Enabling Distributed and Hybrid Digital Twins in the Industry5.0 Cloud Continuum

Paolo Bellavista(✉)

Alma Mater Studiorum University of Bologna, 40122 Bologna, Italy
paolo.bellavista@unibo.it

Abstract. The efficient exploitation of the cloud continuum (interworking of industrial cloud, edge cloud, 5G/B5G base stations, and fog node technologies) is a key factor for future Industry4.0 and beyond applications. For instance, with their ability to work more locally to data sources and to controllable actuators, cloud continuum-based solutions are considered very promising for dynamic manufacturing line control/reconfiguration and for sustainability optimizations (reduction of power and materials consumption). In these contexts, privacy, data sovereignty/control, latency, reliability, and scalability are crucial. This short paper, summarizing some key concepts from the associated keynote speech, aims at offering an overview of the concept of Digital Twins (DTs) for Industry5.0 [1] and at showing some primary and state-of-the-art solution guidelines that are emerging in the field. In particular, in our H2020 IoTwins project [2], we promote, design, implement, and evaluate Industry5.0 DTs that are both distributed because they are able to run over differentiated cloud continuum virtualized resources, also by changing their location dynamically, and hybrid because they combine data-driven machine learning models and simulations based on mathematical-physical modeling of the cyber-physical systems they represent. Distributed and hybrid DTs based on cloud continuum technologies are showing their effectiveness and efficiency in different application cases and scenarios, as demonstrated by our experience of IoTwins testbeds development.

Keywords: Digital Twins · Industry 5.0 · Cloud continuum · Edge cloud computing · IoTwins EU project · Testbed development and evaluation

1 Architectural Guidelines for Distributed and Hybrid DTs

Fulfilling the Industrial Internet of Things (IIoT) promise of autonomous interoperability, solution agility, and flexible reconfiguration of production chains, it is expected not only that the number of interconnected physical devices in the manufacturing domain will increase drastically, but also that there will be a strong need to interact with global/local cloud/edge services to act intelligently and flexibly. This introduces numerous challenges to industrial networked computing environments, which are traditionally quite static and isolated. First, IIoT environments will have to serve a wide range of industrial applications with different Quality-of-Service (QoS) requirements, ranging from traditional, time-sensitive, and reliable closed-loop control systems to event-driven, delay-bounded,

A. Cerone et al. (Eds.): SEFM 2021 Workshops, LNCS 13230, pp. 139–142, 2022.
https://doi.org/10.1007/978-3-031-12429-7_10

or best-effort sensor traffic. Second, these applications will have to fully grasp the potential benefits of the Cloud-to-Things Continuum (C2TC) [3], by exploiting resource virtualization and dynamic management of support/service components over industrial edge gateways, Multi-access Edge Computing (MEC) telco nodes (as standardized in ETSI specifications), and the global cloud. Third, in the reconfigurable factories of the future, the end-to-end QoS of co-existing applications will require to be managed, even when new C2TC industrial applications are dynamically introduced/removed. Fourth, several connectivity options (5G/6G, TSCH, WiFi7, …) will gradually be introduced into production lines, by opening new opportunities of QoS specification and control, but also new challenges for heterogeneous QoS-constrained management.

In these very challenging scenarios and application cases, Digital Twins (DTs) are gaining relevance as comprehensive, actionable, digital representations of industrial physical systems and their behavior [4, 5]. DTs, in fact, provide a software copy of a physical asset by reflecting its properties, behaviors, and relationships according to the operational context. The physical and software counterparts mutually cooperate and co-evolve for enabling features such as device control, simulation, analytics, and more generally, the ability to dynamically enhance the functionality associated with physical objects and the optimization of its operation, management, and maintenance. For instance, DTs are a crucial enabling technology to detect and diagnose anomalies, to determine an optimal set of actions that maximize key performance metrics, and to effectively and efficiently enforce on-line quality management of production processes under latency and reliability constraints.

We strongly believe that future, effective, and efficient **DTs for Industry5.0 have to be distributed, hybrid, and based on cloud continuum technologies**. On the perspective of their distribution, for enhancing flexibility and data/control locality, DTs should be organized in a hierarchy of three layers: i) directly at IoT devices, whenever possible, **IoT twins** are lightweight models of specific components, performing big-data processing and local control for quality management operations; ii) **edge twins** have to be deployed at plant gateways and/or at emerging MEC nodes, thus providing higher level control knobs and orchestrating Internet of Things (IoT) sensors and actuators in a production locality; and iii) **cloud twins** are in charge of most time-consuming and resource-greedy operations, such as (typically off-line) parallel simulation and deep-learning, feeding edge twins with pre-elaborated or reduced order models to be efficiently executed at industry premises for monitoring/control/tuning purposes.

These distributed and interworking DTs should be **cloud continuum-oriented**, i.e., run at and use local virtualized resources (based on full virtual machines or on different types of possible heterogeneous containers). This opens up excellent opportunities of dynamic deployment and offloading/in-loading of needed behavior at runtime [6], but also poses very tough and challenging technical issues in terms of performance isolation, quality control, and middleware/application interference [7]. Finally, on the perspective of simulation models and technologies, DTs have to make use of **hybrid** modeling approaches, i.e., combining pure data-driven machine learning techniques with more traditional simulation models and tools, typically based on the physical modelling of the behavior of the twinned cyber-physical system.

2 The IoTwins H2020 Project

By following the above design principles of distributed and hybrid DTs based on cloud continuum technologies, the H2020 IoTwins project [2] aims to lower the barriers for building C2TC systems and services based on big data for the domains of manufacturing and facility management, by harmonizing standards to enable interoperability, and by developing an easy-to-use service layer that facilitates and decreases the cost of integration and deployment. To this purpose, IoTwins proposes a framework for a seamless, straightforward, and loose integration of already developed and deployed industrial software components running in typical long-lived industrial test-beds. The high-level distributed architecture of the IoTwins platform is depicted in Fig. 1, where you can notice, for example, that traditional machine learning training is first performed at cloud resources (integrated with agent-based and physical simulations), then refined at edge nodes (also with Federated Learning approaches [8]), with the possibility to run already trained anomaly detection algorithms also at IoT nodes.

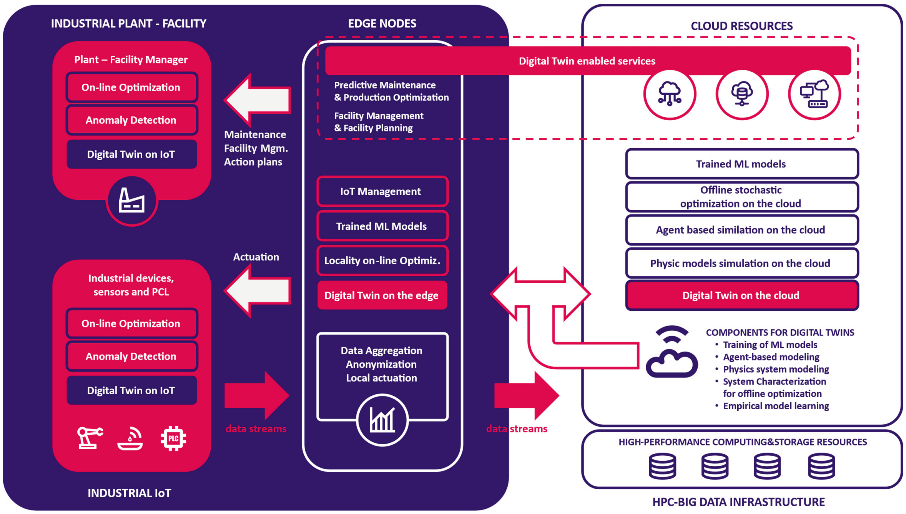

Fig. 1. The concept of distributed and hybrid DTs for the cloud continuum in IoTwins.

The IoTwins architecture above has been implemented, prototyped, and evaluated into different possible incarnations, suitable for different application cases and deployment environments. For example, one incarnation is a novel platform based on the extension of Indigo PaaS and exploiting edge cloud virtualized resources via Mesos; the platform originally implements distributed orchestration based on enhanced TOSCA templates for the C2TC. Another incarnation of the IoTwins architecture is a platform deriving from the integration of more proprietary Siemens solutions for the distributed monitoring and control of smart grids, which are deployed at wide-scale in the Wien metropolitan environment together with network and control equipment by TTTech. Let us also note that, for maximum interoperability, openness, and future extensibility, the

IoTwins architecture is compliant with the most widespread standard Industry4.0 architectural specifications, such as the Reference Architectural Model Industrie 4.0 (RAMI) and the Industrial Internet Reference Architecture (IIRA) [9].

To practically and tangibly show the advantages of the adoption of distributed and hybrid DTs designed according to the IoTwins guidelines, we have already developed a series of testbeds and associated Industry5.0 applications; these tangible testbeds are crucial to promote the adoption of our C2TC techniques in large companies and SMEs, by showing the feasibility and efficiency of the proposed approach in practical cases. To this purpose, twelve testbeds have been implemented, some of them in the manufacturing domain (prescriptive maintenance for wind turbines, accurate modelling of machine tool spindle behavior, optimization of the crankshaft manufacturing production process, and defect reduction in closure manufacturing lines), in the facility management domain (the Barcelona Camp-Nou stadium during crowd upload/download, energy optimization of the Cineca HPC site, and predictive management of electric smart grid in Wien), and in other cases similar to the previous ones but at a different scale. The performance results and economic advantages of the IoTwins adoption have been already measured in those testbeds and are showing the effectiveness, efficiency, and flexibility of dynamically distributing our Industry5.0 platform on every node belonging to the C2TC chain. For finer technical insights about the technologies, middleware, algorithms, and protocols integrated in the IoTwins architecture, please see [2], where a significant part of the IoTwins project results are made publicly available for the community of researchers and practitioners in the field.

References

1. Zong, L., et al.: End-to-end transmission control for cross-regional industrial internet of things in industry 5.0. IEEE Trans. Industr. Inf. **18**(6), 4215–4223 (2022)
2. The H2020 IoTwins Project. https://www.iotwins.eu/
3. Samie, F., Bauer, L., Henkel, J.: From cloud down to things: an overview of machine learning in internet of things. IEEE Internet Things J. **6**(3), 4921–4934 (2019)
4. Minerva, R., Lee, G.M., Crespi, N.: Digital twin in the IoT context: a survey on technical features, scenarios, and architectural models. Proc. IEEE **108**(10), 1785–1824 (2020)
5. Vukovic, M., Mazzei, D., Chessa, S., Fantoni, G.: Digital twins in industrial IoT: a survey of the state of the art and of relevant standards. In: Proceedings of IEEE International Conference on Communications (ICC). IEEE, New York (2021)
6. Zanni, A., et al.: Automated selection of offloadable tasks for mobile computation offloading in edge computing In: CNSM Conference Proceedings. IEEE, New York (2017)
7. Li, Y., Zhang, J., Jiang, C., Wan, J., Ren, Z.: PINE: optimizing performance isolation in container environments. IEEE Access **7**, 30410–30422 (2019)
8. Bellavista, P., Foschini, L., Mora, A.: Decentralised learning in federated deployment environments: a system-level survey. ACM Comput. Surv. **54**(1), 15:1–15:38 (2021)
9. Borghesi, A., et al.: IoTwins: design and implementation of a platform for the management of digital twins in industrial scenarios In: CCGRID Conference Proceedings. IEEE, New York (2021)