EUROPEAN JOURNAL OF EMERGING REAL-TIME IOT AND EDGE INFRASTRUCTURES

THE 6G CONTINUUM: A PLATFORM ARCHITECTURE FOR REAL-TIME INDUSTRIAL DIGITAL TWINS

Dr. Elena J. Foster
Department of Communication, University of Montana, Missoula, MT, USA

Dr. Marcus T. Delgado
Department of Political Science, University of South Alabama, Mobile, AL, USA

V0LUME01 ISSUE01 (2024)
Published Date: 30 December 2024 // Page no.: - 69-76

ABSTRACT

The convergence of next-generation wireless technologies, particularly 6G, with the Internet of Things (IoT), edge computing, and cloud infrastructure is set to revolutionize industries by enabling ultra-responsive and intelligent applications. A key application in this domain is the Digital Twin (DT), which requires a seamless and powerful computational continuum to operate in real-time. This paper presents a novel, flexible, and hyper-distributed platform that spans the IoT-Edge-Cloud continuum, designed specifically to support real-time DT applications for the logistics and industrial sectors. We detail the architecture of this platform, which leverages a multi-tiered approach to computation and data processing, and its implementation on a 6G-intended testbed. The platform's design addresses the critical challenges of latency, data throughput, and scalability inherent in large-scale industrial environments. We validate our approach through two specific use cases: a smart logistics scenario and an industrial automation process. The results from our testbed demonstrate the platform's capability to meet the stringent Key Performance Indicators (KPIs) required for real-time DT operations, such as ultra-low latency and high reliability, paving the way for the next generation of cyber-physical systems.

Keywords: 6G; Digital Twin; IoT; Edge Computing; Cloud Computing; Hyper-Distributed Systems; Real-Time Systems; Logistics; Industry 4.0; URLLC; Testbed.

INTRODUCTION

The 6G Vision and the Rise of Cyber-Physical Systems

The global vision for the sixth generation (6G) of wireless communications extends far beyond simple connectivity, aiming to create a fabric that seamlessly integrates the physical, digital, and human worlds [1, 2]. This future is predicated on the ability to support novel services and applications that demand unprecedented levels of performance, including ultra-reliable low-latency communication (URLLC), massive machine-type communications (mMTC), and enhanced mobile broadband (eMBB) [3, 31, 32]. This paradigm shift moves beyond the human-centric focus of 5G to a more holistic, machine-centered ecosystem where trillions of interconnected devices form the backbone of our society [4]. Central to this vision is the proliferation of the Internet of Things (IoT), which will generate vast amounts of data, creating immense economic value and enabling intelligent automation across all sectors.

To manage this data deluge and enable the real-time responsiveness required by applications like autonomous systems [5], industrial robotics [6], and immersive augmented/virtual reality [7], a fundamental shift in computing architecture is necessary. The traditional centralized cloud computing model, while powerful, is often inadequate for time-sensitive operations due to the inherent latency introduced by transmitting data over long distances and the potential for network bottlenecks [8, 9, 33, 34]. Consequently, a new model, the IoT-Edge-Cloud continuum, has emerged. This model distributes computation, storage, and intelligence from the centralized cloud to the network's edge, placing resources closer to where data is generated and consumed [20, 48]. This hyper-distributed environment [11] promises to deliver the low-latency and high-bandwidth services essential for the next wave of innovation.

1.2. Digital Twins as a Cornerstone of Industrial Transformation

Among the most promising applications enabled by this new paradigm is the Digital Twin (DT). A DT is a high-fidelity virtual representation of a physical object, process, or system that is dynamically updated with real-world data from its physical counterpart [10]. This creates a closed loop between the physical and digital worlds, allowing for advanced monitoring, simulation, prediction, and control [12, 35]. The concept of DTs is being applied across numerous domains, proving to be a transformative technology in manufacturing (cyber-physical production systems) [13], personalized healthcare (human digital twins) [14], smart cities (intelligent traffic management) [15], and autonomous robotics [16].

The successful implementation of a real-time DT is fundamentally dependent on the underlying communication and computation infrastructure. The platform must be capable of synchronizing the physical and virtual worlds with minimal delay, processing massive streams of sensor data, and executing complex simulations. While many proprietary DT solutions exist, developed by major industrial players like Siemens [17], General Electric [18], and IBM [21], there is a growing need for open, flexible, and scalable platforms that can be adapted to a wide range of use cases and avoid vendor lock-in [22].

1.3. The Need for Open, Hyper-Distributed Platforms

Open-source frameworks and technologies are beginning to fill this gap, offering modular and interoperable solutions. Platforms like Eclipse Ditto [23], FIWARE [24, 25], and OpenTwins [26] provide the building blocks for creating customizable DT applications. However, integrating these components into a cohesive, high-performance platform that meets the stringent requirements of 6G and supports complex, real-time DTs remains a significant challenge [27, 28, 29]. The orchestration of resources across the highly heterogeneous and geographically distributed IoT-Edge-Cloud continuum is a particularly complex problem that requires sophisticated management and automation [46]. Existing testbeds have begun to explore these integrations [30], but there is a clear need for comprehensive platforms designed for real-world industrial and logistics scenarios.

This paper addresses these challenges by proposing and validating a flexible, hyper-distributed platform designed specifically to support real-time DT applications. We present the detailed architecture and implementation of this platform on a 6G-intended testbed, demonstrating its capabilities through practical logistics and industrial use cases. Our work makes the following key contributions:

- A Novel Architecture: We propose a modular and open architecture for a hyper-distributed IoT-Edge-Cloud platform that leverages a multi-tiered design for optimal performance and scalability.
- 6G Testbed Integration: We detail the integration of this platform with a 6G-ready testbed, validating its performance against the demanding communication requirements of next-generation networks.
- Real-World Validation: We deploy and evaluate two real-time DT applications for logistics and industry, showcasing the platform's effectiveness and tangible benefits in practical scenarios.
- Performance Analysis: We provide a comprehensive performance analysis that demonstrates the platform's ability to achieve the ultra-low latency, high throughput, and high reliability required for time-sensitive applications.

The remainder of this paper is structured as follows: Section 2 describes the materials and methods used to design and build the platform and testbed. Section 3 presents the results of our experimental evaluation. Section 4 discusses the implications of these results and compares them to the state-of-the-art. Finally, Section 5 concludes the paper and outlines critical open issues and directions for future work.

2. Materials and Methods

The development and validation of our platform involved a multi-faceted approach, encompassing the detailed architectural design of the hyper-distributed system, the strategic selection and integration of enabling technologies, the setup of a comprehensive testbed for validation, and the design of specific use cases to test the platform's capabilities.

2.1. System Architecture Design

The proposed platform is built upon a multi-tiered architecture that mirrors the IoT-Edge-Cloud continuum. This architecture is designed to be modular, scalable, and flexible, allowing for the dynamic allocation of computational tasks based on the specific requirements of the application, such as latency, processing power, and data privacy.

2.1.1. IoT-Device Layer

This is the foundational layer where data originates and physical actions are executed. It consists of a diverse and heterogeneous range of physical sensors, actuators, and connected devices deployed in the target environment (e.g., a logistics warehouse or an industrial production line).

- Sensing and Actuation: Devices include high-definition cameras, LiDAR sensors for 3D mapping, Inertial Measurement Units (IMUs), GPS for localization, temperature and vibration sensors for condition monitoring, industrial robotic arms, and autonomous mobile robots (AMRs).
- Communication Protocols: To ensure interoperability in this heterogeneous environment, we employ standardized communication protocols. For low-power, constrained devices, we utilize protocols like Constrained Application Protocol (CoAP) [39], which is optimized for efficiency. For devices requiring higher throughput and reliable messaging for critical data, we use Message Queuing Telemetry Transport (MQTT) [37]. For service-to-service communication and API access, HTTP is also utilized where appropriate [38].
- Data Abstraction: Data is structured using a modular, ontology-based framework to ensure semantic interoperability [36]. This allows data from different types of devices to be understood and processed uniformly by

the higher layers, which is crucial for building cohesive Digital Twin models.

2.1.2. Edge Computing Layer

The edge layer acts as a powerful intermediary between the IoT devices and the central cloud, bringing computation and intelligence closer to the data source.

- Distributed Nodes: It is composed of multiple edge nodes with varying computational capacities, from lightweight single-board computers (e.g., Raspberry Pi, NVIDIA Jetson) for simple tasks to more powerful edge servers for complex analytics. These nodes are strategically deployed to minimize latency.

- Core Functions: The primary function of the edge layer is to perform low-latency data processing, real-time analytics, data filtering and aggregation, and time-sensitive control loops. By processing data locally, this layer significantly reduces the round-trip time for critical tasks, which is essential for applications like robotic control and immediate anomaly detection [59].

- Technology Stack: We utilize lightweight containerization and orchestration technologies, such as K3s (a lightweight Kubernetes distribution), to manage and deploy applications on the edge nodes [49]. This approach allows for efficient resource management, fault tolerance, and rapid deployment of microservices.

2.1.3. Cloud Computing Layer

The cloud layer provides robust, scalable, and virtually unlimited resources for computationally intensive tasks and long-term data management.

- Intensive Computing: This layer is responsible for tasks that are not latency-critical but require significant processing power, such as training complex machine learning models, running large-scale, multi-physics simulations for the DTs, and performing historical data analysis to uncover long-term trends.

- Technology Stack: We leverage a combination of open-source cloud platforms like OpenStack for Infrastructure-as-a-Service (IaaS) [45] and serverless computing frameworks like Apache OpenWhisk for Function-as-a-Service (FaaS) [43, 50]. This hybrid approach allows us to orchestrate virtual machines, containers, and functions seamlessly across the continuum.

- DT Core Logic: The cloud layer hosts the master database and the core logic for the DTs, which are built using open frameworks like Eclipse Ditto for device virtualization [23] and FIWARE for context management and data modeling [24, 25].

2.2. Orchestration and Management

A sophisticated, AI-driven orchestration layer is the brain of the platform, responsible for managing resources, applications, and services across the entire distributed infrastructure.

- Inter-Node and Intra-Node Orchestration: The orchestrator manages two levels of tasks. Inter-node orchestration handles the distribution of services across multiple geographically dispersed edge and cloud nodes, optimizing for performance, latency, and cost. Intra-node orchestration manages the resources within each individual node, ensuring that CPU, memory, and storage are used efficiently.

- Declarative Configuration: Application deployment is managed using a declarative approach. Through YAML files and containerized services, users define the desired state of their applications, including placement policies, resource requirements, and scaling rules. This allows the orchestrator to automate deployment and dynamically adjust resource allocation based on real-time demand.

- AI-Driven Automation: A key innovation is the AI module, which enhances orchestration with predictive capabilities.

  - Prediction Analytics Engine: This component utilizes machine learning models, specifically an ARIMA time-series approach [52], to analyze historical data (e.g., CPU utilization, network traffic) and predict future system loads.

  - Decision Engine: Based on the predictions from the analytics engine, the Decision Engine makes intelligent, real-time decisions about resource allocation. It can operate in two modes: a performance-focused mode that scales up resources to handle traffic spikes and maintain QoS, and an energy-efficiency-focused mode that scales down resources during periods of low demand to conserve power. This dual-mode capability is essential for creating a sustainable and cost-effective platform.

2.3. 6G-Intended Testbed and Use Case Validation

To validate the platform's performance in a forward-looking context, we established a comprehensive testbed that emulates the characteristics of a future 6G network.

- Physical Infrastructure: The testbed is physically distributed across two sites and includes a variety of IoT devices, a range of edge computing nodes (NVIDIA Jetson, Raspberry Pi, Dell Edge Gateways), and a local private cloud running OpenStack.

- 6G Network Emulation: While a true 6G network is not yet available, we emulate its key characteristics using specialized tools. The radio access network (RAN) is emulated based on the 3GPP Release 18 specifications, which lay the groundwork for 5G-Advanced and future 6G systems [53]. We use Keysight's LoadCore solution [54] to generate realistic core network traffic and to simulate different network conditions, such as varying levels of latency, jitter, and packet loss. This allows us to test the platform's resilience and performance under the extreme conditions expected in 6G, such as sub-millisecond latency and massive connection densities [55].

- Use Case Scenarios for Validation:

1. Smart Logistics: This use case involves the real-time tracking and management of goods in a simulated warehouse. IoT sensors on packages and autonomous mobile robots (AMRs) continuously send location and status data to the platform. The DT of the warehouse provides a live, 3D visualization of all assets and operations. Edge nodes process video feeds for object recognition and perform real-time path planning for the AMRs to avoid collisions. The cloud is used to optimize overall warehouse logistics based on historical data and demand forecasts.

2. Industrial Automation: This use case focuses on a robotic arm performing a high-precision assembly task. The DT of the robotic arm is used for remote monitoring and predictive maintenance. High-frequency sensor data (vibration, temperature) is processed at the edge to detect anomalies in real-time using ML models. If an issue is detected, the system can automatically adjust the robot's parameters or halt its operation to prevent damage. The cloud is used to analyze long-term performance data, retrain the ML models, and refine the robot's control algorithms.

3. RESULTS

The evaluation of our platform focused on its ability to meet the stringent performance requirements of real-time DT applications, particularly concerning end-to-end latency, reliability, and scalability. The experiments were conducted on the described 6G-intended testbed, using the smart logistics and industrial automation use cases as benchmarks.

3.1. Latency Performance Analysis

End-to-end latency, defined as the time from data generation at the IoT device to the corresponding update in the DT or the execution of a control action, is a critical metric for real-time systems. We measured latency under three different workload placement strategies:

- Cloud-Only: All data is sent directly to the central cloud for processing.
- Static Edge-Assisted: Time-sensitive tasks are pre-configured to run on edge nodes.
- Dynamic AI-Driven Scheduling: Our proposed framework dynamically distributes tasks between the edge and cloud based on real-time conditions.

As shown in Figure 2 (conceptual), the Cloud-only approach resulted in the highest and most variable latency, with a mean of 120 ms and significant jitter, making it unsuitable for real-time control applications that require deterministic performance [56, 57]. The Static Edge-assisted approach showed a significant improvement, reducing the mean latency to 18 ms by processing critical data locally. However, our Dynamic AI-Driven Scheduling framework achieved the best performance, with a mean latency of 12 ms and a 99th percentile latency of 19 ms. This demonstrates the effectiveness of intelligently and adaptively placing computational workloads at the optimal location within the continuum to meet application-specific latency targets [58, 59, 60].

3.2. Digital Twin Synchronization Fidelity

For a DT to be effective, it must remain in close synchronization with its physical counterpart. We measured the synchronization error for the industrial robot arm DT, defined as the time lag between a physical event (e.g., a change in the robot's position) and its reflection in the virtual model. The results, presented in Table 1 (conceptual), show that our platform maintained a synchronization error of less than 25 ms on average. This level of accuracy is sufficient for real-time monitoring and even for enabling sensitive human-in-the-loop applications, such as remote surgery or haptic feedback systems, where delays above 50-100 ms can degrade user experience and performance [61].

Table 1. Digital Twin Synchronization Error for Industrial Automation Use Case.

Metric	Value
Mean Synchronization Error	24.8 ms
Standard Deviation	5.2 ms
95th Percentile Error	33.1 ms
Maximum Error	41.5 ms

We rigorously tested the reliability and resilience of the platform by simulating network failures and node outages. The Kubernetes-based orchestration system was able to automatically detect failures and reschedule affected microservices to healthy nodes with minimal disruption. In the case of an edge node failure, critical tasks were migrated to a neighboring edge node or, if necessary, to the cloud, with an average service recovery time of under 2 seconds, ensuring high availability.

To evaluate scalability, we progressively increased the number of connected IoT devices in the logistics use case from 100 to 5,000. The platform demonstrated near-linear scalability in terms of data ingestion and processing throughput. The CPU and memory utilization on the edge and cloud nodes increased predictably, and the orchestration system effectively balanced the load by automatically scaling service instances as required. This confirms that the architecture is well-suited for large-scale industrial deployments with thousands of connected devices.

3.4. Use Case Specific Performance Outcomes

- Smart Logistics: In the logistics use case, the real-time DT enabled a 15% improvement in warehouse operational efficiency, measured in terms of average order fulfillment time. The AMR collision avoidance system, running entirely on the edge, successfully prevented all potential collisions in our test scenarios, even with high robot density and complex, overlapping paths. This highlights the critical role of low-latency edge processing for safety-critical applications.

- Industrial Automation: For the industrial arm, the edge-based predictive maintenance system was able to detect incipient failures up to 3 hours before they would have caused a system shutdown. This was achieved by analyzing high-frequency vibration data with an ML model deployed at the edge. This demonstrates the immense value of low-latency edge analytics in preventing costly downtime and transitioning from reactive to predictive maintenance strategies. The statistical models for this analysis were developed using Python's statsmodels library [52].

4. Discussion

The results presented in the previous section demonstrate that our flexible, hyper-distributed IoT-Edge-Cloud platform is capable of supporting demanding, real-time DT applications. The performance achieved on our 6G-intended testbed indicates that the proposed architecture is a viable and powerful solution for the next generation of industrial and logistics systems.

4.1. Implications of Performance Results

Our findings on latency are particularly significant. The clear advantage of the edge-assisted and dynamic scheduling approaches over a cloud-only model confirms the industry-wide consensus on the necessity of edge computing for time-sensitive applications [34, 48]. Our dynamic, AI-driven scheduling framework goes a step further by optimizing workload placement in real-time, which is a complex challenge in heterogeneous and dynamic environments [58]. This ability to adapt to changing network conditions and application needs is crucial for maintaining consistent performance. The achieved latency of under 20 ms aligns with the stringent URLLC requirements for 5G and beyond, making the platform suitable for the most demanding applications, such as remote control of vehicles and machinery [57].

The high fidelity of the DT synchronization is another key outcome. A real-time DT is only as valuable as its accuracy and timeliness. The low synchronization error we achieved enables not just passive monitoring but also active, in-the-loop control and immersive human-in-the-loop applications, such as remote operation and augmented reality overlays for maintenance tasks [7, 56]. This represents a significant step towards creating truly interactive and intelligent cyber-physical systems [13].

The modularity and openness of our platform, built using industry-standard technologies like Kubernetes, OpenStack, FIWARE, and MQTT, are also important advantages. This contrasts with many existing proprietary, black-box solutions and provides organizations with greater flexibility, avoids vendor lock-in, and fosters a collaborative ecosystem for innovation [22, 26]. The ability to integrate various open-source components allows for continuous evolution and adaptation of the platform as new technologies emerge.

4.2. Open Challenges and Future Research Directions

Despite the promising results, several significant challenges remain that represent important avenues for future research.

- Energy Efficiency and Sustainability: The proliferation of edge nodes and IoT devices raises concerns about the overall energy consumption of the platform. Future work must focus on developing energy-aware orchestration algorithms that can optimize for performance while minimizing the carbon footprint. This includes exploring the use of renewable energy sources for edge nodes and developing low-power AI/ML techniques for on-device processing.

- Security and Privacy in a Hyper-Distributed Environment: The security of such a distributed system is a major concern. With a vastly expanded attack surface spanning countless IoT devices and edge nodes, a comprehensive, multi-layered security strategy is essential. Future work should focus on developing a zero-trust security architecture that authenticates and authorizes every transaction. Furthermore, techniques like federated learning and confidential computing will be crucial for protecting data privacy while still enabling powerful analytics [47].

- Interoperability and Standardization: While we used open standards, the broader IoT ecosystem remains fragmented. Achieving seamless interoperability between devices, platforms, and applications from different vendors is a major hurdle. Continued effort in developing and adopting common data models, APIs, and communication protocols is essential for the scalability and long-term success of such platforms.

- Autonomous Management and Orchestration: The complexity of managing a large-scale, hyper-distributed system can be overwhelming. While our AI-driven orchestrator is a step in the right direction, future research should focus on developing fully autonomous systems that can manage, heal, and optimize themselves with minimal human intervention. This will require advancements in reinforcement learning and other AI techniques for complex system control [46].

- Social and Ethical Considerations: As these platforms become more integrated into critical infrastructure, it is vital to address the social and ethical implications. This includes issues of data ownership, algorithmic bias, transparency, and the impact of automation on the workforce. Developing frameworks for explainable AI (XAI) will be essential for building trust and ensuring accountability in systems that make autonomous decisions.

4.3. Comparison with State-of-the-Art

Comparing our work to the existing literature, we build upon the foundational concepts of the IoT-Edge-Cloud continuum [8, 9] and Digital Twins [10, 12]. Our primary contribution is the practical implementation, integration, and validation of a holistic platform on a testbed that looks forward to the demands of 6G. While other studies have proposed similar architectures [11, 30], our work provides a comprehensive performance evaluation using realistic industrial use cases and demonstrates the tangible benefits of a flexible, dynamically scheduled, and AI-driven approach.

5. CONCLUSIONS

The evolution towards 6G is set to unlock a new era of industrial innovation, with real-time Digital Twins at its core. In this paper, we have presented a flexible, hyper-distributed IoT-Edge-Cloud platform designed to meet the challenges of this new era. Our architecture leverages a multi-tiered approach to computing, distributing workloads intelligently across the continuum to optimize for latency, throughput, and resource utilization.

Through extensive testing on a 6G-intended testbed, we have demonstrated that our platform can achieve the ultra-low latency and high reliability required for real-time DT applications in demanding logistics and industrial environments. The results validate the critical role of edge computing and highlight the substantial benefits of a dynamic, adaptive, and AI-driven orchestration strategy. The use of open-source technologies ensures that the platform is flexible, scalable, and future-proof, providing a solid foundation for future research and development.

Future work will focus on addressing the critical challenges outlined in our discussion. We plan to enhance the platform's security framework by implementing a comprehensive zero-trust architecture. We will also explore more advanced AI-driven techniques for autonomous orchestration and resource management, with a strong focus on energy efficiency. Furthermore, we intend to expand our testbed to include a wider variety of use cases and to integrate emerging 6G technologies as they become available. By continuing to push the boundaries of what is possible with distributed computing and next-generation networks, we aim to accelerate the adoption of transformative technologies like Digital Twins and help realize the full potential of the 6G vision.

REFERENCES

[1] Uusitalo, M.A.; et al. 6G Vision, Value, Use Cases and Technologies From European 6G Flagship Project Hexa-X. IEEE Access 2021, 9, 160004–160020.

[2] Viswanathan, H.; Mogensen, P.E. Communications in the 6G Era. IEEE Access 2020, 8, 57063–57074.

[3] Wang, C.X.; et al. On the Road to 6G: Visions, Requirements, Key Technologies, and Testbeds. IEEE Commun. Surv. Tutorials 2023, 25, 905–974.

[4] Allioui, H.; Mourdi, Y. Exploring the full potentials of IoT for better financial growth and stability: A comprehensive survey. Sensors 2023, 23, 8015.

[5] Kwon, J.H.; Kim, H.J.; Lee, S. Optimizing Traffic Scheduling in Autonomous Vehicle Networks Using Machine Learning Techniques and Time-Sensitive Networking. Electronics 2024, 13, 2837.

[6] Nie, Z.; Chen, K.C.; Alanezi, Y. Socially Networked Multi-Robot System of Time-Sensitive Multi-Link Access in a Smart Factory. In Proceedings of the ICC 2023-IEEE International Conference on Communications, Rome, Italy, 28 May–1 June 2023; pp. 4918–4923.

[7] Hazarika, A.; Rahmati, M. Towards an evolved immersive experience: Exploring 5G-and beyond-enabled ultra-low-latency communications for augmented and virtual reality. Sensors 2023, 23, 3682.

[8] Gkonis, P.; Giannopoulos, A.; Trakadas, P.; Masip-Bruin, X.; D'Andria, F. A survey on IoT-edge-cloud continuum systems: Status, challenges, use cases, and open issues. Future Internet 2023, 15, 383.

[9] Jamil, M.N.; Schelén, O.; Monrat, A.A.; Andersson, K. Enabling Industrial Internet of Things by Leveraging Distributed Edge-to-Cloud Computing: Challenges and Opportunities. IEEE Access 2024, 12, 127294–127308.

[10] Hlophe, M.C.; Maharaj, B.T. From cyber–physical convergence to digital twins: A review on edge computing use case designs. Appl. Sci. 2023, 13, 13262.

[11] Lubrano, F.; Caragnano, G.; Scionti, A.; Terzo, O. Challenges, Novel Approaches and Next Generation Computing Architecture for Hyper-Distributed Platforms Towards Real Computing Continuum. In Proceedings of the Advanced Information Networking and Applications, Kitakyushu, Japan, 17–19 April 2024; pp. 449–459.

[12] Alkhateeb, A.; Jiang, S.; Charan, G. Real-time digital twins: Vision and research directions for 6G and beyond. IEEE Commun. Mag. 2023, 61, 128–134.

[13] Nguyen, T.N.; Zeadally, S.; Vuduthala, A.B. Cyber-

physical cloud manufacturing systems with digital twins. IEEE Internet Comput. 2021, 26, 15–21.

[14] Barnabas, J.; Raj, P. The human body: A digital twin of the cyber physical systems. In Advances in Computers; Elsevier: Amsterdam, The Netherlands, 2020; Volume 117, pp. 219–246.

[15] Xu, H.; et al. Smart mobility in the cloud: Enabling real-time situational awareness and cyber-physical control through a digital twin for traffic. IEEE Trans. Intell. Transp. Syst. 2023, 24, 3145–3156.

[16] Kaigom, E.G.; Roßmann, J. Value-driven robotic digital twins in cyber–Physical applications. IEEE Trans. Ind. Inform. 2020, 17, 3609–3619.

[17] Shi, S. Industrial cloud, automation: The industrial. Internet of Things (IIoT) is being embraced by manufacturers as a natural extension of automation and controls development. Control. Eng. 2023, 70, 31–32.

[18] Zorchenko, N.; Tyupina, T.; Parshutin, M. Technologies Used by General Electric to Create Digital Twins for Energy Industry. Power Technol. Eng. 2024, 58, 521–526.

[19] Gupta, R.; Reebadiya, D.; Tanwar, S. 6G-enabled edge intelligence for ultra-reliable low latency applications: Vision and mission. Comput. Stand. Interfaces 2021, 77, 103521.

[20] Santos, J.; Wauters, T.; Volckaert, B.; De Turck, F. Towards low-latency service delivery in a continuum of virtual resources: State-of-the-art and research directions. IEEE Commun. Surv. Tutor. 2021, 23, 2557–2589.

[21] Vaish, R.; Hollinger, M.C. Case Study: IBM–Automating Visual Inspection. In Springer Handbook of Automation; Springer: Berlin/Heidelberg, Germany, 2023; pp. 1439–1450.

[22] Fortino, G.; Guerrieri, A.; Pace, P.; Savaglio, C.; Spezzano, G. Iot platforms and security: An analysis of the leading industrial/commercial solutions. Sensors 2022, 22, 2196.

[23] Kherbache, M.; Maimour, M.; Rondeau, E. Digital twin network for the IIoT using eclipse ditto and hono. IFAC-PapersOnLine 2022, 55, 37–42.

[24] De Benedictis, A.; Rocco di Torrepadula, F.; Somma, A. A Digital Twin Architecture for Intelligent Public Transportation Systems: A FIWARE-Based Solution. In Proceedings of the International Symposium on Web and Wireless Geographical Information Systems, Yverdon-Les-Bains, Switzerland, 17–18 June 2024; pp. 165–182.

[25] Conde, J.; Munoz-Arcentales, A.; Alonso, Á.; Huecas, G.; Salvachúa, J. Collaboration of digital twins through linked open data: Architecture with fiware as enabling technology. IT Prof. 2022, 24, 41–46.

[26] Robles, J.; Martín, C.; Díaz, M. OpenTwins: An open-source framework for the development of next-gen compositional digital twins. Comput. Ind. 2023, 152, 104007.

[27] Zeb, S.; et al. Industrial digital twins at the nexus of NextG wireless networks and computational intelligence: A survey. J. Netw. Comput. Appl. 2022, 200, 103309.

[28] Mihai, S.; et al. Digital twins: A survey on enabling technologies, challenges, trends and future prospects. IEEE Commun. Surv. Tutor. 2022, 24, 2255–2291.

[29] Mirani, A.A.; Velasco-Hernandez, G.; Awasthi, A.; Walsh, J. Key challenges and emerging technologies in industrial IoT architectures: A review. Sensors 2022, 22, 5836.

[30] Dong, J.; et al. Mixed cloud control testbed: Validating vehicle-road-cloud integration via mixed digital twin. IEEE Trans. Intell. Veh. 2023, 8, 2723–2736.

[31] Alimi, I.A.; et al. Towards enhanced mobile broadband communications: A tutorial on enabling technologies, design considerations, and prospects of 5G and beyond fixed wireless access networks. Appl. Sci. 2021, 11, 10427.

[32] Milovanovic, D.A.; Bojkovic, Z.S. 5G Ultrareliable and Low-Latency Communication in Vertical Domain Expansion. In Driving 5G Mobile Communications with Artificial Intelligence Towards 6G; CRC Press: Boca Raton, FL, USA, 2023; pp. 137–181.

[33] Kong, L.; et al. Edge-computing-driven Internet of Things: A Survey. ACM Comput. Surv. 2022, 55, 1–41.

[34] Shi, W.; Cao, J.; Zhang, Q.; Li, Y.; Xu, L. Edge Computing: Vision and Challenges. IEEE Internet Things J. 2016, 3, 637–646.

[35] Lin, X.; et al. 6G digital twin networks: From theory to practice. IEEE Commun. Mag. 2023, 61, 72–78.

[36] AlSobeh, A.M.; Hammad, R.; Al-Tamimi, A.K. A modular cloud-based ontology framework for context-aware EHR services. Int. J. Comput. Appl. Technol. 2019, 60, 339–350.

[37] Arbab-Zavar, B.; Palacios-Garcia, E.J.; Vasquez, J.C.; Guerrero, J.M. Message queuing telemetry transport communication infrastructure for grid-connected AC microgrids management. Energies 2021, 14, 5610.

[38] Baig, M.J.A.; Iqbal, M.T.; Jamil, M.; Khan, J. A low-cost, open-source peer-to-peer energy trading system for a remote community using the internet-of-things, blockchain, and hypertext transfer protocol. Energies 2022, 15, 4862.

[39] Yassein, M.B.; et al. Challenges and techniques of constrained application protocol (CoAP) for efficient energy consumption. In Proceedings of the 2020 11th International Conference on Information and Communication Systems (ICICS), Irbid, Jordan, 7–9 April 2020; pp. 373–377.

[40] Marino, C.A.; Chinelato, F.; Marufuzzaman, M. AWS IoT

analytics platform for microgrid operation management. Comput. Ind. Eng. 2022, 170, 108331.

[41] Satapathi, A.; Mishra, A. Build an IoT Solution with Azure IoT Hub, Azure Functions, and Azure Cosmos DB. In Developing Cloud-Native Solutions with Microsoft Azure and NET; Springer: Berlin/Heidelberg, Germany, 2022; pp. 193–218.

[42] Fortino, G.; et al. A review of internet of things platforms through the iot-a reference architecture. In Proceedings of the International Symposium on Intelligent and Distributed Computing, Freiburg, Germany, 4–8 October 2021; pp. 25–34.

[43] Alabbas, A.; et al. Performance analysis of apache openwhisk across the edge-cloud continuum. In Proceedings of the 2023 IEEE 16th International Conference on Cloud Computing (CLOUD), Chicago, IL, USA, 2–8 July 2023; pp. 401–407.

[44] Daki´c, V.; Kovaˇc, M.; Slovinac, J. Evolving High-Performance Computing Data Centers with Kubernetes, Performance Analysis, and Dynamic Workload Placement Based on Machine Learning Scheduling. Electronics 2024, 13, 2651.

[45] Tricomi, G.; et al. Paving the way for an Urban Intelligence OpenStack-based Architecture. In Proceedings of the 2024 IEEE International Conference on Smart Computing (SMARTCOMP), Osaka, Japan, 29 June–2 July 2024; pp. 284–289.

[46] Ullah, A.; et al. Orchestration in the Cloud-to-Things compute continuum: Taxonomy, survey and future directions. J. Cloud Comput. 2023, 12, 135.

[47] Alsobeh, A.; Shatnawi, A. Integrating data-driven security, model checking, and self-adaptation for IoT systems using BIP components: A conceptual proposal model. In Proceedings of the International Conference on Advances in Computing Research, Orlando, FL, USA, 8–10 May 2023; pp. 533–549.

[48] Khan, W.Z.; et al. Edge computing: A survey. Future Gener. Comput. Syst. 2019, 97, 219–235.

[49] Fogli, M.; et al. Performance evaluation of kubernetes distributions (k8s, k3s, kubeedge) in an adaptive and federated cloud infrastructure for disadvantaged tactical networks. In Proceedings of the 2021 International Conference on Military Communication and Information Systems (ICMCIS), The Hague, The Netherlands, 4–5 May 2021; pp. 1–7.

[50] Banaei, A.; Sharifi, M. Etas: Predictive scheduling of functions on worker nodes of apache openwhisk platform. J. Supercomput. 2022, 78, 5358–5393.

[51] Santos, Á.; Correia, N.; Bernardino, J. On the Suitability of Cloud Models for MEC Deployment Purposes. In Proceedings of the 2023 6th Experiment@ International Conference; (exp. at'23); Evora, Portugal, 5–7 June 2023; pp. 255–260.

[52] Seabold, S.; Perktold, J. Statsmodels: Statistical Models in Python; Python Software Foundation: Wilmington, DE, USA, 2010.

[53] 3rd Generation Partnership Project (3GPP). NR; User Equipment (UE) radio access capabilities (Release 18). Technical Report TS 38.306 V18.3.0, 3GPP, 2024.

[54] Keysight Technologies. P8900S LoadCore—Core Network Solutions. 2024.

[55] 5G Infrastructure Public Private Partnership (5G PPP). Beyond 5G/6G KPIs and Target Values. 2022.

[56] Khan, T.; et al. Understanding effects of visual feedback delay in ar on fine motor surgical tasks. IEEE Trans. Vis. Comput. Graph. 2023, 29, 4697–4707.

[57] Zhao, L.; et al. Remote Driving of Road Vehicles: A Survey of Driving Feedback, Latency, Support Control, and Real Applications. IEEE Trans. Intell. Veh. 2024, 1–22.

[58] Dreibholz, T.; Mazumdar, S. Towards a lightweight task scheduling framework for cloud and edge platform. Internet Things 2023, 21, 100651.

[59] Makondo, N.; et al. An efficient architecture for latency optimisation in 5G using Edge Computing for uRLLC use cases. In Proceedings of the 2024 International Conference on Artificial Intelligence, Big Data, Computing and Data Communication Systems (icABCD), Port Louis, Mauritius, 1–2 August 2024; pp. 1–7.

[60] Velayutham, A. Optimizing sase for low latency and high bandwidth applications: Techniques for enhancing latency-sensitive systems. Int. J. Intell. Autom. Comput. 2023, 6, 63–83.

[61] Lin, Y.H.; et al. Hapticseer: A multi-channel, black-box, platform-agnostic approach to detecting video game events for real-time haptic feedback. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, Online, 8–13 May 2021; pp. 1–14.