U.S. DEPARTMENT OF ENERGY | Office of Science

View from Germantown
Advanced Scientific Computing Research

Presented to the
ASCAC

by

Barbara Helland
Associate Director

September 24, 2020

Outline

- FY21 Budget actions
- ASCR Personnel Changes
- ASCR Year in Review
- Recognitions

House Mark

The Advanced Scientific Computing Research program develops and hosts some of the world's fastest computing and network capabilities to enable science and energy modeling, simulation, and research.

- *High Performance Computing and Network Facilities*.—In addition to the long-term exascale initiative, the Committee supports continued upgrade and operation of the Leadership Computing Facilities at Argonne National Laboratory and Oak Ridge National Laboratory and of the High Performance Production Computing capabilities at Lawrence Berkeley National Laboratory. The recommendation includes not less than **$150,000,000 for the Argonne Leadership Computing Facility**, **not less than $225,000,000 for the Oak Ridge Leadership Computing Facility, and not less than $110,000,000 for the National Energy Research Scientific Computing** Center at Lawrence Berkeley National Laboratory. Within available funds, the recommendation includes **$90,000,000 to support necessary infrastructure upgrades and operations for ESnet.**
- *Mathematical, Computational, and Computer Sciences Research*.— The Committee notes the importance of a strong research program in applied and computational mathematics to the Department's mission. Maintaining international leadership in high performance computing requires a long-term and sustained commitment to basic research in computing and computational sciences, including applied math, software development, networking science, and computing competency among scientific fields. Within available funds, the recommendation provides **not less than $250,000,000 for Mathematical, Computational, and Computer Sciences Research, including not less than $10,000,000 for the Computational Science Graduate Fellowship program**. Within available funds, the recommendation includes **up to $40,000,000 for the development of AI-optimized emerging memory technology** for AI-specialized hardware to drive national competitiveness.
- The Committee is aware that the Department of Defense, National Institutes of Health, universities, and private entities have formed partnerships to undertake two of the largest national research investigations of Traumatic Brain Injury (TBI) ever conducted in the Unites States. While leadership and investigators from the TRACK–TBI have developed productive relationships with national laboratory scientists, more collaboration may be needed to develop advanced statistical methods and leverage advances in data science and deep analytics. The Committee encourages the Department to collaborate with the CARE and TRACK–TBI investigators to utilize the national laboratory capabilities to apply the most advanced statistical methods to interrogate the highly complex CARE and TRACK–TBI data sets. The Department is directed to conduct such collaborations on a reimbursable basis through Strategic Partnership Projects.

A closer look at the House Mark

($ in thousands)

	FY2020 Enacted	FY2021 Request	House Mark	Difference FY2021 and House Mark
ASCR	**980,000**	**988,051**	**1,015,000**	**+26,949**
ALCF	150,000	150,000	150,000	--
OLCF	225,000	220,000	225,000	+ 5,000
NERSC	110,000	85,000	110,000	+25,000
Esnet*	90,000	80,000	90,000	+10,000
CSGF	10,000	10,000	10,000	--
Total Research	195,100	264,000	>=250,000	
*AI-optimized Memory**	*10,000-15,000*	*<=40,000*		

* Within available funds

- *Artificial Intelligence:* The recommendation includes up to $125,000,000 for Artificial Intelligence and Machine Learning, and the Committee directs the Department to apply those tools to the Department's mission.
- *Bioscience:* The Department is encouraged to expand its relationships with NIH in order to work together more strategically to leverage the Department's research capabilities, including instrumentation, materials, modeling and simulation, and data science…The recommendation includes not less than $1,000,000 for collaboration with NIH within the Department's data and computational mission space.
- *Exascale Computing Initiative*.—The recommendation includes $474,945,000 for exascale activities.
- *Quantum Information Sciences*.—The recommendation provides $235,000,000 for quantum information science, including not less than $120,000,000 for research and not less than $100,000,000 for up to five National Quantum Information Science Research Centers.

ASCR Staffing

- **Division Directors**
  - Facilities: Candidate's package at Office of Personnel Management for final review
  - Research: External search to begin soon
  - ACT: On hold

- **Backfills**
  - Lucy Nowell filled and candidate starts October 12
  - Betsy Riley – in progress
  - ACT Program manager – waiting on HC
  - Carolyn Lauzon – waiting on HC

- **Two new AAAS fellows on board**

Please Welcome New CS Hire: Hal Finkel

Hal Finkel graduated from Yale University in 2011 with a Ph.D. in theoretical physics focusing on numerical simulation of early-universe cosmology.

- **Hal's the Lead for Compiler Technology and Programming Languages at the ALCF. Hal has contributed to the LLVM compiler infrastructure project for many years and is currently the code owner of the PowerPC backend and the pointer-aliasing-analysis subsystem, among others. As part of DOE's Exascale Computing Project (ECP), Hal is a PathForward technical lead, Co-PI for the PROTEAS-TUNE, Flang, Kokkos, and Proxy Apps projects, and a member of several other ECP-funded projects. Hal represents Argonne on the C++ Standards Committee and serves as vice-chair of the committee. He was the lead developer on the bgclang project, which provided LLVM/Clang on IBM Blue Gene/Q supercomputers. Hal also helps develop the Hardware/Hybrid Accelerated Cosmology Code (HACC), a two-time IEEE/ACM Gordon Bell Prize finalist. He has designed and implemented a tree-based force evaluation scheme and the I/O subsystem and contributed to many other HACC components.**

Hal will start in ASCR on October 12, 2020

Please Welcome our New AAAS Fellows

Saswata (Sash) Hier-Majumder:
A computational scientist with expertise in solid Earth and planetary geophysics. Most recently, Sash was a Reader at the University of London, Royal Holloway.
Sash holds a PhD in Geophysics from the University of Minnesota.

Jordan Thomas:
A data scientist with expertise in physical climate dynamics and modeling. Most recently, Jordan served as a Data Analyst for Innovative Decisions, Inc.
Jordan holds a PhD in Chemical and Physical Oceanography from the Johns Hopkins University.

Saying Good Bye

AAAS Fellow Dr. James Ricci left in May for Schmidt Futures

AAAS Fellow Dr. Michael W. Nestor is leaving in October to take a job with Johnson & Johnson

Carolyn Vea Lauzon left DOE for NSA in July

Laura Biven is leaving DOE for NIH in November

Computer Science / Data Management Analysis and Visualization Portfolio

- In Situ Data Management – Rich Carlson
- Visualization – Rich Carlson
- SSIO – Robinson Pino
- FAIR Data and Models for AI and ML – Bill Spotz
- AI Co-Design – Steven Lee

A Memorial for Lucy Nowell

- On June 14, 2020, the visualization community lost one of its best, Lucy Nowell. Lucy was a computer scientist and program manager at U.S. government agencies including the National Science Foundation and the Department of Energy.
- Lucy's passion and enthusiasm for the field of visualization is an important reason for the success of the field today. She supported scientific visualization, the formation of the field of visual analytics and was instrumental in the creation of the Large Data Analysis and Visualization (LDAV) symposium.
- For many in the HPC and visualization communities, Lucy was a skilled mentor, manager, and friend.
- There is a memorial web page for Lucy.
- Members of the research community are organizing a memorial session for Lucy at the IEEE Visualization 2020 conference, held virtually this year from October 25-30, 2020. Details about this session will be posted on the memorial web page.
  - For questions: Wes Bethel (ewbethel@lbl.gov) or James Ahrens (ahrens@lanl.gov).

Memorial web page: http://sites.google.com/view/a-memorial-for-lucy-nowell/home
IEEE Visualization conference web page: http://ieeevis.org

ASCR Year in Review

2020 Early Career Research Program Awardees





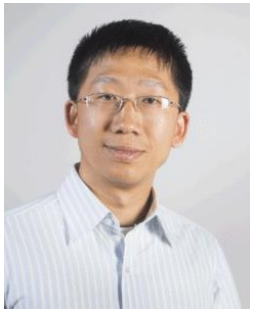

PI Name	Institution	Award Title
Bogdan Nicolae	Argonne National Lab	DataStates: Scalable Versioning for Scientific Data
Bei Phillips	University of Utah	Topology-Preserving Data Sketching for Scientific Visualization
Nathaniel Trask	Sandia National Labs	Physics-informed graph neural networks for data-driven multiscale modeling
Willem Van Rees	Massachusetts Institute of Technology	A multiresolution sharp-interface framework for tightly-coupled Multiphysics simulations
Chen Wang	University of Massachusetts	Enhancing performance of bosonic qubits in circuit QED with reservoir engineering.
Stefan Wild	Argonne National Lab	Structure-Exploiting, Adaptive, Zero-Order Optimization to Improve Efficiency
Kevin Young	Sandia National Lab	Quantum Performance Enhancement

New ASCR Solicitations and awards

	Description	Call Info	Funding
SciDAC Institutes	To provide resources in applied mathematics and computer science, expertise in algorithms and methods, and scientific software tools to advance scientific discovery through modeling and simulation in areas of strategic importance to SC and DOE.	Open: Mar. 17 Closed: May 12	2 teams awarded Total FY2020 Funding $11,500,000
FAIR Data and Models for AI and ML	Proposals that will advance the state of the art in AI and machine learning for science through broad engagement around FAIR data and models.	Open: April 2 Closed: May 15	5 awards Total FY2020 Funding: $8,499,000
Scientific ML for Modeling and Simulations*	Research needed to accelerate the development of artificial intelligence and machine learning for predictive scientific modeling and simulations.	Open: April 8 Closed: May 29	5 awards Total FY2020 Funding: $3,000,000
Artificial Intelligence and Decision Support for Complex Systems	Research in the development and use of artificial intelligence and machine learning in the context of computational decision support for complex systems such as autonomous experiments, resilient cyber-physical systems, and related areas.	Open: April 8 Closed: June 5	6 teams awarded Total FY2020 Funding: $7,015,000
Neuromorphic Computing For Accelerating Scientific Discovery*	Focused on high-risk, high-reward basic research to explore how neuromorphic computing could address emerging scientific computational challenges in energy efficient computing primarily.	Open: April 30 Closed: July 1	5 awards Total FY2020 Funding: $2,000,000

* EXPRESS Calls for high risk, high reward projects

Base Program awards

	Description	Call Info	# Awards
Lab Base Math Program	Numerical methods for solving ordinary and partial differential equations; Computational meshing; Numerical methods for linear and nonlinear equations; Optimization; Multiscale mathematics; Multiphysics computations; and Mathematical software and libraries	Open: Feb. 1 Closed: April 8	9 awards (2 new, 7 renewal) Total FY 2020 funding $6,225,000
Lab Computer Science Machine Learning Research	Enable and identify basic fundamental research challenges to enable extreme scale machine learning and understanding focusing specifically on high performance computing challenges.	Through SC Open Call	4 renewal awards Total FY 2020 funding $3,000,000
Data Management, Analysis and Visualization	Solicited research across data management, analysis, and visualization including *in situ* data management and SSIO with a focus on SC mission needs around complex data.	Through SC Open Call	8 awards (4new, 4 renewal) Total FY 2020 funding $4,482,000

White House Office of Technology Policy, National Science Foundation and Department of Energy Announce Over $1 Billion in Awards for Artificial Intelligence and Quantum Information Science Research Institutes

AUGUST 26, 2020

Home » White House Office of Technology Policy, National Science Foundation and Department of Energy Announce Over $1 Billion in Awards for Artificial Intelligence and Quantum Information Science Research Institutes

WASHINGTON, D.C. – Today, the White House Office of Science and Technology Policy, the National Science Foundation (NSF), and the U.S. Department of Energy (DOE) announced over $1 billion in

See Ceren Susut's presentation at 2:30 PM

Covid-19 rapid response leverages DOE ECRP algorithms

Scientific Achievement

- Quickly created computational methods to screen abnormalities in **lung scans** by repurposing DOE ECRP-funded algorithms developed to tackle structure analysis of samples from energy-related materials.

Significance and impact

- Lung lesions are pathological findings that support characterization of COVID-19. Our algorithms run at NERSC using public datasets to reconstruct lungs (B) from CT scans (A), and estimate regions of lesion (C) for further exploration. Major challenge is to distinguish COVID-19 from other respiratory issues. Preliminary results show ability to transfer learning to chest x-ray images.

Research details

- Predicted lung lesions using deep learning, e.g. Mixed Scale Densenet, with accuracy over 90% when compared with manual ground-truth provided by radiologists;
- Showed that algorithms previously designed to segment multiphase materials (3D synchrotron-based micro-tomography data) can be assembled into complex routines to solve new inspection tasks that might be needed for emergency responses.

BERKELEY LAB · CRD · NERSC · University of California San Francisco UCSF · San Francisco VA Medical Center · Jefferson Thomas Jefferson University

U.S. DEPARTMENT OF ENERGY | Office of Science | ASCR ECRP

[1] Ushizima, ACTS: Accelerating COVID-19 Testing with Screening, *Academic Data Science Alliance*, Oct 2020.
[2] Ushizima, Higa, Saddre, Pelt, Crane, Tosun, Manjumdar, Sundaram, "Computed tomography, chest x-ray and deep learning for lung screening", in preparation.

3D Simulations Slash the Energy Costs of Painting Cars

LBNL mathematicians are partnering with PPG Industries to develop high-performance computational (HPC) models and create energy-efficient coating systems for the automotive industry.

Advanced Computing and Modeling

- ❑ Rotary bell "atomization" is a key mechanism in industrial spray painting operations, a complex process that accounts for **50-70%** of the total energy cost for a factory.
- ❑ A state of the art HPC approach – using discontinuous Galerkin (DG) and level set methods – has been customized & extended to model rotary bell atomization at unprecedented scale & resolution.

Significance and Impact

- ❑ The multi-year collaboration between the LBNL Math Group and PPG Industries is developing ways to increase process efficiency, quality, robustness, and to reduce waste.
- ❑ Insights are expected to yield new ways to reduce the temperature it takes to set the paint, accelerate the onset of multi-layer coating systems, and further reduce energy costs by up to 30%.

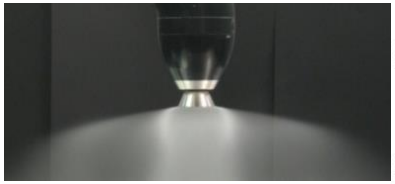

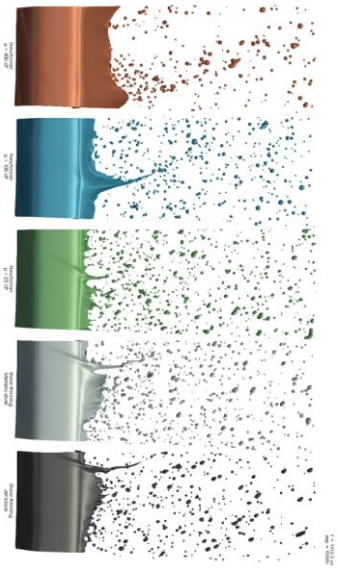

High-resolution peta-scale modeling (right) of liquid atomization (above) comparing different liquid rheologies, computed using HPC implicit mesh DG frameworks.

Smooth, shiny, durable paint films, as thick as a human hair, are design objectives of energy-efficient coating.

Work was performed at LBNL
For more information contact: Robert Saye, rsaye@lbl.gov

Supported by ASCR Applied Math, ALCC, NERSC, & joint PPG-LBNL HPC4Mfg grant

"Physics-Constrained" Neural Network to replace Expensive Fokker-Planck Solver

Collaboration between SciDAC FASTMath and High-fidelity Boundary Plasma Simulation SciDAC Center

Scientific Achievement

Developed a fast approximation to the collision operator for XGC and other fusion gyrokinetic codes, while conserving mass, momentum, and energy, via an augmented Lagrangian method to train the encoder-decoder neural network

Significance and Impact

- XGC is a massively parallel particle-in-cell code studying edge turbulence in magnetic fusion energy devices
- ITER plasma will be affected by **many tungsten species** but the F-P collision operator does not scale well with number of species and **could swamp the ITER simulation even with an exascale computer**
- Our new, **fast** physics-constrained neural network approximates the collision operator while conserving required properties and makes simulation of the ITER with many tungsten species **feasible**

Research Details

The augmented Lagrangian formulation extends conventional neural network training to incorporate nonlinear constraints with minimal intrusion:

- The training time with the augmented Lagrangian method is short, requires only a small number of passes through the data
- Subproblems are solved using stochastic gradient descent method
- Applied a heuristic updates for the multipliers and penalty factor using out-of-sample validation information

arXiv: 2009.06534, 2009.07330

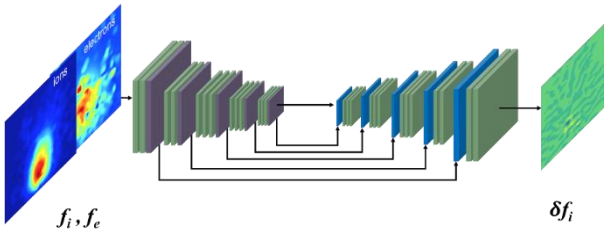

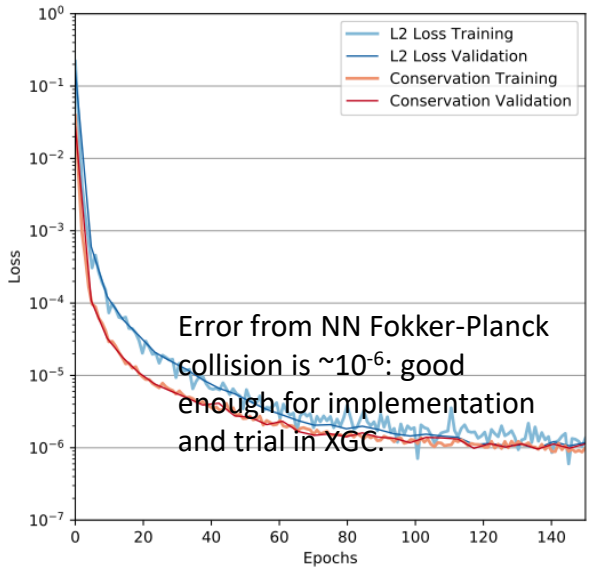

A. Dener (ANL), A. Miller (PPPL), R.M. Churchill (PPPL), **T. Munson (ANL)**, CS Chang (PPPL)

Measuring the Capabilities of Quantum Computers

Scientific Achievement

We introduce the first scalable techniques that can efficiently test the capabilities of any quantum computer, and we apply them to twelve processors.

Significance and Impact

Our methods can be used to learn what programs a particular quantum computer can successfully run. They scale to devices of 1000s of qubits, solving a long-standing problem with going beyond ~50 qubits.

Research Details

We developed a technique that can convert any set of quantum programs into a benchmarking suite of efficiently verifiable programs. We tested them in experiments, revealing that the current hardware suffers complex errors that cause structured programs to fail up to an order of magnitude earlier - as measured by program size - than disordered ones.

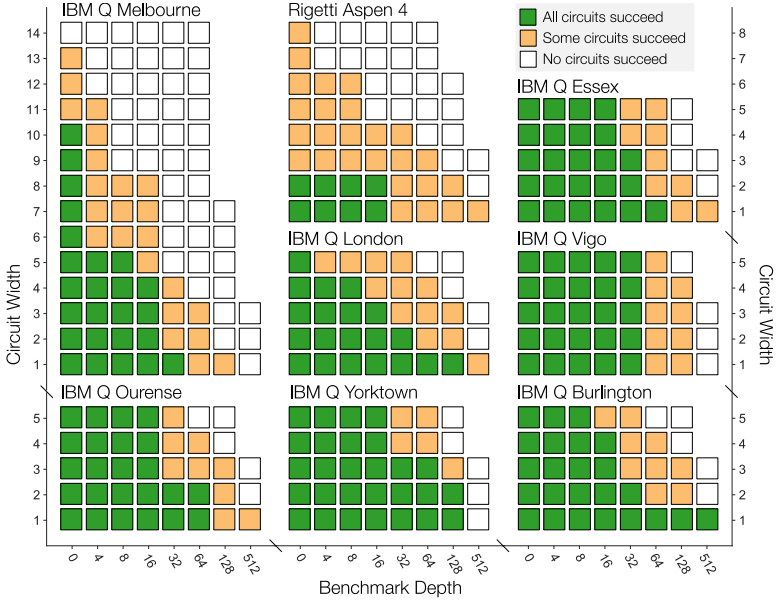

Caption: The measured performance of twelve quantum computers, summarized using the novel concept of "capability regions".

"Measuring the Capabilities of Quantum Computers", Proctor, Rudinger, Young, Nielsen, and Blume-Kohout, arXiv:2008.11294 (2020).

Work was conducted at Sandia National Labs

Sandia National Laboratories

QPL

AIDE|QC

COVID-19 and ASCR

- **CARES Act provided funding**
  - To acquire additional computing resources at ALCF, NERSC, OLCF, LANL and LLNL to augment current capabilities in the fight against COVID-19, in support of NVBL computation efforts. Majority of resources have been delivered, installed, accepted and are in use.
  - To Support Tiger team lead by Kathy Yelick and Fred Streitz to connect DOE's HPC with other agency efforts such as FEMA and HHS
- **COVID-19 HPC Consortium -- Primary beneficiary of new computing resources**
  - As of Wednesday September 23, 2020, the Consortium had received 171 proposals and approved 90, declined 60 and 2 projects have completed. Out of the 90 approved projects, 86 were enabled.
- **ESnet**
  - Traffic related to remote access work resources has tripled since COVID.
  - Science traffic to and from the National Laboratories has decreased ~30% during COVID. Our conjecture is that the decline relates to in-person operations impacts at User facilities and the Laboratories.

For more about ESnet, see Ben Brown's and Inder Monga's talk at 12:30 pm September 25

NERSC Science Highlights

Precision Studies of SARS-COV-2 Binding

With an award from the COVID-19 HPC Consortium, a U. of Missouri, Kansas City-led team is studying how the virus enters and infects human cells at the molecular and atomic level. They published their first results in September as a featured article in the journal PCCP. Using advanced algorithms running on Cori the team determined the precise locations of atoms involved in protein-spike binding to ACE2, improving the precision to less than 0.1 Angstrom.

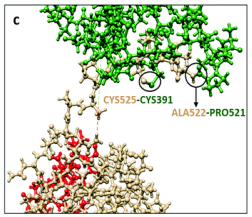

W.Y. Ching, U. Mo-KC

COVID Scholar

Backed by Cori, the center's "Spin" edge-services cloud, and NERSC staff, the COVID Scholar literature search portal was created in record time to help researchers find the latest research results. COVID Scholar, developed by a team of materials scientists, uses natural language processing-techniques to deliver the latest results to 15K users and is incorporated into efforts at MIT, PNNL, and KG-COVID knowledge graph project.

K. Persson, Berkeley Lab

Detailed U.S. Epidemiology Studies

A U.S. Health and Human Services team is performing epidemiology simulations of the entire U.S. using an agent-based epidemiological model to inform policymakers about the number of COVID-19 cases that are projected to occur in various regions, the need for medical resources, and to understand the impact of social distancing and other interventions.

Advances in Carbon Capture Technology

A big advance in carbon capture technology could provide an efficient and inexpensive way for natural gas power plants to remove CO_2 from flue emissions. Researchers from UCB, LBNL and ExxonMobil developed a technique using a metal-organic framework, or MOF, and steam to flush out CO_2 that can be used or stored. The researchers studied and confirmed the structure of the MOF using the Cori supercomputer and the Advanced Light Source at LBNL.

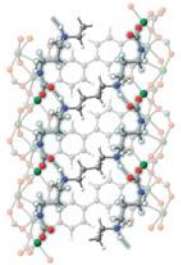

J. Neaton, Berkeley Lab

First 3D Simulations of Superluminous Supernovae

For the first time ever, an international team of astrophysicists simulated the 3D physics of superluminous supernovae—which are about a hundred times more luminous than typical supernovae. They achieved this milestone using the CASTRO code and supercomputers at NERSC.

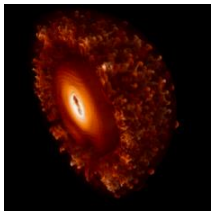

Ken Chen, UCSC

New Catalyst for Chemical Manufacturing

Scientists at Ames Lab and the U. of Oklahoma used Cori to understand the properties of a newly discovered metal-free carbon-based catalyst that has the potential to be much less expensive and more efficient for many industrial concerns, including manufacturing of bio- and fossil fuels, electrocatalysis, and fuel cells.

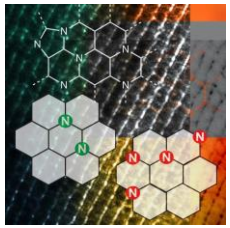

B. Wang, U. Oklahoma

ALCF Science Highlights

COVID-19: Modeling Spread to Inform Policy

Researchers are using the CityCOVID model to test new policies, calibrate unobserved issues, and simulate a variety of interventions and future scenarios. Model results are being used by the City of Chicago and Cook County Public Health Departments, and the Illinois COVID-19 Modeling Task Force.

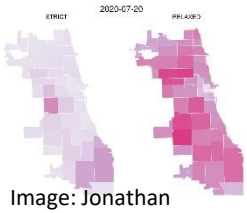

Image: Jonathan Ozik (ANL)

Data-Driven Design of Solar Cells

Light-absorbing dyes are promising, low-cost materials for organic dye-sensitized solar cells that can passively supply energy through tinted windows. Using data mining, machine learning, and computational modeling techniques, researchers identified two high-performing dyes, and then produced a solar cell competitive with common industry materials.

Image: Jacqueline Cole (University of Cambridge)

Improved Radiation Treatment for Cancer

Fast-moving ions have shown promise in clinical radiation oncology, offering more precise targeting of tumor sites and increased preservation of surrounding healthy tissue. The team is improving the understanding of DNA-damaging electronic excitation to help inform improved therapies.

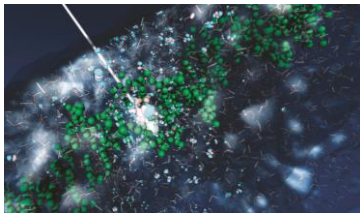

Image: Yosuke Kanai (UNC)

Mira's Last Journey

In the final months of Mira's operation, researchers ran one of the largest cosmological simulations using cutting-edge observational advances from satellites and telescopes. Evolving a massive number of particles, the simulation was designed to help resolve mysteries of dark energy and dark matter. Results will form the basis for sky maps used by numerous surveys.

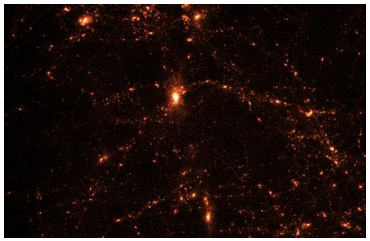

Image: Katrin Heitmann (ANL)

First Tests of DOE Coupled Model E3SM

Researchers carried out the first E3SM simulations in its high-resolution configuration. This higher-resolution version was able to capture the most energetic motions in the ocean as well as the largest of storms in the atmosphere. Such simulations provide more accurate predictions of future changes and better localized information about climate impacts.

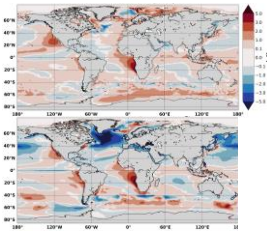

Image: P. M. Caldwell, et al

AI-Driven Drug Discovery for SARS-CoV-2 Proteome

Using AI techniques to screen over 6M small molecules, researchers identified at least 20 partially active molecules that can potentially inhibit viral function in wet lab experiments. The 20 candidates are being validated in labs for activity against the virus.

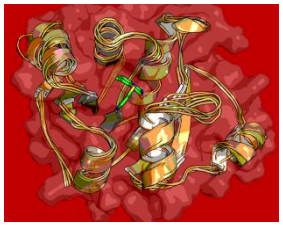

Image: Arvind Ramanathan (ANL)

OLCF 2020 Science Highlights

Identification of Drug Compounds for Treatment of Covid-19

An ORNL and Univ. of Tenn. team used Summit to simulate different compounds docking to the SARS-CoV-2 coronavirus' spike protein to determine if any might prevent it from sticking to human cells. Initially, the team computationally screened more than 8,000 drug compounds—including medications and natural compounds—and identified 77. They were later able to screen 1.5 billion chemical compounds in 24 hours on Summit—the largest drug screening study ever undertaken on a supercomputer.

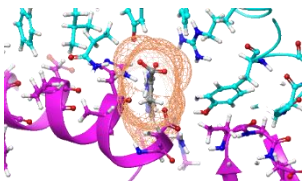

Image above: Compound (gray) calculated to bind to the SARS-CoV-2 Spike Protein (cyan) to prevent it from docking to the Human Angiotensin-Converting Enzyme 2 (ACE2) receptor (purple).

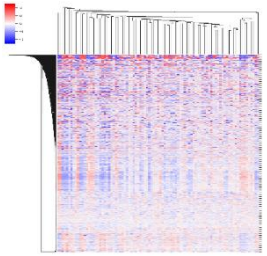

Prediction of Synergistic Drug Combinations for Treatment of Covid-19

A team led by the Icahn School of Medicine at Mount Sinai has developed a machine learning classifier to predict whether certain drug combinations might be effective against the virus. They are using the classifier on Summit to predict gene expression patterns for more than 700 000 combinations of existing drugs.

Image left: A visual representation of the machine learning classifier learning from drug synergy RNAseq dataset and successfully predicting viability of other drug combinations.

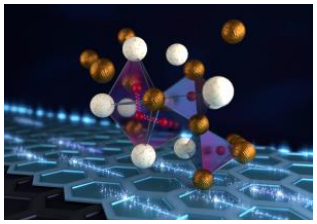

Demonstration of Superconductivity

Researchers used neutrons and computational analysis on Titan to discover the hydrogen atoms in a metal hydride material are much more tightly spaced than had been predicted for decades—a feature that could possibly facilitate superconductivity at or near room temperature and pressure.

Image above: Illustration of a zirconium vanadium hydride atomic structure at near ambient conditions as determined using neutron vibrational spectroscopy and Titan.

Simulation of Quantum Transport

(Gordon Bell Prize winner) An ETH Zurich team studied transistors by simulating quantum transport, or the transport of electric charge carriers through nanoscale materials such as those in transistors. They performed a 10,000-atom simulation of a 2D slice of a transistor on Summit and developed a map of where heat is produced in a single transistor.

Image left: As electrons flow through transistors, such as the ones on circuit boards, they generate heat that dissipates into the environment around them. This project could inform production of new semiconductors with optimal heat-evacuating properties.

Detection of Cancer

An ORNL team updated their MENNDL code and created a multi-objective neural network that can speed up cancer pathology research on a huge scale. Running the updated code on Summit, tumors in biopsy images were identified at a rate of 1.5 minutes per slide—16 times faster—with nearly the same accuracy as previous versions of the code.

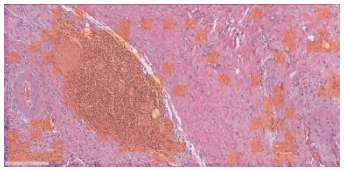

Image above: A portion of a whole-slide cancer biopsy image where tumors appear in the orange overlay. It was labeled using the MENNDL neural network.

Discovery of Quantum Materials

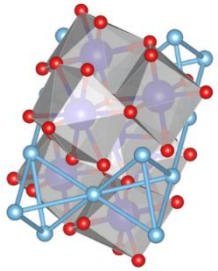

An ORNL-led team used artificial intelligence (AI) to find patterns in neutron scattering data that can lead to an understanding of the physics inside quantum or complex magnetic materials. In more than 50 billion calculations on the OLCF's Titan, Eos, and Rhea systems, a team performed scattering simulations that they then used to train an artificial neural network. They are now training deeper neural networks on Summit to further understand glass-like quantum materials.

Image left: Atomic structure of $Dy_2Ti_2O_7$, a spin ice that is glass-like at low temperatures.

Key ECP Milestones and Schedule

As of March 2020

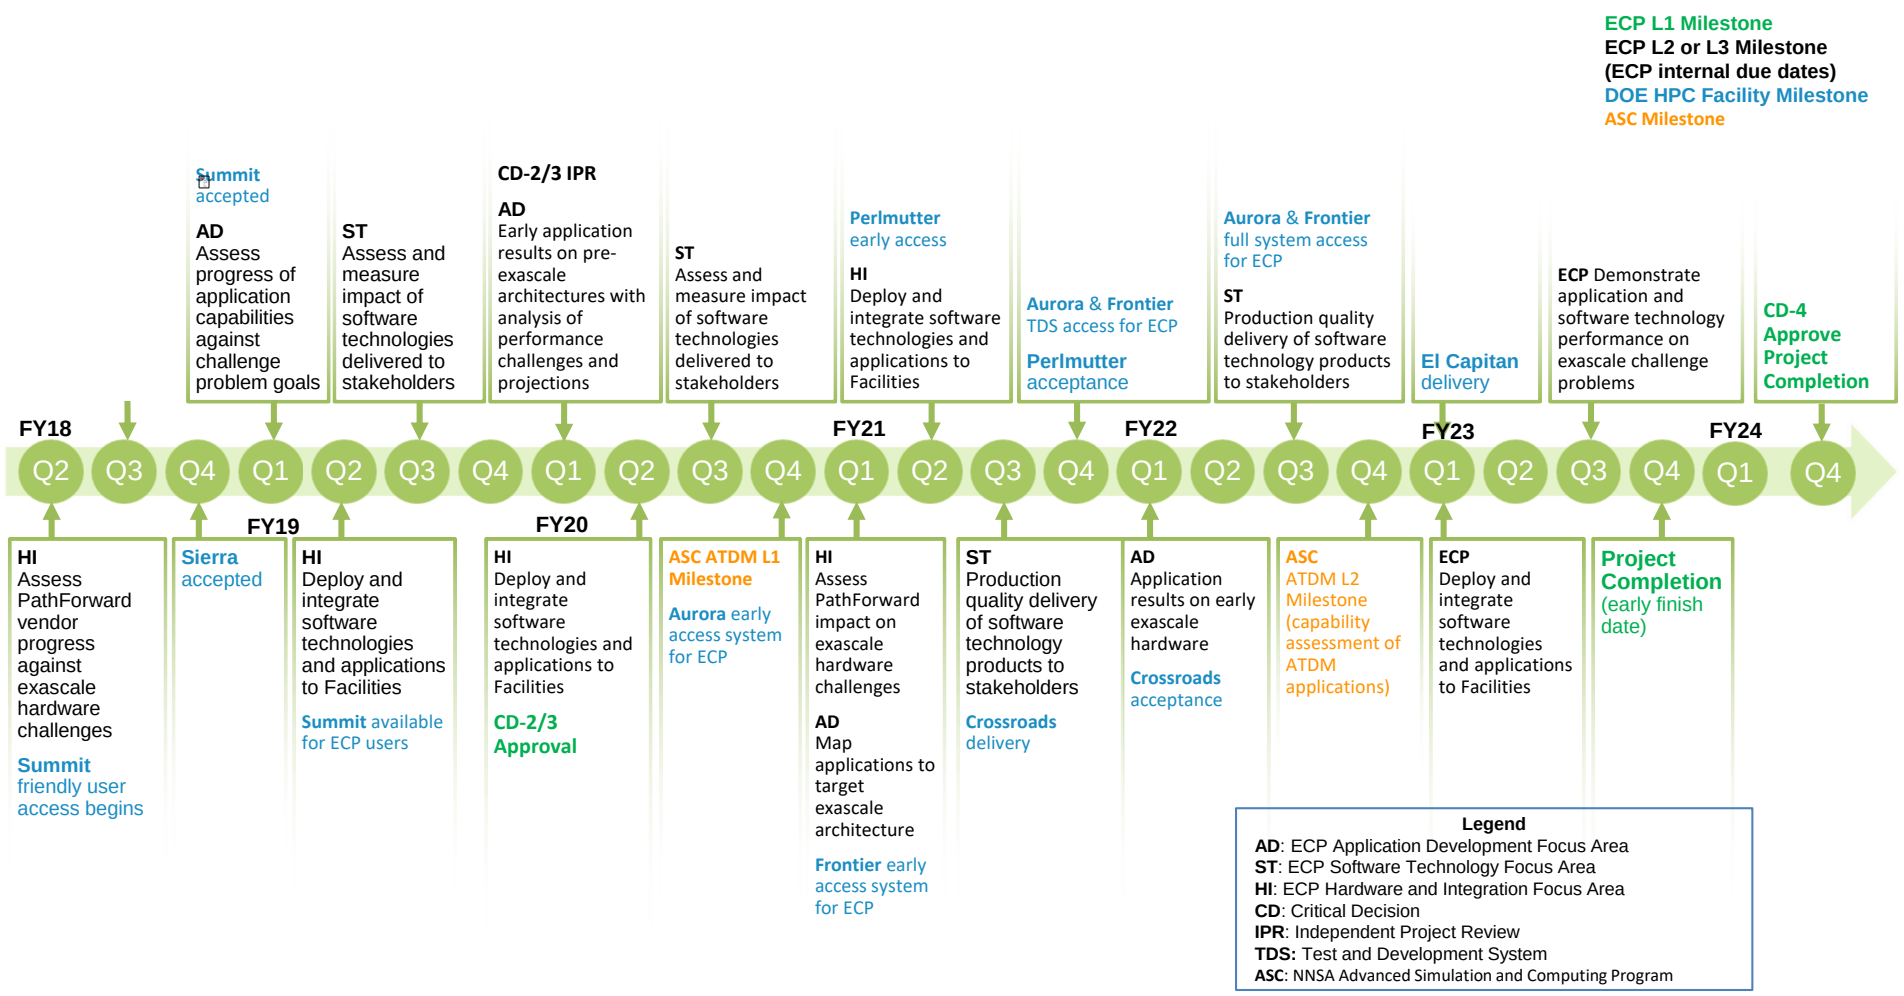

See Mike Heroux's talk at 11:30 on September 25

- **Workshops: Past and Future**

Community of Interest (on Future Scientific Methodologies)

Virtual Workshop [Nov 2, 5, 10, 2020]

- ASCR-funded workshop to explore how the future of computation within the National Lab complex may shape, and be shaped by, advances in scientific methodologies over the next 10-30 years
- Co-Chairs: Ian Foster and Amber Boehnlein
- The curated* unconference** format will be used to foster in-depth discussions on five topics.
  - Tomorrowland Vision – what technological advances may cause a dramatic shift in how science is done (e.g., tabletop accelerators; human-level AI; automated laboratories)
  - Interfaces of the Future – how scientists and agents will interact and communicate
  - Computing Facilities of the Future – how scientists/agents will generate, process, move, find, repurpose, and store data in a global computational fabric
  - Future Missions – what new tasks may be pursued by the National Labs
  - Future Methodologies – the broader implications of these and other developments for how, where, and what science is done at the National Labs

* The workshop organizers will identify the major topics, while attendees determine how to manage the discussions

** An unconference replaces a formal workshop structure with an ad-hoc structure that maximizes attendee interactions

Pioneering the Future of Federally Supported Data Repositories

- Federally supported data repositories play a strategically important role in increasing access to the results of federally funded research and advancing agency missions, with wide-ranging benefits to research communities, industry, and the general public.

- The data science ecosystem is changing rapidly: Large scale open science cloud infrastructure; new expectations from publishers; Machine Learned models as data; AI/ML tools for managing data; convergence of HPC, Big Data, and ML; increasing demand for FAIR data; increasing expectations around privacy, security, and integrity.

- This virtual workshop will bring together representatives from federally supported data repositories and thought-leaders in data science to
  - imagine future opportunities and challenges
  - build and strengthen the community of federally supported data repositories, and
  - Identify areas for cross-agency coordination

A virtual workshop scheduled for January 13 -15 2021

Organized by the NITRD Big Data Interagency Working Group:

- Chaitan Baru (NSF) Co-chair
- Laura Biven (DOE) Co-chair
- Rajeev Agrawal (DOD)
- Frances Carter-Johnson (NSF)
- Ishwar Chandramouli (NIH)
- Wo Chang (NIST)
- Jody Crisp (ORISE)
- Ji Lee (NCO)
- Angela Lester (ORISE)

With invited advisors:

- Deb Agarwal (LBNL)
- Phil Bourne (UVA)
- Julia Lane (NYU)

POC: Ji Lee lee@nitrd.gov

Data Reduction for Science Workshop

Scope

- Data produced by experiments, observations and simulations can overwhelm our capacities to store, preserve and analyze
- Data reduction (DR) includes triggering in an experimental detector, filtering, sampling, feature detection, compression, and reduced order modeling
- Workflows and trust

Purpose

- Priority research directions to enhance DR techniques
- PRDs leading to greater trust in and adoption of DR techniques
- Bring together disparate communities of practice in the DR space

Executive Committee

- Mathematics: Habib Najm, Sandia
- Computer Science: Scott Klasky, Oak Ridge
- Facilities: Jana Thayer, SLAC

Details

- Where: Virtual
- When: Winter, 2020-21

Workshop on Randomized Scientific Computing: Algorithms for AI and Data Science at Scale

Purpose: Explore the **use of randomness** as a foundation & key strategy for high-performance scientific computing

Virtual meetings Jan/Feb 2021

Randomized algorithms are transforming scientific computing in

- **AI & Deep Learning**: Stochastic Gradient Descent
- **Data reduction**: Compressive Sensing, Randomized Projections
- **Massive & streaming data analysis**: Randomized Numerical Linear Algebra

Fundamental properties of randomness can be harnessed for other massive data & post-Moore **computational grand challenges**

- High computational complexity and the development of efficient algorithms
- High data dimensionality and finding sparse representations for **data from user facilities**
- Better algorithm scalability for low-power, high-performance **edge computing**
- Reduced ill-conditioning and sensitivity for inverse problems
- Improved algorithm reliability and robustness to noise

Foundational long-term research & plans are needed for **hybrid algorithms** that anticipate massive data & post-Moore computing challenges over the next decade

Chair: Tammy Kolda (Sandia)
Co-Chairs: Aydin Buluc (LBNL), Stefan Wild (ANL)

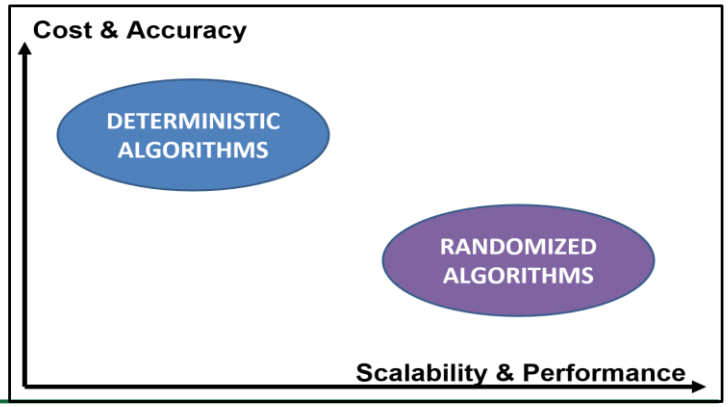

- **Recognitions**

Tammy Kolda (Sandia) elected ACM Fellow and Member of the National Academy of Engineering

2019 Association for Computing Machinery Fellow
Citation: For innovations in algorithms for tensor decompositions, contributions to data science, and community leadership.

2020 Member of the National Academy of Engineering
Citation: For contributions to the design of scientific software, including tensor decompositions and multilinear algebra.

ASCR Applied Math PI

Jacqueline Chen: 2020 DOE/SC Distinguished Scientist Fellow

For advancing frontiers in the fields of combustion and high-performance computing through petascale direct numerical simulations and for mentoring and inspiring generations of researchers.

The Office of Science's recognition honors Jacqueline Chen's career pushing the limits of supercomputers and applied math research to make engines more efficient while minimizing harmful emissions. Chen, a senior scientist in the Chemistry, Combustion and Materials Division of Sandia National Laboratories, is a pioneer in the field of advanced computational methods to understand combustion and chemical reactions relevant to engines.

See Jackie's talk at 1:30 pm on Friday September 25

2020 DOE Sustainability Awards

NERSC Efficiency Optimization Team

Top, left to right: Walker Johnson, Brent Draney, Jeff Broughton, Norm Bourassa, Steve Greenberg, Sadie Joy, Raphael Vitti. Middle including group picture, left to right: Cary Whitney, Deirdre Carter, Norm Bourassa, Jeff Broughton, John Elliott, Deirdre Carter, Walker Johnson, Ernie Jew, Jeff Grounds. Bottom, left to right: Mark Friedrich, John Elliott Tom Davis, Jingjing Liu. Not shown: Mary Gross, Ben Maxwell.

The NERSC / Berkeley Lab Energy Efficiency Team

NERSC's Efficiency Optimization Team was honored with a 2020 DOE Sustainability Award. This honor recognizes work that is "essential in ensuring DOE's continued success as a leader in sustainability."

Savings to date total 560,000 gallons of water per year and 1,800,000 kWh per year, which is equivalent to a reduction of 37% of the non-compute energy of the center.

Some ASCAC Agenda Details

- **REPORT FROM SUBCOMMITTEE ON EXASCALE TRANSITION --** *Roscoe Giles, Boston University*
- **REPORT FROM SUBCOMMITTEE ON 40th ANNIVERSARY ACCOMPLISHMENTS --** *Bruce Hendrickson, Lawrence Livermore National Laboratory*
- **DISTINGUISHED SCIENTIST TALK –** *Jackie Chen Sandia National Laboratories and ASCAC*
- **INCORPORATING GPUS INTO EARTH SYSTEM SCIENCE --** *Mark Taylor, Sandia National Laboratories*
- **APPROACHING THE ESNET-6 ERA –** *Ben Brown, ASCR, Inder Monga, Lawrence Berkeley National Laboratory*
- **A FEW THOUGHTS ON HPC –** *Buddy Bland, Oak Ridge National Laboratory*
- **EXASCALE UPDATE –** *Mike Heroux, Sandia National Laboratories, Rajeev Thakur, Argonne National Laboratory, Jeff Vetter, Oak Ridge National Laboratory*
- **WORKFORCE –** *Valerie Taylor, Argonne National Laboratory*
- **QUANTUM INFORMATION SCIENCE CENTERS –** *Ceren Susut-Bennett, ASCR*

BACKUP

Scientific Discovery through Advanced Computing (SciDAC) Institutes
DE-FOA-0002223 and LAB 20-2223

Solicitation Scope:

- SciDAC Institutes' mission is to provide intellectual resources in applied mathematics and computer science, expertise in algorithms and methods, and scientific software tools to advance scientific discovery through modeling and simulation in areas of strategic importance to SC and DOE.

Context/Coordination:

- SciDAC-5 Institutes will support, complement or develop: 1) Mechanisms for engaging computational grand challenges across the SciDAC Partnership projects that are co-funded by ASCR and its partners. 2) Tools and resources for lowering the barriers to effectively use state-of-the-art computational systems such as those existing and planned for at OLCF, ALCF, NERSC and similar world-class computing facilities over the next 5 years. 3) Mechanisms for incorporating and demonstrating the value of basic research results from ASCR investments. 4) Plans for building up and engaging our nation's computational science research communities.
- SciDAC Institutes are not co-funded by SciDAC partners which are all SC core programs and NE. Since the Institutes are expected to support the SciDAC Partnerships that are co-funded, the solicitation, review and selection of the Institutes are closely coordinated.

Application requirements:

- Eligible Institutions: Universities/colleges, non-profit & for-profit organizations, DOE Labs
- Budgets for 5 year Awards: $3M to $8M per multi-institutional collaboration per year, approximately 2 collaborations
- Cooperative agreements, field work authorizations or interagency agreements
- **Letters of Intent due on April 14th, Proposals due on May 12th**
- An individual may participate in no more than two applications and may be the Institute Director on no more than one application. An institution may be the Lead Institution on no more than one application. There is no limitation on the number of applications in which an institution may participate.
- The following proposals will be deemed unresponsive:
  - Proposals from single institutional collaborations
  - Proposals requesting support of less than $3 million per year for the collaboration or more than $8 million per year for the collaboration
  - Proposals on research within the mission space of other SC and DOE programs
  - Proposals on research or technology funded by other ASCR sub-programs or duplicative of any active SC awards and projects
  - Proposals on research geared towards a specific application
  - Proposals on research and engineering for hardware and architecture development

PM: Ceren Susut

New Awards in Data Management, Analysis and Visualization

PI Name / Co-PIs	Institution	Proposal Title
Bethel, Edward	**Lawrence Berkeley National Laboratory (LBNL),**	**Scalable Analysis Methods and In Situ Infrastructure for Extreme Scale Knowledge Discovery (SENSEI)**
Co-PIs: Duque (Intelligent Light), Ferrier (ANL), O'Leary (Kitware)		
Byna, Surendra	**Lawrence Berkeley National Laboratory (LBNL)**	**Autonomous Object-centric Data Management for HPC Storage**
Co-PI: Snir (UIUC)		
Foster, Ian	**Argonne National Laboratory (ANL)**	**Braid: Data Flow Automation for Scalable and FAIR Science**
Klasky, Scott	**Oak Ridge National Laboratory (ORNL)**	**SIRIUS 2: Science-driven Data Management for Multitier Storage 2.0**
Co PIs: Ainsworth (Brown University), Liu (NJIT)		
Peterka, Tom	**Argonne National Laboratory (ANL)**	**Triple Convergence of HPC, BD, and AI through ASCR In Situ Workflow Tools**
Co-PI: Morozov (LBNL)		
Ramakrishnan, Lavanya	**Lawrence Berkeley National Laboratory (LBNL)**	**Supporting Collaborative Interactive Reproducible Analytics on HPC Systems.**
Co PI: Perez (UC Berkeley)		
Ulmer, Craig	**Sandia National Laboratories, California (SNL-CA)**	**Offloading Data Management Services to Smart NICs**
Co-PI: Maltzahn (UCSC)		
Ushizima, Daniela	**Lawrence Berkeley National Laboratory (LBNL)**	**AMLXD: Analysis and Machine Learning Across Domains**

FAIR Data and Models for AI and ML (DE-FOA-0002306, LAB 20-2306)

PI Name	Institution	Proposal Title
Bouchard, Kristofer	**LBNL**	**ENDURABLE: Benchmark datasets for AI with queryable metadata**
Brinson, Lynda Catherine	**Duke University, Durham, NC**	**FAIR Data and Interpretable AI Framework for Architectured Metamaterials**
Co PI: Daraio (California Institute of Technology)		
Fox, Geoffrey	**Indiana University, Bloomington, IN**	**FAIR Surrogate Benchmarks Supporting AI and Simulation Research**
CoPIs: Beckman (ANL), Dongarra (UTK), Jha (Rutgers)		
Huerta, Eliu	**University of Illinois, Champaign, IL**	**FAIR Framework for Physics-Inspired Artificial Intelligence in High Energy Physics**
Co Pis: Duarte(UCSD), Harris (MIT), Rusack (U Minnesota)		
Liao, Chunhua	**LLNL**	**HPC-FAIR: A Framework Managing Data and AI Models for Analyzing and Optimizing Scientific Applications**
Co PIs: Emani (ANL), Shen (NCSU)		

ENDURABLE: Benchmark datasets for AI with queryable metadata.

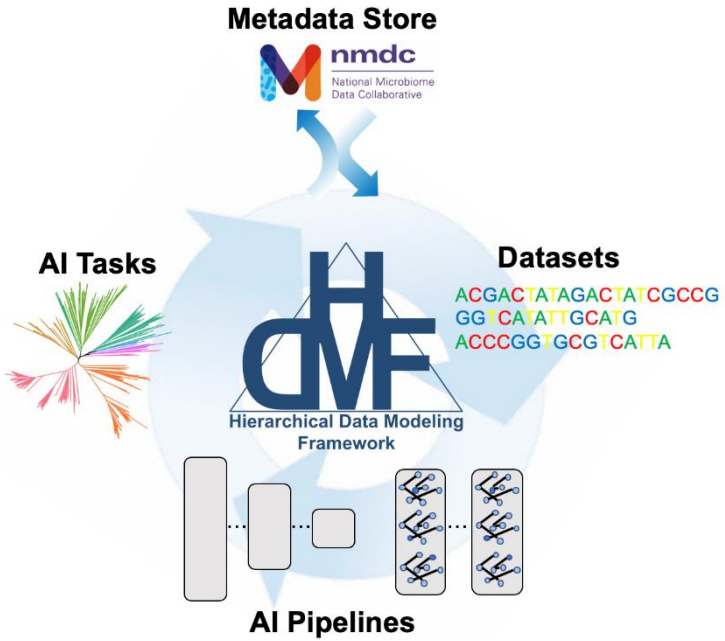

Goal: Provide the scientific and DL communities with robust, scalable, and extensible tools to share and rigorously aggregate diverse scientific data sets for training state-of-the-art DL models.

Thrust 1: Develop standard for creating benchmark datasets and relating deep learning models.

Thrust 2: Develop advanced API for reading benchmark data and saving deep learning models.

Thrust 3: Aggregate and curate data to create accessible benchmark datasets and associated deep learning models.

Science Application Focus: Microbiome data

Impact: The creation of standard description languages and scalable interfaces enables researchers to query the data and models, both individually and with respect to one another, greatly enhances the FAIRness of the data and models.

PI: K.E. Bouchard, LBNL

BERKELEY LAB

nmdc National Microbiome Data Collaborative

Motivation: The integration of AI into materials science remains limited due to lack of sufficient consistent datasets and frameworks for AI models that can readily advance significant fundamental understanding.

Goal 1: FAIR data for material challenges

- ➢ Datasets on metamaterials used for AI frameworks and benchmarking
- ➢ A simulator to verify predictions, create new hypotheses, and create new data
- ➢ A repository of data that is easy for users, including those that are not material scientists, to access and use.
- ➢ Benchmark challenges based on the data to be contributed to UCI Machine Learning Repository

Goal 2: A coarse-to-fine interpretable AI framework for material science data

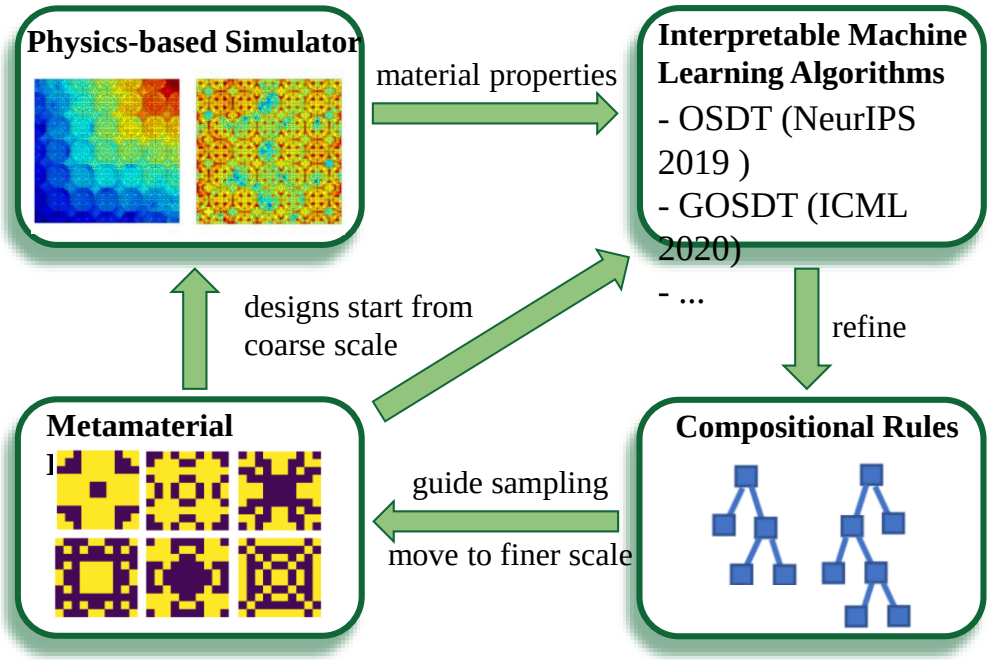

PIs: Cate Brinson & Cynthia Rudin @ Duke
Chiara Daraio @ Caltech

Duke UNIVERSITY Caltech

40

FAIR Framework for Physics-Inspired Artificial Intelligence in High Energy Physics

Interdisciplinary and multi-institutional project Led by the Center for Artificial Intelligence Innovation at the University of Illinois at Urbana-Champaign, this project brings together AI, high energy physics (HEP), software and data science experts from Illinois, MIT, UC San Diego and the University of Minnesota to advance DOE's objectives in data-driven discovery

Motivation Lead the definition and adoption of FAIR principles for AI models and data using HEP as the science driver, spearheading the construction of a rigorous theoretical framework that makes the best use of AI for DOE's strategic projects in science and engineering

Data
- ✓ Procure, develop and share benchmark datasets that adhere to FAIR principles
- ✓ Curate datasets guided by HEP, AI and information science experts

Rigorous AI
- ✓ Use FAIR benchmark datasets to produce physics-inspired AI models
- ✓ Released AI models through the Data and Learning Hub for Science

AI and HEP Discovery
Combine scientific visualizations, accelerated computing and domain-inspired methodologies to gain new insights between the interplay of AI data and models

Community Building
Define and implement FAIR principles for AI models and best practices for sharing data and AI models

PI: Eliu Huerta, NCSA/University of Illinois at Urbana-Champaign

ILLINOIS
NCSA | National Center for Supercomputing Applications

MIT

UC San Diego

HPC-FAIR: A Framework Managing Data and AI Models for Analyzing and Optimizing Scientific Applications

Motivation:

Artificial intelligence (AI) and machine learning (ML) techniques:

- Widely studied to address various challenges of analyzing and optimizing large-scale scientific applications
- Limited, however, because it is extremely difficult to generate, access, and maintain high-quality training data needed to drive ML-based research.

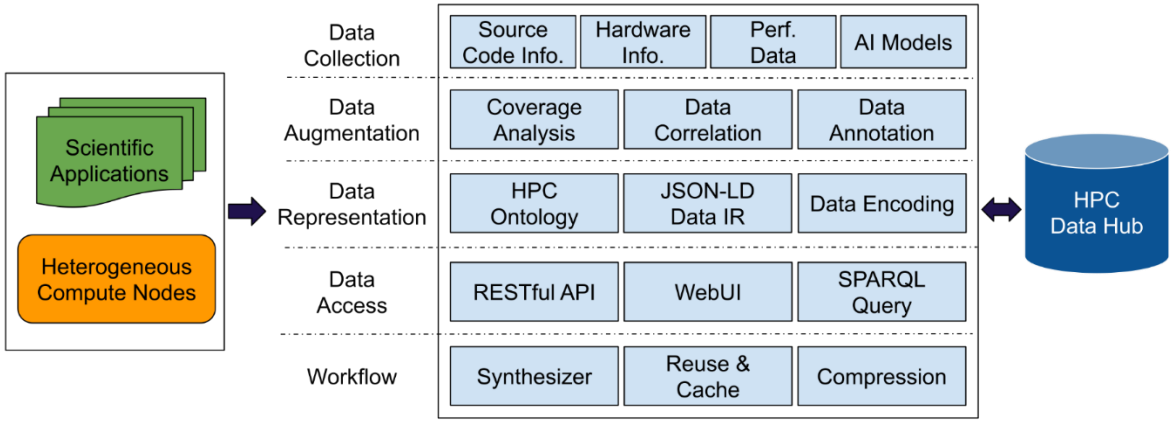

Objectives:

- Developing a generic high-performance computing (HPC) data management framework
- Making both training data and AI models of scientific applications Findable, Accessible, Interoperable, and Reusable (FAIR)

Research Thrusts:

- Collecting and generating representative training data and AI models
- Developing program analysis and translation tools to augment, correlate and annotate data from multiple layers of software/hardware stack
- Designing an internal representation (called Data IR) for representing both training data and AI models
- Providing easy access and user interfaces
- Synthesizing workflows to automatically convert and optimize user queries

ASCAC 09/23/2020

Lawrence Livermore National Laboratory | NC STATE UNIVERSITY | Argonne NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY | Office of Science

SBI: Surrogate Benchmark Initiative
FAIR Surrogate Benchmarks Supporting AI and Simulation Research

Replacing traditional HPC computations with Deep Learning surrogates can improve the performance of simulations and make optimal use of diverse architectures

- Fitting of **hardware** to surrogates
- **Uncertainty Quantification** of the surrogate estimates
- Minimize **Training Data Size** needed to get reliable surrogates for a given accuracy choice.
- Develop and test surrogate **Performance Models**

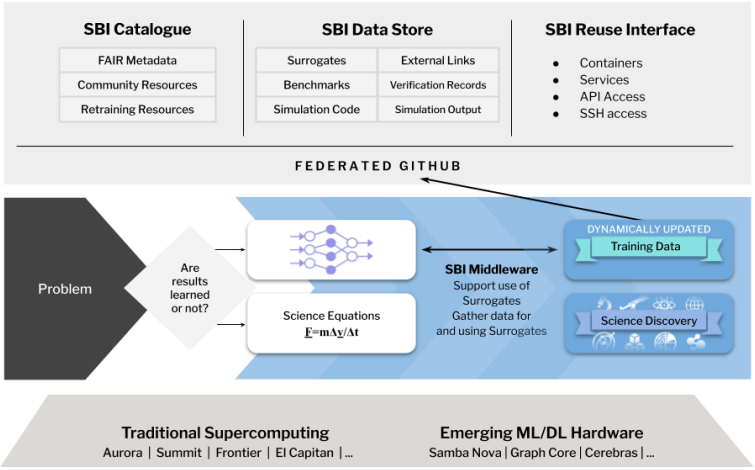

GOAL: Accelerate and better understand **Deep Learning Surrogate** models that can replace all or part or of traditional large-scale HPC computations with major performance increases.

Findable, Accessible, Interoperable, and Reusable **FAIR data** ecosystem for HPC surrogates

SBI collaborates with Industry and a leading machine learning benchmarking activity -- **MLPerf**

Software Research: SBI will design and build general middleware to support the generation and the use of surrogates.

Application Benefits: SBI will also make it easier for general users to develop new surrogates and help make their major performance increases pervasive across DoE computational science.

PI: Geoffrey Fox, IU

Argonne NATIONAL LABORATORY · RUTGERS · THE UNIVERSITY OF TENNESSEE KNOXVILLE

AI & Decision Support for Complex Systems Awardees

Lead PI	Institutions	Proposal Title
Thomas Catanach	**SNL**, Michigan	Goal-oriented Bayesian Machine Learning for Closed-Loop Systems
Xun (Ryan) Huan	**Michigan**	Practical and Optimal Sequential Bayesian Experimental Design for Complex Systems Incorporating Human Experimenter Preferences
Frank Liu	**ORNL**, PNNL, Arizona, UC Santa Barbara	Decision and Control of Complex Systems: a Data-Driven Framework
Adrian Sandu	**Virginia Tech**	Fast outer loop solution algorithms based on a coherent coupling of physical and machine learning models
Karen Willcox	**Texas**, SNL	RISE of the Machines: Robust, Interpretable, Scalable, Efficient Decision Support
Zheng Zhang	**UC Santa Barbara**	Quantum-Inspired Bayesian Sampling for Uncertainty Quantification and Machine Learning

FY2020 Funding Opportunity Announcement: Scientific Machine Learning for Modeling and Simulations

Description:

- Foundational research to accelerate & improve the development of AI and Machine Learning for predictive scientific modeling and simulations
- Received 348 preproposals, Encouraged 32 proposals, Reviewed in Virtual Panel
- 5 Project Awards fully-funded for 2 years
- Total amount of awards: $3,000,000

PI Name	Institution	Proposal Title
Guang Lin	Purdue University	Trustworthy Physics-informed Deep Learning for Predictive Scientific Computing
Mauro Maggioni	Johns Hopkins University	Statistical Learning for Nonlinear Model Reduction from Local Simulations of Stochastic and Particle- and Agent-Based Systems
Reese Jones	Sandia National Labs	Graph Neural Network Models of Complex Initial Boundary Value Problems that embed Physical Invariances
Ryan King	National Renewable Energy Lab	Scientific Machine Learning Frameworks for Multifidelity Computational Stacks
Guannan Zhang	Oak Ridge National Lab	Black-box Training for Scientific Machine Learning Models

The RAPIDS(2) Institute

Enabling scientific breakthroughs using DOE supercomputers by assisting SciDAC and DOE scientists and engineers to solve computer science, data, and artificial intelligence (AI) challenges.

- **Five primary thrust areas, including engagement (top), covering key needs**
- **New emphasis on AI: enabling use by science teams and leveraging within RAPIDS tools**
- **Coordination with facilities and ECP**

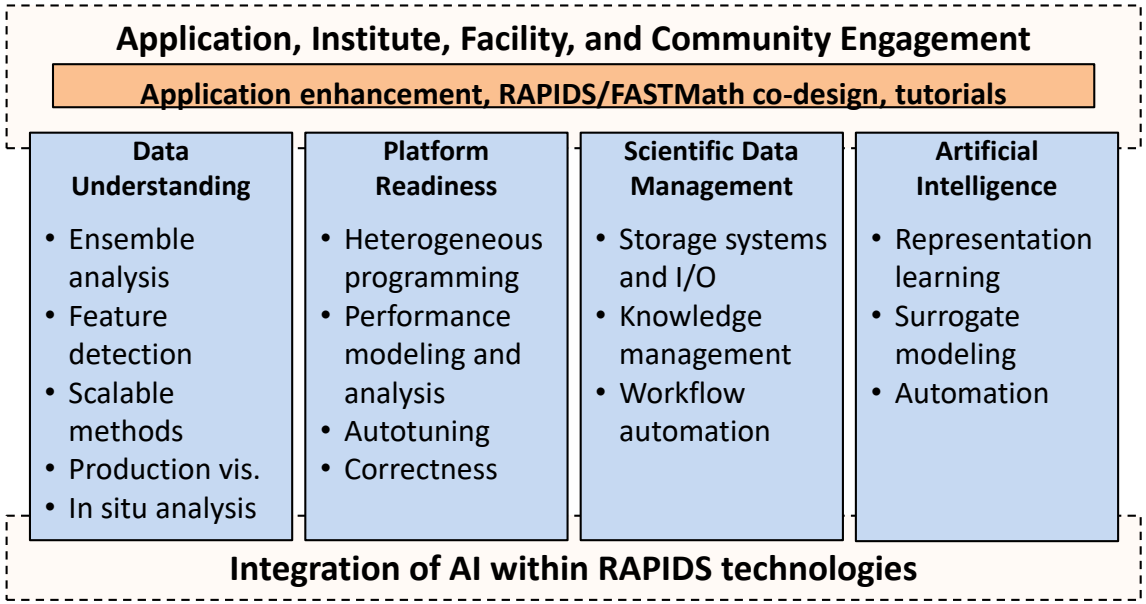

https://rapids.lbl.gov/

Argonne NATIONAL LABORATORY · BERKELEY LAB · University of Delaware · Kitware · Lawrence Livermore National Laboratory · Los Alamos National Laboratory EST. 1943 · Oak Ridge National Laboratory · Northwestern University · The Ohio State University · Rutgers · UF · University of Oregon · Brookhaven National Laboratory

SciDAC-5 FASTMath Institute: Frameworks, Algorithms and Scalable Technologies for Mathematics (http://scidac5-fastmath.lbl.gov)

Eight technical areas

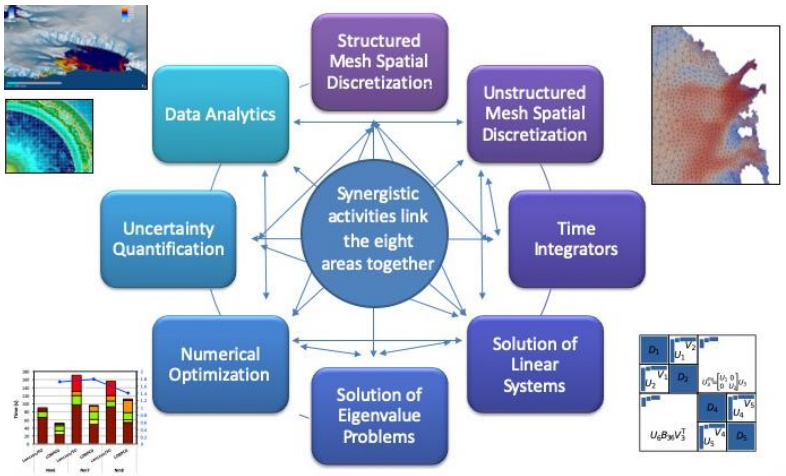

FASTMath Goals:

- Develop robust mathematical techniques and numerical algorithms for DOE applications
- Deliver high-performance software on current and next-generation DOE Office of Science supercomputers
- Work with domain scientists to leverage our math and ML expertise and deploy our software in large-scale modeling
- Engage and support the computational science community across the DOE complex

100s of person years of experience building robust & efficient software

SuperLU AMReX sundials ZOLTAN mfem PETSc UQTk symPACK hypre PUMI Trilinos Albany DAKOTA

60 researchers from 5 DOE labs & 5 universities

Argonne NATIONAL LABORATORY, BERKELEY LAB, OAK RIDGE National Laboratory, Sandia National Laboratories, Boulder, USC, MIT, Rensselaer, SMU

Director: Esmond G. Ng, LBNL (EGNg@lbl.gov)
Deputy: Karen Devine, SNL (kddevin@sandia.gov)

Quantum Testbeds for Science

ASCR's quantum testbeds provide the research community with fully transparent access to novel quantum computing hardware, enabling foundational research to explore high-risk, high-reward approaches.

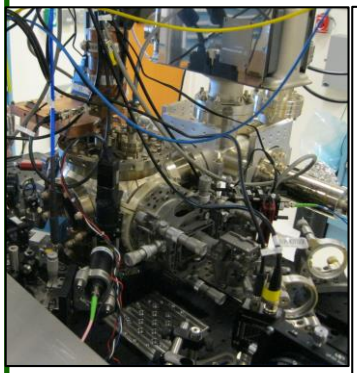

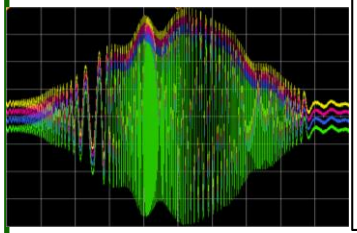

The **Quantum Science Open User Testbed (QSCOUT)** at SNL is the world's first publicly-available trapped ion quantum computer. Its custom assembly language gives users access to low level control and is compatible with most higher-level languages. State-of-the-art coherent pulse generation and a new ion trap design push the envelope for ion trap quantum computing.

The **Advanced Quantum Testbed** at LBNL (**AQT@LBNL**) offers access to ternary logic (qutrits) as well as binary logic (qubits) in a unique superconducting platform. Second-generation hardware will include increased connectivity on a 3D chip.

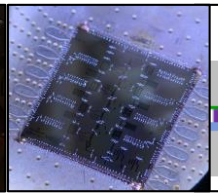

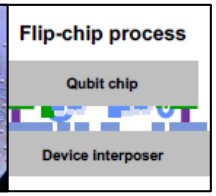

Despite COVID-related delays, both testbeds will be open to external collaborators in CY 2020. For details and to sign up for email updates, go to: https://qscout.sandia.gov and https://aqt.lbl.gov

Scientific Machine Learning for Modeling and Simulations
DE-FOA-0002319 and LAB 20-2319

Solicitation Scope:

- FOA and companion Lab Announcement will solicit research needed to accelerate and greatly improve the development of artificial intelligence and machine learning for predictive scientific modeling and simulations.

Context/Planning:

- The purpose of this solicitation is to develop the foundations of artificial intelligence and machine learning for predictive modeling and simulations in such DOE-relevant areas as physics, chemistry, biology, cosmology, and basic energy research.
- Solicitation stems for Priority Research Directions identified in the SC/ASCR workshop report on Basic Research Needs for Scientific Machine Learning: Core Technologies for Artificial Intelligence and AI for Science Town Hall meetings.
- Solicitation is aligned with key strategies identified in the NITRD National Artificial Intelligence Research and Development Strategic Plan.
- Scientific computing within DOE traditionally has been dominated by complex resource-intensive numerical simulations. The combination of traditional scientific computing knowledge, coupled with massive data and machine-learning based adaptivity, has the potential to greatly improve predictive scientific modeling and simulations.

Application requirements:

- Eligible Institutions: Universities/colleges, non-profit & for-profit organizations, DOE Labs
- Budgets for 2 year Awards: Labs at $400K/year; Universities at $150K/year
- Grants or field work authorizations
- **Pre-proposals due on May 1, with encourage/discourage decisions by May 12**
- Encouraged proposals due by May 29
- An individual may participate in no more than two applications
- Preproposals will be evaluated to determine their competitiveness (Responsiveness to the objectives of the solicitation, Scientific/technical merit, Appropriateness of research approaches, Likelihood of scientific impact)

PM: Steve Lee

Artificial Intelligence and Decision Support for Complex Systems
DE-FOA-0002321 and LAB 20-2321

Solicitation Scope:

- FOA and companion Lab Announcement solicits research in the development and use of artificial intelligence and machine learning in the context of computational decision support for complex systems such as autonomous experiments, resilient cyber-physical systems, and related areas.

Context/Planning:

- Solicitation stems for Priority Research Directions identified in the SC/ASCR workshop report on Basic Research Needs for Scientific Machine Learning: Core Technologies for Artificial Intelligence and AI for Science Town Hall meetings.
- Solicitation is aligned with key strategies identified in the NITRD National Artificial Intelligence Research and Development Strategic Plan.
- Advanced modeling applications implement an expensive forward model of a complex system at the center of a larger algorithm: these include optimization, inverse problems, uncertainty quantification, data assimilation, and control. These advanced applications are the type that are typically most useful to decision makers, but are often too expensive to be practical, even on projected exascale computers. This solicitation seeks AI/ML solutions that address the drawbacks of these "outer loop" applications.

Application requirements:

- Eligible Institutions: Universities/colleges, non-profit & for-profit organizations, DOE Labs
- Budgets for 3 year Awards: Labs at $400K - $1,200K per year; Universities at $150K - $400K per year
- Grants or field work authorizations
- **Pre-proposals due on May 6th, Proposals due on June 5th**
- An individual may participate in no more than two applications
- Preproposals will be evaluated to determine their competitiveness (Responsiveness to the objectives of the solicitation, Scientific/technical merit, Appropriateness of research approaches, Likelihood of scientific impact)

PM: Bill Spotz

NERSC-9

NERSC 9 Project Continues on Pace

Site preparation activities are still on pace for Phase I of the Perlmutter supercomputer, which is arriving later this year.

The NERSC application readiness team has made outstanding progress preparing key Office of Science applications for Perlmutter's GPUs.

A set of power distribution units for Perlmutter are in place (photo from the week of 9/7/2020).

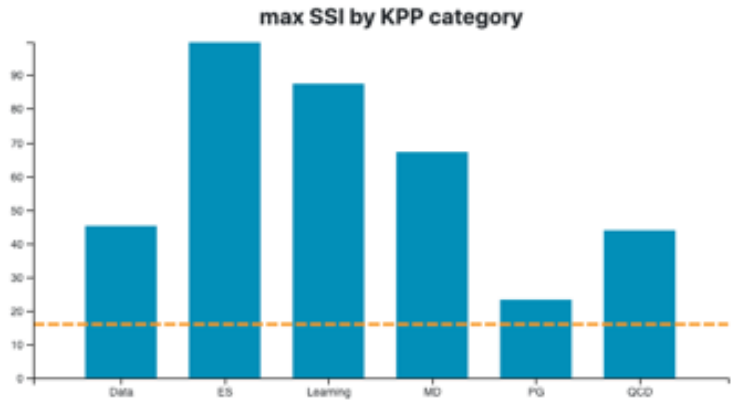

The projected GPU performance in each of the NERSC 9 project's key application categories* already exceeds targets (dashed orange line).

* (L to R): Data, Electronic Structure, Learning, Molecular Dynamics, Particles & Grids, QCD

ALCF Outreach Summary – FY20

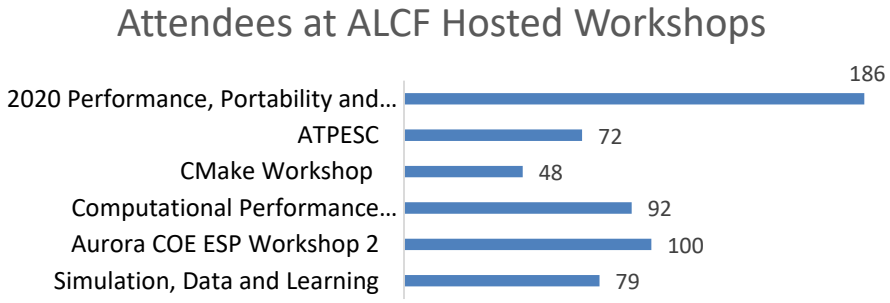

- Ten topics covered at the ALCF Monthly Webinar Series with over 900 registered total
- Aurora Early Adopter series topics included in the monthly webinars:
  - An Overview of Aurora, Argonne's Upcoming Exascale System (Dec-19)
  - DAOS: Next-Generation Data Management for Exascale (Mar-20)
  - OpenMP Offload Capabilities in the oneAPI HPC Toolkit (Jun-20)
- Two Aurora ESP training webinars – Dec-19 and Jun-20 – 133 registered
- Second Aurora COE ESP Workshop held in-person Feb 25-27 – 115 registered

OLCF Strategic Highlights

- **Delivered Summit upgrade project ahead of schedule and under budget - retaking US leadership in HPC and Computational science.**
  - #1 on Top 500 list, June 2018-June 2020
  - Supported Gordon Bell nominated and winning teams, 2018, 2019 and 2020
- **Successfully transitioned facility and project leadership to junior staff, overlapping senior staff to ensure transfer of expertise.**
- **A focus on reducing technical risks has put the Frontier project on path to deliver exascale capabilities in calendar year 2021.**
  - Negotiated to provide a custom CPU to improve the schedule and lower technical risk
  - Synergistic SPP project enabled early access and debugging of Frontier Software
  - NRE contract modification to enable "plan B" software option for DOE complex
  - Vendor and ORNL analysis of alternative HPE software solution for Frontier
- **Quickly delivered CAREs funded hardware customized for COVID related research and supported more than a dozen teams awarded nearly two million hours of computing time.**
- **Added electrical safety officer to improve project and facility focus on safety.**
- **Maintained uninterrupted Summit operations and Frontier site preparation during the pandemic.**