

Applications

Martin Bauer*, Flavio Cirillo, Jonathan Fürst, Gürkan Solmaz and Ernö Kovacs

Urban Digital Twins – A FIWARE-based model

Städtische digitale Zwillinge – Ein FIWARE-basiertes Modell

<https://doi.org/10.1515/auto-2021-0083>

Received June 9, 2021; accepted September 27, 2021

Abstract: This article describes the use of digital twins for smart cities, i. e., the *Urban Digital Twin (UDTw)* concept. It shows how UDTws can be realized using the open source components from the FIWARE ecosystem that are already used in more than 200 cities worldwide. The used NGSI-LD standard is supported by the European Connecting Europe Facility, the Open and Agile Smart City community, the Indian Urban Data Exchange platform, and the Japanese Smart City Reference Model. Unlike digital twins in other domains, e. g., manufacturing, where digital twins are co-developed with their physical counterparts, UDTws often evolve driven by different stakeholders, on different time scales, as well as by utilizing many different data sources from the city. This article builds on a well-established lifecycle model for Digital Twins and combines this with a conceptual model for digital twins consisting of data, reactive, predictive and forecasting (“what if”) digital twin functionalities. The article also describes how AI-based technologies can be used to extract knowledge to build the UDTws from the IoT-based infrastructure of a smart city.

Keywords: Digital Twin, Internet of Things, data sharing, Smart City

Zusammenfassung: Dieser Artikel beschreibt die Verwendung von digitalen Zwillingen für Smart Cities, d.h. das Konzept des städtischen digitalen Zwilling (Urban Digital Twin). Es zeigt, wie städtische digitale Zwillinge mit den Open Source Softwarekomponenten aus dem FIWARE-Ökosystem realisiert werden können, die bereits in mehr als 200 Städten weltweit im Einsatz sind. Der benutzte NGSI-LD-Standard wird von der Connecting Europe Facility, der Open Agile Smart City Community, der Indian Urban Data Exchange Plattform und dem Japanischen Smart

City-Referenzmodell unterstützt. Im Gegensatz zu digitalen Zwillingen in anderen Domänen, z. B. in der Fertigung, wo digitale Zwillinge gemeinsam mit ihren physischen Gegenstücken entwickelt werden, werden städtische digitale Zwillinge von unterschiedlichen Akteuren zu unterschiedlichen Zeiten sowie auf Basis vieler verschiedener Datenquellen aus der Stadt entwickelt. Dieser Artikel baut auf einem bewährten Lebenszyklusmodell für digitale Zwillinge auf und verbindet dies mit einem konzeptionellen Modell für digitale Zwillinge. Dieses besteht aus Daten-, reagierenden, voraussagenden und prognostizierenden („was wäre wenn“) Funktionalitäten. Der Artikel beschreibt auch, wie KI-basierte Technologien verwendet werden können, um Wissen zu extrahieren, und damit aus der IoT-basierten Infrastruktur der Smart City städtische digitale Zwillinge zu bauen.

Schlagwörter: Digitaler Zwilling, Internet der Dinge, Datenaustausch, Smart City

1 Introduction

Digital twins are an essential concept that has already been developed in the manufacturing industry [1, 2, 3, 4], and is now getting more attention in other applications areas as well [5]. In this paper, we explore the application of the digital twin concept in cities, i. e., creating Urban Digital Twins (UDTw).

As in many other areas, the digital transformation of smart cities needs to break down the barriers of disconnected information silos. One reason for the lack of sharing is the low-level representation and tight coupling of applications and system components making any reuse difficult. To break these silos and enable sharing knowledge and gaining relevant insights, suitable concepts, representations, and interfaces are needed. This is especially true for cities with their scale, heterogeneous structures, and a multitude of stakeholders involved. The Internet-of-Things (IoT) and the related Cyber-Physical Systems (CPS) build the technology foundation on top of which the UDTws will be built.

*Corresponding author: Martin Bauer, NEC Laboratories Europe, Heidelberg, Germany, e-mail: martin.bauer@neclab.eu

Flavio Cirillo, Jonathan Fürst, Gürkan Solmaz, Ernö Kovacs, NEC Laboratories Europe, Heidelberg, Germany, e-mails: flavio.cirillo@neclab.eu, jonathan.fuerst@neclab.eu, gurkan.solmaz@neclab.eu, ernoe.kovacs@neclab.eu

UDTws create value for smart cities through the total representation of cities (and even their citizens) in cyberspace – from city planning to end-of-life (Digital Twin Lifecycle). UDTws can have different purposes ranging from purely representing the real-world assets in cyberspace to providing various services, e. g., through reactive optimization, prediction, and simulation.

In the operational phase, the UDTw enables monitoring the city, its buildings, and the activities happening in the city. City objects in a Smart City are utilized by many different processes, e. g., energy management, cleaning, maintenance, operation. Many stakeholders are involved, thus making UDTws shared resources through which a dynamic set of stakeholders can interact, share data, and collaborate. UDTws become the centerpiece of understanding and managing the city as a system-of-system.

The contributions of this paper are (a) use cases for smart cities creating value through using UDTws, (b) conceptualizing the DTw lifecycle, as well as the DTw usage spectrum, (3) concrete realizations of the UDTw concept utilizing the NGSI-LD standard and components from the FIWARE [6] open source ecosystem, and (4) AI technologies supporting the creation and maintenance of the various digital twin models. Finally, we point to open research and technology questions in connection with UDTws.

2 Use cases

UDTws can be applied to a wide range of use cases in urban areas: from mobility to energy, resilience and improving the life of citizens. For example, in the energy domain, digital twins can be used to model the interplay of energy production, energy storage and energy usage [7]. In general, the UDTw can be applied to improve the resilience of the different critical infrastructures in cities like energy, water, communication and traffic, as they are investigated in projects like emergenCITY [8] or Future Resilient Systems [9].

In the following, we discuss two example use cases of urban digital twins in the mobility area: (1) creation of UDTws from heterogeneous data sources (the City Data Lake) and (2) using UDTws for supporting autonomous driving.

Creating Urban Digital Twins from heterogeneous data sources

The Connecting Europe Facility (CEF) funded project ODALA develops and deploys a unified ‘Hybrid Data Lake’

for cities. This data lake will facilitate the development of advanced AI applications and the digital twin transition of cities at large, demonstrated with use cases from the mobility and environmental sensing domain. ODALA follows the principles established in SynchroniCity, a large-scale European pilot that implemented 35 city services in 27 smart cities in Europe and South Korea [10]. In particular, ODALA contributes to the management of heterogeneous data lakes and the AI-based creation of UDTws.

Figure 1 depicts how data from heterogeneous, silo-ed data sources is ingested and enriched to connected digital twins in the ODALA data lake. Data acquired by sensors are used to represent a digital twin and to infer current insights, e. g., air quality status for the “city district twin” or occupancy for the “bus twin”. These insights are used for predicting situations based on historical records. Further, a “city district twin” infers traffic status using cameras, interacting with the “bus twin”(s) and vehicle positions and speed. In return, the “bus twin” can estimate the time of arrival. The historical data can then be used to predict traffic. The “cyclist twin” might use the insights of the other twins to simulate the best route for the best cycling experience. For example, the application might simulate the usage of a bus for the longest segments of the journey. In some cities, public transportation regulations might not allow bikes to be carried on a bus when a bus is crowded. Thus, the estimation of occupancy of the “bus twin” helps to create an optimal route for the users or suggest to the user to start the travel some minutes later. Urban planning might use all these UDTws and their interaction to support human decisions on planning an urban environment more friendly to cyclists and the environment.

Finally, UDTws altogether are a live snapshot of an urban environment with always up-to-date information and historical records of past events. This information might be of vital importance when faced with disaster. For example, the real-time estimation of traffic and the continuous tracking of the bus fleet might be crucial to identify the areas more threatened by ongoing flash floods to quickly deploy a rescue crew. Also, bus crowd estimation might be used for epidemic simulations [11].

Urban digital twins in autonomous driving

Figure 2 shows the concept of digital twins in an autonomous driving scenario. The physical car is driving in the real environment. It has a limited view of this environment, which is based on information from its own sensors. This physical car has a digital twin. The digital twin is up-

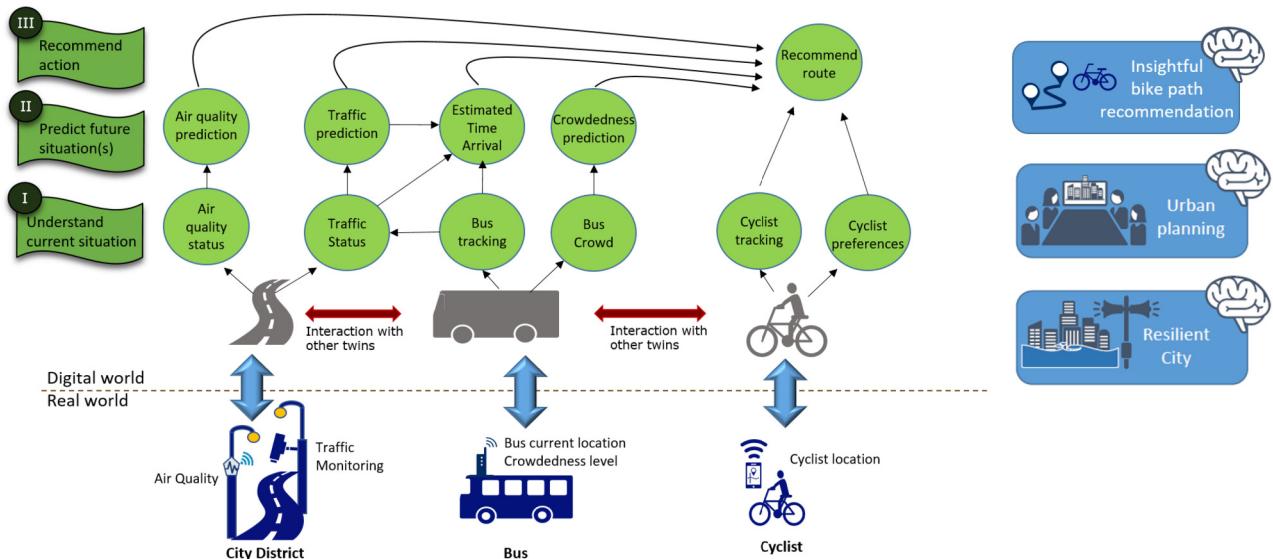


Figure 1: ODALA Use Case. Data from heterogeneous sources is enriched to UDTws.

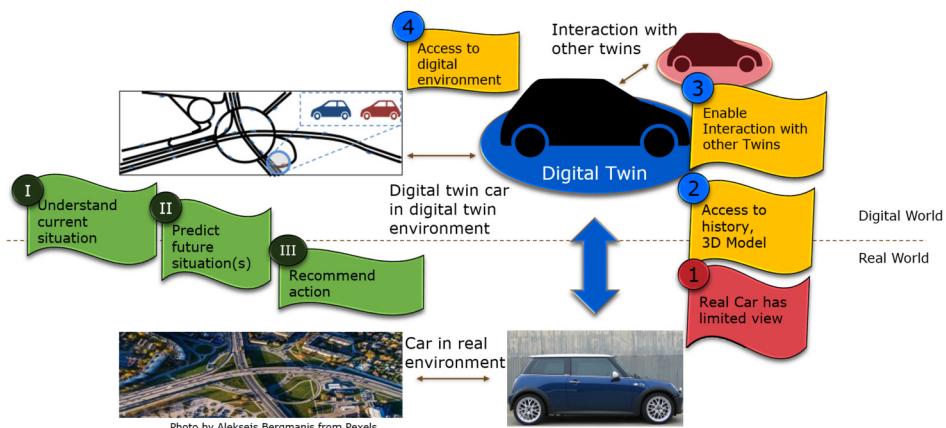


Figure 2: Digital twins in autonomous driving.

dated with the current sensor information from the real car, but also has access to historical information and models of relevant aspects of the car, e. g., a 3D model. Furthermore, the digital twin can interact with other digital twins – the digital twins of other cars, but also digital twins representing the environment, e. g., a quarter of the city or an intersection with all the roads connected to it. The sum of all this information would be too large to transfer to the car in real time and in the car there might not be enough processing capacity.

While the digital twin of the car can be tightly coupled and developed together with the physical car (DTw lifecycle or DTw continuity), the digital twins of the environment are more independent constructs, which need to be fed from multiple, independent sources, not necessarily

deployed in connection with the creation of the digital twin.

Through the interaction with other digital twins, the digital twin of the car has access to a wealth of information, which is not limited to only the sensors of the car. This information can be used to better understand the current situation, predict future situations and, on this basis, recommend the best possible actions [12]. For example, as the digital twin knows what is happening around the bend, it can choose the optimal path or recommend slowing down early, if the traffic ahead has come to a stop.

Challenges

For the above use cases we identify several challenges. (C1 – Federation): *UDTws are distributed across a federation*

of systems and handled by different organizations. Applications and UDTws need to seamlessly work regardless of the complexity and fragmentation of the systems. (C2 – Heterogeneity): *data are highly heterogeneous* in terms of formats and models – formats ranging from raw sensors reading to a 3D model or camera images. Further, the data models are coming from many independently engineered devices, each providing data with a different data schema. (C3 – Dynamics) *Dynamic characteristics of the UDTws realm*: The set of digital twins changes continuously, such as vehicles entering and exiting the city. Applications and services should be automatically orchestrated to serve and interact with new UDTws. (C4 – Programming Complexity): Service developers need to discover the available UDTws and available data to establish data flows to/from the service. (C5 – Regulations) *Privacy and Data Usage Control*: regulations (such as GDPR) for personal data management and data licenses need to be respected.

3 Digital Twin concept

In the previous chapters, we have already established the Digital Twin lifecycle: In the product planning phase, the DTw Prototype represents the product and can be used, e.g., for prediction and simulation of the real product. When a product is leaving the assembly line, the DTw prototype is transformed into a DTw instance, partially containing the product planning data (like the 3D model), partially storing the real-time data coming from the real asset. Advanced DTw Instances make use of the DTw world by interacting with other DTws. At the end of life, the DTw will be archived and re-used for improving the product.

Figure 3 shows the different digital twin functionalities. At the core is the data twin functionality and – from becoming a DTw instance – the interaction between the physical twin and the data twin functionality. A predictive twin functionality¹ can be used for planning purposes, i.e., the predictive twin functionality can be created from the DTw prototype. Such a predictive twin functionality can be used for simulations and the desired properties can be optimized, so the physical twin will perfectly fulfil the requirements and fit into its environment. For example, the predictive twin functionality can be used when planning a smart factory or when planning a new city quarter.

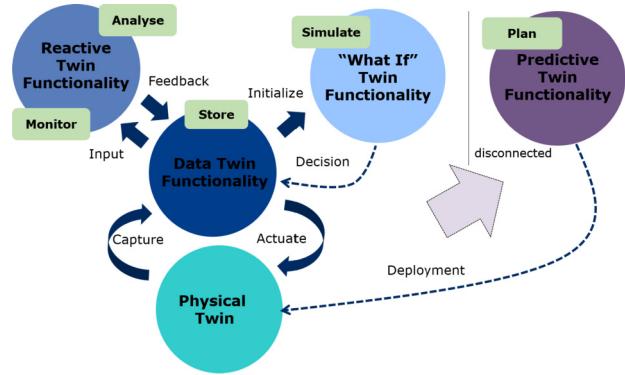


Figure 3: Digital twin functionalities.

During the construction process, the predictive twin functionality can be used for monitoring the progress against plan and detect any discrepancies, so they can be immediately addressed.

The physical twin (a real-world entity like a car) has sensors of its own, capturing important aspects of its state, e.g., the speed. Especially in urban scenarios, external sensors like traffic cameras can also capture status information. The data twin functionality stores the captured information of the physical twin for later use. A reactive twin functionality² uses this data stream for monitoring the current behaviour or historical analysis. This derived insights can be fed back to the data twin functionality. In addition, the “what if” twin functionality can be initialized using the current information from the DTw world, combined with simulation model assumptions. It then simulates what can happen in the future. Reactive analysis, prediction, as well as “what-if” simulation can be used to decide on which actuation should be executed on the physical twin [12].

4 Digital Twin modelling with NGSI-LD

The data twin functionality handles all data, which models the physical twin. The NGSI-LD Information Model [13] provides the basis for this. Its core concept, the *entity* can be used as a basis for the data twin functionality. The NGSI-LD Information Model provides the following information:

¹ <https://www.youtube.com/watch?v=TwOWAISeS68>, by Brian Matthews (accessed 20/09/2021).

² <https://www.youtube.com/watch?v=TwOWAISeS68>, by Brian Matthews (accessed 20/09/2021).

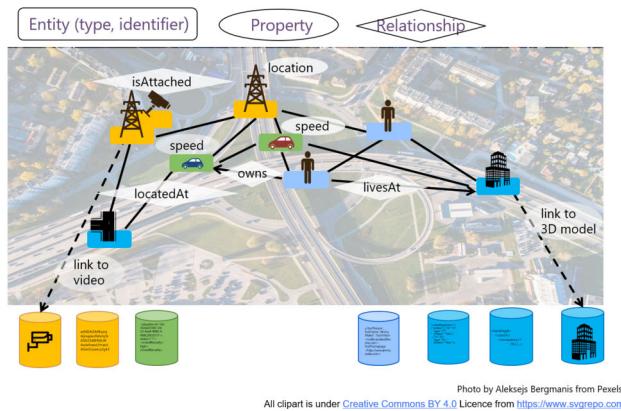


Figure 4: NGSI-LD Information Model Instance.

- an identifier, represented as a URI to uniquely identify the entity
- a known type (defined in an ontology), also represented as a URI, which allows discovering relevant entities and defines which information is required or may optionally be available for the entity
- properties representing relevant aspects of the entity, e.g., sensor information or characteristics of the entity
- relationships to other entities

NGSI-LD is an Information Model [13] and API [14] specified by the ETSI Industry Specification Group on Context Information Management (ETSI ISG CIM).³ ETSI is the European Telecommunications Standards Institute, one of three bodies recognized by the European Commission as a European Standards Organization. NGSI-LD is the latest step in the evolution of the NGSI Context Interfaces [15], originally standardized by the Open Mobile Alliance (OMA)⁴ in 2010. NGSI-LD information is represented in the JSON-LD format, where LD stands for Linked Data.

Figure 4 shows an example with digital twins modelled as entities using the NGSI-LD Information Model. Entities of type building, car, person, intersection, power pole and camera are shown as rectangles with rounded corners (addressing challenge (C2 – Heterogeneity) and partially (C1 – Federation)). The color represents the data source to which they belong or from which they have been extracted. These data sources have often originally been set up as information silos for specific use cases and have later been integrated into the cross-domain digital twin model.

Entities can have properties. In Figure 4, a few example properties are shown as ovals, e.g., the cars have a property *speed* and a power pole has a property *location*. The black arrows with diamond-shaped labels show relationships between entities. For example, one person *owns* one of the cars and *lives at* a building, one car is *located at* an intersection and a camera is *attached to* a power pole.

The underlying model is a property graph with entities as nodes and relationships as edges. As not all kinds of information can easily be represented in property graph format, there can be links to other information sources (C2 – Heterogeneity). The camera has a link to its video stream, the building has a link to its 3D model. Meta information, e.g., the data format of the 3D model and the unit of car speed, can be modelled as a property of properties.

Modeling digital twins with a property-graph-based model allows sharing information among many different stakeholders on a suitable abstraction level. Information can easily be found and enhanced with new insights (C3 – Dynamics).

5 Infrastructure for Digital Twins

NGSI-LD information can be stored and made available using Context Brokers implementing the NGSI-LD API [14]. Data twins can be implemented on the basis of NGSI-LD. The NGSI-LD API enables access to information about a specific digital twin and the search for and discovery of relevant other digital twins (C2 – Heterogeneity) (C3 – Dynamics). The API allows filtering according to properties, e.g., only cars that are moving with a certain speed; all vehicles within an area specified by geographic coordinates. The NGSI-LD API provides operations to synchronously request information (i.e., query-response) and operations for subscribing to information resulting in asynchronous notifications, triggered either by changes in the information itself or the expiration of a time interval (C4 – Programming Complexity).

The data twins are fed with information from different information sources, in particular sensors. These sensors can be attached directly to the respective physical twin. They can be located in the environment or even come from other physical twins whose sensors provide relevant information about this particular digital twin. In the other direction, actuators can be triggered by information updated, e.g., as provided by the reactive twin functionality.

³ <https://www.etsi.org/committee/cim> (accessed 20/09/2021).

⁴ <https://omaspecworks.org/> (accessed 20/09/2021).

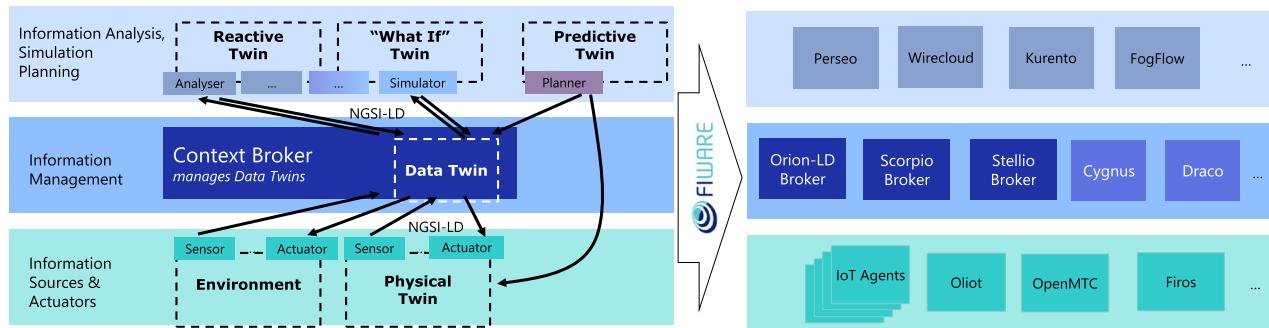


Figure 5: Digital twin infrastructure (left) and respective FIWARE components (right).

Reactive twin functionalities can consist of analyzer components, “what if” twin functionalities of simulator components, and predictive twin functionalities of planner components. Of course, actual components can also combine different functionalities, e. g., integrate a combination of analyzer and simulator components. Analyzer, simulator, and planner components access information from the data twin and feed their results back, which can ultimately trigger actuation.

The NGSI-LD API is the current version of the core API of the FIWARE⁵ open source ecosystem. FIWARE has evolved from the platform project of the European Future Internet Public-Private-Partnership (PPP) to a community driving curated framework of open source software components (see FIWARE Foundation). As shown in Figure 5 (right side), FIWARE offers a number of components (available through the FIWARE Catalogue⁶) that can be utilized to build the digital twin infrastructure sketched in Figure 5.

For information management, FIWARE offers three alternative Context Broker implementations – Orion-LD, Scorpio and Stellio – that all implement the NGSI-LD API and cater for different deployment environment, e. g., using limited resources on the edge, scaling in the cloud, or federating multiple systems together (see below for details). In addition, there are specialized components like Cygnus, which connects to different databases, and Draco, which supports the flow of information between systems. The FIWARE Context Broker is also a building block of the Connecting Europe Facility (CEF)⁷ that provides standard-

based building blocks to facilitate delivery of digital public services across borders.

As Information Sources and Actuators, FIWARE provides a number of different IoT Agents for connecting to IoT devices, particularly sensors and actuators. For example, there is support for Ultralight, LoRaWAN, LWM2M, OPC UA, and simple JSON sources. In addition, Oliot [16] connects to GS1 standards, OpenMTC [17] is a oneM2M implementation and Firos [18] enables the connection to ROS-based robots (C2 – Heterogeneity).

Finally, for information analysis, simulation and planning, there are a number of processing and visualization components. Perseo is a complex event processing tool (currently still based on the predecessor of NGSI-LD), Wirecloud is a web application mashup platform, Kurento [19] provides stream-oriented support for multimedia applications, and FogFlow [20] is an IoT edge computing framework that automatically orchestrates data processing flows over cloud and edges (C4 – Programming Complexity).

As mentioned before and shown in the use cases, urban digital twins are not always centralized components that are only connected to their physical counterparts. Instead, different (sub-)systems have their own view of the digital twin, e. g., storing partial information. If these systems are connected in a system-of-systems setup, the different views can be federated, resulting in an aggregated digital twin view as visualized in Figure 6.

The NGSI-LD API enables such a federated architecture (C1 – Federation), where individual brokers can register the information they can provide at a federation broker. On a request, the federation broker queries the brokers that can contribute information, aggregates the information returned by the brokers and provides it to the requester. In another configuration option, instead of the hierarchy of Figure 6, the federation brokers might be configured as peers.

5 <https://www.fiware.org/> (accessed 20/09/2021).

6 <https://www.fiware.org/developers/catalogue/> (accessed 20/09/2021).

7 <https://ec.europa.eu/cefdigital/wiki/display/CEFDIGITAL/Context+Broker> (accessed 20/09/2021).

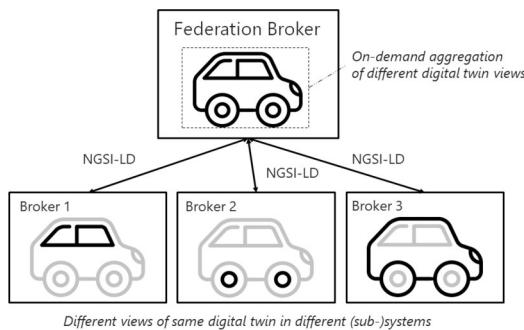


Figure 6: Federating different digital twin views.

Developing UDTws as well as applications and services that interact with UDTws are not a trivial task due to the dynamic nature of UDTws (C3 – Dynamics) (C4 – Programming Complexity). Typically, DTw services are implemented using REST APIs (or other RPC-like mechanisms). The FogFlow framework [21] enables a distributed programming model and the dynamic orchestration of information processing across a cloud/edge environment, which can be used to deploy reactive, “what if” and predictive twin functionalities.

Finally, a recent experimental implementation of FogFlow, namely IntentKeeper [22], permits to automatically handle usage control policies of data owners (e.g., GDPR policies) (C5 – Regulations). It orchestrates service instances, forcing them to comply with the policies (e.g., anonymize data before use) without any involvement of either the service developer or service consumer in the process.

of the captured information is only encoded in the application with which the sensor was originally deployed and cannot easily be reused.

On the target side, the data twin functionality requires that the NGSI-LD model is followed, i. e., all information is represented as entities with properties and relationships. An additional data model or ontology is needed to know what types of entities are modeled and which properties and relationships they can have. FIWARE, in collaboration with IUDX and TM Forum, has created the Smart Data Model⁸ initiative, which defines NGSI-LD compatible data models. The goal is to develop a homogeneous set of models across different IoT domains.

To integrate heterogeneous information in a data twin functionality, a mapping between the often implicit source model and the NGSI-LD compliant digital twin data model is needed. Figure 7 shows the different steps needed for extracting knowledge, i. e., semi-automatically translating information from the source information to the digital twin model.

In step 0, if not explicitly given, the underlying source model has to be extracted from the instance data, e. g., taking into account the terms and the structure of the data. In step 1, the source model has to be matched to the data twin model. We apply a weakly supervised machine-learning approach called knowledge infusion [23] to combine various matching heuristics for programmatic labeling [24]. The labels are then used to train a machine learning model that performs the final predictions needed to match source data model to the data twin model [25]. In step 2, example data is annotated with the candidate data twin concepts identified in step 1. This provides the basis for a human expert to check and complete the concept mapping in step 3. Based on this concept mapping and the required syntax adaptation, the translation is configured. A fully automated translation would be desirable but is (not yet) realistic given the relatively low-level information representations used by information sources.

⁸ <https://github.com/smart-data-models> (accessed 20/09/2021).

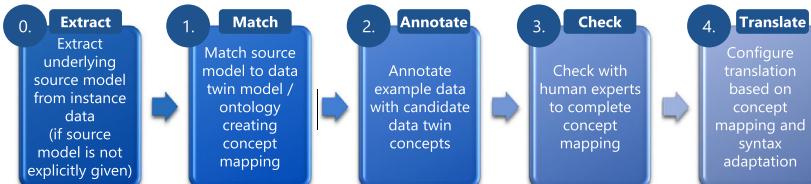


Figure 7: Translation of source information to digital twin model.

7 Conclusions

The concept of digital twins has evolved and broadened its scope to include urban digital twins as a concept for implementing smart cities. UDTws are not necessarily co-developed together with their physical counterparts but may be developed independently – sometimes later, to add smart functionality to existing city elements, and sometimes even earlier, for planning purposes. The NGSI-LD information model can be used for modeling the data of digital twins and the NGSI-LD API for discovering, accessing, and managing digital twin data. Open source components from the FIWARE ecosystem, based on NGSI-LD, can be used to implement the digital twin infrastructure, including the reactive, “what if” and predictive twin functionalities.

The challenge when developing UDTws as part of pre-existing smart city deployments is to extract and integrate information from existing low-level information sources such as sensors into the UDTw. AI techniques such as weakly-supervised machine learning provide a promising basis for this task and continue being the focus of our ongoing research activities.

We believe UDTws will open up new venues for improvements and collaboration with the control community, where UDTws and autonomous systems interact with each other for recognition and future prediction of the environment and the states of dynamic systems using techniques of AI and simulations. Autonomous systems include autonomous vehicles, drones, and robots in smart cities. Computing and communication systems can also dynamically adapt based on the UDTw concept.

Funding: This work is supported in part by the CEF ODALA project and has received funding from the European Union’s Connecting Europe Facility programme under the grant agreement No. INEA/CEF/ICT/A2019/2063604.

Conflict of interest: The content of this paper does not reflect the official opinion of the European Union. Responsibility for the information and views expressed therein lies entirely with the authors.

References

1. F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang and F. Sui, “Digital twin-driven product design, manufacturing and service with big data,” *The International Journal of Advanced Manufacturing Technology*, vol. 94, no. 9, pp. 3563–3576, 2018.
2. Q. Qi and F. Tao, “Digital twin and big data towards smart manufacturing and Industry 4.0: 360 degree comparison,” *IEEE Access*, vol. 6, pp. 3585–3593, 2018.
3. G. Shao and M. Helu, “Framework for a digital twin in manufacturing: Scope and requirements,” *Manufacturing Letters*, vol. 24, pp. 105–107, 2020.
4. C. Zhang, G. Zhou, H. Li and Y. Cao, “Manufacturing blockchain of things for the configuration of a data-and knowledge-driven digital twin manufacturing cell,” *IEEE Internet of Things Journal*, vol. 7, no. 12, pp. 11884–11894, 2020.
5. J. Conde, A. Munoz-Arcentales, A. Alonso, S. Lopez-Pernas and J. Salvachua, “Modeling digital twin data and architecture: A building guide with fiware as enabling technology,” *IEEE Internet Computing*, 2021. doi:10.1109/MIC.2021.3056923.
6. F. Cirillo, G. Solmaz, E. L. Berz, M. Bauer, B. Cheng and E. Kovacs, “A standard-based open source iot platform: Fiware,” *IEEE Internet of Things Magazine*, vol. 2, no. 3, pp. 12–18, 2019.
7. S. Velickov, “Energy 4.0: Digital Twin for Electric Utilities, Grid Edge and Internet of Electricity,” April 2020.
8. M. Hollick, B. Freisleben, M. Knodt, G. Hornung, M. Mühlhäuser, C. Fraune and J. I. Engels, “Resiliente informations- und kommunikationstechnologie für ein krisenfestes deutschland,” Oktober 2020.
9. J. Tang and H. Heinemann, “A resilience-oriented approach for quantitatively assessing recurrent spatial-temporal congestion on urban roads,” *PLOS ONE*, vol. 13, e0190616, 01 2018.
10. F. Cirillo, D. Gómez, L. Diez, I. Elicegui Maestro, T. B. J. Gilbert and R. Akhavan, “Smart city iot services creation through large-scale collaboration,” *IEEE Internet of Things Journal*, vol. 7, no. 6, pp. 5267–5275, 2020.
11. T. Meuser, L. Baumgärtner and P. Lieser, “Pandemic skylines: Digital twins for more realism in epidemic simulations,” in *ISCRAM 2021 Conference Proceedings*, pp. 133–145, May 2021.
12. G. Solmaz, E. L. Berz, M. F. Dolatabadi, S. Aytaç, J. Fürst, B. Cheng and J. den Ouden, “Learn from iot: pedestrian detection and intention prediction for autonomous driving,” in *Proc. of ACM SMAS’19*, pp. 27–32, 2019.
13. ETSI ISG CIM, “ETSI GS CIM 006 Context Information Management (CIM); Information Model,” 2019.
14. ETSI ISG CIM, “ETSI GS CIM 009 Context Information Management (CIM); NGSI-LD API,” 2021.
15. M. Bauer, E. Kovacs, A. Schülke, N. Ito, C. Criminisi, L.-W. Goix and M. Valla, “The context api in the oma next generation service interface,” in *2010 14th International Conference on Intelligence in Next Generation Networks*, pp. 1–5, IEEE, 2010.
16. J. Byun and D. Kim, “Oliot epcis: New epc information service and challenges towards the internet of things,” in *2015 IEEE International Conference on RFID (RFID)*, pp. 70–77, IEEE, 2015.
17. S. Wahle, T. Magedanz and F. Schulze, “The openm2t framework—m2m solutions for smart cities and the internet of things,” in *Proc. of IEEE WoWMoM’12*, pp. 1–3, IEEE, 2012.
18. R. Limosani, A. Manzi, L. Fiorini, P. Dario and F. Cavallo, “Connecting ros and fiware: Concepts and tutorial,” in *Robot Operating System (ROS)*, pp. 449–475, Springer, 2019.
19. L. L. Fernández, M. P. Díaz, R. B. Mejías, F. J. López and J. A. Santos, “Kurento: a media server technology for convergent www/mobile real-time multimedia communications supporting webrtc,” in *Proc. of IEEE WoWMoM’13*, pp. 1–6, IEEE, 2013.
20. B. Cheng, G. Solmaz, F. Cirillo, E. Kovacs, K. Terasawa and A. Kitazawa, “Fogflow: Easy programming of iot services over

cloud and edges for smart cities," *IEEE Internet of Things Journal*, vol. 5, no. 2, pp. 696–707, 2018.

21. B. Cheng, J. Fuerst, G. Solmaz and T. Sanada, "Fog function: Serverless fog computing for data intensive iot services," in *2019 IEEE International Conference on Services Computing (SCC)*, pp. 28–35, 2019.
22. F. Cirillo, B. Cheng, R. Porcellana, M. Russo, G. Solmaz, H. Sakamoto and S. P. Romano, "Intentkeeper: Intent-oriented data usage control for federated data analytics," in *2020 IEEE 45th Conference on Local Computer Networks (LCN)*, pp. 204–215, 2020.
23. J. Fürst, M. Fadel Argerich, B. Cheng and E. Kovacs, "Towards knowledge infusion for robust and transferable machine learning in IoT," *Open Journal of Internet of Things (OJIOT)*, vol. 6, no. 1, pp. 24–34, 2020.
24. A. Ratner, S. H. Bach, H. Ehrenberg, J. Fries, S. Wu and C. Ré, "Snorkel: Rapid training data creation with weak supervision," *The VLDB Journal*, vol. 29, no. 2, pp. 709–730, 2020.
25. J. Fürst, B. Cheng and B. Hebgen, "Realizing the digital twin transition for smart cities," *Open Journal of Internet Of Things (OJIOT)*, vol. 7, no. 1, pp. 32–42, 2021.

Bionotes



Martin Bauer
NEC Laboratories Europe, Heidelberg,
Germany
martin.bauer@neclab.eu

Dr. Martin Bauer is a Senior Researcher in the IoT Platform research group at the NEC Laboratories Europe (NLE) in Heidelberg, Germany. He received his MSc in Computer and Information Science from the University of Oregon, USA, in 1998, and both his Dipl. Inf. and his doctorate degree from the University of Stuttgart, Germany, in 2000 and 2007 respectively. His research interests are in the area of IoT, context information management and semantic interoperability.



Flavio Cirillo
NEC Laboratories Europe, Heidelberg,
Germany
flavio.cirillo@neclab.eu

Dr. Flavio Cirillo is a Senior Researcher in the IoT Platform research group at the NEC Laboratories Europe (NLE) in Heidelberg, Germany. His research topics include AI/ML applied to IoT, and IoT analytics platforms focusing on cloud-edge, federation, privacy, data usage control aspects, mainly in the scenario of Smart Cities. He has worked on many IoT related European research projects. He has ob-

tained a PhD in Information Technology and Electric Engineering at the University Federico II of Naples (Italy).



Jonathan Fürst
NEC Laboratories Europe, Heidelberg,
Germany
jonathan.fuerst@neclab.eu

Jonathan Fürst is a Senior Researcher in the IoT Platform research group at NEC Laboratories Europe (NLE) in Heidelberg, Germany. Previously he was a PostDoc at IT University of Copenhagen (ITU) in Denmark. He holds a Ph. D. and a M. Sc. from ITU and a B. Eng. in Industrial Engineering at University of Aalen. During his Ph. D., he spent six months as a visiting researcher at UC Berkeley in AMPLab (Software Defined Buildings group). His interests are in the area of IoT systems and practical applications such as smart buildings & cities, indoor localization and augmented reality. His current research focuses on automated data integration and knowledge extraction using machine learning.



Gürkan Solmaz
NEC Laboratories Europe, Heidelberg,
Germany
gurkan.solmaz@neclab.eu

Dr. Gürkan Solmaz is a Senior Researcher in the IoT Platform research group at NEC Laboratories Europe (NLE) in Heidelberg, Germany. His research interests include AI/ML, mobile computing, and cloud-edge systems for novel internet of things (IoT) applications to smart cities and smart mobility. He received his BS degree in Computer Engineering from Middle East Technical University (METU) in Turkey and his MS and Ph. D. degrees in Computer Science from the University of Central Florida (UCF) in the USA. At NEC, he has contributed to EU research projects in areas such as IoT, cloud/edge computing, localization, and automated driving.



Ernö Kovacs
NEC Laboratories Europe, Heidelberg,
Germany
ernoe.kovacs@neclab.eu

Dr. Ernö Kovacs received his PhD from the University of Stuttgart. He worked in ICT research various positions (at IBM from 1986–90, at

Sony from 1997–2004) on topics like multimedia e-mail, multimedia documents, distributed hypermedia systems, mobile networks, and mobile services. In 2005 he joined NEC’s “Networking Laboratories” as Senior Manager for the IoT Platform research group. His group works on Internet of Things, cloud/edge computing, weakly supervised machine learning, ontology matching, digital twins and context management.