Vrije Universiteit Amsterdam
https://www.overleaf.com/read/wjwrqfvtnjfx

Honours Programme, Problem Statement

Exploring Cloud Simulation at Scale through Digital Twinning as a Service

Author: Daniel Ştefan Halasz (2804542)

1st supervisor: Prof. Dr. Ir. Alexandru Iosup
Daily supervisor: MSc. Dante Niewenhuis

A report submitted in fulfillment of the requirements for the Honours Programme, which is an excellence annotation to the VU Bachelor of Science degree in Computer Science/Artificial Intelligence/Information Sciences

May 15, 2026

Abstract

Our society is experiencing the golden age of distributed systems. Large-scale digitalization efforts have pushed the demand for efficient, cost-effective, and sustainable datacenters to new extremes. With their increasing complexity, datacenters often face risks addressable solely by predicting future system behavior. Simulators emerged to approach the impending problem of capacity planning, but fail to be adopted in real-world operations due to their contextual limitations and scalability constraints. However useful, current simulators are not mature enough to be evolved into a datacenter digital twin, enabling their use in massive computer infrastructure. In this work, we address the problem of integrating simulation in datacenter design processes by proposing a novel, state-of-the-art reference architecture for an ICT digital twin service, incorporated in a unified tool, capable of serving real-world scenarios. We envision our contribution to benefit the entire ICT community, by proving suitable for small, medium and large use cases alike.

Keywords

ICT, datacenters, digital twin, SaaS

1 Introduction

Cloud infrastructure is a fundamental building block of digital society. The demand for distributed computing has grown proportionally to our increasing reliance on computer systems. Sectors ranging from banking and government to healthcare and academia have shifted their data-driven activity to the cloud. The economic implications of this paradigm shift are significant: in 2022, more than half of the Netherlands' GDP and over 3.3 million jobs relied on distributed computer systems, with a 15-fold return on investment [1]. Large-scale digitalization introduces a severe risk factor: datacenters are a critical point of failure. The repercussions of outages are echoed to all stakeholders, possibly leading to catastrophic consequences [2]. Faced with this responsibility, datacenter architects must provide persistent, reliable and performant infrastructure, while simultaneously minimizing the costs of operation and environmental impact. Inaccurate configurations are the underlying cause of concerning yearly energy consumption assessments. In 2020, the ICT sector's energy consumption was estimated at a staggering 205 TWh, leading to a carbon footprint as high as 2.8% of global greenhouse emissions [3].

Datacenter simulators emerged to address the complex challenges in designing cloud infrastructure. These tools enable sandbox experimentation and advanced operation analysis at datacenter level [4]. Simulators address the risks of inaccurate configurations by enabling informed decision making, allowing experts to carefully observe the effects of each design choice. *Digital twinning* is the process of seamlessly integrating the digital and physical spaces [5], through real-time simulation, long-term planning, and discrete short-term observation. A datacenter digital twin would open the possibility of investigating and measuring large-scale, non-deterministic system configurations, otherwise comparable only by real-world experimentation. Such analytic procedures would be significantly faster and less expensive than their real-world counterpart [6], greatly reducing risks involved in datacenter design when applied. Numerous advances in cloud infrastructure simulation, such as the development of OpenDC, CloudSim, and SimGrid [4, 7, 8] bridged the gap to a state-of-the-art datacenter digital twin. Despite the academic effort going towards datacenter simulation, the ICT industry still relies on rule-of-thumb decisions in its design processes.

We argue that current state-of-the-art ICT digital twinning is not ready to be applied in real-world use cases. Thus, datacenter architects rely on overprovisioning and underutilization instead of adopting simulation to combat the risks of resource scarcity and possible outages, leading to a substantial waste of energy and lack of system effectiveness [9, 10]. Although simulators are capable of producing accurate predictions, they are often limited by their narrow context, lack of scalability and the absence of a unified ecosystem. These limitations often hinder the applicability of predictive models in real-world scenarios [6]. Computer systems require more than a core computational engine to produce their desired functionality. Therefore, a digital twin lacking in context generality and scaling policies is incomplete by design and unable to fulfill its primary *purpose* [11]. We envision ICT digital twinning as a holistic system that is readily available to serve simulation requests under any datacenter configuration, with great accuracy and scalable throughput. In this work, we address the lack of a unified simulation ecosystem and the absence of a standardized interface, currently preventing the integration of simulators into a single digital twin. We explore development methodologies for the software as a service model to identify the best-suited approach for simulation. We then analyze industry-standard system development patterns, with the goal of designing a reference architecture for a *pioneering, state-of-the-art ICT digital twin service*, capable of serving real-world use cases. Furthermore, we integrate the reference architecture into a prototype with one or more simulation engines, serving as a scientific instrument and proof of concept for the proposed system. Lastly, we evaluate the prototype's performance by assessing its real-world applicability as a factor of scalability and reliability.

2 Background

In this section, we analyze the key concepts related to ICT digital twinning and simulation. We map each concept to its corresponding challenges and present its societal and economic implications in the context of large-scale distributed systems.

2.1 Datacenter Simulation

Adapting and reacting to non-deterministic system behavior is a challenging task. Predicting system capacity based on current and past states is essential to designing and operating large-scale infrastructure. Simulators address this problem by capturing digital models of real-world systems, enabling fine-grained analysis and prediction-based comparison between different configurations [4]. Simulation is widely used in many industries, including computer systems, healthcare, and aerospace engineering [12, 13, 14].

The benefits of simulation in designing distributed systems are evident: simulation-based experimentation exhibits unparalleled scalability and speed compared to its physical counterpart. Reproducing datacenter simulation experiments on real hardware would require as much energy as 3400 hospitals throughout a year [2]. The sheer cost of real-world experimentation is hardly justified and often impossible to cover relative to simulation.

Simulation can only achieve as much as the assumptions built into it and the technical limitations of its implementation. Thereby, prediction accuracy is dictated by the model employed and the instrument's technical specification [15]. Simulation models are classified into three main categories: *static or dynamic*, *continuous or discrete* and *stochastic or deterministic* [16].

2.2 Software as a Service

The rapid evolution of cloud infrastructure in the past decade enabled reliable deployment of end-user applications with little to no regard to hardware and operational concerns [17]. The *Software as a Service* (SaaS) model implies delivering applications over the internet, deployed in a cloud platform [18]. The success of SaaS lies in abstracting the execution environment, no longer dependent on the client's machine architecture. Before its emergence, developers relied on local deployment of applications, heavily restricted by architecture, hardware capacity and software compatibility, and only invoked cloud services for granular tasks, such as licensing or storage.

Deploying applications in the cloud, instead of providing a local executable provides one key advantage relevant to our work: computational heavy workflows are shifted to resourceful datacenters, leaving the application to render the user interface. This approach significantly reduces the hardware requirements specific to certain classes of resource-intensive applications. Simulation is inherently a computationally demanding task that could benefit from integration with the SaaS model.

2.3 ICT Digital Twinning

Connecting the physical and digital spaces has been a long-standing goal in many fields. One of the key benefits of digital twinning is the ability to predict, observe and analyze with high degree of granularity the behavior of a system before it is built or deployed. Over the past two decades, the concept of digital twinning evolved in many industries and applications, including computer systems, healthcare, and product design [5].

An ICT digital twin is a "virtual model that replicates the structure, context, and behavior of a data center". It is **more** than datacenter simulation: digital twinning also implies a bidirectional interaction between the virtual and physical spaces, facilitating a feedback loop, data-driven updating, and optimal decision-making. This complex system does not replace human oversight, but rather aims to augment it, by providing a high-fidelity, interactive representation of the datacenter's behavior [19, 20]. In our work, we analyze the missing components in the current state-of-the-art datacenter digital twin solutions based on the holistic definition of digital twinning.

A robust datacenter digital twin would contribute significantly to the distributed systems community [12, 19]. At a massive scale, digital twinning must meet the community's performance and reliability requirements: accuracy, speed, reproducibility, and cost efficiency are all equally important. Stakeholders across the education sector, scientific research, and medium to large datacenter operators uphold rigorous standards for the instrument's quality assurance [1].

3 Problem

Operating datacenters is a complex, non-trivial task associated with high risks and severe penalties in case of failure. To address the challenges related to datacenter design and operation, simulation tools have been developed to model and predict system behavior. We argue that existing simulators are not advanced enough to be adopted in real-world scenarios, due to their lacking features and limited context. In this work, we identify and aim to address three main problems in the field of datacenter simulation.

Firstly, the current state-of-the-art datacenter design process does not include simulation, leading to severe losses in energy efficiency and cost. Datacenter design planning suffers from guesswork, chronic underutilization, and overprovisioning of resources, following the absence of dedicated and specialized simulation tools [21]. Thus, we identify a *lack of integration* of simulators in the field of datacenter design due to their limited specialization and missing features.

Secondly, the limited capabilities of currently available simulators prevent their use as a holistic datacenter digital twin. ICT digital twinning is more than just simulation and modeling: it is a comprehensive, real-time representation of a datacenter, receiving real-time feedback from the physical environment and providing discrete, granular insights. Albeit still valuable, current simulators are not mature enough to be used as a service in real-world scenarios. We observe the need for a *thorough reference architecture for an ICT digital twin* that establishes the requirements, capabilities, and components of a scalable, reliable, and readily accessible datacenter simulation service.

Lastly, operating a simulation service should be intuitive and equitable for both novice users and expert datacenter operators alike. Existing tools are either too complex for novice users or overly simplistic for professional operators, leading to an unwieldy experience for both parties. Therefore, we identify the need for a *flexible, publicly available ICT simulation service* that can handle complex operations at a reasonable cost, and deliver fast and reliable results. We also acknowledge the need for a *holistic evaluation* of the proposed system to ensure it meets the stakeholders' requirements and expectations, aiming to identify potential shortcomings and threats to validity.

As of the moment of writing, there does not exist a comprehensive reference architecture for an ICT digital twin service, prioritizing accuracy, reliability, scalability and performance. Also, no unified instrument has been built for datacenter design and operation planning, that facilitates real-world adoption.

4 Related Work

The concept of digital twinning as a service (DTaaS) has been explored before as a paradigm to be elaborated in specific domains. In this section, we examine the related work that serves as a foundation for our research.

4.1 Domain-agnostic DTaaS Reference Architecture

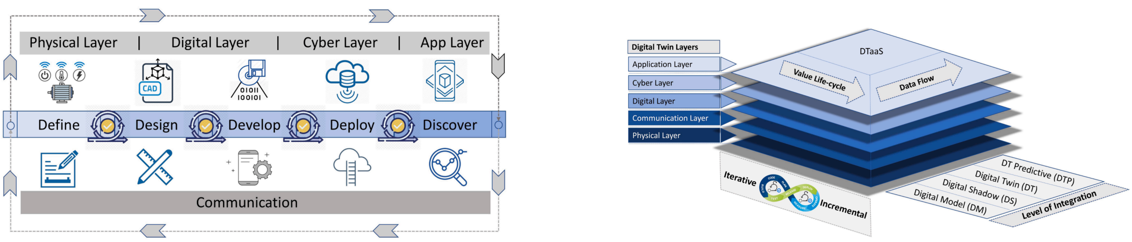

(a) 5D product design process

(b) DTaaS Vertical Layers

Figure 1: Domain-agnostic DTaaS reference architecture. (Source: [22].)

Aheleroff et al. proposes a holistic DTaaS reference architecture, generalizable to any domain. The architecture is based on the 5D product design process: *Define, Design, Develop, Deploy, and Discover* (see fig. 1a). The communication between the cyber and physical layer is essential: continuous data exchange is required to maintain the feedback loop between the digital twin and the real-world system. Five vertical layers define the DTaaS architecture: *physical layer, communication layer, digital layer, cyber layer, and application layer*, mapped one-to-one to the 5D product design process (see fig. 1b).

Each layer is significant and contributes modularly to the overall architecture. The physical layer represents the real-world system, providing traces of the system's state. The communication layer is responsible for the data exchange between the physical and digital layers. The digital layer represents the simulation core, responsible for modeling the real-world system and predicting its behavior. The cyber layer represents the digital twin's environment, a backend system that manages its lifecycle. The application layer is the frontend system that provides the user interface for interacting with the digital twin. Multiple layers of integration are possible, serving different physical layer contexts. We differentiate between two main classes of digital twinning, characterized by the directionality of the dataflow: bidirectional and unidirectional. Not all systems benefit from a bidirectional feedback loop; thus, the reference architecture opens the possibility of unidirectional dataflow systems. Beyond digital twin integration stands the predictive behavior: *Digital Twin Predictive* (DT_p), able to fully replicate and anticipate system behavior through a two-dimensional real-time data communication over cyberspace [22].

In our work, the DT_p model aligns most closely with the ICT digital twin requirements, which we aim to further develop and elaborate on.

4.2 Datacenter Digital Twin Architecture and Use Cases

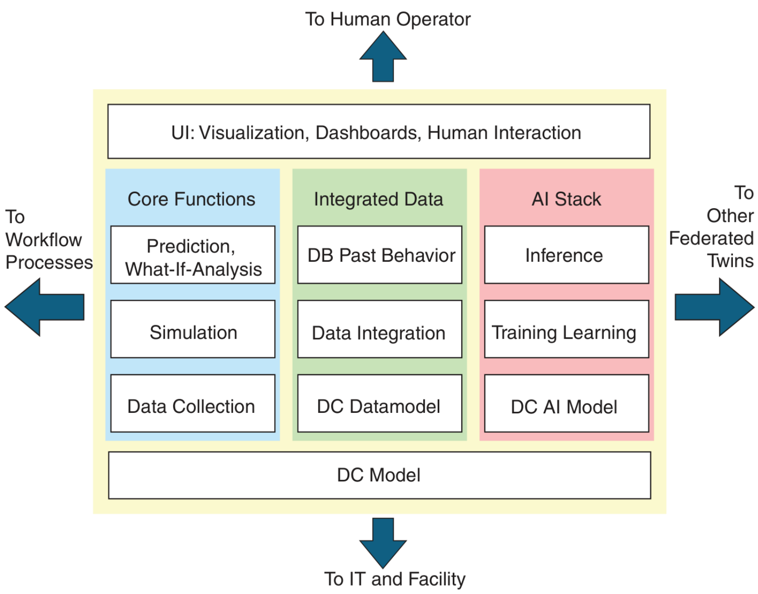

Figure 2: ICT Digital Twin architecture. (Source: [19].)

Athavale et al. defines a datacenter digital twin architecture as a continuous model replicating the state of a physical datacenter (see fig. 2). The replicated state can often be coupled with an external state to boost predictive capabilities. The simulation component combines past behavior with real-time data to predict future states. Uniquely, datacenter digital twins can be *federated*, opening the possibility of navigation between different system configurations under the same parameters and offering significant insights into the system's behavior in relation to the same workload. The authors envision, based on the referenced architecture numerous use cases for datacenter digital twins, each with notable benefits. Some of the most relevant are *capacity planning*, *energy optimization* and *scenario planning*, each shown to be effective in improving datacenter performance and efficiency [2, 4, 21].

Our work aims to expand the datacenter digital twin architecture to a DTaaS model, focusing on the ICT domain.

5 Research Questions

We address the problem statement by raising the main research question, which is further divided into three research questions **(RQ)**.

Main Research Question

How to design, prototype and evaluate a datacenter digital twin under the software as a service model, as a unified tool for datacenter simulation?

RQ1 How to design and integrate the architecture of a datacenter digital twin under the software as a service model?

Many medium-scale sub-fields of Computer Systems already apply digital twin services in their practices [5, 23]. The ICT field has yet to establish a standard for digital twinning, resulting in imprecise methodologies for datacenter operations. It is necessary that we design a *holistic reference architecture* for a datacenter digital twin, which can be integrated into a unified tool for datacenter simulation as a service.

This is a very challenging task, because the overarching nature of the ICT field imposes numerous fine-grained operational and technical requirements for the proposed system. Ideally, the architecture should exhaustively address datacenter topologies, and establish a comprehensive model of interaction between the real-world system and the digital twin.

To answer **RQ1**, we must first define a list of system requirements consistent with the state-of-the-art for datacenter design, and identify the trade-offs between scalability, reliability and system performance.

RQ2 How to design the backend infrastructure for a datacenter digital twin service?

Current simulators resemble static execution models, with no regard to task scheduling, data pipelining, and resource allocation. Our work aims to identify the best-suited execution model for a backend infrastructure that can serve a datacenter digital twin. The design must facilitate large-scale requests, carefully route data, invoke multiple simulation engines when appropriate, and seamlessly integrate with the frontend.

The key objective of the proposed system is wide adoption. Thus, its legitimacy is contingent on the ability to serve all classes of stakeholders, and integrate into a state-of-the-art unified tool. To answer **RQ2**, we analyze the industry-standard backend development patterns, and refine our system architecture in accordance with the backend infrastructure requirements.

RQ3 **How to evaluate the design of a datacenter digital twin service in relation to the system requirements?**

To understand the effectiveness of the proposed system and assess its real-world adoption potential, we must quantify its performance. We must implement a prototype of the datacenter digital twin service, and identify a set of metrics appropriate for measuring the system's scalability, reliability and accuracy. Conducting a sound evaluation methodology is challenging; we identify possible threats to validity in the prototype's reproducibility and the metrics' relevance to the ever-evolving ICT state-of-the-art. Answering **RQ3** requires validating the metrics used for performance evaluation, and ensuring their significance to the stakeholders' requirements.

6 Approach

We approach the problem statement by addressing the subsequent research questions, following the distributed systems research methodology, and guided by the state-of-the-art AtLarge Design Process [24]. The fundamental design principle crucial for architecting a datacenter digital twin service stands in *massivizing computer systems*. Primarily, our goal is to foster a verifiable design pattern that innovates at a large scale and facilitates reproducibility.

To answer **RQ1**, we conduct a literature survey on datacenter simulators, with the objective of identifying integration limitations. Further, we conduct a methodical survey of the state-of-the-art in datacenter design processes and operations, and observe the key components that benefit from capacity planning. Then, we examine primary stakeholders for an ICT digital twin service, and derive a list of functional system requirements. Having a good understanding of the ICT industry, we proceed with defining the system's non-functional requirements, capabilities, and modules, followed by a thorough trade-off analysis. Respecting the established methodical model, we explore potential designs for the system's architecture and continuously integrate new findings into a final, state-of-the-art reference architecture for a datacenter digital twin service.

To answer **RQ2**, we analyze industry-standard backend development patterns by consulting the reference architecture, and identifying the best-suited approach for datacenter simulation. Additionally, we conduct a literature review on the SaaS development model, and collect our findings into a technical specification for the system's backend infrastructure. We develop our specification with regard to *scale*; performance under high stress conditions and efficient resource allocation hold significant importance in providing a cost-effective, robust solution. Notably, we co-evolve our reference architecture to reflect the technical specification, ensuring a feasible system design.

To answer **RQ3**, we integrate our reference architecture into a unified tool, constituting an ICT digital twin service. We analyze industry-standard software engineering processes and adopt them in the development process. Then, we provide the instrument's structure and documentation, and open-source the artifacts. Furthermore, we conduct an exhaustive review of the currently known metrics for measuring simulation performance and select the ones that suit our system architecture. Lastly, we conduct experiments to evaluate the digital twin service and assess its adoption potential into real-world scenarios.

Our contribution stands threefold. We propose a novel reference architecture for datacenter digital twinning as a service, suitable for integration into real-world datacenter operations. We design and validate a technical specification for the backend infrastructure of this service, focusing on compatibility and scale. Finally, we integrate and evaluate the proposed system into a prototype able to demonstrate the real-world applicability of datacenter digital twinning.

7 Plan

In this section we document the preliminary planning for the ongoing work and establish important milestones **(MS)**. The project should span over the next 10 months, with a self-imposed deadline of 1st of October 2025.

7.1 Planning RQ1

Firstly, I need to document the currently existing datacenter simulators, to observe the shortcomings of simulation at scale. I have completed this step with OpenDC, however it is essential to investigate at least two more simulators. Afterwards, I must identify the current state-of-the-art in datacenter design. I assume this will be a crucial, non-trivial step, consisting of both a literature survey and real interviews of a few datacenter operators.

MS1 Define a list of system requirements and essential capabilities for the proposed system architecture.

The final step for this research question is designing the reference architecture. This step is divided into three parts, forming a feedback loop:

- Analyze a potential design and describe its trade-offs
- Discuss the design with the daily supervisor
- Update the design to cover shortcomings or identify the need for restructuring

MS2 Define the system reference architecture and develop an exhaustive system block diagram.

7.2 Planning RQ2

I begin by consulting the industry-standard for backend development and identify the key components of a SaaS backend infrastructure. At the same time, I conduct a literature review on the SaaS model, and ensure that the development pattern is adequate for the reference architecture. At this point it is possible to notice a discrepancy between the architecture and the backend infrastructure, in which case I will update either of them to embrace the other. Having a preliminary backed infrastructure, I shift my focus towards scalability: ensuring the design can be applied in large-scale settings.

MS3 Develop a technical specification for the backend infrastructure, and its system block diagram.

A small but essential step is validating the technical specification with the reference architecture.

7.3 Planning RQ3

It is time to finally implement the system into a prototype. I will start by choosing the most appropriate programming language for the project, and then develop the system in an iterative manner. At the same time, I conduct a literature review on software engineering processes, ensuring the best quality in the development process.

MS4 Develop a prototype of the system, and document the structure and the system's documentation.

Following with the system evaluation, I first need to list relevant metrics for measuring the system's performance. Of interest is the system's scalability and potential for real-world adoption. Finally, I conduct experiments to evaluate the system. I aim for one experiment per metric, with a minimum of three metrics.

MS5 System evaluation and multiple diagrams of the system's performance.

7.4 Results

In the end, we should be able to assess whether the proposed system is feasible and can be integrated into real-world datacenter operations. The results will be documented in the thesis, and the prototype will be open-sourced. We finish the project by creating a presentation of the results.

MS6 Completing any missing parts of the thesis and preparing a presentation of the results.

8 Conclusion

In this work, we identified an absence of simulation adoption in the design and operations of datacenters. We address three main research questions concerning the design of an ICT digital twin, and propose a novel reference architecture for datacenter simulation at scale, under the software as a service model. We integrate the proposed system into a state-of-the-art, unified, scientific instrument that demonstrates the grand potential of datacenter digital twinning. Lastly, we evaluate the system experimentally based on a set of validated metrics, relevant to datacenter operation performance and large-scale distributed system design.

References

[1] A. Iosup, F. Kuipers, A. L. Varbanescu, P. Grosso, A. Trivedi, J. Rellermeyer, L. Wang, A. Uta, and F. Regazzoni, "Future computer systems and networking research in the netherlands: A manifesto," 2022.

[2] F. S. Mastenbroek, "Data-driven risk analysis of sustainable cloud infrastructure using simulation," Master's thesis, Delft University of Technology, March 2022.

[3] C. Freitag, M. Berners-Lee, K. Widdicks, B. Knowles, G. Blair, and A. Friday, "The climate impact of ict: A review of estimates, trends and regulations," 02 2021.

[4] F. Mastenbroek, G. Andreadis, S. Jounaid, W. Lai, J. Burley, J. Bosch, E. van Eyk, L. Versluis, V. van Beek, and A. Iosup, "OpenDC 2.0: Convenient modeling and simulation of emerging technologies in cloud datacenters," in *CCGRID*, 2021.

[5] F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 4, pp. 2405–2415, 2019.

[6] A. Iosup, A. Uta, L. Versluis, G. Andreadis, E. van Eyk, T. Hegeman, S. Talluri, V. van Beek, and L. Toader, "Massivizing computer systems: a vision to understand, design, and engineer computer ecosystems through and beyond modern distributed systems," 2018.

[7] R. N. Calheiros, R. Ranjan, C. A. F. D. Rose, and R. Buyya, "Cloudsim: A novel framework for modeling and simulation of cloud computing infrastructures and services," *CoRR*, vol. abs/0903.2525, 2009.

[8] H. Casanova, A. Legrand, and M. Quinson, "Simgrid: A generic framework for large-scale distributed experiments," in *Tenth International Conference on Computer Modeling and Simulation (uksim 2008)*, pp. 126–131, 2008.

[9] M. Dayarathna, Y. Wen, and R. Fan, "Data center energy consumption modeling: A survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 732–794, 2016.

[10] Y. Ran, H. Hu, Y. Wen, and X. Zhou, "Optimizing energy efficiency for data center via parameterized deep reinforcement learning," *IEEE Transactions on Services Computing*, vol. 16, no. 2, pp. 1310–1323, 2023.

[11] D. H. Meadows, "Thinking in systems: A primer," *Sustainability Institute*, 2008.

[12] M. Molan, J. A. Khan, A. Bartolini, R. Turra, G. Pedrazzi, M. Cochez, A. Iosup, D. Roman, J. Rožanec, A. L. Vărbănescu, and R. Prodan, "The graph-massivizer approach toward a european sustainable data center digital twin," in *2023 IEEE 47th Annual Computers, Software, and Applications Conference (COMPSAC)*, pp. 1459–1464, 2023.

[13] L. M. Sutherland, P. F. Middleton, A. Anthony, J. Hamdorf, P. Cregan, D. Scott, and G. J. Maddern, "Surgical simulation: a systematic review," *Annals of surgery*, vol. 243, no. 3, pp. 291–300, 2006.

[14] R. K. Phanden, P. Sharma, and A. Dubey, "A review on simulation in digital twin for aerospace, manufacturing and robotics," *Materials today: proceedings*, vol. 38, pp. 174–178, 2021.

[15] H. A. Simon, *The Sciences of the Artificial, reissue of the third edition with a new introduction by John Laird.* MIT press, 2019.

[16] N. Kheir, *Systems modeling and computer simulation.* Routledge, 2018.

[17] E. Van Eyk, A. Iosup, S. Seif, and M. Thömmes, "The spec cloud group's research vision on faas and serverless architectures," in *Proceedings of the 2nd international workshop on serverless computing*, pp. 1–4, 2017.

[18] S. Kounev, N. Herbst, C. L. Abad, A. Iosup, I. Foster, P. Shenoy, O. Rana, and A. A. Chien, "Serverless computing: What it is, and what it is not?," *Communications of the ACM*, vol. 66, no. 9, pp. 80–92, 2023.

[19] J. Athavale, C. Bash, W. Brewer, M. Maiterth, D. Milojicic, H. Petty, and S. Sarkar, "Digital twins for data centers," *Computer*, vol. 57, no. 10, pp. 151–158, 2024.

[20] National Academies of Sciences, Engineering, and Medicine, *Foundational research gaps and future directions for digital twins.* 2023.

[21] G. Andreadis, F. Mastenbroek, V. van Beek, and A. Iosup, "Capelin: Data-driven capacity procurement for cloud datacenters using portfolios of scenarios–extended technical report," *arXiv preprint arXiv:2103.02060*, 2021.

[22] S. Aheleroff, X. Xu, R. Y. Zhong, and Y. Lu, "Digital twin as a service (dtaas) in industry 4.0: an architecture reference model," *Advanced Engineering Informatics*, vol. 47, p. 101225, 2021.

[23] W. Brewer, M. Maiterth, S. Bouknight, J. Hines, and T. J. Webb, "Exadigit/raps." [Computer Software] `https://doi.org/10.11578/dc.20240627.4`, jun 2024.

[24] A. Iosup, L. Versluis, A. Trivedi, E. van Eyk, L. Toader, V. van Beek, G. Frascaria, A. Musaafir, and S. Talluri, "The atlarge vision on the design of distributed systems and ecosystems," in *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pp. 1765–1776, 2019.