U.S. DEPARTMENT OF ENERGY | Office of Science

View from Germantown

Advanced Scientific Computing Research Update

Barbara Helland, Associate Director
Ceren Susut, Computational Partnerships and Research Division Director
Ben Brown, Facilities Division Director

September 29, 2021

Some ASCAC Agenda Details

- A VISION FOR ASCR FACILITIES – *Ben Brown, ASCR Facilities Division Director*
- LAWRENCE AWARD TALK – *Rob Ross, Argonne National Laboratory*
- COVID-19 HPC Consortium -- *Bronson Messer Oak Ridge National Laboratory*
- EXASCALE UPDATE – *Andrew Siegal, Argonne National Laboratory*
- REPORT FROM COV – *Alexandra Landsberg, ASCAC*
- RANDOMIZED ALGORITHMS FOR SCIENTIFIC COMPUTING WORKSHOP REPORT – *Steven Lee, ASCR*
- TAULBEE, TALENT, AND TRENDS, *Betsy Bizot, Sandhya Dwarkadas, Erik Russell, Amanda Stent, and Burçin Tamer; Computing Research Association*
- THE LEGION PROGRAMMING SYSTEM– *Alex Aiken, Stanford National Accelerator Laboratory and Pat McCormick, Los Alamos National Laboratory*
- PERLMUTTER EARLY SCIENCE, *Jack Deslippe, Lawrence Berkeley National Laboratory*

ASCR's Leadership Computing Challenge Program Manager

Jordan Thomas

On September 26, 2021, Dr. Jordan Thomas joined ASCR as the new ASCR Leadership Computing Challenge (ALCC) program manager. Jordan has worked in ASCR for the last year as an AAAS Science & Technology Policy Fellow, assisting with ALCC, science communication, and the launch of the Office of Science's Public Reusable Research Data Resources (SC PuRe) initiative. Prior to her time at DOE, Jordan worked as a data analyst for a small consulting company where she specialized in data visualization, database architecture, and machine learning. Additionally, she previously served as an adjunct professor at George Mason University.

Jordan is a physical climate scientist with a background in global climate modelling and simulation. She obtained her PhD from Johns Hopkins University in oceanography and her B.S. in meteorology from Penn State University.

New Senior Technical Advisor, on Detail from NSF

Bill Miller

Dr. William (Bill) Miller joined ASCR in June 2021 as a visiting Senior Technical Advisor to the Facilities Division, on detail from the National Science Foundation.

At NSF, Bill is the Science Advisor for the Office of Advanced Cyberinfrastructure (OAC) in the Computer Information Science and Engineering (CISE) Directorate, where he leads efforts to foster new CI-enabled discovery pathways in science and engineering research. He has also been deeply involved in planning and oversight of NSF's large-scale scientific facilities, federal-wide initiatives, and international coordination on Open Science.

Prior to his government service, Bill conducted fundamental research in experimental neuroscience for several decades following a first career in space mission development. He holds a B.S.E. in Aerospace Engineering from the University of Michigan and a Ph.D. in Neuroscience from the University of California – Davis.

ADVANCED SCIENTIFIC COMPUTING RESEARCH

Functional Organization Chart

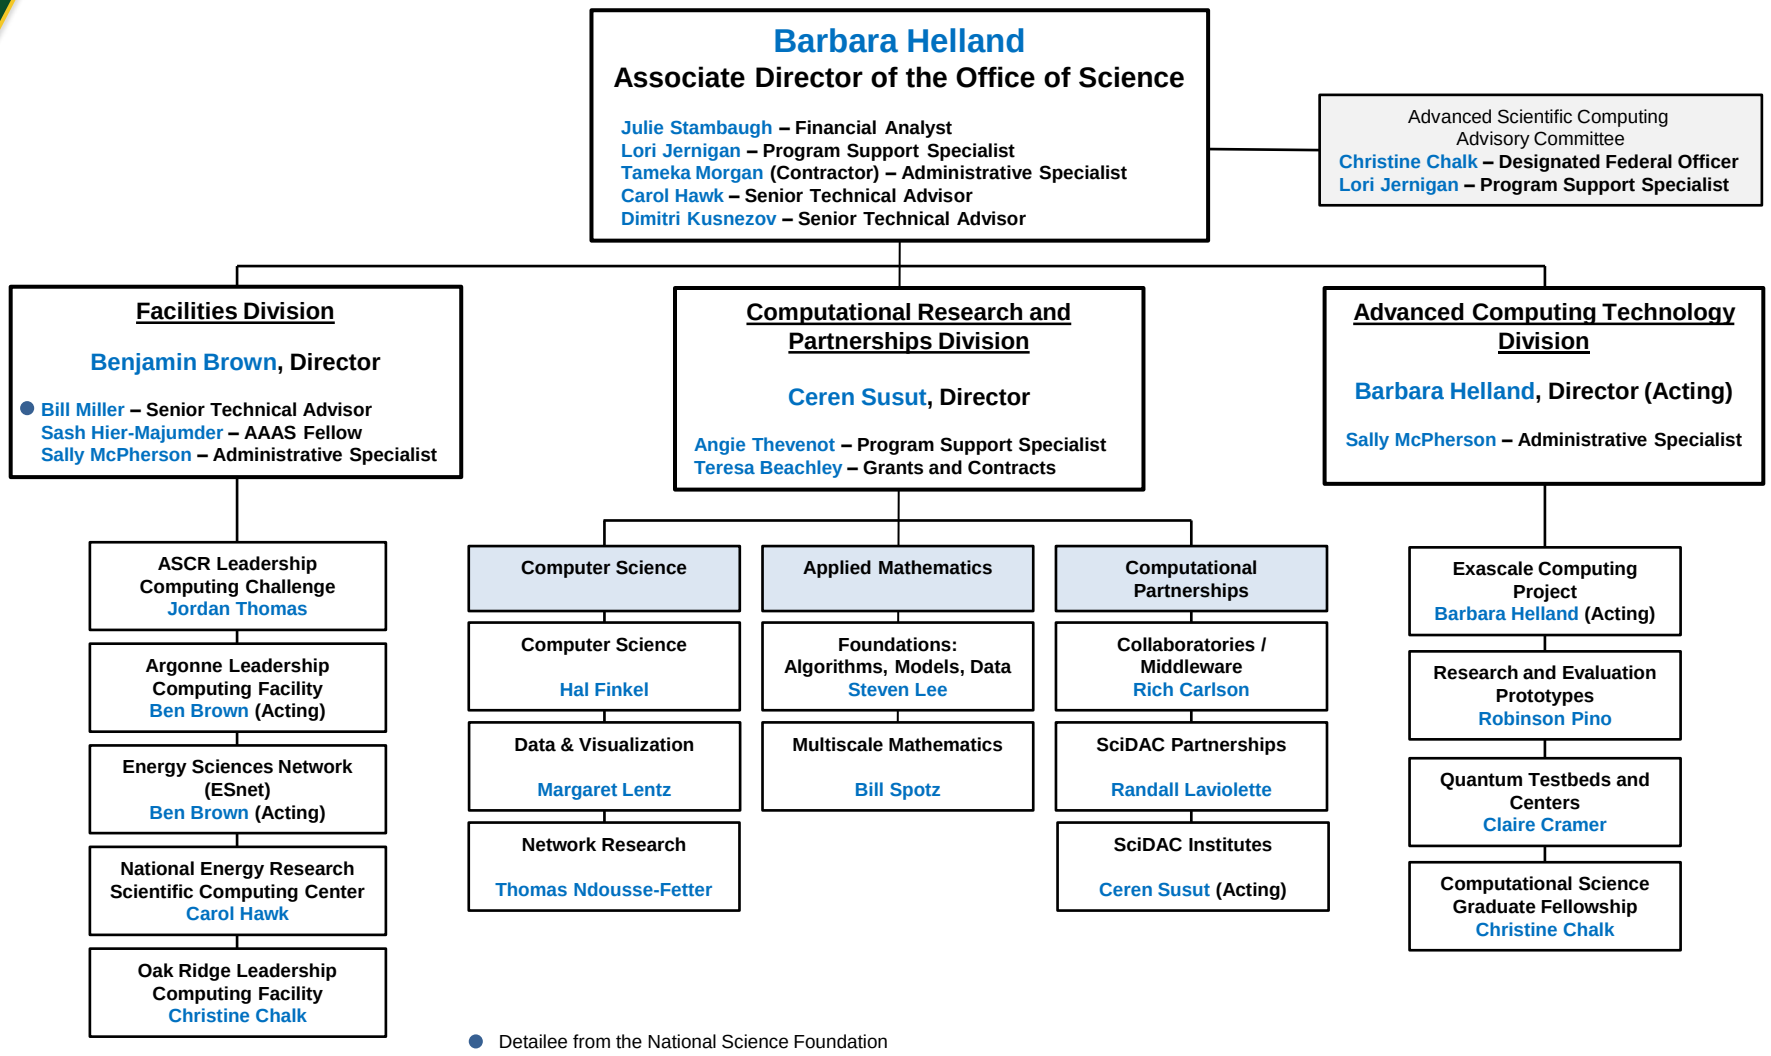

Congressional Actions: Reconciliation and Infrastructure

- Reconciliation: DOE Office of Science for fiscal year 2022 $10,310,804,000 to remain available until September 30, 2026
  - $7,780,566,000 for Constructions projects
    - Contains $1,335,016,000 for ECP, NERSC, ALCF, OLCF and Esnet
  - $1,470,238,000 for Major Items of Equipment
    - Contains $302,000,000 for the High Performance Data Facility
  - $2,000,000,000 for Research, Development and Demonstration activities
    - Contains $116,000,000 for the Computational Science Graduate Fellowship
    - Contains $340,000,000 for activities to facilitate access of researchers to US quantum computing facilities

https://science.house.gov/imo/media/doc/Science%20Committee%20Print.pdf

- Infrastructure Investment and Jobs Act-- $1T
  - Passed in Senate August 10, includes $550B in new federal investments in American Infrastructure over the next 5 years and tens of billions for applied R&D and technology demonstration programs at the Department of Energy and nearly $3 billion for the National Oceanic and Atmospheric Administration

https://www.congress.gov/117/bills/hr3684/BILLS-117hr3684eas.pdf

Continuing Resolution

- House passed a temporary, stop-gap funding measure on September 22, 2021 to keep the Government running through December 3. The CR includes $28.6B in additional disaster relief, $6.3B to support resettlement of Afghan evacuees and suspend the debt limit through December 16, 2022.
- Senate blocked the CR bill Monday night.

Other Congressional Actions

- Senate: United States Innovation and Competition of 2021 (USICA) (was Endless Frontier Act) passed June 8, 2021 -$200B, includes funding for the CHIPs Act and NSF research to compete globally in science.
- House:
  - the National Science Foundation for the Future Act – passed 345-67, June 28
  - the Department of Energy Science for the Future Act – passed 351-68 on June 28;
    - Bill provides funding for investment in renewable energies and research into emerging tech, and, in all, invests $50 billion over five years in the DoE Office of Science and National Labs. It would boost the DoE Office of Science Budget to $8.8 billion in FY2022, a $1.8 billion increase from its enacted FY2021 levels. By 2026, the office's budget would reach $11 billion if the bill is passed.

Other Updates

FY 2022 Funding Opportunity Announcements

- New Collaborator Template
  - Excel file submitted by the lead applicant listing collaborators for all individuals listed as senior/key personnel on the proposal.
  - Information used to manage reviewer selection.
  - Requested for all FY 2022 funding opportunities and lab calls.
- Biosketches and Current and Pending Support
  - Office of Science is working with SciENcv to develop a DOE SC template for biosketches and current and pending support.
  - Expected to be available on or before January 2022.
  - SciENcv can be linked to ORCID accounts to leverage existing data.

DOE SC Data Management Overview

SC data management principles		
Enable discovery	Share, preserve, validate	Cost management

Data Management Plan (DMP) requirements			
Share, preserve, validate	Make data associated with publications accessible	Availability of data management resources	Privacy, security, confidentiality

- DMPs are reviewed as part of the overall SC research proposal merit review process
  - Additional requirements and review criteria for the DMP may be identified in a solicitation
  - Proposals may include requested funding to implement a DMP, which will be considered during merit review

Complete information available at: https://science.osti.gov/Funding-Opportunities/Digital-Data-Management

Updates to Digital Data Management Guidance

- Office of Science is updating the **Suggested Elements of a Data Management Plan** (DMP) and adding **Guidance for Reviewers of Data Management Plans**
  - Current guidance will remain in effect for all solicitations issued through December 31, 2021
  - Updated guidance will be effective for all solicitations issued after **January 1, 2022**
  - There are **no changes to formal DMP requirements** that are part of solicitations

Suggested Elements of a DMP	Guidance for DMP Reviewers	PAMS Updates for Reviewers
• Suggested Elements offer guidance to researchers about what to include in a DMP • Provide a framework for planning a DMP that satisfies requirements • Tool to aid in aligning with best practices in data management	• Reviewers are asked if the DMP is suitable and supports validation of the proposed research • Reviewer guidance connects suggested elements to DMP requirements • Encourages constructive feedback to continue improving future DMPs	• PAMS emails to reviewers will include a link to the Guidance for Reviewers of DMPs • Reviewers will need to certify once a year that they have read the Guidance for Reviewers of DMPs

Complete information available at: https://science.osti.gov/Funding-Opportunities/Digital-Data-Management

The OLCF is ready for Frontier

Frontier is currently being installed at ORNL

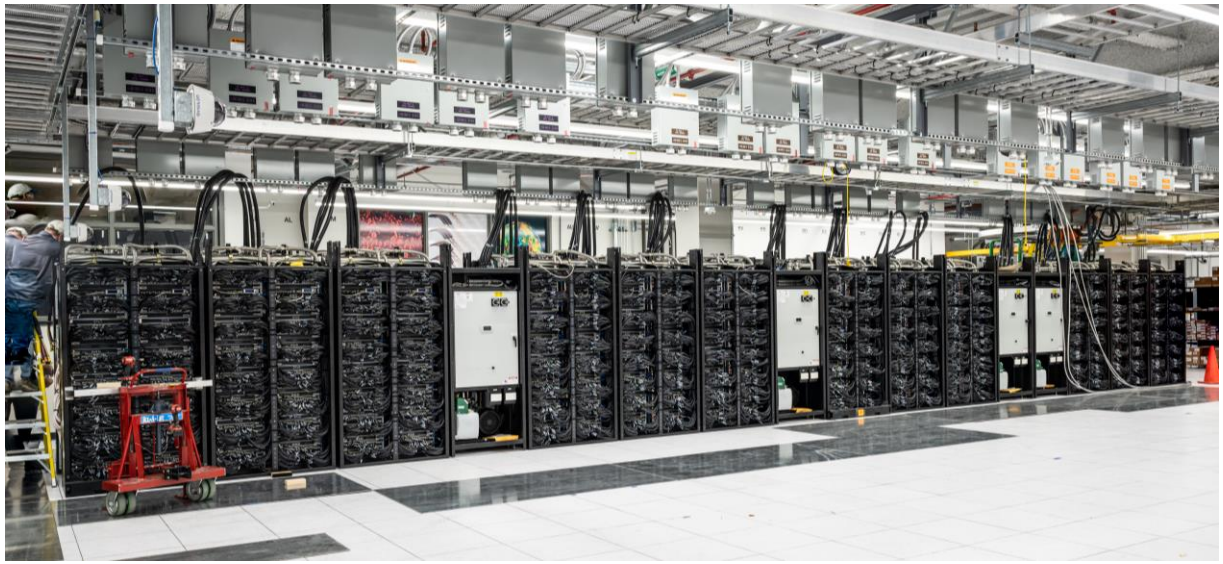

"Engage in research and development to create a capable exascale computing ecosystem that integrates hardware and software capability delivering at least 100 times the performance of current 10 petaflops (10^{15} floating-point operations per second) systems across a range of applications representing government needs and societal priorities such as Artificial Intelligence (AI) technologies. **By September 30, 2021, begin deployment (receiving and installing hardware) of at least one Exascale Computing system."** *DOE Agency Priority Goal*

The COVID-19 High Performance Computing Consortium

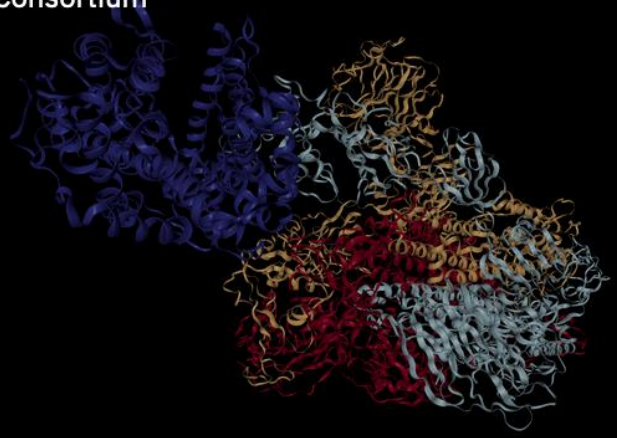

Bringing together the Federal government, industry, and academic leaders to provide access to the world's most powerful high-performance computing resources in support of COVID-19 research.

100

Projects

600

Petaflops

Active

Fighting COVID-19 will require extensive research in areas like bioinformatics, epidemiology, and molecular modeling to understand the threat we're facing and to develop strategies to address it.

Here are some of our projects.

See all

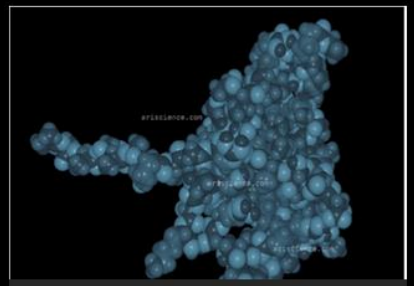

SARSCOV2/COVID19 protein interruption search with existing active compounds using quasi-quantum simulation

Therapeutics

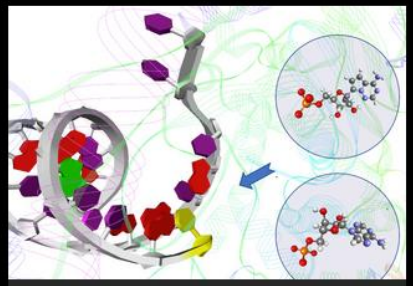

The Competition of Antiviral Drugs with ATP to Inhibit the SARS-CoV-2 RNA-dependent RNA Polymerase: A Key to Enhanced Drug Screening

Therapeutics

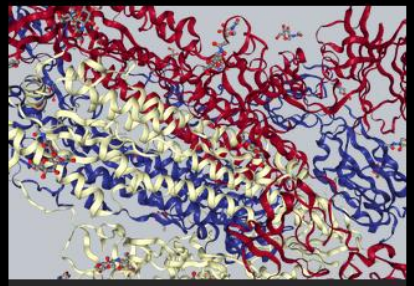

High-Performance Causal Inference for COVID-19 Mitigation and Response

Patients

Members and Affiliates

Academia

- Massachusetts Institute of Technology
- MGHPCC
- Rensselaer Polytechnic Institute
- University of Illinois
- University of Texas at Austin
- University of California - San Diego
- Carnegie Mellon University
- University of Pittsburgh
- Indiana University
- University of Wisconsin-Madison
- Ohio Supercomputing Center
- UK Digital Research Infrastructure
- CSCS – Swiss National Supercomputing Centre
- SNIC PDC – Swedish National Infrastructure for Computing, Center for High Performance Computing

Federal Agencies

- NASA
- National Science Foundation
  - XSEDE
  - Pittsburgh Supercomputing Center
  - Texas Advanced Computing Center (TACC)
  - San Diego Supercomputer Center (SDSC)
  - National Center for Supercomputing Applications (NCSA)
  - Indiana University Pervasive Technology Institute (IUPTI)
  - Open Science Grid (OSG)
  - National Center for Atmospheric Research (NCAR)

Industry

- IBM
- Amazon Web Services
- AMD
- BP
- D.E.Shaw Research
- Dell
- Google Cloud
- Hewlett Packard Enterprise
- Intel
- Microsoft
- NVIDIA

Affiliates

- Atrio
- Data Expedition
- Flatiron
- Fluid Numerics
- Immortal Hyperscale InterPlanetary Fabrics
- MathWorks
- Raptor Computer Systems
- SAS
- The HDF Group

Department of Energy National Laboratories

- Argonne National Laboratory
- Idaho National Laboratory
- Lawrence Berkeley National Laboratory
- Oak Ridge National Laboratory
- Lawrence Livermore National Laboratory
- Los Alamos National Laboratory
- Sandia National Laboratories

International

- Korea Institute of Science and Technology Information (KISTI)
- Ministry of Education, Culture, Sports, Science and Technology (MEXT) - Japan
  - RIKEN Center for Computational Science (R-CCS)

Timeline through May, 2020

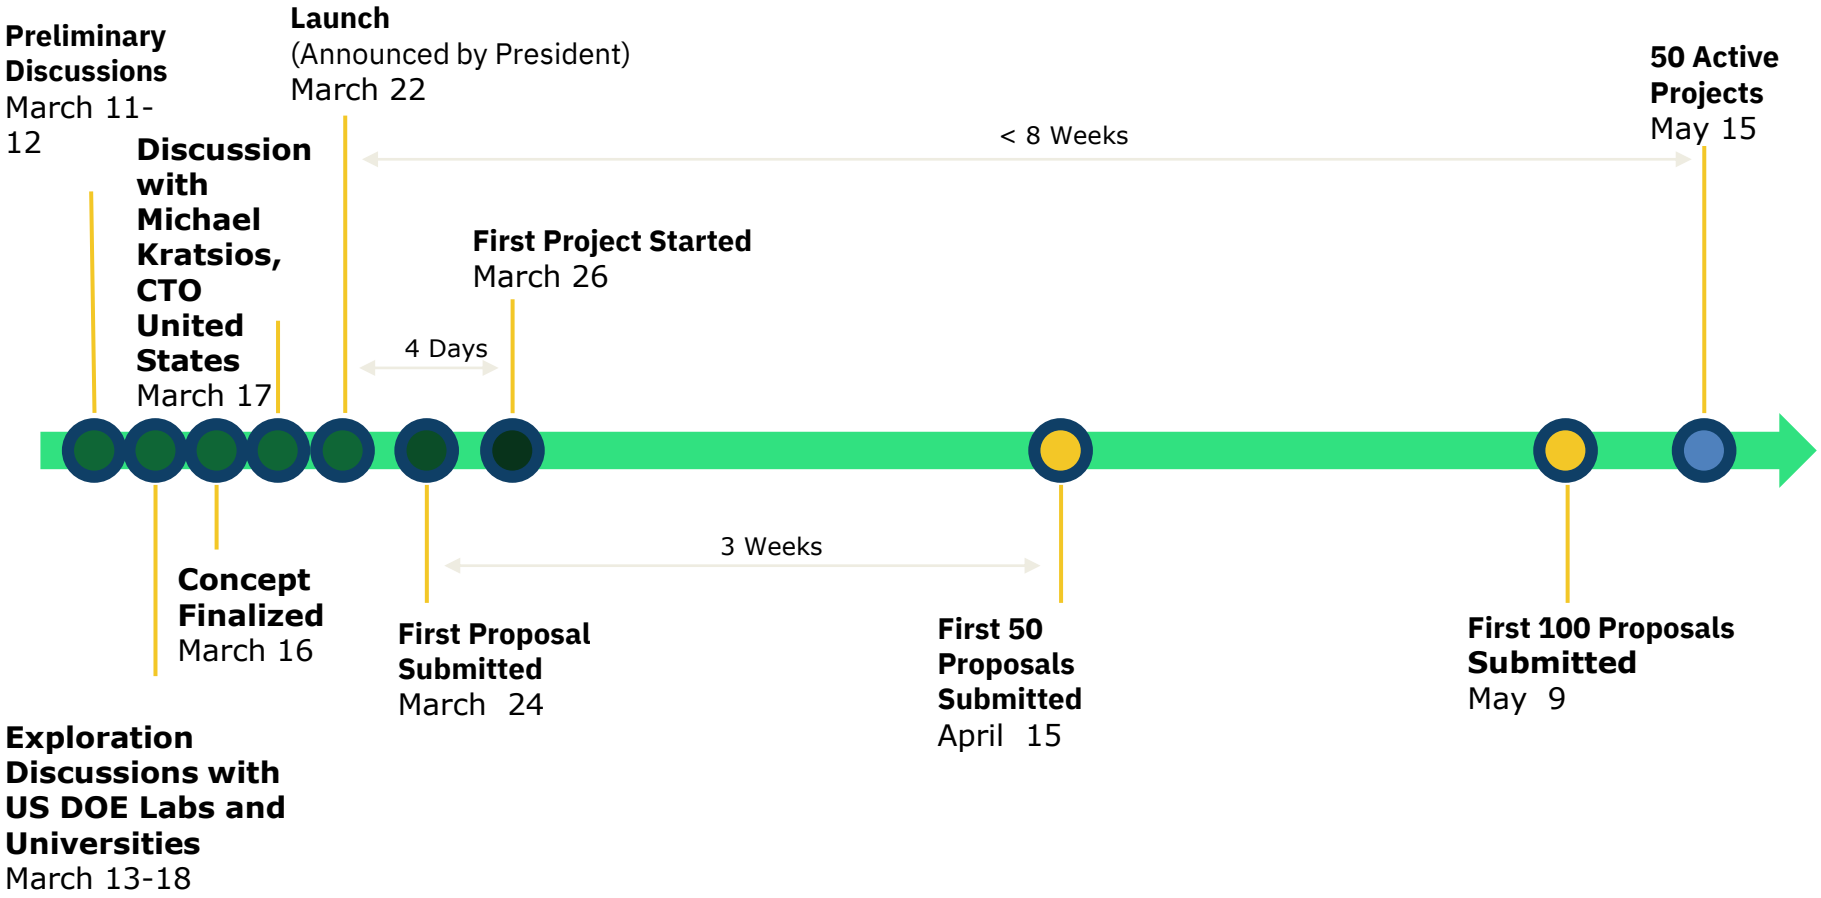

Research projects

QUANTITY	
Approved	**112**
Total projects submitted	213
Eligible for review	192
Declined	80
Returned without review	21
Completed	42

ORIGINATING ORGANIZATION	US	Non-US
Total	**82**	**19**
Company	7	4
Medical organization	2	1
R&D organization	10	7
University	63	7

32% of projects with experimental, clinical, or policy transition plans in place

Check out Bronson Messers' talk at 11:05 am ET Sept 30, 2021

Current focus

Given the continued impacts from COVID-19, as seen most recently with the prevalence of the Delta variant, the Consortium is still active with a focus on projects which can have near-term impacts in areas including, but, not limited to:

- Understanding and modeling patient response to the virus using large clinical datasets
- Learning and validating vaccine response models from multiple clinical trials
- Evaluating combination therapies using repurposed molecules
- Epidemiological models driven by large multi-modal datasets
- Mutation understanding and mitigation methods

ASCR Software-Stewardship Taskforce Update

- The task force has already met with a number of key groups, including:
  - The former ASCAC subcommittee on *Transitioning ASCR after ECP*.
  - ECP leadership
  - ASCR Facilities leadership
  - The Computational Research Leadership Council (CRLC), which is composed of Lab Division Directors
- Meetings with other science funding agencies and stakeholders are upcoming.
- The task force is preparing an RFI to seek broad feedback from the wider community.

Advanced Computing Technology Accomplishments

Advances in Machine Learning to Improve Scientific Discovery at Exascale and Beyond
PI: Thomas Potok, ORNL, MENNDL

- R&D 100 award winning research for (AI) software systems, Multinode Evolutionary Neural Networks for Deep Learning (MENNDL)
- Licensed to General Motors for use in autonomous vehicles.
- MENNDL could be used to generate and train the AI models that enable the self-driving technology in next generation autonomous cars.
- First commercial license for MENNDL.
- First AI technology to be commercially licensed by ORNL.
- Technology challenge: How can cars quickly and accurately perceive their surroundings to navigate safely through them?
- MENNDL offers the potential to quickly and accurately develop AI models to enable the next generation of self-driving cars.

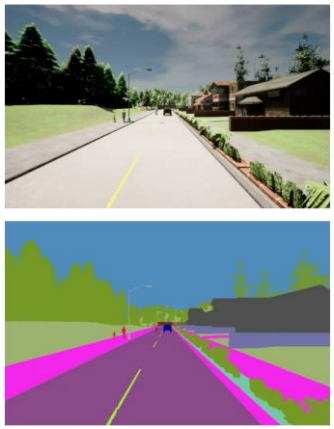

Automatically generating labelled data from simulation

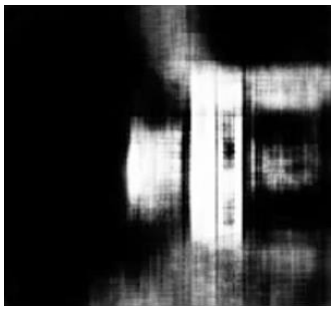

What the "driver" agent sees when driving

Mark Coletti et al., "**Diagnosing Autonomous Vehicle Driving Criteria with an Adversarial Evolutional Algorithm**," *Genetic and Evolutionary Computation Conference Companion*, 2021. DOI: 10.1145/3449726.3459573
Catherine D. Schuman et al., "**Neuromorphic Computing for Autonomous Racing**," *International Conference on Neuromorphic Systems*, 2021. DOI: TBD

Coordinate Transformations from Dragonflies to Neuromorphic Hardware:

- PI: Frances Chance, SNL
- Chance, F.S. (2021) Dragonfly AI: The computer bug you want. IEEE Spectrum, August 2021 issue.

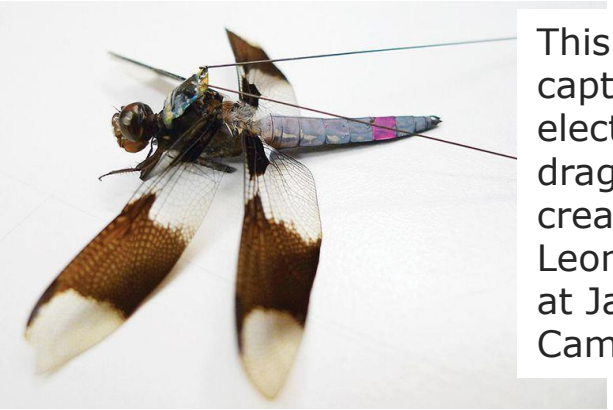

This backpack that captures signals from electrodes inserted in a dragonfly's brain was created by Anthony Leonardo, a group leader at Janelia Research Campus.

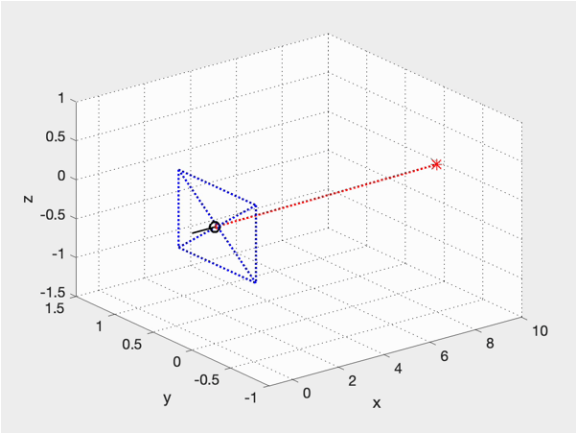

The model dragonfly reorients in response to the prey's turning. The smaller black circle is the dragonfly's head, held at its initial position. The solid black line indicates the direction of the dragonfly's flight; the dotted blue lines are the plane of the model dragonfly's eye. The red star is the prey's position relative to the dragonfly, with the dotted red line indicating the dragonfly's line of sight.

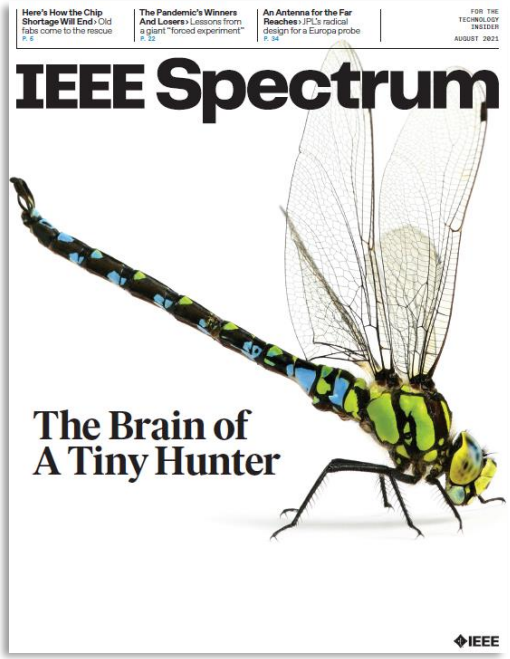

ASCR Quantum Computing Testbeds

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.

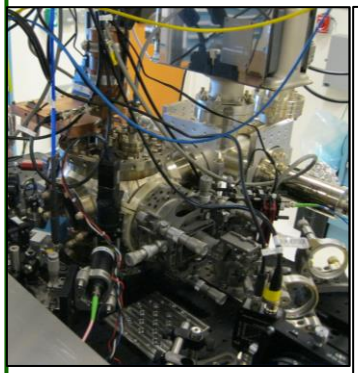

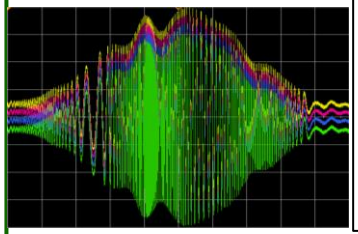

The **Quantum Science Open User Testbed (QSCOUT)** at SNL is the world's first publicly-available trapped ion quantum computer. Its custom assembly language gives users access to low level control and is compatible with most higher-level languages. State-of-the-art coherent pulse generation and a new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

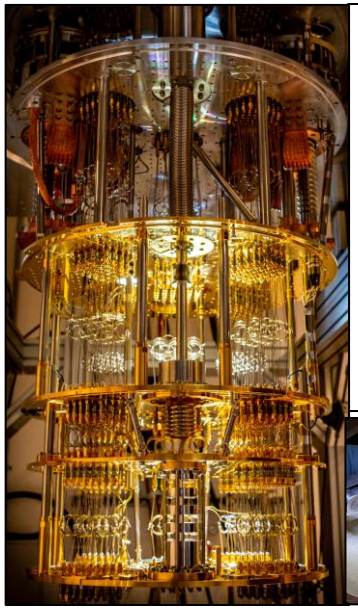

The **Advanced Quantum Testbed** at LBNL (**AQT@LBNL**) offers access to ternary logic (qutrits) as well as binary logic (qubits) in a unique superconducting platform. Second-generation hardware will include increased connectivity on a 3D chip.

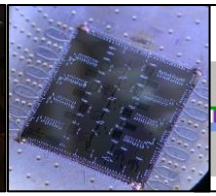

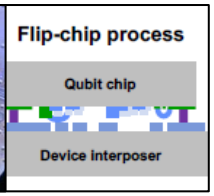

Less formal than a user facility while maintaining a transparent selection process for access: Both testbeds are open to external collaborators, subject to merit review. https://qscout.sandia.gov, https://aqt.lbl.gov

Increased Fidelity by Decomposition to AQT Native Gates

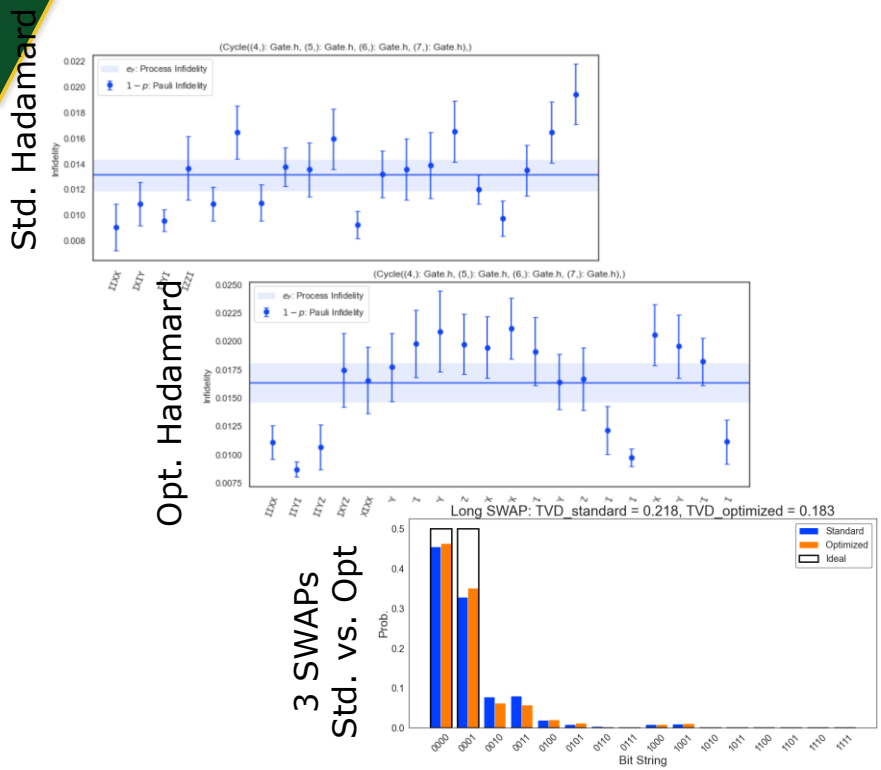

Top and middle: optimized Hadamard gate reduces infidelity from 1.3% to 1.15% → a 2x improvement since reference (idle) infidelity is 1%
Bottom: the error of a chain of 3 opt. SWAPs is reduced by 20%

Scientific Achievement

Improvements in gate speeds and fidelity by optimizing the decompositions of quantum circuits to Berkeley Advanced Quantum Testbed's native gates (CZ + X90 + Virtual-RZ).

Significance and Impact

2x reduction in latency and 2x error reduction for a Hadamard gate. 15% reduction in runtime of a SWAP gate.

Research Details

- Instead of decomposing a Hadamard gate into two physical X90 gates plus one virtual RZ gate, we decompose into only one physical X90 gate plus two virtual RZ gates. This doubles speed and halves error.
- We also introduce an optimized SWAP that both cancels out unnecessary gates in the prevailing decomposition and leverages our optimized SWAP.
- For a chain of three SWAPs, we observe a 20% reduction in infidelity.