OAK RIDGE National Laboratory

High Performance Data-centre Digital Twins

BoF Session 1B @ CUG2024

2024-05-06 - **16:35-18:00pm**

Dr. Matthias Maiterth

Dr. Tim Dykes

Dr. Jess Jones

Adrian Jackson

Prof. Michele Weiland

Dr. Wes Brewer

ORNL is managed by UT-Battelle, LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

Agenda:

- Intro and Welcome
- Technical Talks: (20 min each)
  1. Matthias Maiterth (ORNL): On the motivation for developing a digital twin framework
  2. Adrian Jackson(EPCC): On the potential of using digital twins to improve overall system and data center efficiency
  3. Jess Jones (HPE): On EX system architectures and networking in the context of digital twins.
- Community outreach / Audience discussion: (30 min)
  perspectives, experiences, and discussion
  to contribute and collaborate on data-center digital twins.

High Performance Data-center Digital Twins

Terms and context:

Capable of modeling:

- Large Scale-Systems
- 3Vs for Telemetry (Volume/Velocity/Variety)
- Including capable Simulations

Goal:

- Understand & Optimize

Relevant for:

- HPC centers
- Data-centers
- Workload Agnostic & System Specific

Goal:

- Abstract/Generalizable Design

"A digital twin is a set of virtual information constructs that

1. mimics the structure, context and behavior of an individual/unique physical asset
2. is dynamically updated with data from its physical twin throughout its lifecycle
3. informs decisions that realize value."

AIAA Digital Engineering Integration Committee (2020)

Project at ORNL: EXADIGIT

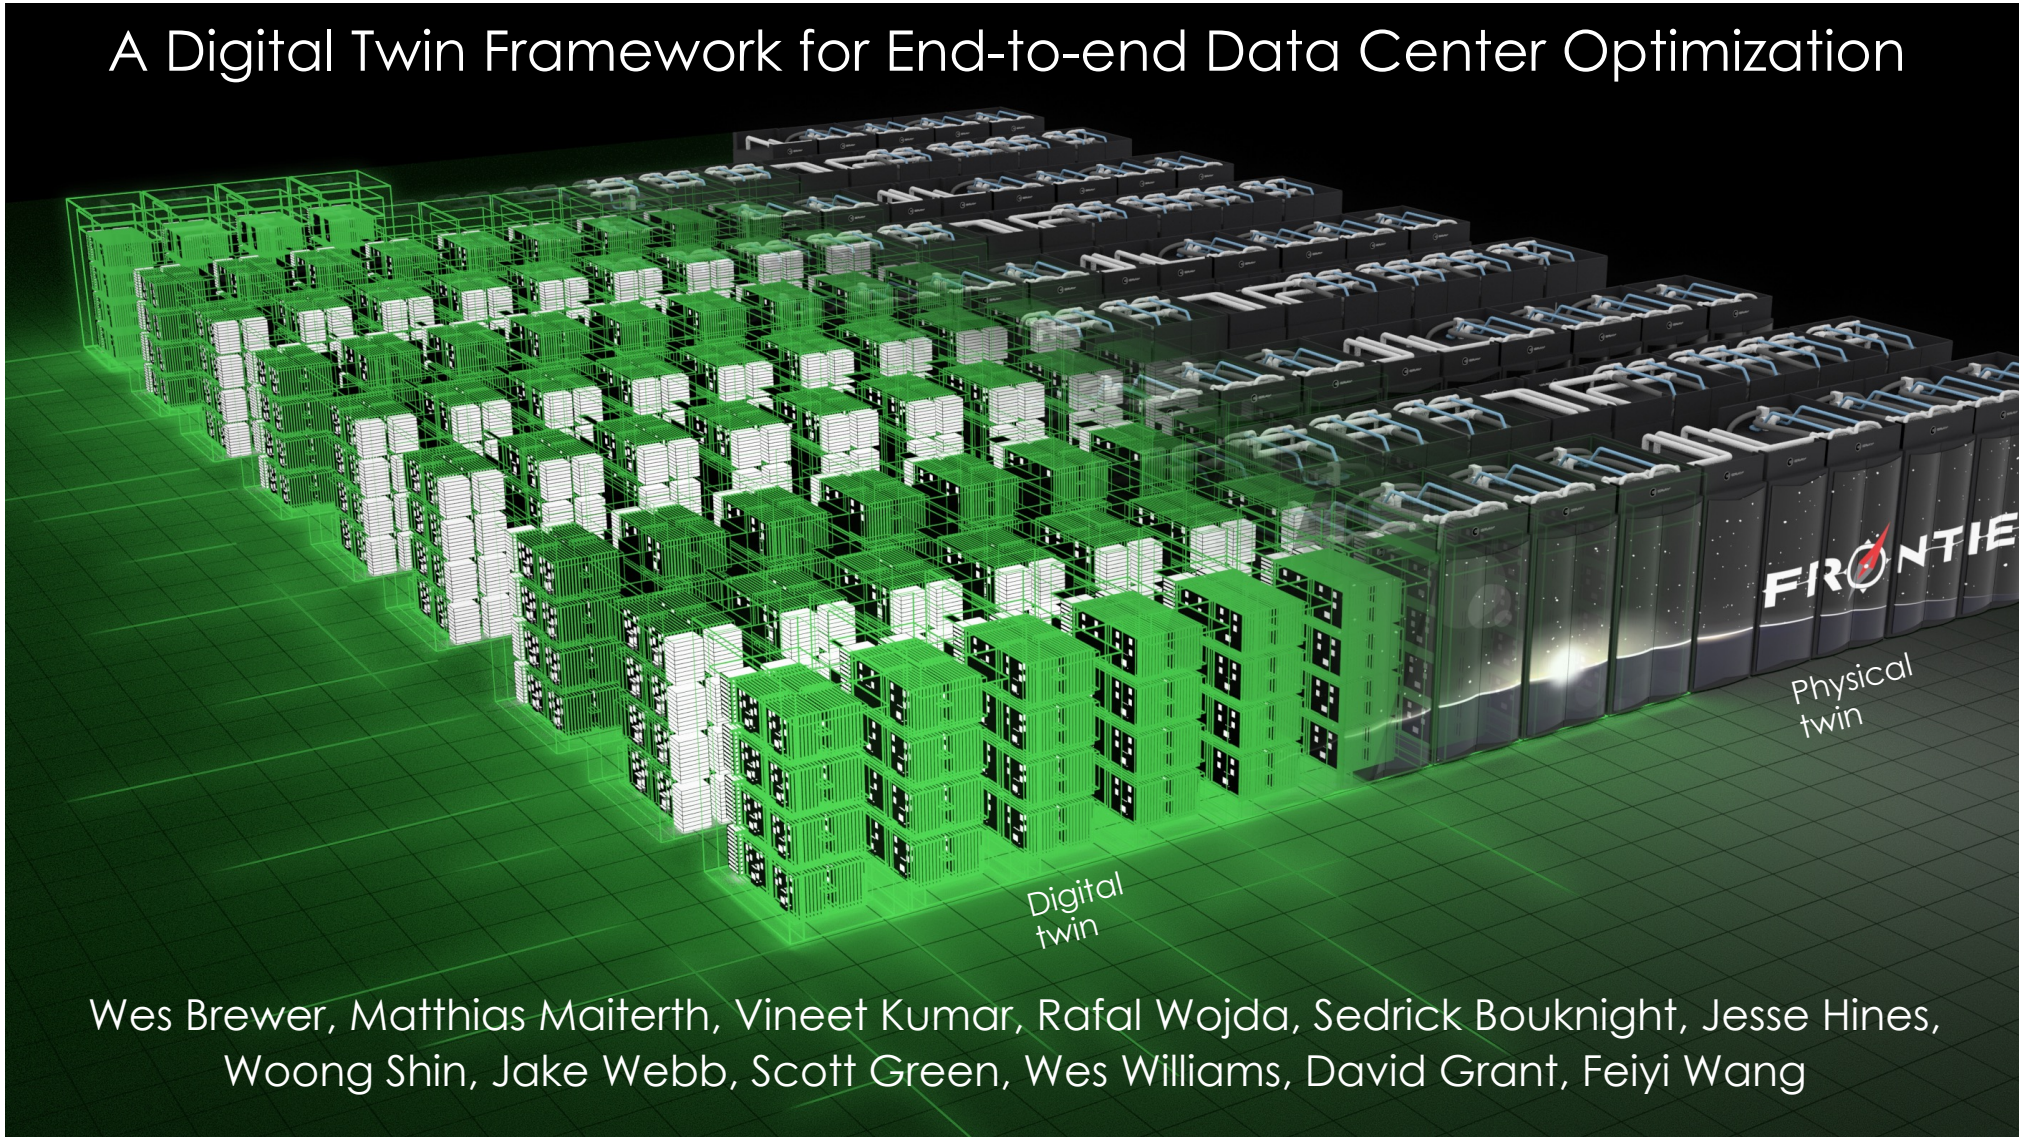

Frontier's Cooling System

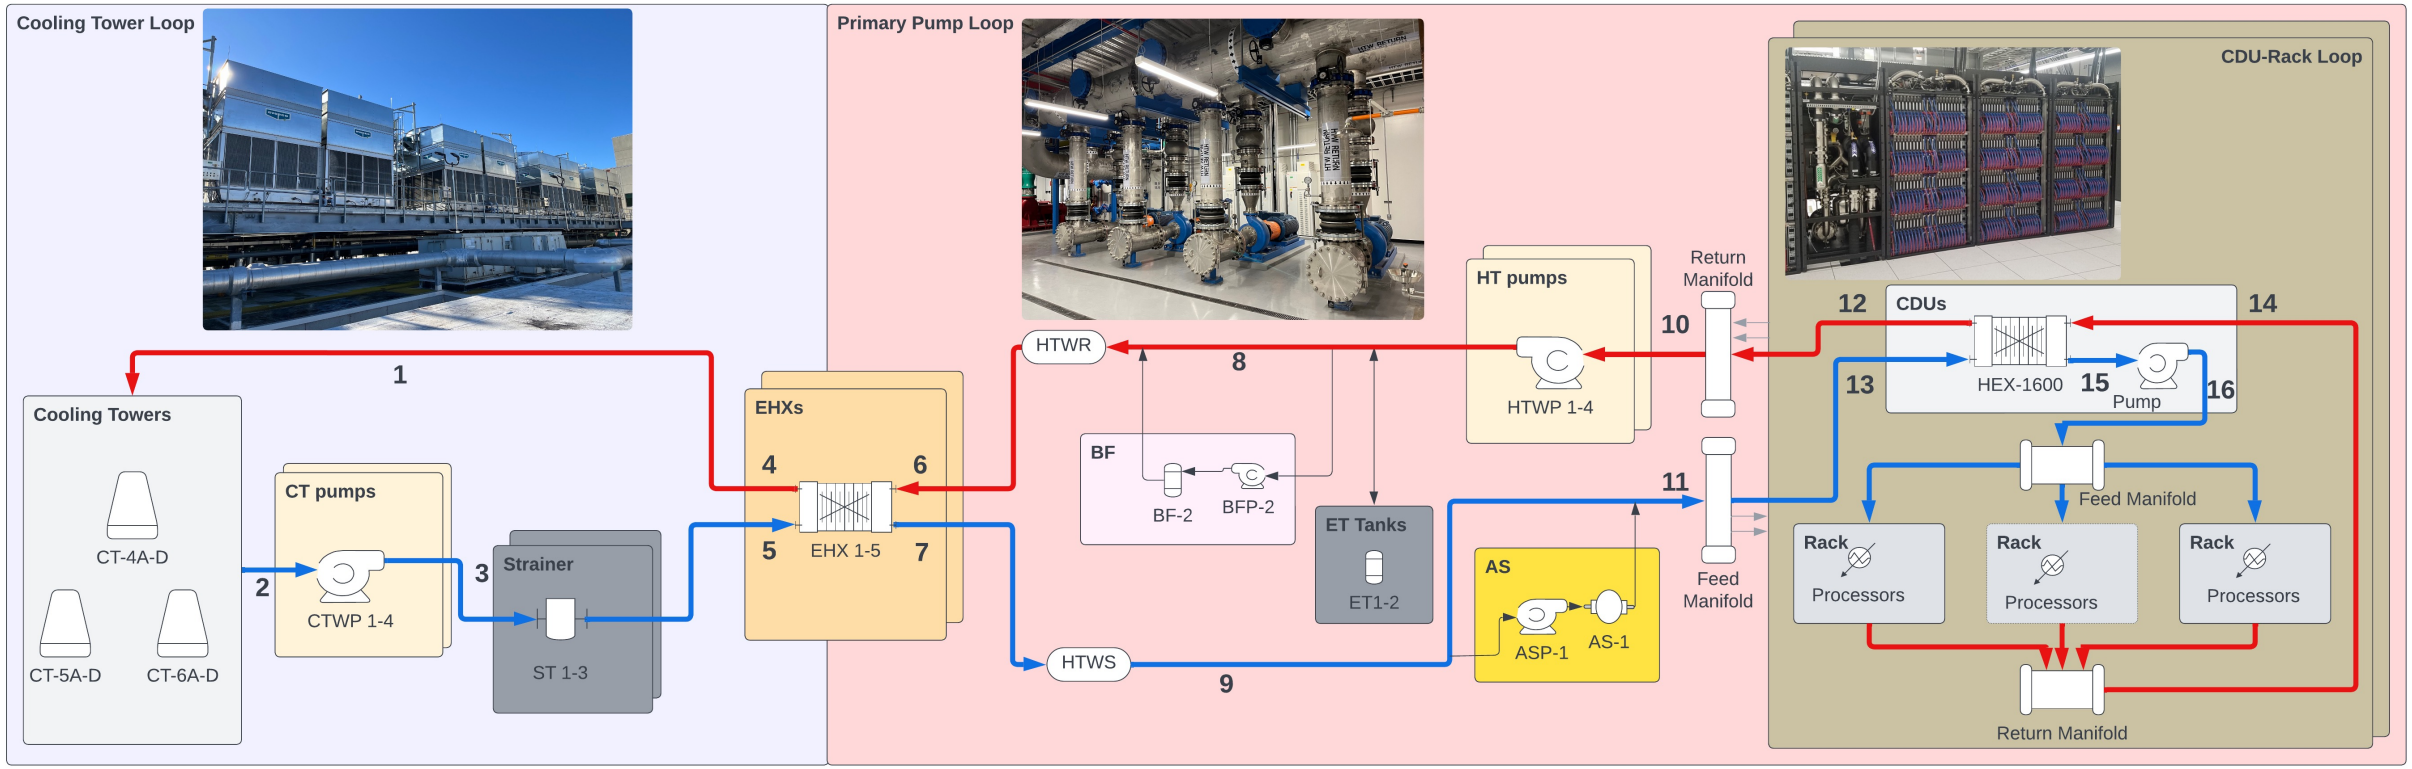

Frontier's Power System

480 kW power supply per cabinet

Power consumption model:

- Every second compute current CPU and GPU power and add up power on all nodes.
- Apply DC-DC voltage conversion losses at node-level.
- Apply AC-DC voltage conversion losses at chassis level.

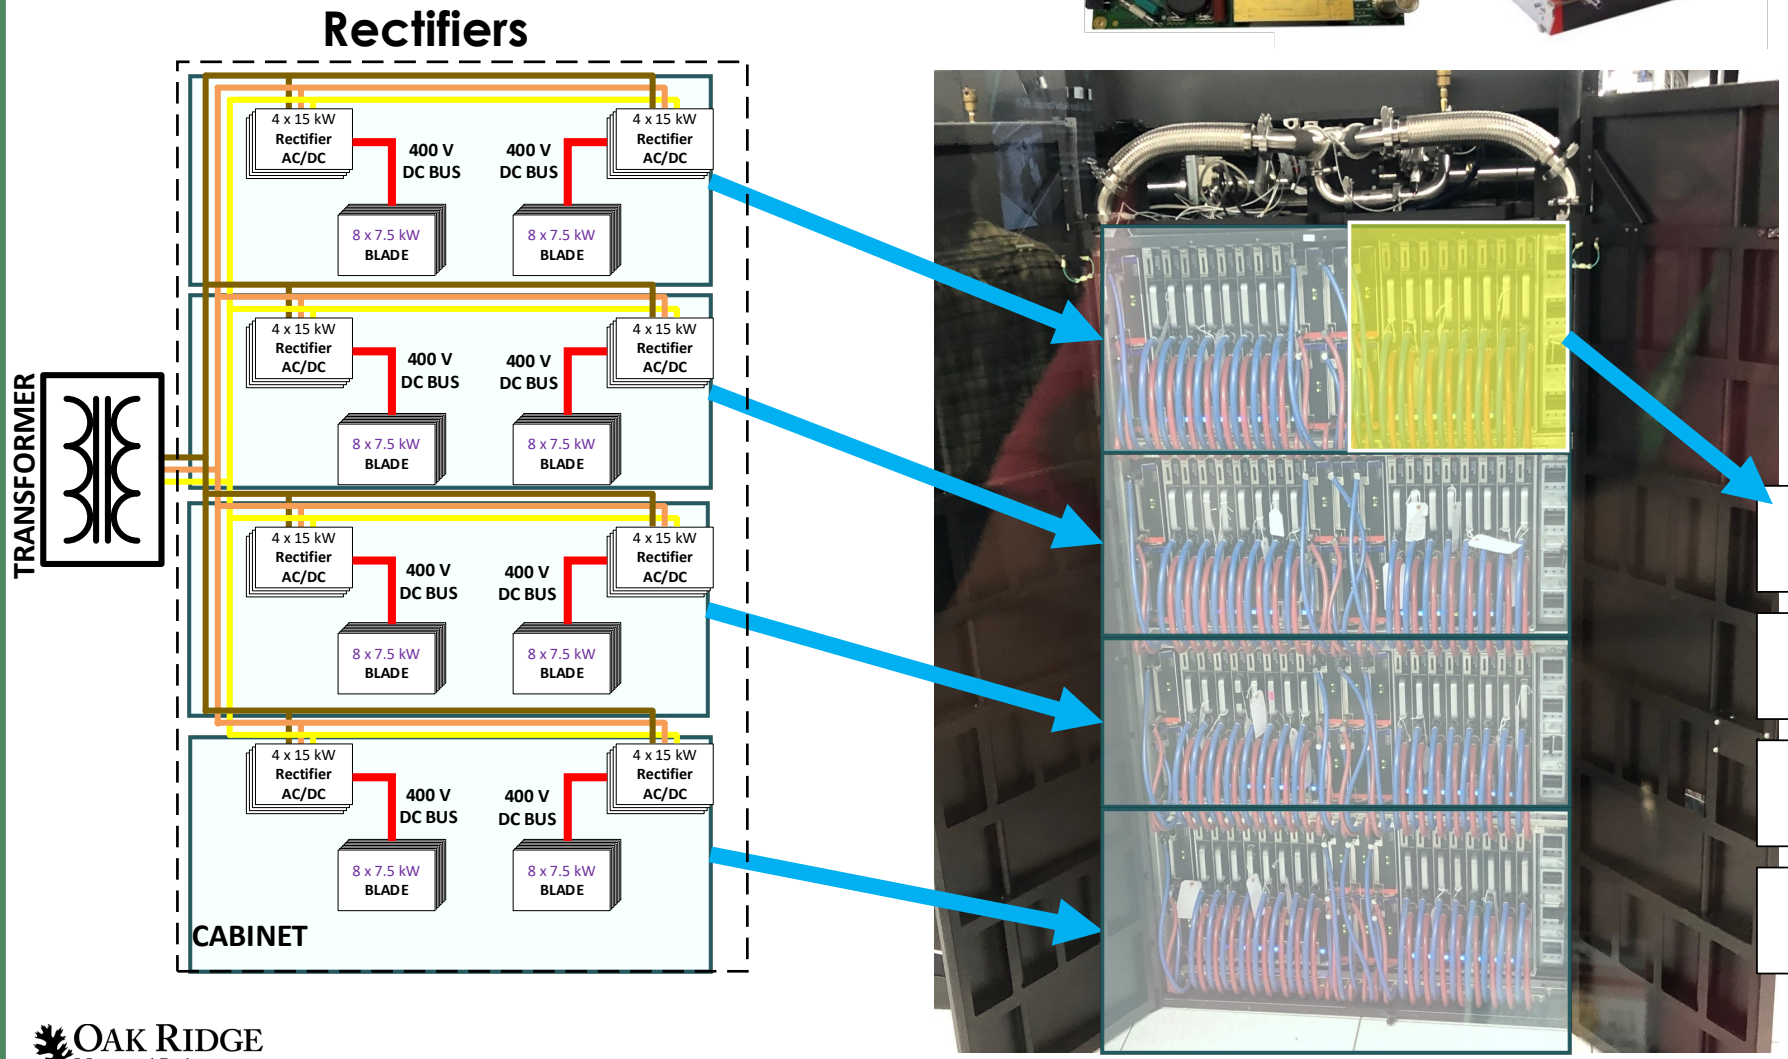

OAK RIDGE National Laboratory

Architectural Overview (as-built)

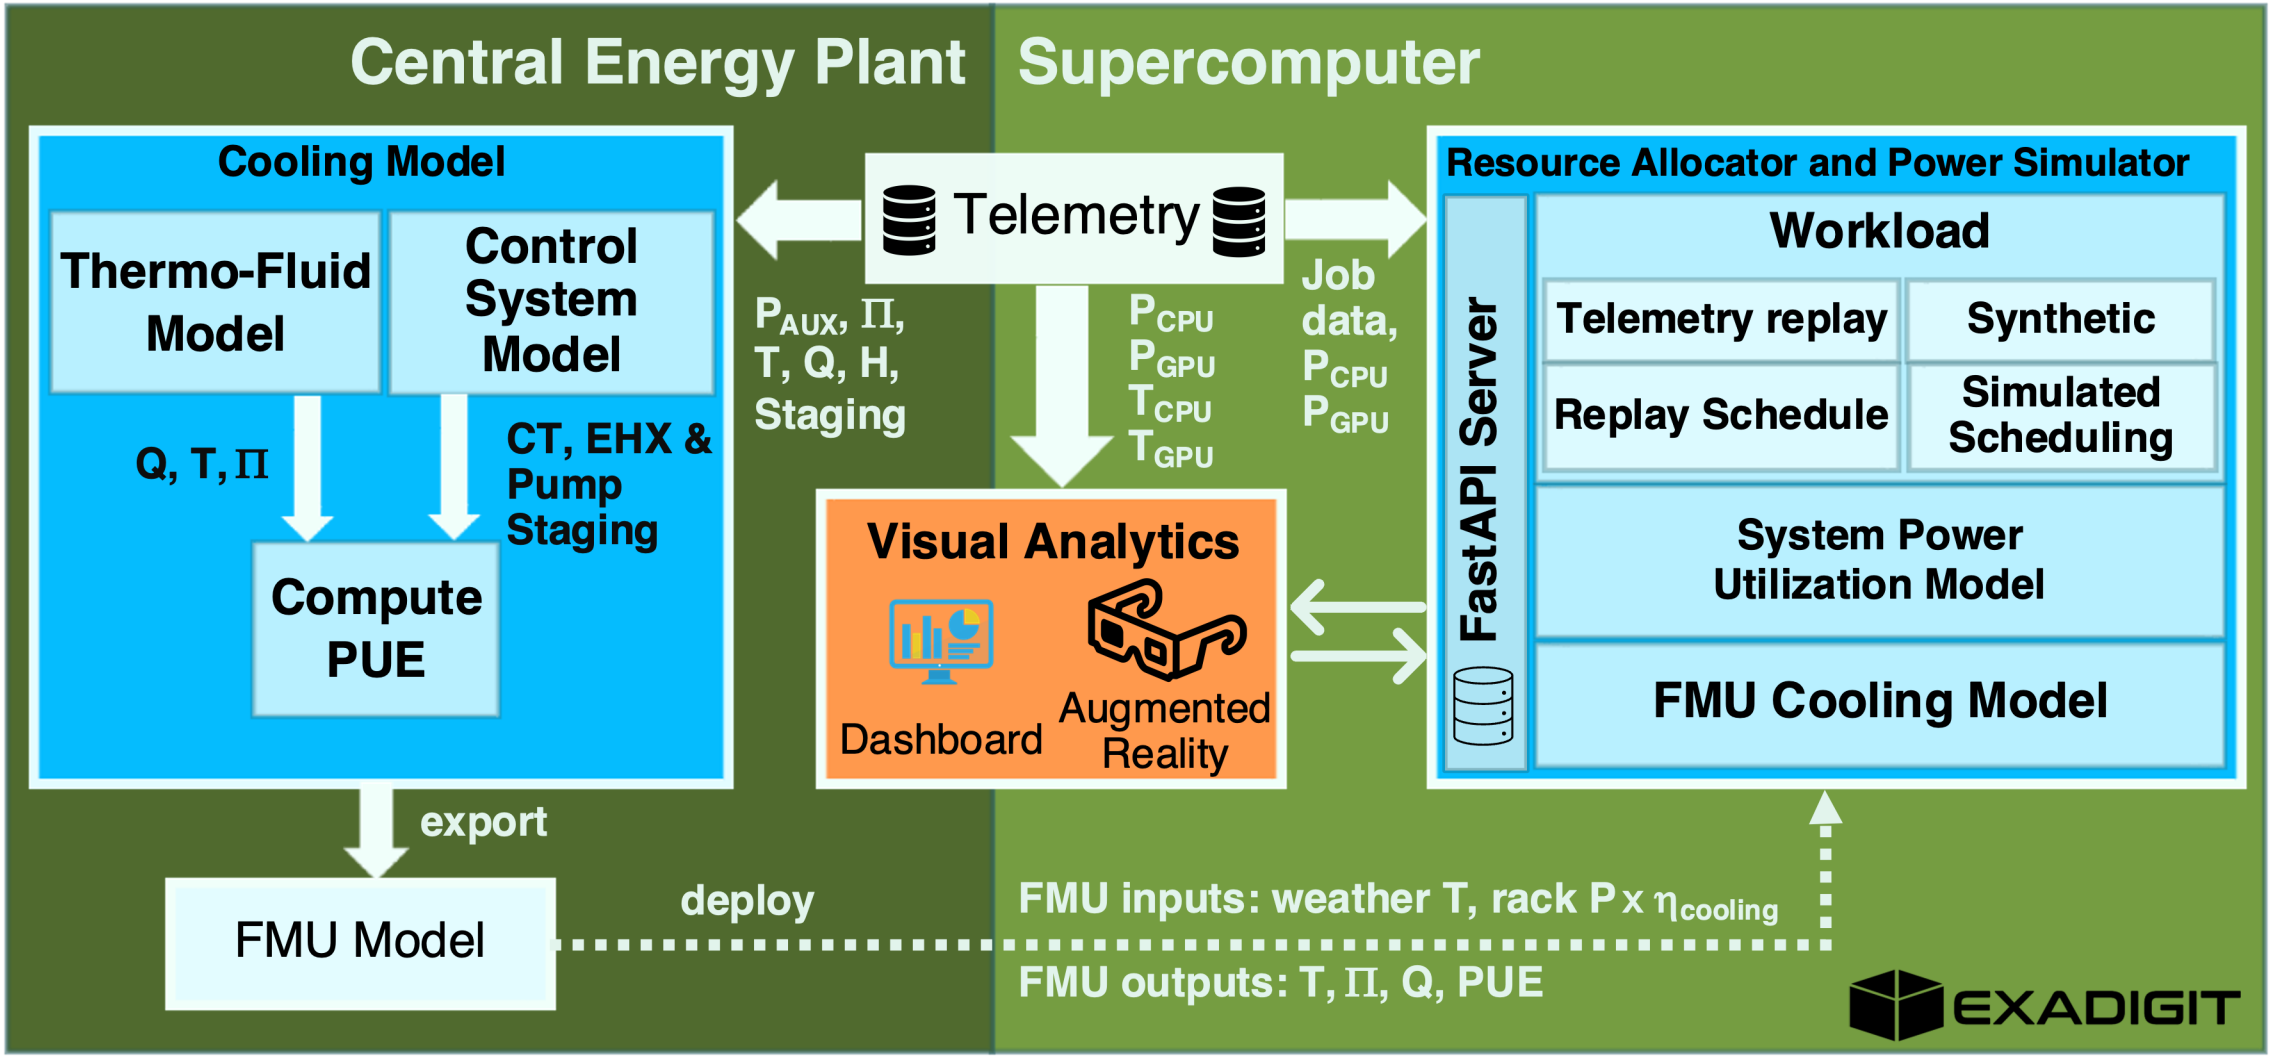

Current state: EXADIGIT

*Initial Paper under review

- Cooling Model: Modelica Based Transient Cooling Simulation
- Resource Allocator and Power Simulator: Job Simulator / Replay
- Visualization: exadigitUE5 (3d-AR/VR Interaction) / Dashboards
- Recent Extension: Network (see John Holmen's Talk on Thu)

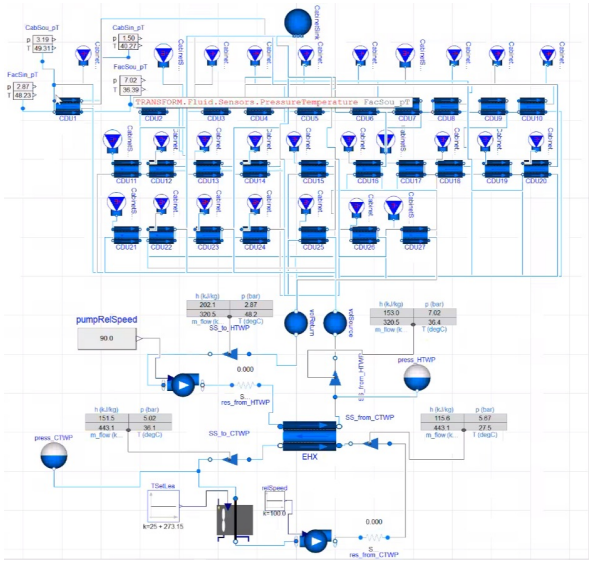

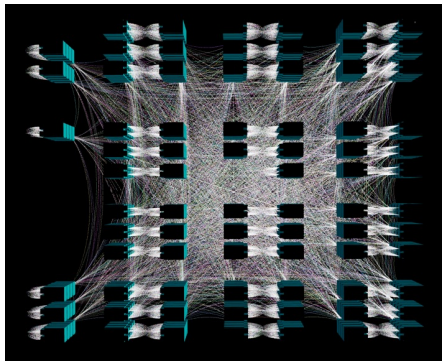

OAK RIDGE National Laboratory

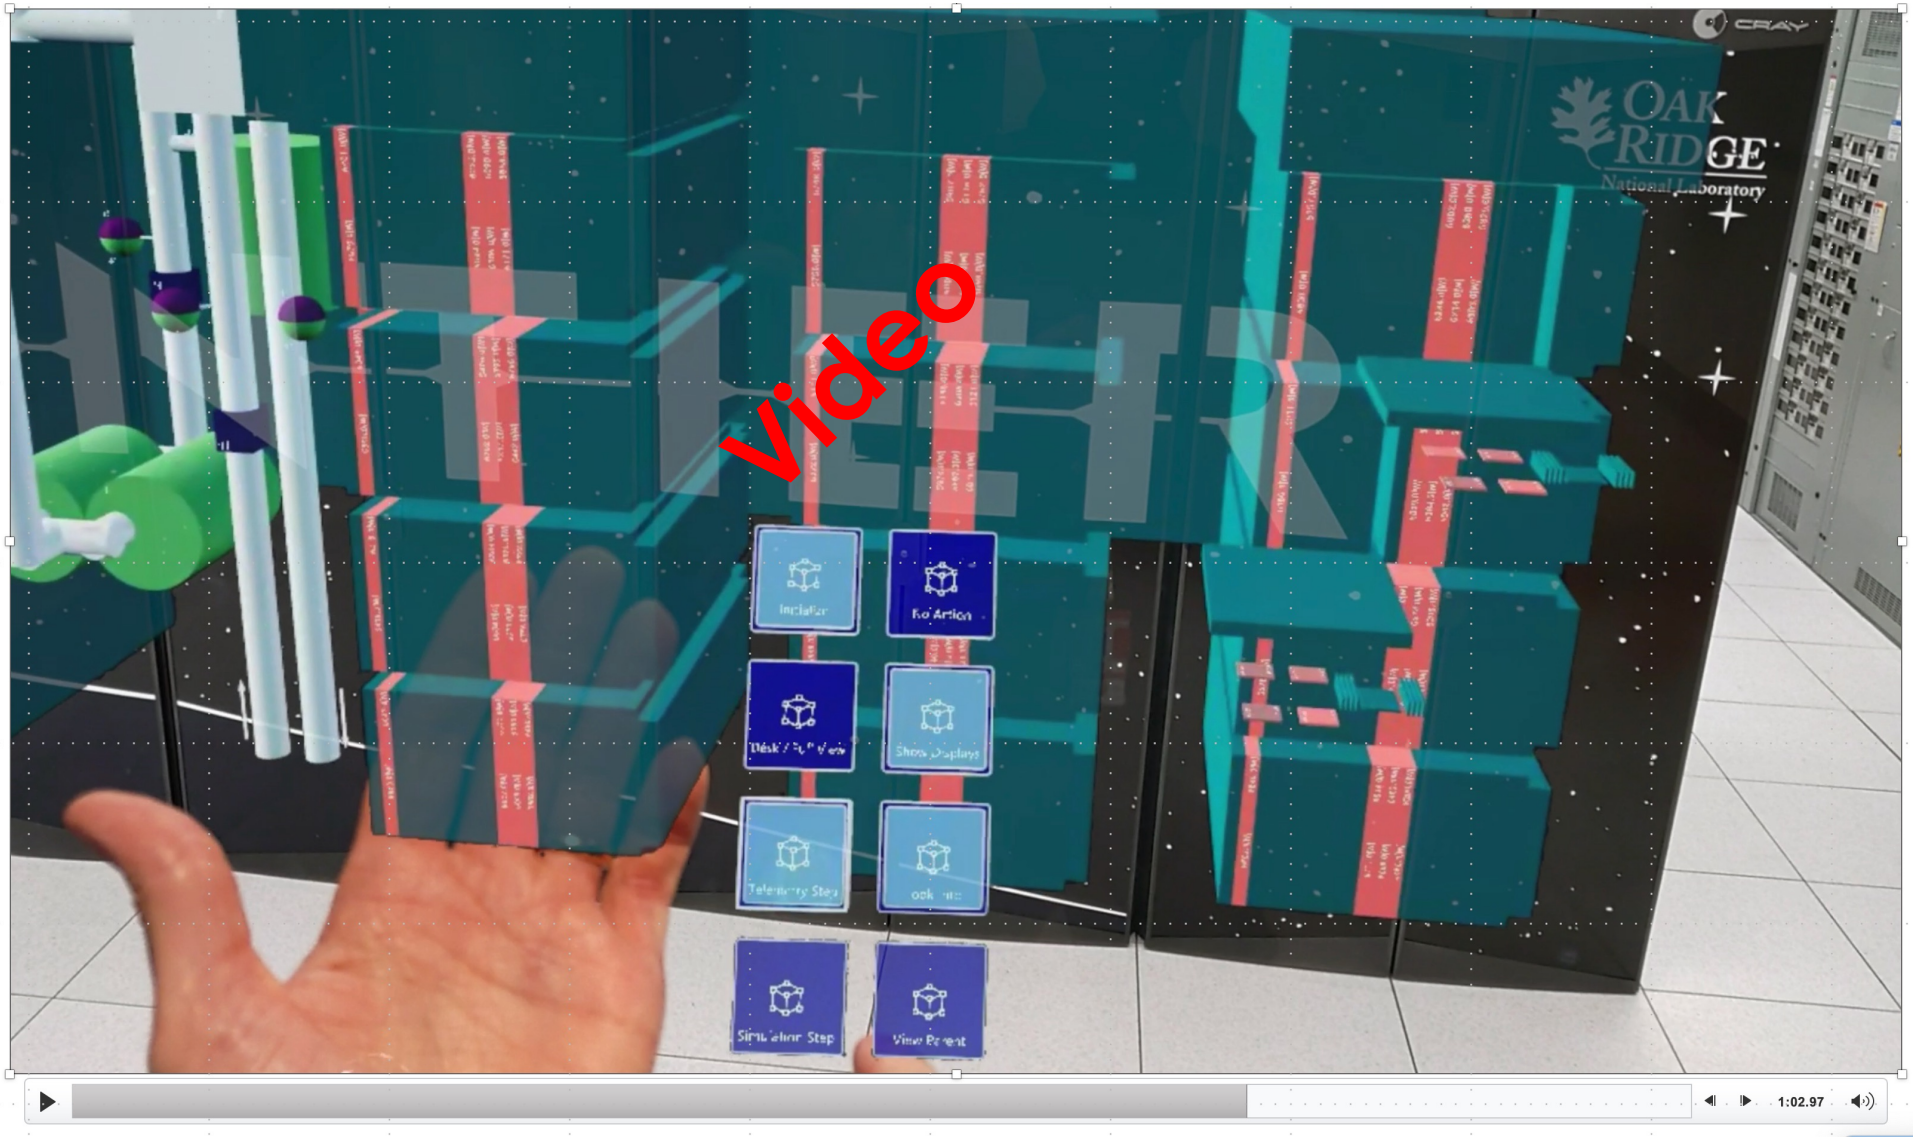
Video
OAK RIDGE
National Laboratory
1:02.97

Digital Twin Use Cases

Insights from Visualization

System-level insights
Inferring reliability
Job information

What-if Scenarios

Pump failure
Grid blackout
Parameter evaluation
Cybersecurity attacks

Forensic Analysis and Diagnostics

Diagnosing network congestion
Diagnosing node health

Operational Optimization

Energy efficiency
Cooling optimization
Optimal job scheduling

Understanding Complex Dynamics

Transient cooling dynamics
Mult-job interaction on network congestion

Virtual Prototyping

Designing future systems

Virtual Training

Synthetic Data Generation
HPC Training Tool

Initial Focus @ ORNL:

- **System-level insights** → understanding transient model of cooling loop
  → understanding energy conversion losses
- **What-if scenarios** → testing ideas for energy-efficiency

Initial Progress and Strong Positive Feedback

- Initial Outreach:
  - Presentation at SC'23 (invite only at the time)
- Community Feedback:
  - Similar Systems / Similar Goals
  - Technology Readiness
  - Initial investigations done, e.g. by IT4I and others w/o simulation as done by our approach
- Need for an Open-Source Framework
  - Different solutions and generational differences of Systems
  - Different emphasis/interest for each participant
  - Emphasis Interoperability (Plug-and-Play)

FINLAND

CSC

UNITED KINGDOM

epcc

CZECH REPUBLIC

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

AUSTRALIA

FRANCE

INES

pawsey

SWEDEN

GERMANY

JÜLICH Forschungszentrum

KTH VETENSKAP OCH KONST

INDUSTRY PARTNERS

Hewlett Packard Enterprise

NVIDIA

EXADIGIT

Community & workgroups

Application Fingerprinting	AI/ML/RL	Visual Analytics	Power & Cooling	Network	VVUQ
• Terry Jones	• Soumyendu Sarkar	• Matthias Maiterth	• Adrian Jackson	• Puneet Sharma	• TBD • Delayed Start
Use Case / Architectures (Tim Dykes)					
Documentation (Gabriel Hautreux)					

Monthly General Meeting + Workgroup Meetings (Online)
Slack: https://exadigit.slack.com/
Email: Wes Brewer - brewerwh@ornl.gov
Repositories: Accessible to collaborators and in the process of open sourcing (stay tuned)

Agenda:

- Intro and Welcome
- Technical Talks: (20 min each)
  1. Matthias Maiterth (ORNL): On the motivation for developing a digital twin framework
  2. **Adrian Jackson (EPCC): On the potential of using digital twins to improve overall system and data center efficiency**
  3. Jess Jones (HPE): On EX system architectures and networking in the context of digital twins.
- Community outreach / Audience discussion: (30 min)
  perspectives, experiences, and discussion
  to contribute and collaborate on data-center digital twins.

WHY BOTHER WITH A DATA CENTRE/SYSTEM DIGITAL TWIN?

Adrian Jackson

a.jackson@epcc.ed.ac.uk

Optimisation work primarily by

Andrew Turner

THE UNIVERSITY of EDINBURGH

epcc

ExaDigiT

- Goal to model the data centre and system inside it sufficiently well to
  - Replicate/replay data collected
  - Investigate the interplay of system jobs, environment, and data centre systems
  - Optimise data centre/system to reduce/increase something (energy, heat, power, performance)
  - Model future systems and system configurations
  - Understand where overheads/issues are
  - Visualise system/data centre data
- ExaDigiT broadly building
  - Thermo-fluid model of the data centre and system
  - Power/energy model of the data centre and system
  - Network performance/capability
  - Node level performance/capability
  - Visualisation
  - Coupling/replay system to bring all this together

Data Centre

- Wide range of facilities/equipment
  - Transformers
  - UPS
  - Heat exchangers
  - Chillers
  - etc...
- Different rooms with different configurations
- Wide range of systems
  - Liquid cooled
  - Air cooled

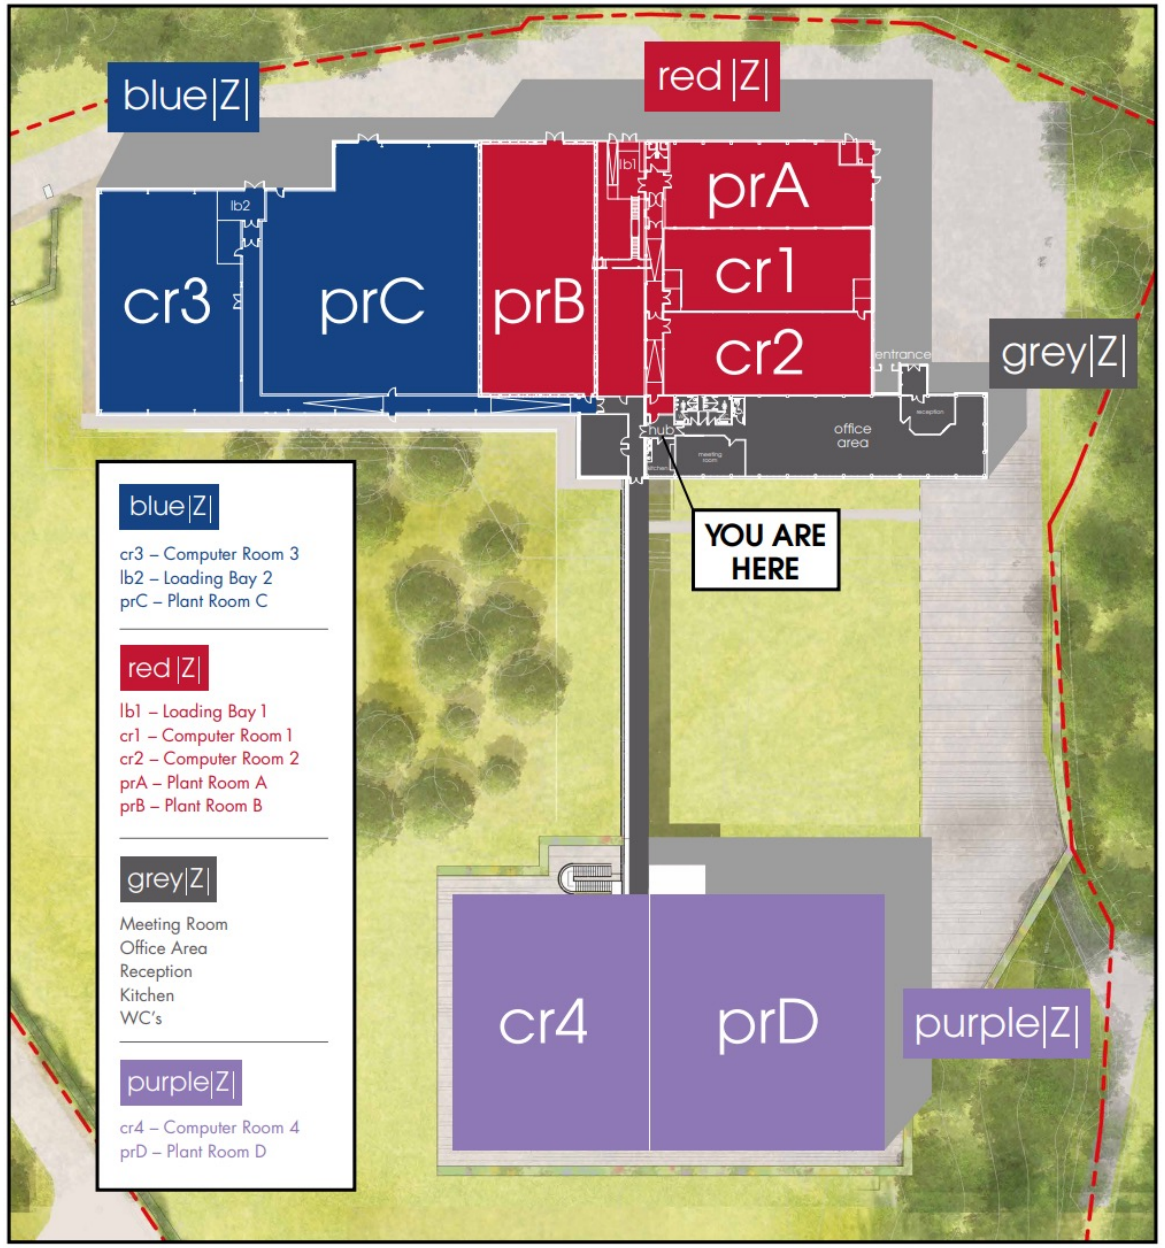

Data Centre

- Specifying new plant/equipment/room
  - Moving from *over provision* to *just enough*
  - Understanding maximum system size/load possible
  - Understanding burst capacity/issues
- Net zero planning/enabling
  - Energy usage a big issue for HPC centres
  - Evaluate the energy impact of changes to the centre
  - Estimate future energy usage

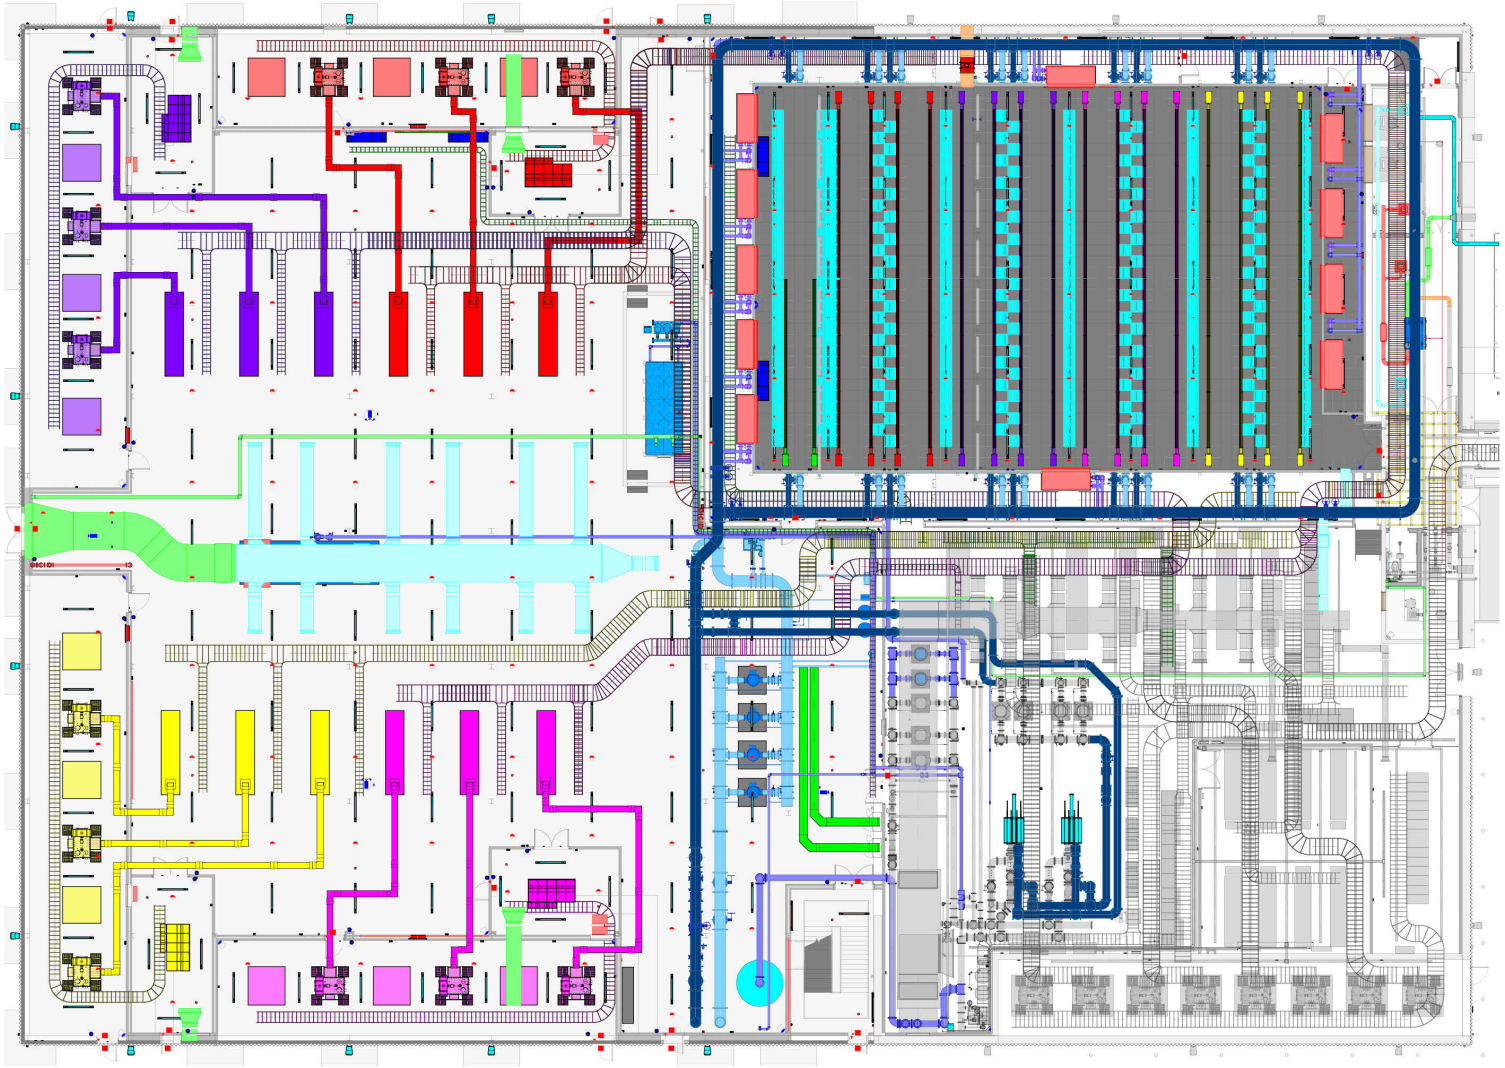

System optimisation

- Processor and system can be configured for different runtime modes
- Original ARCHER2 configuration called *Power Determinism*
- *Performance Determinism* keeps processor performance more consistent
  - Performance of multi-node parallel applications is determined by slowest node
  - Any extra power draw for performance above the slowest node is wasted power

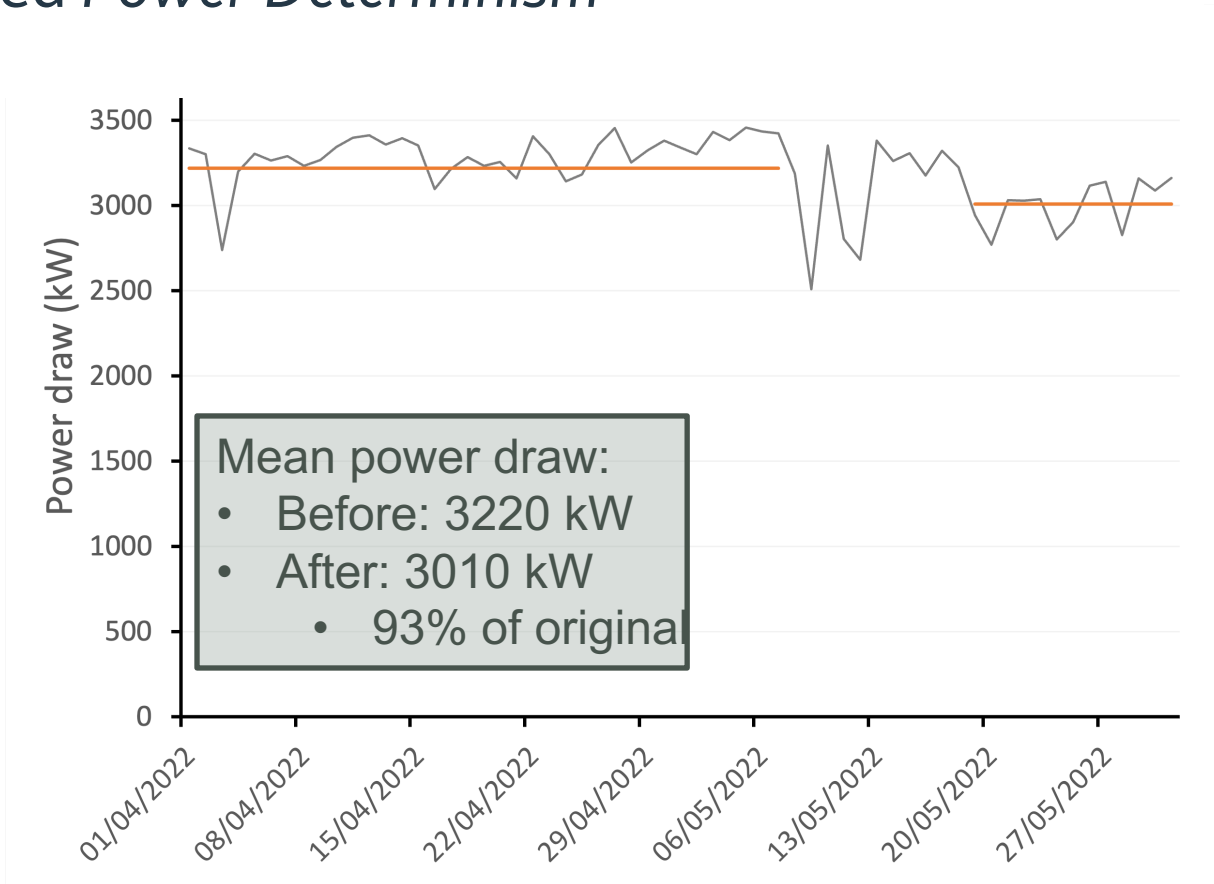

System optimisation

- As well as configuring the processor/node overall, can modify processor behaviour on the fly
- ARCHER2 configuration:
  - 2.25GHz processor frequency
  - turbo boost enabled
- New configuration
  - 2.00 GHz (no turbo boost)
- Can be done on a per application/per job basis

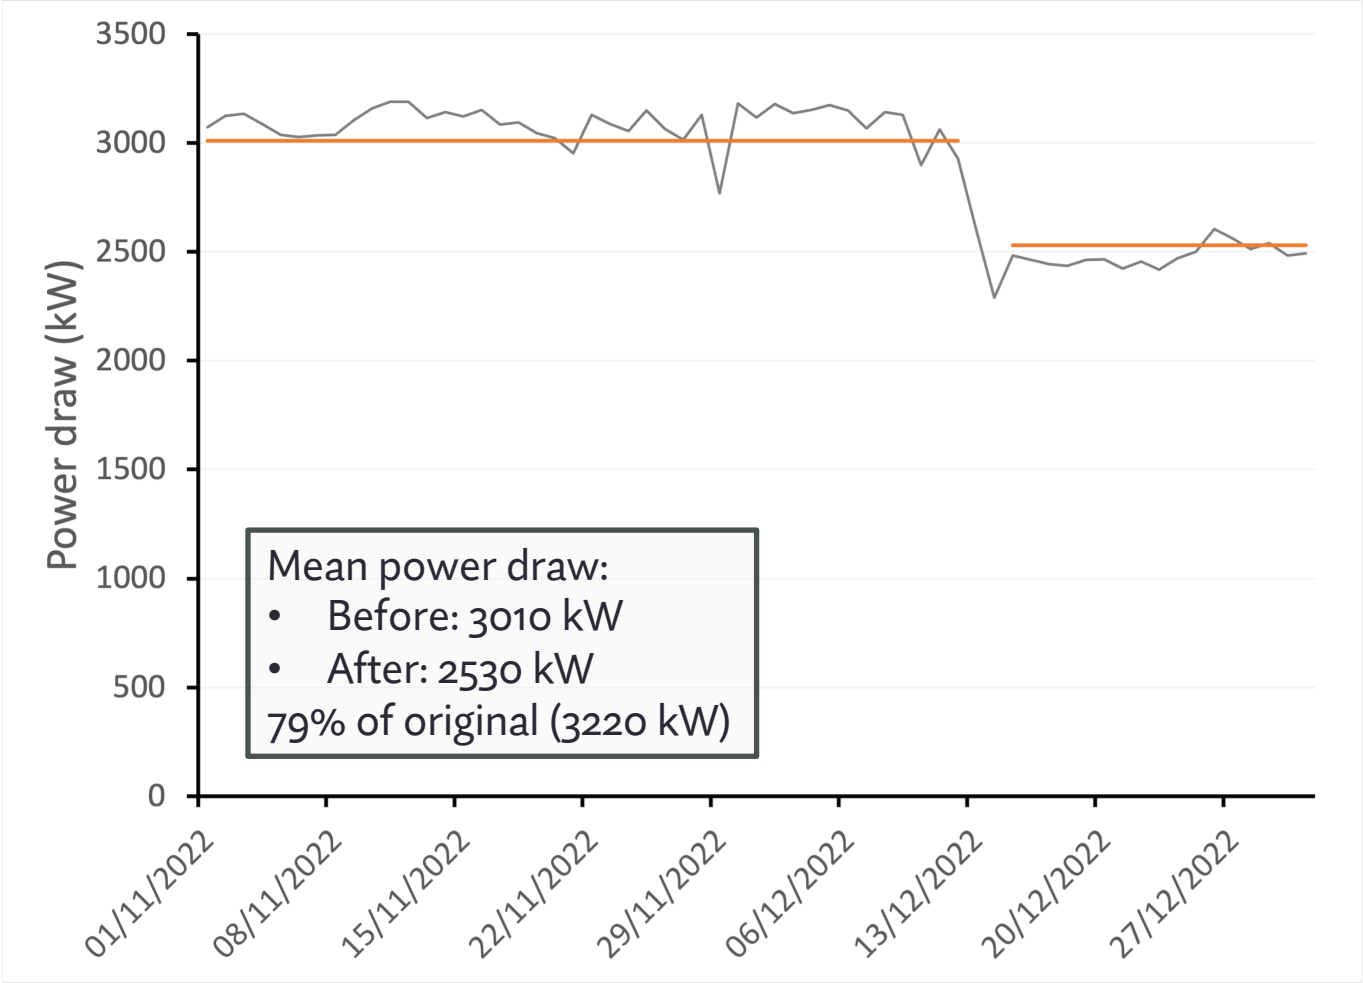

System optimisation

Application benchmark	Performance ratio	Energy usage ratio	Energy to performance ratio
VASP CdTe	0.95	0.88	1.08
GROMACS 1400k atoms	0.83	0.92	0.9
CP2K H2O 2048	0.91	0.93	0.98
LAMMPS Ethanol	0.74	0.92	0.8
CASTEP Al Slab	0.93	0.88	1.05
ONETEP hBN-BP-hBN	0.92	0.82	1.12
Nektar++ TGV 128 DoF	0.80	0.80	1

- All applications are more energy efficient at 2.0 GHz
- Looking at cost-efficiency would suggest:
  - Frequency set to 2.25 GHz: GROMACS and LAMMPS, Nektar++ [due to increased residency costs]
  - Frequency set to 2.0 GHz: VASP, CASTEP, ONETEP, CP2K
- Default frequency: 2.0 GHz with strong advice to users to test impact on their software

ExaDigiT use case work

General Categories of Use Cases

Insights from Visualization

System-level insights
Inferring reliability
Job information

What-if Scenarios

Pump failure
Grid blackout
Parameter evaluation
Cybersecurity attacks

Forensic Analysis and Diagnostics

Diagnosing network congestion
Diagnosing node health

Operational Optimization

Energy efficiency
Cooling optimization
Optimal job scheduling

Understanding Complex Dynamics

Transient cooling dynamics
Mult-job interaction on network congestion

Virtual Prototyping

Designing future systems

Synthetic Data Generation

Fault identification

ExaDigiT use case work

More concrete use cases

- Optimization
  - Scheduling
    - Optimize for network congestion
    - Optimize for power consumption
  - Cooling
    - Optimize cooling system efficiency
    - Use DT for design of PID controllers for minimizing setpoint overshoot
- Virtual prototyping studies
  - Smart load-sharing rectifiers
  - What size HPC system can cooling system support?
- SCADA cybersecurity
  - Study implications of a SCADA cybersecurity attack
- Virtual training
  - Use DT for training neural networks
  - Use DT for training new HPC users, operators, etc.

ExaDigiT use case work

(some of the) HPE Use Cases

- **System User Use Cases**
  - Assist customers (datacenter operator; datacenter user) with their use cases through CoE/ACC's
  - Focus work (Tim): LUMI visual model & workload visualization (visual analytics WG); Architectures & use cases (this WG)
- **System support and deep troubleshooting**
  - Gaining additional insight into complex issues
  - Deep dive into system performance issues
    - I/O: LASSi, Network: SST & SS monitoring; Energy: job sim)
  - Visualising job placement across the system (esp. across network groups)
  - Focus work (Tim) – job scheduler simulation & workload visualization (in app fingerprinting & visual analytics WG)
- **Customer system pre-sales**
  - Build virtual model of to-be-installed systems
  - Visualisation of data-center in-situ
  - Demonstrate prospective job load performance
  - Focus work (Tim) – system configurator & visual models (in visual analytics WG)

Range of digital twin functionality

- Level 0
  - Visualise current plant/system data in real time
  - Reporting on collected data
- Level 1
  - Replay previous data and correlate between data sources
- Level 2
  - Simulate possible configurations of the system to predict behaviour
- Level 3
  - Control system operation with feedback from the digital twin simulation in real time

Use cases

- What other use cases should we be considering?
- What benefits could a digital twin of a data centre/HPC system provide you?

Agenda:

- Intro and Welcome
- Technical Talks: (20 min each)
  1. Matthias Maiterth (ORNL): On the motivation for developing a digital twin framework
  2. Adrian Jackson(EPCC): On the potential of using digital twins to improve overall system and data center efficiency
  3. **Jess Jones (HPE): On EX system architectures and networking in the context of digital twins.**
- Community outreach / Audience discussion: (30 min)
  perspectives, experiences, and discussion
  to contribute and collaborate on data-center digital twins.

Hewlett Packard
Enterprise

DIGITAL TWINS FOR DATACENTERS AT HPE

Dr Jessica R Jones, HPC & AI EMEA Research Lab

May 6, 2024

WHY IS HPE INTERESTED IN DIGITAL TWINS?

- HPE has ambitious sustainability targets
  - Net-zero across entire value chain by 2040 (or sooner)
  - Science Based Targets Initiative (SBTI) approved
  - Digital twins allow us to optimise our datacenters from sustainability perspective

- They also allow us to…
  - Demonstrate prospective system installations to customers
  - Collaborative remote system insight
  - Improve system operation for existing customers
  - Design future systems
  - Develop, debug, predict..

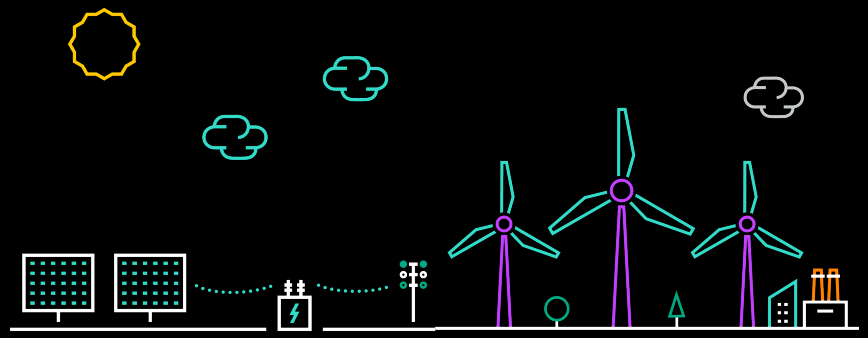

A SNAPSHOT OF DIGITAL TWIN RESEARCH AT HPE

ExaDigiT

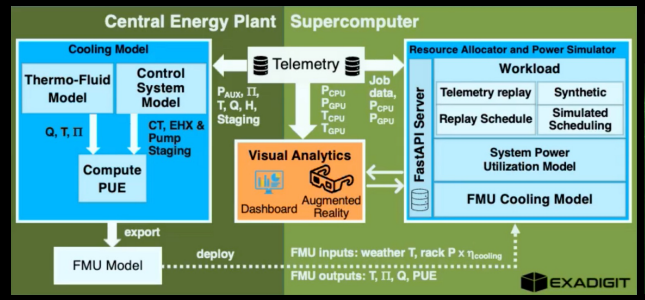

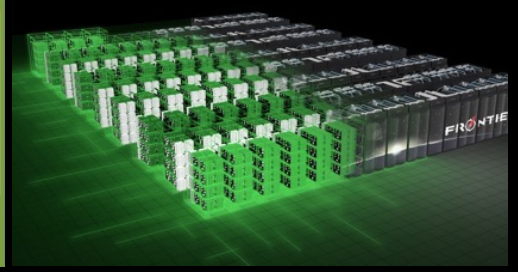

AI Research & Systems Architecture Labs with Oak Ridge National Lab

Digital Twin for Datacenter

- Empowers IT and facilities management teams to collaborate on datacenter energy efficiency and performance optimization
  - Walk through physical space
  - Visualize Power & Cooling
- Cadence DC Insight demo available at CIC
- Additional initiative complementing HPE Labs' digital twin research

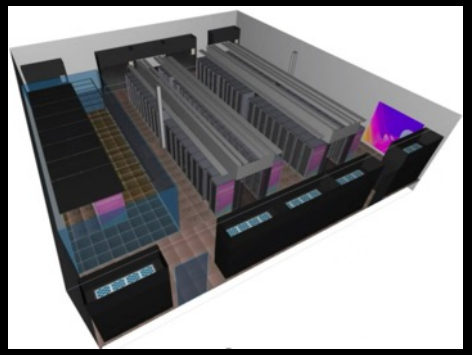

Workload Metering

- Provide visibility on the energy being consumed at the workload level by hardware components (processor, memory, accelerators, …)
- Bare Metal, Virtual Machine and Containers
- Converge and compare energy data monitored in-band and out of band to take account of shared components and losses
- Provides insightful data for future decision making
- Alignment with HSC HPE Labs and HPC Business Unit

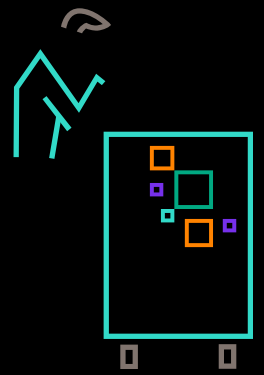

Digital Twin of the Earth

- Destination Earth (DestinE)
- Unprecedented resolution < 5 km
- Use cases span cities, farming, energy and climate
- Consortium of 12 European partners
- Deployed on HPE's LUMI supercomputer
- HPE proposing Maestro HPC streaming middleware

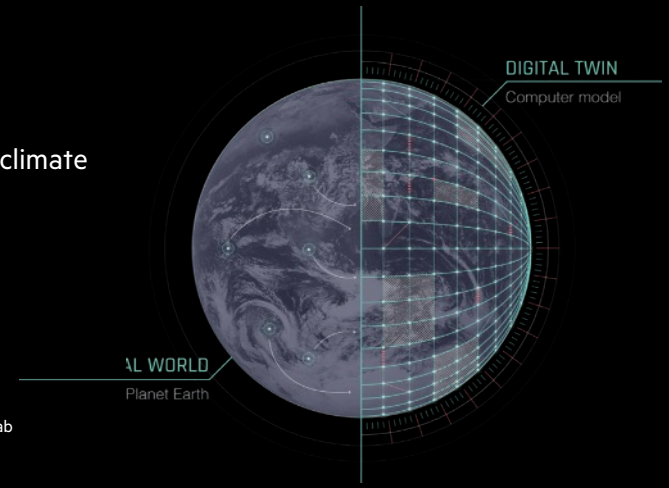

HPC & AI Europe, Middle East and Africa Research Lab

DIGITAL TWINS OF CRAY EX SYSTEMS IN EXADIGIT

- Working with ORNL, LUMI and others on the ExaDigit DT project
- We are particularly interested in:
  - system configuration and optimisation
  - monitoring & workload visualisation
  - network twins
- Working with LUMI to build LUMI variant of Frontier visual model & integrating EX models
- CDU 3D model still under dev.

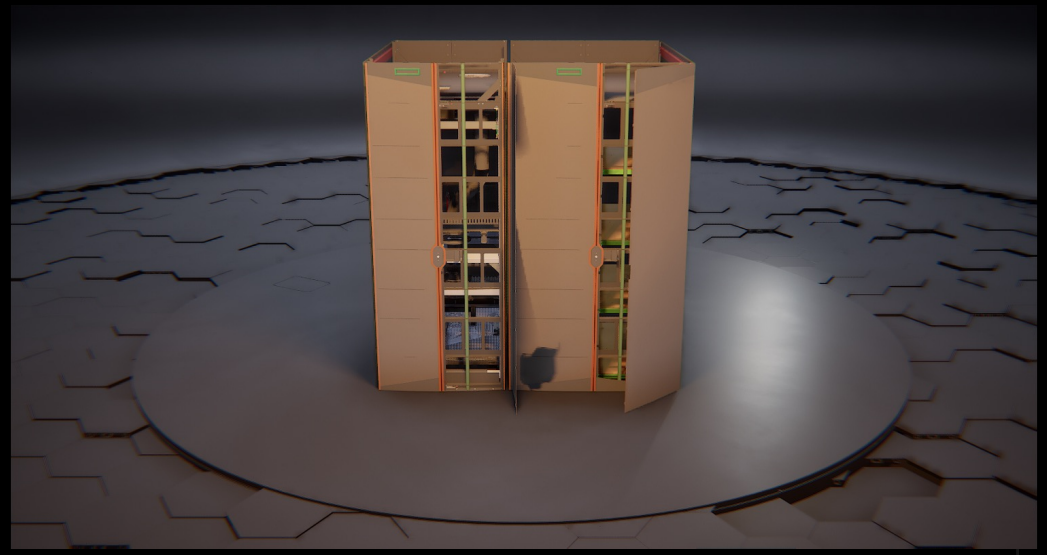

Cray EX4000 CDU & Compute Cabinet in UE5

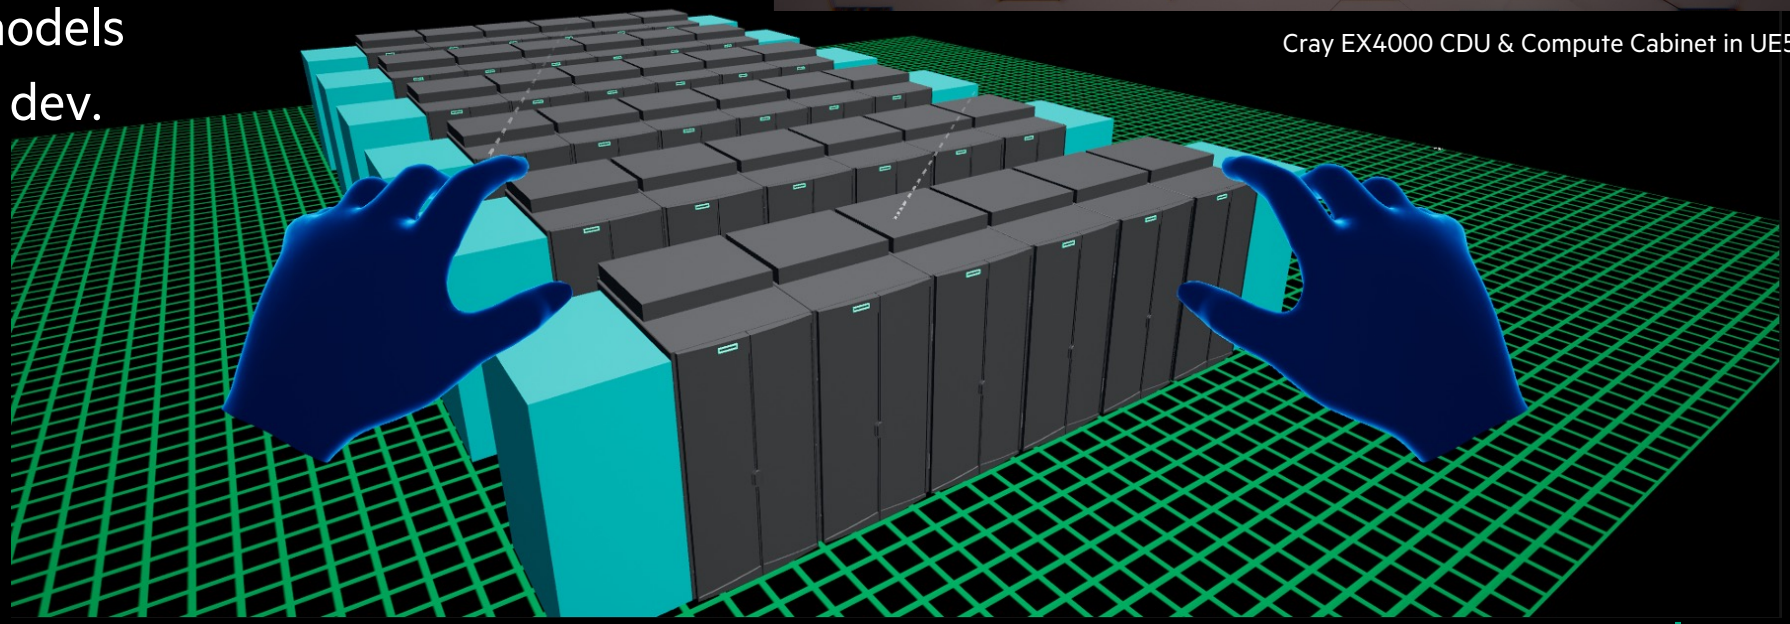

LUMI System layout demo in Exadigit visual analytics module

PHYSICAL DIGITAL TWINS

- Built for customer ACC
- Based on CAD models of EX4000, 1/20th scale
- 3D printed in black & green PLA filament
- 4x Raspberry Pi Zero 2 w per cabinet
  - 4 core 64 bit, with hw FP
- Cut down Raspbian
  - Custom build using RPi foundation toolchain
- Currently ad-hoc wifi network
- Modules, Slurm with cgroups, MPICH, OpenMPI
- Distributed PyTorch & TensorFlow
- JupyterHub on ‘head’ node

WORKLOAD MANAGEMENT SIMULATION

- Used real data from ARCHER2 and LUMI
- Predicted the impact of scheduler configuration changes on workload throughput
- Focus on power management
- Used to provide insights on how CPU frequency changes could impact on job runtime

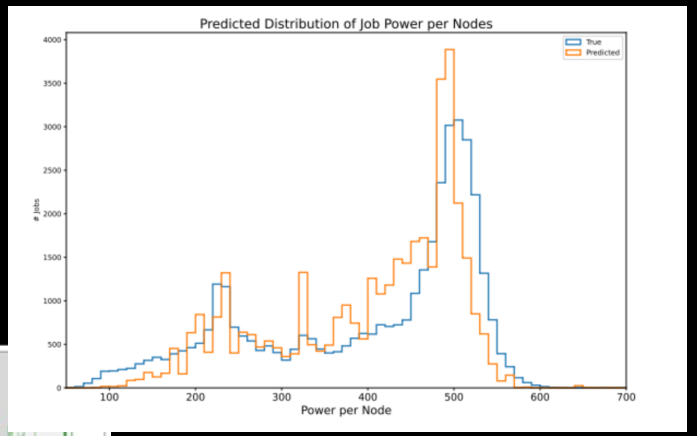

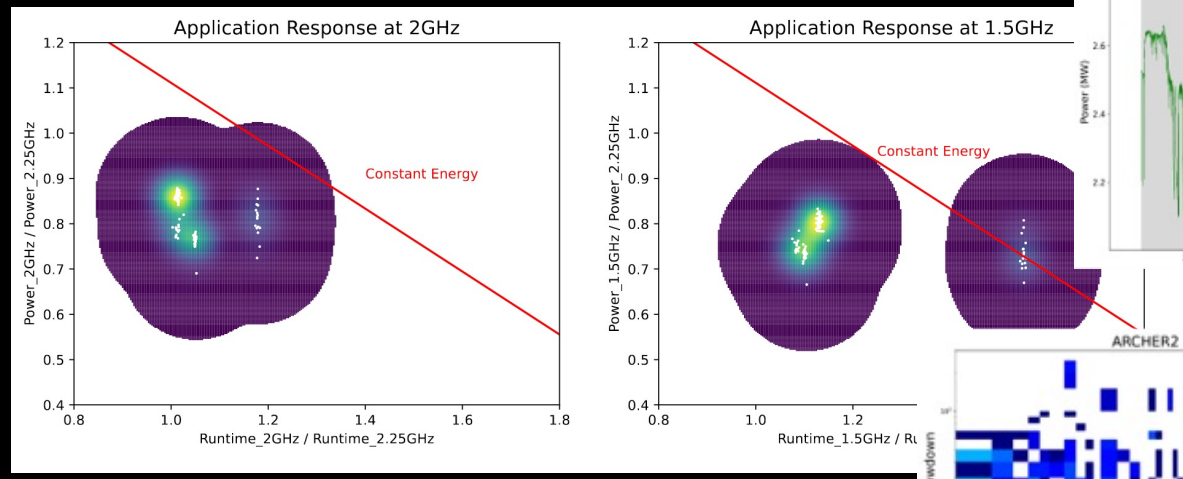

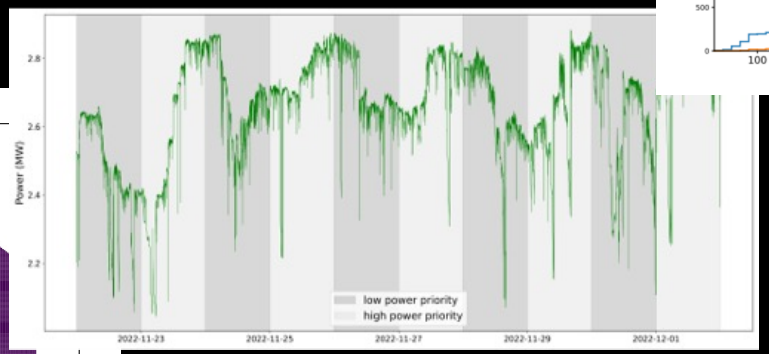

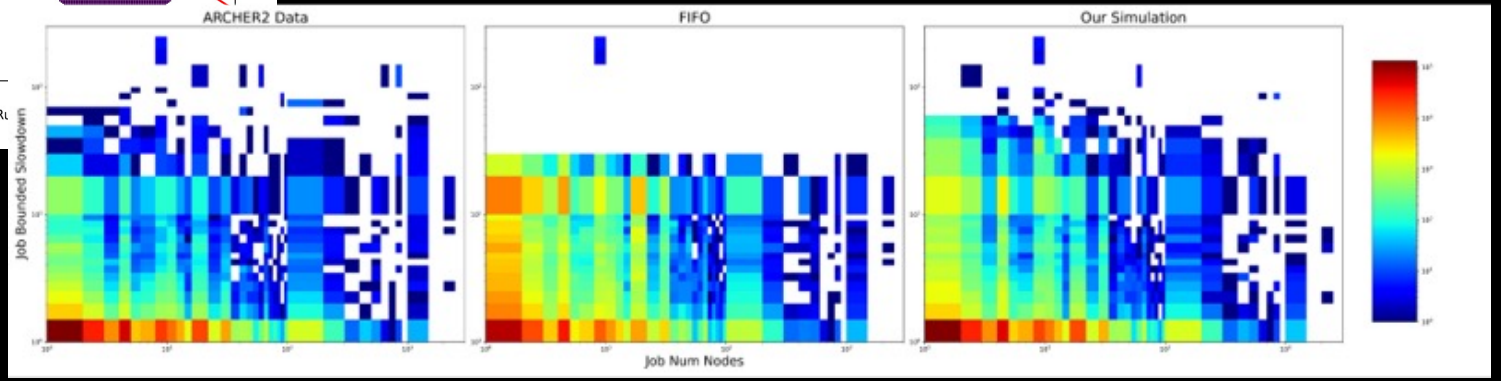

https://doi.org/10.1007/978-3-031-40843-4_24

SCHEDULING FOR ENERGY EFFICIENCY

Shifting power across the schedule...

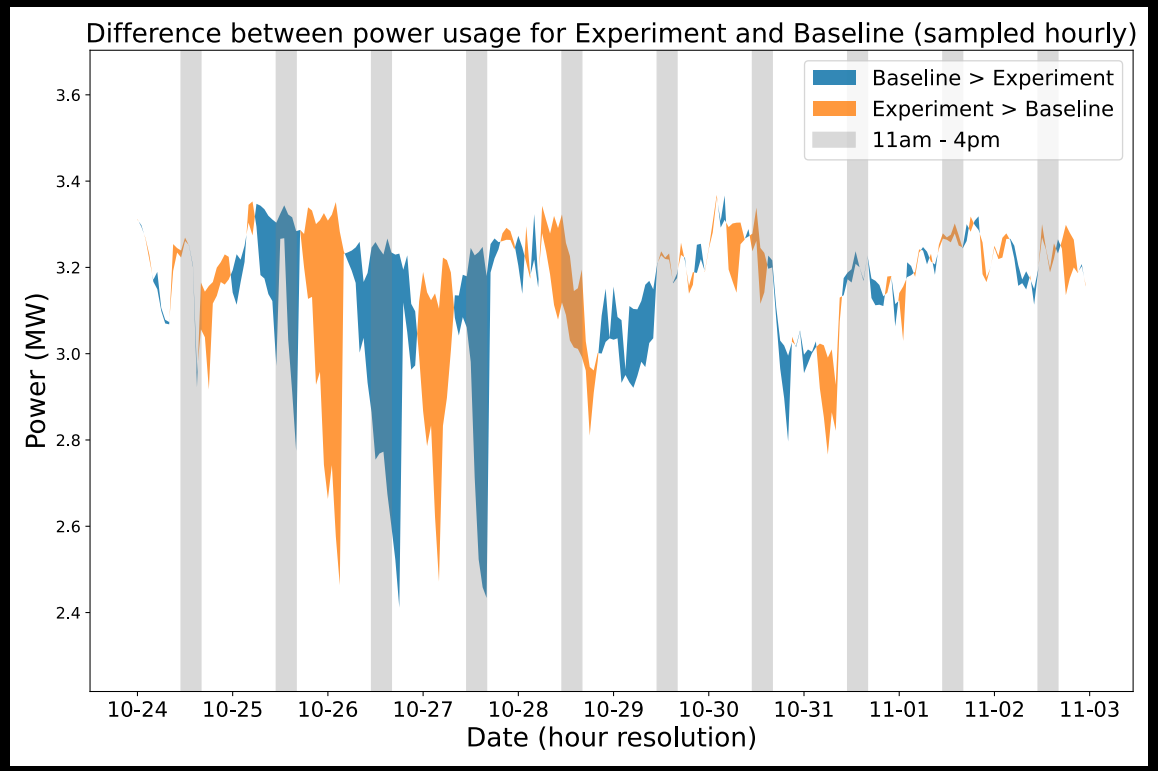

... with low impact on users

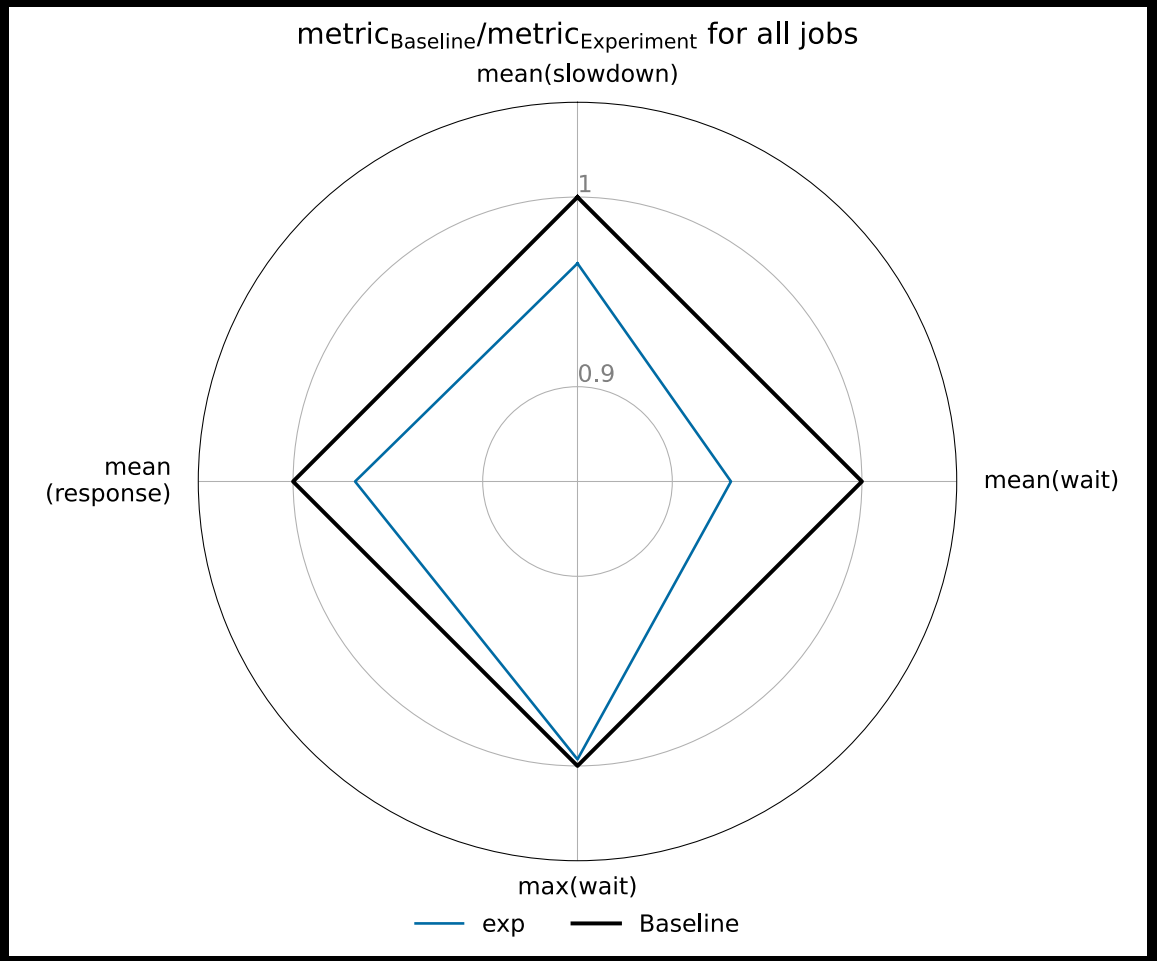

Alex Wilkinson et al, MODA4@ISC23

WORKLOAD MANAGEMENT / APPLICATION FINGERPRINTING

- Visualising the workload on LUMI
  - Visualisation of system load
  - Job distribution, overlap, congestion
  - Highlight potential spaciotemporal aspects of complex technical issues

- Working on module to support multiple interfaces in various projects
  - CLI
  - Web
  - 3D

- Other projects on-going at HPE
  - Performance modelling
  - Application fingerprinting

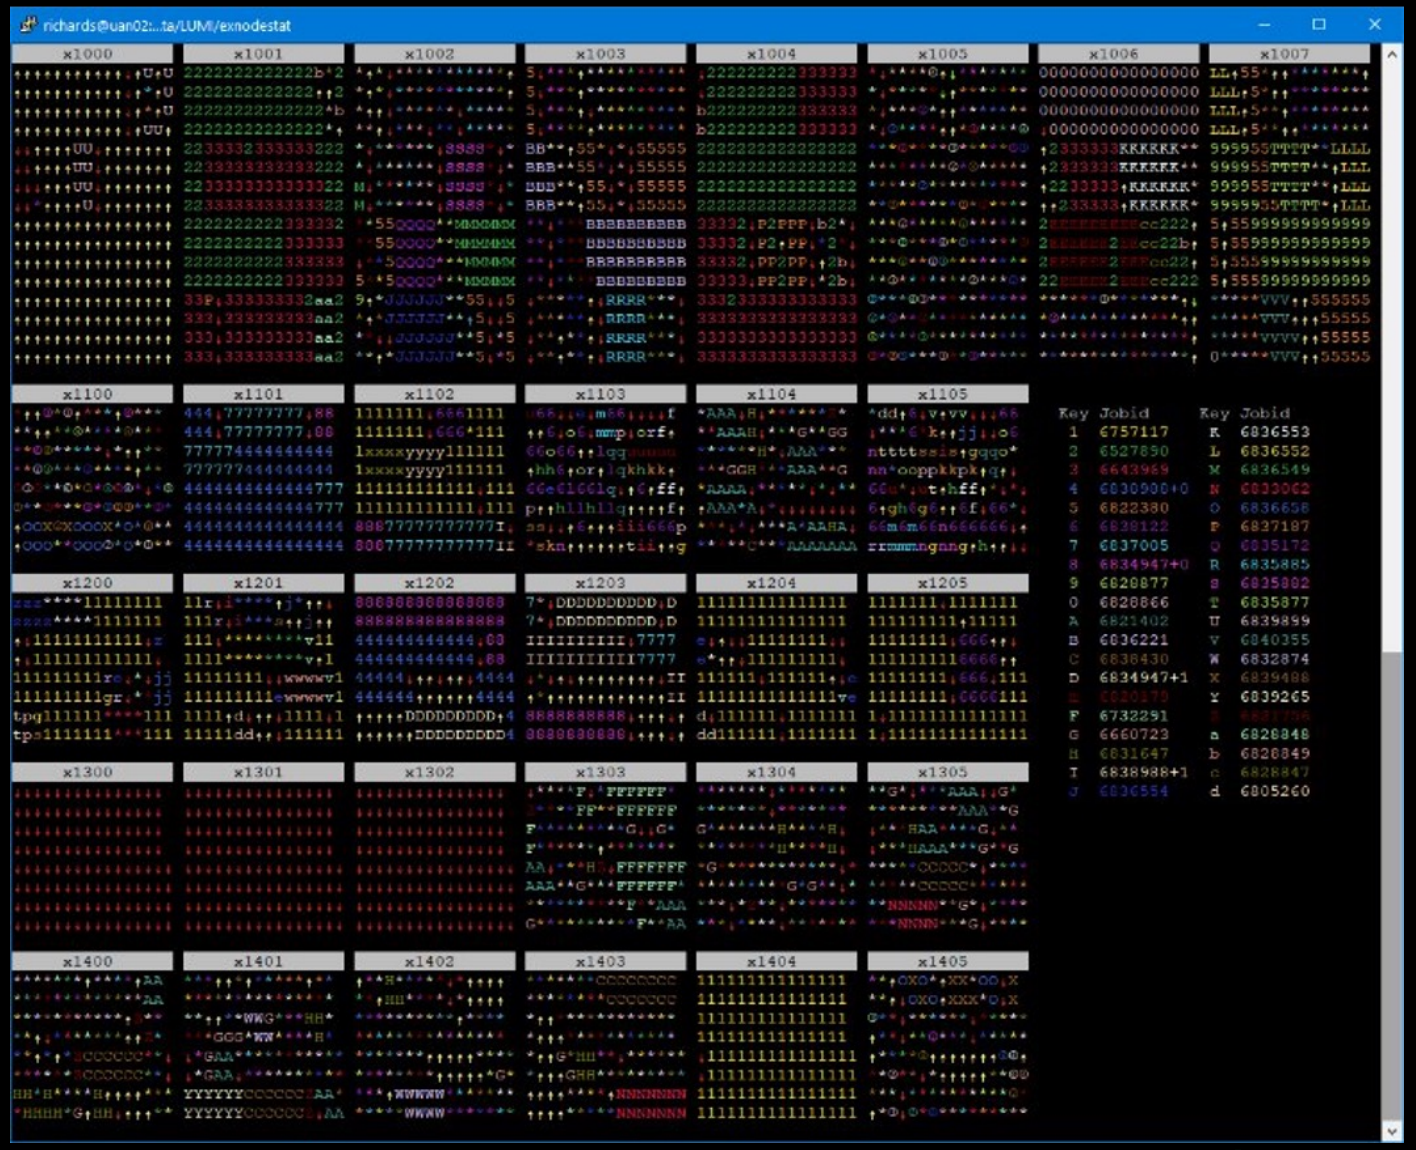

NETWORK DIGITAL TWINS

- Challenges?
  - Data capture from real network – vast quantities, etc
  - Slow?
- Why bother?
  - The network is one of the most important components
- A digital copy of a real system gives opportunities for:
  - Prediction
  - Debugging
  - Exploration

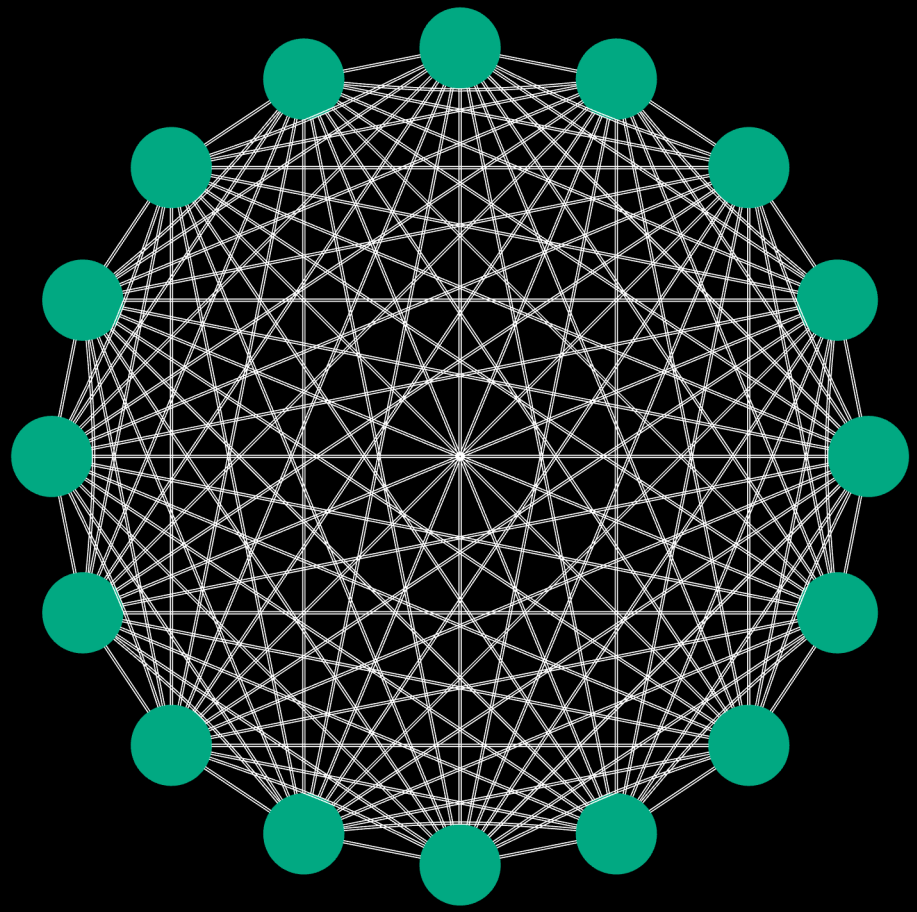

WHAT TOOLS ARE ALREADY BEING USED BY HPE?

- NetSim:
  - QEMU with custom modules to represent Rosetta and Cassini devices
  - Pros: Allows us to test the full stack on different hardware platforms, such as RISC-V
  - Cons: Very slow

- SST:
  - Used for Slingshot performance prediction
  - Pros: Allows us to experiment with different network topologies; comparing dragonfly and fat tree, for example.
  - Cons: Has some numerical problems that we are still tracking down. Not very accurate, only indicative.

- FPGAs:
  - Used internally to simulate individual hardware devices
  - Pros: Allows for rapid prototyping
  - Cons: Expensive, slow

NETWORK DIGITAL TWINS

- Existing tools are good, but limited

- Need something new
  - Ingest data from live systems
  - Build on existing tools?
  - Visual representation of communication overlayed on the 3D models

- Broader interest from across HPE:
  - Slingshot
  - 5G/6G
  - Photonics
  - Next Gen Ethernet & WiFi
  - Quantum networks

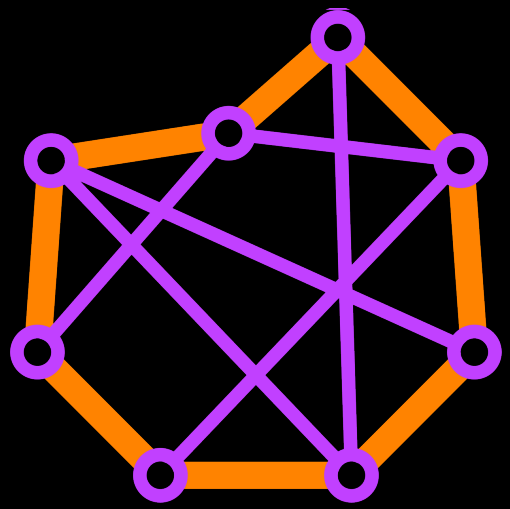

THANK YOU

HPE HPC & AI EMEA Research Lab | emearesearchlab@hpe.com

Agenda:

- Intro and Welcome
- Technical Talks: (20 min each)
  1. Matthias Maiterth (ORNL): On the motivation for developing a digital twin framework
  2. Adrian Jackson (EPCC): On the potential of using digital twins to improve overall system and data center efficiency
  3. Jess Jones (HPE): On EX system architectures and networking in the context of digital twins.
- **Community outreach / Audience discussion: (30 min)**
  perspectives, experiences, and discussion
  to contribute and collaborate on data-center digital twins.

Discussion Points:

- Show of Hands:
  - Whose organization has started a Digital Twin Project in the last year(s)?
  - Are you already involved in ExaDigiT / Would you like to be?
- Questions:
  - Use-cases at your organization?
  - Technology interest / Recent developments?
  - Usability / Users at your organization
  - Research Opportunities (by combining traditional work)
- Continued:
  - Analysis of past / present / future (simulation)
  - Collaborative Analysis
  - General Collaboration with other efforts
  - Standards?
  - Are we missing anything?
  - Open Points?

To Join our community:
Email: Wes Brewer - brewerwh@ornl.gov

Slack: https://exadigit.slack.com/