Designing a Future-Proof Reference Architecture for Network Digital Twins

Roberto Verdecchia[0000−0001−9206−6637], Leonardo Scommegna[0000−0002−7293−0210], Enrico Vicario[0000−0002−4983−4386], and Tommaso Pecorella[0000−0002−0009−8154]

Department of Information Engineering, University of Florence, Italy
roberto.verdecchia@unifi.it, leonardo.scommegna@unifi.it, enrico.vicario@unifi.it, tommaso.pecorella@unifi.it

Abstract. As the complexity, distribution, and heterogeneity of networks continue to grow, how to architect and monitor of these networking environments is becoming an increasingly critical open issue. Digital twins, which can replicate the structure and behavior of a physical network, are seen as potential solution to address the problem. While reference architectures for digital twins exist in other fields, a comprehensive reference architecture for the networking context has yet to be developed. This paper discusses the need for such a reference architecture and outlines the key elements necessary for its design. We present the findings of a preliminary survey that explores the need for a network digital twin reference architecture, the crucial information it should include, and practical insights into its design. The survey results confirm that existing standards are inadequate for modeling network digital twins, outlining the necessity of a new reference architecture.

We then articulate our position on the need for a reference architecture for network digital twins, focusing on three main aspects, namely: (i) digital twins *of what*, (ii) *for what*, and (iii) *how to deploy them*.

We then proceed to delineate the fundamental obstacles that a reference architecture must confront, in tandem with the essential characteristics it needs to embody to successfully navigate these challenges.

As conclusion, we present our vision for the reference architecture and outline the main research steps we plan to take to address this open problem. Our ultimate goal is to tightly collaborate both with the networking and digital twin software architecture communities to jointly establish a sound network digital twin architecture of the future.

Keywords: Reference architecture · Digital twin · Networking.

1 Introduction

In recent years, Digital Twins (DTs) gained and increasing popularity, and year after year are becoming more adopted in different and new industrial contexts. A digital twin is a virtual representation of a system, facilitating bidirectional communication between the system and its digital representation [18]. Such virtual representation is used, among other goals, for designing, modeling, and monitoring physical asses [11]. DTs enable to mimic the structure, context, and behavior of a single or groups of physical assets, supporting both design and runtime decision making processes of the physical counterparts. By collecting and analyzing data from multiple sources, DTs can be used to digitally gain information on various attributes, such as performance and related inefficiencies, to identify and design solutions to improve their physical counterparts.

In networking environments, DTs are commonly used to represent physical networking assets such as routers, switches, controllers, and communication channels [32]. Network DTs (NDTs) usually include information regarding operational status, performance data, and environmental conditions of their physical twins. By exchanging network data and control messages with a network of DT instances through twin-physical interfaces, network engineers can rely on DT representations to design, test, assess security, and improve the maintenance of physical networks. This allows for efficient and intelligent management of networks, with the ultimate goal of supporting the improvement of network performance, reliability, and accelerate network innovation.

The concept of DT has been largely developed in the context of Cyber Physical Systems, much promoted by the agenda of Industry 4.0 where it was also addressed and formalized in standardization initiatives [15]. In the context of future generation networks, the growing level of softwarization demands architectural paradigms that can drive the organization of functional responsibilities, their connection with data collection and intelligent processing, and their deployment and composition across network computing, storage, and connectivity resources.

While various concepts can inherit results consolidated in contexts where the DT paradigm has already reached higher maturity and readiness, application in software driven networks raises several and hard new challenges, notably including: distribution across a large-scale network, with sustainable footprint on communication and storage resources; critical need for high levels of interoperability among heterogenous resources and services managed by multiple operators; autonomic orchestration capability supporting efficient and self-adaptive placement of network functions and applications across edge-to-cloud levels and localities. The relevance of these challenges, and their scientific and technological perception is clearly testified by the level of standardization initiatives and the growing number of scientific works.

In order to architect NDTs, comprehensively model their characteristics, and manage the high complexity such systems entail, a reference architecture, *i.e.,* a template solution for an architecture of a particular domain [3], could be used. Such solution was recently introduced for manufacturing environments, with the establishment of the ISO Standard 23247 [15] (which was also picked up in recent software architecture literature [10]). To the best of our knowledge, a reference architecture to model NDTs, covering both functional and non-functional aspects of NDT architectures, is still missing in the current body of knowledge.

As note, while reference architectures and standards might have similar properties, they convey different concepts. Specifically, a reference architecture is a template architectural solution for a particular domain and context. Software engineering standards instead are a set of guidelines for the process, quality, and documentation of software development and maintenance, usually developed by industry organizations or governing bodies, *e.g.,* IEEE and ISO. Therefore while a standard can document a reference architecture (*e.g.,* in the case of the ISO Standard 23247 [15]), the opposite is not always true. Moreover, a reference architecture can leverage standards (*e.g.,* data interchange standards), but its is often independent from the particular standard being used.

The NDTs architectures are foreseen to play a critical role in the RESTART Foundation (RESearch and innovation on future Telecommunications systems and networks, to make Italy more smART)[1], funded by the European Union (EU), under the Next Generation EU (NGEU) program.[2] The RESTART Foundation is a partnership between 25 Italian universities (*e.g.,* the Sant'Anna School and the University of Rome La Sapienza.), research centers (*e.g.,* the Italian National Research Council), and companies (*e.g.,* Vodafone and Ericsson). The goal of the RESTART project is to leverage DTs to provide a structural improvement of telecommunications research and development in Italy, supporting the digital transformation of industries, and growth of related research and professional communities. Within the RESTART Foundation, the COHERENT project *"Shaping a Digital Twins future proof network architecture"* focuses explicitly on integrating the outcomes of all RESTART research activities in a comprehensive network architecture considering both a technical and a business point of view. The research project, founded for a total of 116 million euros, aims to fill a current gap in networking, namely the lack of an extensible and evolvable NDT reference architecture. Current standards and documentation related to a NDTs reference architecture result to either be too generic to effortlessly incorporate the specifics of the networking domain, *e.g.,* consider the DT framework of Josifovska *et al.* [16] or the DT archetypes of van der Valk *et al.* [30], or result to be deeply grounded in current technologies, and are therefore inherently hard to evolve according to future emerging technologies. As documented by the funding body, realizing a future proof DT network architecture and documenting its related design rationale allows to establish a set of best-practices to fully harness the potential of the implementation of projects in the networking domain.

As part of COHERENT, in this position paper we outline how, in order to comprehensively consider and integrate the various facets of NDTs, a future proof *reference architecture for network digital twins* needs to be established.

The contributions of the paper are (i) an opening survey empirically investigating the need of a reference architecture for NDTs, (ii) a grounding problem statement outlining the need of such reference architecture, and (iii) our vision on a future proof reference architecture of NDTs.

This research builds upon the initial position paper presented at the second International Workshop on Digital Twin Architecture (Twin-Arch) [31] by (i) discussing the main challenges of designing a network digital twin reference architecture, (ii) reporting the key features the reference architecture must provide, and (iii) providing a stepping stone towards the concrete implementation of the architecture.

2 Opening Survey

In order to gain introductory empirical insights into the need for a NDT reference architecture, independent of the statements and goal set by the RESEARCH funding

[1] `https://www.fondazione-restart.it/`. Accessed 18th June 2023.

[2] `https://next-generation-eu.europa.eu/index_en`. Accessed 2 August 2023.

body (see Section 1), we conducted a survey involving researchers and practitioners working in the field of networking. Participants were recruited *via* convenience sampling starting from the RESTART Foundation participant list and the personal network of the authors, followed by a subsequent snowballing sampling. Survey invitation target networking experts, belonging either to academic entities, renowned large scale industrial companies, or networking standardization entities. Under the human ethics guidelines governing this study, we cannot disclose affiliations of participants to preserve their anonymity.

In total, 16 participants took part in the survey.

The survey comprised a mix of close-ended 5-point Likert scale questions (CE) and free form open-ended questions (OE). Each CE was accompanied by a OE, where respondents could further clarify their answer.[3] The survey was composed of three main parts, namely:

1. Participant demographic questions: Current job position (OE), years of experience (OE), familiarity with networking and digital twins (CE);
2. On need of a NDT reference architecture: Degree to which the ISO 23247 can be used to represent NDTs (CE), degree to which the ISO 23247 needs to be modified to represent NDTs (CE), and perceived usefulness of a NDT reference architecture (CE);
3. Further advice to establish a NDT reference architecture: expected networking components modeled (OE), expected grouping of networking components (OE), degree to which elements of standardisation groups (*e.g.,* ETSI or IETF) should appear in the NDT reference architecture (CE).

To ensure respondents have enough knowledge on DT to answer the survey, a definition of DT is provided at the beginning of the survey. Similarly, an overview of the ISO 23247 standard provided by Bucaioni *et al.* [9] is provided in the survey. Participants who acknowledge not being familiar with networking and/or DT concepts are discarded from the respondents.

From the demographic answers, the vast majority of participants resulted to work in academia (11/16), possess an average of 10 years of experience, be highly familiar with networking concepts, and moderately familiar with DT.

Regarding the ISO 23247, most participants noted that it can be applied to networking concepts only to a moderate extent (6/16) or low extent (5/16). From the supporting OE answers, we note that this is primarily due to a perceived lack of generalizability of the ISO 23247 standard. By considering the extent to which the ISO 23247 standard needs to be modified in order to be used for NDTs, respondents primarily indicated a medium, or medium-high degree (13/16). Accompanying OE questions clarified that this is mostly due to the need to model concepts specific to NDTs, *e.g.,* details regarding network virtualization functions, and other networking-related attributes, which require new abstraction levels. All participants agreed on a medium-high, or high usefulness of a NDT reference architecture (15/16).

[3] To support replicability and scrutiny, the survey and received answers are made available online at: `https://github.com/STLab-UniFI/twinarch-2023-reference-architecture-rep-pkg`

When considering the further advice provided by participants to establish a NDT reference architecture, respondents mostly indicated basic hardware networking components, *e.g.,* routers, switches, and hubs (8/16). In contrast, virtual elements, *e.g.,* virtual machines, VPNs, and firewalls, were mentioned far less frequently (3/16). Only seldom, communication-related elements, *e.g.,* physical channels, were mentioned (3/16). Only few respondents described the expected grouping of networking components, providing heterogeneous answers, *e.g.,* "physical layer; security; services; hardware; software; protocols" and "SDN control plane; 5G-oriented data plane". Finally, respondents indicated that networking elements of presented by standardisation groups (*e.g.,* the the European Telecommunications Standards Institute (ETSI).[4]) could be used between a medium and medium-high extent to model a NDT reference architecture (15/16).

Overall, as main takeaways of the opening survey conducted for this position paper, we can conclude that, based on the opinion of mostly academic researchers experienced in networking:

1. The ISO 23247 does not fit completely the networking context, and would need to be considerably modified;
2. A NDT reference architecture is perceived as highly useful;
3. Elements to be covered in the NDT reference architecture should primarily focus on hardware networking components, could use to a moderate extent elements of existing standards.

3 On the Need of a Reference Architecture for Network Digital Twins

Albeit extensive literature considered network DT [1,19,26,27,32], the topic has been primarily addressed from a purely networking point of view. As such, aspects related to a reference architecture NDTs, *i.e.,* a reusable metamodel that can applied to heterogeneous contexts, considering disciplines such as software engineering and software architecture, seem to have been almost completely neglected in current literature [5]. To address this point, in this position paper, we take a software engineering stance by reviewing the topic of NDTs reference architectures through the lens of software architects.

As emerges from recent reviews [1,11,32], when considering NDTs, three main aspects can be taken into account, namely *NDTs for what?*, *NDTs of what?*, and *how to deploy NDTs?*. In the following, we detail our position on these three aspects, building towards our vision on the main proprieties an NDT reference architecture needs to possess.

3.1 Network digital twins *for what?*

As one of the most consolidated aspects of NDTs, the related body of literature extensively describes the different application scenarios of NDT, *e.g.,* their use for

[4] `https://www.etsi.org`. Accessed 18th June 2023.

network function virtualization, controlled orchestration, and reliability/security monitoring and assurance processes. For example, NDTs can be used to facilitate service placement, allowing for the efficient streaming of data from one point to another within a network [4].

Reference architectures for NDTs are available (*e.g.,* the NDTs architecture presented by the Telecommunication Standardization Sector of the International Telecommunication Union [28]). Nevertheless, such reference architectures considered primarily, if not exclusively, the functional nature of NDTs, *i.e.,* do not consider aspects related to the characteristics of the entities that have to be represented, or their concrete use / deployment (see also following sections).

Similarly, standards regarding functional aspects of DTs are widespread knowledge within the industry, as documented for example by the industry-driven effort in the Internet Research Task Force (IRTF) [35], as well as the evolution of standards relative to the network devices management planes (see for example the standards issued by the IETF NETCONF Working Group[5])

Overall, it appears as if the "*NDTs of what?*" field is a quite consolidated in the networking community. For example, the field of network function virtualization experienced a growing interest through the years, and can now be regarded as a mature, consolidated, and standardized area [34].

As more recent example of NDTs functional viewpoints, current research investigates the use of NDTs for AI model lifecycle management [17]. This topic, currently under investigation, opens for new challenges of functional NDT aspects, *e.g.,* controlling responsibilities, management of AI model lifecycle within NDTs, and consistency between models distributed *via* federated learning.

3.2 Network digital twins *of what?*

As less explored area, we note that often the literature on NDTs does not appear to predicate in detail and precision on the specific network elements that are required to be modeled in the NDT context. As a matter of fact, frequently the nature of network components which need to be modeled within NDT architectures seem to be reported at a rather high level, with auxiliary elements left implicit, or not regarded at all. This more often than not seem to cause the unsystematic documentation of incomplete or vague NDT reference architectures, that, due to their abstract and at times speculative nature, cannot be ported into practice without making considerable assumptions.

Even in the rare cases in which the most important elements of NDTs architectures are explicitly documented, their description often lacks basic details regarding property characteristics and attributes NDTs must posses. Therefore, theoretical or even simulation results are hardly portable into practice by implementing a concrete NDT architecture. In fact, the development process would imply a considerable upfront conceptual effort, which would require *per se* an independent study and verification prior the concrete development can take place.

As a possible solution to address this issue, the information to model NDTs could be derived from standard network architecture documentation, *e.g.,* the documen-

[5] `https://datatracker.ietf.org/wg/netconf`. Accessed 18th June 2023.

tation provided by ETSI. Similarly, the necessary information could be identified by porting the modeling information of network simulators (*e.g.,* `ns-3`[6] and `OMNET++`[7]) to a NDT reference architecture, documenting *via* a metamodel the NDT elements, their attributes, and relations.

3.3 *How to deploy* network digital twins?

Another area that appears to be only marginally considered from a practical standpoint is the concrete deployment of NDTs over a network.

As for DT in general, one of the challenges in the use of NDTs within a network is the distributed deployment of these virtual representations. To date, standards do not appear to provide a clear guidance on how NDTs should be deployed, distributed, and relocated. From an architectural standpoint, one approach could be to consider network elements, such as Media Access Control (MAC) addresses, as monolithic entities. Nevertheless, given the growing functional complexity of NDTs, this approach might be considered as too simplistic. As alternative, network elements could be factored into bounded contexts. This strategy would lead to the production of microservices, allowing, albeit their potential complexity increase, to take advantage of the benefits of the microservice architecture style, *e.g.,* fault tolerance and fault isolation.

By considering the adoption of a microservice architecture in the context of NDTs however, there is a special emphasis on enabling deployment and placement at different levels of the edge-to-cloud continuum at different localities.

As a double-edged sword, on one hand NDTs are responsible for resolving placement problems through their state and associated computational power (or by delegating the task at hand to other NDTs to obtain states and/or delegate the processing). However, NDTs also rise novel issues associated to how to place these responsibilities on physical and virtual resources within a network. Therefore, while DTs can resolve placement problems, they also open up new challenges in terms of the placement of NDT themselves. The challenges associated to the deployment of NDTs must be carefully managed, in order to optimize the performance of DTs within a network environment. To date, this problem appears to be marginally addressed in the literature, lacking to provide concrete guidance and reference on how NDTs should be deployed.

4 Main Challenges for a Network Digital Twin Reference Architecture

In this section, we identify the key challenges a network digital twin architecture must be capable of addressing to be considered future-proof. The ensuing list is derived from discussions and a workshop event with sector-leading experts involved in the COHERENT project, focusing on the examination of proposed standards on the topic [28,35].

[6] `https://www.nsnam.org`. Accessed 18th June 2023.
[7] `https://omnetpp.org`. Accessed 18th June 2023.

4.1 Large Scale Data Collection

Considering the inherent characteristics of the network and the potential of digital twin networks, we posit that a digital twin network architecture should be equipped to manage substantial volumes of data efficiently. It is expected that the architecture will be capable of collecting and managing data of various natures originating from multiple sources.

- Recordings and event logs from all elements of the network;
- Statistics-related data like traffic throughput, latency, and packet loss;
- Data related to service usage and users;
- Monitoring data of observable entities;
- Operational and provisional data;
- Simulation and emulation results.

This information generates a continuous flow of large data sets. The digital twin should use this data to represent the current state of the entity. Additionally, it's crucial to preserve this data to maintain a historical record of the information The inherent heterogeneity of data, which encompasses both variety and volume, presents a significant challenge in its effective management.

The exponential growth of data from mobile devices and IoT applications will make this problem a central concern in network management. One of the challenges that Network Digital Twin architectures will face is managing massive network data collection from network infrastructures.

4.2 Scalability

Given the expected surge in network components and participants (*e.g.,* more sensors, clients, and applications), a Network Digital Twin architecture must effectively accommodate this growth and ensure consistent performance. As the number of network elements increases, it is crucial for a network digital twin architecture to effectively handle virtual representations of real networks, regardless of their scale. Furthermore, as network size gradually expands, the features offered by the architecture must remain efficient and effective. Therefore, features such as data collection from the network, reconfiguration, and simulations should always be available, ensuring consistent functionality irrespective of the network's size, the number of data sources, or the number of network applications. In addition to functionalities, the performance provided by the architecture should be scalable. For example, latency, the accuracy of prediction, and simulation algorithms should maintain the expected performance without depending on the size of the network that the digital twins compose.

4.3 Flexibility and Autonomous Reconfiguration

It is expected that the network will maintain a dynamic behavior, with network components, applications, and clients evolving over time and the load fluctuating in intensity. For this reason, a digital twin network should be able to cope with the network variation. Therefore, it is expected to possess the ability to execute on-demand behaviors and reconfigure itself (possibly automatically without the intervention of human operators) in response to events while maintaining Quality of Service constraints. These behaviors are moving towards the implementation of networks that are often identified with the term Zero Touch Network [8], which are envisioned to be highly autonomous networks capable of self-configuration, self-healing, and self-optimization with minimal to no human intervention. It is also expected that many of the features and constraints can be specified dynamically, thus allowing for greater system availability without taking the system online. Last but not least, given the heterogeneity of the network, the architecture must be able to provide adequate adaptability to new elements, provide functionality to new network applications, and collect and store data of various nature and format.

4.4 Heterogeneous Performance Requirements

In the near future, the network will be used to support many challenging applications. Some of these will require near real-time response requirements, *e.g.,* in the context of autonomous driving, virtual reality, gaming, and healthcare. Other applications will instead necessitate elevated levels of parallelism or data storage to execute algorithms with efficiency. Such applications typically encompass simulations, algorithms pertaining to artificial intelligence and machine learning, and, in recent developments, practices of federated learning.

The complex requirements required by network applications, sometimes in contrast with each other, will have to coexist and be natively supported in an architecture that aims to be future-proof.

4.5 Interfaces

In the context of a network digital twin architecture, we deem the selection of interfaces as a fundamental. A Network Digital Twin requires an interface to interact with the physical network. This interface is commonly called *south bound* and is responsible for establishing the ways in which digital twins exchange information with the corresponding physical entities. To ensure the architecture is accessible and open to extension, we posit that it is also crucial to establish a standard interface through which the architecture can expose its services and features externally. Beyond digital representation an control of entities, digital twins offer the significant ability to gather information and conduct analysis, emulations, and simulations based on real-world data. To execute these types of functionalities and ensure the scalability of the network architecture capabilities over time, it is crucial in the architectural design process to identify the methods through which these functions will be invoked. In this sense, we consider the definition of internal interfaces (sometimes referred to as *side bound*) to be equally fundamental to the definition of the architecture.

4.6 Digital Twin Security

Even though NDTs can be used to improve the network security of a system, *e.g.,* by analyzing and quickly applying changes related to adverse events, they have to be also resistant to attacks targeted to the NDT itself.

As a matter of fact, NDTs can be seen as a particular case of a Cyber-Physical system, and their security depends not only on the security of the NDT itself (the software component), but also on the capability to have a 'useful' knowledge of the physical counterpart, and to control it. It is fairly evident that the attack surface for NDTs is larger and more complex than the one of a traditional network, and even larger than the one of a Software-Defined Network (SDN).

In our opinion the technologies needed to address the security requirements should not be part of the NDT architecture. However, the security analysis should be part of the architecture, at least to highlight the attack surfaces and the possible threats.

For what concerns the threat agents and their capabilities, these can be considered as "normal" threat actors targeting traditional networks, as the goals and means to perform an attack are the same. The normal approach used by IETF is to define the possible attacks based on the threat actor capabilities (see [24]), *i.e.,* splitting the attacks into passive or active.

Passive attacks to NDTs are not to be underestimated, as an attacker might gather a very precise (and timely) knowledge of the network status. As a matter of fact, intercepting the data collection traffic between the physical and digital parts, an attacker can not only understand the network operational status, and perform targeted (active) attacks, but also understand the existence of particular network statuses, and correlate them with user data. Again, this can be useful to perform further actions.

Hence, the loss of confidentiality in the data collection can create serious risks, not only to the NDT itself, but also to the user data.

Active attacks are peculiar as well, as they can be targeted to the NDT software components, to the data collection mechanism, or to the network configuration elements, used by the NDT to modify the actual network setup.

The NDT components and the network elements configuration are well understood, are almost universally considered as sensitive elements, and secure-by-design architectures and network configuration protocols are generally available. On the contrary the data collection has only recently evaluated, and several protocols have been, or are in the process of being updated to add confidentiality and integrity features (see for example [7]).

Hence, we believe that a NDT architecture should contain the threat models and possible attacks to the NDT, in order to guide the implementations to use the proper security models, both in the components, and in the protocols used by the NDT.

5 Key Features of a Network Digital Twin Architecture

After identifying the challenges that the design of a network digital twin architecture will have to face (Section 4), we now provide the key features that, in our vision, are necessary for the architecture to adequately address these challenges.

5.1 Edge to Cloud Continuum Deployment Awareness

The increasingly pervasive spread of edge devices capable of collecting and computing data will necessitate a paradigm shift in the use of network infrastructures. The emergence of Mobile Edge Computing (MEC) represents a transformative paradigm shift [13], which is anticipated to rapidly gain prominence. This advancement is poised to facilitate the execution of algorithms and the provisioning of services in closer proximity to the end-user. MEC will enable a reduction in latency that is unthinkable in architectures that rely solely on cloud computing. Thus, MEC will enable all those applications that require near real-time requirements. In addition, by processing requests locally, the cloud and the backhaul network [25], *i.e.,* the network that connects the edge to the cloud, will be relieved of a large number of requests.

Although mobile edge computing represents a turning point for many future applications, its use introduces complexities that must be considered. Indeed, an MEC node is characterized by a limited amount of resources, so it is not possible to deploy all functionalities on these nodes. Moreover, some particularly resource-demanding tasks, such as machine learning algorithms or simulations, are better suited to execution on the cloud where there is no problem of resource scarcity.

Another intrinsic complexity in the MEC paradigm and more generally in edge computing, is the local nature of data and services in an environment characterized by clients who typically move in space and vary over time [25]. This implies the implementation of strategies such as service placement, handover, and service offloading, should be native in a network digital twin architecture [23].

It is necessary also to consider that the duality between edge and cloud cannot be defined in a clear-cut fashion and that there is continuum between the two entities in which intermediate nodes are able to provide functionalities and data storage halfway between the two extremes. This is known as the Edge-to-Cloud continuum and is also closely related to the concept of Fog Computing [29].

Ultimately, we believe that a reference architecture for network digital twins must necessarily take into account such aspects of deployment and resource management natively. In our view, this is necessary to best manage challenges such as *large-scale data collection*, *scalability*, and *heterogeneous performance requirements*.

5.2 Digital Twin Interoperability

In our opinion, one aspect that is not adequately highlighted in the presented standards (ITU Y.3090 [28] and IRTF [35]) and which we deem central to a network digital twin architecture is the interoperability that should exist between network digital twins.

In our view, the architecture does not manage a single network digital twin but an entire ecosystem. This ecosystem, like the actual network, is not static but changes over time in terms of both elements and functionalities, and collaborates to achieve the objectives required by the network applications it uses.

The definition of an ecosystem of digital twins that collaborate with each other enables the possibility to dynamically extend the digital representation of the network. A change in the physical network does not necessitate to change the entire digital twin, but rather the addition, removal, or update of a single element. Such change does not affect the entire representation of the network but only a single part. Furthermore, this modularity allows for a more granular control of the elements.

An ecosystem of digital twin allows for extremely *flexible* management of the network with the possibility of implementing policies of distributed *auto-reconfiguration*, management of *heterogeneous requirements*, especially if combined with a deployment awareness as described in Section 5.1, and even *security*.

5.3 Distributed Network Digital Twin

A Network Digital Twin is anticipated to deliver a diverse array of functionalities, reflective not only of the spectrum of performance requisites but also of their inherent complexity. Considering the imperative for deployment awareness (see Section 5.1) in conjunction with the diversity of functionalities each NDT furnishes, we believe that the network digital twin ought to exhibit a distributed rather than a monolithic structure.

The decentralization of the NDT is poised to facilitate performance optimization: for instance, functionalities with low latency requirements could be deployed at the edge, whereas those demanding substantial computational resources might be more aptly positioned within the cloud infrastructure. This approach is also likely to enhance network utilization, preventing the potential for data flow congestion within both the network and the cloud.

We advocate for the adoption of a distributed model for Network Digital Twins, as it promises augmented *flexibility* within the network. This model enables the dynamic reallocation of digital twin components to nodes that are optimally suited, through the implementation of strategic placement and offloading. A distributed model also permits the precise management of the *heterogeneous performance requirements* that characterize digital twins. Ultimately, this strategic framework is expected to support extensive *data collection* and *scalability*, while simultaneously reinforcing security measures, as sensitive data and functionalities can be securely housed within more secure nodes.

5.4 Composite (Hierarchical) Digital Twin

The concept of network is often associated with a composition of other sub-networks (*e.g.,*, the internet is a network of networks). We therefore naturally consider the digital twin network to be characterized by a composite and hierarchical nature. This concept, moreover, is not new to digital twins in the field of manufacturing and also appears in the seminal paper by Grieves *et al.* [12] under the term "*Digital Twin Aggregate*". A hierarchical representation of digital twin networks enables the representation of the network at different levels of granularity, allowing for different views and functionalities depending on the type of granularity required.

At the lowest rung of the hierarchy is the digital representation of the atomic elements of the network (*e.g.,*, routers and switches). Digital twin networks with this level of granularity will therefore provide information on the status of network elements and expose functionalities aimed at their configuration. However, there are situations where it is necessary to interact with a large network like Metropolitan Area Networks (MANs) or Wide Area Networks (WANs). In such cases, it is plausible that the pertinent information extends beyond the scope of singular network components and pertains instead to an aggregated construct. Similarly, the type of actions to be performed on the network, reasonably, will not concern the individual network element but the network as a whole. In such a case, the Digital Twin Network will represent a network with aggregated information derived from Digital Twin Networks of lower granularity levels and will provide high-level functionalities that could cascade into the NDTs that compose it.

This type of representation allows bringing to the different levels of abstraction only the necessary aggregated information, ignoring data and functionalities that are too fine or too coarse. This type of hierarchical representation of digital twin networks is beneficial for many reasons and helps to address various challenges identified in this paper. The use of data of the right granularity for the level of abstraction of the digital twin network allows for *efficient data collection* and excellent *scalability*. These advantages are further enhanced when combined with deployment awareness (Section 5.1). Moreover, a hierarchical strategy allows for a high level of network *flexibility*. Finally, the ability to represent aggregated data combined with the presentation of high-level functionalities will greatly simplify the *interfaces* to be presented to a client (North Bound).

5.5 Prototyping

One of the main features of digital twins is to execute simulations based on real data collected in the field. Within a network context, this enables *what-if* analysis simulations and predictions, wherein hypothetical scenarios are studied based on present data. In such instances, it becomes necessary to have prototypes of network elements readily available. That is, having at disposal virtualized representations of network elements not directly associated with a physical element. A prototype is indeed a stereotypical representation of a specific type of network element and as such encapsulates default functionalities and statistical data common to the category of element it represents. The concept of a prototype is widely used in the industrial and manufacturing context, enabling procedures for product lifecycle management [12].

A prototype is a stereotypical representation of a specific type of network element and, as such, encapsulates default functionalities and statistical data common to the category of element it represents, which are based on real-world observations collected in the field. The prototype thus permits an optimized and aggregated use of data gathered during the monitoring that NDTs continuously perform.

We therefore posit that a reference architecture for network digital twins should take into account this type of representation of components. Especially, we contend that the employment and management of prototypes are fundamental, particularly for addressing challenges such as *efficient data collection*, *flexibility and autonomous reconfiguration* capabilities, and *security* through the possibility of conducting appropriate simulations.

5.6 Digital Twin of Anything

Nowadays, many network elements are digitized and programmable. The primary elements include software defined networks, virtualized network functions, and network slices [22]. However, even elements such as virtual machines and containers can be considered as integral parts of the network and contribute to its operation. We therefore believe that for a satisfactory architectural representation of the network,

network digital twin should not be limited to representing physical network elements, but also those elements that inherently possess a digital nature.

Even the existing standards acknowledge this possibility [35] but without identifying its main advantages. Below we list the reasons why we consider it useful to establish a digital twin for all digital elements of the network digital twins reference architecture. The digital twin of a digital component can act as a wrapper for the component and enrich it with additional functionalities that are not natively supported by the original component. For example, a component could collect information from monitoring and perform simulations or prediction algorithms. Through the digital twin, better management of interfaces is allowed, thus improving the management by the user of the element (north bound) and the management of interoperability between different digital twins (side bound). The establishment of an interface relieves the other architectural components of the responsibility of having to know all possible elements and their specific interfaces in advance, greatly increasing the flexibility and maintainability of the network.

5.7 An Intent-Based Architecture

A network must be capable of fulfilling various requirements specified by clients and applications external to the network. These requirements change over time and are often in conflict with one another. To address the evolving nature of network requirements and ensure flexibility, self-configuration, and user-friendly interface interaction, it is crucial to adopt the concept of *Intent* as the foundation for the network configuration and specification mechanism. The concept of Intent has become increasingly prominent in recent standards [8,28,33,35]. according to IRTF NMRG [14], an intent is "*a set of operational goals (that a network should meet) and outcomes (that a network is supposed to deliver) defined in a declarative manner without specifying how to achieve or implement them*". An intent-based network aligns with a user-centric perspective, simplifying the expression of requirements and enhancing the architecture's usability and *flexibility*. An intent can also be expressed in human language and then easily translated into machine-readable language. This simplicity in defining an intent dramatically increases the usability of the architecture and the definition of the North Bound *interface*. In addition to offering high simplicity at interface level, intent-based architectures allow for a high degree of *autonomy* and *flexibility* and *auto reconfiguration capabilities*.

In our vision, the adoption of intent-based networking techniques could be highly beneficial. Managing through intents provides several advantages, particularly when various applications are concurrently utilizing the network. The network will be capable of defining a space of possible configurations that satisfy all the intents specified by the applications. The exploration and identification of an optimal configuration can then be entrusted to algorithms and programmatic policies.

6 A Step Towards an Architecture Implementation

After identifying the set of key features deemed as essential to address the challenges posed by a network digital twin architecture, we now provide a perspective on how the architecture can be developed with a more implementational outlook. In doing so, we will also specify which key features, in our opinion, will be implemented in this manner.

In an architecture where service decoupling and high interoperability are required, as outlined in this article, a microservices ecosystem is certainly the most suitable [21]. Indeed, each microservice natively has the ability to be deployed on various network nodes while maintaining its independence and still being able to collaborate with other microservices. This characteristic thus ensures strong *interoperability between digital twins*. Furthermore, a network digital twin does not necessarily have to be confined within a single microservice. Various functionalities pertaining to the same digital twin could be represented by a set of microservices, thus implementing a *distributed digital twin*. Regarding prototyping, in our vision, just as for the network digital twin, we imagine a microservice or a set of microservices dedicated to individual prototypes. This makes the architecture both maintainable and extendable over time.

In a microservices architecture, it is simple to identify a set of microservices as "front-end" microservices, *i.e.*,, microservices that collect information and expose high-level functionalities. These microservices will therefore be the ones used by the client and will act as entry points for network functionalities. Microservices are particularly suited to defining flexible interfaces, and it would thus be possible to specify *Intents* through such entities as well (see Section 5.7). In networks of considerable complexity, it is conceivable that a collection of front-end microservices alone may not suffice to fulfill client intents. Under such circumstances, the front-end microservices would engage in collaboration with other microservices that operate at a more granular level of abstraction, thereby delineating a workflow of microservices: the invocation of a front-end service triggers a sequential activation of other microservice instances [2]. This cooperative mechanism thus establishes a hierarchical structure of microservices, which, in turn, facilitates the representation of *composite digital twins*.

Through specific actors and technologies commonly implied in microservices architectures (*e.g.*,, Kubernetes and its container orchestration platform [20]) it is possible to programmatically manage the scaling of individual services and also their deployment. This makes it possible to have a *deployment-aware* architecture capable of defining edge-to-cloud deployment policies and also implement dynamic service placement strategies [6].

7 Conclusions, Our Vision, and Future Work

Despite the growing adoption and complexity of network digital twins, a reference architecture for this context, which considers both functional and non-functional aspects, appears to date to be missing in the literature. From the preliminary motivating survey conducted for this position paper, we noted that (i) such reference could be highly helpful, (ii) existing standards do not totally fit the networking context, and would need to be considerably modified, and (iii) elements to be considered would be primarily of hardware nature, and could to a certain extent be modeled by leveraging existing network standards and tools.

We documented our position on the current state of the art, and what is needed to move towards a future proof reference architecture for NDTs. By considering current trends and advancements, we reasoned on the key aspects of architecting NDTs, which we formulated in terms of *NDTs for what*, *NDTs of what*, and *How to deploy NDTs*. Based on these three facets, we note that research and development endeavors primarily focused on the functional "*for what*" aspects of NDTs. As such, albeit crucial, which elements to be represent with NDT, and how / where to deploy NDTs, are aspects that are only marginally considered in the current state of the art. In an effort to crystallize the scenarios that a reference architecture for NDT must manage, we have pinpointed the most intricate challenges, subsequently identifying the essential key features we believe the architecture requires to address the identified issues. Ultimately, we have also provided our perspective with a more implementational outlook.

To move towards a standardized modeling of NDTs architectures, we posit that all three aspects, digital twins *of what*, *for what*, and how to *deploy them*, need to be considered. To do so, a reference architecture covering all three of these aspects needs to be established. Providing a standardized framework for NDTs would allow the community to move with a unified effort towards consolidated new abstractions of networking attributes, supporting the design and development of the next-generation wireless networks.

As future work, we plan to proactively build upon the position outlined in this document, by working towards the establishment of a future proof reference architecture for network digital twins. As first research step, we plan to conduct (i) a comprehensive qualitative empirical research involving network researchers and practitioners, and (ii) a systematic literature review on network digital twins. With this first step, we aim at gaining a deep and systematic understanding of the state of the art and practice of NDTs. In a second phase, we plan to design a reference architecture that comprehensively covers aspects related to NDTs *of what*, *for what*, and *how to deploy them*. Data and inspiration could be drawn from existing concrete artifacts to model networks, *e.g.,* the elements and attributes used by widespread simulation tools such as `NS3` and `OMNET++`. Finally, we plan to evaluate and refine the established NDT reference architecture in a design science fashion, by gathering feedback from researchers and practitioners in the field *via* qualitative assessments and concrete case studies.

The task of establishing a NDT reference architecture is ambitious, and requires by definition interdisciplinary knowledge coming from the areas of software architecture, digital twin modeling, and networking. For this reason, we more than welcome feedback, insights, and collaboration with researchers and practitioners of any of these areas who are interested in jointly progress towards a holistic, standardized reference architecture for NDT.

With this position paper, we aim to reach out to both the networking and digital twin software architecture research and practitioners communities, in order to jointly progress towards the end goal of the RESTART mission, namely the establishment of a future proof digital twin network architecture.

Acknowledgments

This work was partially supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on "Telecommunications of the Future" (PE0000001 - program "RESTART").

References

1. Almasan, P., Ferriol-Galmés, M., Paillisse, J., Suárez-Varela, J., Perino, D., López, D., Perales, A.A.P., Harvey, P., Ciavaglia, L., Wong, L., et al.: Network digital twin: Context, enabling technologies, and opportunities. IEEE Communications Magazine **60**(11), 22–27 (2022)
2. Alshuqayran, N., Ali, N., Evans, R.: A systematic mapping study in microservice architecture. In: IEEE Int. Conf. on SO Comp. and App. pp. 44–51. IEEE (2016)
3. Bass, L., Clements, P., Kazman, R.: Software architecture in practice. Addison-Wesley Professional (2003)
4. Bellavista, P., Giannelli, C., Mamei, M., Mendula, M., Picone, M.: Application-driven network-aware digital twin management in industrial edge environments. IEEE Transactions on Industrial Informatics **17**(11), 7791–7801 (2021)
5. Dalibor, M., Jansen, N., Rumpe, B., Schmalzing, D., Wachtmeister, L., Wimmer, M., Wortmann, A.: A cross-domain systematic mapping study on software engineering for digital twins. Journal of Systems and Software p. 111361 (2022)
6. Detti, A.: Microservices from cloud to edge: an analytical discussion on risks, opportunities and enablers. IEEE Access (2023)
7. Elkins, N., Ackermann, M., Deshpande, A., Pecorella, T., Rashid, A.: IPv6 Performance and Diagnostic Metrics Version 2 (PDMv2) Destination Option. Internet-Draft draft-ietf-ippm-encrypted-pdmv2-05, Internet Engineering Task Force (Oct 2023), `https://datatracker.ietf.org/doc/draft-ietf-ippm-encrypted-pdmv2/05/`, work in Progress
8. Zero-touch network and Service Management (ZSM); Intent-driven autonomous networks; Generic aspects. Standard, European Telecommunications Standards Institute (Feb 2023)
9. Ferko, E., Bucaioni, A., Behnam, M.: Architecting digital twins. IEEE Access **10**, 50335–50350 (2022)
10. Ferko, E., Bucaioni, A., Pelliccione, P., Behnam, M.: Standardisation in digital twin architectures in manufacturing. In: 2023 IEEE 20th International Conference on Software Architecture (ICSA). pp. 70–81. IEEE (2023)
11. Fuller, A., Fan, Z., Day, C., Barlow, C.: Digital twin: Enabling technologies, challenges and open research. IEEE access **8**, 108952–108971 (2020)
12. Grieves, M., Vickers, J.: Origins of the digital twin concept. Florida Institute of Technology **8**, 3–20 (2016)
13. Haibeh, L.A., Yagoub, M.C., Jarray, A.: A survey on mobile edge computing infrastructure: Design, resource management, and optimization approaches. IEEE Access **10**, 27591–27610 (2022)
14. Intent-Based Networking - Concepts and Definitions. Standard, Internet Research Task Force (Dec 2022)
15. ISO/IEC/IEEE: Automation systems and integration — digital twin framework for manufacturing — part 2: Reference architecture. ISO/IEC/IEEE ISO 23247-2:2021 pp. 1 –9 (10 2021)

16. Josifovska, K., Yigitbas, E., Engels, G.: Reference framework for digital twins within cyber-physical systems. In: 2019 IEEE/ACM 5th International Workshop on Software Engineering for Smart Cyber-Physical Systems (SEsCPS). pp. 25–31. IEEE (2019)
17. Kaur, M.J., Mishra, V.P., Maheshwari, P.: The convergence of digital twin, iot, and machine learning: transforming data into action. Digital twin technologies and smart cities pp. 3–17 (2020)
18. Kritzinger, W., Karner, M., Traar, G., Henjes, J., Sihn, W.: Digital twin in manufacturing: A categorical literature review and classification. Ifac-PapersOnline **51**(11), 1016–1022 (2018)
19. Kuruvatti, N.P., Habibi, M.A., Partani, S., Han, B., Fellan, A., Schotten, H.D.: Empowering 6g communication systems with digital twin technology: A comprehensive survey. IEEE Access (2022)
20. Li, W., Lemieux, Y., Gao, J., Zhao, Z., Han, Y.: Service mesh: Challenges, state of the art, and future research opportunities. In: 2019 IEEE International Conference on Service-Oriented System Engineering (SOSE). pp. 122–1225. IEEE (2019)
21. Lombardo, A., Morabito, G., Quattropani, S., Ricci, C.: Design, implementation, and testing of a microservices-based digital twins framework for network management and control. In: 2022 IEEE 23rd International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM). pp. 590–595 (2022). https://doi.org/10.1109/WoWMoM54355.2022.00092
22. Long, Q., Chen, Y., Zhang, H., Lei, X.: Software defined 5g and 6g networks: A survey. Mobile networks and applications **27**(5), 1792–1812 (2022)
23. Malazi, H.T., Chaudhry, S.R., Kazmi, A., Palade, A., Cabrera, C., White, G., Clarke, S.: Dynamic service placement in multi-access edge computing: A systematic literature review. IEEE Access **10**, 32639–32688 (2022)
24. Rescorla, E., Korver, B.: Guidelines for Writing RFC Text on Security Considerations. RFC 3552 (Jul 2003). https://doi.org/10.17487/RFC3552, `https://www.rfc-editor.org/info/rfc3552`
25. Singh, R., Sukapuram, R., Chakraborty, S.: A survey of mobility-aware multi-access edge computing: Challenges, use cases and future directions. Ad Hoc Networks **140**, 103044 (2023)
26. Tang, F., Chen, X., Rodrigues, T.K., Zhao, M., Kato, N.: Survey on digital twin edge networks (diten) toward 6g. IEEE Open Journal of the Communications Society **3**, 1360–1381 (2022)
27. Tao, F., Zhang, H., Liu, A., Nee, A.Y.: Digital twin in industry: State-of-the-art. IEEE Transactions on industrial informatics **15**(4), 2405–2415 (2018)
28. Telecommunication Standardization Sector - International Telecommunication Union: Digital twin network - requirements and architecture. Series Y: Global Information Infrastructure, Internet Protocol Aspects, Next-Generation Networks, Internet of Things and Smart Cities. (Future Networks), 1–26 (2022)
29. Tuli, S., Mirhakimi, F., Pallewatta, S., Zawad, S., Casale, G., Javadi, B., Yan, F., Buyya, R., Jennings, N.R.: Ai augmented edge and fog computing: Trends and challenges. Journal of Network and Computer Applications p. 103648 (2023)
30. van der Valk, H., Haße, H., Möller, F., Otto, B.: Archetypes of digital twins. Business & Information Systems Engineering pp. 1–17 (2021)
31. Verdecchia, R., Scommegna, L., Vicario, E., Pecorella, T.: Network Digital Twins: Towards a Future Proof Reference Architecture. International Workshop on Digital Twin Architecture (2023)
32. Wu, Y., Zhang, K., Zhang, Y.: Digital twin networks: A survey. IEEE Internet of Things Journal **8**(18), 13789–13804 (2021)

33. Yastrebova, A., Kirichek, R., Koucheryavy, Y., Borodin, A., Koucheryavy, A.: Future networks 2030: Architecture & requirements. In: 2018 10th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT). pp. 1–8. IEEE (2018)
34. Yi, B., Wang, X., Li, K., Huang, M., et al.: A comprehensive survey of network function virtualization. Computer Networks **133**, 212–262 (2018)
35. Zhou, C., Yang, H., Duan, X., Lopez, D., Pastor, A., Wu, Q., Boucadair, M., Jacquenet, C.: Digital Twin Network: Concepts and Reference Architecture. Internet-Draft draft-irtf-nmrg-network-digital-twin-arch-03, Internet Engineering Task Force (Apr 2023), `https://datatracker.ietf.org/doc/draft-irtf-nmrg-network-digital-twin-arch/03/`, work in Progress