Towards an Urban Digital Twins Continuum Architecture

Sergio Laso[1](✉), Lorenzo Toro-Gálvez[2], Javier Berrocal[1], Javier Troya[2], Carlos Canal[2], and Juan Manuel Murillo[1]

[1] University of Extremadura, Badajoz, Spain
{slasom,jberolm,juanmamu}@unex.es
[2] ITIS Software, Universidad de Málaga, Málaga, Spain
{lorenzotoro,jtroya,carloscanal}@uma.es

Abstract. In the era of smart cities, where the integration of Internet of Things devices and the need to efficiently manage urban environments have generated considerable interest, the Digital Twin concept emerges as a key solution. This technology allows us to study and simulate the behavior of complex urban dynamics. However, conventional Digital Twin architectures face significant challenges, such as limited scalability, inherent latency, and data privacy concerns stemming mainly from their centralized nature. In response to these challenges, this paper proposes an innovative distributed architecture for the so-called Urban Digital Twins, implemented on top of the Computing continuum. The main objective is to establish a more efficient and scalable framework, specifically designed for the demands of smart cities. To support the feasibility of this proposal, two case studies are presented: one focused on urban public transportation systems, and the other focused on a pollution monitoring system. These case studies illustrate how a distributed architecture can effectively address existing challenges, providing a solid foundation for the smart and sustainable management of urban environments.

Keywords: Urban Digital Twins · Architecture · Smart City · Computing continuum

1 Introduction

In recent years, we have witnessed a significant increase in the use of Internet-connected devices, which has led to the emergence of the IoT paradigm [20]. This approach has been applied in various fields, such as smart homes, healthcare, industry 4.0, and smart cities. However, making adjustments to the operation in these largely distributed systems can lead to unexpected problems. For example, in a smart city scenario, managing traffic based on air quality values to alleviate congestion in one area may inadvertently lead to traffic jams in other places.

To try to solve this kind of problems, Digital Twins [22] have recently emerged as a powerful technology to replicate the behavior of real systems. A Digital Twin

B. Tekinerdoğan et al. (Eds.): ECSA 2023, LNCS 14590, pp. 272–286, 2024.
https://doi.org/10.1007/978-3-031-66326-0_17

(DT) is a digital representation of a system, service, or product. To create such a digital representation, it must receive the status of the physical system, which allows it to replicate the real environment. In a DT system, there is a bidirectional communication between the two twins that allows keeping the digital twin updated and proposing changes on the physical twin [13]. DTs become fundamental tools for monitoring, predicting, and integrating data from IoT devices.

However, conventional DT implementations have several limitations on scalability [18], latency [12], and privacy [6]. These limitations become evident when we want to represent highly-distributed systems such as smart cities, with a large variety of data sources from IoT devices, people, etc. Conventional DTs are highly-coupled systems where all collected data is centralized in one single place. Creating a so-called Urban Digital Twin (UDT) this way is practically unfeasible due to its size and complexity.

To try to overcome these challenges, the Computing continuum paradigm [2] has been gaining importance in recent years. The Computing continuum represents an evolution of Cloud computing that extends from cloud environments to IoT devices, located closer to people, for the purpose of storing and processing information. The Fog and Edge computing paradigms have brought cloud computing environments and data processing closer to information sources. By processing information closer to its source, it is possible to reduce infrastructure load, improve the quality of service, and preserve privacy.

In this paper, we propose a distributed architecture on the low-coupling Computing Continuum [3,7] composed of different DTs that simulate the behavior of different entities in a highly-distributed system, so that altogether they compose a UDT. Due to their low coupling, the DTs work independently but can interact with each other to enrich their models allowing greater scalability and flexibility. Their distributed nature and proximity to users allow for improved data privacy as well as faster responsiveness, as it minimizes data transmission. To demonstrate the feasibility of our proposal, two case studies focused on smart cities are considered in this work, one focused on an urban public transportation system, and the other focused on a pollution monitoring system.

The paper is organized as follows. Section 2 provides an introductory context on the current state and need for an Urban Digital Twin (UDT) in the computing continuum. In Sect. 3, some related work is examined in detail. Section 4 provides a detailed explanation of the proposed Digital Twins architecture for the Computing Continuum. Section 5 presents in detail the two case studies that addressed transportation and pollution in a smart city. Section 6 discusses how our proposal addresses the current challenges and also its limitations. Finally, in Sect. 7, the conclusions derived from the study are presented and some directions for future research are outlined.

2 Motivation

In the context of smart cities, it encompasses a diversity of systems and applications aimed at optimizing their governance. These include fundamental areas such as transportation [14], whose main objective is to offer an efficient service of the arrival times of public transport, as well as in the identification of traffic patterns to improve urban mobility. Also, in the field of pollution control [5], new regulations are emerging due to the problems generated by air quality, intending to monitor them to avoid penalties and improve air quality.

The implementation of Digital Twins in these strategic areas would not only improve operational efficiency but also contribute to the sustainability and resilience of cities. By providing a real-time digital representation of the infrastructure, Digital Twins facilitate a detailed understanding and analysis of urban systems, enabling data-driven decision-making to address critical challenges such as traffic congestion, environmental pollution, and public transportation planning.

Digital Twins represent essential tools [1] in monitoring, predicting, and integrating data from IoT devices. They continuously exchange data, including dynamic physical twin data and environmental data, and store it in a data storage system (the so-called *data lake*). They use ontologies for data comprehension, high-dimensional data analysis, and data fusion algorithms to integrate multiple data sources. With artificial intelligence (AI) algorithms, DTs can perform feature selection, pattern recognition, and optimization. In addition, they enable closed-loop optimization, allowing the physical entity to respond to changes based on the DT's analysis and optimization. DTs have self-adaptation and self-parametrization capabilities, allowing them to resemble the physical twin throughout their lifecycle. They employ predictive analytics to forecast future statuses and use prescriptive analytics to make data-driven decisions.

The Digital Twins in the context of smart cities are referred to as Urban Digital Twins (UDTs) [4]. An Urban Digital Twin (UDT) is a Digital Twin that acts as a virtual representation of a city's physical resources, making use of data, analytics, and AI to generate real-time, adaptive simulation models. This digital twin captures both the present state and historical context of various aspects of a smart city. Moreover, additional applications provide practical intelligence, contributing to the construction of a collective picture of urban reality. UDTs facilitate more informed decision-making, foster collaborative governance, and improve urban planning by providing a safe environment for environmental sustainability [26].

To achieve the development of a UDT, it is essential to collect data from a variety of sources that are produced by a smart city, ranging from sensors to contextual information and the experience of users interacting with different services such as the transportation system. In previous work [9,15], we proposed innovative models such as *People-as-a-Service (PeaaS)* and *Human Microservices* designed specifically for IoT environments. The purpose of these models lies in the collection and detailed analysis of user habits and routines. By implementing these models, we can effectively integrate citizen information into UDT models. This approach would give us the opportunity to enrich the UDT to make it more accurate and adapted to the context of the citizens who make use of it.

However, a traditional approach would involve implementing a centralized and monolithic UDT, usually deployed in the cloud. Inevitably, this comes with significant drawbacks. First, there is the complexity and lack of scalability of the system, as all information and processes are concentrated in a single location. Storing data generated by the entire city would require extremely large data transmission and storage capacities, generating considerable costs. As an example, consider the Los Angeles Department of Transportation [23], which processed more than 7 Terabytes of real-time data daily in 2022 from different sources like traffic lights signals, traffic sensors, bus signals, etc. In addition, there are concerns about user privacy when storing sensitive data about individuals, which compromises the integration of data from various sources. By concentrating all information at a single point, considerable computing power would be required to integrate all sources of information. Since the processing capabilities of a centralized model may be limited, this would adversely affect system performance and the quality of results, or greatly increase operational costs.

To overcome these drawbacks, a solution is to use distributed architectures such as Computing continuum [17] to deploy a UDT. With this type of architecture, different entities (citizens, buses, air monitoring stations, etc.) deploy their own DT providing low coupling, allowing them to work independently but interact with each other to enrich their models allowing for greater scalability and flexibility. By adopting this distributed architecture, concerns associated with centralized storage and processing on a single node, as well as high coupling between entities, are overcome. In addition, this strategy addresses privacy concerns, as personal information is handled and processed in a decentralized manner on individual devices, ensuring anonymization and preservation of data privacy.

To carry out this initiative, in this paper, we present a detailed proposal of an architecture designed in the framework of the Computing continuum to implement a UDT. To demonstrate the feasibility of this proposal, we provide a description of its application in two specific contexts: a transportation system and a pollution control system. We explain in detail how this architecture would be implemented in each of these scenarios, highlighting its adaptability to complex and dynamic urban environments, materializing the usefulness and practical applicability of a distributed UDT in concrete cases of great relevance for the improvement of the quality of life in urban environments.

3 Related Work

In this section, we present and analyze some related works on Urban Digital Twins, and we detail the requirements of this kind of system.

Lehtola et al. [11] study the impact of digital twins in smart cities. They argue that UDTs must address the specific needs of a city, offering high-fidelity content. In addition, continuous updating of the UDT—using devices like IoT sensors—is essential to reflect the constant urban changes. The authors emphasize the necessity of taking humans into consideration to ensure successful implementation in order to improve decision-making, which is a concern that we also share. They also highlight the incorporation of AI techniques for automatically updating models through the utilization of sensor data.

Schrotter and Hurzeler [21] present a UDT for the city of Zurich, which is defined as a digital and spatial model of the city that integrates 3D spatial data and models for different themes. We have chosen this work because it shows diverse applications of digital twins in smart cities: analyze city growth, visualize construction projects, assess the impact on urban climate, and enable active public participation in planning, among others. The authors also highlight the availability of open data, which is an essential component for the success and usefulness of the UDT.

Ruiz et al. propose BODIT [19], a UDT of the public transportation system in the city of Badalona (Spain). They use a traffic simulator and a genetic algorithm to reproduce the city's traffic and adapt to different situations. Bus schedules are used to predict and detect a lack of punctuality at bus stops, enabling informed decision-making as a response to unusual situations such as accidents.

Although these works present interesting UDT initiatives, there are several significant concerns that they do not properly address. In particular, all of them describe a monolithic and centralized UDT architecture. A smart city is a complex system composed of a large set of subsystems, devices, people, etc. Therefore, centralized UDTs may be unmanageable and unfeasible due to its sheer scale and complexity.

One example of a distributed architecture is given in Villalonga et al. [25], in which the authors present a distributed DT framework that improves local decision-making in the manufacturing industry. The integration of local and global digital twins enables more accurate fault detection, notifying the system for reconfiguration and scheduling actions. This distributed approach offers the advantage of increasing efficiency in decision-making by using improved predictive models and performing simulations at different levels.

After reviewing and analyzing these and other related works, we have identified several limitations that we aim to address with our approach. Some of the proposals [11,19,21] primarily focus on monolithic and centralized architectures for UDTs, ignoring the potential advantages of a distributed architecture as discussed in [25]. We can go further and consider this distribution of UDTs over the Continuum to solve the problems that centralized architectures present, such as scalability, response time, and structural complexity, among others. Taking these concerns into consideration, we advocate that DT proposals for smart cities and other complex systems follow a hierarchical and distributed architecture over the Continuum, addressing the following requirements:

- **Scalability and flexibility.** The DT architecture must rest on weakly coupled systems, which operate independently and interact and coordinate with each other. A distributed architecture allows for the storage and processing of information in the entities where it is generated or consumed. This provides greater scalability and flexibility, avoiding the complexity associated with centralized systems. Each entity manages its own information and resources, contributing to a more adaptable system.

- **Data privacy.** Personal data privacy must be prioritized by keeping sensitive information stored and processed locally on citizens' devices. This decentralized approach minimizes the need for data transmission, reducing the risks associated with centralizing sensitive data.
- **Reduction of duplicity.** By distributing computation, the system optimizes resource allocation and avoids duplicating computational tasks. Component reuse and modularization minimize duplicity by designing components that prevent redundant functionalities in different layers.
- **Reactivity and responsiveness.** By enabling local-level responsiveness, where data are computed in proximity, the system will achieve faster response times, enabling the system to react promptly to changes and events, ensuring quick processing, and providing almost instantaneous responses.
- **Adaptability to complex systems.** Smart cities are complex systems that integrate technology, interconnected infrastructures, and citizen participation to address urban challenges. Monolithic and single-deployment DTs do not address these challenges correctly. A distributed DT handles the complexity in a better way, by adapting and scaling components according to the needs. This adaptability will ensure that the DT adapts to the specific requirements and complexity of smart cities, offering flexibility and scalability as mentioned above.
- **Collaboration.** A loosely-coupled distributed DT architecture facilitates collaboration between systems or applications by providing a platform for data and knowledge exchange. Multiple applications can connect to shared twins, enabling real-time collaboration and decision-making. By leveraging a distributed architecture, DTs can be shared by different systems, leading to interoperability and resource optimization.

4 Architecture

This section presents our proposal for an Urban Digital Twin architecture in the Continuum. We will also describe which modules and components form the different layers of the architecture.

Figure 1 illustrates the proposed architecture, where the Physical Twin (PT) and the Digital Twin (DT) establish communication through a Distributed Data Lake (DDL), following the architectural conceptualization for DT systems suggested by Muñoz et al. [13]—only two Edge nodes and one Fog node have been represented for simplicity, although this structure could be expanded. The particularity of this architecture in the Continuum lies in the flexible coupling between components, managed by the Cloud, Fog, Edge, and Things layers. Figure 1 illustrates the proposed architecture, where the Physical Twin (PT) and the Digital Twin (DT) establish communication through a Distributed Data Lake (DDL), following the architectural conceptualization for DT systems suggested by Muñoz et al. [13]—only two Edge nodes and one Fog node have been represented for simplicity, although this structure could be expanded. The particularity of this architecture in the Continuum lies in the flexible coupling between components, managed by the Cloud, Fog, Edge, and Things layers.

In this context, the DT adopts a distributed architecture in which each instance of the twin operates autonomously, avoiding dependence on a centralized replica of the entire system. This approach provides flexibility and efficiency in data management, offering different levels of information and knowledge, and avoiding unnecessary data transmission and infrastructure overload. The key to achieving this objective lies in the distribution of the data lake along the Continuum so that each layer can write to and read from the corresponding DDL module. This distribution favors system adaptability, allowing each component to contribute to knowledge generation in an autonomous and collaborative way.

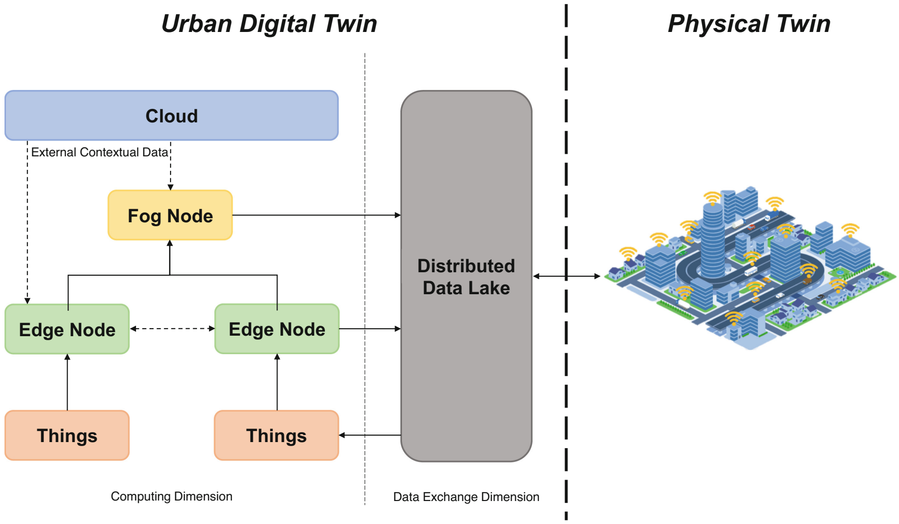

Fig. 1. Architecture proposed.

– **Things layer:** This layer focuses on the sensors that provide information to the Edge nodes. These sensors are physical devices that capture real-time data about the environment in the PT and write it in the DDL. This data is fundamental to the operation of the UDT, as it feeds the monitoring, analysis, and decision-making process. Each DT module in the Things layer gets the corresponding data from the DDL and sends it to an Edge node. This is represented in Fig. 1 with the arrow from the DDL to a Things module—only the connection with one Thing module is depicted for simplicity. Unlike Edge nodes or the Fog layer, the Things layer generally does not perform computations or data processing itself but focuses on capturing the corresponding data from the DDL and transmitting it to the Edge nodes.

- **Edge layer:** The Edge layer is composed of devices located at the periphery of the network that perform the function of collecting information from the Things layer or from other services. These Edge nodes also host DT modules that store the information collected from the Things layer that allows it to analyze and predict the behavior of the elements it monitors. Therefore, these devices must have sufficient capabilities for data collection and the provision of a DT.

  The results obtained through these analyses are written in the DDL for the PT to have access to them as well as the Fog nodes for further processing. Additionally, Edge nodes also exchange information with each other. This collaboration and information sharing contribute to improving and enriching the results generated by the DTs present in each of the nodes. By working together, the Edge nodes achieve greater accuracy and quality in the results obtained. They also increase their capacity to adapt to changes and unexpected situations, as they can benefit from the experience and knowledge shared among them.
- **Fog layer:** The Fog layer serves as a crucial intermediary, tasked with gathering and retaining data sourced from the Edge nodes. This collection of data from the Edge nodes is instrumental in providing a more expansive and nuanced understanding of the system. Within each Fog node, Digital Twin (DT) modules play a pivotal role-they not only store the information acquired from the Edge nodes but also conduct in-depth analyses, unraveling the intricacies of the data. This involves identifying patterns, discerning trends, pinpointing emerging issues, and distilling this raw data into valuable knowledge and insights.

  The knowledge thus amassed proves to be indispensable for anticipating events and making well-informed decisions. This reservoir of insights is not left untapped; instead, it is meticulously recorded in the DDL. This strategic move ensures that the Physical Twin (PT) has direct and immediate access to this wealth of information. By leveraging the capabilities of the Fog layer, the system not only enhances its capacity to comprehend the intricacies of the data but also empowers decision-makers with the foresight and understanding needed to navigate complex scenarios and make proactive, informed choices.
- **Cloud layer:** The Cloud layer plays a pivotal role in the system by serving as a repository for diverse external contextual data obtained from a multitude of sources including databases, APIs, and more. This layer is essential for consolidating and storing this contextual information, which is subsequently utilized to feed and enhance the functionality of the other layers within the architecture. By gathering and integrating data from various external sources, the Cloud layer ensures that the entire system has access to a comprehensive and up-to-date pool of information, facilitating more accurate analyses, predictions, and decision-making across the distributed digital twin architecture.

5 Case Studies

The proposed architecture demonstrates adaptability, allowing it to be applied to various case studies or applications. Depending on the scenario, different elements within the architecture can assume different roles, tailored to the requirements of each application. This versatility enables the architecture to be customized based on the problem being addressed. In this section, we explore two case studies that exemplify the adaptability of our architecture. In the first one, the government of a smart city wants to collect data about its public transportation system and how its citizens use it, with the goal of deploying a UDT to optimize and improve service to citizens. In the second one, the government has to comply with current pollution regulations and wants to implement a UDT pollution monitoring system to improve the accuracy of air quality prediction and avoid non-compliance.

5.1 Urban Transportation System

In Fig. 2, we apply the proposed architecture to a bus transportation system, focusing on improving the accuracy of predicting bus arrival times, detecting potential skipped stops, and/or knowing how citizens move around the city. To achieve this, our architecture takes into consideration both buses and passengers. Next, we present a description of the architecture implementation:

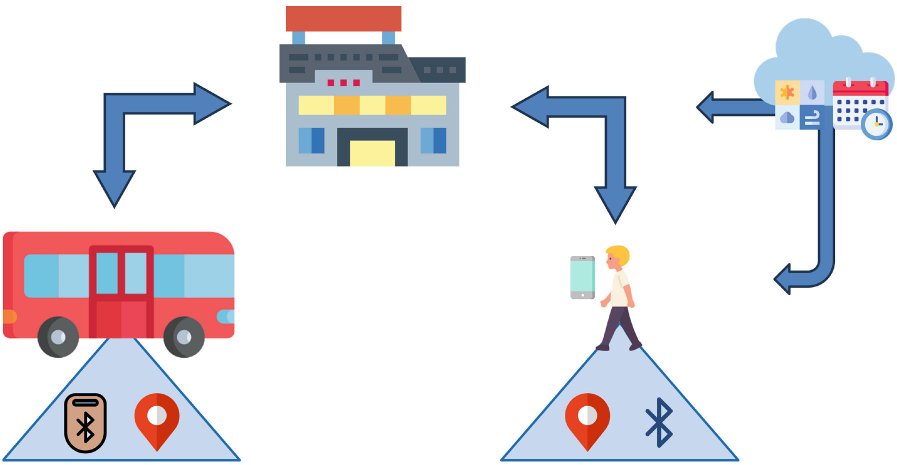

Fig. 2. Case study: Urban transportation system.

Things Layer: This layer contains Bluetooth beacons and GPS devices installed on each bus of the transportation system, as well as the sensors (GPS and Bluetooth) of passengers' smartphones. The bus' GPS tracks and locates the current position, while beacons are used to detect passengers' presence through Bluetooth and count the number of people on board. These data are essential to analyze pedestrian movement patterns and calculate whether the bus will be able to make a stop at the following destinations, considering the maximum capacity allowed.

Edge Layer: The Edge nodes represent both buses and passengers, each with their own DT. Regarding the buses, each one is represented by an Edge node consisting of a DT that simulates the bus behavior. These Edge nodes are fed with the data of each bus, coming from the information of the Things layer. The information gathered allows the DTs to more accurately predict bus arrival times and identify situations where a stop may be skipped due to capacity limitations.

Similarly, each passenger is represented by an Edge node consisting of a DT that represents and simulates the user's Digital Avatar [16] that refers to the virtual representation of a person residing on their smartphone, collecting information about their habits and preferences, and allowing them to interact with the other DTs. These Edge nodes capture the information of individual passengers within the transportation system. By integrating the data from both the buses' DTs and Digital Avatars' DTs, the Edge layer improves the accuracy of arrival time predictions and overall system optimization.

Apart from their individual roles, the Edge layer facilitates horizontal communication between buses and passengers within the transportation system. Passengers can receive real-time updates and notifications about bus schedules or delays. Furthermore, buses can also communicate with other buses within the Edge layer, promoting collaboration and coordination for better efficiency.

Fog Layer: The Fog layer is essential in the bus transportation system. Each bus line is governed by a Fog node, which contains a DT representing and simulating the buses' behavior on that specific line. These Fog nodes serve as data collection and processing points for the Edge nodes, enabling the creation of an overall model for the entire line. In addition, the Fog layer is responsible for periodically distributing the federated model to other buses to update their respective DTs. This update occurs regularly or when a new bus joins the line, ensuring all buses benefit from the latest updates in the prediction model.

Moreover, the Fog nodes may be interested in incorporating information about people's habitual travel patterns to determine if a person who usually takes a specific bus line will be using it on a particular day. This information can be leveraged to provide targeted recommendations for specific buses, ensuring that the transportation system adjusts to the individual's needs and preferences.

Cloud Layer: This layer plays a crucial role in providing external contextual information, such as weather conditions, event calendars, and vacation dates. This contextual information is used to further enrich the prediction and decision-making model. For instance, considering weather conditions allows anticipating possible delays due to rain or adverse weather factors. Likewise, by taking into account the calendar of events and vacation dates, the transportation demand prediction can be adjusted and resource allocation can be optimized.

5.2 Pollution Monitoring System

In Fig. 3, we apply the proposed architecture to a pollution control system, focusing on analyzing and predicting pollution levels of CO2 and allergenic particles to improve the accuracy of air quality prediction. Next, we present a description of the architecture implementation:

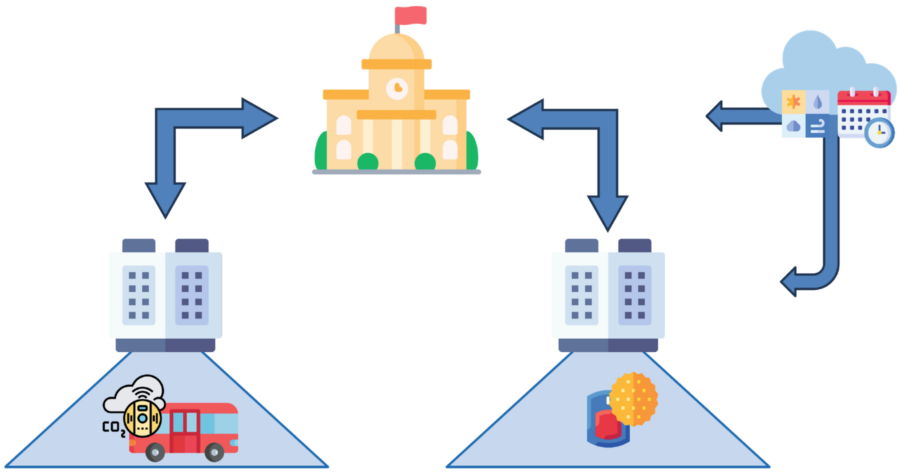

Fig. 3. Case study: Pollution monitoring system.

Things Layer: In this layer, there are sensors distributed throughout the city to measure air pollution levels, such as CO2 and suspended particles (pollutants, pollen, and other allergenic particles). CO2 sensors are located on the buses moving around the city. The data from the suspended particle sensors are strategically distributed at different points in the city.

Edge Layer: In the Edge layer, there are pollution analysis control centers distributed in different neighborhoods in the city. These centers are responsible for capturing the data coming from the Things layer. Each control center has its own DT, which represents and simulates the behavior of the center in question.

The pollution analysis control centers receive the data collected through their DTs and perform a comprehensive analysis of these data. The DTs are able to predict accurately CO2 pollution levels and suspended particle concentrations. This information is used to generate notifications and alerts to citizens through information panels. These panels display real-time and predictive information on air quality, providing relevant data on CO2 levels and allergenic particles. In this way, citizens can be informed about the air quality in their environment and receive alerts in case of risk situations or high levels of allergenic particles or activate action protocols to address the problem.

Fog Layer: The Fog layer is represented with one node by the city's air management system. Its main function is to collect and process the information coming from the different control centers. The information collected by the Fog node is essential to feed and enrich the DT present in this layer. The DT uses historical data to analyze and evaluate air quality in the city. By analyzing long-term patterns and trends, the DT can provide a complete and detailed picture of pollution and allergy levels.

The results are critical to making informed decisions. For example, if an area with high CO2 levels is detected, authorities can take measures to regulate traffic in that area, reduce pollutant emissions, and comply with regulations set by the European Union. Regarding the concentration of allergenic particles, if high levels of allergenic particles are predicted in certain areas of the city, this will imply an increase in allergy cases in the local population. The health system can use this information to take proactive measures and provide better resource allocation and better care to affected citizens.

Cloud Layer: The Cloud layer is essential for integrating data sources from meteorological portals and event calendars to enrich the models of the different DTs. For instance, weather conditions, such as wind speed and precipitation, along with temperature and humidity, impact the dispersion of pollutants and concentration of allergenic particles in the air. Additionally, integrating event calendar data helps identify activities like sports events, concerts, or festivals that can lead to increased traffic and crowds, directly affecting air quality.

6 Discussion

In this section, we analyze the proposed architecture for the Urban Digital Twin (UDT) over the Continuum. We analyze how this proposal stands out for its scalability, flexibility, and privacy enhancement in complex urban environments. In addition, we carefully examine challenging considerations, such as implementation complexity, infrastructure requirements, and security-related issues. This critical assessment provides insight into the suitability and feasibility of the UDT architecture in the Continuum, offering a holistic perspective for consideration in smart city management. The issues addressed position it as a robust and adaptable solution for the efficient management of urban environments. However, like any proposal, it also presents certain limitations.

Firstly, the distributed architecture in the Continuum allows for efficient scalability. Each entity deploys its own Digital Twin, which facilitates the incorporation of new urban elements without affecting the existing infrastructure. This ensures exceptional adaptability as the city evolves and expands, enabling the seamless integration of new services and devices. Nevertheless, the implementation of a distributed architecture composed of different entities can be complex, requiring careful planning and coordination. This can lead to challenges in terms of integration and initial configuration, especially in urban environments already established with pre-existing systems. It is therefore important to consider economic and technological feasibility, especially for those cities with limited resources.

Secondly, the low coupling between the DTs of the different entities enables their independent operation. In addition, the architecture enables collaboration between them, thus enriching their models and improving joint decision-making. This optimizes operational efficiency while fostering collaboration and synergy between the different urban components. This low coupling brings inherent privacy improvements as information is stored and processed at the data source. This feature is fundamental to gaining the trust of citizens and ensuring compliance with privacy regulations. This low coupling in data management also poses additional challenges in terms of security and protection against threats. Careful implementation of security measures is required to ensure data integrity and confidentiality, as well as resilience to potential attacks.

The Urban Digital Twin architecture in the Continuum offers numerous advantages, from its flexibility to its privacy enhancement. However, its effective implementation requires carefully addressing the associated constraints, ensuring a smooth transition and long-term benefits for smart city management.

7 Conclusions and Future Work

Recently, Digital Twins have emerged as a powerful enabling technology for the virtualization of systems, products, and services. They represent fundamental tools for monitoring, prediction, and integration of data from IoT devices. Despite their benefits, conventional implementations face challenges such as scalability, coupling, or privacy, which can affect their ability to adapt optimally in scenarios with a high distribution of resources. To try to overcome these challenges, there are architectures such as the Computing continuum that enable the flexibility and distribution capabilities inherent to Digital Twins, thus ensuring their usefulness in diverse environments and applications.

In this paper, we have presented an architecture for Urban Digital Twins deployed in the Continuum, which addresses the limitations of conventional DTs in terms of scalability, latency, and privacy, among others. Through the distribution of DTs in the Continuum, we achieve greater flexibility and responsiveness, as well as more efficient data and privacy management. In addition, through citizen participation, we can ensure personalized data-driven decision-making. The case studies presented demonstrate the applicability and effectiveness of our architecture in the context of smart cities, enabling more accurate and participatory management of urban environments.

We are currently implementing the architecture for both case studies, setting up the necessary infrastructure, and deploying the different components across the layers. We have already obtained some preliminary implementation results, presented in [10,24]. Furthermore, in [8] we introduce Perses, a tool that emulates different parts of the Computing Continuum, including user smartphones.

Acknowledgements. This work has been partially funded by grant DIN2020-011586, funded by MCIN/AEI/10.13039/501100011033 and by the European Union "Next

GenerationEU/PRTR", by the Spanish Ministry of Science, Innovation, and Universities (projects PID2021-125527NB-I00, TED2021-130913B-I00, TED2021-130523B-I00, and PD C2022-133465-I00), by the Regional Ministry of Economy, Science and Digital Agenda of the Regional Government of Extremadura (GR21133) and the European Regional Development Fund.

References

1. Barricelli, B.R., Casiraghi, E., Fogli, D.: A survey on digital twin: definitions, characteristics, applications, and design implications. IEEE Access **7**, 167653–167671 (2019)
2. Bendechache, M., Svorobej, S., Takako Endo, P., Lynn, T.: Simulating resource management across the cloud-to-thing continuum: a survey and future directions. Future Internet **12**(6), 95 (2020)
3. Costantini, A., et al.: Iotwins: Toward implementation of distributed digital twins in industry 4.0 settings. Computers **11**(5), 67 (2022)
4. Deng, T., Zhang, K., Shen, Z.: A systematic review of a digital twin city: a new pattern of urban governance toward smart cities. J. Manag. Sci. Eng. **6**(2), 125–134 (2021)
5. Fadda, M., Anedda, M., Girau, R., Pau, G., Giusto, D.D.: A social internet of things smart city solution for traffic and pollution monitoring in Cagliari. IEEE Internet Things J. **10**(3), 2373–2390 (2022)
6. Fuller, A., Fan, Z., Day, C., Barlow, C.: Digital twin: enabling technologies, challenges and open research. IEEE Access **8**, 108952–108971 (2020)
7. Laso, S., et al.: Elastic data analytics for the cloud-to-things continuum. IEEE Internet Comput. **26**(6), 42–49 (2022)
8. Laso, S., Berrocal, J., Fernández, P., Ruiz-Cortés, A., Murillo, J.M.: Perses: a framework for the continuous evaluation of the QoS of distributed mobile applications. Pervasive Mob. Comput. **84**, 101627 (2022)
9. Laso, S., Berrocal, J., García-Alonso, J., Canal, C., Manuel Murillo, J.: Human microservices: a framework for turning humans into service providers. Software: Practice Exp. **51**(9), 1910–1935 (2021)
10. Laso, S., Toro-Gálvez, L., Berrocal, J., Canal, C., Murillo, J.M.: Deploying digital twins over the cloud-to-thing continuum. In: 2023 IEEE Symposium on Computers and Communications (ISCC), IEEE (2023)
11. Lehtola, V.V., et al.: Digital twin of a city: review of technology serving city needs. Int. J. Appl. Earth Obs. Geoinf. **114**, 102915 (2022)
12. Lei, B., Janssen, P., Stoter, J., Biljecki, F.: Challenges of urban digital twins: a systematic review and a Delphi expert survey. Autom. Constr. **147**, 104716 (2023)
13. Muñoz, P., Troya, J., Vallecillo, A.: A conceptual architecture for building digital twins. In: Proceedings of the 3rd International Workshop on MDE for Smart IoT Systems (MeSS 2023) (2023)
14. Nikitas, A., Michalakopoulou, K., Njoya, E.T., Karampatzakis, D.: Artificial intelligence, transport and the smart city: definitions and dimensions of a new mobility era. Sustainability **12**(7), 2789 (2020)
15. Pérez-Vereda, A., Canal, C.: A people-oriented paradigm for smart cities. In: Cabot, J., De Virgilio, R., Torlone, R. (eds.) ICWE 2017. LNCS, vol. 10360, pp. 584–591. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-60131-1_46
16. Pérez-Vereda, A., Canal, C., Pimentel, E.: Modelling digital avatars: a tuple space approach. Sci. Comput. Program. **203**, 102583 (2021)

17. Pujol, V.C., Raith, P., Dustdar, S.: Towards a new paradigm for managing computing continuum applications. In: 2021 IEEE Third International Conference on Cognitive Machine Intelligence (CogMI), pp. 180–188. IEEE (2021)
18. Ramu, S.P., et al.: Federated learning enabled digital twins for smart cities: concepts, recent advances, and future directions. Sustain. Urban Areas **79**, 103663 (2022)
19. Ruiz, P., Seredynski, M., Torné, Á., Dorronsoro, B.: A digital twin for bus operation in public urban transportation systems. In: Big Data Intelligence and Computing, pp. 40–52 (2023)
20. Samih, H.: Smart cities and internet of things. J. Inf. Technol. Case Appl. Res. **21**(1), 3–12 (2019)
21. Schrotter, G., Hürzeler, C.: The digital twin of the city of Zurich for urban planning. PFG - J. Photogrammetry Remote Sens. Geoinf. Sci. **88**, 99–112 (2020)
22. Singh, M., Fuenmayor, E., Hinchy, E.P., Qiao, Y., Murray, N., Devine, D.: Digital twin: origin to future. Appl. Syst. Innov. **4**(2), 36 (2021)
23. Technology, L.: LADOT. Los Angeles. Technology Action Plan; 2022. https://ladot.lacity.org/sites/default/files/documents/annual-report-2022_2023.pdf. Accessed 30 Nov 2023
24. Toro-Gálvez, L., García-Luque, R., Troya, J., Canal, C., Pimentel, E.: Towards the integration of digital avatars in urban digital twins on the cloud-to-thing continuum. In: Proceedings of the 3rd International Workshop on Big data driven Edge Cloud Services (BECS 2023) (2023)
25. Villalonga, A., Negri, E., Fumagalli, L., Macchi, M., Castaño, F., Haber, R.: Local decision making based on distributed digital twin framework. IFAC-PapersOnLine **53**(2), 10568–10573 (2020)
26. Weil, C., Bibri, S.E., Longchamp, R., Golay, F., Alahi, A.: A systemic review of urban digital twin challenges, and perspectives for sustainable smart cities. Sustain. Cities Soc. 104862 (2023)