Realising Distributed Digital Twins within Federated Digital Infrastructures

Dylan Kierans[a], Dirk Pleiter[a]

[a] *KTH Royal Institute of Technology, Division of Computational Science and Technology, 100 44 Stockholm, Sweden*

Abstract

Digital twins are a concept that has initially become popular in an industrial context to support product life cycle management. Over time, the number of domains where this concept is applied has grown significantly. This includes, in particular, domains where distributed digital infrastructures become mandatory for operating digital twins. A prominent example is the earth systems, weather, and climate domain.

In the area of digital infrastructures, we observe significant efforts towards the federation of computing, storage, and data management services. In the context of digital twins, we consider efforts of particular interest that aim for distributed digital infrastructures based on geographically distributed resources with services operated by a diversity of organisations.

In this paper, we review a selection of use cases for distributed digital twins. Analysing these use cases and their implementation leads us to a set of common features and requirements. The key idea of this paper is to link these to earlier identified research and development challenges in the area of federated digital infrastructures.

Keywords

Digital twins, federated digital infrastructures, high-performance computing (HPC),

1. Introduction

One can observe a significant boost in interest in the concept of Digital Twins (DTs). This can be concluded from an almost exponential growth in the number of industry-related publications using DT in their title [1]. Also in the academic context, this concept receives increased interest in rather diverse of research domains, ranging from earth system modelling to medicine and biodiversity.

Over time, a variety of definitions of DTs have been proposed (for a review, see [2]). They have in common that a real entity is connected to one or more virtual entities. Most of the research is focussed on how to design and realise the virtual entities, i.e. the creation of suitable models. In the context of this paper, we consider it important to highlight three distinctive features of DTs with respect to established modelling approaches. Firstly, the real and virtual entities are connected, i.e. twinned, through virtual-to-real and real-to-virtual couplings. Secondly, the virtual entities are dynamic models that are regularly updated through the real-to-virtual coupling and are designed with the purpose of impacting the real entity. Finally, in many cases it is foreseen to have a human in this loop, i.e. interactivity can be an essential feature.

Several recent efforts aim for complex DTs which are expected to be used by a large and diverse set of users. Both operations and the use of these DTs will require a system comprising different services based on a diversity of computing and storage resources. Some of these services, like services facilitating analysis of DT output data products, may be rather loosely connected with the DT itself. In the following, we will refer to this system as the *digital twin system*. Note that this terminology does

---

DiDiT 2024 – 1st International Workshop on Distributed Digital Twins, June 17, 2024, Groningen, Netherlands

dylanki@kth.se (D. Kierans); pleiter@kth.se (D. Pleiter)
0009-0006-9861-0190 (D. Kierans); 0000-0001-7296-7817 (D. Pleiter)

CEUR Workshop Proceedings (CEUR-WS.org)

not imply this being a monolithic and/or closed system. In particular, for DTs in the area of earth system or biodiversity modelling, the openness of the system has been identified as a key feature, which, e.g., can facilitate the involvement of a larger diversity of end-users including citizen scientists. This extended view of DTs is important as it allows for a more comprehensive view of opportunities and challenges when using federated digital infrastructures.

Here, we aim to document the importance of federated digital infrastructures for realising complex DTs. The importance has also been recognized in a recent joint report of several national academies in the US: “Digital twins must seamlessly operate in a heterogeneous and distributed infrastructure” [3]. Still, only a limited amount of research on digital infrastructure requirements has been published.

In the academic context, there have since many years been various efforts towards realising federated digital infrastructures. We use this term to refer to a set of digital services based on a diversity of computing, storage, and network resources that are operated by a set of geographically distributed organisations that may be located in different states, i.e. be subject to different legislation. Federation of services including the underlying resources may happen at different levels and may result in a loose or tight integration of these services. One example is the integration of services within a single Identity Access Management (IAM) framework to ensure coherent access of users to all relevant services, independently of the organisation that operates those services.

We will base our analysis on an abstract view of digital infrastructure design based on two layers. A lower layer, which we call the *digital infrastructure services layer*, comprises all sufficiently generic services to serve multiple user communities. Examples of such services are services that facilitate the execution of compute jobs on High-Performance Computing (HPC) systems, the spawning of services deployed in containers, or the access to storage resources through established object-store interfaces like S3. We refer to the upper layer as the *digital platform services layer*. The latter comprises domain-specific services like metadata services or web portals designed for specific DT systems.

This paper makes the following contributions: Firstly, an analysis of selected DTs that are being deployed in or would benefit from a distributed environment and the assessment of common features and requirements. Secondly, an overview of key benefits and research and development challenges related to the design of federated digital infrastructures suitable for operating distributed DT systems.

The paper is organised as follows: We start with an overview of related work in the next section. Based on the documentation of selected use cases in section 3, we perform an analysis to discuss common features and requirements relevant from a digital infrastructure design perspective in section 4. On this basis, we discuss the needs with respect to federated digital infrastructures and identify key opportunities as well as research and development challenges in section 5. Finally, we present our summary and conclusions in section 6.

2. Related Work

Related research can be divided into two aspects of this paper, i) DT reviews, and ii) federated digital infrastructures.

Firstly, focusing on the former research into DTs in industrial applications. Systematic reviews of DT-related publications show exponential growth in the term, particularly after an initial review in 2017 [1]. Recent work attempts to port pre-existing knowledge in industry to bridge gaps in the manufacturing sector [4], with a focus on use-cases for DTs in a range of processes from production, to predictive maintenance and also after-sales services. An important area for DTs in the industrial area is management and optimisation of maintenance. A good overview is provided by a systematic literature review covering over 800 research papers relating to “Digital Twins” and “Predictive

Maintenance" [5]. The study identifies different design patterns, including "Digital Monitoring" and "Digital Control". One conclusion of the study is that by far most of the studies deal with monitoring, while far fewer studies consider the use of a DT for control. This may indicate that real-to-virtual coupling is considered to be more critical than virtual-to-real coupling. In this study, challenges related to digital infrastructures are not covered. The term "platform" is not used to describe a services layer, but rather to refer to modelling frameworks.

Focusing now on efforts towards federated digital infrastructures. Chervenak et al. [6] proposed the definition of the data grid, in order to deploy the federation of heterogeneous storage resources. The grid is defined from fundamental principles of distributed data management with respect to storage systems and metadata management. The paper suggested an abstraction from the underlying storage systems such that users could have a uniform view of data and access mechanisms. This has later been extended to compute grids. The grid approach to federated digital infrastructures has been realised at a global scale in the context of the Large Hadron Colider (LHC) [7]. The grid has become a common concept for federated infrastructure design for various research domains.

Nowadays, there are various technologies which support abstractions of distributed infrastructures. The LEXIS platform is an integrated tool for workflows on cloud and HPC resources. The platform provides Distributed Data Infrastructure (DDI), Orchestration (YORC/HEAppe), and IAM in a distributed environment in order to enable complex HPC workflows [8]. LEXIS DDI is based on the middleware iRODS designed for federating distributed storage. Importantly to bridge the HPC-cloud divide, the middleware supports a diverse range of common back-end storage solutions and includes support for object storage interfaces such as S3, which has been developed by AWS, a major public cloud provider, and is increasingly frequently used at traditional HPC sites. LEXIS Orchestration can enable hybrid use of HPC and cloud resources.

A shortcoming of the aforementioned approaches is the reliance on a small number of technical solutions developed by academic groups. These efforts could in part not catch up with the extremely fast development of technologies in the context of the emergence of public cloud providers. Several projects that aimed at the realisation of federated digital infrastructures tried to overcome these shortcomings by focussing on cloud technologies with good commercial support, while trying to avoid vendor and technology locks. One example is Fenix, which emerged from a large EU-funded project relating to brain science [9]. The project was driven by the diverse requirements of that community in terms of both small- and large-scale simulations, management, sharing, and processing of extreme-scale data sets, and support of collaborative research. Fenix initially comprised 5 and later 6 supercomputing centres around Europe. In the US, the NFS funded similar efforts towards a federated digital infrastructure including HPC and resources based on cloud technologies in the Jetstream and Jetstream2 projects [10]. During the second term of the project, a particular focus was on bridging the HPC gap towards Continuous Integration and Continuous Development (CI/CD) environments typically dealt with in the cloud. This required also organisational changes relevant to DTs, namely a change of the resource allocation policies. In this case, more flexibility was required with the goal of hiding details of the resource allocation process from end-users.

An overview of various challenges related to the realisation of federated digital infrastructures is provided in a white paper by ETP4HPC as a collaborative effort of representatives from industry (in particular, providers of HPC solutions) and academia [11]. The goal of this white paper is similar to this contribution, namely identify topical research areas for future research and innovation activities in the context of federated digital infrastructures.

3. Use Cases

In this section, we review a selected set of DT use cases and implementation examples, with a focus on complex DTs that are likely to depend on distributed digital infrastructures. While we do not claim a comprehensive overview, we claim that our choice ensures a suitable level of diversity for identifying relevant design patterns and challenges for designing federated digital infrastructures.

In the context of standardisation efforts, the US National Institute of Standards and Technology (NIST) published a set of use cases in the context of smart manufacturing [12]. These include DTs for machine health for monitoring and schedule adjustment, scheduling and routing manufacturing processes, and machining system commissioning. These use cases have been defined in a too generic manner to serve as input for digital infrastructure requirements identification but provide an overview of key industry use cases.

This concerns, e.g., the aerospace industry. Here, DTs are used more specifically for aircraft design and manufacturing [13]. These efforts involve large amounts of data, as modern aeroplanes generate vast amounts of sensor data in the $O(1)$ TByte/day/plane range. DTs in this area rely on infrastructures for data collection, data transformation, and data processing. The latter includes HPC capabilities for model simulations [14].

Digital twins have been used for more than 20 years in the electrical industry, where various benefits have been identified, including managing expensive assets as well as designing and managing electrical grids [15, 16]. Electrical grid management involves monitoring the energy status of a largely distributed infrastructure. Use cases for DTs in the area of electrical grids include management of the demand side and component lifecycle management. Digital twins in this area require the ability to integrate data from distributed sensors and databases, near-real-time responses, and support of high security levels.

Another example of DTs for large-scale infrastructures is DTs for ports [17]. These are complex as they should integrate multiple twins to model a port as a smart city and to model supply chains. Such a system of multiple DTs can serve a broad variety of goals, including supply chain optimisation, work safety improvement, and smart decision-making, e.g., in the context of berth allocation, quay crane assignment, and quay crane scheduling. Users of such a system involve various stakeholders including terminal operators, vessel operators, land transport operators, industry associations, municipalities, and government agencies.

A very different set of use cases arise from medicine. Here, DTs have been proposed as an approach for realising personalised medicine [18]. One specific case is the use of DTs for clinical oncology [19]. Wherein an approach is to combine patient-specific imaging with mechanism-based modelling to form a patient-specific DT. The DT is continuously updated in the course of a treatment based on images created during pre-treatment imaging and later monitoring sessions. One challenge is that repeatedly solving such models can be computationally demanding, resulting in the need for HPC resources. This challenge can be mitigated by using data-driven or hybrid data-driven and mechanistic approaches. Another challenge is the handling of highly sensitive patient data.

In recent years, various initiatives have been started to establish DTs in the area of earth system and ecosystem modelling. The most prominent example is Destination Earth (DestinE)[1], which has the goal to develop on a global scale a digital model of the earth to monitor and predict the interaction between natural phenomena and humans. The initiative was started by the European Commission [20]. DestinE will be a very complex digital twin system comprising multiple DTs, e.g. a DT for extreme weather events [21]. The latter will be based on a weather simulation model that requires

---

[1] https://destination-earth.eu/

large-scale HPC systems for the timely generation of data products. A second DT that is currently being developed focuses on climate change adaptation. Data is expected to be in the future produced at a rate of 1 PByte/day. Similar to the previous use case, a large number and diverse set of users are foreseen to interact with the DT or the derived data products and models.

In close relation to the DestinE efforts, a set of DTs of the ocean are being developed within the ILIAD project[2] [22]. The goal is to combine in a DT setup the sensing of ocean parameters, forecasting models, and data analysis (including pattern recognition) algorithms. One challenge is the integration of a vast set of data sources.

Another related project, called BioDT[3], is developing a variety of DTs in the area of biodiversity. One example is a DT for honey bees [23], which is based on a mechanism-based model, namely BEEHAVE. This high-resolution ecological model supports only a relatively small spatial extent. Covering the area of a country like Germany requires thousands of runs and is suitable for an HPC system providing $O(100)$ CPU cores or more. Also in this case, the ability to aggregate various types of data from multiple sources is important, including land cover data, weather data, bee monitoring data, model parameters, and flower resource parameters. The DT will be suitable for different types of end-users, who mainly will interact through a web-based interface to generate and explore bee vitality maps based on available model data. The model does support analysis of climate change effects and, therefore, coupling of this DT to the relevant DTs developed in the context of DestinE is of interest. It is furthermore foreseen, to facilitate adapted model executions by advanced users.

As a final area of use cases, we consider an emerging use case where digital infrastructures are connected to DTs. There are various efforts towards developing DTs for HPC systems (see, e.g., [24]) as well as an effort for creating a DT for a compute continuum infrastructure [25].

4. Use Cases Analysis

In this section, we analyse the previously documented use cases with a focus on common features and requirements that impact federated infrastructure design.

The DTs listed in the previous step have the following workflow steps in common (see Fig. 1 for a graphical representation):

1. Real-to-virtual coupling: Pre-processing and aggregation of raw data from multiple sources, resulting in model input data.
2. Virtual entity update: DT model building or update resulting in an update of the model state parameters and output data products.
3. Virtual-to-real coupling: Model evaluation based on model state parameters, and analysis of the output data products.

The real-to-virtual coupling is realised by injecting data into the digital twin system. Most of the DTs use raw input data from a broad range of sources. In some cases, this data is retrieved from databases and, therefore, is in principle persistently identifiable input data. In other cases, the raw input data might be data streams, e.g. sensor data. In either case, this data typically needs to be converted, transformed, and aggregated before it can be used as input for a DT. The pre-processed data may be stored within the digital twin system, e.g. for making DT model updates reproducible.

The previous section 3 revealed a broad range of modelling approaches used during the virtual entity update. In many cases, data-driven modelling is used, e.g. modelling based on machine learning.

---

[2] https://ocean-twin.eu/

[3] https://biodt.eu/

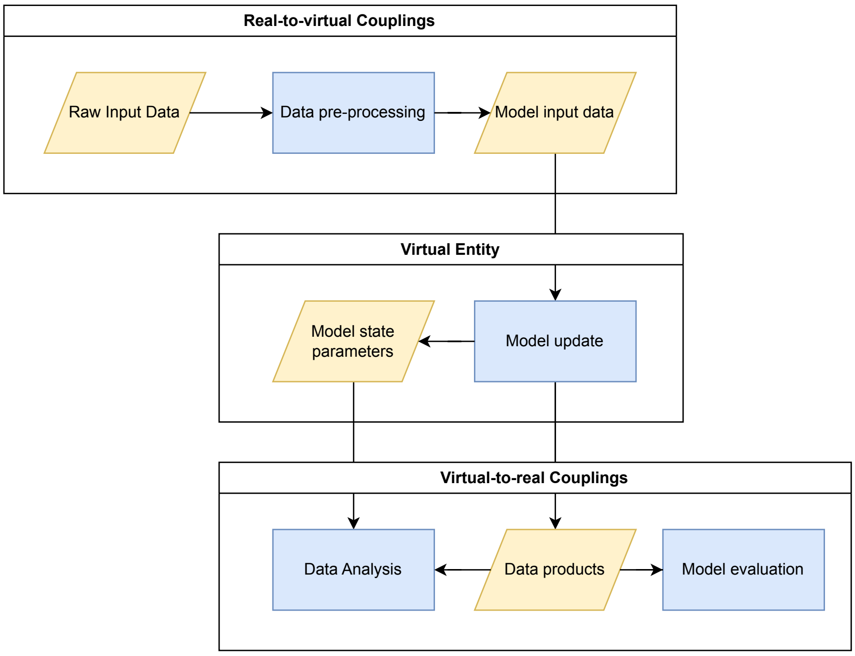

Figure 1: Abstracted workflow

Also, statistical methods are used to parameterise models based on relatively simple mathematical expressions that typically require limited computational resources. But there are also mechanism-based and other simulation models involved, or some combination with data-driven approaches. The justification for large computing capabilities during modelling may vary based on the use case. In some cases, a single large-scale simulation is performed or a deep neuronal network (DNN) with a very large number of parameters is trained. In other cases, it may be necessary to perform computations within a given time limit, e.g. to facilitate fast decision-making. The computations might, therefore, require either a significant compute capability, which can only be provided by HPC systems, or a large compute capacity, which could also be realised through more loosely coupled computing resources.

The model building or update step results in a new set of DT state parameters as well as various data products. These model parameters may facilitate fast and interactive model evaluation, an approach that is reported for ecological models. Data produced by DTs is generally of interest for further processing and data analysis (including visualisation). Based on the information found in the literature, it is difficult to assess the amount of data that is produced during virtual entity updates, but in a few cases, this may reach 1 PByte/day.

It is important to emphasize that the virtual-to-real coupling step in many cases foresees a human being in the loop, i.e. the digital twin system needs to support interactivity. Most of the use cases presented in the previous section do foresee either a large number of users interacting with the system or the set of users to be at least very diverse. User types may include, but are not limited to, governing bodies, industry, academia, citizen scientists, or recreationalists.

An emerging topic is the coupling of different DTs with examples mentioned in the previous section. In some cases, the coupling happens in the real world at long timescales, which relaxes time constraints for coupling at the virtual entities level. In other cases, time constraints are tight, e.g. to connect DTs in the area of smart cities with those in the area of weather. Here the state of the real entities may change at a timescale of $O(1)$ hour and decision-making processes may require the ability to respond even faster.

A final commonality of several use cases is the need for security mechanisms, as they involve sensitive data. This includes the case of DTs for ports, which in various countries are classified as critical infrastructure, as well as the medicine-related DTs, which involve the processing of patient data.

5. Federated Infrastructure: Opportunities and Challenges

In this section, we investigate based on the preceding analysis opportunities from using federated digital infrastructures for realising digital twin systems. Furthermore, we explore various challenges for the future design of such infrastructures. We leverage earlier work in the context of ETP4HPC [11].

5.1. Opportunities

Despite additional complexities and various challenges, the use of federated digital infrastructures creates various opportunities that can be leveraged for distributed digital twin systems.

Firstly, these types of digital infrastructures make it easier to integrate a diverse set of services based on a variety of computing and storage resources. Federated digital infrastructures allow realising a compute continuum based on edge devices for aggregating sensor data and HPC systems for large-scale simulations or model training.

The flexibility to distribute computing and storage geographically can be exploited to optimise for data locality and to optimise for latency. The latter is relevant in cases where low-latency responses are mandatory. The former is relevant in cases where data volumes are too large to be easily transferred, such that it becomes mandatory to bring computational tasks close to the data. Data locality might also become mandatory for protecting sensitive data. Exporting, e.g., patient data outside certain organisational boundaries may result in excessive costs for extending security measures or be prohibited by regulatory boundary conditions, in particular when federated digital infrastructures extend beyond state borders.

Thirdly, pooling of resources within a federated digital infrastructure increases the chances of resource availability within short waiting times. This can be exploited for new computing paradigms like function-as-a-service, which makes it easier to relocate computations in cases where the amount of processed data is small.

Geographically distributed resources facilitate the replication of services and data in a manner that minimum availability can be significantly improved even in case of significant site incidents. The TIA-942 standard foresees 99.671 % as basic and 99.995 % as the highest availability level [26]. The latter translates into an annual downtime of 0.4 hours. With DTs becoming used for time-critical decision-making, the highest availability level may become a necessary target, although we have not found availability requirements specifications in the literature, yet. Achieving such high levels of availability at the digital infrastructure services layer is often very expensive and, therefore, replication at a digital platform services layer may be preferable to reach a high availability level.

Furthermore, federated digital infrastructures may make it also easier to extend the infrastructure when needed without having to rely on the capabilities or to depend on the boundary conditions of a single organisation to extend the provisioned resources and services. Extensibility can have multiple aspects, e.g. extending the available capacity or diversity of services and resources. There is also the aspect of extending services towards more users, in particular at a global level including users in the so-called Global South [27], which is relevant to initiatives like DestinE.

Finally, there are examples like the LHC grid [7] where a federated digital infrastructure led to more openness to a broader range of users. Concepts like the creation of virtual organisations [6], which have been established in the context of federated digital infrastructures like the LHC grid [7], shifted the control from organisations that provide services and resources to user communities.

5.2. Challenges

Various challenges remain to realise federated digital infrastructures that fully support the operation of digital twin systems, as documented in section 3.

Identity and Access Management: A basic level of integration of services within a federated digital infrastructure is realised through an IAM. An IAM is necessary to ensure coherent access of users to all relevant services, independently of the organisation that operates a given service. For the latter, it has the benefit of reducing the efforts related to user management. Note that it is not required to use a single IAM framework for a digital twin system. For instance, different frameworks may be used for different types of services.

There have been various initiatives in Europe that are working towards federated digital infrastructures with an IAM based on the architecture proposed by the AARC project [28]. In conjunction with the global network of Identity Providers (IdPs) organised in eduGAIN[4], this allows to cover many of the target users at a global scale, but gaps remain. Given the breadth of user communities for the earlier documented digital twin systems, there is a high risk that users (e.g. citizen scientists) cannot obtain a virtual identity from any of the supported IdPs. Even in the case of an IdP being available, part of the security requirements of some of the digital twin systems may mandate a level of assurance that the given IdP does not guarantee.

Integration of HPC computing resources: The integration of such resources remains challenging for various reasons. One of the main reasons is the high security level that needs to be enforced for HPC systems. Users of such systems will have to fulfil specific requirements. One illustrative example is compliance with current EU embargo regulations concerning persons linked to institutions in Belarus and Russia released in the context of the war in the Ukraine[5], which requires information about users that is not easily available. This challenge can be mitigated by not providing users with full access to an HPC system. Instead, services can be implemented that allow users to start only specific workloads. While such solutions have been realised, these implementations are often site-specific, both in terms of technical implementation as well as organisational setup.

Access to diverse storage resources and services: In the area of storage services, the diversity of deployed solutions with different interfaces makes data management in a distributed environment generally challenging. The diversity ranges from large parallel file systems attached to HPC systems with a (near-)POSIX interface to object stores like Ceph with an S3 interface. Today, typically, the term data lake is used to refer to a pool of storage resources that is organised through metadata [29]. There are different strategies to overcome this challenge. One is to mandate a single interface for federated storage services like WebDAV in the case of the current LHC grid [30], Swift in the case of Fenix [31], or iRODS in the case of LEXIS [8]. This strategy bears the risk of vendor or technology locks. Another strategy is the implementation of dedicated services that can cope with different types

---

[4] https://edugain.org/

[5] http://data.europa.eu/eli/reg/2022/328/oj

of storage service endpoints. One example of this strategy is DestinE where so-called “bridges” are implemented and deployed to interface with the different storage resources of the Destination Earth Data Lake (DEDL).

The differences in the interfaces make it difficult or even impossible to realise a coherent access control. In the future, approaches like signed URLs, as they are used by Amazon CloudFront, could be a possible solution. Providing such user-specific URLs for one object or a set of objects allows, in principle, to realise fine-grained, attribute-based access control.

Alternatively, one can address the challenge of diversity of interfaces to generic storage services and coherent access control by introducing a data management middleware layer like Rucio [32]. This solution is increasingly widely used by the LHC community and has been extended to the astroparticle and radio-astronomy communities [33].

Resources allocation in a federated context: With services being deployed on top of resources provided by different organisations, resource allocation and resource consumption monitoring become a challenge. This is both an organisational as well as a technical challenge. In academia, there are currently very few mechanisms that foresee resource allocation at an international level. In the context of HPC, private cloud, and related resources, recently promising solutions like FURMS [34] and Puhuri[6] have been developed. Both solutions support the allocation of different types of resources provided by different organisations, which are features relevant to the distributed operation of digital twin systems. FURMS and Puhuri are, however, not yet widely supported.

Trust federation: As the distributed resources are assumed to be operated by different organisations, they will be operated in different security domains and data will be transferred over organisational boundaries. To use such federated digital infrastructures for workflows, for which confidentiality and security are particularly critical, requires reconsideration of existing security measures, more active security management targeting harmonised security levels at different sites as well as suitable mechanisms for establishing trust. While some aspects need to be addressed at a policy level, others require suitable technical approaches and solutions. Examples of the latter are the application of security design principles like Zero Trust [35], the provisioning of security mechanisms like trusted execution environments (TEE), and the support of data encryption both while it is in flight or at rest. The latter is in current federated digital infrastructures often not or not well-supported as important components like key vaults are either not available or not well integrated.

Short response times and interactivity: Some of the digital twin systems presented earlier do have requirements for low response times. They mainly result from the requirement of interactivity, where time scales are determined by the human perception of a responsive system. Here, the upper limit is at about 100 ms from the start of the interaction until the system’s reaction arrives back at the user. Another source of response time constraints is reality, i.e. cases where events in the real environment happen at relatively short time scales. Examples are extreme weather or Tsunami events. In the latter case, response times for running statistical models (and possibly also faster-than-real-time simulations) have to be $O(1)$ minute [36].

Support of a compute continuum: Few of the explored use cases mention real-time or near-real-time requirements, although quantitative details are lacking. There are, however, scenarios where DTs might be connected to processes with known requirements. One example is intersection control,

---

[6] https://puhuri.io/

where image data has to be translated into actions within 300 ms, coupled with smart city or digital infrastructure DTs. This requires an extension of federated digital infrastructure towards the edge, i.e. a realisation of a compute continuum. While there are many efforts towards such realisations, practical implementations are still rare.

Federation of Service Level Agreements (SLAs): Interactivity, short response times, and security are several aspects of digital twin systems where SLAs can be important. Today, there is, however, a severe lack of (de facto) standards that could facilitate a coherent implementation of SLAs in a distributed environment involving various organisations that act as service providers. In the area of security, this concerns common standards on topics like physical data centre security, (storage) asset management, security incident management, or implementation of data protection regulations. In practice, this may result in the need for multi-party agreements for each federated digital infrastructure, where requirements are defined on a case-by-case basis.

Network connectivity: Network connectivity from any user to the geographically distributed services as well as between the different services can become a limiting factor. Tight coupling of digital twins may result in the need for guaranteed bandwidth and, therefore, dedicated service contracts with network providers. Based on experience with large-scale distributed digital infrastructures like the LHC grid, there is likely sufficient connectivity available in Western countries, North America, Japan, and Australia. However, this may not be the case for countries in the Global South that are relevant to some of the earlier presented digital twin systems.

6. Summary and Conclusions

In this paper, we reviewed various digital twin systems from various application areas including industrial cases, urban infrastructures, medicine, earth systems and ecosystems as well as digital infrastructures. We argue that to fully exploit the digital twins, these need to be embedded in a variety of services to form a digital twin system. For most of the considered cases, we believe a federated digital infrastructure to be beneficial or even required in order to fulfil the diverse set of requirements. We identified an abstract workflow description that matches many of the considered cases and, furthermore, documented several common requirements like interactive access involving a diverse set of stakeholders, the upcoming need for enabling coupling of DTs, and requirements related to security. On this basis, we assessed both opportunities and challenges for realising suitable federated digital infrastructures.

The list of challenges includes known problems like federation of identity and access management, integration of HPC resources, and coping with a diversity of generic services for accessing storage resources. Promising approaches for addressing the need for support of resource allocation and resource consumption monitoring within federated digital infrastructures are ongoing. In the future, more efforts are needed to address the challenge of establishing a trust federation, provide the necessary levels of security, guarantee short response times, and improve interactivity.

The identified challenges require both changes at an organisational and policy level as well as further technical research and innovation efforts. Addressing interesting challenges from an academic standpoint like pushing new compute paradigms as the compute continuum will help advance the concept of DTs and apply it in more areas.

Acknowledgements

This study has received funding from the European Union's (EU's) Horizon Europe research and innovation programme under grant agreements No 101057437 (BioDT, https://doi.org/10.3030/101057437) and 101092582 (DECICE, https://doi.org/10.3030/101092582). Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the EU or the European Commission (EC). Neither the EU nor the EC can be held responsible for them.

References

[1] M. Sjarov, et al., The Digital Twin Concept in Industry - A Review and Systematization, in: 2020 25th IEEE ETFA Conference, 2020, pp. 1789–1796. doi:`10.1109/ETFA46521.2020.9212089`.

[2] D. Jones, et al., Characterising the Digital Twin: A systematic literature review, CIRP Journal of Manufacturing Science and Techn. 29 (2020) 36–52. doi:`10.1016/j.cirpj.2020.02.002`.

[3] National Academy of Sciences, National Academies of Engineering, National Academies of Medicine, Foundational Research Gaps and Future Directions for Digital Twins, The National Academies Press, Washington, DC, 2024. doi:`10.17226/26894`.

[4] Y. Jiang, et al., Industrial applications of digital twins, Philosophical Transactions of the Royal Society A 379 (2021) 20200360. doi:`10.1098/rsta.2020.0360`.

[5] R. van Dinter, et al., Predictive maintenance using digital twins: A systematic literature review, Information and Software Techn. 151 (2022) 107008. doi:`10.1016/j.infsof.2022.107008`.

[6] A. Chervenak, et al., The data grid: Towards an architecture for the distributed management and analysis of large scientific datasets, Journal of Network and Computer Applications 23 (2000) 187–200. doi:`10.1006/jnca.2000.0110`.

[7] I. Bird, et al., The Organization and Management of Grid Infrastructures, Computer 42 (2009) 36–46. doi:`10.1109/MC.2009.28`.

[8] S. Hachinger, et al., Leveraging High-Performance Computing and Cloud Computing with Unified Big-Data Workflows: The LEXIS Project, Springer International Publishing, Cham, 2022, pp. 159–180. doi:`10.1007/978-3-030-78307-5_8`.

[9] S. Alam, et al., Fenix: Distributed e-Infrastructure Services for EBRAINS, in: K. Amunts, L. Grandinetti, T. Lippert, N. Petkov (Eds.), Brain-Inspired Computing, Springer International Publishing, Cham, 2021, pp. 81–89. doi:`10.1007/978-3-030-82427-3_6`.

[10] D. Y. Hancock, et al., Jetstream2: Accelerating Cloud Computing via Jetstream, in: Practice and Experience in Advanced Research Computing, PEARC 2021, ACM, New York, NY, USA, 2021. doi:`10.1145/3437359.3465565`.

[11] U.-U. Haus, S. Narasimharmurthy, M. S. Perez, D. Pleiter, A. Wierse, Federated HPC, Cloud and Data infrastructures, 2022. doi:`10.5281/zenodo.6451288`.

[12] G. Shao, et al., Use Case Scenarios for Digital Twin Implementation Based on ISO 23247, NIST, Gaithersburg, MD, USA (2021). doi:`10.6028/NIST.AMS.400-2`.

[13] L. Li, et al., Digital Twin in Aerospace Industry: A Gentle Introduction, IEEE Access 10 (2022) 9543–9562. doi:`10.1109/ACCESS.2021.3136458`.

[14] E. Glaessgen, D. Stargel, The digital twin paradigm for future NASA and US Air Force vehicles, 2012. URL: https://ntrs.nasa.gov/citations/20120008178.

[15] G. D. M. Serugendo, et al., Digital Twins: From Conceptual Views to Industrial Applications in the Electrical Domain, Computer 55 (2022) 16–25. doi:`10.1109/MC.2022.3156847`.

[16] M. M. H. Sifat, et al., Towards electric digital twin grid: Technology and framework review, Energy and AI 11 (2023) 100213. doi:10.1016/j.egyai.2022.100213.

[17] R. Klar, et al., Digital Twins for Ports: Derived From Smart City and Supply Chain Twinning Experience, IEEE Access 11 (2023) 71777–71799. doi:10.1109/ACCESS.2023.3295495.

[18] R. Laubenbacher, et al., Digital twins in medicine, Nature Computational Science 4 (2024) 184–191. doi:10.1038/s43588-024-00607-6.

[19] C. Wu, et al., Integrating mechanism-based modeling with biomedical imaging to build practical digital twins for clinical oncology, Biophysics Reviews 3 (2022) 021304. doi:10.1063/5.0086789.

[20] S. Nativi, et al., Digital Ecosystems for Developing Digital Twins of the Earth: The Destination Earth Case, Remote Sensing 13 (2021). doi:10.3390/rs13112119.

[21] N. Wedi, et al., Destination Earth: High-Performance Computing for Weather and Climate, Computing in Science and Engineering 24 (2022) 29–37. doi:10.1109/MCSE.2023.3260519.

[22] M. Majidi Nezhad, et al., Marine energy digitalization digital twin's approaches, Renewable and Sustainable Energy Reviews 191 (2024) 114065. doi:10.1016/j.rser.2023.114065.

[23] J. Groeneveld, et al., Prototype Biodiversity Digital Twin: Honey Bees in Agricultural Landscapes 5 (2024). doi:10.3897/arphapreprints.e124639.

[24] T. Ohmura, et al., Toward Building a Digital Twin of Job Scheduling and Power Management on an HPC System, Springer, 2023, pp. 47–67. doi:10.1007/978-3-031-22698-4_3.

[25] J. M. Kunkel, et al., DECICE: Device-Edge-Cloud Intelligent Collaboration Framework, in: Proceedings of the 20th ACM CF Conference, ACM, New York, NY, USA, 2023, pp. 266–271. doi:10.1145/3587135.3592179.

[26] ADC Telecommunications, TIA-942 Data Center Standards Overview, 2006.

[27] R. Heeks, Digital inequality beyond the digital divide: conceptualizing adverse digital incorporation in the global South, Information Technology for Development 28 (2022) 688–704. doi:10.1080/02681102.2022.2068492.

[28] N. Liampotis, et al., AARC Blueprint Architecture 2019 (AARC-G045), 2020. doi:10.5281/zenodo.3672785.

[29] C. Madera, A. Laurent, The Next Information Architecture Evolution: The Data Lake Wave, in: Proceedings of the 8th International MEDES Conference, ACM, New York, NY, USA, 2016, pp. 174–180. doi:10.1145/3012071.3012077.

[30] Bockelman, Brian, et al., Bootstrapping a New LHC Data Transfer Ecosystem, EPJ Web Conf. 214 (2019) 04045. doi:10.1051/epjconf/201921404045.

[31] S. R. Alam, et al., Archival Data Repository Services to Enable HPC and Cloud Workflows in a Federated Research e-Infrastructure, in: 2020 IEEE/ACM SuperCompCloud Workshop, 2020, pp. 39–44. doi:10.1109/SuperCompCloud51944.2020.00012.

[32] M. Barisits, et al., Rucio: Scientific Data Management, Computing and Software for Big Science 3 (2019) 11. doi:10.1007/s41781-019-0026-3.

[33] R. Dona, R. Di Maria, The ESCAPE Data Lake: The machinery behind testing, monitoring and supporting a unified federated storage infrastructure of the exabyte-scale, EPJ Web Conf. 251 (2021) 02060. doi:10.1051/epjconf/202125102060.

[34] B. Hagemeier, 17th Fenix Infrastructure Webinar: Federated Resource Management using FURMS, 2022. URL: https://juser.fz-juelich.de/record/912053.

[35] S. Rose, et al., Zero Trust Architecture, Technical Report, National Institute of Standards and Technology, 2020. doi:10.6028/NIST.SP.800-207.

[36] F. Løvholt, et al., Urgent Tsunami Computing, in: 2019 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC), 2019, pp. 45–50. doi:10.1109/UrgentHPC49580.2019.00011.