Advanced Scientific Computing Research (ASCR): Applied Mathematics

ASCR Applied Mathematics Team:

Steven Lee & William Spotz

https://science.osti.gov/ascr/officehours

Office of Science Statement of Commitment & other Guidance

- **SC Statement of Commitment** – SC is fully and unconditionally committed to fostering safe, diverse, equitable, inclusive, and accessible work, research, and funding environments that value mutual respect and personal integrity. https://science.osti.gov/SW-DEI/SC-Statement-of-Commitment
- **Expectations for Professional Behaviors** –SC's expectations of all participants to positively contribute to a professional, inclusive meeting that fosters a safe and welcoming environment for conducting scientific business, as well as outlines behaviors that are unacceptable and potential ramifications for unprofessional behavior. https://science.osti.gov/SW-DEI/DOE-Diversity-Equity-and-Inclusion-Policies/Harassment
- **How to Address or Report Behaviors of Concern**– Process on how and who to report issues, including the distinction between reporting on unprofessional, disrespectful, or disruptive behaviors, and behaviors that constitute a violation of Federal civil rights statutes. https://science.osti.gov/SW-DEI/DOE-Diversity-Equity-and-Inclusion-Policies/How-to-Report-a-Complaint
- **Implicit Bias** – Be aware of implicit bias, understand its nature – everyone has them - and implicit bias if not mitigated can negatively impact the quality and inclusiveness of scientific discussions that contribute to a successful meeting. https://kirwaninstitute.osu.edu/article/understanding-implicit-bias

U.S. DEPARTMENT OF ENERGY | Office of Science

Our Mission:

Deliver scientific discoveries and major scientific tools to transform our understanding of nature and advance the energy, economic, and national security of the United States.

More than **34,000 r**esearchers supported at more than **300** institutions and **17** DOE national laboratories

Steward **10** of the 17 DOE national laboratories

More than **37,000** users of **28** Office of Science scientific user facilities

$8.1B
(FY 23 enacted)

U.S. Department of Energy Office of Science User Facilities

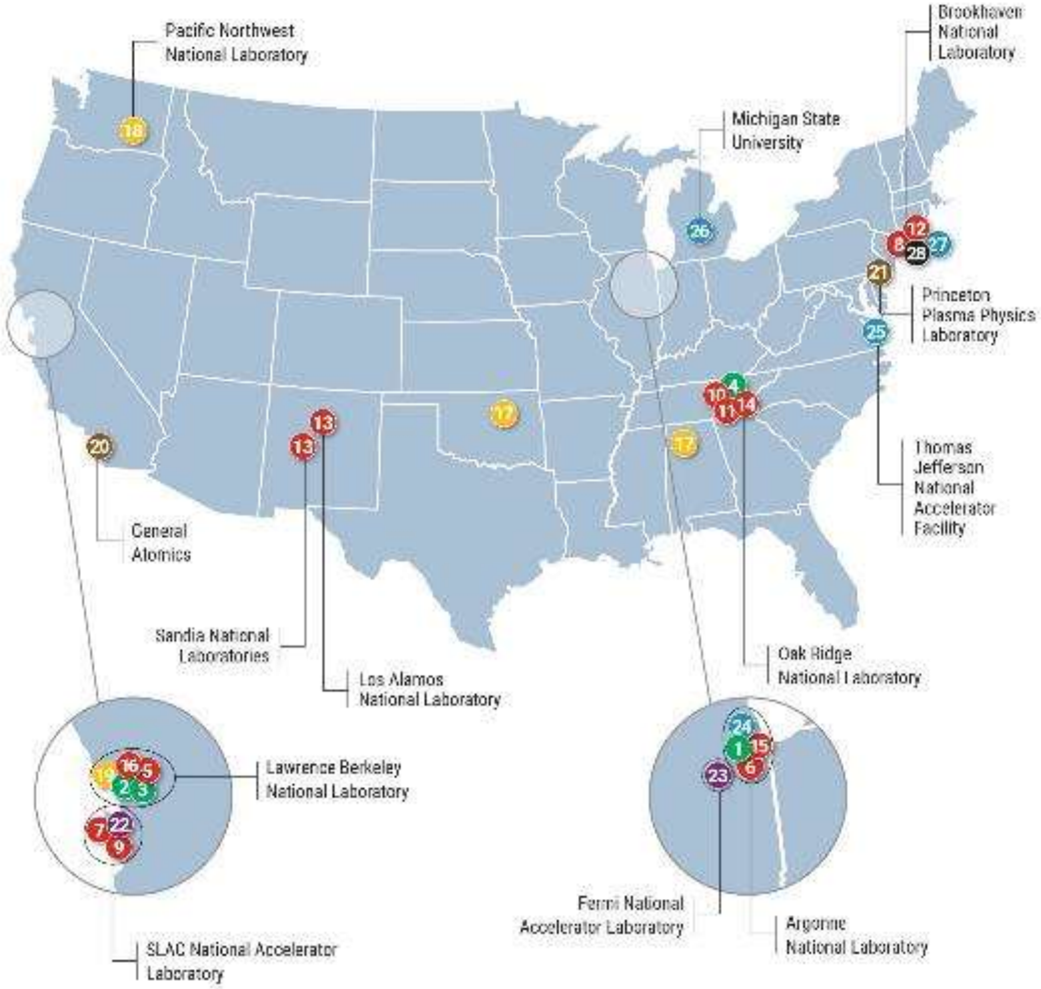

Advanced Scientific Computing Research (ASCR)

1. **Argonne Leadership Computing Facility (ALCF)**
   Argonne National Laboratory
2. **Energy Sciences Network (ESnet)**
   Lawrence Berkeley National Laboratory
3. **National Energy Research Scientific Computing Center (NERSC)**
   Lawrence Berkeley National Laboratory
4. **Oak Ridge Leadership Computing Facility (OLCF)**
   Oak Ridge National Laboratory

Basic Energy Sciences (BES)

LIGHT SOURCES

5. **Advanced Light Source (ALS)**
   Lawrence Berkeley National Laboratory
6. **Advanced Photon Source (APS)**
   Argonne National Laboratory
7. **Linac Coherent Light Source (LCLS)**
   SLAC National Accelerator Laboratory
8. **National Synchrotron Light Source II (NSLS-II)**
   Brookhaven National Laboratory
9. **Stanford Synchrotron Radiation Lightsource (SSRL)**
   SLAC National Accelerator Laboratory

NEUTRON SOURCES

10. **High Flux Isotope Reactor (HFIR)**
    Oak Ridge National Laboratory
11. **Spallation Neutron Source (SNS)**
    Oak Ridge National Laboratory

NANOSCALE SCIENCE RESEARCH CENTERS

12. **Center for Functional Nanomaterials (CFN)**
    Brookhaven National Laboratory
13. **Center for Integrated Nanotechnologies (CINT)**
    Sandia National Laboratories and Los Alamos National Laboratory
14. **Center for Nanophase Materials Sciences (CNMS)**
    Oak Ridge National Laboratory
15. **Center for Nanoscale Materials (CNM)**
    Argonne National Laboratory
16. **The Molecular Foundry (TMF)**
    Lawrence Berkeley National Laboratory

Biological and Environmental Research (BER)

17. **Atmospheric Radiation Measurement (ARM) User Facility**
    Fixed and Mobile Sites Across the Globe
18. **Environmental Molecular Sciences Laboratory (EMSL)**
    Pacific Northwest National Laboratory
19. **Joint Genome Institute (JGI)**
    Lawrence Berkeley National Laboratory

Fusion Energy Sciences (FES)

20. **DIII-D National Fusion Facility**
    General Atomics
21. **National Spherical Torus Experiment Upgrade (NSTX-U)**
    Princeton Plasma Physics Laboratory

High Energy Physics (HEP)

22. **Facility for Advanced Accelerator Experimental Tests (FACET)**
    SLAC National Accelerator Laboratory
23. **Fermilab Accelerator Complex**
    Fermi National Accelerator Laboratory

Nuclear Physics (NP)

24. **Argonne Tandem Linac Accelerator System (ATLAS)**
    Argonne National Laboratory
25. **Continuous Electron Beam Accelerator Facility (CEBAF)**
    Thomas Jefferson National Accelerator Facility
26. **Facility for Rare Isotope Beams (FRIB)**
    Michigan State University
27. **Relativistic Heavy Ion Collider (RHIC)**
    Brookhaven National Laboratory

Accelerator R&D and Production (ARDAP)

28. **Accelerator Test Facility (ATF)**
    Brookhaven National Laboratory

OFFICE OF SCIENCE BY THE NUMBERS

Delivering scientific discoveries and major scientific tools to transform our understanding of nature and advance the energy, economic, and national security of the United States

FY23

6 CORE SCIENCE PROGRAMS

- Advanced Scientific Computing Research
- Basic Energy Sciences
- Biological and Environmental Research
- Fusion Energy Sciences
- High Energy Physics
- Nuclear Physics

3 ENGINEERING AND TECHNOLOGY OFFICES

- Accelerator Research and Development and Production
- Isotope Research and Development and Production
- Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR)

5 NATIONAL QUANTUM INFORMATION SCIENCE RESEARCH CENTERS

ACROSS ITS 10 NATIONAL LABS, OFFICE OF SCIENCE MAINTAINS APPROXIMATELY

24 MILLION SQUARE FEET OF SPACE

1,600 BUILDINGS | **38,000** ACRES OF LAND OWNED

SUPPORTS RESEARCH SPANNING

16 DOE NATIONAL LABS

50 STATES, GUAM, PUERTO RICO, AND WASHINGTON, D.C.

>310 UNIVERSITIES AND HIGHER-LEARNING INSTITUTIONS

4 BIOENERGY RESEARCH CENTERS

2 ENERGY INNOVATION HUB PROGRAMS

51 ENERGY FRONTIER RESEARCH CENTERS

STEWARDS **10** DOE NATIONAL LABORATORIES

3 World-Leading Supercomputers

ESTIMATED RESEARCHERS SUPPORTED

- **11,100** Permanent PhDs
- **3,400** Postdoctoral Associates
- **5,200** Graduate Students
- **9,700** Other Scientific Personnel

OVER **39,500** USERS AT **28** OFFICE OF SCIENCE FACILITIES

10 SITE OFFICES

1 CONSOLIDATED SERVICE CENTER

OVER **100** NOBEL PRIZES

$8.1 BILLION OVERALL OFFICE OF SCIENCE BUDGET

$918 MILLION USER FACILITY CONSTRUCTION

$281 MILLION SCIENCE LABORATORIES INFRASTRUCTURE

The Office of Science Research Portfolio

★ Advanced Scientific Computing Research	• Delivering world leading computational and networking capabilities to extend the frontiers of science and technology
Basic Energy Sciences	• Understanding, predicting, and ultimately controlling matter and energy flow at the electronic, atomic, and molecular levels
Biological and Environmental Research	• Understanding complex biological, earth, and environmental systems
Fusion Energy Sciences	• Supporting the development of a fusion energy source and supporting research in plasma science
High Energy Physics	• Understanding how the universe works at its most fundamental level
Nuclear Physics	• Discovering, exploring, and understanding all forms of nuclear matter
Isotope R&D and Production	• Supporting isotope research, development, production, processing and distribution to meet the needs of the Nation
Accelerator R&D and Production	• Supporting new technologies for use in SC's scientific facilities and in commercial products

ASCR – over 70 years of Advancing Computational Science

Beginnings: During the Manhattan Project, John Von Neumann advocated for the creation of a Mathematics program to support the continued development of applications of digital computing

Over 40+ years, ASCR has a rich history of investment in computational science and applied mathematics research, and revolutionary computational and network infrastructure.

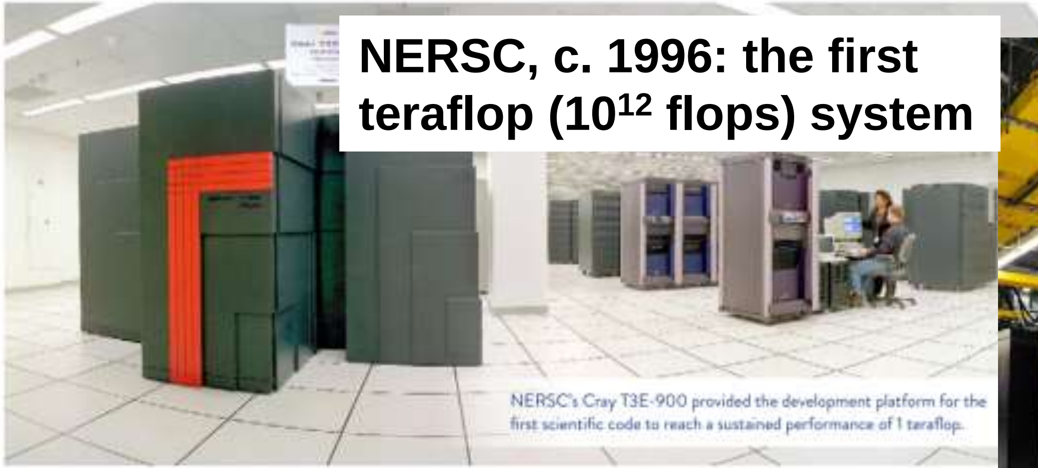

NERSC, c. 1996: the first teraflop (10^{12} flops) system

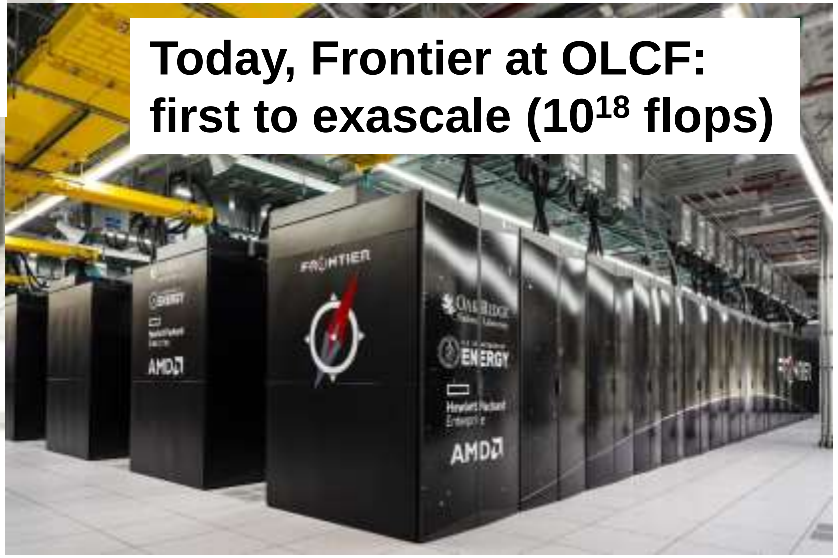

Today, Frontier at OLCF: first to exascale (10^{18} flops)

WHY COMPUTATIONAL SCIENCE?

- Computational science adds a third pillar to researcher's toolkit along side theory and experiments
- Computational science is essential when experiments are too expensive, dangerous, time-consuming or impossible
- Computational science facilitates idea-to-discovery that leads from equations to algorithms
- Virtually every discipline in science and engineering has benefited from DOE's sustained investments in computational science

Scientific Data at Extreme Scale

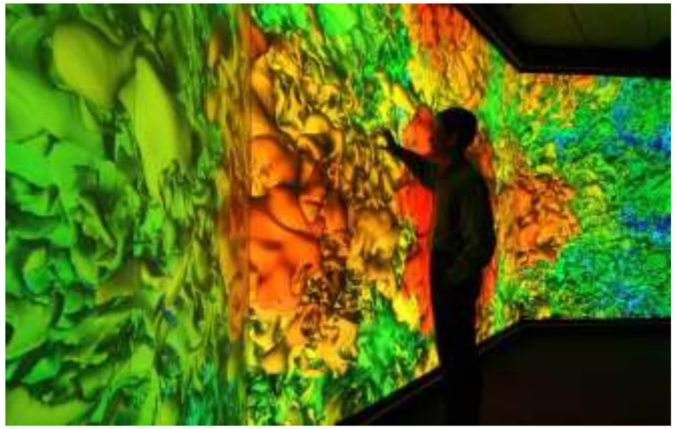

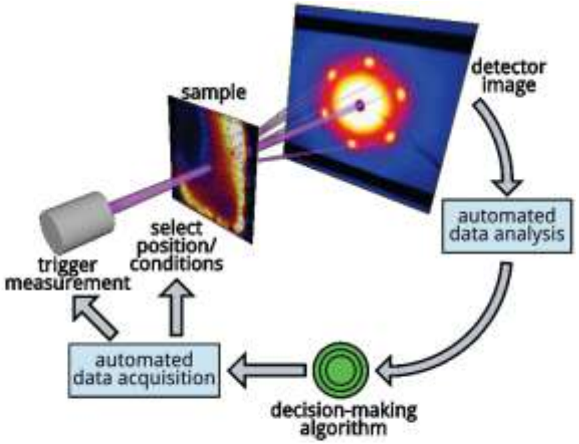

- Scientific computations and experiments produce terabytes or petabytes of data that must be efficiently stored.
- That data is stored on collections of disk drives and archive systems at ASCR computing facilities.
- As with ASCR's computing capabilities, high-performance data management requires performing many operations in parallel.
- ASCR invests in innovative ways to store, compress, search, and analyze data that maximizes parallelism and performance.
- ASCR also invests in advancements in streaming data and federated learning, allowing data in geographically-separated places to contribute to scientific modeling without needing to store all of the data in once place.

Moore’s Law

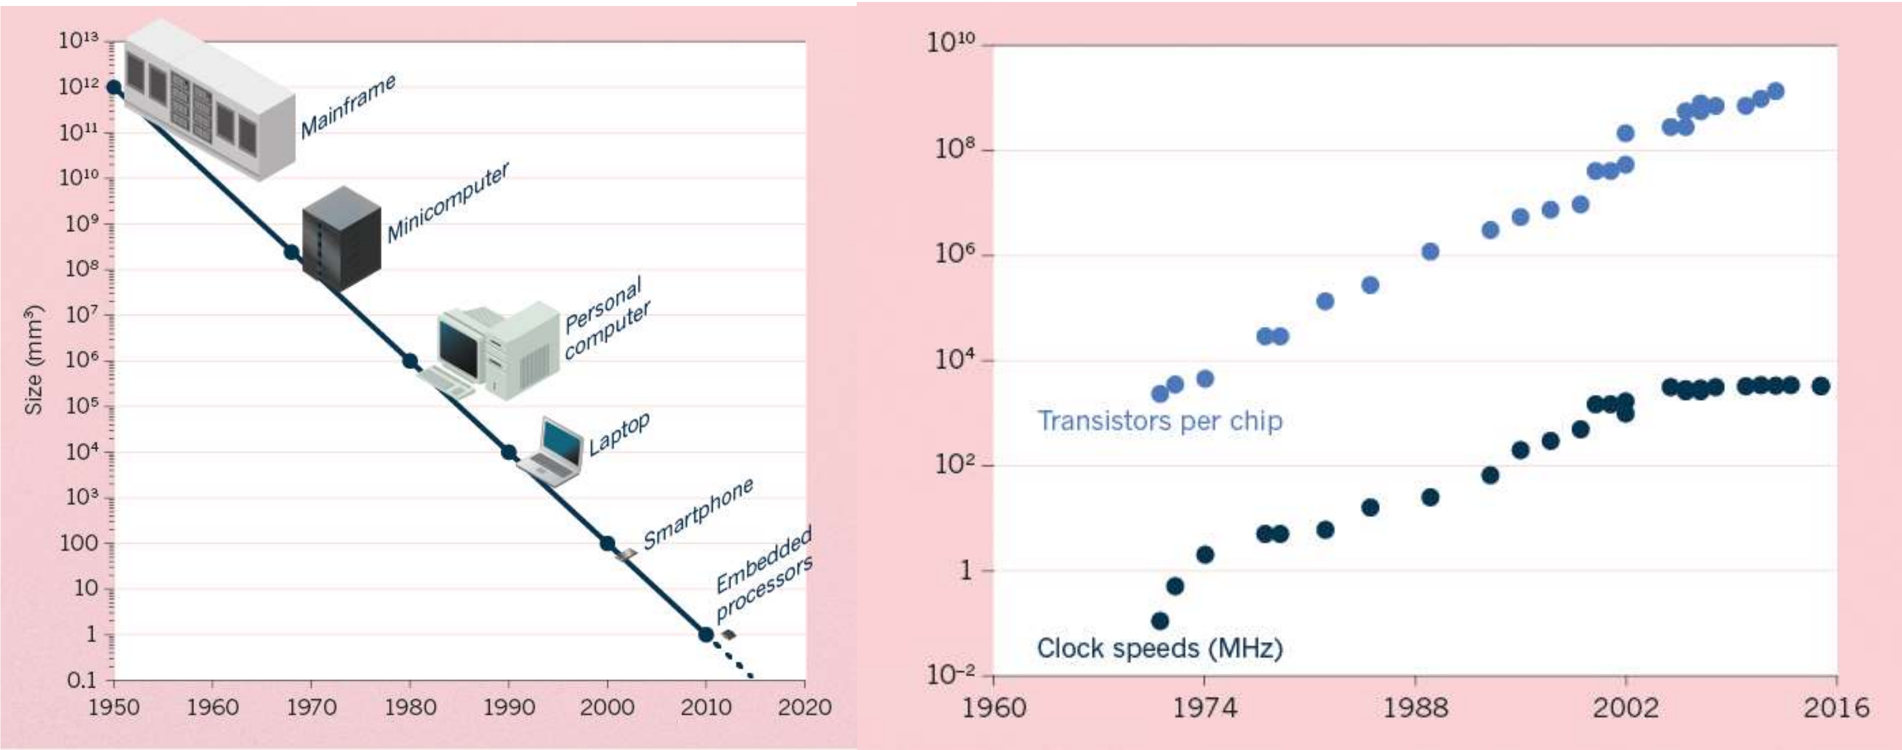

https://www.nature.com/news/the-chips-are-down-for-moore-s-law-1.19338

- Moore's law is the observation that the number of transistors in an integrated circuit (IC) doubles about every two years.
- As Moore’s law has continued computers have continued to shrink *and* become more capable.
- However, the clock speed of energy-efficient computers stopped increasing some time ago – this is why parallel computing, doing more simultaneously, is critical to modern computing including ASCR’s supercomputers.

Quantum Computer Simulation of Physical Systems

Richard P. Feynman
Simulating Physics with Computers,
Int. J. Theor. Phys. (1982)

Power of Quantum Superposition

Classical	Quantum
• Classical bit can represent 0 or 1 • Eight classical bits can represent one of 256 integers (2^8) • N classical bits can represent one of 2^N integers	• Quantum bit can be in a superposition of 0 and 1 • Eight quantum bits can represent all of 256 integers (2^8) • N quantum bits can represent all of 2^N integers

- Qubit states are fragile.
- Wiring qubits together into a functional architecture is hard.

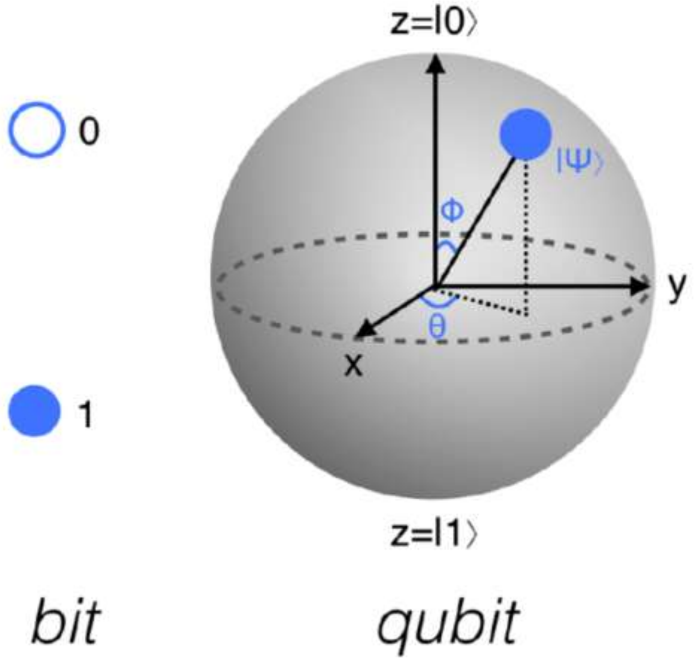

Emani, P.S., Warrell, J., Anticevic, A. et al. Quantum computing at the frontiers of biological sciences, Nat. Methods (2021)

Emerging Technology Trends for Scientific Computing

Advanced Modeling, Simulation, and Visualization

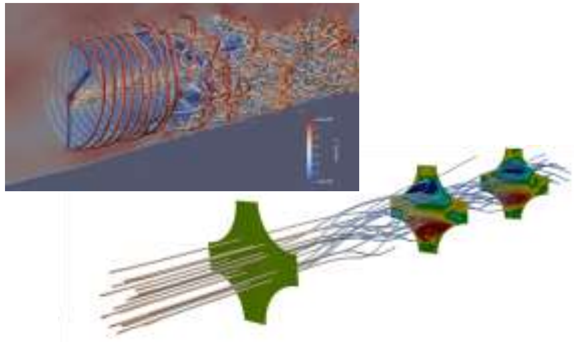

Trustworthy Artificial Intelligence and Data

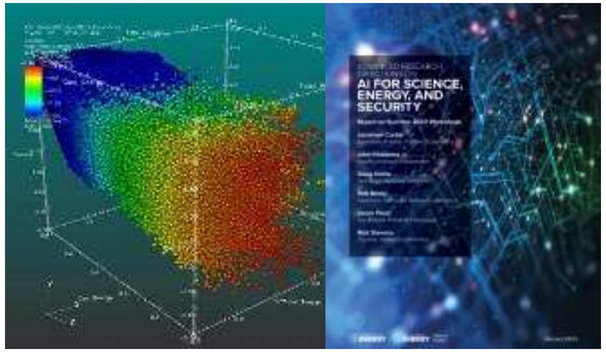

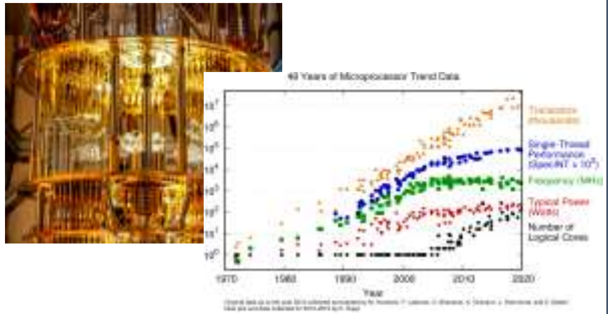

Heterogeneous, Distributed, Co-Designed, Energy-Efficient Computing and Algorithms

Software Complexity for Increased Versatility

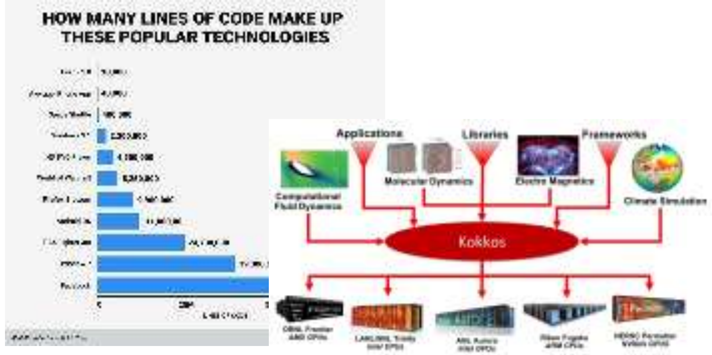

High-Performance Computing and Networking across Experiments, Exascale, and the Edge

ASCR R&D Funding (**)

Funding Opportunity Announcements (FOAs)

- https://science.osti.gov/ascr/Funding-Opportunities
- Announced on grants.gov (hint: sign up for email notifications for ‘ASCR’)
- Read each announcement carefully to understand who can apply and other restrictions/requirements
- Depending on the announcement, supports 2–5-year projects
- University researchers can apply directly (please coordinate with your organization’s sponsored-research office)
- Subcontracting is often permitted, and sometimes collaborative applications are permitted

Early Career Research Program

- https://science.osti.gov/early-career
- Research grants for five years
- Stays with PI if PI changes institutions
- Eligible within 10 years of Ph.D. (can apply up to three times)
- University-based researchers receive about $175,000/year
- Topics released in the summer, pre-applications generally due in the fall

DOE National Laboratory Announcements

- https://science.osti.gov/ascr/Funding-Opportunities (bottom of the page)
- Open only to DOE Laboratories
- Often allow subcontracts to support collaborators at other organizations

Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR)

- https://science.osti.gov/sbir
- Grants to for-profit US businesses with 500 or fewer employees (including affiliates)
- Phase I: ~$200k for 6-12 months, Phase II: ~$1M for 2 years
- Subcontracting is permitted, STTR: requires collaboration with a research Institution
- Topics released in the summer, pre-applications generally due in the fall

Computational Science Graduate Fellowship (CSGF)

http://www.krellinst.org/csgf/

Transforming the Fundamentals of Computing

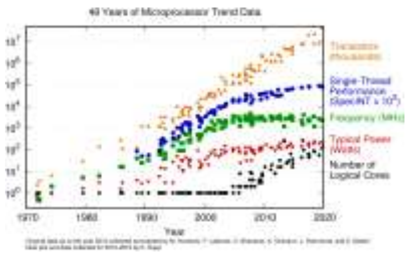

Heterogeneous, Distributed, Co-Designed, Energy-Efficient Computing and Algorithms

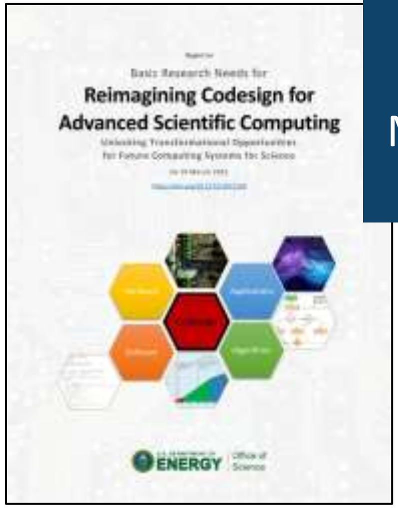

ASCR Workshop on Reimagining Codesign, March 2021: https://doi.org/10.2172/1822199

Quantum Computing for Biomedical Computational and Data Sciences Joint DOE NIH Quantum Roundtable March 2023: https://doi.org/10.2172/2228574

ASCR Basic Research Needs in Quantum Computing and Networking, July 2023: https://doi.org/10.2172/2001045

Transforming the Fundamentals of Computing

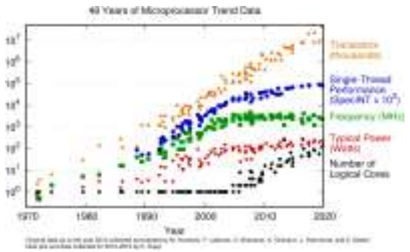

Heterogeneous, Distributed,
Co-Designed, Energy-Efficient
Computing

Past Solicitations (FY 2021 – 2022):

- Entanglement Management and Control in Transparent Optical Quantum Networks, 2021.
- Microelectronics Co-Design Research, 2021.
- Quantum Internet to Accelerate Scientific Discovery, 2021.
- Quantum Algorithms and Mathematical Methods, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2022.
- Quantum Computing at the Edge, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2022.

Transforming the Fundamentals of Computing

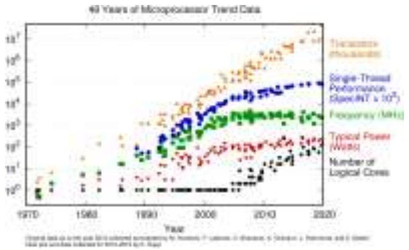

Heterogeneous, Distributed, Co-Designed, Energy-Efficient Computing

Past Solicitations (FY 2023):

- Accelerate Innovations in Emerging Technologies, 2023.
- Modeling Future Supercomputing Systems, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2023.
- Programming Techniques for Computational Physical Systems, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2023.
- Quantum Algorithms across Models, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2023.
- Quantum Testbed Pathfinder, 2023.
- Scientific Enablers of Scalable Quantum Communications, 2023.

Transforming the Fundamentals of Computing

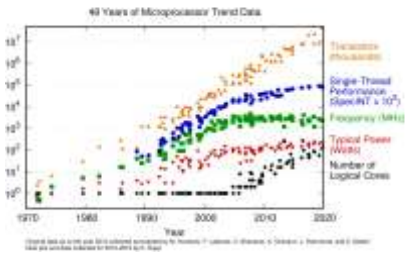

Heterogeneous, Distributed, Co-Designed, Energy-Efficient Computing

Solicitations (FY 2024, to date):

- Accelerated Research in Quantum Computing, 2024
- Neuromorphic Computing, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2024.
- Quantum Hardware Emulation, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2024.
- The Co-Design of Energy-Efficient AI Algorithms and Hardware Architectures, in Advancements in Artificial Intelligence for Science, 2024.

Empowering Science Through Data Innovations

Data, Privacy, and Scientific Integrity

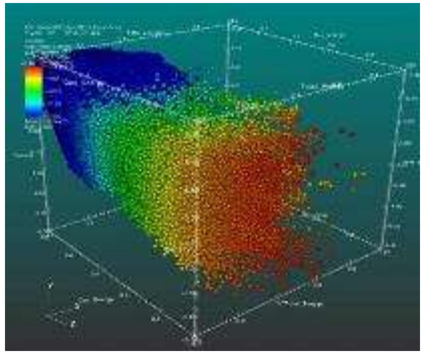

ASCR Workshop on Basic Research Needs for Management and Storage of Scientific Data, January 2022: https://doi.org/10.2172/1845707


Data Reduction for Science: Brochure from the Advanced Scientific Computing Research Workshop

Introduction

Data Reduction for Science, January 2021: https://doi.org/10.2172/1770192

ASCR Basic Research Needs Visualization for Scientific Discovery, Decision-Making, and Communication, January 2022: https://doi.org/10.2172/1845708 (brochure; report forthcoming)

Empowering Science Through Data Innovations

Data, Privacy, and Scientific Integrity

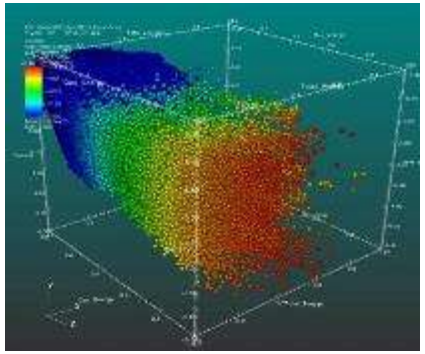

Past Solicitations (FY 2021 – 2022):

- Data Reduction for Science, 2021.
- Management and Storage of Scientific Data, 2022.
- Data Visualization for Scientific Discovery, Decision-Making, and Communication, 2022.

Solicitations (FY 2024, to date):

- Federated and Privacy-Preserving Machine Learning and Synthetic Data Creation, in Advancements in Artificial Intelligence for Science, 2024.
- Harnessing Technology Innovations to Accelerate Science through Visualization, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2024.
- Data Reduction for Science, 2024.

ECP Exascale Computing Project

Many projects are building on ECP investments!

Enhancing Scientific Programming

Exploding Software Complexity

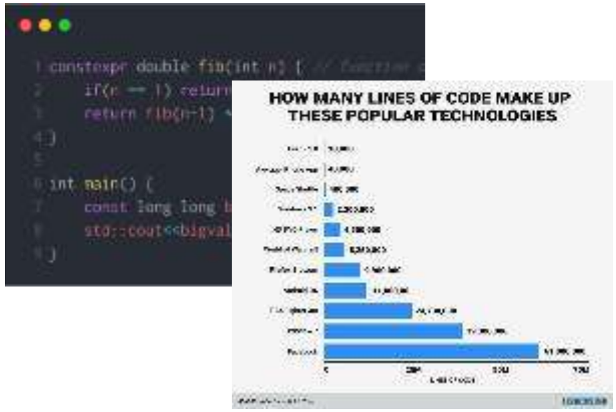

ASCR Workshop on Basic Research Needs in The Science of Scientific Software Development and Use, December 2021: https://doi.org/10.2172/1846009

DOE/NSF Workshop on Correctness in Scientific Computing, June 2023: https://arxiv.org/abs/2312.15640

Enhancing Scientific Programming

Exploding Software Complexity

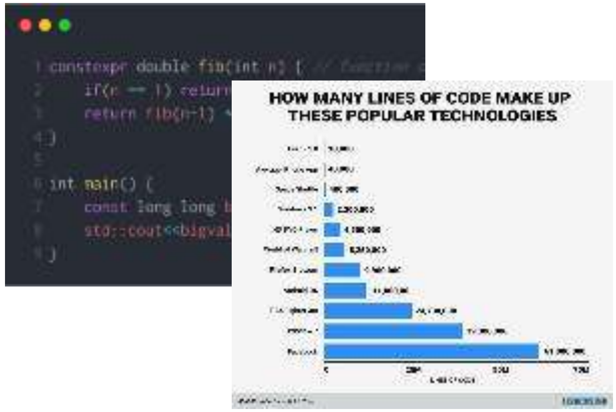

Past Solicitations (FY 2021 – 2022):

- X-STACK: Programming Environments for Scientific Computing, 2021.
- Differentiable Programming, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2022.

Solicitations (FY 2024, to date):

- AI Innovations for Scientific Knowledge Synthesis and Software Development, in Advancements in Artificial Intelligence for Science, 2024.

ECP EXASCALE COMPUTING PROJECT

Many projects are building on ECP investments!

Accelerating Science from Exascale to the Edge

Scientific Computing and Networking: from Exascale to the Edge

Integrated Research Infrastructure Architecture Blueprint Activity, 2023: https://doi.org/10.2172/1984466

Roundtable on Foundational Science for Biopreparedness and Response, March 2022: Report available from https://science.osti.gov/ascr/Community-Resources/Program-Documents

Accelerating Science from Exascale to the Edge

Scientific Computing and Networking: from Exascale to the Edge

Past Solicitations (FY 2021 – 2022):

- 5G Enabled Energy Innovation Advanced Wireless Networks for Science, 2021.
- SciDAC: Partnerships in Basic Energy Sciences, 2021.
- Integrated Computational and Data Infrastructure for Scientific Discovery, 2021.
- SciDAC: Partnerships in Earth System Model Development, 2022.
- SciDAC: Partnership in Nuclear Energy, 2022.

ECP EXASCALE COMPUTING PROJECT

Many projects are building on ECP investments!

Accelerating Science from Exascale to the Edge

Scientific Computing and Networking: from Exascale to the Edge

- SciDAC: Partnership in Nuclear Physics, 2022.

Past Solicitations (FY 2021 – 2022):

- SciDAC: High Energy Physics, 2022.
- Advancing Computer Modeling and Epidemiology for Biopreparedness and Response, 2022.

ECP Exascale Computing Project

Many projects are building on ECP investments!

Accelerating Science from Exascale to the Edge

Scientific Computing and Networking: from Exascale to the Edge

Past Solicitations (FY 2023):

- Energy Earthshot Research Centers, 2023.
- Science Foundations for Energy Earthshots, 2023.
- Biopreparedness Research Virtual Environment (BRaVE), 2023.
- Advanced Scientific Computing Research for DOE User Facilities, 2023.
- Accelerate Innovations in Emerging Technologies, 2023.
- Distributed Resilient Systems, 2023.
- SciDAC: FES Partnerships, 2023

ECP EXASCALE COMPUTING PROJECT

Many projects are building on ECP investments!

Accelerating Science from Exascale to the Edge

Scientific Computing and Networking: from Exascale to the Edge

Solicitations (FY 2024, to date):

- Advanced Wireless, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2024.

ECP EXASCALE COMPUTING PROJECT

Many projects are building on ECP investments!

Innovating in Algorithms and Mathematics

Advanced Modeling, Simulation, and Visualization

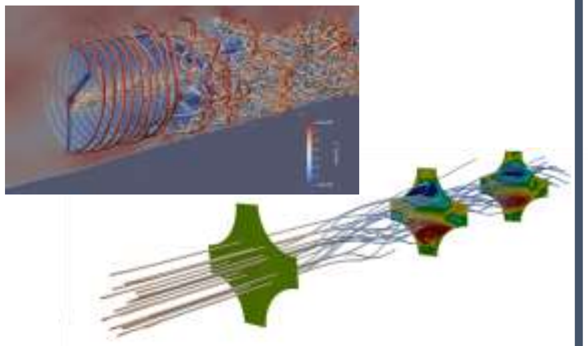

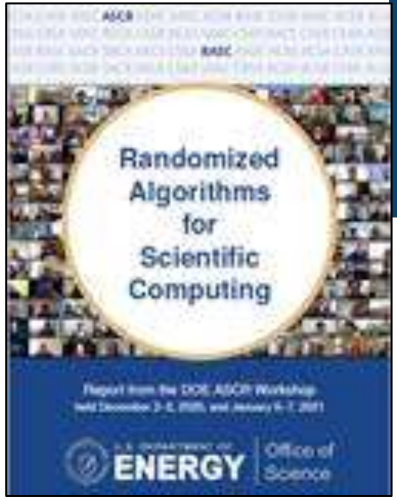

ASCR Workshop on Randomized Algorithms for Scientific Computing, January 2021: https://doi.org/10.2172/1807223

Roundtable on Computer Science Research Needs for Parallel Discrete Event Simulation, 2022: https://doi.org/10.2172/1855247

Innovating in Algorithms and Mathematics

Advanced Modeling, Simulation, and Visualization

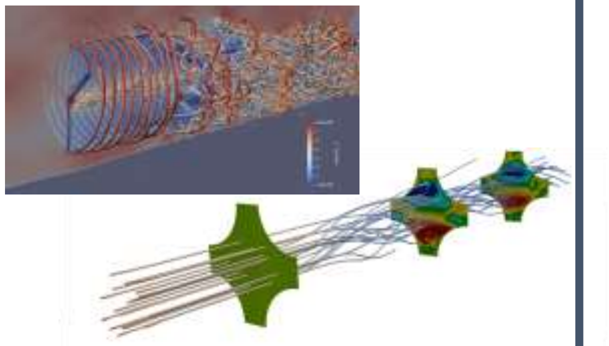

Past Solicitations (FY 2021 – 2022):

- EXPRESS: Randomized Algorithms for Extreme-Scale Science, 2021.
- Mathematical Multifaceted Integrated Capability Centers (MMICCS), 2022.
- Randomized Algorithms for Combinatorial Scientific Computing, 2022.
- Parallel Discrete-Event Simulation, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2022.

Solicitations (FY 2024, to date):

- Scalable Space-Time Memories for Large Discrete/Agent-Based Models in Exploratory Research for Extreme-Scale Science (EXPRESS), 2024.

Scientific Machine Learning: 2018 Workshop & Report

Foundations Themes – Priority Research Directions

1. **Domain-Aware**: Leverage & respects scientific domain knowledge. Physics principles, symmetries, constraints, uncertainties & structure-exploiting models.
2. **Interpretable:** Explainable and understandable results. Model selection, exploiting structure in high-dimensional data, use of uncertainty quantification with machine learning.
3. **Robust:** Stable, well-posed & reliable formulations. Probabilistic modeling in machine learning, quantifying well-posedness, reliable hyperparameter estimation.

Capabilities Themes – Priority Research Directions

4. Data-Intensive: Scientific inference & data analysis. Machine learning methods for multimodal data, in situ data analysis & optimally guide data acquisition.
5. Machine Learning-Enhanced Simulations: Hybrid algorithms & models for predictive scientific computing. ML-enabled adaptive algorithms, parameter tuning & multiscale surrogate models.
6. Intelligent Automation & Decision Support: Adaptivity, automation, resilience, control. Exploration of decision space with ML, ML-based resource management, optimal decisions for complex systems.

AI4SES Report

- AI for Science, Energy, and Security Report, released May 2023: https://www.anl.gov/ai-for-science-report

- Created by a confederation of laboratories, informed by a series of workshops held in 2022.

- Covers AI approaches:
  - AI and Surrogate Models for Scientific Computing
  - AI Foundation Models for Scientific Knowledge Discovery, Integration, and Synthesis
  - AI for Advanced Property Inference and Inverse Design
  - AI-Based Design, Prediction, and Control of Complex Engineered Systems
  - AI and Robotics for Autonomous Discovery
  - AI for Programming and Software Engineering
- Also covers crosscuts, including workflows, data, AI hardware, computing infrastructure, and workforce

Creating Trustworthy and Efficient AI For Science

Artificial Intelligence

Past Solicitations (FY 2021 – 2022):

- Bridge2AI And Privacy-Preserving Artificial Intelligence Research, 2021.
- Data-Intensive Scientific Machine Learning and Analysis, 2021.
- Federated Scientific Machine Learning, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2022.
- Explainable Artificial Intelligence, in Exploratory Research for Extreme-Scale Science (EXPRESS), 2022.

Creating Trustworthy and Efficient AI For Science

Artificial Intelligence

Past Solicitations (FY 2023):

- Scientific Machine Learning for Complex Systems, 2023.

Solicitations (FY 2024, to date):

- Advancements in Artificial Intelligence for Science, 2024.
  - Extreme-Scale Foundation Models for Computational Science.
  - AI Innovations for Scientific Knowledge Synthesis and Software Development.
  - AI Innovations for Computational Decision Support of Complex Systems.
  - Federated and Privacy-Preserving Machine Learning and Synthetic Data Creation.
  - The Co-Design of Energy-Efficient AI Algorithms and Hardware Architectures.

Building on ECP investments!

Growing and Diversifying Our Research Community

Past Solicitations (FY 2023):

- Early Career Research Program, 2023.
- FY 2023 Funding for Accelerated, Inclusive Research (FAIR), 2023.
- Reaching A New Energy Sciences Workforce (RENEW), 2023.

Solicitations (FY 2024, to date):

- Early Career Research Program, 2024.
- FY 2023 Funding for Accelerated, Inclusive Research (FAIR), 2024.
- Reaching A New Energy Sciences Workforce (RENEW), 2024.

Additional Information on ASCR's Website

https://science.osti.gov/ascr/Funding-Opportunities

https://science.osti.gov/ascr/Community-Resources/Program-Documents

Look at past opportunity announcements

Look at abstracts for current awards

Look at recent reports from ASCR-sponsored workshops. These discuss priority research directions, as identified by the research community, along with relevant background information, in various areas.

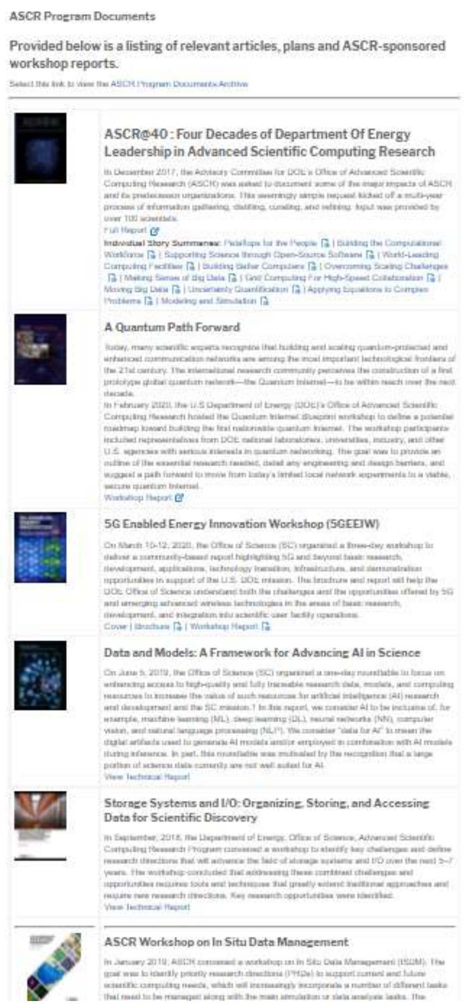

Exascale Today Enables the AI of Tomorrow

Long-term investments in applied mathematics and computer science enabled exascale.

MPICH2

R&D 100

2005

TOP500 #1

GREEN500 #2

HPL-MxP #1

Frontier, #1 on the Top500, **leads the world in computational capability**, and is also **#2 in the world in energy efficiency**, and is **#1 in the world for AI capability**.

The exascale and AI-enabled science era will lead to dramatic capabilities to predict extreme events and their impacts on the electric grid across weather and climate time scales…

and will accelerate the design and deployment of clean-energy technologies to create a better future.

Finding Out More About ASCR – ASCAC

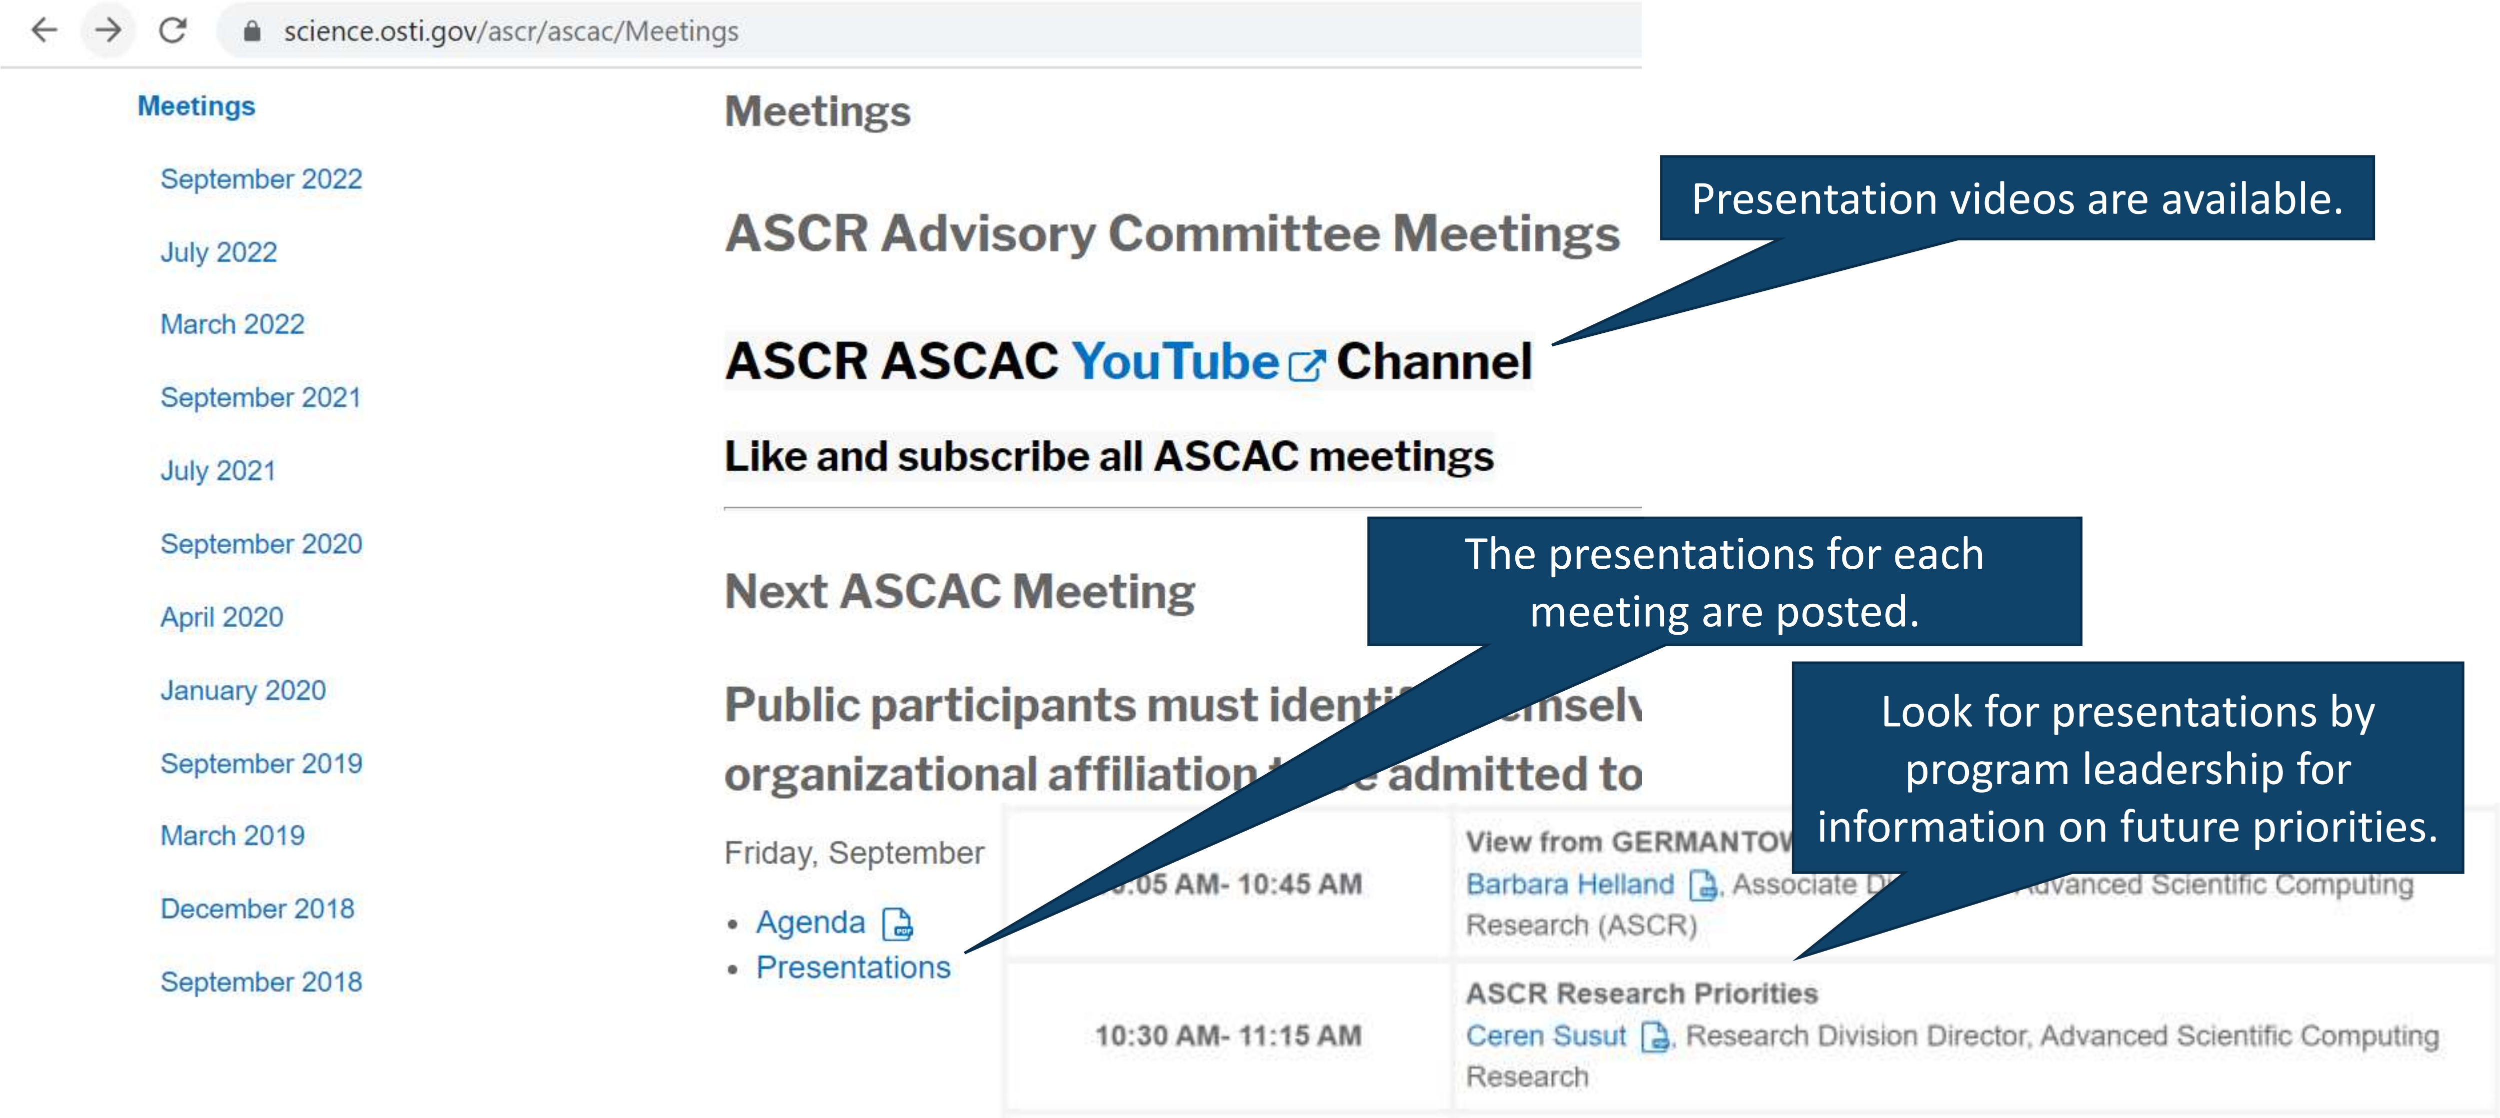

Presentation videos are available.

The presentations for each meeting are posted.

Look for presentations by program leadership for information on future priorities.

Award Lists – A New Website Location

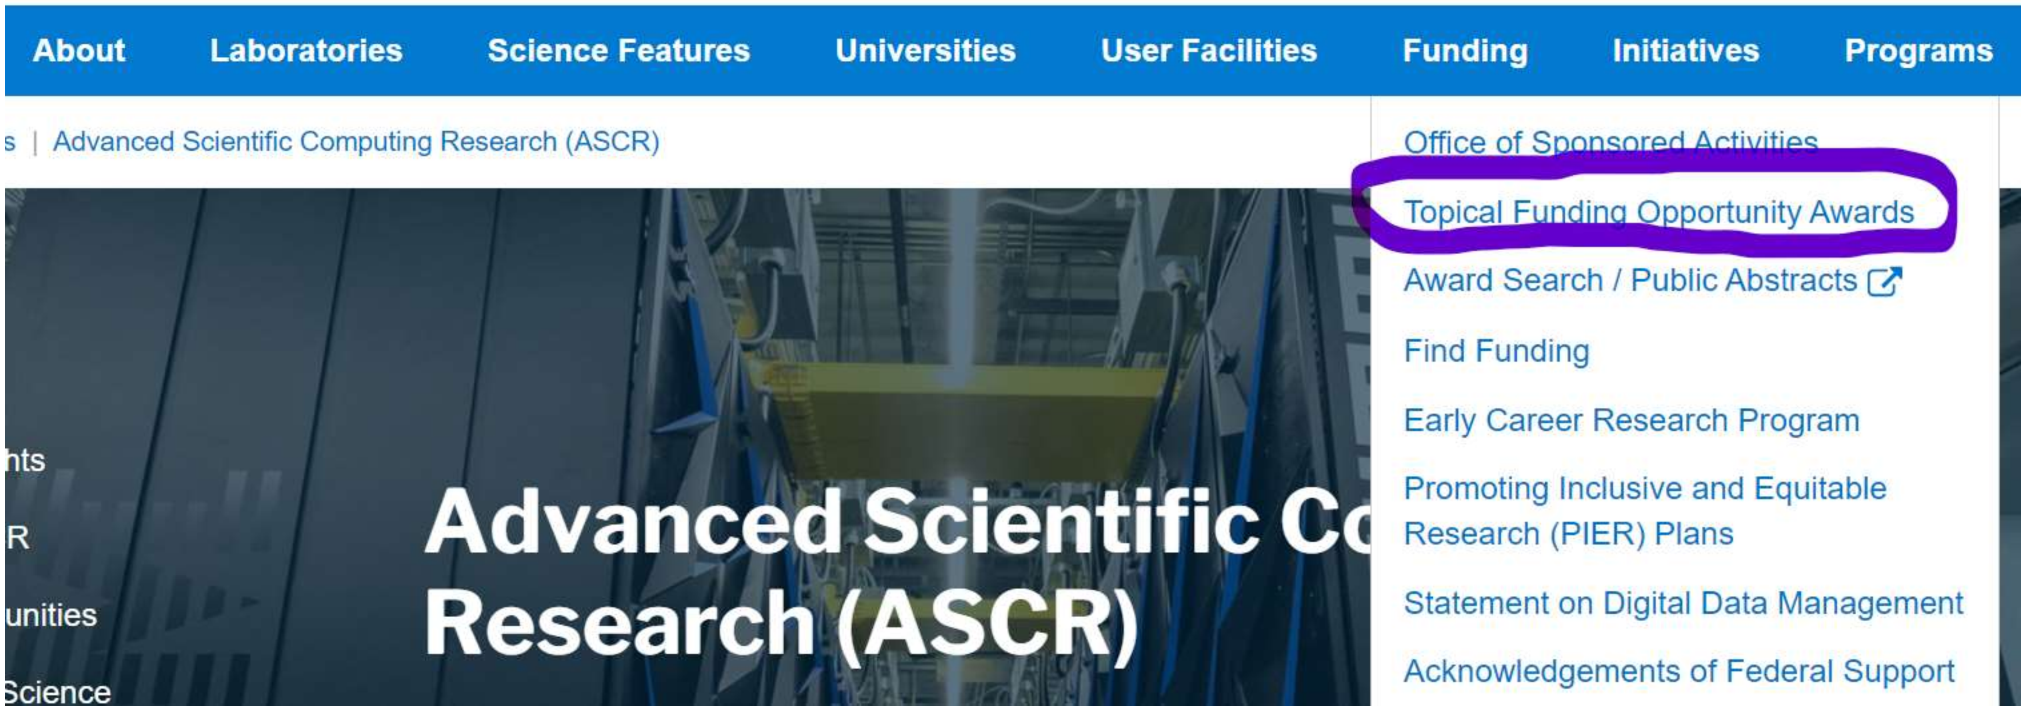

Award lists are now posted to https://science.osti.gov/Funding-Opportunities/Award along with other awards from the Office of Science. To receive award and solicitation announcements, and other ASCR-related news, signup for the Office of Science's GovDelivery email service, and check the box for the Advanced Scientific Computing Research Program in your subscriber preferences:

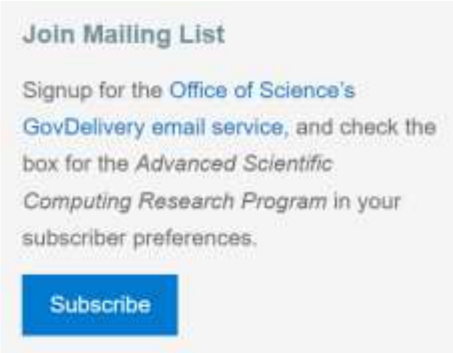

ASCR Office Hours

- Starting in March, ASCR started holding virtual office hours on the second Tuesday of the month, 2 PM ET
- Researchers, educators, and leaders within research administration from all institutional types are encouraged to join
- A primary goal of the virtual office hours is to broaden awareness of our programs; no prior history of funding from DOE is required to join
- Program managers will be available to answer questions
- Upcoming topics include:
  - *Tuesday, June 11, 2024, at 2pm ET – Introduction to ASCR's user facilities and their allocation programs*
  - *Tuesday, July 9, 2024, at 2pm ET – Overview of the ASCR research proposal and review process*

Check the ASCR website (https://science.osti.gov/ascr/) for Zoom registration links.

A Selection of Highlights and Backup Slides

Submodular Matchings for Balancing Data and Computations

A scalable parallel algorithm and a case study in Chemistry

Scalable Quantum Chemistry via Submodular Matching

- Computing electronic properties of molecules via density functional theory involves the data intensive and compute intensive Fock matrix, whose elements consist of multidimensional integrals. The computation scales as $O(n^4)$, where n is the number of basis functions.
- We provide a scalable parallel algorithm for computing the Fock matrix within the NWChemEx software from Pacific Northwest National Lab.
- The algorithm assigns blocks of Fock submatrix computations to processors in order to balance the data and work load among the processors, and also the number of messages each processor is involved in.
- This is accomplished by computing a *b*-matching in the block-processor graph, with a nonlinear (submodular) objective function, to satisfy both objectives mentioned above.
- A submodular function balances the load on the processors, whereas a linear function cannot distinguish between unbalanced and balanced task assignments.
- Although the submodular *b*-matching problem is computationally intractable, we design fast approximation algorithms that provide constant-factor approximations to the optimal matching.

Performance of NWChemEx on Summit

- We designed a submodular matching algorithm and incorporated it with the NWChemEx library.
- The code speeded up the Fock matrix computation for the ubiquitin protein molecule by a factor of four over the current task assignment.
- It also scaled the NWChemEx code to 14000 processors on Summit, from 4000 processors.
- More work could be done to reduce the size of the data even further by means of matrix factorizations.
- We collaborated with colleagues at PNNL from the ExaGraph and NWChemEX projects.

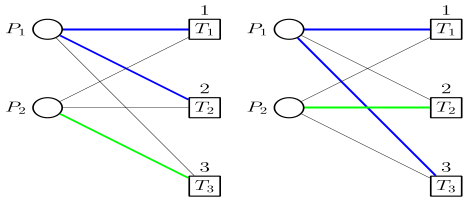

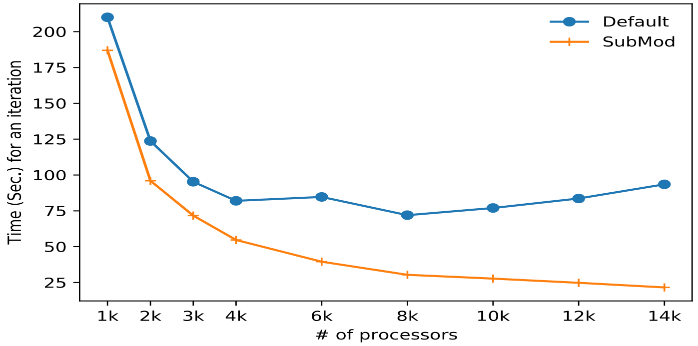

Top Fig. : A submodular matching balances the work in assigning tasks T to processors P (left), while a linear matching does not (right). Bottom Fig.: Submodular assignment balances the load in computing the energy levels of the Ubiquitin protein, reducing the time on 14K Summit processors four-fold over the default.

PI: Alex Pothen
Collaborating Institutions: Purdue University, PNNL
ASCR Program: Computer Science
ASCR PM: Hal Finkel
Publication(s) for this work: S M Ferdous et al., “A parallel approximation algorithm for submodular b-matching,” Proceedings SIAM Applied Computational Discrete Algorithms, (2021): pp. 45-56, . Doi: 10.1137/1.9781611976830.5

Giant leap toward quantum internet realized with Bell state analyzer

The Science

A multi-institutional team featuring ORNL's Joe Lukens has made strides toward a fully quantum internet by designing and demonstrating the first ever Bell state analyzer for frequency bin coding. Measuring Bell states is critical to performing many of the protocols necessary to perform quantum communication and distribute entanglement across a quantum network. The team's method represents the first Bell state analyzer developed specifically for frequency bin coding, a quantum communications method that harnesses single photons residing in two different frequencies simultaneously.

The Impact

The analyzer was designed with simulations and has experimentally demonstrated 98% fidelity for distinguishing between two distinct frequency bin Bell states. This incredible accuracy is expected to enable new fundamental communication protocols necessary for frequency bins.

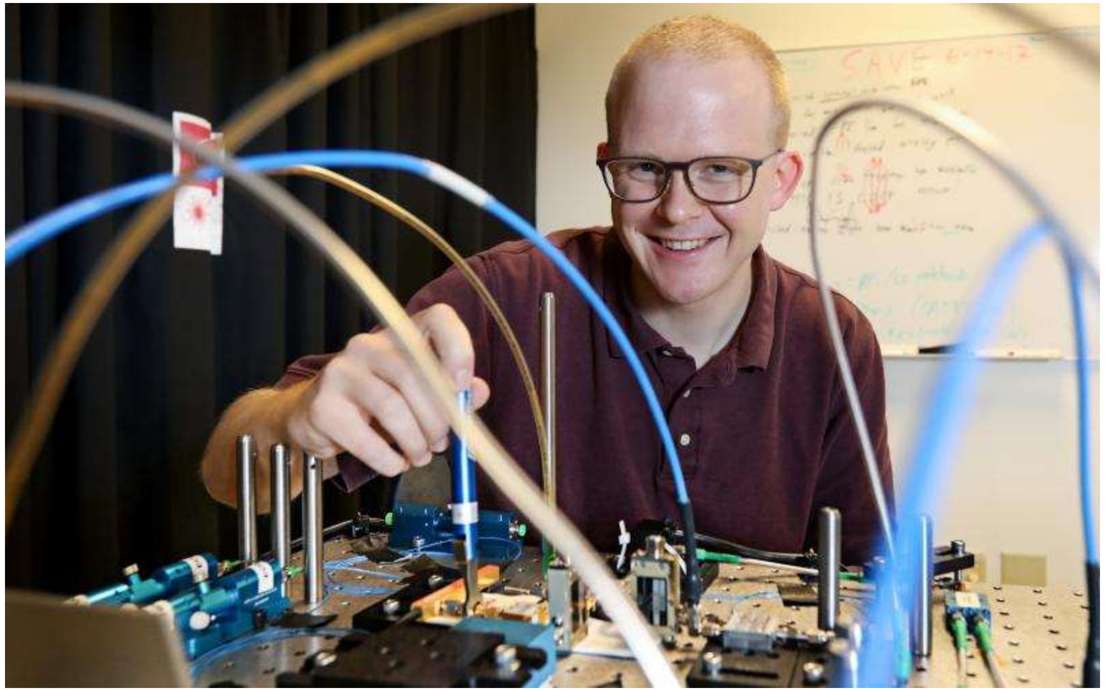

ORNL's Joseph Lukens runs experiments in an optics lab. Credit: Jason Richards/ORNL, U.S. Dept. of Energy

PI(s)/Facility Lead(s): Joe Lukens (ORNL)
ASCR Program/Facility: N/A
ASCR PM: Lali Chatterjee
Funding: Office of Science through the Early Career Research Program
Publication for this work: Navin B. Lingaraju, Hsuan-Hao Lu, Daniel E. Leaird, Steven Estrella, Joseph M. Lukens, and Andrew M. Weiner. "Bell state analyzer for spectrally distinct photons," Optica Vol. 9, Issue 3, pp. 280-283 (2022).
Date submitted to ASCR: Spring 2022

Scalable Transformers on Frontier for Real-Time Experiment Steering

Scientific Achievement

ORNL developed a scalable transformer on OLCF Frontier for real-time decision-making in neutron diffraction experiments at the TOPAZ beamline of SNS. This work:

- Develops a stochastic process model for the time-of-flight neutron scattering data and exploits a temporal fusion transformer to **help reduce the experiment time**.
- Demonstrates outstanding scalability of the ML model on Frontier, which is necessary to synchronize neutron diffraction experiments, data analysis, and decision making.

Significance and Impact

- The ML algorithm could help neutron scientists to **reduce the over-counting beamtime by around 30%** at TOPAZ, while achieving the similar data quality.
- This effort proves the concept of connecting BES's neutron facilities and ASCR's HPC facilities through AI/ML, **forming an integrated research infrastructure**.

Technical Approach

- The developed stochastic process model provides a novel and effective approach to describe the time-of-flight neutron scattering data.
- The hierarchical parallelization approach effectively uses ~60% of Frontier's computing power to keep up with the neutron experiment speed.

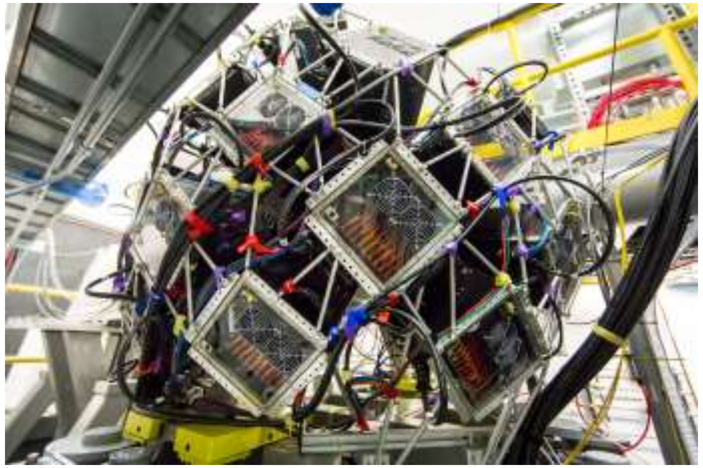

A single-crystal diffractometer on the TOPAZ beamline at SNS

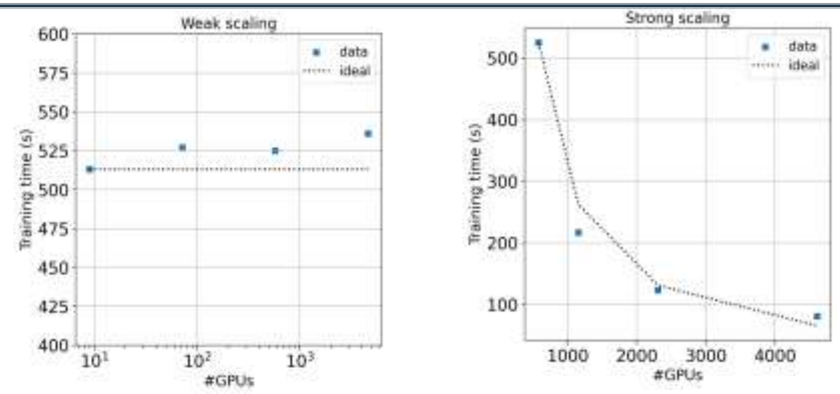

Outstanding weak and strong scalability on Frontier with up to 4608 GPUs.

Oak Ridge National Laboratory

PI : Guannan Zhang (ORNL); ASCR Program: Data-Intensive Scientific Machine Learning and Analysis; ASCR PM: Steve Lee
Publication: J. Yin, S. Liu, V. Reshniak, X. Wang, and G. Zhang, *A scalable transformer model for real-time decision making in neutron scattering experiments*, ***Journal of Machine Learning for Modelling and Computing***, Vol 4 (1), pp. 95-107, 2023

SuperNeuro: An Accelerated Neuromorphic Computing Simulator

Scientific Achievement

ORNL scientists have developed SuperNeuro, the world's fastest simulator for neuromorphic computing. It was designed for speed and scalability, and is capable of running **300 times faster** than its competitors, garnering the team the **2023 R&D 100 Award in the Software/Services Category**.

Significance and Impact

Neuromorphic architectures have the potential to increase computing power and efficiency, as well as advance AI applications. SuperNeuro provides an indispensable capability for this effort via the leveraging of GPU computing to provide superior performance for neuroscience, increased adaptability, spiking neural networks (SNNs), and general-purpose computing workloads.

Technical Approach

Two novel approaches used: matrix computation (MAT) and agent-based modeling (ABM).

- MAT Mode: Homogeneous simulations, built-in learning, CPU execution
- ABM Mode: Heterogeneous simulations, GPU acceleration

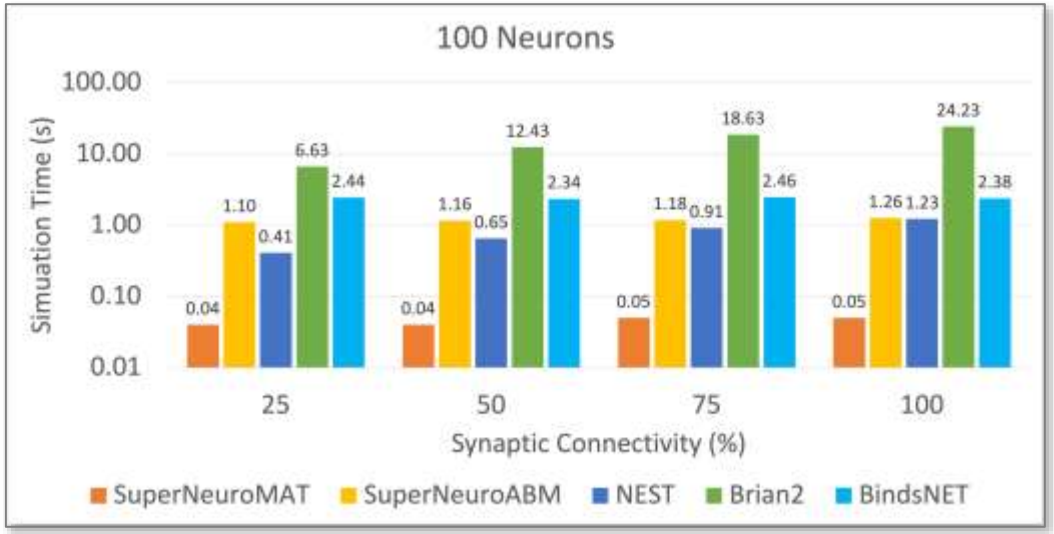

Simulating 100 neurons on 5 neuromorphic simulators with 4 different synaptic connectivites. SuperNeuroMAT performs 300 times faster than other neuromorphic simulators.

https://github.com/ORNL/superneuromat

PI(s): Prasanna Date, Chathika Gunaratne, Shruti Kulkarni, Robert Patton, Mark Coletti, and Thomas Potok
Collaborating Institutions: Oak Ridge National Laboratory
ASCR Program: Neuromorphic Computing for Accelerating Scientific Discovery
ASCR PM: Robinson Pino
Publication(s) for this work: Date, Prasanna, Chathika Gunaratne, Shruti R. Kulkarni, Robert Patton, Mark Coletti, and Thomas Potok. "SuperNeuro: A Fast and Scalable Simulator for Neuromorphic Computing." In Proceedings of the 2023 International Conference on Neuromorphic Systems, pp. 1-4. 2023.

Dehallucination of LLMs for High-Level Planning

Scientific Achievement

- Large language models can generate plans for solving high-level planning problems, such as the operation of robots in DOE national laboratories.
- While the plans may appear to be of high quality, it is not uncommon for the produced plans to contain actions that cannot be executed in reality.
- We have developed a framework that mitigates hallucinations (generated errors) in LLM generated plans.

Significance and Impact

The project provides a solution to specifying scientific problems in natural languages (or text) while solving them using neuro symbolic methods. This is a step towards lowering technical barriers for future engineers and scientists.

Technical Approach

- The code generation capabilities of the LLM is used to specify logical constraints that every generated plan must satisfy.
- A solver is used to automatically check the adherence to the constraints and provide feedback to the AI model regarding unsatisfied constraints.
- The feedback allows the LLM to generate a new provably correct plan.

PI(s)/Facility Lead(s): Rickard Ewetz, Sumit Kumar Jha; University of Central Florida
Collaborating Institutions: Florida International University
ASCR Program: EXPRESS, Explainable AI
ASCR PM: Margaret Lentz
Publication(s) for this work: S. Jha, et al., "Counterexample Guided Inductive Synthesis Using Large Language Models and Satisfiability Solving," MILCOM, November, (2023). (to appear).

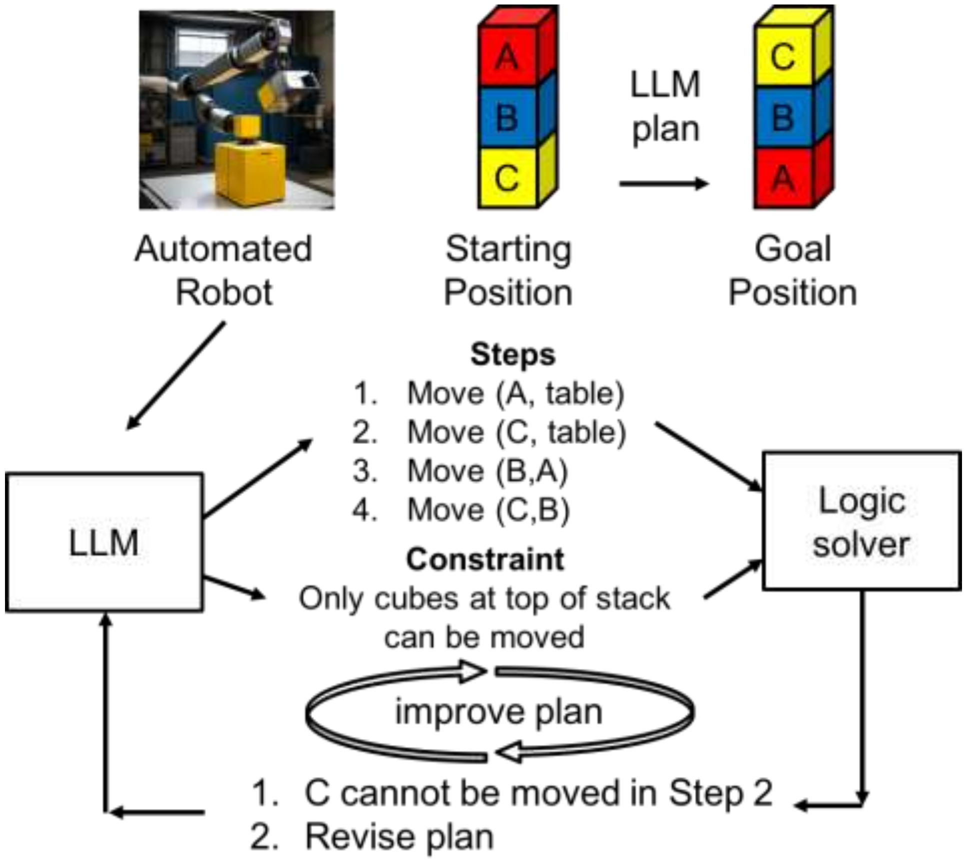

The LLM generates a high-level plan for moving the starting position to the goal position. The LLM also generates mathematical constrains describing how cubes are allowed to be moved. The plan and the constraints are fed into a logic solver, which determines that a constraint is violated in step 2. The C cube is attempted to be moved while cube B is on top. The solver provides feedback to the LLM why the plan is infeasible such that a new legal plan can be generated.