Sunfish: Enabling Predictive Analytics for Datacenters Through Digital Twinning

M. Kwiatkowski Drs. D. Niewenhuis[1] Prof. A. Iosup[2]

[1]Daily Supervisor

[2]Main Supervisor

Vrije Universiteit Amsterdam

June 26, 2026

Online slideshow: `mjkw.pl/vu/bsc`

Motivation

VRIJE UNIVERSITEIT AMSTERDAM

Context

21st century datacenters (DC) are mostly heterogeneous [10] and modern computational needs of AI drive managers to diversify datacenters even more [1]. In result datacenters become extremely complex and hard to operate with millions of CPU's, GPU's etc.

Figure 1.1: Society depends on datacenters to keep running, and therefore we cannot afford to let these systems break down or experience significant performance-related issues. With millions of servers in the largest datacenters, real-time management becomes very difficult. Left to right: a Google datacenter, server racks, Ada Lovelace AD102 GPU architecture.

Problem Statement

VU VRIJE UNIVERSITEIT AMSTERDAM

DCDT's lack predictive analytics

We need Datacenter Digital Twins (DCDT) to be better able to detect and solve issues in critical ICT infrastructure [1]. However, DCDT's are still actively developed and lack crucial features such as predictive analytics [11] to *e.g.,* prevent unexpected failures. With predictive analysis (*e.g.,* simulation) DCDT's could save millions of lost $USD [13].

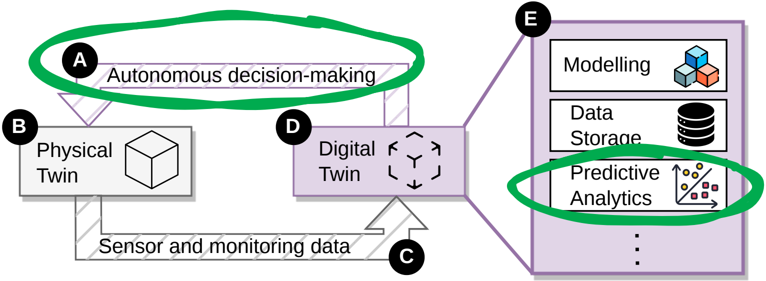

Figure 1.2: Where does our work fit within the field of datacenter digital twinning? There are 5 core elements to any Digital Twin: (A) The Digital → Physical Twin link, (B) the Physical Twin (*e.g.,* the datacenter), (C) the Physical → Digital Twin link, (D) the Digital Twin, (E) the features necessary to any Digital Twin. Highlighted areas are the contributions from this thesis, which include the autonomous actions resulting from predictive insights (A) and the predictive analysis itself within (E).

Research Questions

Main Research Question

How to enable predictive analytics for datacenters through digital twinning?

Research Question 1

How to asses the current state-of-the-art of digital twinning for datacenters?

Research Question 2

How to design a reference architecture for a predictive datacenter digital twin using discrete-event simulation?

Research Question 3

How to validate and evaluate a datacenter digital twin architecture in relation to system requirements?

RQ1: Literature Review I

Main Finding I

The literature on DCDTs is scarce. Some systems barely classify as DTs (*e.g.,* Kalibre [16], ChatTwin [8]). Existing deployments specialize in Cooling and Heat Modelling, together with 3D visualizations. Most lack predictive modelling of DC operations.

Project	Simulation Technique	Focus	Stakeholders	Modelling Capability
ExaDigiT [2]	CFD/HT, AI/ML	Energy Loss Prediction, Heat Modelling	HPC Engineers and Operators	3D⋆, CH✻, VP⋆, PE✻, RA, SE‡
SmartDC [17]	CFD/HT, BIM, AI/ML	Heat Modelling, PUE optimization	Cloud Datacenter Engineers	CH‡, PE, 3D⋆
DyTwin [12]	Gaussian Process Regression, AI/ML	Anomaly Detection	Cloud Datacenter Operators	A⋆, FD, VP⋆, SE‡
ChatTwin [8]	**?**	Digital Twin Definition Language	Cloud Datacenter Engineers	3D⋆
Reducio [3]	POD, Gaussian Process Modelling (ML)	Heat Modelling	Edge and Hyper-scale Datacenter Operators	CH✻, 3D⋆, SE
NetGraph [5]	Graphs	Network Management	Cloud Datacenter Operators	VP⋆, RA⋆, N⋆, SE‡
Kalibre [16]	CFD/HT, ML	Heat Modelling	Cloud Datacenter Engineers	CH✻, 3D⋆

Table 1.1: Comparison of selected datacenter digital twins. **Modelling capability:** 3D = Visualizations; CH = Cooling/Heat, PE = Power/Energy Consumption, A = Anomaly Detection, N = Network Modelling, SE = Scenario Exploration, VP = Virtual Prototyping, FD = Federation, RA = Resource Allocation; **Data Analytics**: ✻ = Predictive Analysis; ⋆ = Descriptive Analysis, ‡ = Prescriptive Analysis.

RQ1: Literature Review II

VRIJE UNIVERSITEIT AMSTERDAM

A holistic DCDT system model

We propose a generic model of datacenter digital twinning that can be mapped to each system from **Table 1.1**. Within this model (see **Fig. 1.3**) we introduce a concept of the *Digital Thread*: a bridge between the DCDT and the physical DC equipment.

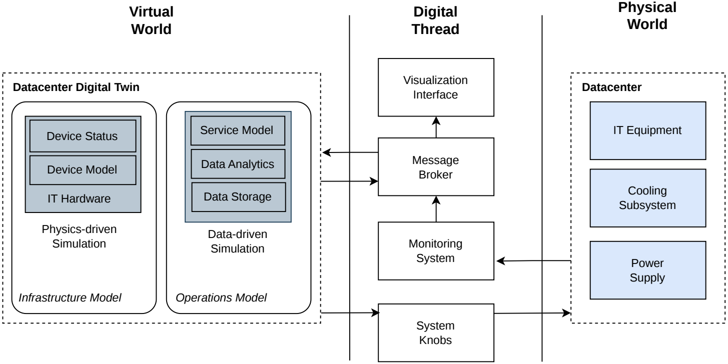

Figure 1.3: To answer **RQ1** we designed a generic datacenter digital twin system model based on a comprehensive literature review and findings from **Table 1.1**. The *Infrastructure Model* simulates the structure of the DC and the *Operations Model* simulates the behaviour of the DC.

RQ2: Reference Architecture

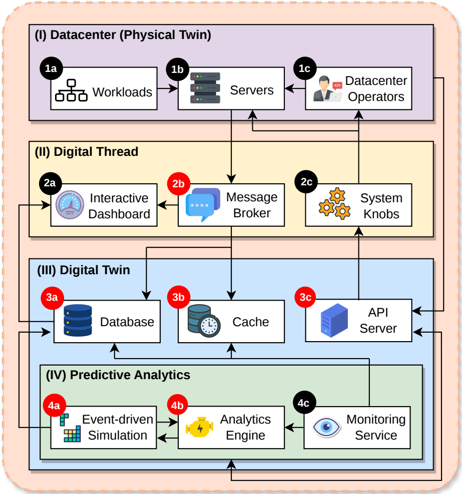

Figure 1.4: The predictive datacenter digital twin reference architecture. The architecture was designed with the *AtLarge Design Process* [7] over several iterations in the past months.

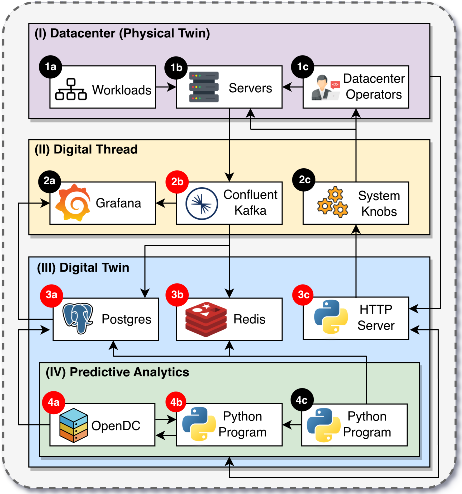

Figure 1.5: The prototype – *Sunfish*, and its components based on **Figure 1.4**. The time-series data flows first to the `Grafana` dashboard, `PostgreSQL` database and `Redis` cache [12].

RQ3: Experimental Setup

Problem

We cannot just go and test digital twins on large systems, because we do not have large systems at hand. Moreover, real-world experimentation is costly and unsustainable in the long run [9].

Solution

They way we test our reference architecture prototype is by using multiple simulators. We use an additional OpenDC process to play the role of a real datacenter.

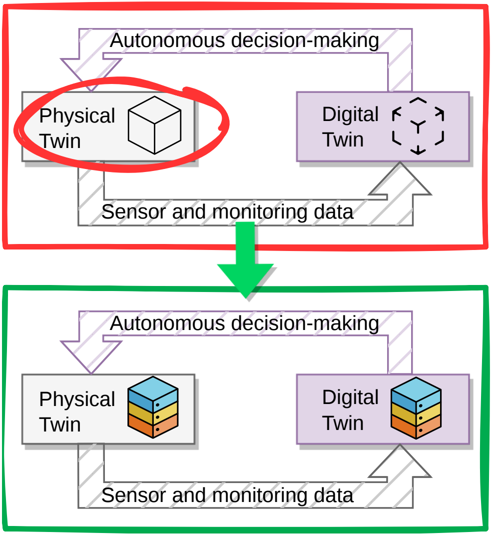

Figure 1.6: The experimental setup. Answering **RQ3** we provide a novel way to evaluate datacenter digital twins through discrete-event simulation.

Main Finding II

We posit digital twinning can be used for failure detection to the benefit of DC operators. We replicate an experiment from DyTwin [12] designed by Milojicic *et al.* to show our system can reliably detect *unexpected* host failures.

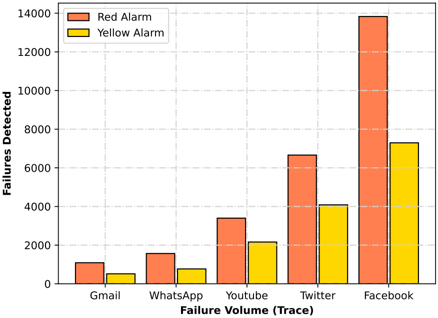

Figure 1.7: Experiment 1a. In this experiment we use red and yellow alarms to notify datacenter operators of unexpected failures. We use a threshold based on predictions done by the simulator and a statistical distribution.

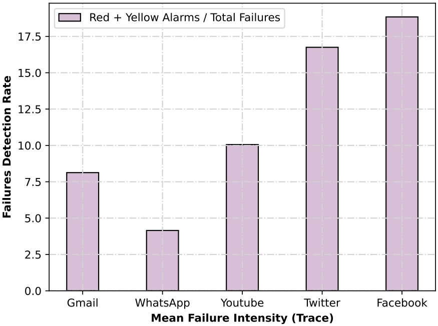

Figure 1.8: Experiment 1b. The mean failure detection rate is around 15%. Even though this seems low, if we look at **Fig. 1.9** (see Extra Slides), this simply means around 15% of failures are unexpected.

Key Takeaways

VRIJE UNIVERSITEIT AMSTERDAM

What is the societal context?

Datacenter manageability is a top-priority for the digital society. Over 3 million jobs in the Netherlands directly depend on cloud services, which are hosted in datacenters [6].

What problem did we solve?

DCDT's, still under development, lack crucial features such as predictive analytics to manage datacenters well. The entire DCDT design space remains largely unexplored.

How did we solve this problem?

Our contributions are: a thorough literature survey with a system model, a DCDT reference architecture, and prototype-based experiments via a novel evaluation method.

What did we find?

Sunfish can reliably detect unexpected failures based on discrete-event predictions, and can serve as a foundation for additional research and future work.

Extra Slides: References I

Jyotika Athavale, Cullen E. Bash, Wesley Brewer, Matthias Maiterth, Dejan S. Milojicic, Harry Petty, and Soumyendu Sarkar.
Digital twins for data centers.
Computer, 57(10):151–158, 2024.
URL https://doi.org/10.1109/MC.2024.3436945.

Wesley Brewer, Matthias Maiterth, Vineet Kumar, Rafal P. Wojda, Sedrick Bouknight, Jesse Hines, Woong Shin, Scott Greenwood, David Grant, Wesley Williams, and Feiyi Wang.
A digital twin framework for liquid-cooled supercomputers as demonstrated at exascale.
In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, SC 2024, Atlanta, GA, USA, November 17-22, 2024*, page 23. IEEE, 2024.
URL https://dl.acm.org/doi/10.1109/SC41406.2024.00029.

Zhiwei Cao, Ruihang Wang, Xin Zhou, and Yonggang Wen.
Reducio: model reduction for data center predictive digital twins via physics-guided machine learning.
In Jorge Ortiz, editor, *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys 2022, Boston, Massachusetts, November 9-10, 2022*, pages 1–10. ACM, 2022.
URL https://doi.org/10.1145/3563357.3564050.

Douglas Donnellan, Andy Lawrence, and Rose Weinshenk.
Executive summary: Annual outage analysis 2025, May 2025.
URL https://uptimeinstitute.com/resources/research-and-reports/annual-outage-analysis-2025.

Hanshu Hong, Qin Wu, Feng Dong, Wei Song, Ronghua Sun, Tao Han, Cheng Zhou, and Hongwei Yang.
Netgraph: An intelligent operated digital twin platform for data center networks.
In *NAI'21: Proceedings of the ACM SIGCOMM 2021 Workshop on Network-Application Integration, Virtual Event, USA, August 27, 2021*, pages 26–32. ACM, 2021.
URL https://doi.org/10.1145/3472727.3472802.

Extra Slides: References II

Alexandru Iosup, Fernando Kuipers, Ana Lucia Varbanescu, Paola Grosso, Animesh Trivedi, Jan S. Rellermeyer, Lin Wang, Alexandru Uta, and Francesco Regazzoni.
Future computer systems and networking research in the netherlands: A manifesto.
CoRR, abs/2206.03259, 2022.
URL `https://doi.org/10.48550/arXiv.2206.03259`.

Alexandru Iosup, Laurens Versluis, Animesh Trivedi, Erwin Van Eyk, Lucian Toader, Vincent van Beek, Giulia Frascaria, Ahmed Musaafir, and Sacheendra Talluri.
The atlarge vision on the design of distributed systems and ecosystems.
In *39th IEEE International Conference on Distributed Computing Systems, ICDCS 2019, Dallas, TX, USA, July 7-10, 2019*, pages 1765–1776. IEEE, 2019.
URL `https://doi.org/10.1109/ICDCS.2019.00175`.

Minghao Li, Ruihang Wang, Xin Zhou, Zhaomeng Zhu, Yonggang Wen, and Rui Tan.
Chattwin: Toward automated digital twin generation for data center via large language models.
In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys 2023, Istanbul, Turkey, November 15-16, 2023*, pages 208–211. ACM, 2023.
URL `https://doi.org/10.1145/3600100.3623719`.

Fabian Mastenbroek, Georgios Andreadis, Soufiane Jounaid, Wenchen Lai, Jacob Burley, Jaro Bosch, Erwin Van Eyk, Laurens Versluis, Vincent van Beek, and Alexandru Iosup.
Opendc 2.0: Convenient modeling and simulation of emerging technologies in cloud datacenters.
In Laurent Lefèvre, Stacy Patterson, Young Choon Lee, Haiying Shen, Shashikant Ilager, Mohammad Goudarzi, Adel Nadjaran Toosi, and Rajkumar Buyya, editors, *21st IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, CCGrid 2021, Melbourne, Australia, May 10-13, 2021*, pages 455–464. IEEE, 2021.
URL `https://doi.org/10.1109/CCGrid51090.2021.00055`.

Extra Slides: References III

Dejan S. Milojicic, Paolo Faraboschi, Nicolas Dubé, and Duncan Roweth.
Future of HPC: diversifying heterogeneity.
In *Design, Automation & Test in Europe Conference & Exhibition, DATE 2021, Grenoble, France, February 1-5, 2021*, pages 276–281. IEEE, 2021.
URL `https://doi.org/10.23919/DATE51398.2021.9474063`.

National Academy of Engineering, National Academies of Sciences Engineering, and Medicine.
Foundational research gaps and future directions for digital twins.
The National Academies Press, Washington, DC, 2024.
ISBN 978-0-309-70042-9.
URL `https://nap.nationalacademies.org/catalog/26894/foundational-research-gaps-and-future-directions-for-digital-twins`.

Ebad Taheri, Pedro Bruel, Pavana Prakash, Gourav Rattihalli, Ninad Hogade, Aditya Dhakal, Rolando P. Hong Enriquez, Torsten Wilde, Leo Popokh, Dejan S. Milojicic, and Cullen E. Bash.
Dytwin: Federated adaptive digital twins for data centers - visualization and anomaly detection.
In *SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis , Atlanta, GA, USA, November 17-22, 2024*, pages 847–852. IEEE, 2024.
URL `https://doi.org/10.1109/SCW63240.2024.00120`.

Sacheendra Talluri, Leon Overweel, Laurens Versluis, Animesh Trivedi, and Alexandru Iosup.
Empirical characterization of user reports about cloud failures.
In Esam El-Araby, Vana Kalogeraki, Danilo Pianini, Fr édéric Lassabe, Barry Porter, Sona Ghahremani, Ingrid Nunes, Mohamed Bakhouya, and Sven Tomforde, editors, *IEEE International Conference on Autonomic Computing and Self-Organizing Systems, ACSOS 2021, Washington, DC, USA, September 27 - Oct. 1, 2021*, pages 158–163. IEEE, 2021.
URL `https://doi.org/10.1109/ACSOS52086.2021.00039`.

Extra Slides: References IV

Fei Tao, Meng Zhang, Yushan Liu, and A.Y.C. Nee.
Digital twin driven prognostics and health management for complex equipment.
CIRP Annals, 67(1):169–172, 2018.
ISSN 0007-8506.
URL https://www.sciencedirect.com/science/article/pii/S0007850618300799.

Laurens Versluis, Mehmet Çetin, Caspar Greeven, Kristian Laursen, Damian Podareanu, Valeriu Codreanu, Alexandru Uta, and Alexandru Iosup.
Less is not more: We need rich datasets to explore.
Future Gener. Comput. Syst., 142:117–130, 2023.
URL https://doi.org/10.1016/j.future.2022.12.022.

Ruihang Wang, Xin Zhou, Linsen Dong, Yonggang Wen, Rui Tan, Li Chen, Guan Wang, and Feng Zeng.
Kalibre: Knowledge-based neural surrogate model calibration for data center digital twins.
In *BuildSys '20: The 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, Virtual Event, Japan, November 18-20, 2020*, pages 200–209. ACM, 2020.
URL https://doi.org/10.1145/3408308.3427982.

Ziting Zhang, Yu Zeng, Haoran Liu, Chaoyue Zhao, Feng Wang, and Yunqing Chen.
Smart DC: an AI and digital twin-based energy-saving solution for data centers.
In *2022 IEEE/IFIP Network Operations and Management Symposium, NOMS 2022, Budapest, Hungary, April 25-29, 2022*, pages 1–6. IEEE, 2022.
URL https://doi.org/10.1109/NOMS54207.2022.9789853.

Technical Setup

What is the simulation workload?

The compute workload is BitBrainsSmall. The failure traces include user reports from Gmail, WhatsApp and Twitter.

What is the experiment environment?

A commodity laptop: Framework Laptop 13, with 32GB of DDR5 RAM and an AMD Ryzen 7840U processor and an ArchLinux OS with Linux 7.0.13-arch1-1 kernel.

How did we adjust OpenDC (Physical Twin)?

We use a SURF [15] datacenter topology with 277 hosts. We wrote a custom Kotlin `ComputeMonitor` to export live-metrics into Kafka, and a custom Kotlin `HTTPClient` to talk to the digital twin. We add a new scheduling mechanism, the `SmartScheduler`.

Which metrics do we measure?

Timestamps, host names, uptime, downtime, CPU utilization *etc.*

Extra Slides: Why Digital Twinning?

Definition

A DCDT mirrors the structure, context and behaviour of a datacenter [1]. The prerequisite to any digital twin is good monitoring and sensing capabilities in the physical entity. Datacenters meet this requirement easily because they already connect hundreds of monitoring sensors.

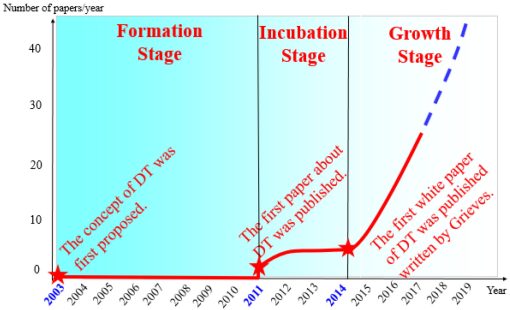

Figure E.2: Due to insufficient technological foundations, little work is available on DTs between 2003 and 2018, and it is only with the rapid growth of cloud computing, Internet-of-Things and Big Data analytics that DTs have reemerged [14]. That is why nobody used digital twins to mirror datacenters earlier.

Extra Slides: Why not pure simulation?

Predicting job failures

Preventing failure-caused outages in advance can reduce huge operational costs, as over 20% of all reported outages amount to more than 1 million US$ [4]. Only a constant bi-directional interaction (digital twin $\Longleftrightarrow$ physical entity) can achieve this.

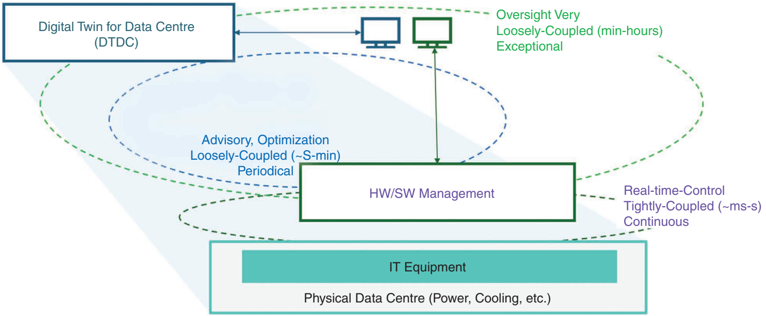

Figure E.3: Real-time control that is tightly-coupled with the IT equipment is a prerequisite for timely predictions within seconds/minutes [1].