U.S. DEPARTMENT OF ENERGY | Office of Science

View from Germantown
Advanced Scientific Computing Research

Presented to the
ASCAC

by

Barbara Helland
Associate Director

April 23, 2020

COVID-19 High Performance Computng Consortium

https://covid19-hpc-consortium.org/

- Unique private-public effort spearheaded by the White House Office of Science and Technology Policy, the U.S. Department of Energy and IBM to bring together federal government, industry, and academic leaders who are volunteering free compute time and resources on their world-class in support of COVID-19 research.
- Announced on March 22, 2020
- Current Members:
  - Industry: IBM, Amazon Web Services, AMD, BP, Dell, Google Cloud, Hewlett Packard Enterprise, Microsoft and NVIDIA
  - DOE National Labs: ANL, LLNL, LANL, ORNL, LBNL and SNL
  - Agencies: NSF, NASA
  - Universities: MIT, RPI, U. of Illinois, U. of Texas, Austin, UCSD, Carnegie Mellon, U. of Pittsburgh, Indiana University, University of Wisconsin-Madison
- Scientific Review Committee reviews proposal in morning and Computing Matching committee matches resources with projects in the afternoon
- Progress as of 4/17/2020: 65 proposals have been submitted, 35 projects have been approved; 27 enabled and started projects; 4 waiting for user accounts and 1 awaiting resource selection by PI.

COVID-19

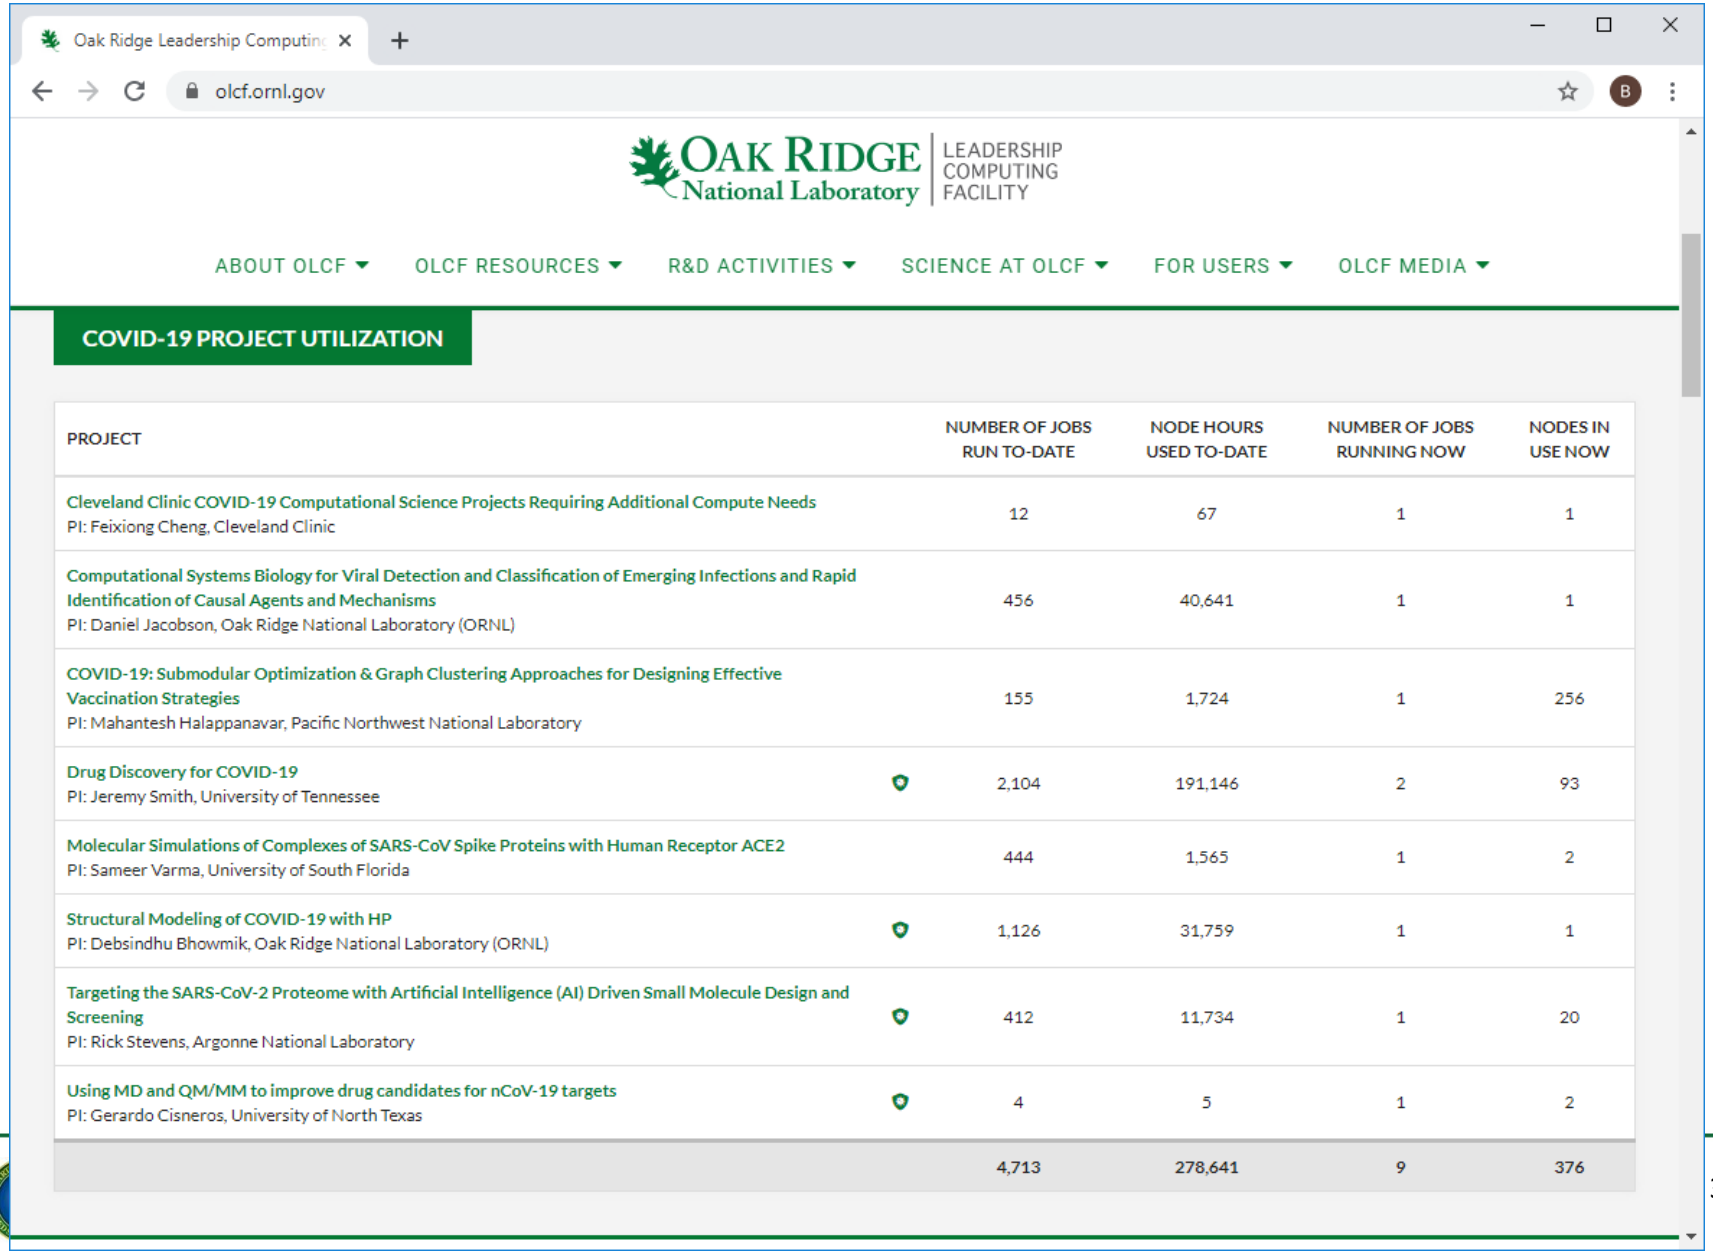

COVID-19 PROJECT UTILIZATION

PROJECT	NUMBER OF JOBS RUN TO-DATE	NODE HOURS USED TO-DATE	NUMBER OF JOBS RUNNING NOW	NODES IN USE NOW
Cleveland Clinic COVID-19 Computational Science Projects Requiring Additional Compute Needs PI: Feixiong Cheng, Cleveland Clinic	12	67	1	1
Computational Systems Biology for Viral Detection and Classification of Emerging Infections and Rapid Identification of Causal Agents and Mechanisms PI: Daniel Jacobson, Oak Ridge National Laboratory (ORNL)	456	40,641	1	1
COVID-19: Submodular Optimization & Graph Clustering Approaches for Designing Effective Vaccination Strategies PI: Mahantesh Halappanavar, Pacific Northwest National Laboratory	155	1,724	1	256
Drug Discovery for COVID-19 PI: Jeremy Smith, University of Tennessee	2,104	191,146	2	93
Molecular Simulations of Complexes of SARS-CoV Spike Proteins with Human Receptor ACE2 PI: Sameer Varma, University of South Florida	444	1,565	1	2
Structural Modeling of COVID-19 with HP PI: Debsindhu Bhowmik, Oak Ridge National Laboratory (ORNL)	1,126	31,759	1	1
Targeting the SARS-CoV-2 Proteome with Artificial Intelligence (AI) Driven Small Molecule Design and Screening PI: Rick Stevens, Argonne National Laboratory	412	11,734	1	20
Using MD and QM/MM to improve drug candidates for nCoV-19 targets PI: Gerardo Cisneros, University of North Texas	4	5	1	2
	4,713	278,641	9	376

ESnet is enabling remote work across the DOE complex in the time of COVID-19

ESnet backbone traffic originating from top home broadband providers
bits per second (bps) [-In / +Out]

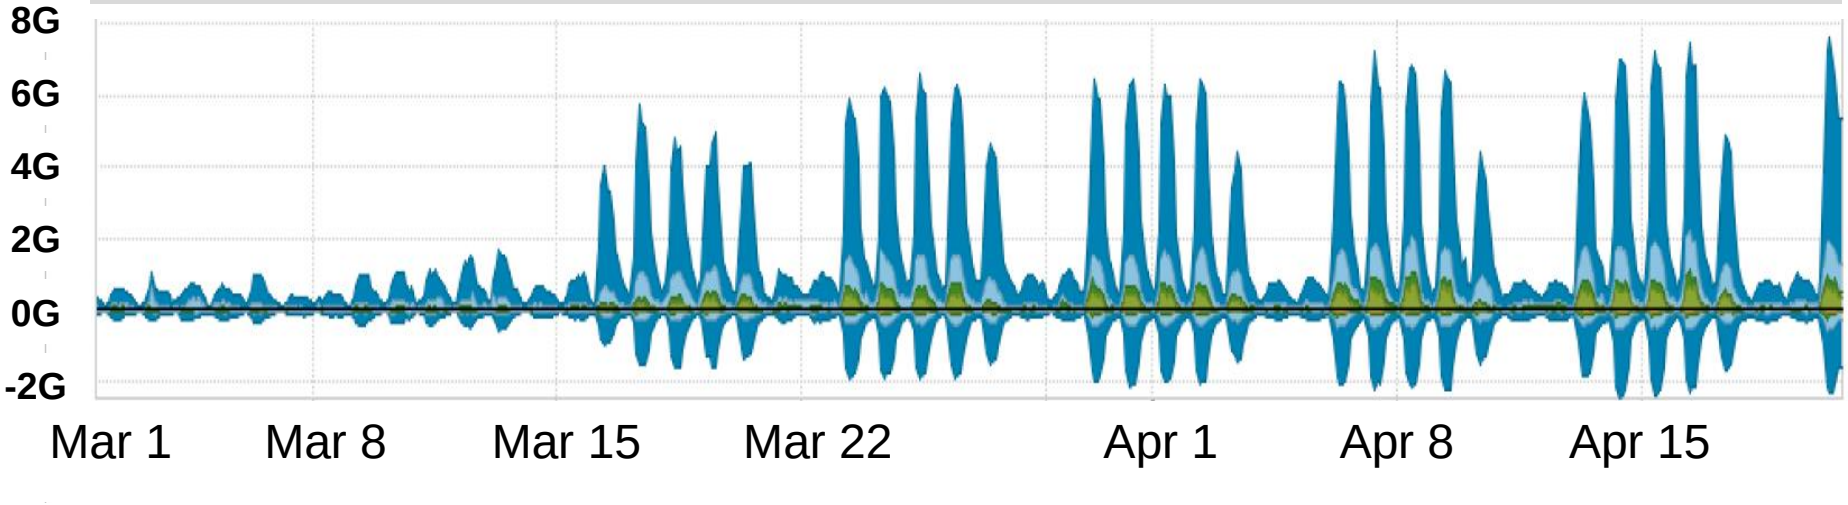

Each color represents a separate broadband provider. The traffic does not include entertainment streaming services.

Preliminary analysis indicates that the traffic increase is likely due to employees working remotely and connecting into their institutional enterprise IT system via VPN or other means.

- Comcast
- AT&T
- CCI (partial)
- Sparklight
- Frontier

ESnet's legacy of network services innovation is helping broadband providers maintain customer connectivity

ESnet developed iperf3 as a tool for measuring available connection bandwidth

- Available either as a component of perfSONAR or as a standalone tool
- Released as an open-source, BSD license: https://github.com/esnet/iperf/

Comcast has integrated iperf3 into millions of its cable modems and many of its servers for performance measurement and testing

- Comcast has measured a 32% increase in peak traffic during COVID-19.
- Largest known iperf3 deployment; iperf3 is used to perform ~700K speed tests per day.
- Comcast has contributed bug reports, fixes, and enhancements via Github, and sponsored an iperf3 developer event in April 2019 for information exchange.

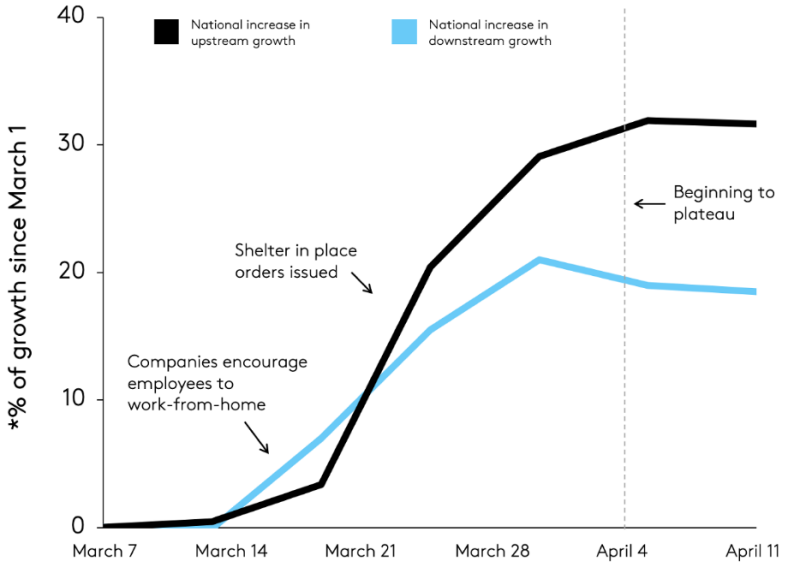

https://corporate.comcast.com/covid-19/network

FY 2020 President's Office of Science Budget

FY 2018 Enacted: $6.260B – ASCR $647.0M

FY 2019 Enacted: $6.585B – ASCR $935.5M

FY 2020 Enacted: $7.000B – ASCR $980.0M

FY 2021 Request: $5.838B – ASCR $988.1M

Office of Science Priorities

- Support High Priority research investments
  - Integrated Computational and Data Infrastructure for Scientific Discovery, Next Generation Biology Initiative, Rare Earth / Separation Science Initiative, Revolutionizing Polymer Upcycling, Strategic Accelerator Technology, and Data and Computational Collaboration with NIH.
- Continue investments in Exascale Computing, Artificial Intelligence/ Machine Learning, Quantum Information Science, Microelectronics, DOE Isotope Initiative, Biosecurity, and U.S. Fusion Program Acceleration.
- Continue operations of the national laboratories
  - SC oversees the operation of 10 DOE national laboratories. SC conducts a formal laboratory strategic planning process annually with its labs to understand future directions, immediate and long-range challenges, and resource needs.
  - Look to strengthen smaller/single purpose laboratories to be more multidiscipline
  - Strengthen/upgrade core laboratory infrastructure, i.e. utilities and laboratory workspace.
- Initiate new Infrastructure Projects
  - Reduce backlog of deferred maintenance
  - Improve obsolete infrastructure at National Laboratories
- Maintain all on-going line-item construction and MIE projects.

ASCR - FY 2021 Highlights

- **New SC Initiatives:**
  - **Integrated Computational and Data Infrastructure for Scientific Discovery (+11.8M) --** In partnership with the other SC programs, ASCR will design and deploy a flexible multi-tiered data and computational management architecture that enables a diverse array of on-demand scientific workflows and simulations for SC mission research—incorporating AI from the beginning to revolutionize the scientific productivity of our facilities and open up new avenues to scientific discoveries.

  - **Data and Computational Collaboration with NIH (+1.0M)**– based on experiences with the Cancer moonshot pilot with NCI – support DOE laboratories role in partnership with NIH to expand the capabilities of DOE's tools and address NIH's rapidly growing data and computational challenges.

ASCR - FY 2021 Highlights (cont …)

- **On-going Administration priorities**
  - **Artificial Intelligence/Machine Learning($56.0M, +$20M from FY2020 enacted)** – increases in the Applied Mathematics, Computer Science and Computational partnerships continue support of foundational research in machine learning and artificial intelligence to improve the reliability, robustness and interpretability of big data and AI technologies; development of new algorithms, methods and software tools for extracting information from scientific and engineering data; development of collaboration tools to ensure the seamless integration of Big Data with computing; and partnerships to broaden the applicability of AI and big data solutions across a range of applications.
  - **Quantum Information Science (QIS) ($86.2M, +$31.2M from FY2020 enacted)** – continue basic research investments in quantum algorithms, applications and networking and in partnership with the other programs in the Office of Science, QIS Centers. *Also initiates early stage research associated with developing a quantum internet.*
  - **Strategic Computing ($428.9M, -$27.8 from FY2020 enacted)** Exascale Initiative in partnership with NNSA which will see the deployment of the first exascale systems in calendar 2021.

ASCR FY 2020 President's Request
in thousands

	FY 2018	FY 2019	FY 2020		FY 2021
	Enacted Approp.	Enacted Approp.	Request	Enacted Approp.	Request
Mathematical, Computational, and Computer Sciences Research					
Applied Mathematics	34,104	28,206	41,500	41,500	53,728
Artificial Intelligence and Big Data (Non Add)	*(…)*	*(…)*	*(14,281)*	*(14,281)*	*(23,473)*
Computer Science	29,508	22,000	38,700	38,700	49,605
Artificial Intelligence and Big Data (Non Add)	*(6.402)*	*(2,000)*	*(9,719)*	*(9,719)*	*(14,875)*
Quantum Information Science (Non Add)	…	(3.000)	(5,000)	(5,000)	(7,000)
Computational Partnerships	49,910	75,667	60,959	69,142	75.051
Artificial Intelligence and Big Data (Non Add)	*(3.500)*	*(13,000)*	*(12,000)*	*(12,000)*	*(17,652)*
Quantum Information Science (Non Add)	*(6,349)*	*(16,214)*	*(16,708)*	*(20,680)*	*(18,709)*
Research and Evaluation Prototypes	**24,260**	**24,452**	**39,453**	**39,000**	**76,007**
CSGF	***(10,000)***	***(10,000)***	***(10,000)***	***(10,000)***	***(10,000)***
Quantum Information Science (Non Add)	***(14,260)***	***(14.452)***	***(29,453)***	***(29,000)***	***(60,453)***
SBIR/STTR	4,301	4,768	5,347	5,658	9,637
Total, Mathematical, Computational, and Computer Sciences Research	117,823	130,641	146,506	155,000	264,028
High Performance Computing and Network Facilities					
High Performance Production Computing (NERSC)	94,000	104,000	85,000	110,000	85,000
Leadership Computing Facility at ANL (ALCF)	110,000	140,000	150,000	150,000	150,000
Exascale (Non Add)	*(110,000)*	*(140,000)*	*(150,000)*	*(150,000)*	*(150,000)*
Leadership Computing Facility at ORNL (OLCF)	162,500	199,000	210,000	225,000	220,000
Exascale (Non Add)	*(62,500)*	*(100,000)*	*(125,000)*	*(125,000)*	*(125,000)*
Total, Leadership Computing Facilities	272,500	339,000	360,000	375,000	370,000
High Performance Network Facilities and Testbeds (ESnet)	79,000	84,000	80,000	90,000	90,000
SBIR/STTR	17,417	20,701	21,194	22,265	20,078
Total, High Performance Computing and Network Facilities	487,177*	572,153*	585,647*	636,265*	555,078
Exascale Computing					
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	205,000	232,706	188,735	188,735	168,945
Total, Advanced Scientific Computing Research	**647,000**	**935,500**	**920,888**	**980,000**	**988,051**

*Includes R & E Prototype funds

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

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

FAIR DATA AND MODELS FOR ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING
DE-FOA-0002306 and LAB 20-2306

Solicitation Scope:

- FOA and companion Lab Announcement will solicit efforts to make research data and artificial intelligence (AI) models findable, accessible, interoperable, and reusable (FAIR) to facilitate the development of new AI applications in SC's congressionally authorized mission space, which includes the advancement of AI research and development. In particular, ASCR will support FAIR benchmark data for AI; and FAIR frameworks for relating data and AI models.
  - **FAIR Benchmark Data**: Publicly available benchmark datasets are important tools to compare, enhance, train, and analyze AI models. Research resulting from this solicitation will facilitate the development of AI tools that are more suitable to science data and science tasks.
  - **FAIR Frameworks for data and models**: Educated trial and error continues to guide advances in science applications of AI. Currently, theory to address critical, foundational questions such as: What information about a dataset can be deduced from a model trained on the data? For a given data set and task, what are the best model, hyper-parameters, and training method? When are more data needed, and how much incremental information will they have? In what circumstances can a model be transferred to new data? Research resulting from this solicitation will facilitate systematic observation of data-model relationships to help answer these questions.

Context/Planning:

- Solicitation stems from key findings of the Office of Science roundtable on Data for AI.
- Solicitation is aligned with key strategies of the Administration's American AI Initiative.

Application requirements:

- Eligible Institutions: Universities/colleges, non-profit & for-profit organizations, DOE Labs
- Budgets for 2-3 year Awards: Labs at $300K - $750K per year; Universities at $150K - $750K per year
- Cooperative Agreements or field work authorizations
- **Letters of Intent due on April 17th. Proposals due on May 15th**
- An individual may participate in no more than three letters of intent.
- Preproposals will be used to expedite the review process only.

PM: Laura Biven

National QIS Research Centers

- First large scale QIS effort that crosses the technical breadth of SC
- Scope built on extensive community-wide RFI inputs– from technical scope to partnership model to management construct
- Seamlessly integrates the S&T innovation chain to accelerate progress in QIS R&D
- Maximizes teaming flexibility and options (TIAs, cooperative agreements, field work authorizations, interagency agreements) to foster direct participation by academics, national/federal labs, and for-profits
- Leverages other federal agency investments such as NSF's Quantum Leap Challenge Institutes and the NIST Quantum Economic Development Consortium (QEDC)

January 10th, 2020: FOA Issued
February 10th, 2020: Pre-apps Due
April 17th, 2020: Proposals Due

2-5 Awards
$10-25M/Year/Center
Up to $625M in 5 Years

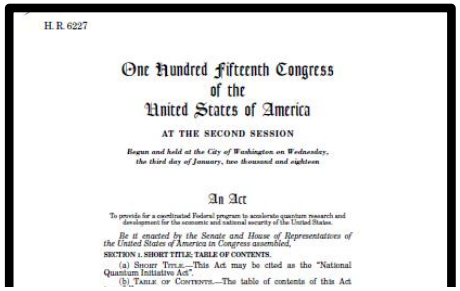
H. R. 6227

One Hundred Fifteenth Congress
of the
United States of America
AT THE SECOND SESSION

Begun and held at the City of Washington on Wednesday, the third day of January, two thousand and eighteen

An Act

To provide for a coordinated Federal program to accelerate quantum research and development for the economic and national security of the United States.

Be it enacted by the Senate and House of Representatives of the United States of America in Congress assembled,

SECTION 1. SHORT TITLE; TABLE OF CONTENTS.

(a) SHORT TITLE.—This Act may be cited as the "National Quantum Initiative Act".

(b) TABLE OF CONTENTS.—The table of contents of this Act

National Quantum Initiative Act

Hybrid quantum-classical approach to chemistry calculations

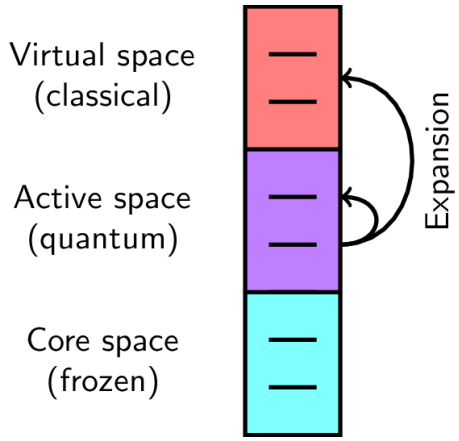

The virtual quantum subspace expansions employs a quantum computer to model the strongly-correlated active space of a molecule and a classical computer to find corrections originating from higher electronic orbitals in the virtual space.

M. Urbanek *et al.*, "Chemistry on quantum computers with virtual quantum subspace expansion", arXiv:2002.12902 [quant-ph]

Scientific Achievement

A novel algorithm improves accuracy of chemistry calculations on current quantum computers.

Significance and Impact

The method divides the work between a quantum and a classical computer in a scalable way. It achieves a higher accuracy than previous approaches without the need for additional quantum resources. It works well on noisy intermediate-scale quantum computers.

Research Details

- Developed a technique to minimize the impact of experimental noise on the solution stability.
- Tested the algorithm on the IBM Q Johannesburg quantum chip and calculated potential energy curves of the H_2 and Li_2 molecules in the ground state.
- Achieved accuracy that would require 20 qubits with previous methods using only 2 qubits, additional measurements, and a classical post-processing.

M. Urbanek, D. Camps, R. Van Beeumen, W. A. de Jong (LBL)

BERKELEY LAB

FASTMath provides machine learning for Neutron Science

Scientific Achievement

To advance existing software, we developed ML to refine peak locations, predict peak shapes and yield more accurate integrated intensities and structural reconstruction in crystallography.

Significance and Impact

ML able to resolve weak Bragg peaks, a challenge for neutron crystallography, leading to increased accuracy of the SNS instrument MaNDI.

Research Details

- Feeding the software need for crystallography tools.
- Parallel implementation on GPUs provide users with near-real-time feedback, distinct advantage for users faced with such an incredible amount of data. Up to 100Gb/s for X-ray crystallography.

Sullivan, Archibald, Azadmanesh, Vandavasi, Langan, Coates, Lynch, and Langan, "BraggNet: integrating Bragg peaks using neural networks", J. Appl. Cryst. (2019). 52, 854-863, 10.1107/S1600576719008665

The Macromolecular neutron diffractometer (MaNDi) is a single crystal diffractometer at the SNS

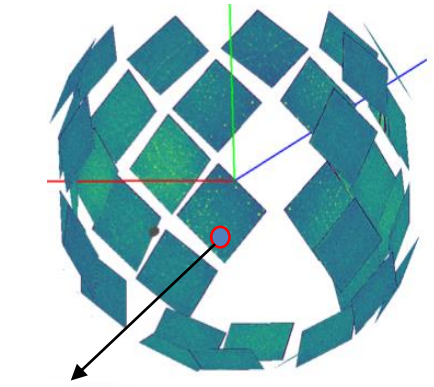

MaNDI detects thousands of Bragg Peaks that position, shape and orientation measure material properties

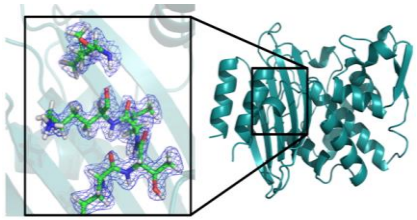

Better characterization of Bragg Peaks by Machine Learning improves inference of nuclear density maps (blue) which improves structure calculations(teal).

POC:
Rick Archibald: archibaldrk@ornl.gov

OAK RIDGE National Laboratory

DOE Applied Math Codes to Aid NRC Licensing Process for Nuclear Technologies

Argonne's fluid-thermal simulation codes, Nek5000 and NekRS, developed as part of the *High-Order Methods for High-Performance Multiphysics Simulations* project supported by the DOE Applied Math Base Program, are tapped to help DOE and the NRC in future licensing processes for nuclear technology.

Paul Fischer, Misun Min, Ananias Tomboulides (ANL)

Exploiting Advanced Computing and Modeling

- Nek5000 is a Gordon Bell and R&D 100 prize-winning code with high accuracy and demonstrated scaling to millions of processors.
- NekRS is a new GPU-oriented version of Nek5000 that is targeting exascale platforms.
- Nek5000/NekRS are based on efficient spectral element discretizations that realize high accuracy with minimal data movement.

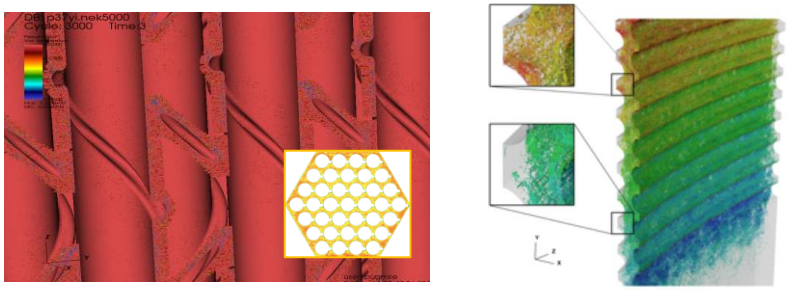

Turbulent flow in a 37-pin rod assembly simulated by NekRS on 6144 GPUs (left) and the flow structure in a helical coil heat exchanger by Nek5000 on CPUs (right)

Gaining Insight into Reactor Performance

- These updated codes can help expedite the review process and can be used to predict expected reactor operations, including fuel and material performance.
- These capabilities will ultimately reduce the time it takes to validate and certify new designs, enabling a faster commercialization process.

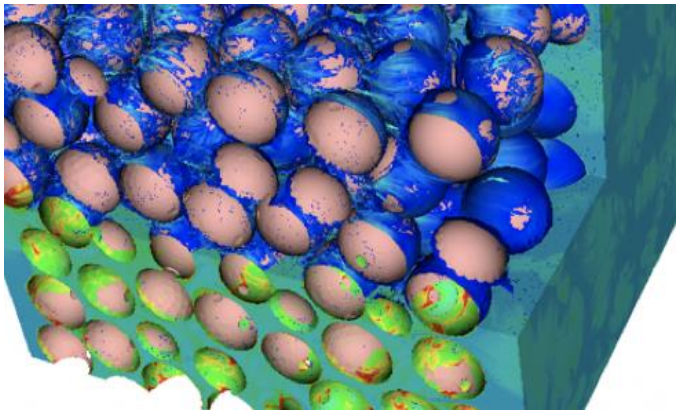

Nek5000 flow and temperature profiles in a pebble bed reactor.

Sharing Technical Expertise

- Nek5000/NekRS put advanced numerical methods into engineering practice to allow efficient simulations of turbulence with high-fidelity on platforms ranging from clusters to the world's fastest supercomputers.
- Through a new initiative, the team will work with the DOE's National Reactor Innovation Center and the NRC to provide state-of-the-art modeling codes to support licensing of advanced reactors.

- **Workshops: Past and Future**

5G-Enabled Energy Innovation Workshop
https://www.orau.gov/5GScience

Held March 10-12, 2020 at the Westin Michigan Avenue Hotel, Chicago, Il.
Goal: Deliver a community-based report highlighting 5G and beyond basic research, development, applications, technology transition, infrastructure, and demonstration opportunities in support of the U.S. DOE mission. The report will help the DOE Office of Science understand both the challenges and the opportunities offered by 5G and emerging advanced wireless technologies in the areas of basic research, development, and integration into scientific user facility operations.

Contributors: Aaron Tremaine, SLAC; Andrew Wiedlea, LBNL; Andy Nonaka, LBNL; Angel Yanguas-Gil, ANL; Arden Warner, FNAL; Arupjyoti Bhuyan, INL; Barney Maccabe, ORNL; Caleb Phillips, NREL; Charlie Catlett, ANL; Draguna Vrabie, PNNL; Elena Peterson, PNNL; Eric Schwegler, LLNL; Greg Tchilinguirian, PPPL; Harinarayan Krishnan, LBNL; Jason Fields, NREL; Jerome Lauret, BNL; Johnathan Cree, PNNL; Keith Tracey, SNL; Kevin Brown, BNL; Klaehn Burkes, SRNL; Kurt Sorensen, SNL; Luke Gosink, PNNL; Mark Bryden, AMES ; Matt Bickley, JLAB; Mike Ritsche, ANL; Pete Beckman, ANL; Peter Barnes, LLNL; Peter Fuhr, ORNL; Prasanna Date, ORNL; Scott Collis, ANL; Tammy Chang, LLNL; Theresa Windus, AMES; Thomas Potok, ORNL ; Valerie Taylor, ANL

Draft 5G for Science ”Priority Research Directions”

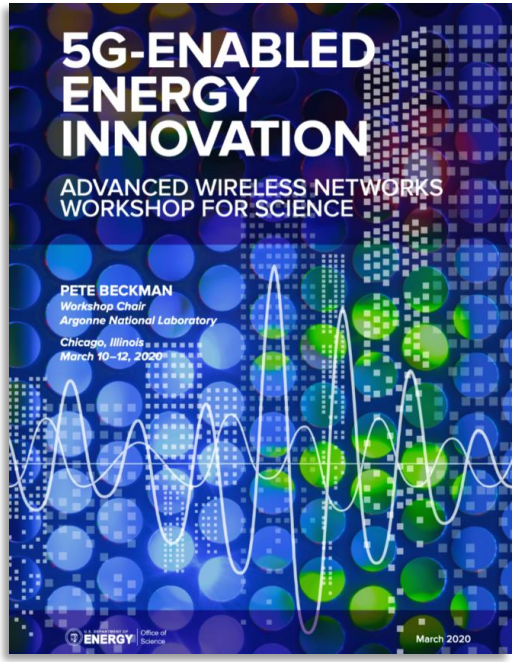

- **Revolutionize wireless communication in extreme** environments through advances in materials science and physics
- **Reinvent scientific instrumentation and critical national infrastructure** with wireless technology to provide rapid, AI-driven adaptation
- **Reinvent the digital continuum linking** the wireless edge to advanced scientific user facilities, data analysis, and high-performance computing
- **Revolutionize AI-enabled edge computing** for advanced wireless
- **Accelerate innovation** through the use community testbeds to explore advanced wireless for science

Community of Interest (on Future Scientific Methodologies) Workshop [Aug 4-6, 2020]

- **ASCR funded workshop to create a vision for how future computational fabrics will shape, and be shaped by, scientific advances that will occur 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 the following topics.**
  - Tomorrowland Vision – what could cause a dramatic shift in how science is done (tabletop accelerators for material science exploration)
  - Interfaces of the Future – how scientists and agents will interact and communicate
  - Computing Facilities of the Future – how will scientists/agents generate, process, move, find, repurpose, and store data in a global computational fabric
  - Future Methodologies – what are the broader implications of this future

* 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

Example Tomorrowland Vision Topic details

A world where small scale accelerators are available for use in individual labs, creates a number of important implications for the conduct of science. Studies of material structures and crystallography would no longer be constrained by the limited availability of beamtime. This could radically change the speed of scientific discovery—provided the next level of constraints were addressed. New methods that can produce large quantities of high quality experimental samples would need to be developed and widely deployed. New data management challenges would arise as curation, storage, analysis, visualization, and sharing of data are needed to drive future discoveries. This would require greater fidelity in computational models, which could, in turn, lead to refinements in which samples are produced and to new lines of inquiry.

The social impact of faster discovery could be the development of beneficial medicines and faster time to market of revolutionary materials.

Data Reduction Workshop

Background:

- Data reduction is a critical capability across the Office of Science. SC-supported experiments, observations, and simulations produce data at volumes and rates that can overwhelm our capabilities in storage, preservation, and analysis.
- Data reduction can take many forms from triggering in an experimental detector, filtering, sampling, feature detection, compression, and reduced order modeling. Currently, these techniques are being developed by distinct communities with relatively little overlap and coordination.
- From creation to analysis, many data reduction methods may be used at various points in a workflow. It's not always clear, however, how individual or combined data reduction steps can affect important properties of the data or how they will affect error and bias in derived quantities.

Details:

- **Date:** October, 2020 (exact dates TBD)

ASCR Organizing PMs: Bill Spotz, James Ricci, Laura Biven

- **FACILITIES**

Operational Assessment Review (OAR)

- The OAR is ASCR's primary management process for assessing the status and quality of operations of its scientific user facilities.
- The OAR focuses on four principal areas:
  - measurement of user results;
  - measurement of business and strategic results;
  - assessment of innovation; and
  - assessment of risk management.
- The OAR process results in an important record of procedure and accomplishment that informs ASCR and SC leadership, and which is used to respond to a variety of external data calls throughout the year. The OAR and review comments are a critical input to the Office of Science's annual Laboratory Appraisal Process.

OAR at a Glance

- **Purpose:** Assess status and quality of ASCR facility operations in previous Calendar Year (CY)
  - Document processes, compliance, metrics and improvements
  - Share best practices across facilities
  - Inform future improvements and reinforce success
- **Format:** Guidance memo from ASCR informs written Operational Reports with 5-7 Science Accomplishments from each facility, which are peer reviewed by staff from other DOE facilities.
- **Timing:** Yearly report and paper peer reviews with one rotating on-site review (approximately once every 4 years for each facility)
- **Result:** Peer feedback is shared with facility staff who respond to recommendations in writing. A summary of the facility OARs informs ASCR management as well as next year's guidance.

An impressive, productive lineup of *accelerated node* systems supporting DOE's mission

First U.S. Exascale Systems

2016	2018	2020	2021-2023

ORNL
IBM/NVIDIA

ORNL
Cray/AMD/AMD

ANL
Cray/Intel KNL

ANL
Intel/Cray

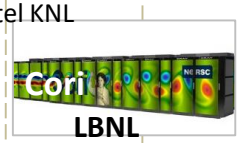

LBNL
Cray/Intel Xeon/KNL

LBNL
Cray/AMD/NVIDIA

LANL/SNL
Cray/Intel Xeon/KNL

LLNL
IBM/NVIDIA

LANL/SNL
TBD

LLNL
Cray

Perlmutter Configuration

- 12 GPU-accelerated cabinets
  - NVIDIA Volta-Next GPUs and Milan CPUs
- 12 CPU cabinets
  - Milan CPUs
- 35 PB All-flash storage
- Cray Slingshot High-performance network

System will support
- Simulation
- Complex workflows
- Data-Intensive workloads

GPU cabinets available in late 2020
CPU cabinets available in 2021

NERSC Science Highlights

Machine Learning for Solar Cell Design

Impurity energy levels in semiconductors determine the characteristics of solar cells. This study ran high-throughput atomistic simulations and generated computational datasets of impurity properties in two classes of semiconductors: Pb-based hybrid perovskites and Cd-based chalcogenides. These datasets led to machine-learned models that accelerate prediction and design for the entire space of materials and impurities in these semiconductor classes.

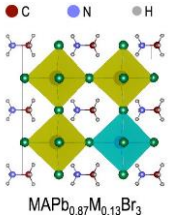

Maria Chan, ANL

Superfacility supports LCLS-II

The LLAna pilot project* was awarded $1M from ASCR & BES to promote the design and deployment of a new computing environment for next-generation free electron lasers. The team is developing tools for composable workflows, data management and analysis for LCLS-II. Jupyter notebooks have been adapted to run on HPC and at SLAC, capability enabled to convert LCLS-II file format into HDF5, and LCLS-II workflows are being profiled and optimized to run at NERSC.

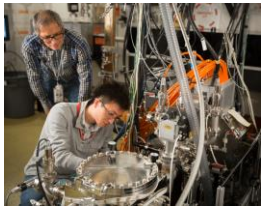

LCLS, SLAC

Uncovering Reasons for Sodium Battery Degradation

Researchers have uncovered a reason for the loss of capacity that occurs over time in sodium batteries: the unintended presence of hydrogen, which leads to degradation of the battery electrode.

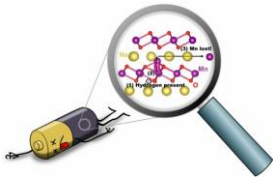

C. Van de Walle, UC Santa Barbara

Insights into the Human Gut Microbiome

Understanding the human gut microbiome through genome sequencing of gut microbes is crucial for human health. But only about 50% of species have a sequenced genome. This project reconstructed 61K genomes from publicly available gut metagenomes, including those of 2,058 previously unknown species, bringing the number of known human gut species to 4,558.

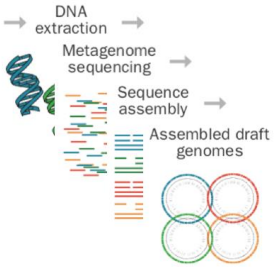

Stephen Nayfach, JGI

Pointing DESI in the Right Direction

The Dark Energy Spectroscopic Instrument (DESI) has just begun creating the largest 3D map of the universe's distribution of "ordinary" matter. To prepare, the DESI team has been using NERSC to build catalogs of the most interesting observational targets, modeling the shapes and colors of over 1.6 billion individual galaxies detected in 4.3 million images collected by three large-scale sky surveys.

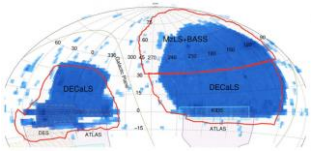

David Schlegel, LBNL

New Understanding of Turbulent Magnetized Plasma Flows

Researchers have discovered two separate conservation laws over a large range of length scales in turbulent magnetized plasma flows. These findings have important application to the study of galaxies and clusters, nebulae and the interstellar medium, stellar evolution, solar winds and space weather, nuclear fusion, and metallurgy.

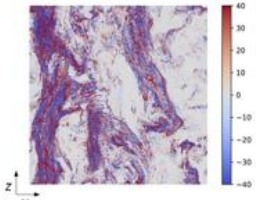

H. Aluie, U. Rochester

*LBNL-LCLS-II collaboration for data Analytics

ALCF 2019 Science Highlights

Cellular Control for Bioenginnering

Researchers using ALCF are exploring engineered pairs of proteins that bind exclusively to one another. These types of protein would enable sophisticated cellular control logic in living systems for bioengineering applications—with potentially large impacts for medicine and biomaterial production.

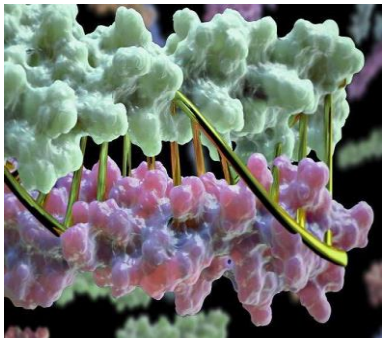

Detecting cancer at early stage

Researchers developed algorithms to create a “microscope in a computer.” They used ALCF resources to obtain more accurate and efficient early-stage cancer detection with minimal false positives.

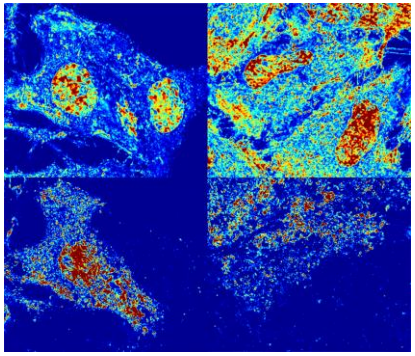

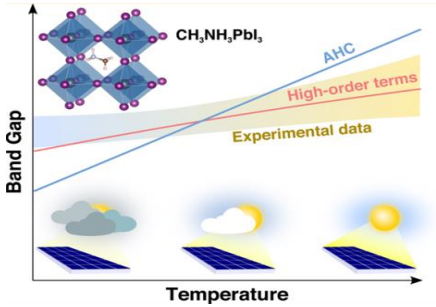

Durable and Reliable Solar Cells

Temperature of solar cell materials is critical to their durability and reliability. A team using ALCF studied electron-phonon coupling, reaching important conclusions on this area..

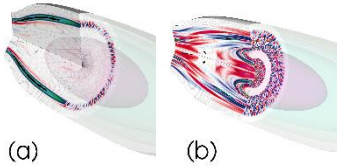

simulation of turbulence in fusion plasma

Commercially Viable Nuclear Fusion

TAE Technologies is designing a commercially viable fusion-based power plant driven by an aneutronic hydrogen-boron (p-B11) fuel cycle. Models ran at ALCF are helping TAE understand fundamental questions of science and technology for this design.

Energy Efficient Devices

Future devices such as computer memory, electric motors, generators, and magnetic bearings, could be much more energy efficient if using a mechanism to control a major magnetic property in these thin-film magnetic metals.. A research team using ALCF studied Interactions between electron spin and an external magnetic field, making important discoveries.

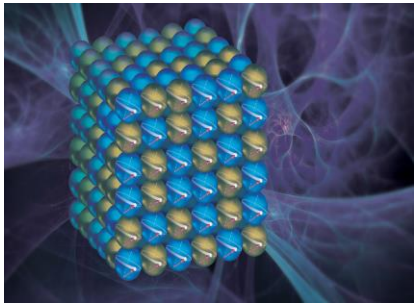

Machine Learning and Water

Using a machine-learning workflow, a team using ALCF resources devised a model that correctly predicts water’s key features, such as the melting point of ice, with a fraction of the computational cost of the best atomistic water models available today.

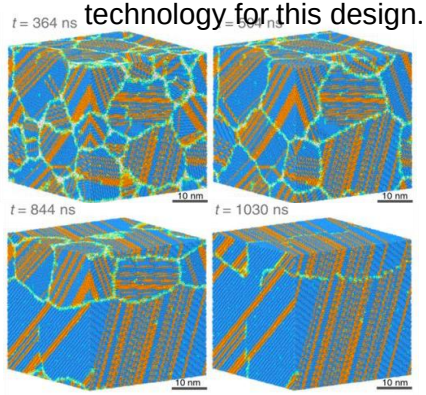

OLCF 2019 Science Highlights

Origins of Genetic Diseases

Researchers create the most complete model of complex protein machinery and use it to accurately identify clusters of gene mutations (spheres), which helped them study the emergence of various genetic diseases by unveiling their distinguishing molecular mechanisms

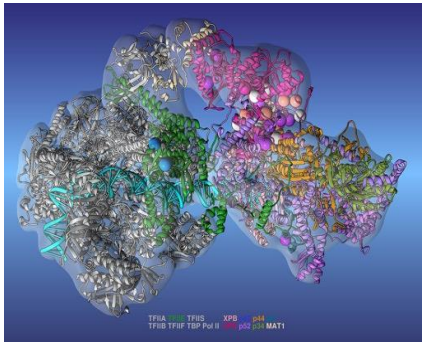

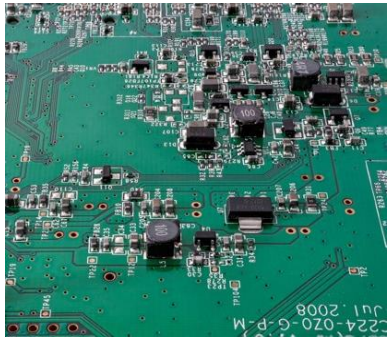

Gordon Bell Finalist

Team tackles transistors with a new programming paradigm to simulate nanoelectronics through data-centric lens. Gaining a 140-fold speedup they developed a map of where heat is produced in a single transistor. This will inform the production of new heat resistant semiconductors

Braking' for a Mars Landing

NASA team simulates a retropropulsion-powered descent to Mars surface, enabling the team to predict how the vehicle's engines should be designed and controlled with 35X increased performance.

GPUs Power Up GE Code

Simulating an entire multipassage, multirow turbine would require years of CPU simulations to complete a design. GENESIS's speedups on GPUs will allow GE to do more simulations in the same amount of time or to perform larger, higher-resolution simulations, turning intractable design challenges to nearer-term simulation capabilities

Quantum supremacy

A joint research team from Google Inc., NASA Ames Research Center, and OLCF has demonstrated that a quantum computer can outperform a classical computer at certain tasks. A software library developed by ORNL allowed the team to take full advantage of Summit to run a quantum benchmark code.

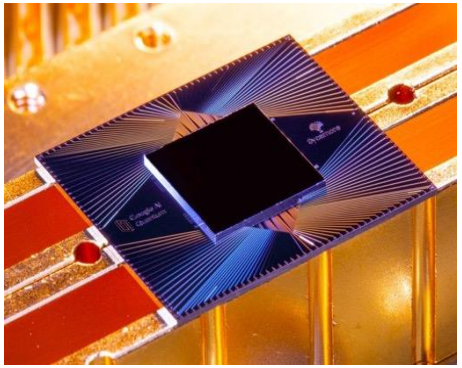

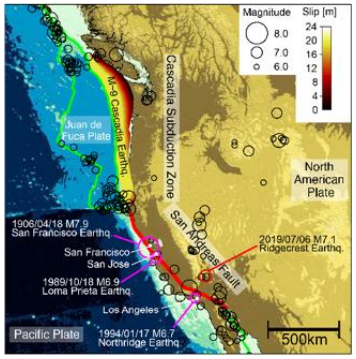

Earthquake Prediction

A team studying the deformation of the Earth's tectonic plates is simulating areas where plate subduction can trigger earthquakes and volcanic activity. Harnessing an AI-like algorithm to gain 75X speedups they simulated a fault slip in the Cascadia Subduction Zone.

ASCR Staff Changes

- **Bill Vanderlinde is returning to DOD, May 1st**
- **Betsy Riley is retiring at the end of April**
- **James Ricci has a new job starting June 1st**

- **The Computer Science Program Manager Position closed on April 15.**
- **Once we receive final approval, the Facilities Division Director**

DOE Leadership Development Program

- A corporate-wide approach to leadership development
- The program targets leadership development opportunities at all levels GS 13-15
- ECQ Driven;
- Competency-based

Program Components

- Strengths Assessment
- Development Sessions
- Professional Coaching
- Action Learning Project

Benefits

- Improved proficiency of leadership competencies
- Ability to collaborate across platforms
- Deeper exposure to the inner workings of the agency
- Increased competence
- Better overall leadership

ASCR's Laura Biven graduated with the first cohort in April, 2020

Congratulations Laura

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*
- **COVID-19 EFFORTS AT THE OLCF --** *Gina Tourassi, Oak Ridge National Laboratory*
- **COVID-19 RESEARCH –** *Rick Stevens, Argonne National Laboratory*
- **SCIENCE AND ENGINEERING INDICATORS –** *Julia Phillips, National Science Foundation*
- **EXASCALE UPDATE –** *Doug Kothe, ORNL and Lori Diachin, LLNL*
- **REPORT FROM AI TOWN HALLS –** *Valerie Taylor, Argonne National Laboratory*
- **WORKSHOP ON QUANTUM NETWORKING –** *Kerstin Kleese Van Dam, Brookhaven National Laboratory*
- **UPDATE ON CURRENT CHARGES**
  - *Artificial Intelligence, Tony Hey ASCAC*