Oak Ridge National Laboratory

So…
You want to install an exscale supercomputer?

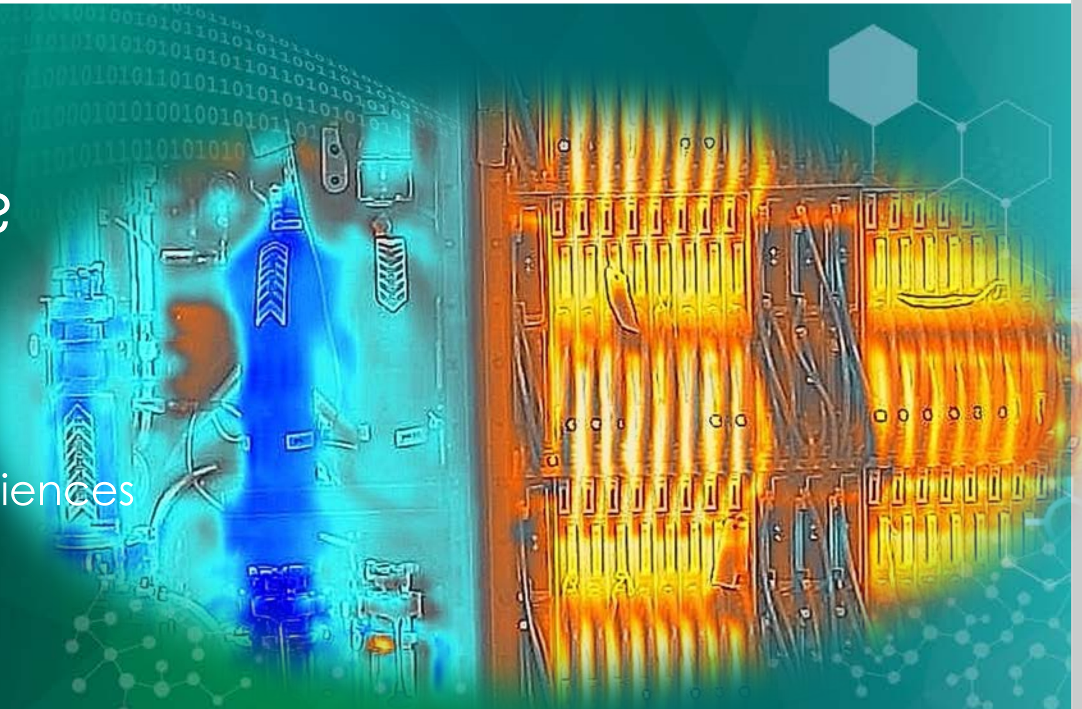

Jim Rogers

Director, Computing and Facilities
National Center for Computational Sciences
Oak Ridge National Laboratory

U.S. DEPARTMENT OF ENERGY

How hard can it be?

Buy exascale computer

Plug computer in

Do science!

Airplane! (1970)

The Exascale Timeline- Five years from RFP to Production

RFP Draft – Dec 2017 *

RFP Release – April 2018

Ordered - Nov 2018

Source Selection Announced
May 2019

Initial Delivery - Sept 2021

HPL – May 2022

Full Production - March 2023

** The RFP defined electrical and mechanical constraints, allowing long-lead facility work to commence in 2018.*

Fun Fact #1: Power at Exascale – New MV Distribution

*12.3.2. **The planned power budget for the CORAL system and associated I/O subsystem is constrained to no more than 40MW.** The electrical distribution method for the Offeror's compute solution should be based on 3-phase wye (Y) 277/480VAC. The maximum size of any circuit supplying a compute rack will be 200A.*

The target data center had 15MW of existing capacity.

Curing a potential 25MW deficit required two new 13,800V distribution feeders

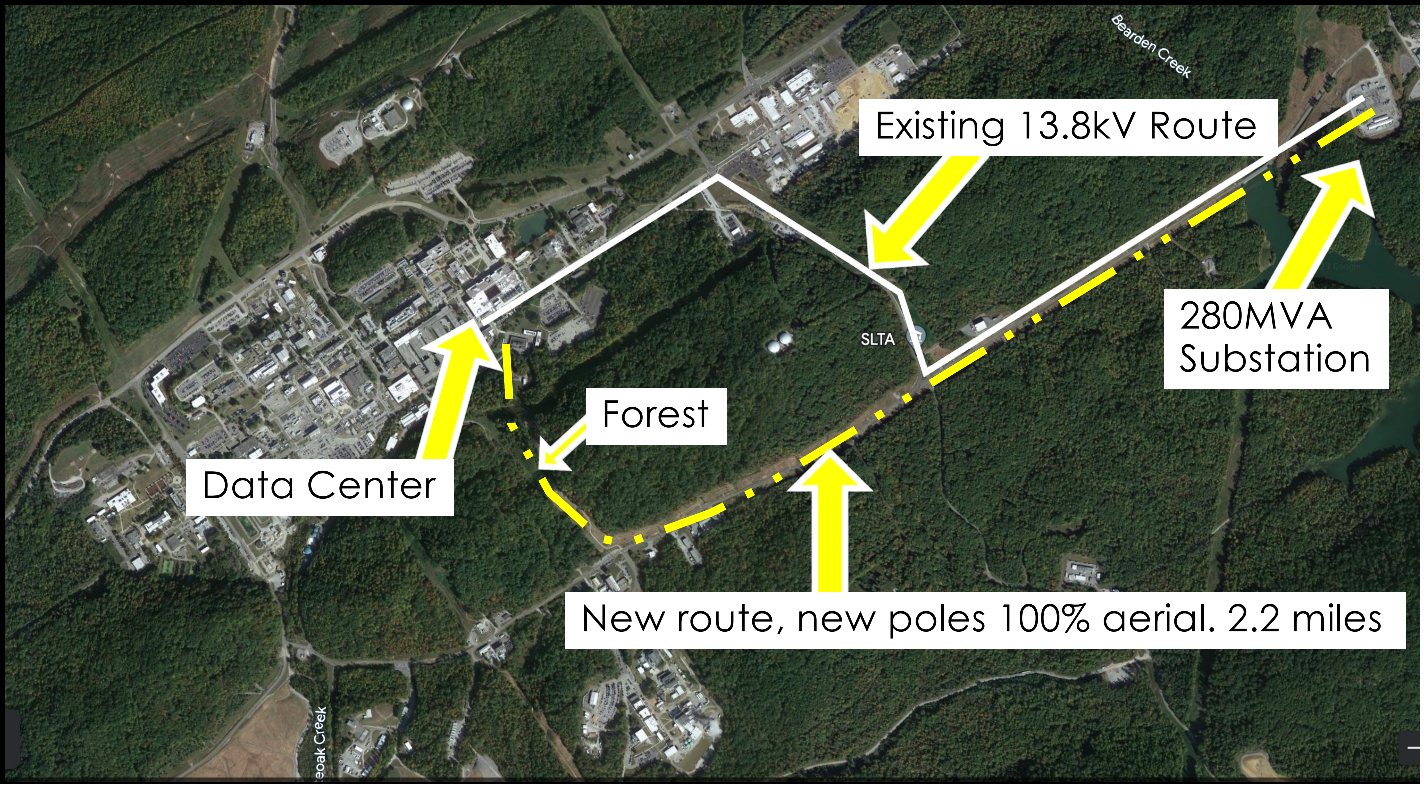

You cannot make this stuff up… Bat Mating Season

Evening Bat, Nycticeius humeralis

Description:
A small bat with thick, dark brown fur and black ears, muzzle, and wing and tail membranes.

Habitat:
Evening Bats mostly use tree cavities during the spring and summer months. Winter habitat is unknown, but accumulations of fat reserves suggest they migrate or hibernate. Nursery colonies are formed under the bark of dead trees or in houses.

Breeding information:
Breeding occurs in the fall, and delayed fertilization allows for the birth of 1-2 young in late May through June.

The 13.8kV power line schedule, with a significant utility corridor through the forest, had to carefully avoid the bat mating and nursery schedule.

New 13.8kV lines emerge at the Bldg 5600 Data Center

- ❑ Bats protected
- ❑ Insects eaten
  - ❑ a 100-bat colony consumes 1.25M insects each season
- ❑ Mud vanquished
- ❑ Power pole supply chain issues overcome
- ❑ Utility schedule saved

Electrical Facility Modifications Supporting Exascale

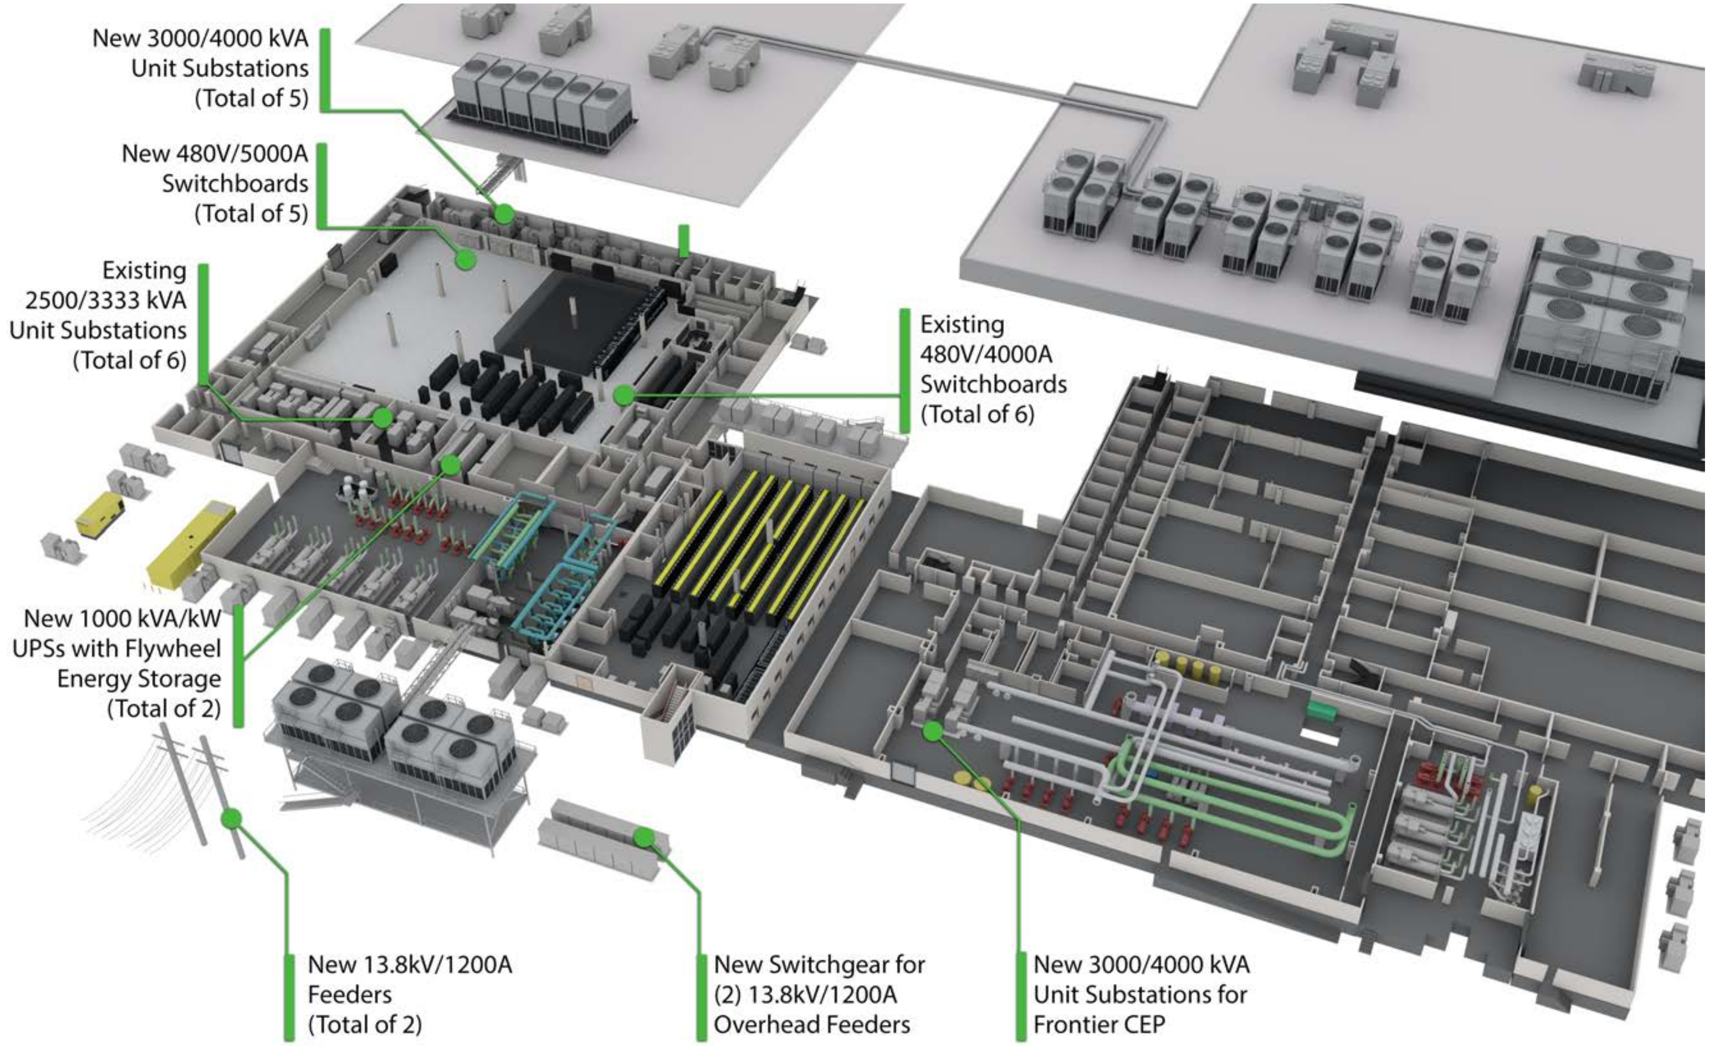

Fun Fact #2: Power at Exascale – Energy Consumption

- Population: 0
- Households: 0
- Land Area: 0.00014348 square miles
- Peak power demand: 29MW
- Average power demand: 18MW
- **Energy use (annual): 158M kW-hours**

- Population: 32,000
- Households: 12,900
- Land Area: 85.25 square miles
- Energy use (annual, per household): 12,200
- **Energy use (annual): 158M kW-hours**

Fun Fact #3: Power at Exascale – Load Swings

ORNL's Frontier HPL Result June 2022: Power Profile

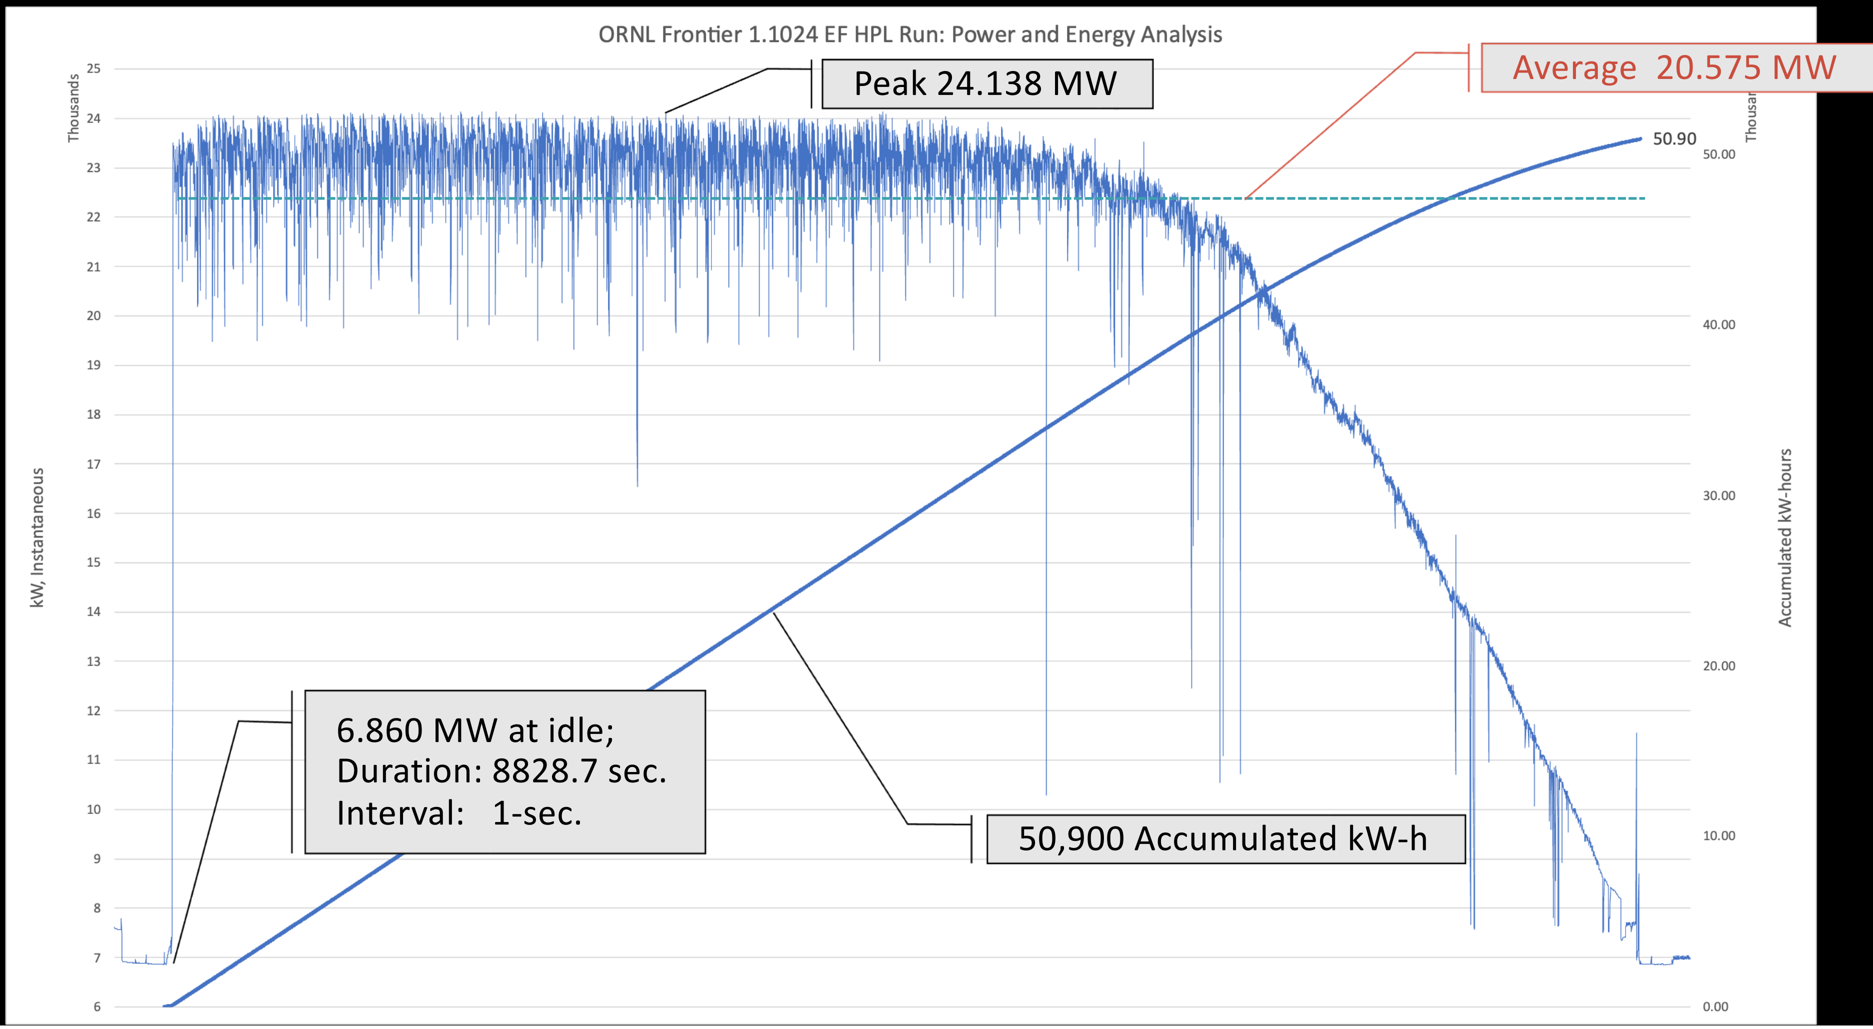

ORNL's Frontier HPL Result June 2022: Power Profile

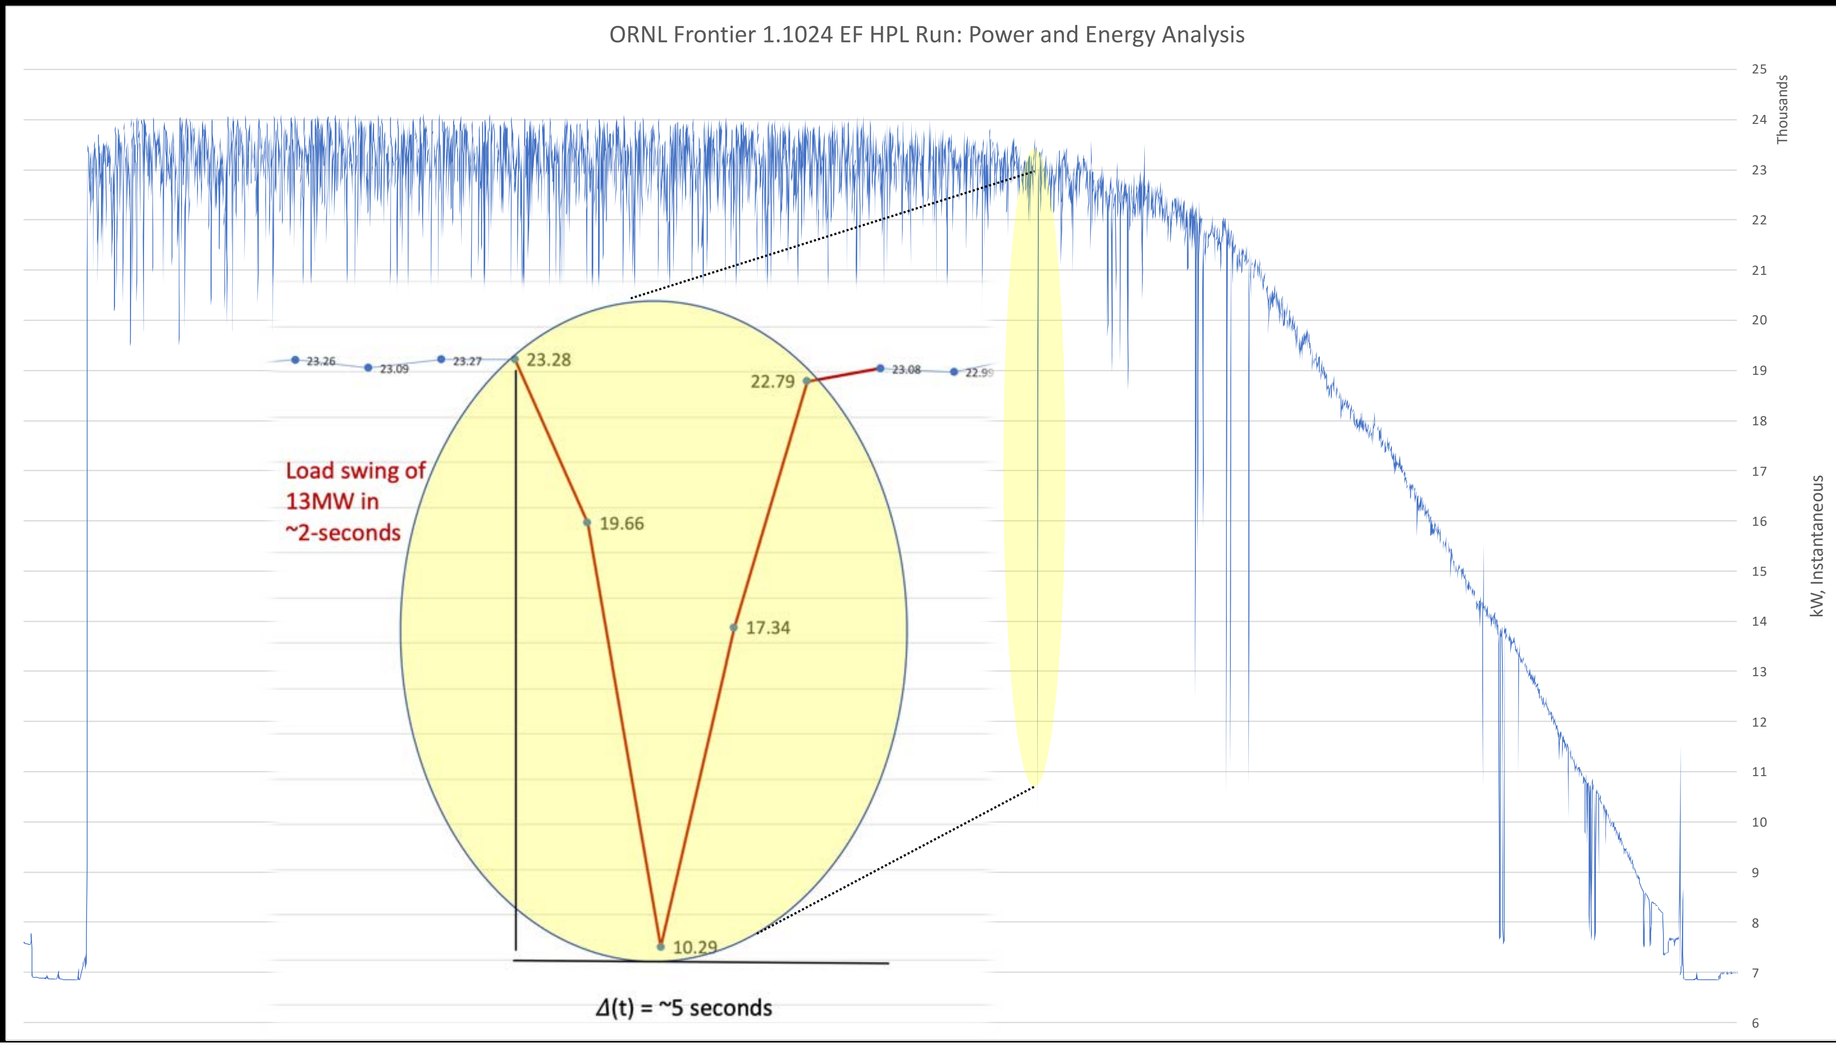

ORNL's Frontier HPL Result June 2022: Power Profile

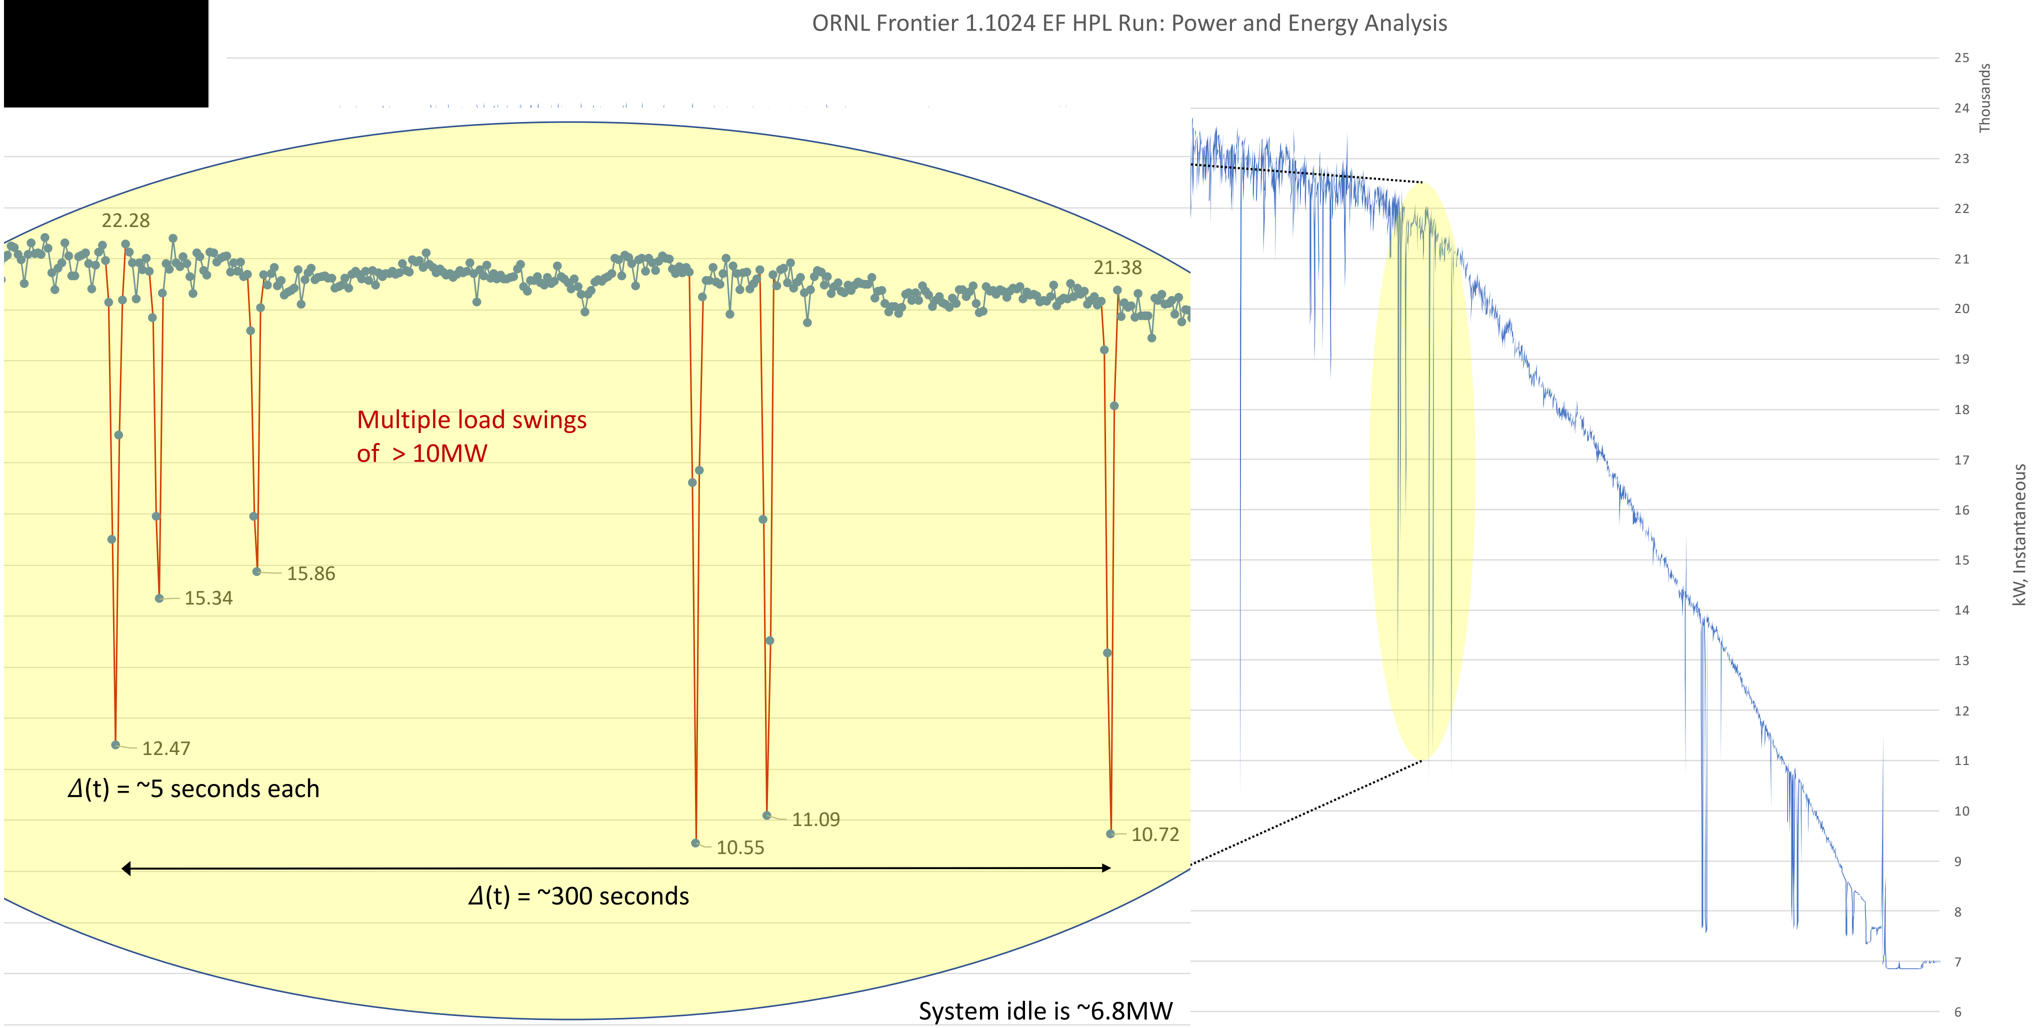

Fun Fact #4: Power at Exascale – Power Factor Correction

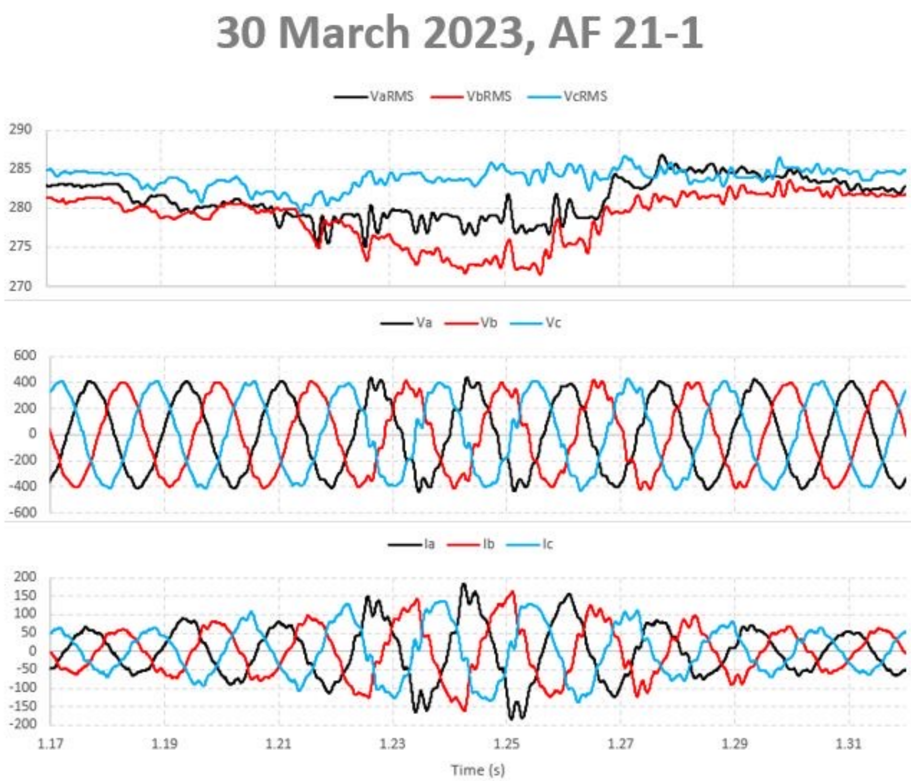

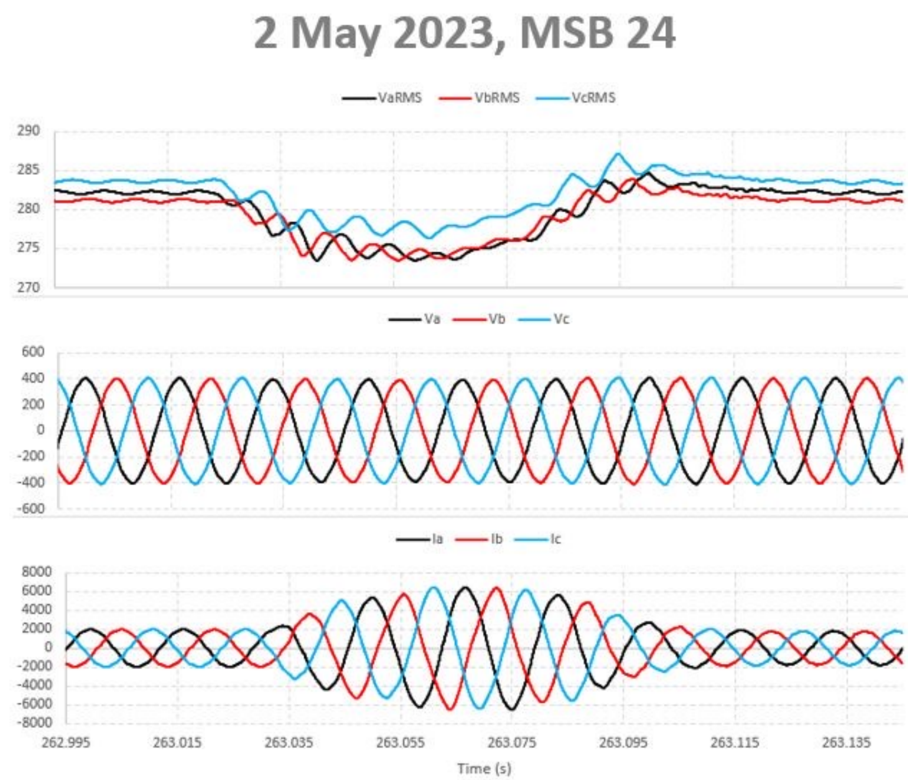

The Cholla application (astrophysical hydrodynamics), running on Frontier. It has sharp and periodic power demands (synchronized GPU workloads) that are a good test case for this issue.

OAK RIDGE
National Laboratory

Fun Fact #4: Power at Exascale – Power Factor Correction

Waveform (single-phase) before and after revised power factor correction

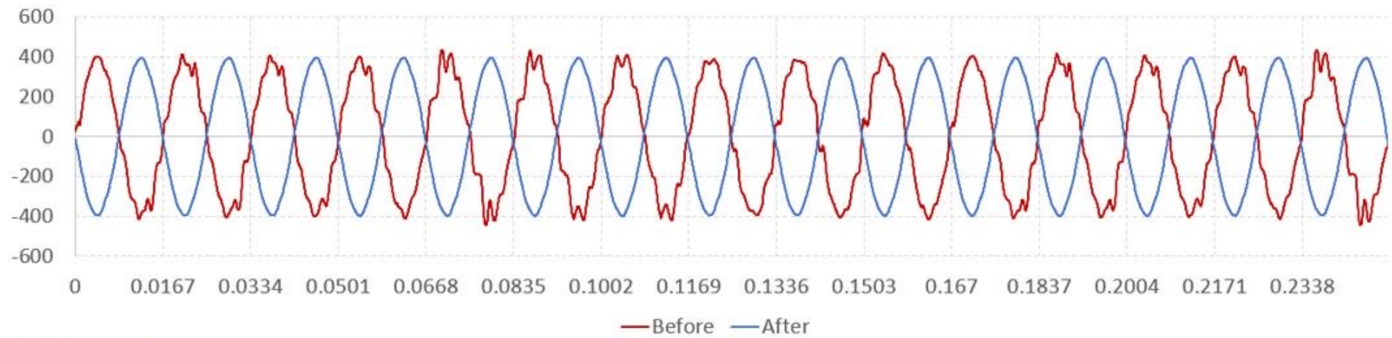

%V_{THD}

Harmonics before and after revised power factor correction

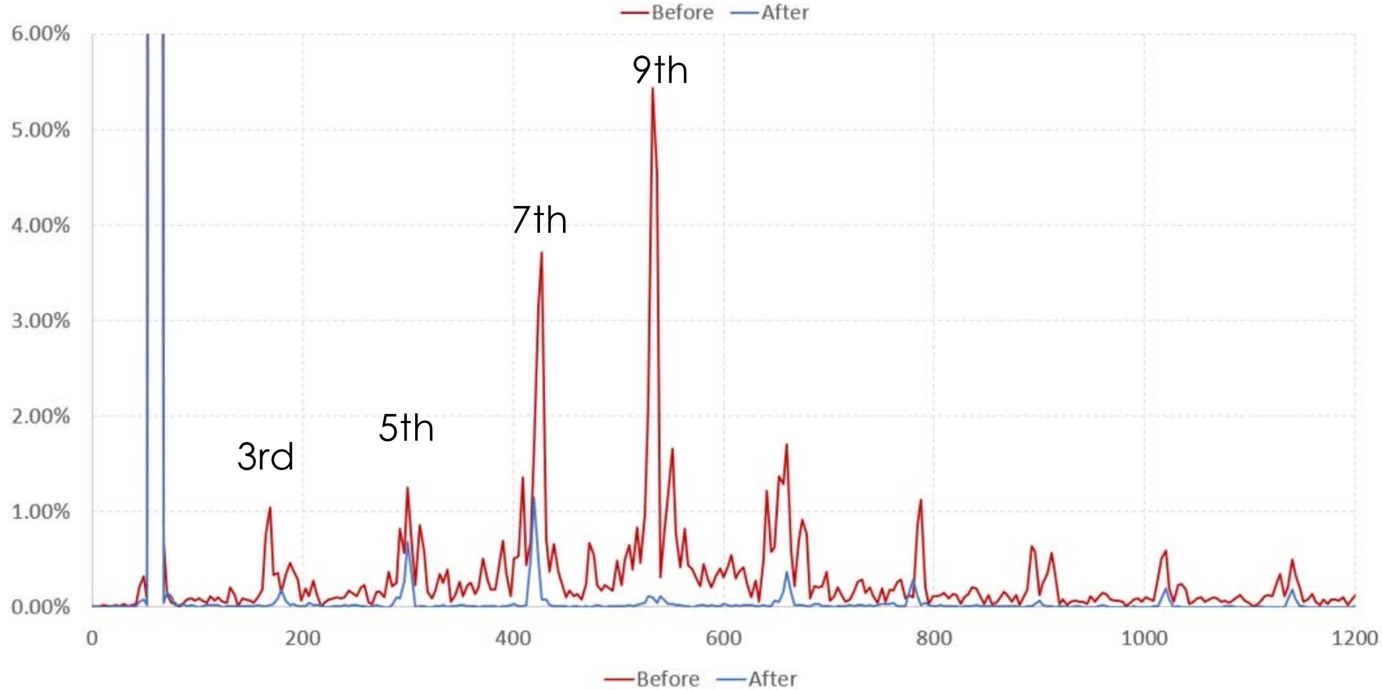

Oak Ridge National Laboratory’s Journey from Petascale to Exascale

National Avg 2011 PUE: 1.8 ▲

10^{15}

2009

Jaguar
2.3 PF
Multi-core CPU
7 MW

Air-cooled System

Room CRAC

Facility: Chilled Water

PUE 1.4

10^{16}

2012

Titan:
27 PF
Hybrid GPU/CPU
9 MW

Hybrid Air/Refrigerant HX w/ Dedicated CDUs

~Room-Neutral

Facility: Chilled Water

PUE 1.29

10^{17}

2017

Summit
200 PF
Hybrid GPU/CPU
13 MW

Hybrid Air/ Cold-Plate

95% Room-Neutral

Facility: Warm Water (21°C) with CW trim

PUE 1.03-1.1

10^{18}

2021

Frontier
2,000 PF
Hybrid GPU/CPU
29 MW

Cold Plate HX w/Dedicated CDUs

99% Room-Neutral

Facility: High-Temp Water to 32°C. No CW trim

PUE ~1.03

Fun Fact #2: Cooling at Exascale

12.3.3. Offeror will use a facility water supply temperature described by ASHRAE Technical Committee 9.9 Liquid Cooling Guidelines, as Liquid Cooling Class W3 (32C), or higher for the compute racks.

- With a 40MW design limit, the worst-case cooling demand is ~12,000 tons.
- ORNL *chose* evaporative cooling for discharging waste heat from the high temperature water loop (that includes Frontier).
  - Makeup Water – 200,000 gallons/day
- ORNL *chose* variable HTW supply, which can fluctuate to 32C based on outside conditions
  - Summit is fixed MTW supply at 21C

12,000 tons of evaporative cooling, perched high above the mechanical plant

Fun Fact #3: Cooling at Exascale –Deliver Steady HTW ST

- Primary-side controls predictively stage evaporative cooling towers up/down based on load
    - Fast response to add ECT capacity
    - Slow response to take ECT capacity away
- MILES of piping and 10,000 gallons of water in-flight significantly affect throttle control / response

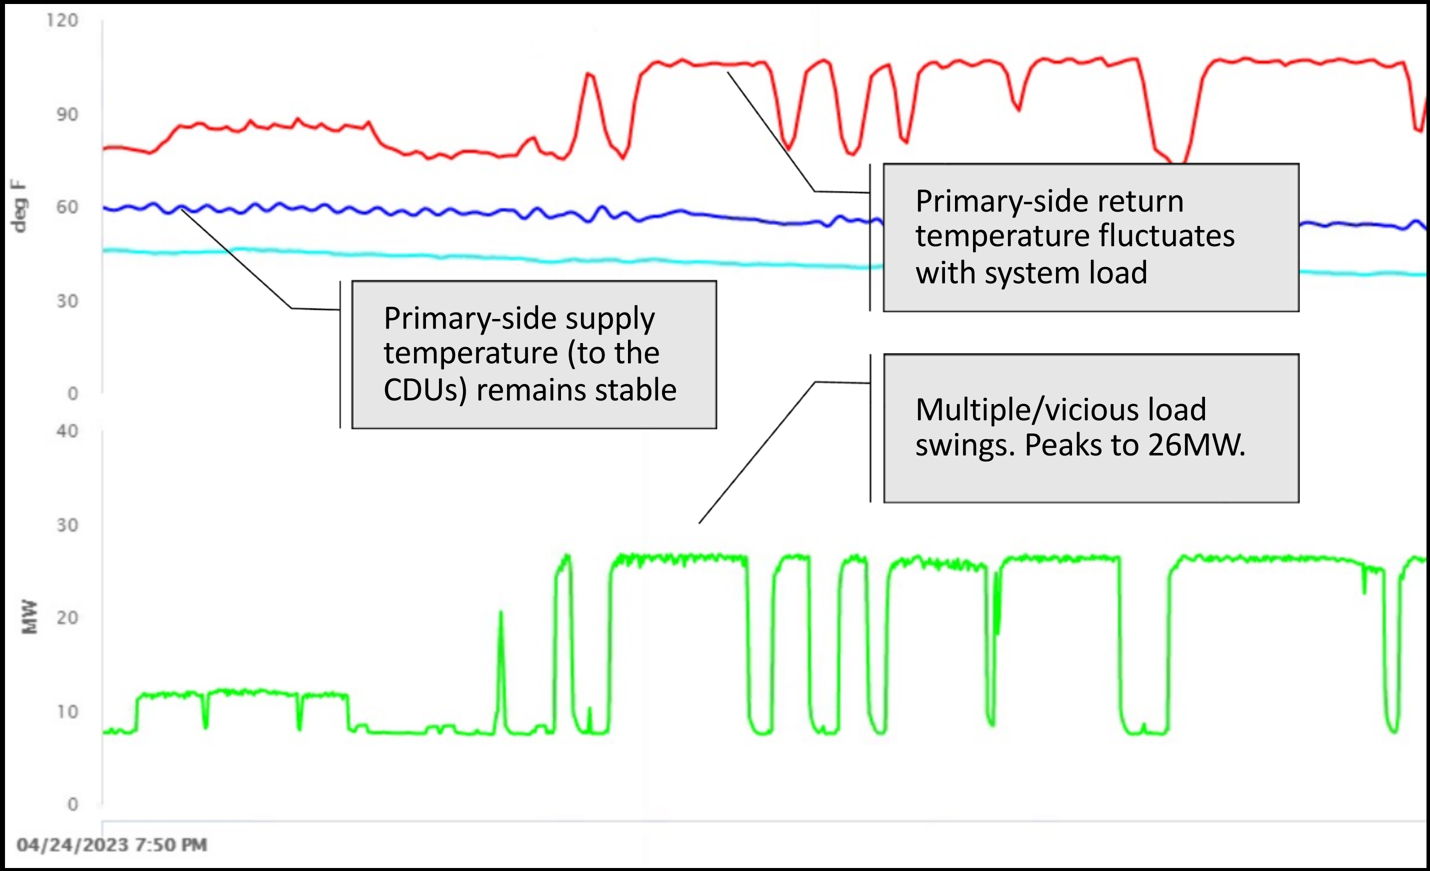

Fun Fact #4: Cooling at Exascale – PUE

- *Sustainability objective: minimize energy use*
- We're reached the point of diminishing returns from a mechanical perspective.
- HPE packaging already returns 95% of the waste heat directly to water.
- HX between the HPE (secondary) and ORNL (primary) eject the heat using just pump and fan energy.
- Frontier's real time PUE, including all of the mechanical systems, is below 1.03 during load.

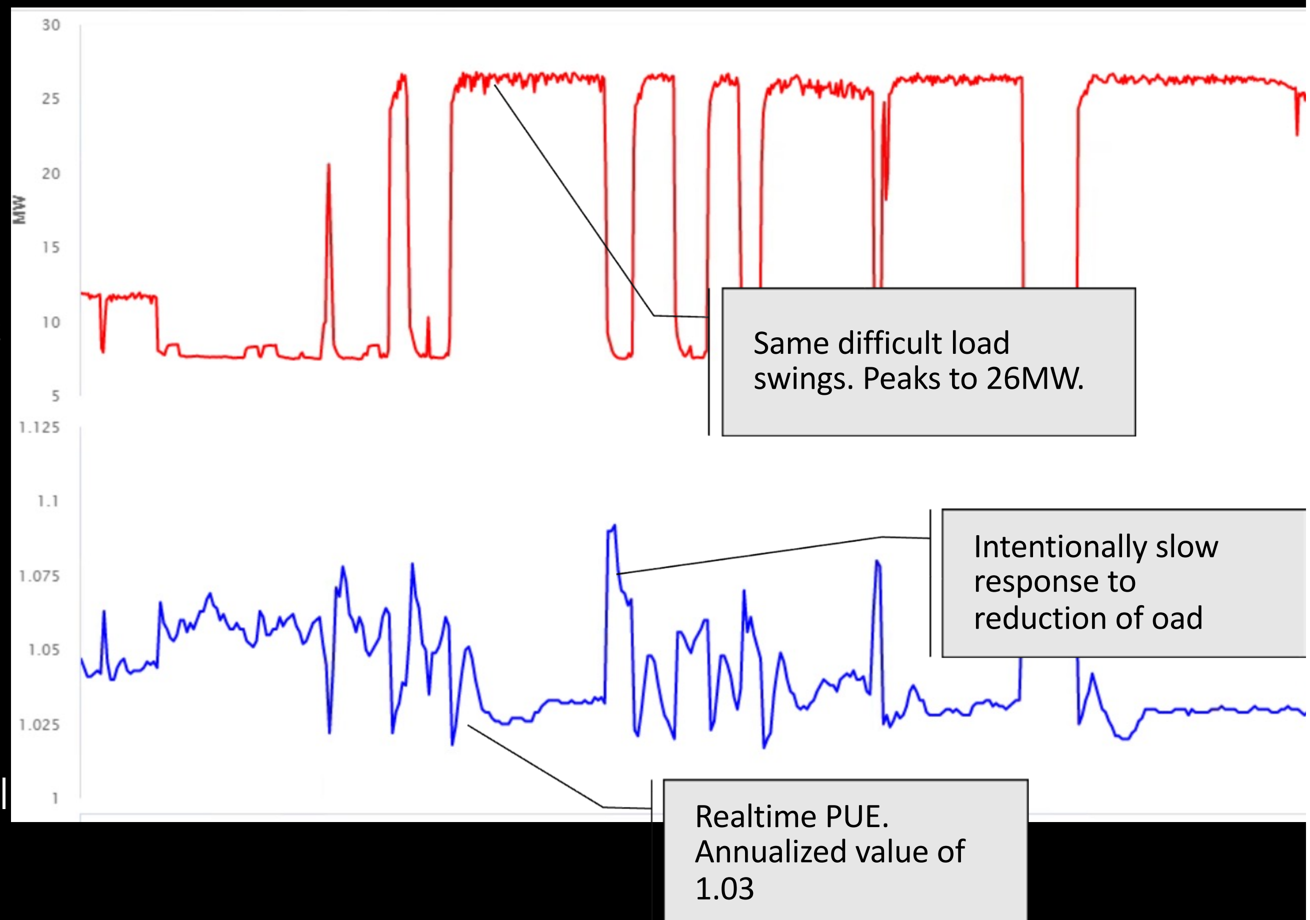

Fun Fact #5: Cooling at Exascale – Pressure Testing

- The primary-side supply and return system was pressure-tested to well-above100psi, >2x the anticipated worst case.
- A flawed weld in the Aquatherm piping failed during testing at ~50psi, spilling several thousand gallons of water on the roof of the mechanical plant.
- New PPE was issued.

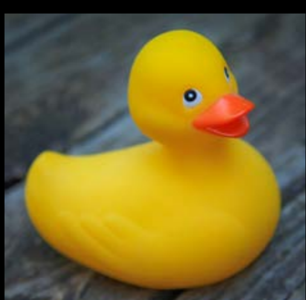

Frontier as a Machine for Science: How To Build an Exascale Computer and Why You Would Want To

Bronson Messer
Director of Science

Leadership Computing Facility
Oak Ridge National Laboratory

U.S. DEPARTMENT OF ENERGY

Doing science is the ultimate aim of Frontier

- Exascale will enable a new classes of simulations across scientific fields.
- This will require work
  - Much has already been done via ECP and other support
  - But, much remains in these and many cases
- GPU acceleration and scalability are requirements
- But, actually *doing* computational science goes beyond these straightforward (if challenging!) needs

Gordon Bell Award Finalist using the First Exascale Supercomputer

April-July 2022: WarpX on **world's largest HPCs**

L. Fedeli, A. Huebl et al., SC'22, 2022

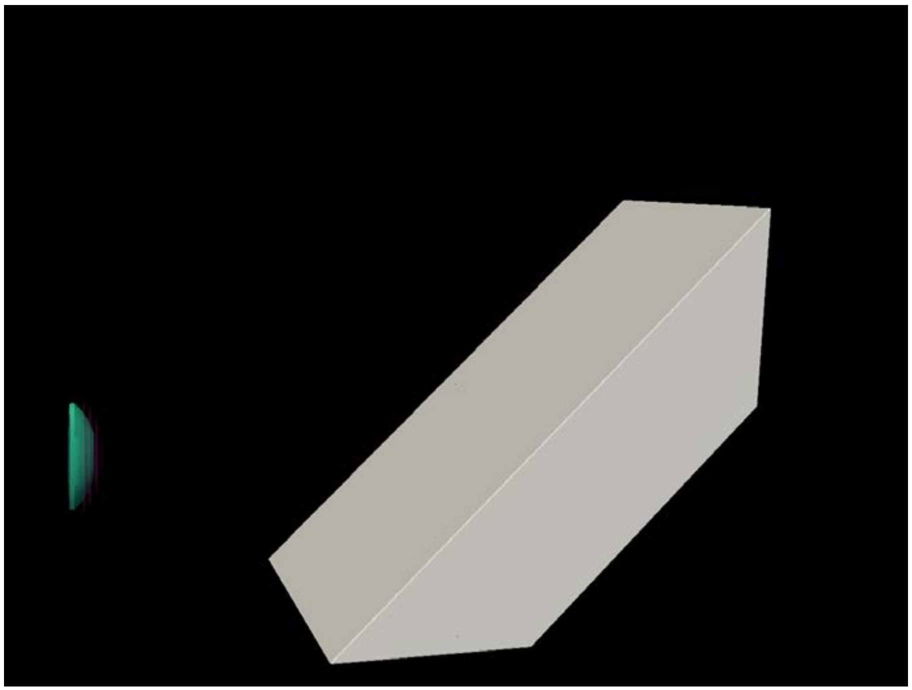

Novel hybrid solid-gas target concept

Our simulations demonstrated that the new concept leads to unprecedented beam quality using a PW-class laserl, and are supporting experiments at LOA (Ecole Polytechnique, France) to validate the new concept.

Success story of a multidisciplinary, international multi-institutional team!

Large Scale Density Functional Theory at the Exascale with LSMS

Workflows and high performance computations to predict materials properties

Research Topics

- Understanding the role of disorder and defects in materials for electronic and mechanical properties
- Complex magnetic order – topological magnetic structures (e. g. Skyrmions) and magnetism beyond ideal crystal

Recent Highlights

- Successful porting of the LSMS code (github.com/mstsuite/lsms) to Frontier for exascale materials simulations.
- **Scaling of first principles calculations to O(100,000) up to O(1,000,000) atoms for the first time.**
- Demonstrated scaling of LSMS on Frontier up to 1,048,576 atom FePt system on 8192 Frontier nodes.
- Speedup of LSMS from Summit to Frontier from combined hardware and software improvements is ~8x

Future work

- Capabilities for non-metallic quantum materials
- Calculation of forces for ab-initio relaxation and first-principles molecular dynamics.

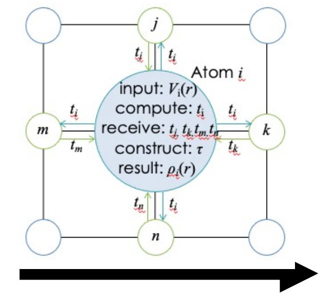

Moving from CUDA to HIP

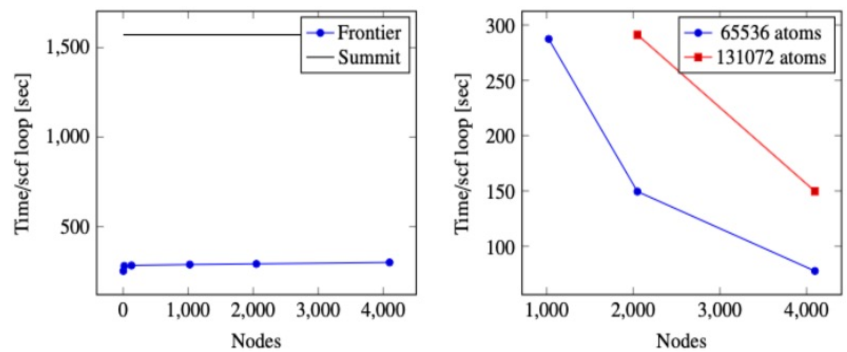

Weak (left) and strong (right) scaling results of LSMS for FePt calculations on Frontier

OAK RIDGE National Laboratory

CoMet for correlation analysis on Frontier

- Comet is used to compute similarity metrics from large datasets in genomics, climate, and other fields
- Currently being used to analyze the geospatial and temporal evolution of SARS-CoV-2 variants
- CoMet has achieved up to **6.6 ExaFlops** mixed precision performance on Frontier (3-way DUO method)

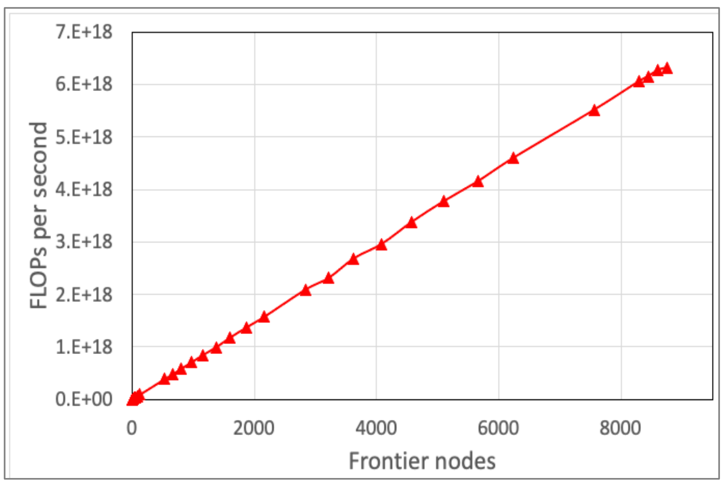

Geospatial 7-day moving average of SARS-CoV-2 genome sequences by strain

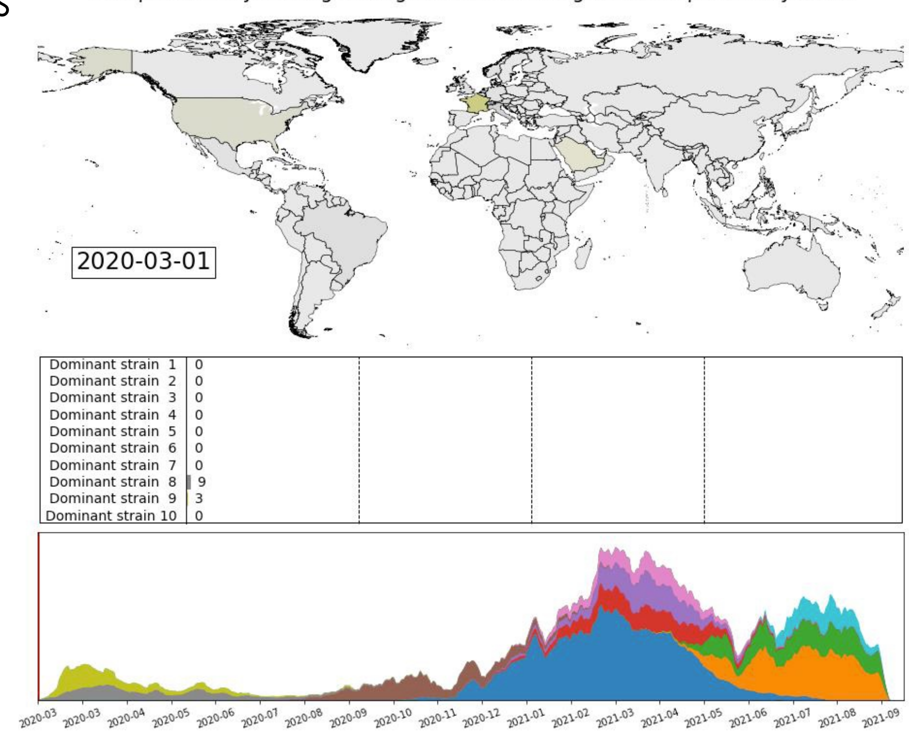

An ill-posed question, but…

- **What is the "killer app" for exascale computing?**
- I don't like this question. A lot of the value of unique supercomputing facilities is the ability to impact a huge variety of scientific problems.
- But, people ask it…
- Maybe there's not a killer app, but there is a ubiquitous physical problem that requires:
  - More memory (i.e. resolution)
  - Faster compute speed
  - Inclusion in multiphysics simulations…

Understanding turbulence

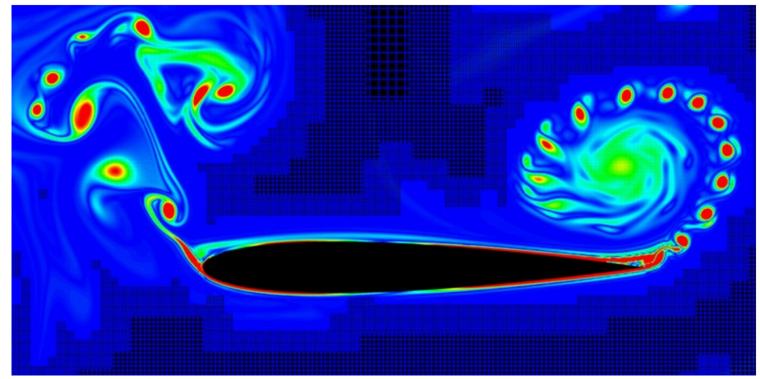

- “The last great unsolved problem in classical physics” (One of the 7 Millennium Problems)
- Werner Heisenberg assuredly never said: “When I meet God, I’m going to ask him two questions: why relativity? And why turbulence? I really believe he’ll have an answer for the first.”
- We remain far away from being able to resolve turbulent physics from the largest scales where it is generated (even in terrestrial settings) to the molecular dissipation scale.
- But, there are many places where turbulence is important where other physics arrests the impact of the turbulent cascade before it gets to the smallest scales.

Reaching New Heights in Weather Forecasting's Exascale Future

ECMWF and ORNL researchers use the power of Summit to simulate the Earth's atmosphere for a full season at 1-square-kilometer grid-spacing

The Science

Using Summit, a team of researchers from ECMWF and ORNL achieved a computational first: a global simulation of the Earth's atmosphere at a 1-square-kilometer average grid-spacing for a full 4-month season. Completed in June, the milestone marks a big improvement in resolution for the "European Model," which currently operates at 9-kilometer grid-spacing for routine weather forecast operations. It also serves as the first step in an effort to create multi-season atmospheric simulations at high resolution, pointing toward the future of weather forecasting—one powered by exascale supercomputers.

The Impact

The team has made the simulation's data available to the international science community. By eliminating some of the fundamental modelling assumptions prevalent in conventional simulations, the high-resolution data may help to improve model simulations at coarser resolutions.

PI(s)/Facility Lead(s): Nils Wedi, ECMWF
ASCR Program/Facility: INCITE/OLCF
ASCR PM: Christine Chalk
Publication(s) related to this work: Wedi, N. P., et al. A baseline for global weather and climate simulations at 1 km resolution. *Journal of Advances in Modeling Earth Systems*, 12 (2020), e2020MS002192. doi: 10.1029/2020MS002192

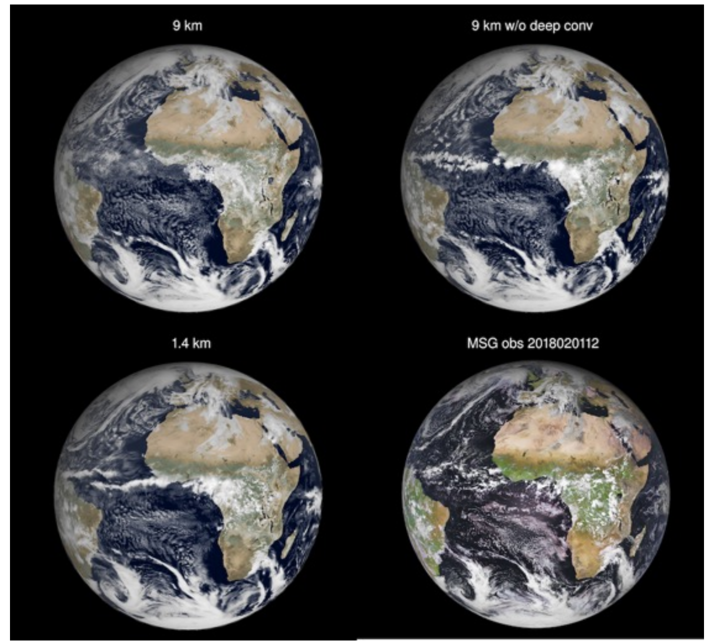

These simulated satellite images of Earth show the improvement in resolution of the ECMWF Integrated Forecasting System from 9-kilometer grid-spacing with parametrized deep convection (top left), to 9-kilometer grid-spacing (top right), and 1-kilometer grid-spacing (bottom left). On the bottom right is a Meteosat Second Generation satellite image at the same verifying time. Image courtesy ECMWF.

OAK RIDGE National Laboratory
ECMWF
NVIDIA

Closing In on Fusion

A team modeled plasma turbulence on the nation's fastest supercomputer to better understand plasma behavior

The Science

The same process that fuels stars could one day be used to generate massive amounts of power here on Earth. Nuclear fusion—in which atomic nuclei fuse to form heavier nuclei and release energy in the process—promises to be a long-term, sustainable, and safe form of energy. But scientists are still trying to fine-tune the process of creating net fusion power. A team led by computational physicist Emily Belli of General Atomics has used Summit supercomputer at the Oak Ridge Leadership Computing Facility to simulate energy loss in fusion plasmas. The team used Summit to model plasma turbulence, the unsteady movement of plasma, in a nuclear fusion device called a tokamak. The team's simulations will help inform the design of next-generation tokamaks like ITER—the world's largest tokamak, which is being built in the south of France—with optimum confinement properties.

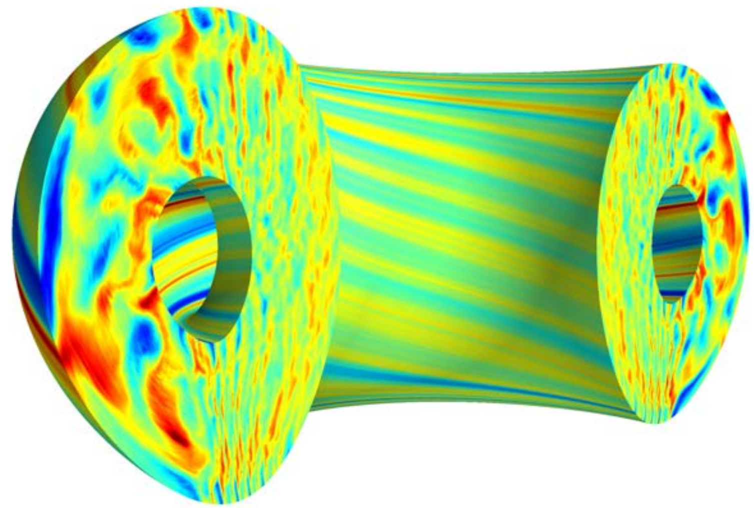

A visualization of deuterium-tritium density fluctuations in a tokamak driven by turbulence. Areas of red are representative of high density and areas of blue are representative of low density. Image Credit: Emily Belli, General Atomics

The Impact

Until now, almost all fusion simulations have only included only deuterium or tritium isotopes, but Summit enabled the team to include both as two separate species, model the full dimensions of the problem, and resolve it at different time and spatial scales. The results provided estimates for the particle and heat losses to be expected in future tokamaks and will help scientists and engineers understand how to achieve the best operating scenarios in real-life tokamaks.

PI(s)/Facility Lead(s): Emily Belli
ASCR Program/Facility: ALCC and INCITE / OLCF
ASCR PM: Christine Chalk
Publication(s) for this work: Emily A. Belli and Jeff Candy, "Asymmetry between Deuterium and Tritium Turbulent Particle Flows," *Physics of Plasmas* 28, no. 6 (2021), doi:10.1063/5.0048620.

ECP EXAALT MD on Frontier

- Damaged surface of Tungsten in conditions relevant to plasma facing materials in fusion reactors
  - 100,000 atoms
  - T=1200K
- ML-based materials model: Spectral Neighborhood Analysis Potential (SNAP) with 205 features per atom
- Used the **Sub-Lattice Parallel Trajectory Splicing** *(SL-ParSplice)* accelerated molecular dynamics method [1,2]
- ParSplice allows for simulations to be parallelized in the *time domain*. Many MD simulations are executed *concurrently* and carefully spliced into a single *dynamically-correct* trajectory [3].
- Unique capability for long-timescale simulation of materials

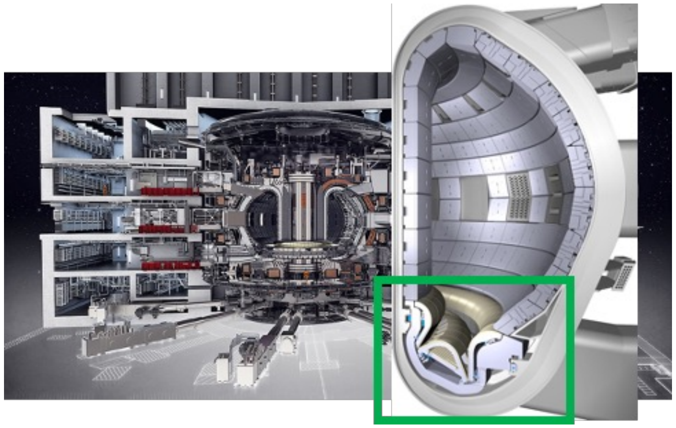

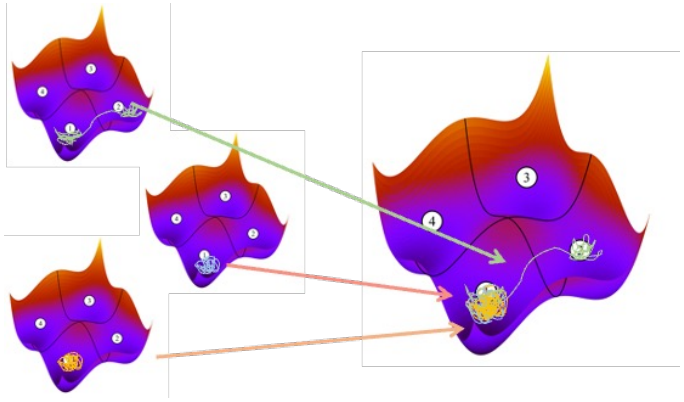

[1] Long-time dynamics through parallel trajectory splicing, *Journal of chemical theory and computation* 12 (1), 18-28
[2] Sublattice parallel replica dynamics. *Physical Review E*, *89*(6), p.063308.
[3] A mathematical formalization of the parallel replica, Monte Carlo Methods and Applications 18 (2), 119-146

Mapping MD onto Frontier

- Simulation executed using the EXAALT code on 7000 Frontier nodes
- EXAALT provides a hierarchical framework for task and data management
  - ~1% of resources for management
  - ~99% of resources to simulation engine
- Infrastructure re-assigns MD tasks to workers every ~7 seconds
- ~170 instances of each sub-domain execute concurrently

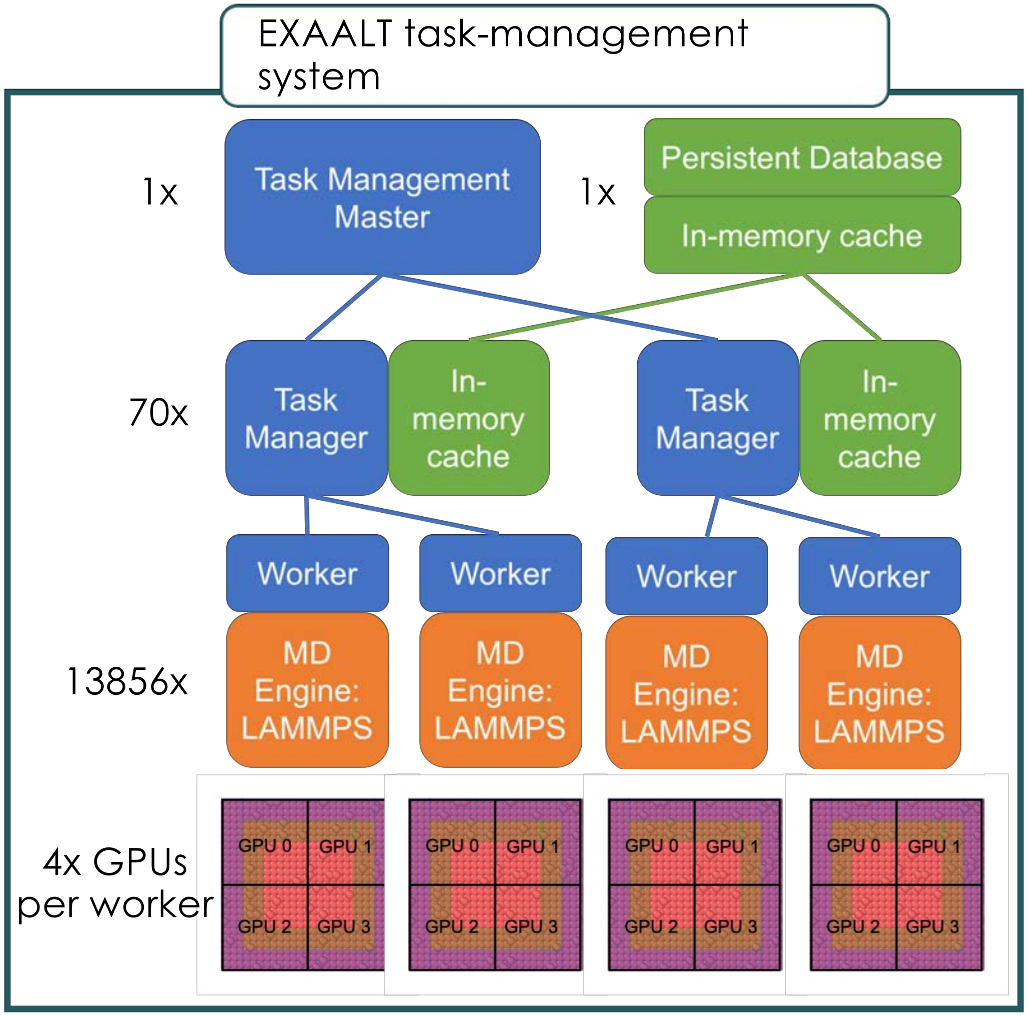

72 nodes for data and task management
6928 nodes for MD simulations

GE Spins up Supercomputer Models to Zero in on Energy Loss in Turbines

A team at GE Aviation and the University of Melbourne is studying turbulent flows on the Summit supercomputer for better engines

The Science

High-pressure turbines are vital components of gas turbines used to propel jet engines. The more efficient these jet engines are, the better they are for the aircraft industry and their customers. But these large, dynamic systems are difficult to study via experiments and physical testing. A team led by scientists at General Electric (GE) Aviation and the University of Melbourne used the Summit supercomputer to run for the first time real-engine cases capturing the largest eddies, or circular fluid movements, down to those that were tens of microns away from the turbine blade surface. From the simulations, the researchers determined which regions near a turbine blade experience a greater loss of energy. For the case with the highest Mach number, which describes the flow's velocity compared with the speed of sound, they discovered an extra loss of energy resulting from strong shock waves, or violent changes in pressure, that interact with the edge and wake of the flow to cause a massive amount of turbulence.

The Impact

More accurate prediction of real-engine conditions will lead to more efficient engines that consume less fuel and other positive derivative effects. A 1 percent reduction in fuel consumption across a fleet of engines is equal to about 1 billion dollars a year in fuel cost savings. Reduced fuel consumption also translates into reduced emissions—a 1 percent reduction in fuel burn reduces CO_2 emissions by roughly 1.5 percent.

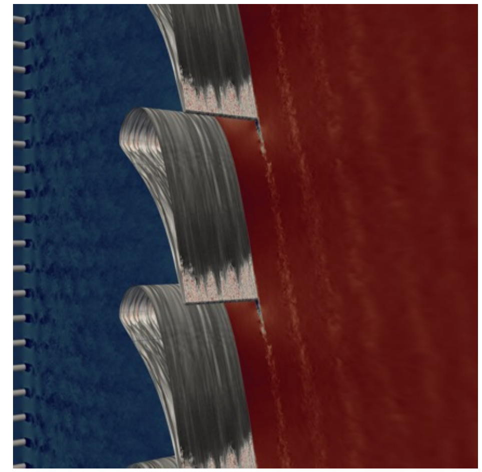

A row of upstream bars produces highly turbulent flow that gets accelerated through a high-pressure turbine blade row and interacts with the blade surface, causing significant temperature variations. Image Credit: Richard Sandberg, University of Melbourne

PI(s)/Facility Lead(s): Richard Sandberg, Univ. Of Melbourne; Sriram Shankaran, GE Aviation
ASCR Program/Facility: INCITE/OLCF
ASCR PM: Christine Chalk
Publication(s) for this work: Y. Zhao and R. D. Sandberg, "High-Fidelity Simulations of a High-Pressure Turbine Vane," *Journal of Turbomachinery* 143, no. 9 (2021).

Y. Zhao and R. D. Sandberg, "Using a New Entropy Loss Analysis," *Journal of Turbomachinery* 142, no. 8 (2020): 081008

Some quotes from early users

- “I’ve been impressed at just how smooth Frontier is. Things are just working. Good job all around.”
- “The machine is a thing of beauty… This is just sick what this machine can do.”


OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE
FOR U.S. DEPARTMENT OF ENERGY