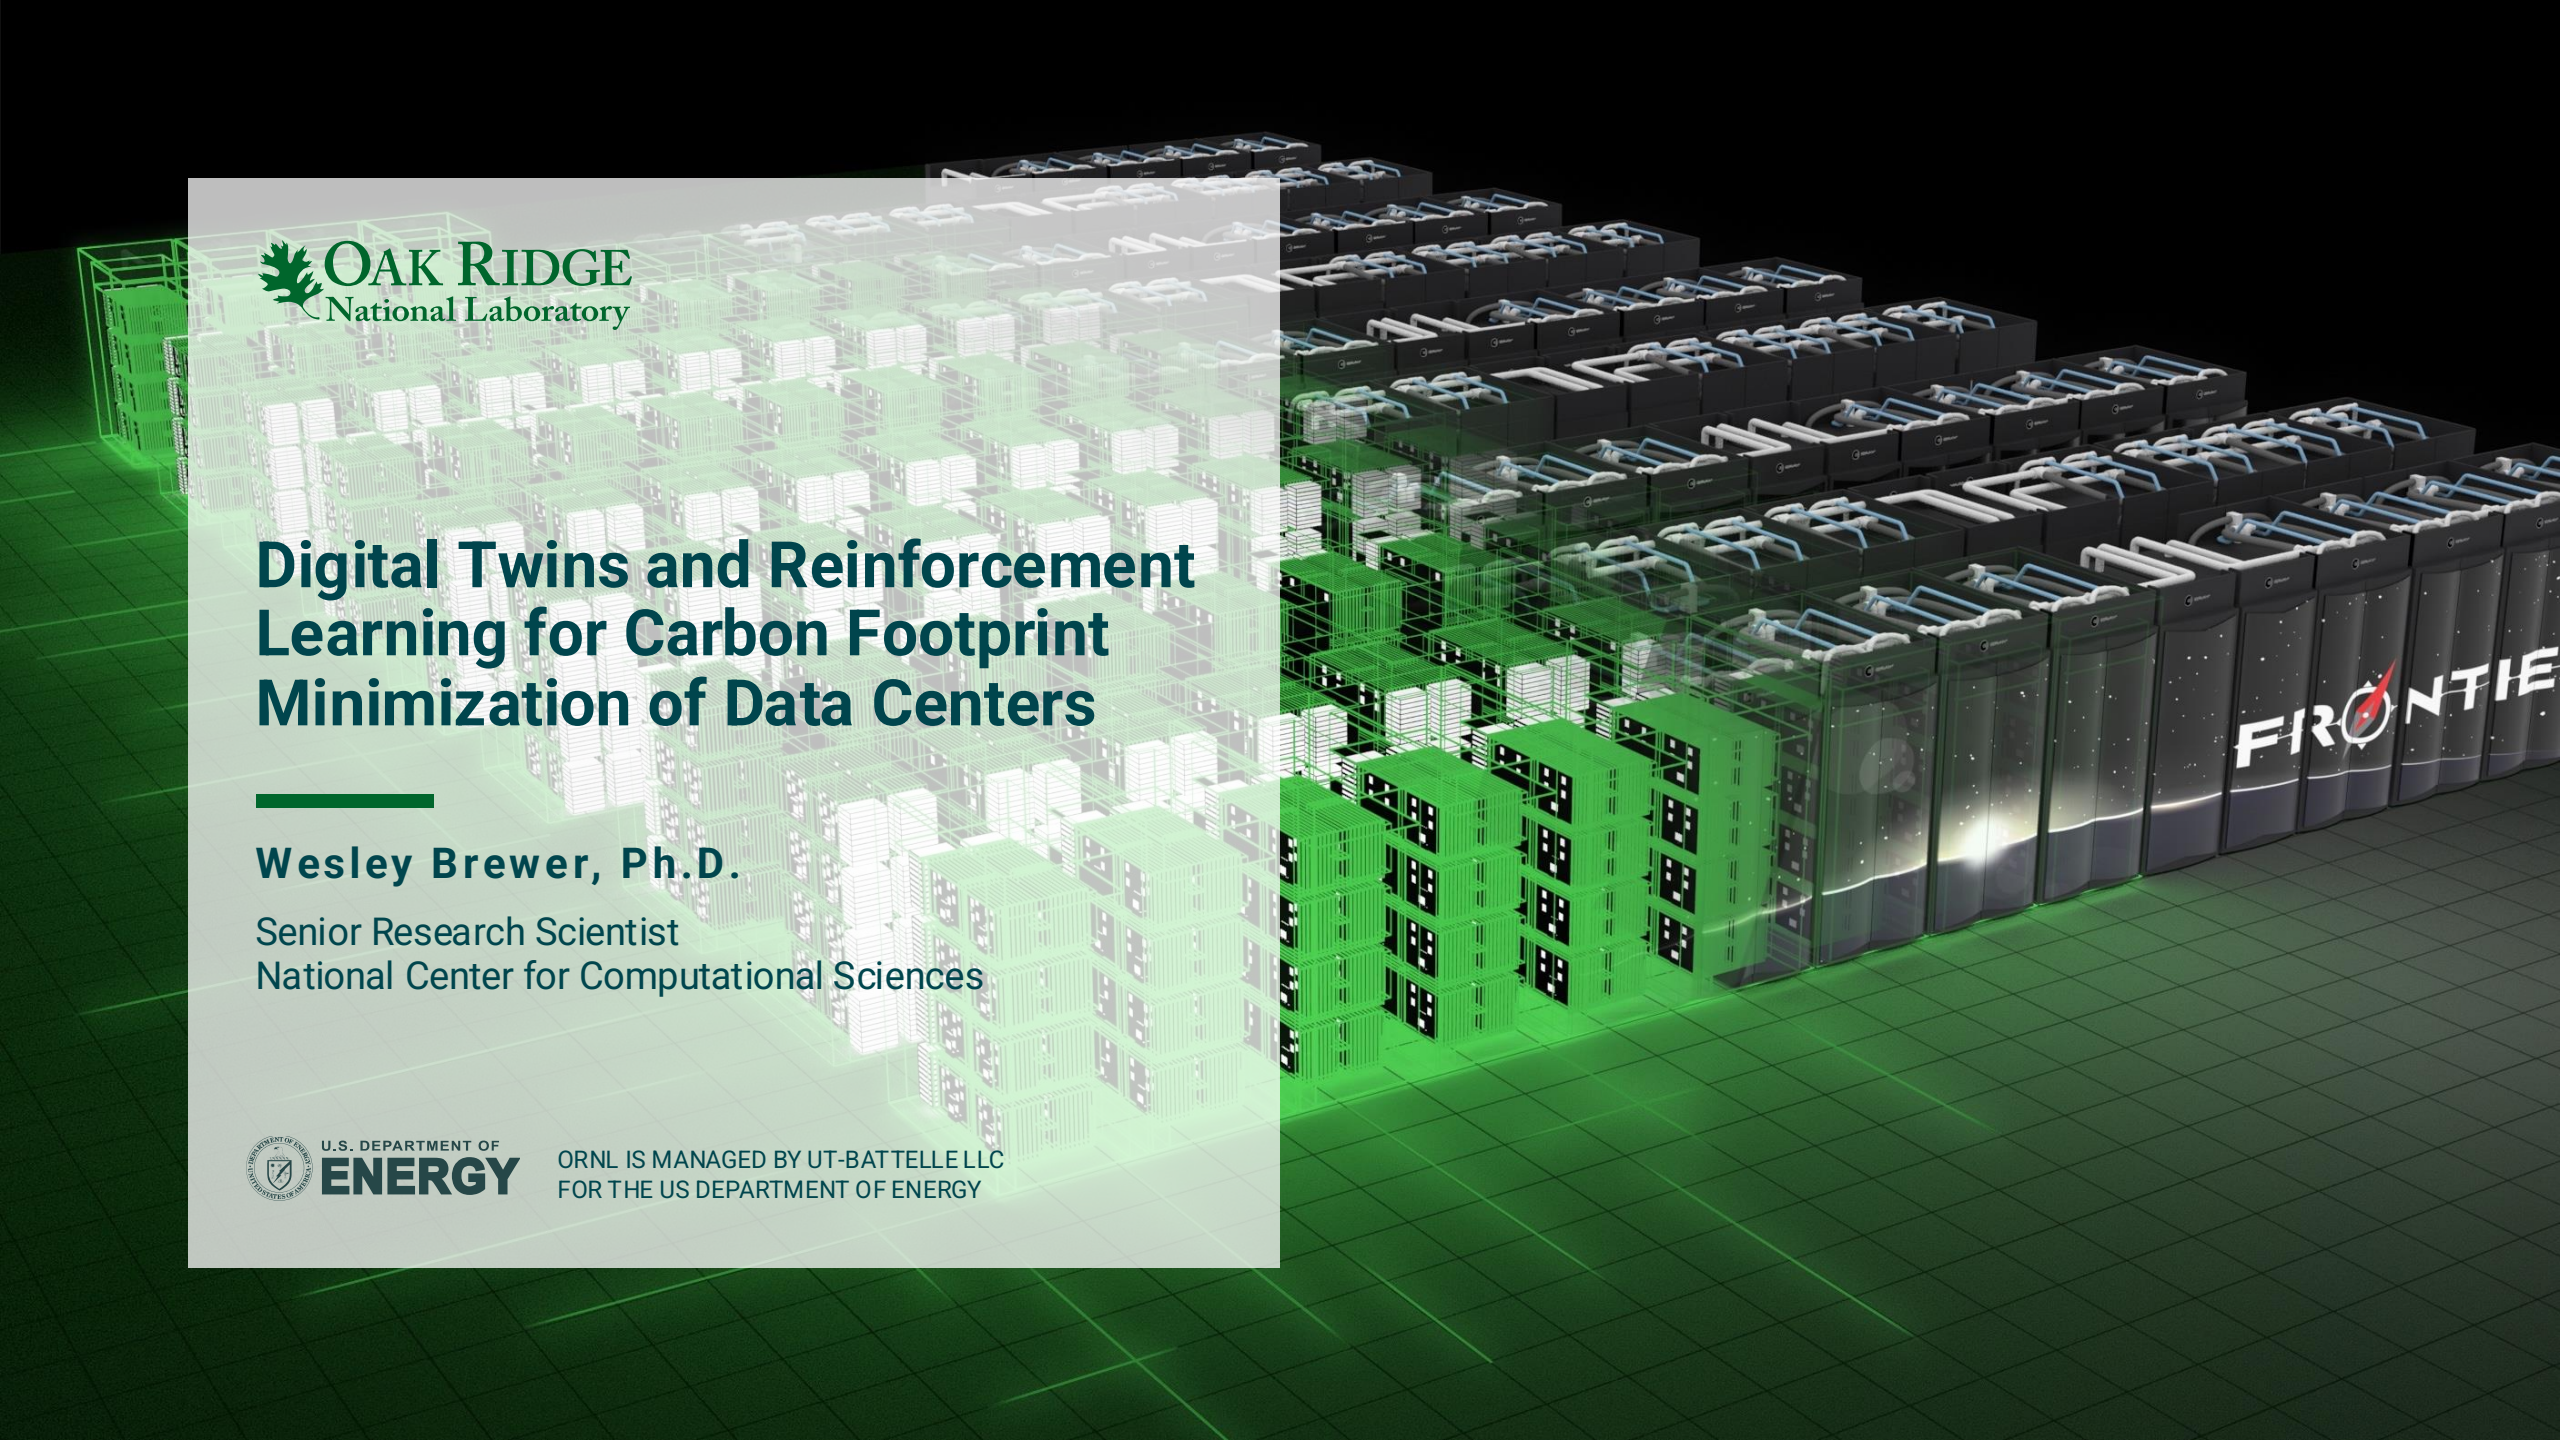

OAK RIDGE National Laboratory

Digital Twins and Reinforcement Learning for Carbon Footprint Minimization of Data Centers

Wesley Brewer, Ph.D.

Senior Research Scientist
National Center for Computational Sciences

U.S. DEPARTMENT OF ENERGY

ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE US DEPARTMENT OF ENERGY

Digital Twin / Carbon Optimization Architecture

A digital twin is a set of virtual information constructs that

mimics the structure, context, and behavior of an individual/unique physical asset,

is dynamically updated with data from its physical twin throughout its lifecycle, and

informs decisions that realize value.

AIAA Digital Engineering Integration Committee (2020)

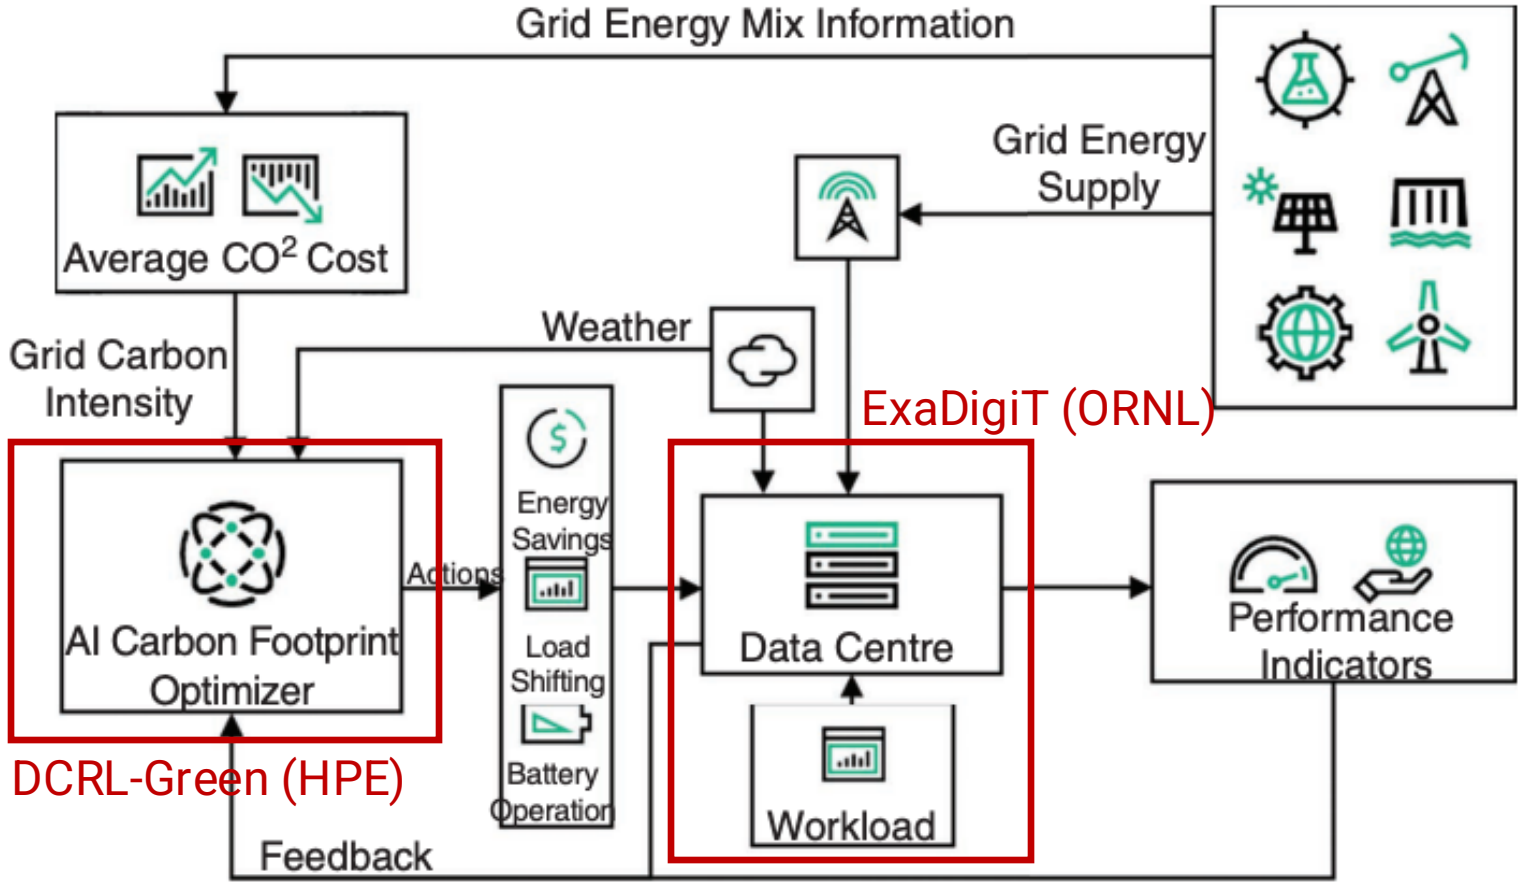

Athavale et al. (2024)

Workloads

Power profile landscape of workloads

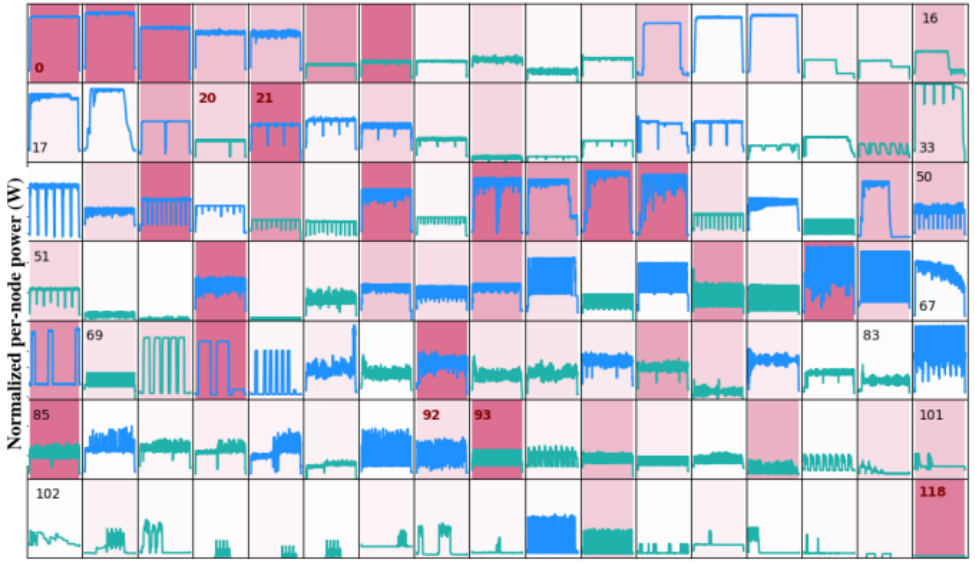

Karimi et al. (2024)

EXADIGIT

Power Model

Dynamically predicts total system power, energy conversion losses.

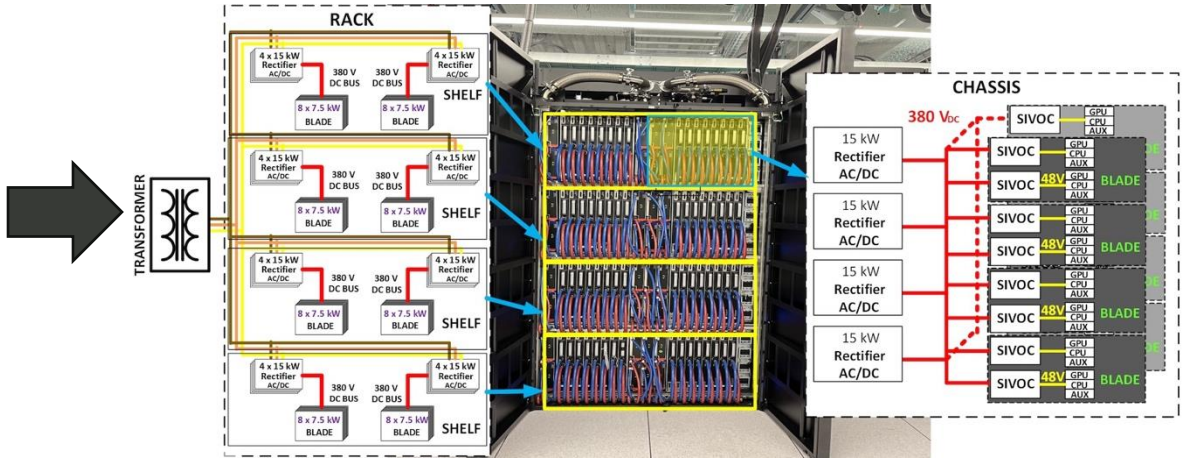

System Power

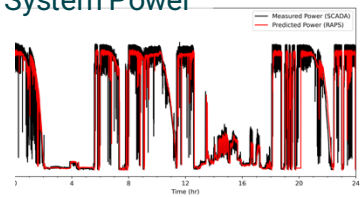

Brewer et al. (2024)

Cooling Model

Predicts pressures, temperatures, flow rates at 87 different locations...

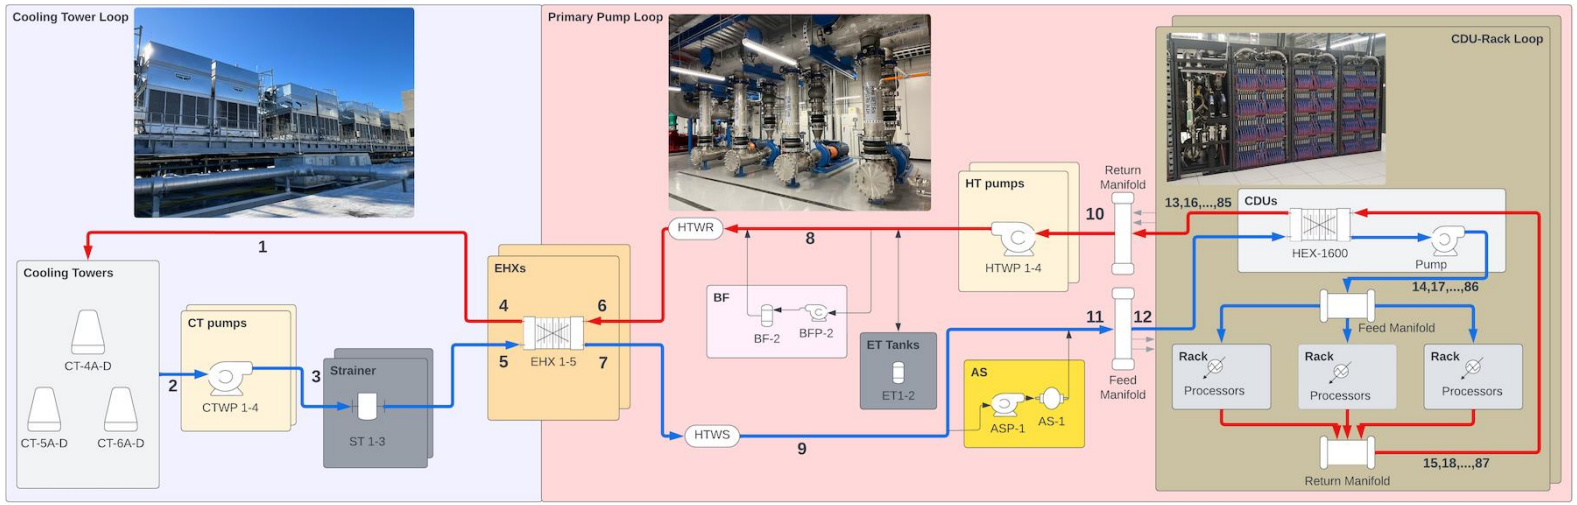

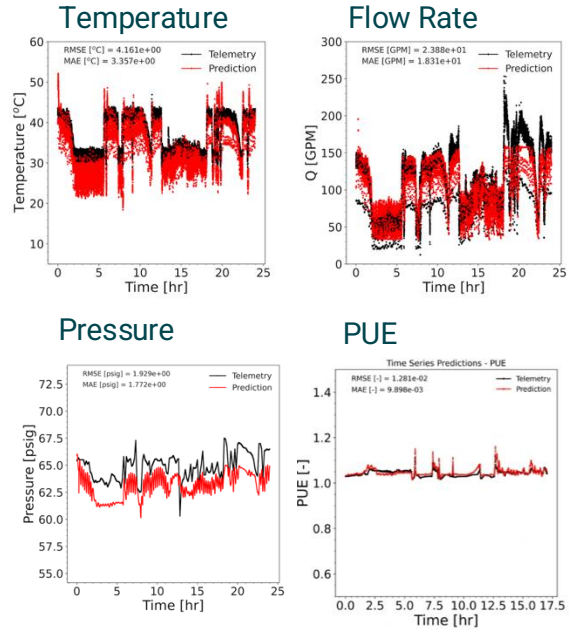

ExaDigiT can be used to create a virtual datacenter

Augmented Reality

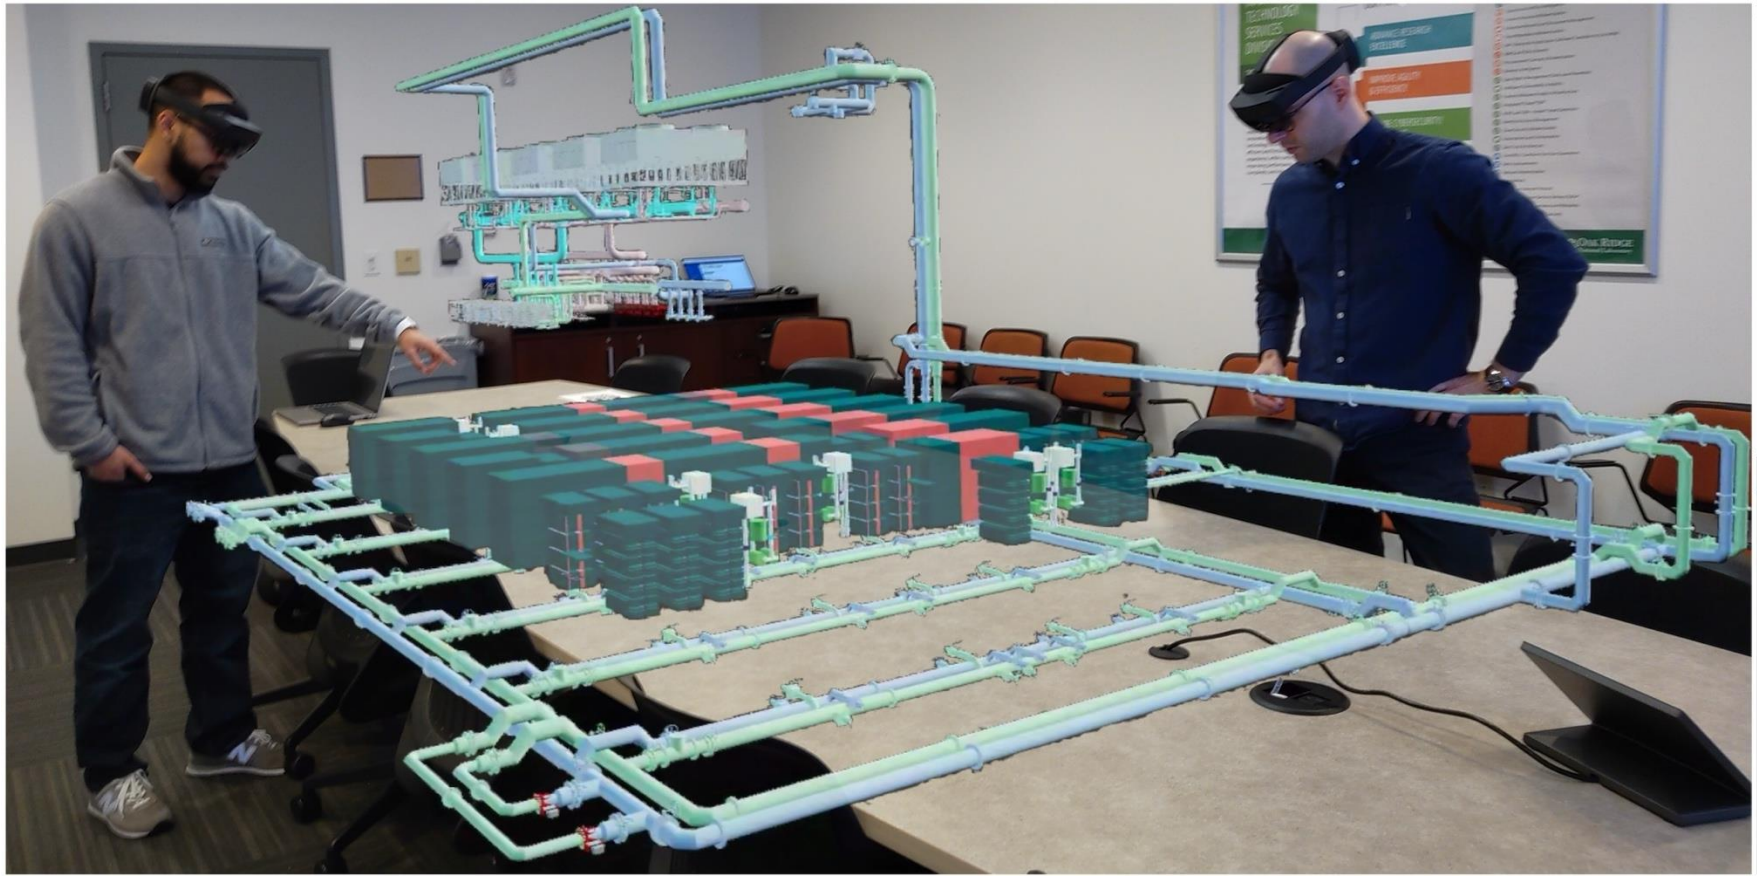

Console Interface

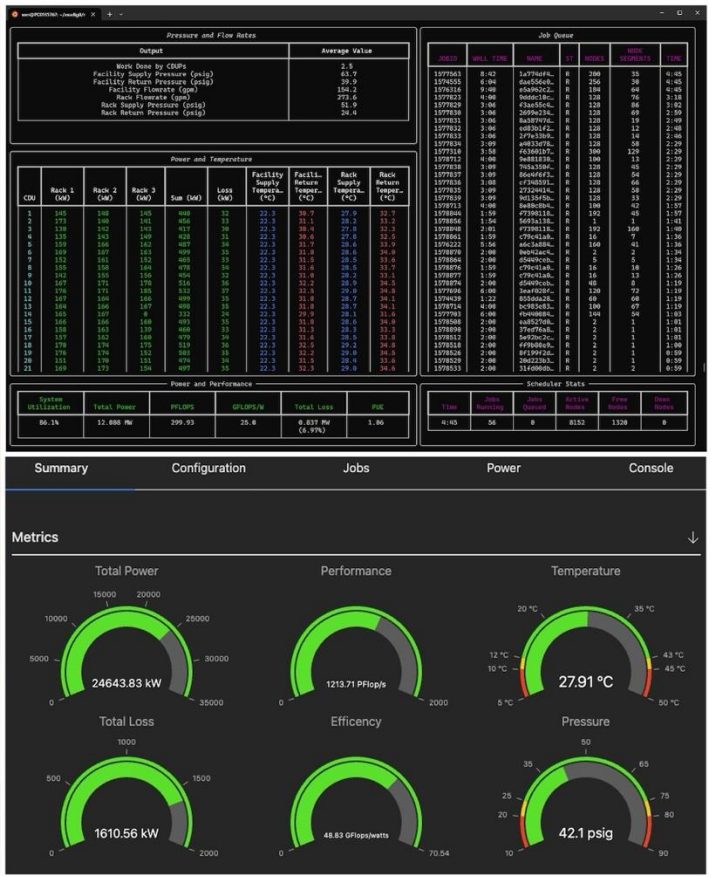

Dashboard interface

System-level insights

DAILY STATISTICS OF DT FROM TELEMETRY REPLAY OF 183 DAYS.

Parameter	Min	Avg	Max	Std
Avg Arrival Rate, $t_{avg}(s)$	17	138	2988	331
Avg Nodes per Job	39	268	5441	626
Avg Runtime (m)	17	39	101	14
Jobs Completed	32	1575	5157	1171
Throughput (jobs/hr)	1.3	66	215	49
Avg Power (MW)	10.2	16.9	23.0	2.4
Loss (MW)	0.52	1.14	1.84	0.15
Loss (%)	6.26	6.74	8.36	0.11
Total Energy Consumed (MW-hr)	129	405	553	64
Carbon Emissions (tons CO_2)	53	168	229	26

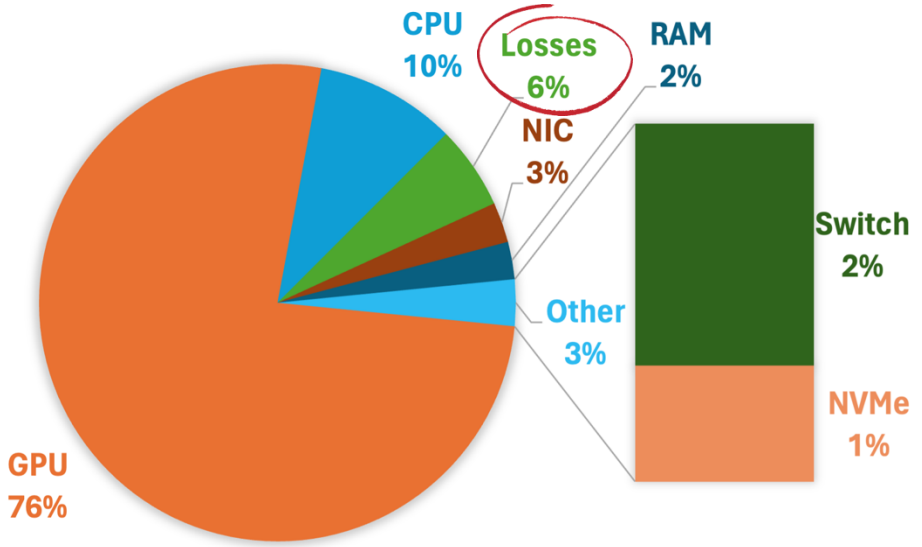

Peak-power breakdown

~$1M per year in losses!

Considerable losses are incurred during both AC-DC rectification and DC-DC voltage conversion.

Exploring energy-efficient solutions

"Smart-load-sharing" rectifiers	Unnecessary rectifiers can be put in standby mode when compute nodes are idle, leading to $120k per year savings for Frontier
DC-direct power	Providing 380V DC-direct power increased the system efficiency from 93.3% to 97.3%, a potential savings of $542k per year.

Carbon Footprint Optimization via Reinforcement Learning (RL)

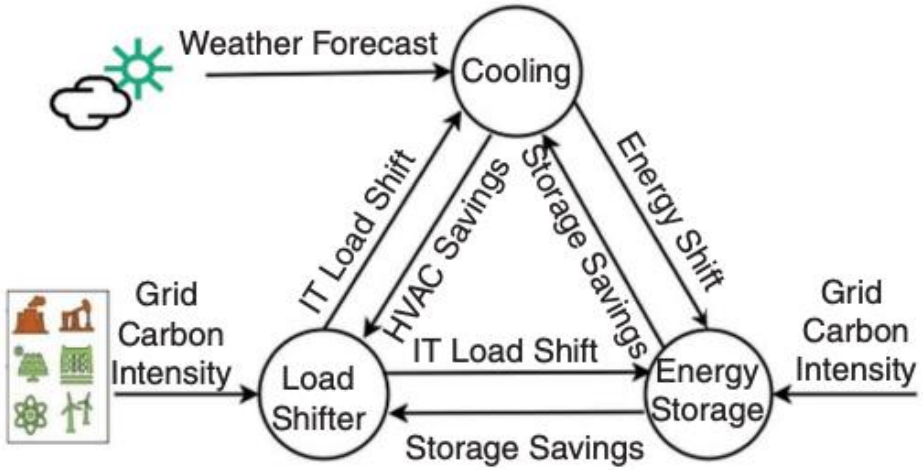

Internal and External Dependencies for the Cooling, Load Shifting, and Battery Agents

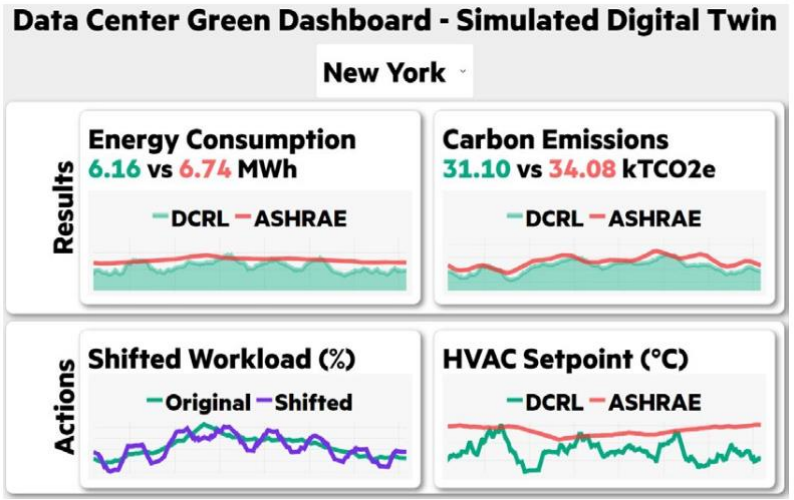

RL seeks to learn a policy which maps states to actions:

$$a = \pi(s)$$

Single-agent RL on HVAC setpoint control yields a 7% reduction in CO_2 vs. ASHRAE G36 baseline; multi-agent RL yields a 13% reduction.

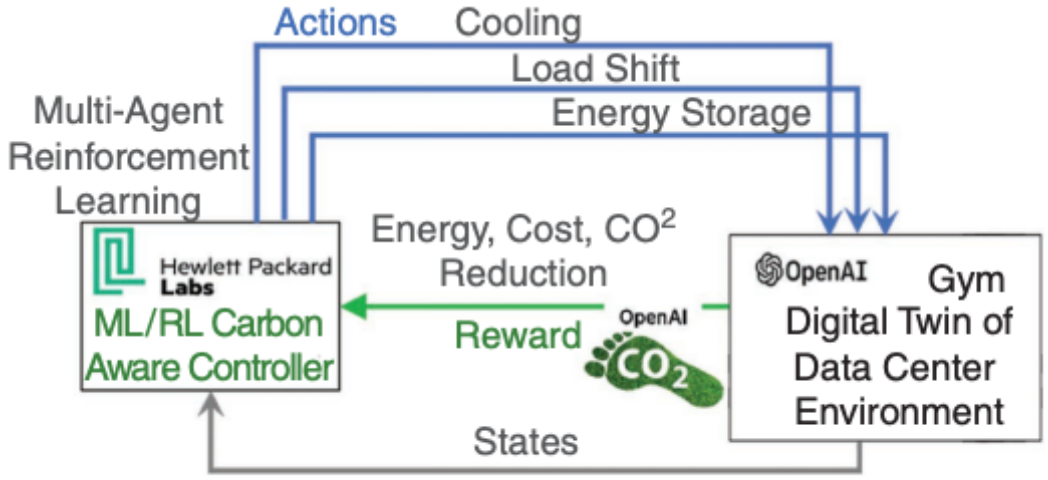

Multi-Agent Reinforcement Learning Controller

Liquid cooling offers significantly higher cooling capacity per unit volume and can reduce cooling energy and carbon footprint by up to 63%.

Khalaj and Halgamuge (2017)

Sarkar et al. (2024)

EXADIGIT COMMUNITY

- 149 members
- 42 organizations
- 10 countries

ExaDigit Member Meetup at SC'24 BoF in Atlanta

FINLAND

CSC

CZECH REPUBLIC

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

AUSTRALIA

Pawsey

UNITED STATES

SDSC SAN DIEGO SUPERCOMPUTER CENTER

Colorado State University

Texas State University

UIC UNIVERSITY OF ILLINOIS CHICAGO

UND UNIVERSITY OF NORTH DAKOTA

UNITED KINGDOM

epcc

Georgia Institute of Technology

Los Alamos NATIONAL LABORATORY

Lawrence Livermore National Laboratory

FRANCE

INES

GERMANY

JÜLICH Forschungszentrum

ITALY

ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

Cornell University

BERKELEY LAB

NREL NATIONAL RENEWABLE ENERGY LABORATORY

OAK RIDGE National Laboratory

Argonne NATIONAL LABORATORY

SWITZERLAND

CERN openlab

UNIVERSITÀ DI TRENTO

SWEDEN

KTH VETENSKAP OCH KONST

DOD HPC MODERNIZATION PROGRAM

INDUSTRY COLLABORATORS

Hewlett Packard Enterprise

NVIDIA

Meta