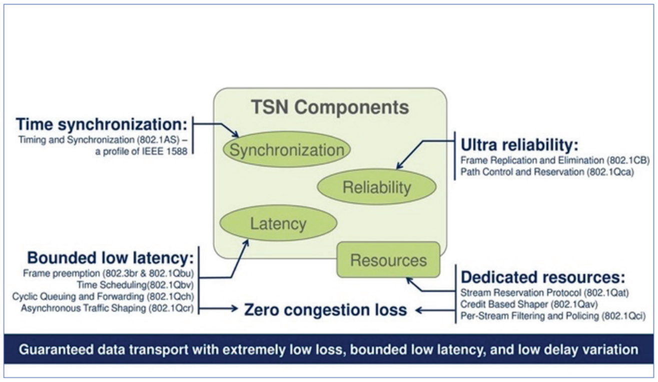
TSN Components
Time synchronization:
Timing and Synchronization (802.1AS) – a profile of IEEE 1588
Synchronization
Ultra reliability:
Frame Replication and Elimination (802.1CB)
Path Control and Reservation (802.1Qca)
Reliability
Latency
Resources
Bounded low latency:
Frame preemption (802.3br & 802.1Qbu)
Time Scheduling(802.1Qbv)
Cyclic Queuing and Forwarding (802.1Qch)
Asynchronous Traffic Shaping (802.1Qcr)
Zero congestion loss
Dedicated resources:
Stream Reservation Protocol (802.1Qat)
Credit Based Shaper (802.1Qav)
Per-Stream Filtering and Policing (802.1Qci)
Guaranteed data transport with extremely low loss, bounded low latency, and low delay variation

Fig. 9. Summary of requirements and standardization efforts for IEEE Time-Sensitive Networking (TSN)

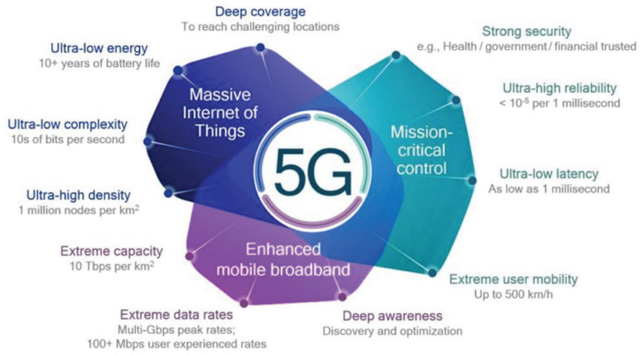

Fig. 10. Summary of requirements and standardization efforts for 5G

Quality of Service (QoS), availability, determinism, security, and (ultimately) Functional Safety (FuSa) listed above with respect to edge computing. These requirements, diverging from the typical requirements applying to IT networking, have led to the proliferation of nonstandard, often regionalized networking technologies deployed in the Operational domain. In the past 20 years, the adoption of Ethernet, WiFi, Bluetooth, and Cellular, although not yet in a fully standard fashion, has moved IT and OT networking closer. In the past years, we have seen this process converge with the ongoing effort towards the adoption of operational technologies which are also based on IT standards, namely Ethernet with IEEE TSN [10] for wired networking and 5G [17] for wireless networking. The two technologies are deeply complementary and are cross-pollinating in their developing standards. Fig. 9 and Fig. 10 summarize the functional objectives of these technologies.

These developments will lead not only to a more open infrastructure, easier interoperability, and a more seamless transition between operational and IT networking domains but also across the layers of the operational domain, as depicted in Fig. 6 and Fig. 7 above. This is most important since the foreseen deployment of applications such as digital twins and other forms of models requires a flexible data distribution from end points to computers hosting the various levels in a hierarchy of models. Also, timing requirements may now need to apply beyond the area of field buses close to the endpoints, the domain of machine control (Level 1 in the Industrial Automation pyramid), and extend to higher levels, such as Level 2 (Supervisor Monitoring and Control) and even Level 3. Also, based on a solid, standard timing distribution architecture, such as that provided by TSN and 5G, it will be possible to create a distributed fabric of applications that work in synchrony with each other and also with the physical systems they control.

2) Data Distribution Middleware at the Edge The next technological element required in the operational infrastructure at the Edge in order to effectively enable, standardize, scale, and operate the deployment of distributed applications at the operational edge is the Data Distribution middleware.

With reference to multiple figures above, we envision a hierarchical distribution of applications synchronized and collaborating through the distribution of data over the underlying network. Data distribution middleware is designed to abstract away from the application designer the task of distributing data, of making sure the data are received in a timely and reliable way, and of assuring the data received are understandable across the interoperating application (semantic interoperability).

Again, Data Distribution at the operational edge needs to satisfy the same type of critical requirements on Quality of Service (QoS), availability, determinism, security, and (ultimately) Functional Safety (FuSa) listed above with respect to edge computing and networking.

Communication middleware has grown, mostly in the IT domain over the years, to address a number of application areas. The choice of appropriate candidates, satisfying the requirements of the operational edge, is extremely limited, and much work remains to be done in this area. Two key candidates are at the forefront of these evolution, OPC UA [18], and DDS [19]. As indicated in Fig. 11 and Fig. 12, they differ in their core architectures, but both are based fundamentally on the Publish- Subscribe data exchange model. In order to address the time sensitivity requirement, they are both in the middle of a standardization effort to integrate over TSN or over 5G.

The models of data distribution required across applications at the operational edge need to go well beyond a publish subscribe model, even when it becomes time sensitive when bolted over TSN. For example, with the current data distribution standards (OPC UA and DDS), there is no way to share the "ownership" of a data object across multiple applications which would continue to modify the shared data in a collaborative and distributed mode while enforcing a fair sequencing policy of updates that would maintain coherency. This model of data sharing would be needed if we want a collaborative activity across distributed systems, looking like "Machine Conferencing," where multiple systems coherently work on the same data, e.g., the position or status of an object multiple machines work on together. Promising advances in the direction of more functional Data Distribution and management at the edge are starting to surface.

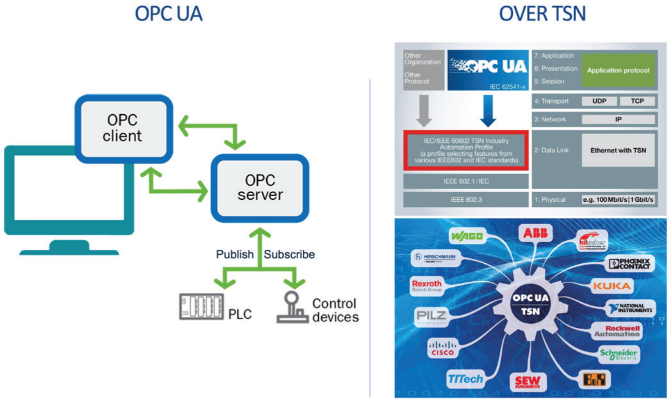

Fig. 11. Industrial Data Middleware today – OPC UA over TSN

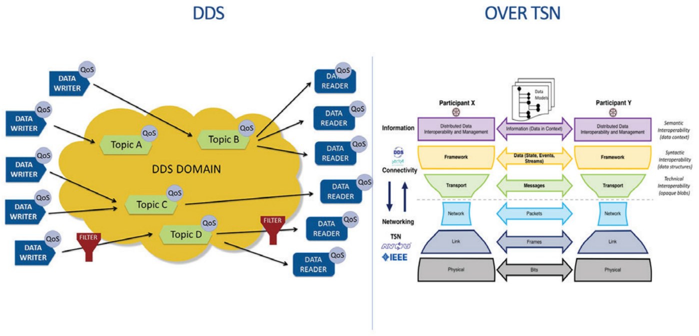
DDS
OVER TSN
DATA WRITER
DATA READER
QoS
Topic A
Topic B
Topic C
Topic D
DDS DOMAIN
FILTER
Participant X
Participant Y
Data Models
Information
Distributed Data Interoperability and Management
Information (Data in Context)
Semantic Interoperability (data context)
DDS
Connectivity
Framework
Data (State, Events, Streams)
Syntactic Interoperability (data structures)
Transport
Messages
Technical Interoperability (opaque blobs)
Networking
Network
Packets
TSN
AVNU
IEEE
Link
Frames
Physical
Bits

Fig. 12. Industrial Data Middleware today – Data Distribution System (DDS) over TSN

Fig. 6 and Fig. 7 illustrate a distribution of applications deployed as VMs or containers/microservices, including a hierarchy of interworking models, hosted on mission critical edge nodes across the operational edge of an industrial floor. In those figures we show a data distribution fabric characterized by a number of data exchange modes, including traditional clientserver and publish-subscribe, but also alternative modes such as "broadcast," "exchange," a transactional mode between two entities, and "coordinate," the coherent conferencing-like data sharing mentioned above. More innovation is needed in this area of technology.

C. Edge Data Management and Storage

One of the key drivers for Edge Computing is the need for storage at the Edge, motivated by the growing volume of traffic generated by the endpoints, by the cost and time penalty involved in moving this data to the Cloud, by privacy concerns, and by the responsiveness advantages coming from storing the data, processing it, and deriving insights and action indications from it close to where the data originated. The requirements on memory bandwidth, latency, determinism, and reliability imposed by streaming analytics, AI, and Digital Twins at the Edge are producing deep innovations in storage technologies. Storage will be deployed with Hyperconverged Computing (HCI) in on-premises Data Centers or even closer to the endpoints, where Industrial servers may be required.

Storage featuring mission critical characteristics will be associated with networking and computing resources deeper in the Operational Edge and in sensors and actuators on low level controllers. These modern, fast, dense, reliable, storage resources will mostly consist in Solid State Drives (SSD), well suited to the environmental requirements at the Operational Edge. Memory technologies faster than today's SSD, such as Intel's Optane [20] will help accelerate analytics, AI, and Digital Twin processing. Large DRAM based designs will lead to further improvements in edge application performance and help support more deterministic workloads.

The envisioned Edge infrastructure, with its distributed, nonhomogeneous nature will lend itself to the application of Software Defined Storage concepts [21] which decouple the software managing distributed non-homogeneous storage elements from the hardware hosting such storage. These innovative storage architectures are stimulating much progress in the areas of distributed streaming Data Bases, SQL and Object Data Bases, and File Systems.

With reference to Fig. 13, two key innovations which will provide important performance and virtualization advances highly relevant for Edge. Computing will be NVMe [22] capable SSD drives that also will provide support for SR-IOV [23]. Together, these technologies are enabling a higher bandwidth to storage over PCIe, compared with legacy SATA drives, providing multiple QoS levels accessing storage necessary to achieve predictable access times, and slicing a single drive into multiple virtual drives.

To be useful across distributed applications, data stored across the Edge as described above will also need to be processed, polished, modeled and secured as it is communicated between endpoints and distributed applications. Data at the Edge is organized and structured in Distributed File Systems, lighter weight Object and Relational Databases, as well as in Streaming Databases.

D. Edge Management and Orchestration

As discussed in previous sections, one of the key advances and benefits brought about by the evolving Software Defined infrastructure, particularly at the edge, is its manageability inherited and developed based on management functionality deployed in IT and, specifically, in Cloud Computing.

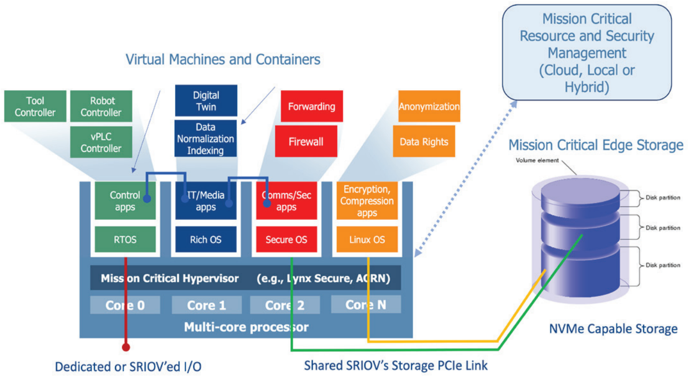
Mission Critical Resource and Security Management (Cloud, Local or Hybrid)
Virtual Machines and Containers
Tool Controller
Robot Controller
vPLC Controller
Digital Twin
Data Normalization Indexing
Forwarding
Firewall
Anonymization
Data Rights
Mission Critical Edge Storage
Volume element
Control apps
IT/Media apps
Comms/Sec apps
Encryption, Compression apps
RTOS
Rich OS
Secure OS
Linux OS
Disk partition
Disk partition
Disk partition
Mission Critical Hypervisor (e.g., Lynx Secure, ACRN)
Core 0
Core 1
Core 2
Core N
Multi-core processor
NVMe Capable Storage
Dedicated or SRIOV'ed I/O
Shared SRIOV's Storage PCIe Link

Fig. 13 illustrates a Mission Critical Edge Computing node complemented by an SSD drive featuring NVMe and SR-IOV. This combination brings to the Operational Edge high performance, reliability, and capacity storage, supporting critical virtualized computing.

Management and orchestration functionality and interfaces at various points in the infrastructure offer remote and aggregate control on resources from the hardware device level to the networking and storage configuration, to the virtualization at the VM and Container level, to the performance and security configuration and monitoring levels. The high level integrated end-to-end management architecture discussed here is illustrated in Fig. 7 above.

Cloud native applications can be deployed, monitored, updated, interconnected, and orchestrated across the end-to-end infrastructure supporting a CI/CD (Continuous Integration/Continuous Delivery) model across the entire distributed infrastructure. These powerful capabilities will reduce operations costs, improve efficiency, performance, and security enabling dynamic upgrades for product-by-product customization or subsystem repurposing.

In the perspective of the deployment of intelligence and other applications at the Edge, this modern infrastructure will facilitate the dynamic deployment, interconnection, monitoring, and continuous refinement of the models which is critical for the successful application of this technologies.

The management layer offered by the new Operational Infrastructure will be an essential enabler towards the implementation of well instrumented, monitored, and supervised complex cyber-physical and industrial structures dynamically evolving as natural organisms and providing visibility and intelligent control on critical systems such as energy production and distribution, manufacturing, and transportation systems.

IV. SUMMARY AND CONCLUSIONS

In this chapter, we have presented a contribution to the understanding of the end-to-end infrastructure required to support evolving cyber-physical systems and the Digital Transformation. We highlighted the importance of decoupling applications from infrastructure hosting, interconnecting, executing, storing, and managing them. This decoupling is consistent with the powerful Software Defined paradigm.

We have also motivated the importance of the component of the infrastructure positioned between physical endpoints and the more traditional cyber space (IT) infrastructure. This is the Edge where IT and OT technologies deeply intertwine in their requirements and functionality. This is the component of the infrastructure that will enable richer intelligence to impact the physical systems and contribute to their efficient control. The successful evolution of the technologies empowering this Edge component of the infrastructure will significantly influence the future of the Digital Transformation.

REFERENCES

[1] National Science Foundation, “Cyber-Physical Systems: Enabling a Smart and Connected World,” National Science Foundation, 2022. Fig. 13. Virtualized Edge storage with NVMe and SRIOV is a natural match for mission critical virtualized computing and networking [Online]. Available: https://www.nsf.gov/news/special_reports/cyberphysical/. [Accessed May 2022].

[2] D. Kreutz, F. Ramos, P. Verissimo, C. Rothenberg, S. Azodolmolkyand S. Uhlig, “Software-Defined Networking: A Comprehensive Survey,” 8 October 2014. [Online]. Available: https://arxiv.org/pdf/1406.0440.pdf. [Accessed May 2022].

[3] Wikipedia, “Software-defined storage,” 2022. [Online]. Available: https://arxiv.org/pdf/1406.0440.pdf. [Accessed May 2022].

[4] M. Raza, “What Is SDI? How Software Defined InfrastructureWorks,” 2018. [Online]. Available: https://arxiv.org/pdf/1406.0440.pdf. [Accessed May 2022].

[5] F. Bonomi, R. Milito, J. Zhu and S. Addepalli, “Fog Computing and Its Role in the Internet of Things,” in Mobile Cloud Computing, Helsinki, 2012.

[6] F. Bonomi, R. Milito, P. Natarajan and J. Zhu, “Fog Computing: A Platform for Internet of Things and Analytics,” in Big Data and Internet of Things: A Roadmap for Smart Environments, Switzerland, Springer International, 2014, pp. 169-186.

[7] Open Fog Consortium, “Industry IoT Consortium,” February 2017. [Online]. Available: https://www.iiconsortium.org/pdf/OpenFog_Reference_Architecture_2_09_17.pdf . [Accessed May 2022].

[8] IBM, “What is edge computing?,” [Online]. Available: https://www.ibm.com/cloud/what-is-edge-computing. [Accessed May 2022].

[9] N. Kochar, “Digital Twins in Automotive,” in Digital Twin, in press Springer Verlag.

[10] IEEE, “IEEE Standard for Local and Metropolitan Area Networks --Timing and Synchronization for Time-Sensitive Applications,” 19 062020. [Online]. Available: https://standards.ieee.org/ieee/8802-1AS/10767/802.1AS/7121/.

[11] Lynx Software Technologies, “Lynx Software Technologies,” 2022. [Online]. Available: https://www.lynx.com/. [Accessed May 2022].

[12] Green Hills Software, “Green Hills Software,” 2022. [Online]. Available: https://www.ghs.com. [Accessed May 2022].

[13] Windriver, “Windriver,” 2022. [Online]. Available: https://windriver.com. [Accessed May 2022].

[14] The Linux Foundation Projects, “ACRN - A Big Little Hypervisor for IoT Development,” 2022. [Online]. Available: https://projectacrn.org. [Accessed 2022].

[15] Texas Instruments, “Jailhouse Hypervisor,” 2020. [Online]. Available: https://software-dl.ti.com/processor-sdklinux/esd/docs/06_03_00_106/linux/Foundational_Components/Virtualization/Jailhouse.html. [Accessed 2022].

[16] T. Loveless, “Lynx Software Technologies: What is a Separation Kernel?,” 2020. [Online]. Available: https://www.lynx.com/embedded-systems-learning-center/what-is-aseparation-kernel. [Accessed May 2022].

[17] ACIA 5G, “Integration of 5G with Time-Sensitive Networking for Industrial Communications,” 02 2021. [Online]. Available: https://5gacia. org/whitepapers/integration-of-5g-with-time-sensitivenetworking-for-industrial-communications/.

[18] OPC Foundation, “Unified Architecture,” [Online]. Available: https://opcfoundation.org/about/opc-technologies/opc-ua/.

[19] DDS Foundation, “What is DDS,” 2021. [Online]. Available:https://www.dds-foundation.org/what-is-dds-3/. [Accessed May 2022].

[20] Intel, “Intel Optane Technology: Revolutionizing Memory and Storage,” 2022. [Online]. Available: https://www.intel.com/content/www/us/en/architecture-andtechnology/intel-optane-technology.html. [Accessed May 2022].

[21] IBM, “Software-defined storage,” 2022. [Online]. Available: https://www.ibm.com/storage/software-definedstorage?utm_content=SRCWW&p1=Search&p4=43700060595209662&p5=e&gclid=Cj0KCQjw0umSBhDrARIsAH7FCoeKzWuV84MEygoXRICXI1KfHOf3FrYFWzAf8RsViWeHW9wKoOtrzEaAgFkEALw_wcB&gclsrc=aw.ds. [Accessed May 2022].

[22] R. Gupta, “Western Digital Blog: What is NVMe and why is it important? A Technical Guide,” 2020. [Online]. Available:https://blog.westerndigital.com/nvme-important-data-drivenbusinesses/. [Accessed May 2022].

[23] T. Loveless, “Lynx Software Technologies: What is SR-IOV and Why is it Important for Embedded Devices,” 2019. [Online]. Available: https://www.lynx.com/embedded-systems-learning-center/ what-is-sr-iov-and-why-is-it-important-for-embedded-devices. [Accessed May 2022].

1.800.255.5969

Lynx Software Technologies, Inc.
855 Embedded Way
San Jose, CA 95138-1018
+1 (800) 255-5969
+1 (408) 979-3900
+1 (408) 9793-920 fax
inside@lynx.com
www.lynx.com

Lynx Software Technologies UK
400 Thames Valley Park Drive
Thames Valley Park
Reading, RG6 1PT
United Kingdom
+44 (0) 118 965 3827
+44 (0) 118 965 3840 fax

Lynx Software Technologies France
38 Avenue Pierre Curie
78210 Saint-Cyr-l’École
France
+33 (0) 1 30 85 06 00
+33 (0) 130 85 06 06 fax