Distributed Modular Digital Twin Network for High-Performance and Reliable Data Centers

Yan Chen, Xing Lu, Cary Faulkner, Alex Vlachokostas, Hanlong Wan, Jeremy Lerond
Pacific Northwest National Laboratory
Richland, WA, USA
{yan.chen,xing.lu,cary.faulkner,alex.vlachokostas,hanlong.wan,jeremy.lerond}@pnnl.gov

ABSTRACT

High-density computing workloads are driving rack power densities beyond 100 kW per rack, creating unprecedented pressure on cooling and power systems in high-performance computing (HPC) data centers. Conventional CFD-based models remain the gold standard for data center *design optimization* but are too computationally expensive and rigid for *operational-stage use*. We present the first physics-constrained Distributed Modular Digital Twin Network (DMDTN), where each subsystem (cooling, power, IT load, and controls) is represented as a surrogate-driven module, interconnected through conservation laws and coordinated via a distributed message bus. This design ensures physical consistency while supporting real-time adaptability and scalability.

Evaluation on synthetic datasets shows that DMDTN supports accurate thermal and power prediction, performance evaluation, and fault diagnostics during operations. Compared to a monolithic model, DMDTN achieved ~60% lower error (RMSE 172 vs. 450), 30% faster training (497 vs. 704 seconds), and greater robustness under stress. These results demonstrate DMDTN as a practical complement to CFD: while CFD is superior for design and optimization, DMDTN enables real-time monitoring, prediction, and resilience in the operational stage. In the poster session, we present the system architecture, validation workflow, and comparative results.

KEYWORDS

Digital Twin, Data Center, Modular Modeling, Surrogates, Energy Efficiency, Cooling, HPC

ACM Reference Format:
Yan Chen, Xing Lu, Cary Faulkner, Alex Vlachokostas, Hanlong Wan, Jeremy Lerond. 2025. Distributed Modular Digital Twin Network for High-Performance and Reliable Data Centers. In . ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

The rise of HPC workloads is pushing rack power densities beyond 100 kW, creating severe stress on cooling and power systems. Data centers already account for 2–3% of global electricity use, with cooling consuming up to 40% of operating costs. Ensuring efficiency and resilience is critical.

Computational fluid dynamics (CFD) tools remain the industry standard for design and optimization, where they provide high-fidelity modeling of airflow and thermal interactions. However, CFD is computationally expensive, inflexible under changing conditions, and not suitable for real-time control. In contrast, operational-stage management requires models that are lightweight, adaptive, and capable of supporting real-time prediction, evaluation, and fault detection.

2 APPROACH

We propose the Distributed Modular Digital Twin Network (DMDTN) to meet operational needs. In this framework, each subsystem of the data center is represented as a surrogate model, such as cooling towers, chillers, pumps, CDUs, PDUs, IT racks, UPS/battery systems, and controls. These surrogates capture subsystem dynamics without requiring detailed proprietary models.

Modules are interconnected through physics-based constraints that enforce conservation of mass, energy, and charge across the facility. A distributed message bus coordinates inter-module communication. This hybrid approach preserves physical consistency while delivering flexibility and scalability. New or upgraded equipment can be incorporated seamlessly without retraining the entire system.

3 NOVEL CONTRIBUTION

The novelty of this work lies in introducing the first modular surrogate twin framework with explicit physics-constrained coupling for HPC data centers. Unlike monolithic surrogate models, DMDTN localizes errors within subsystems, improves robustness, and enhances scalability. Importantly, it is designed for operational use, complementing CFD by addressing the real-time stage of monitoring, prediction, and diagnostics.

4 EXPERIMENTAL SETUP

We evaluated DMDTN using synthetic datasets representing interactions of IT load, cooling loops, and power distribution systems. Minute-level data from a data center facility was used to generate the dataset, which was then split into 75% for model development and 25% for testing. Use case demonstrations included thermal & power prediction, capacity analysis, fault diagnostics, and storage/retrofit recommendation.

5 RESULTS

DMDTN achieved strong quantitative improvements compared to a monolithic model:

Conference'17, July 2017, Washington, DC, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
https://doi.org/10.1145/nnnnnnn.nnnnnnn

- ∼60% lower error (RMSE 172 vs. 450).
- 30% faster training (497 vs. 704 seconds).
- Comparable R^2 with improved robustness under stress conditions.

Thermal and power predictions closely tracked ground truth, enabling capacity analysis and anomaly detection. Fault events were identified reliably, supporting diagnostics. These results validate DMDTN as an effective operational twin: fast, accurate, and resilient.

6 STATUS OF WORK

This research is in the early evaluation stage. Synthetic datasets confirm feasibility, and next steps include validation with real facility data, higher-fidelity module development, and integration with real-time control systems. Future extensions will explore fleet-level coordination and adaptive online learning.

7 POSTER PRESENTATION

The poster will present the DMDTN architecture, surrogate modules, and physics-based coupling. Comparative evaluation with monolithic models will be shown through scatter plots, time-series predictions, and error metrics. A conceptual figure will contrast CFD-based design twins with modular operational twins, highlighting the complementary roles.

8 CONCLUSION

The Distributed Modular Digital Twin Network offers a scalable, physics-constrained alternative to monolithic surrogate models for the operational stage of data center management. While CFD remains superior for design optimization, DMDTN enables accurate performance evaluation, real-time load prediction, and fault detection. Results show lower error, faster training, and stronger robustness, positioning DMDTN as a foundation for resilient and sustainable HPC operations. Beyond HPC, this framework generalizes to smart grids, industrial plants, and district energy systems, where modular real-time twins are equally critical.

ACKNOWLEDGMENTS

This work was supported by Pacific Northwest National Laboratory under the U.S. Department of Energy.

The authors acknowledge the use of AI-assisted text generation (OpenAI ChatGPT, GPT-5) for writing support in preparing portions of the abstract and summary text. All scientific content, claims, and conclusions are the sole responsibility of the authors. LLMs were not used for data analysis, figure generation, or scientific results, and are not eligible for authorship.