OAK RIDGE National Laboratory

Enabling Biomedical Science with Leadership Computing

Georgia Tourassi, PhD
National Center for Computational Sciences
Oak Ridge National Laboratory

Advancing Medical Care through Discovery in the Physical Sciences
Joint DOE-NIH Workshop
July 12-13, 2021

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

U.S. DEPARTMENT OF ENERGY

What is a Leadership Computing Facility?

- Established in 2004 as part of the Department of Energy High-End Computing Revitalization Act of 2004 (Public Law 108-423)
- A collaborative, multi-lab initiative funded by DOE's ***Advanced Scientific Computing Research*** program
- Operates as **one facility** with two centers, at Argonne and at Oak Ridge National Laboratories
- **Mission** is to provide an ecosystem, including partnering opportunities, that enables unsurpassed capability computing opportunities and the associated science and engineering breakthroughs
- Deploys and operates **advanced architectures** that are 10-100 times more powerful than other systems typically available for open scientific research

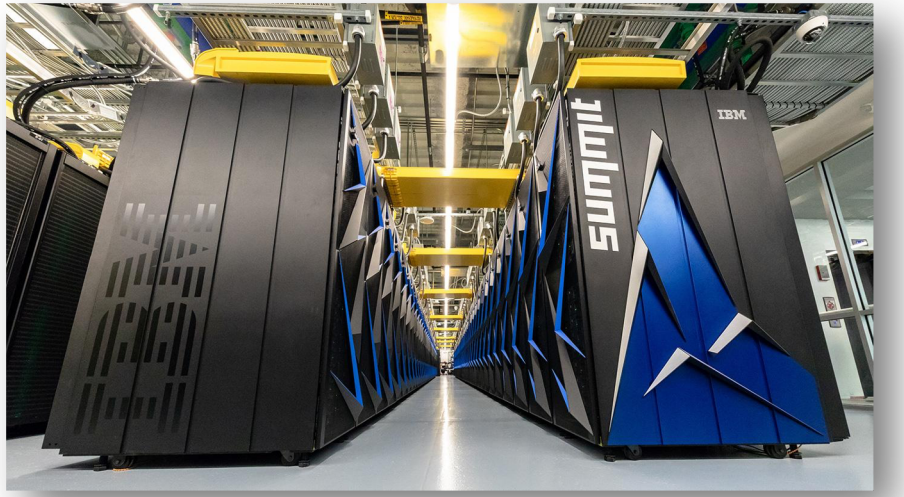

200 PF Peak Performance

OAK RIDGE National Laboratory

14 PF Peak Performance

Argonne NATIONAL LABORATORY

OAK RIDGE National Laboratory

Oak Ridge Leadership Computing Facility: Advancing Science at Leadership Scale

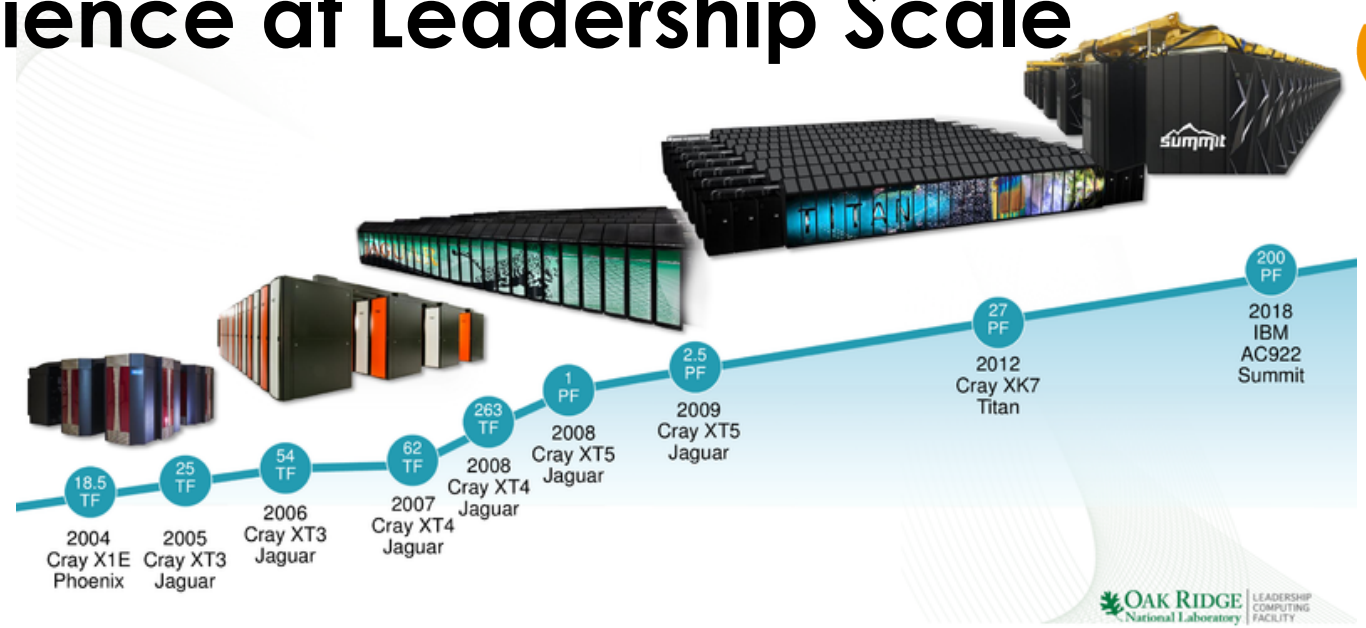

TOP 500 The List.

#1 through mid-2020

Frontier – forthcoming Fall 2021

OLCF has systematically delivered a series of leadership-class systems over the years.

New science challenges for a smart Supercomputer will exploit data and workflows.

Identifying Next-generation Materials
By training AI algorithms to predict material properties from experimental data, longstanding questions about material behavior at atomic scales could be answered for better batteries, more resilient building materials, and more efficient semiconductors.

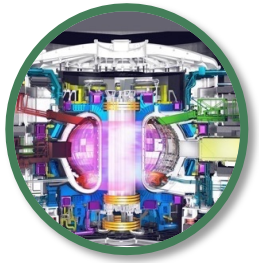

Predicting Fusion Energy
Predictive AI software is already helping scientists anticipate disruptions to the volatile plasmas inside experimental reactors. Summit's arrival allows researchers to take this work to the next level and further integrate AI with fusion technology.

Deciphering High-energy Physics Data
With AI supercomputing, physicists can lean on machines to identify important pieces of information—data that's too massive for any single human to handle and that could change our understanding of the universe.

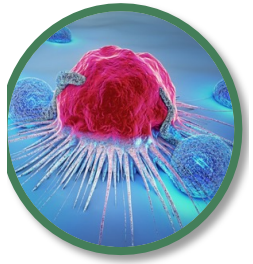

Combating Cancer
Through the development of scalable deep neural networks, scientists at the US Department of Energy and the National Cancer Institute are making strides in improving cancer diagnosis and treatment.

ORNL Summit System Overview

System Performance

- Peak of 200 Petaflops (FP_{64}) for modeling & simulation
- Peak of 3.3 ExaOps (FP_{16}) for data analytics and artificial intelligence

The system includes

- 4,608 nodes
- Dual-rail Mellanox EDR InfiniBand network
- 250 PB IBM file system transferring data at 2.5 TB/s

Each node has

- 2 IBM POWER9 processors
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory (96 GB HBM2 + 512 GB DDR4)
- 1.6 TB of NV memory

Core Capabilities Beyond Compute Technology

- **Data Management, Sharing, Analysis, and Movement**
  - Data repositories
  - Data portals
  - Campaign and community file systems
  - Supporting data analysis clusters
  - High-speed networking via ESNet
- **Sensitive Data**
  - Host and work with ITAR, HIPAA, and other sensitive data sets on Summit via the CITADEL framework
- **Computational Science Expertise**
  - Experts in domain science, performance engineering, AI/ML/DL, computer architectures
  - Testbed systems to explore and plan for the future
- **Strategic Partnership Programs**
  - Capability to procure and host HPC systems for other federal agencies

Diverse Application Portfolio of Strategic Importance

National security

Next-generation, **stockpile stewardship** codes

Reentry-vehicle-environment simulation

Multi-physics science simulations **of high-energy density physics** conditions

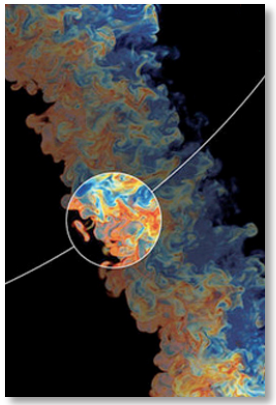

Energy security

Turbine **wind plant** efficiency

Design and commercialization of **SMR**s

Nuclear fission and fusion reactor **materials design**

Subsurface use for **carbon capture**, petroleum extraction, waste disposal

High-efficiency, low-emission **combustion engine** and gas turbine design

Scale up of **clean fossil fuel** combustion

Biofuel catalyst design

Economic security

Additive manufacturing of qualifiable metal parts

Reliable and efficient planning of the **power grid**

Seismic hazard risk assessment

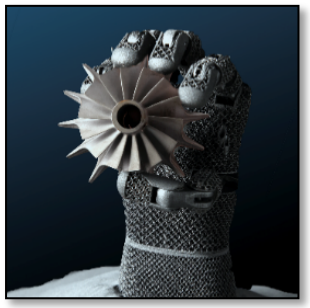

Scientific discovery

Cosmological probe of the standard model of particle physics

Validate fundamental laws of nature

Plasma wakefield accelerator design

Light source-enabled **analysis of protein and molecular structure** and design

Find, predict, and control materials and properties

Predict and control **magnetically confined fusion plasmas**

Demystify **origin of chemical elements**

Earth system

Accurate regional impact assessments in **Earth system models**

Stress-resistant crop analysis and catalytic conversion of **biomass-derived alcohols**

Metagenomics for analysis of biogeochemical cycles, climate change, environmental remediation

Biomedicine

Accelerate and translate **cancer research** (partnership with NCI)

The COVID-19 High Performance Computing Consortium

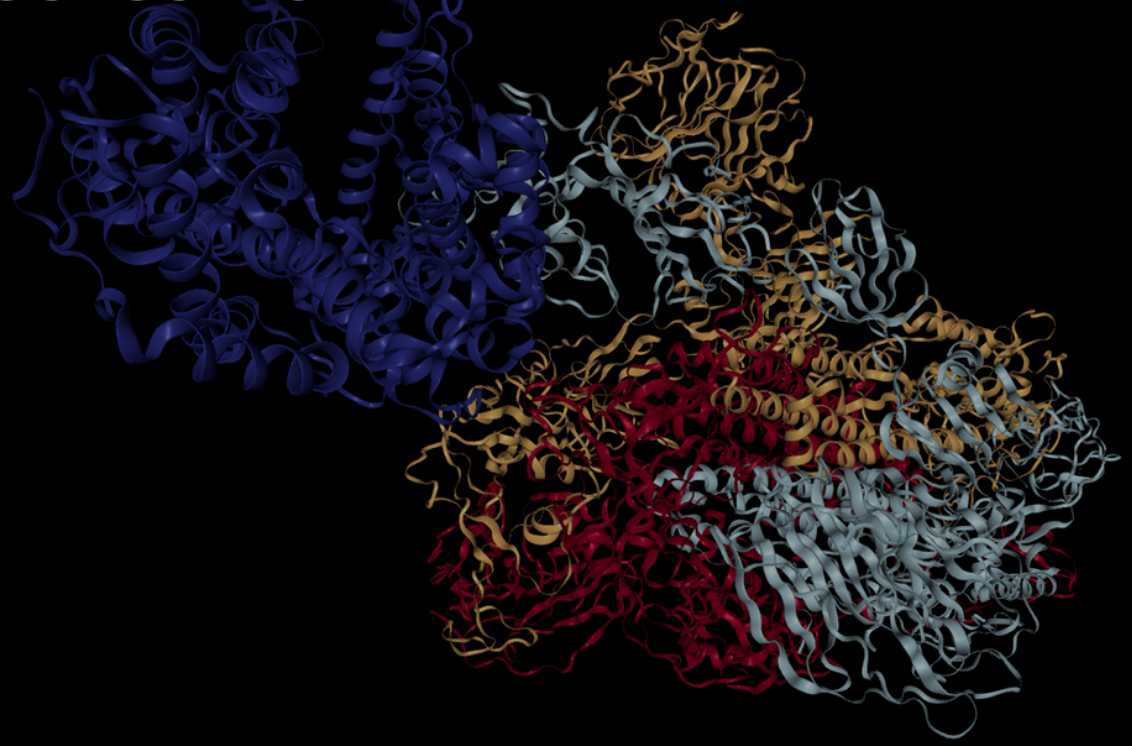

Bringing together the Federal government, industry, and academic leaders to provide access to the world's most powerful high-performance computing resources in support of COVID-19 research.

97

Projects

6.8m

CPU cores

2020 Gordon Bell Special Prize

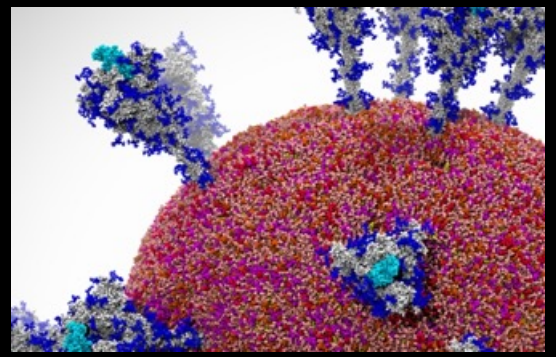

Rommie Amaro,
UC San Diego

"AI-Driven Multiscale Simulations Illuminate Mechanisms of SARS-CoV-2 Spike Dynamics,"

Quantum Computing User Program Resources

IBM

- General-purpose gate system provides
  - 22 systems with up to 65 qubits available

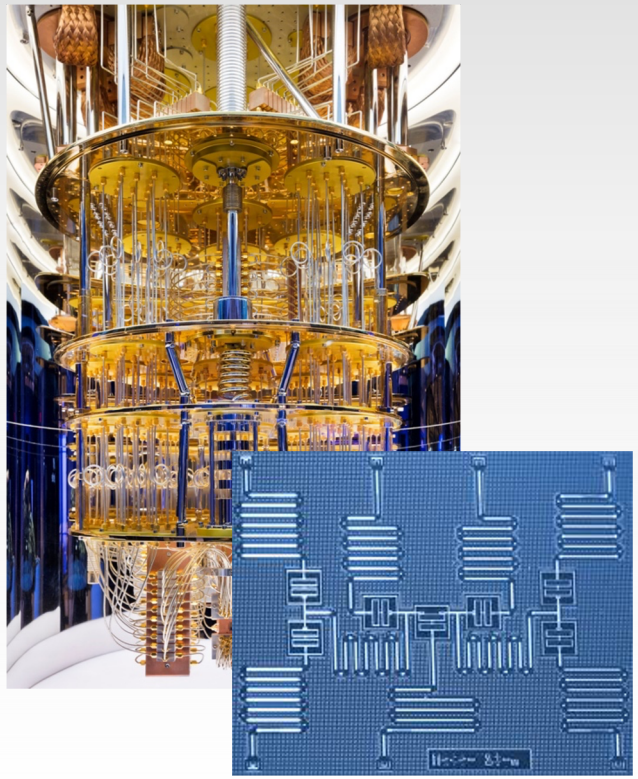

Rigetti

- General-purpose gate system
  - One system available with up to 31 qubits

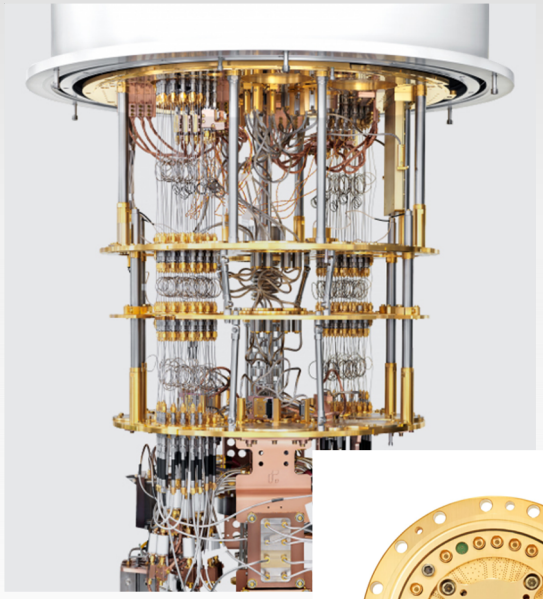

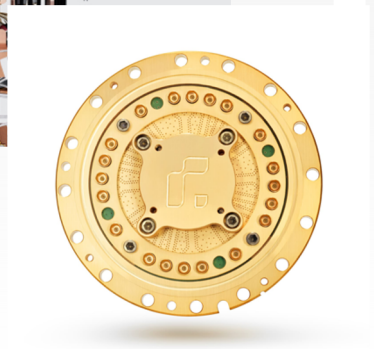

Xanadu

- General-purpose gate system
  - One system with 8 qumodes available

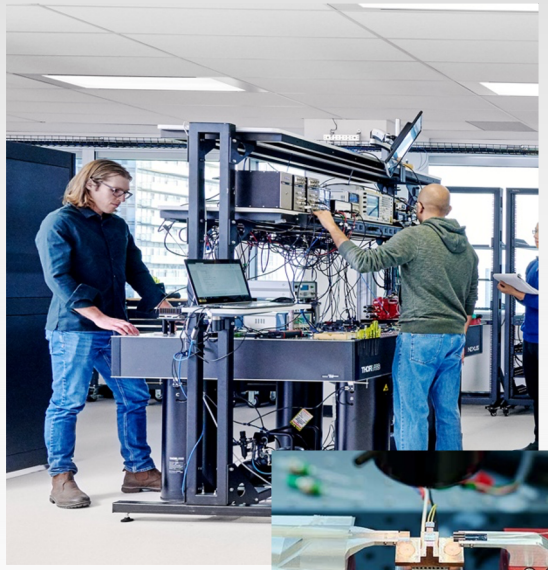

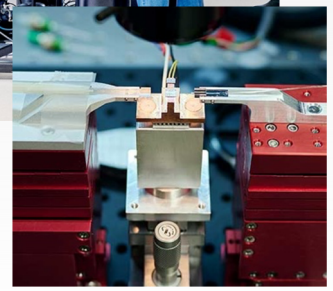

Intersection of High-Performance Computation and Large-Scale Experimental Science and Data Analysis

Large-Scale Computation

Large-Scale Experiment/ Observation

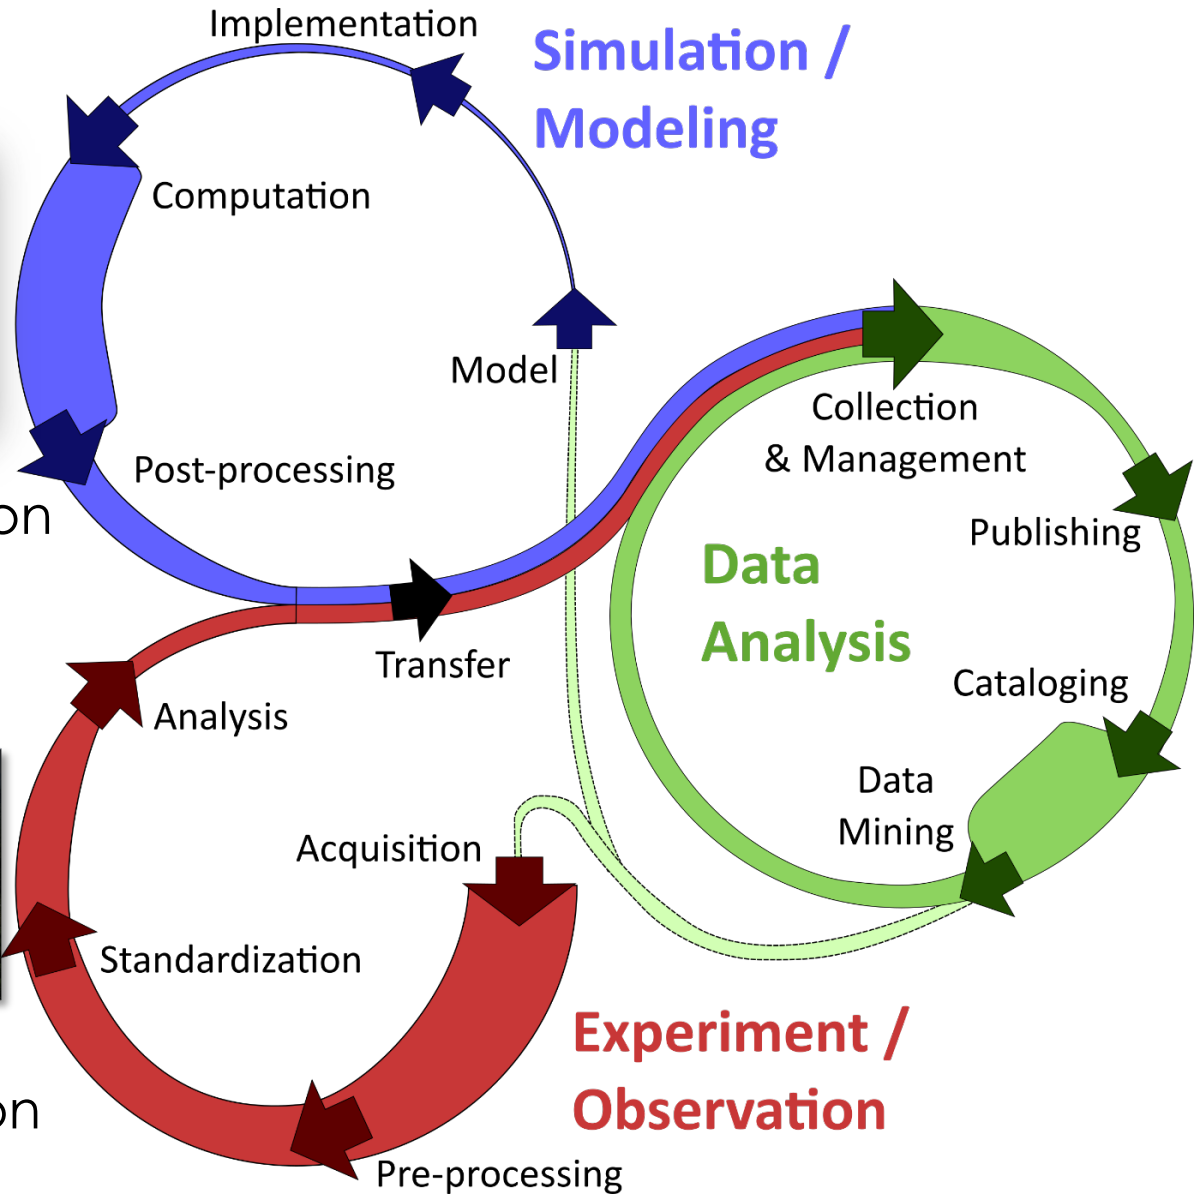

Data Lifecycle Figure Credit: Dr. Suhas Somnath

New operating modalities!

Smart Labs of the Future (2040+)

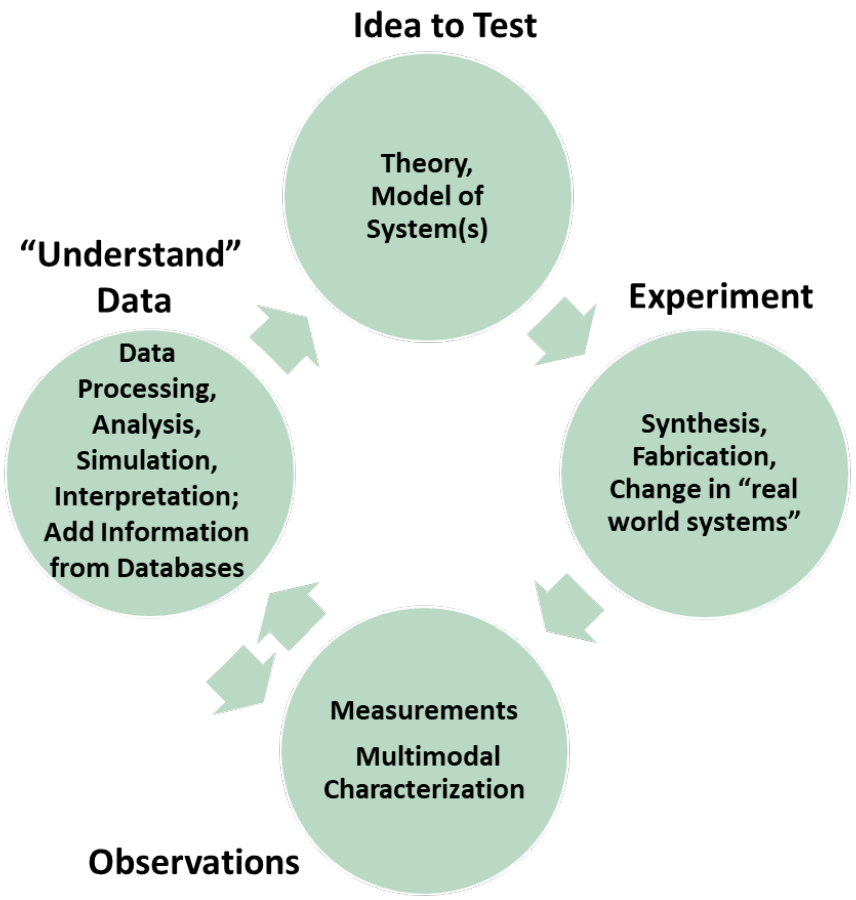

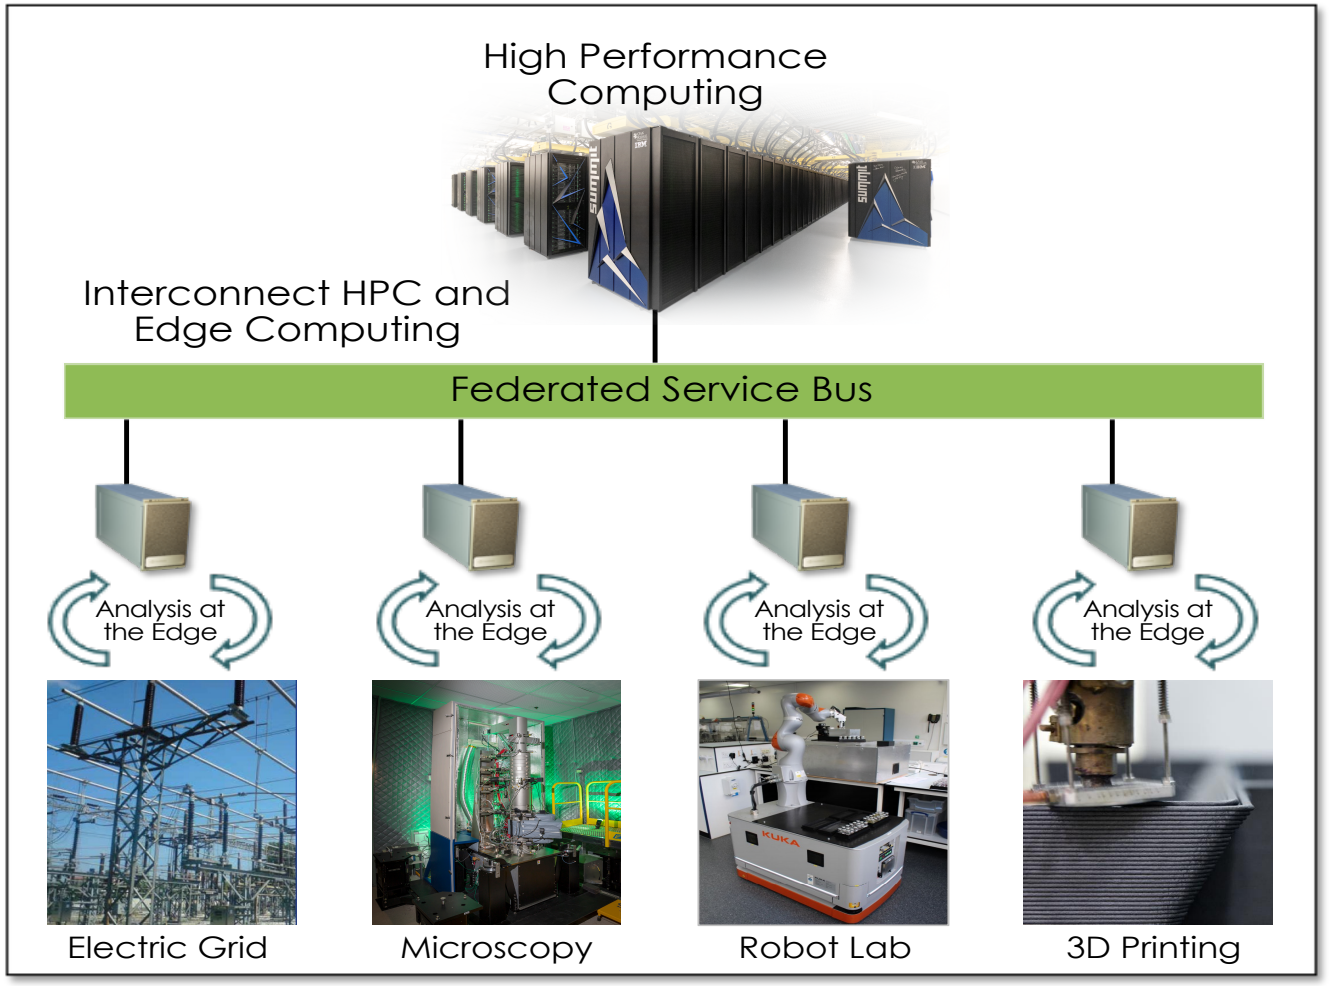

Thank you!

https://www.olcf.ornl.gov/