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

Yan Chen, Xing Lu, Cary Faulkner, Alex Vlachokostas, Hanlong Wan, Jeremy Lerond

Primary Challenge: Data Center Energy & Cooling Constraints

Global data center demand is rising quickly, driven by high-performance computing workloads that can push rack power densities above 100 kW/rack. This rapid growth places increasing stress on cooling and power infrastructure, which must achieve high efficiency, resilience, and scalability under dynamic operating conditions.

The industry standard for thermal design, CFD-based simulation, provides accuracy but remains too computationally intensive for real-time operational use. It is difficult to adapt quickly to scaling or retrofit scenarios, leaving operators without flexible tools to manage capacity and efficiency in live facilities.

As shown in Figure ***1*** below, the challenge spans multiple domains. Data centers operate at the intersection of electric grids, cooling plants, storage systems, and IT workloads. Electricity supply, thermal storage, and cooling loops interact continuously to meet demand. Effective management requires an integrated framework that couples power, cooling, and IT loads into a coordinated system.

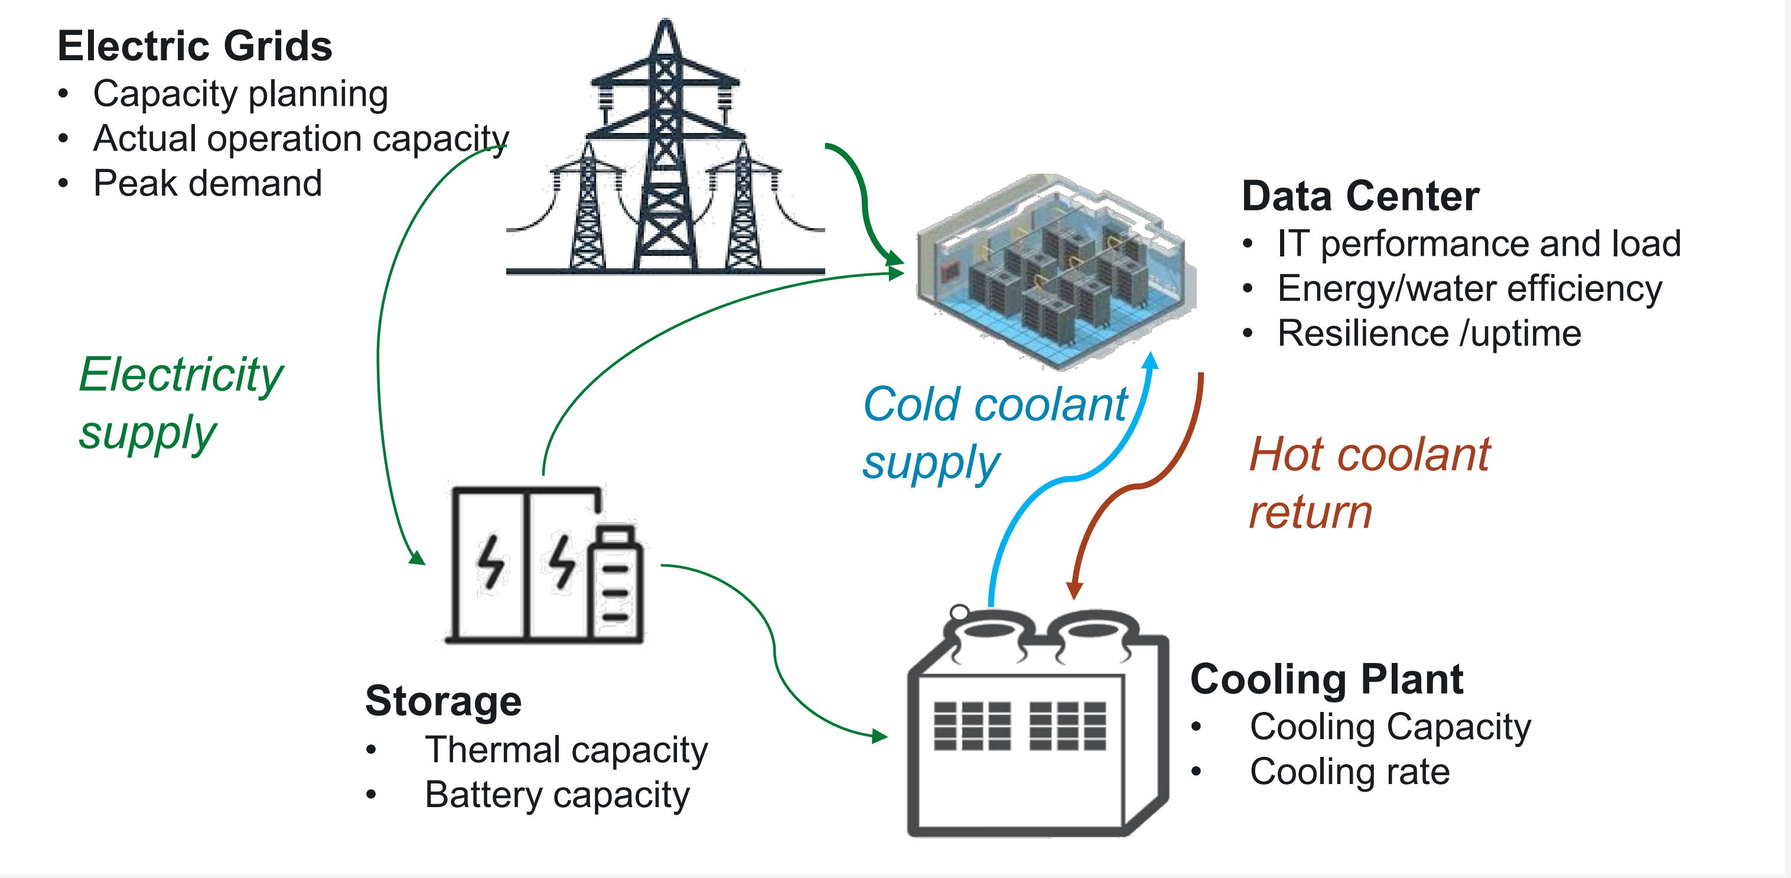

Electric Grids
- Capacity planning
- Actual operation capacity
- Peak demand

Data Center
- IT performance and load
- Energy/water efficiency
- Resilience /uptime

Storage
- Thermal capacity
- Battery capacity

Cooling Plant
- Cooling Capacity
- Cooling rate

Figure 1 *Energy and cooling ecosystem linking grids, storage, cooling plants, and data center demand.*

Objective and Impact

Our objective is to develop a modular digital twin approach that enables real-time, scalable management of data center power, cooling, and IT workloads. The framework links simplified models with physical principles to deliver accurate insights under changing operating conditions. This matters because it enables data centers to:

- **Predict and optimize** thermal and power behavior under dynamic workloads.
- **Enhance efficiency** in energy and water use while maintaining uptime.
- **Detect faults early and evaluate retrofit options** with reduced risk.
- **Strengthen resilience** against stress and uncertainty at scale.

Through this work, we move toward smarter, sustainable, and more reliable digital infrastructure.

PNNL is operated by Battelle for the U.S. Department of Energy

For additional information, contact:
Yan Chen | yan.chen@pnnl.gov

Our Approach: Distributed Modular Digital Twin Network (DMDTN)

This is the first demonstration of a modular, physics-constrained surrogate twin framework for HPC data centers. We represent the data center as a network of surrogate-driven modules interconnected through physics-based constraints as shown in ***Fig. 2***. Each subsystem is modeled as a black-box surrogate, while conservation laws enforce global consistency (mass, energy, and charge balances). This hybrid structure provides flexibility, scalability, and physical rigor, ensuring fidelity without exposing proprietary modeling details.

$$\dot{\boldsymbol{x}}(\boldsymbol{t}) = \boldsymbol{S}(\boldsymbol{x}(\boldsymbol{t}),\boldsymbol{u}(\boldsymbol{t}),\boldsymbol{d}(\boldsymbol{t})), \mathbf{0} = \begin{bmatrix} \{\boldsymbol{\Psi}_{e}(\boldsymbol{p}_{i}(\boldsymbol{t}),\boldsymbol{p}_{j}(\boldsymbol{t}),\boldsymbol{q}_{e}(\boldsymbol{t}))\}_{e\in\varepsilon} \\ \boldsymbol{Bq}(\boldsymbol{t}) - \boldsymbol{d}(\boldsymbol{x}(\boldsymbol{t}),\boldsymbol{u}(\boldsymbol{t})) \end{bmatrix}$$

Here, $x(t)$ denotes the system states, $u(t)$ the control inputs, and $d(t)$ the external disturbances. The surrogate dynamics are represented by $S(\cdot)$, while Ψe captures the physics of interconnections such as airflow, power transfer, or fluid transport. Network-level conservation is enforced through the operator Bq, which ensures mass, energy, and charge balance across the system.

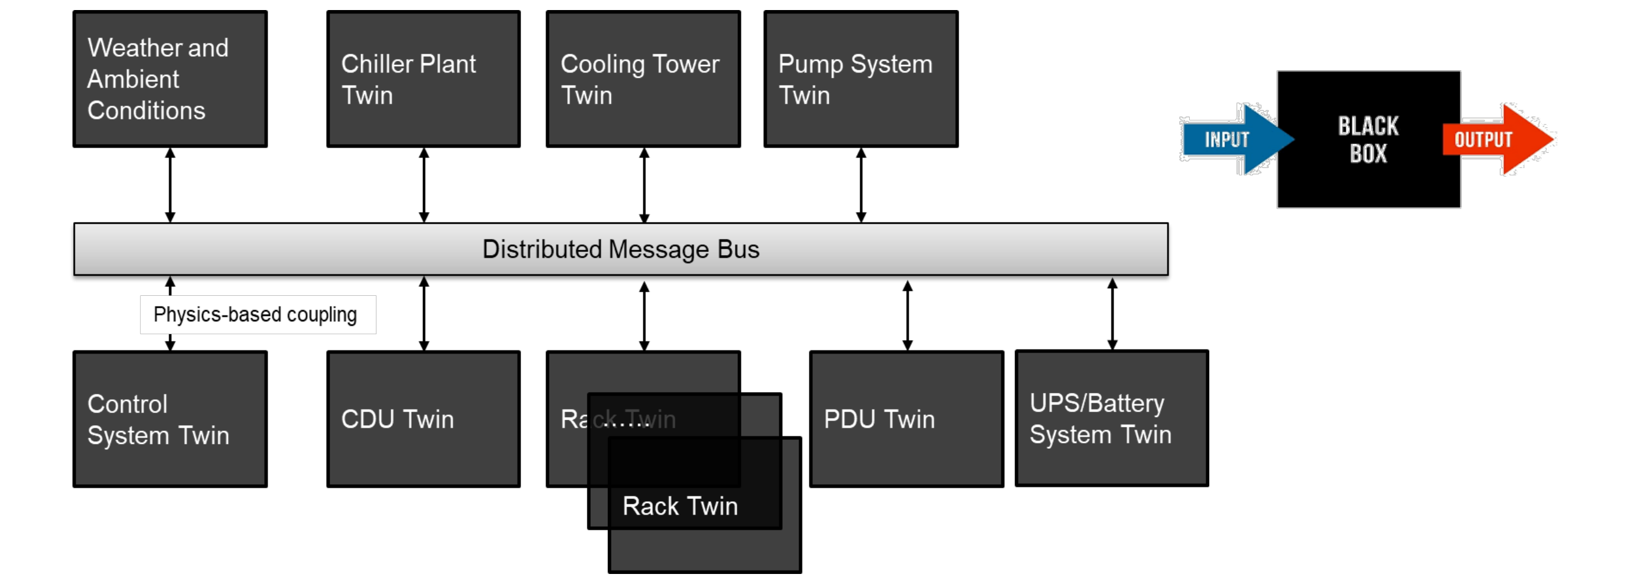

Figure 2: *Distributed Modular Digital Twin Network (DMDTN) concept. Each subsystem is modeled as a black-box surrogate, coupled by physics, and coordinated via a distributed message bus.*

Key Features

- Fully data-driven framework trained on operational data for real-time inference & control.
- Modular twin design for each subsystem/equipment, such as *IT equipment & workloads, rack, fan, cooling towers, chillers, pumps, CDUs, Power distribution units & UPS/battery storage, and controls.*
- Distributed intelligence with a message bus for subsystem/equipment coordination.

Advantages and Validation

Our approach eliminates reliance on heavy CFD models and delivers a more adaptive, scalable solution. Validation is built into the workflow to ensure accuracy and resilience before live deployment.

- **Fast and stable training** with lightweight modular components
- **Rapid adaptability** to changing loads, equipment, and climate conditions
- **Self-calibration** with continuous data streams
- **Supports retrofits** and multi-site fleet management
- **Simulation-based testing** (e.g., Modelica) used to:
  - Generate synthetic data for rare-event training
  - Validate predictions before deployment
  - Stress-test algorithms under extreme scenarios

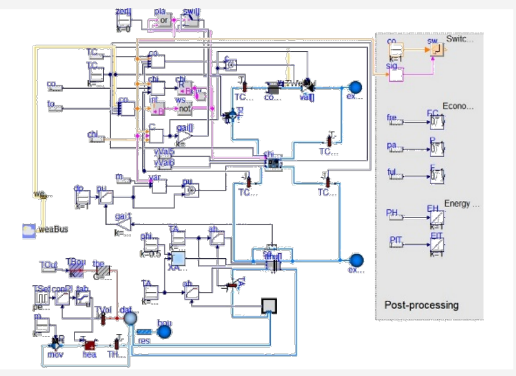

Case Study: Test Case Design and Performance Evaluation

To demonstrate feasibility, we applied the framework to a synthetic test case representative of high-performance computing (HPC) chilled-door cooling (***Fig. 3***) and power dynamics. 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.

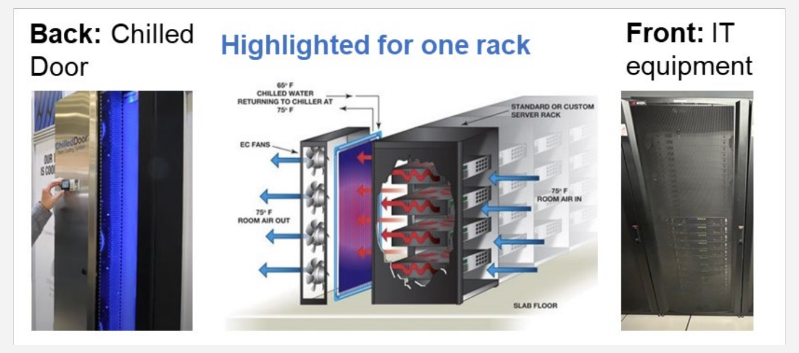

Figure 3: *Chilled Door cooling configuration*

Use Cases Demonstrations

- **Thermal & power prediction:** Accurate prediction of load and cooling trends (***Fig. 4***)
- **Fault diagnostics:** Early detection of anomalies under stress (***Fig. 5***)
- **Storage/retrofit** recommendation (***Fig. 6***)
- **Capacity analysis:** Assess headroom vs. CPU and cooling design limits

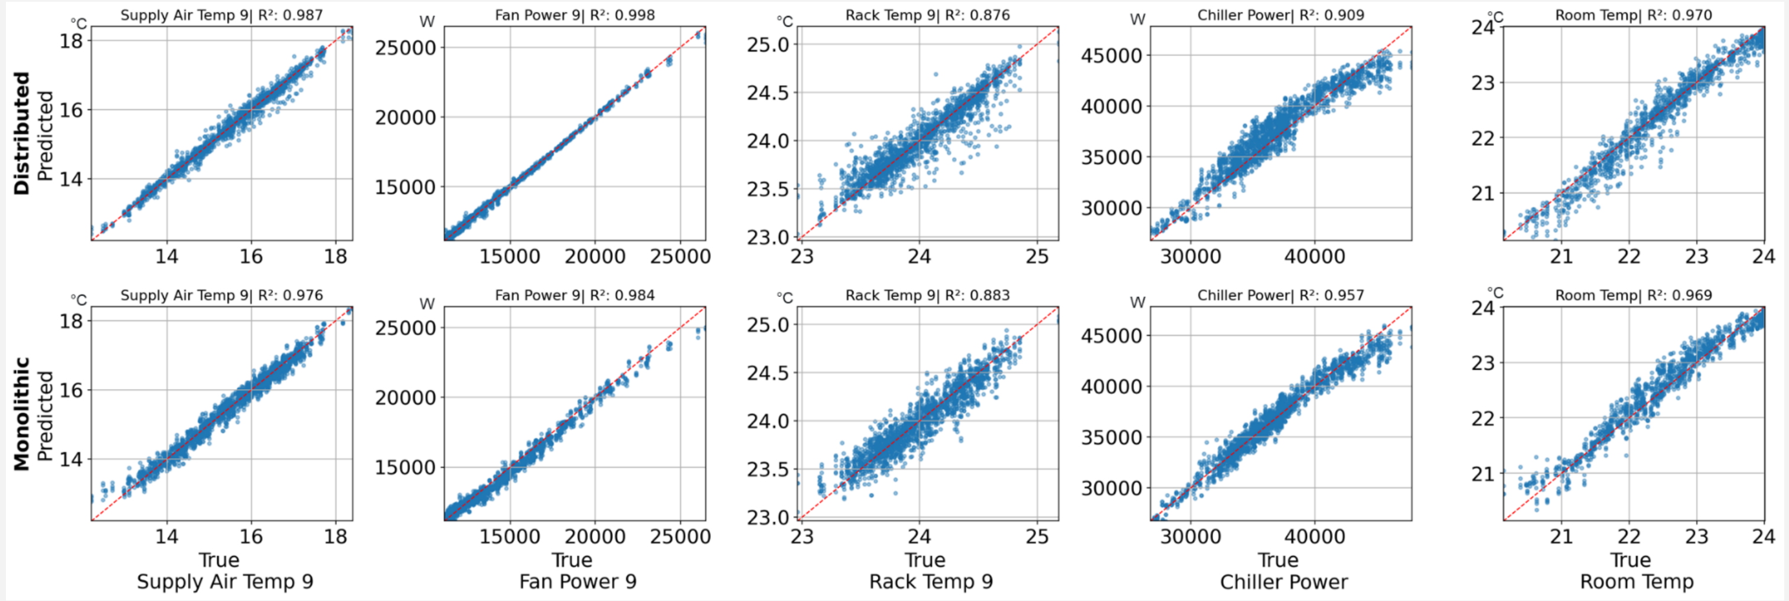

Figure 4: *Thermal and power prediction comparison (selective)*

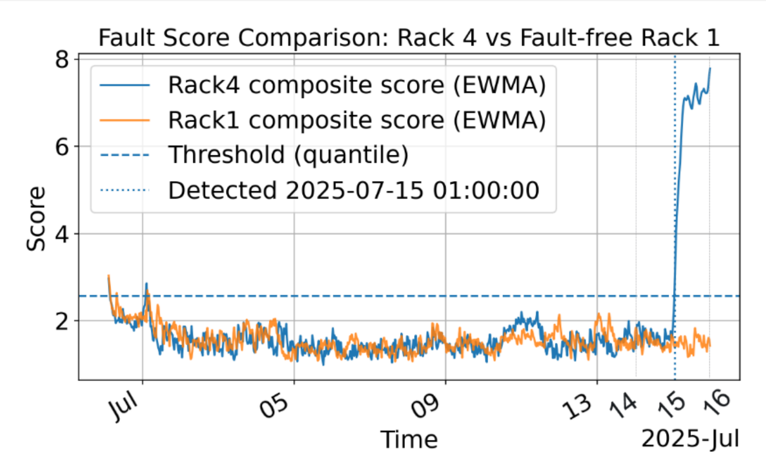

Figure 5: *Fault diagnostics example with stuck valve*

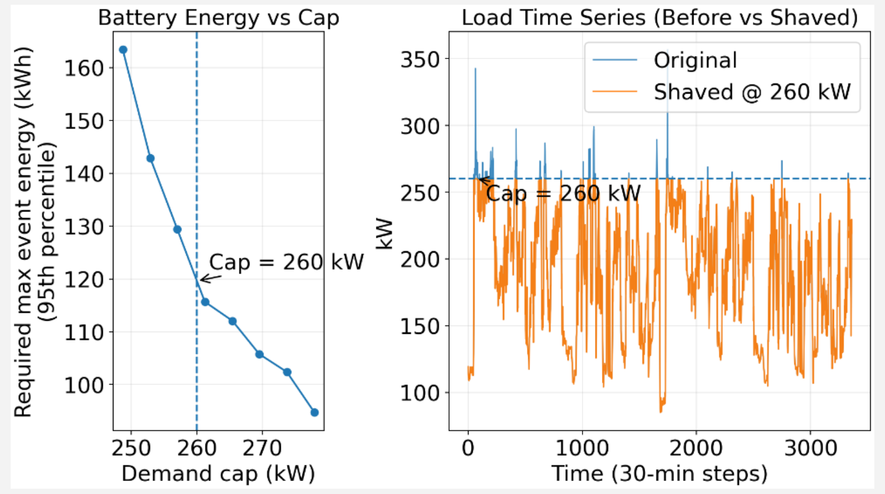

Figure 6: *Battery sizing for demand flexibility*

Distributed vs. Monolithic Evaluation

Compared to a single complex monolithic model, the distributed modular digital twin approach achieved lower error (RMSE 172 vs. 450), and faster training (497s vs. 704s) with comparable R².

Approach	# of models	Mean RMSE	Mean R2	Training time (sec)
Distributed	42	172	0.84	497
Monolithic	1	450	0.87	704

Conclusion

- First physics-constrained modular digital twin network for HPC data centers
- Achieved ~60% lower error and 30% faster training than a monolithic model
- Enabled accurate prediction, capacity analysis, and fault diagnostics
- Robust under stress, maintaining stable performance under workload shifts
- Paves the way for scalable, real-time data center management

Acknowledgment*: Portions of the text were prepared with assistance from an AI language model (OpenAI ChatGPT, GPT-5). The authors retain full responsibility for all content.*

Pacific Northwest
NATIONAL LABORATORY