U.S. DEPARTMENT OF ENERGY | Office of Science

Updates from ASCR Computational Science Research and Partnerships Division

September 29, 2021

FY2021 At-A-Glance: 12 Solicitations, 129 New Awards

Early Career Research Program	10 new projects
SciDAC Partnerships in Basic Energy Sciences	5 new projects*
X-Stack: Programming Environments for Scientific Computing	5 new projects
Integrated Computational and Data Infrastructure (ICDI) for Scientific Discovery	3 new projects
Entanglement Management and Control in Transparent Optical Quantum Networks	5 new projects
5G Enabled Energy Innovation Advanced Wireless Networks for Science	5 new projects
Microelectronics Co-Design Research	10 new projects*
EXPRESS: Randomized Algorithms for Extreme-Scale Science	6 new projects
Data-Intensive Scientific Machine Learning and Analysis	5 new projects
Quantum Internet to Accelerate Scientific Discovery	3 new projects
Data Reduction for Science	9 new projects
Bridge2AI and Privacy-Preserving Artificial Intelligence Research	1 new project

* Indicates partnerships with other SC offices

FY2021 At-A-Glance: 8 Community Events

THZ-6G Wireless Communications Roundtable	October 1st, 2020
DOE SC Community of Interest Workshop	November 2nd, 5th, 10th, 2020
RASC: Randomized Algorithms for Scientific Computing	December 2-3, 2020; January 6-7, 2021
Roundtable Discussion on Operating-Systems Research	January 25th, 2021
Data Reduction for Science Workshop	January 25-26 and 28th, 2021
SciDAC-4 Virtual Get Together	July 28th, 2021
ASCR Workshop on Reimagining Co-Design	March 16-18, 2021
ASCR Roundtable on Parallel Discrete Event Simulation	September 20th, 2021

NUCLEI SciDAC-4 Collaboration a Catalyst for Early Career Researchers

- Six early-career NUCLEI scientists awarded DOE Office of Science Early Career Research Programs (ECRPs)
- Their ECRPs range from machine learning to dynamics in nuclei to mathematical optimization to ab initio methods
- NUCLEI is a multidisciplinary SciDAC collaboration bringing together applied mathematicians, computer scientists, and nuclear physicists to advance understanding of atomic nuclei

DOE ASCR

Prasanna Balaprakash
2018

Stefan Wild
2020

NUCLEI
Nuclear Computational Low-Energy Initiative

SciDAC
Scientific Discovery through Advanced Computing

DOE NP

Heiko Hergert
2017

Stefano Gandolfi
2018

Alessandro Lovato
2020

Maria Piarulli
2021

Projected Land Ice Contribution to 21st Century Sea Level Rise

Scientific Achievement

Estimate the range of future sea-level rise from glaciers and ice sheets (land ice) using the most recent models and emissions scenarios.

Significance and Impact

- Greatly improved models and methods provide the most comprehensive projections of sea-level rise from land ice to date.
- Results confirm that Antarctica remains a critical focus for reducing future sea level uncertainty.
- Limiting global warming to 1.5°C reduces 21st century land ice contribution to sea-level rise from 25 to 13 cm.

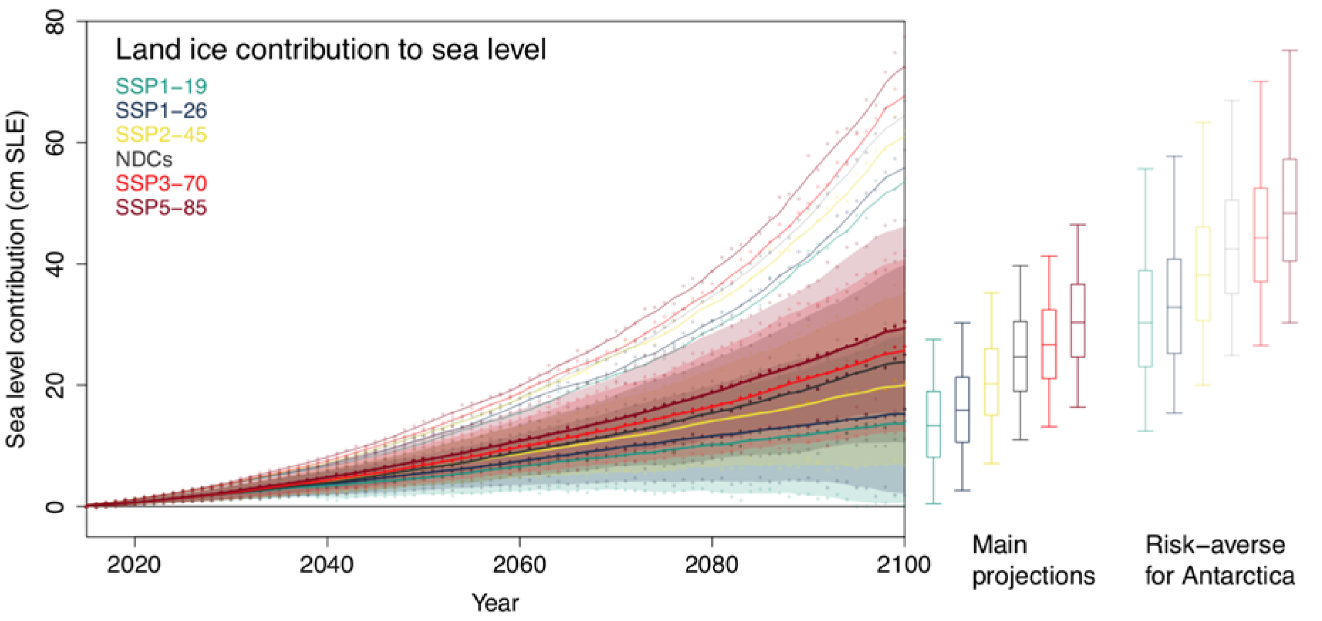

Projected 2015-2100 land ice contribution to sea level for a range of emissions scenarios. Solid lines and shaded regions show median and 5-95th percentile estimates, respectively. Pale solid lines denote 95th percentiles for risk-averse projections. Box-and-whiskers indicate 5, 25, 50, 75, 95th percentiles at 2100.

Research Details

- Apply multiple glacier models, ice sheet models, climate models, and emissions scenarios to characterize the impact of land ice evolution on future sea level rise.
- Using these multi-model ensembles and statistical emulation, build probabilistic projections of future sea-level rise from land ice.
- DOE ASCR- and BER-supported researchers contributed at multiple levels including selecting appropriate climate models (HiLAT), developing model parameterizations (ProSPect), and conducting high-fidelity, high-resolution ice sheet model simulations using MALI and BISICLES (ProSPect).

Edwards, T. and 84 others (including 7 DOE), *Nature* **593**, doi:10.1038/s41586-021-03302-y

Work was performed at Lawrence Berkeley and Los Alamos National Laboratories, using NERSC resources.

Cost-Function-Dependent Barren Plateaus in Shallow Quantum Neural Networks

Scientific Achievement

We refined the understanding of barren plateaus, or the trainability, or quantum circuits by showing the dependence on the output that is measured.

Significance and Impact

This constitutes one of the few rigorous scaling results proven for variational quantum algorithms (VQAs). VQAs are viewed as the best hope for near-term quantum advantage, but our work shows that care must taken when defining the cost function. Global cost functions are untrainable for all circuit depths. Local cost functions are trainable for O(log n) depth.

Research Details

Assuming a hardware-efficient ansatz composed of blocks forming local 2-designs, we prove universal bounds on the variance of the gradient, which allow us to study how the gradient scales with circuit depth *D* and system size *n*.

"Cost-Function-Dependent Barren Plateaus in Shallow Quantum Neural Networks", Cerezo, Sone, Volkoff, **Cincio**, **Coles**. *Nature Comm.* **12**, 1791 (2021)

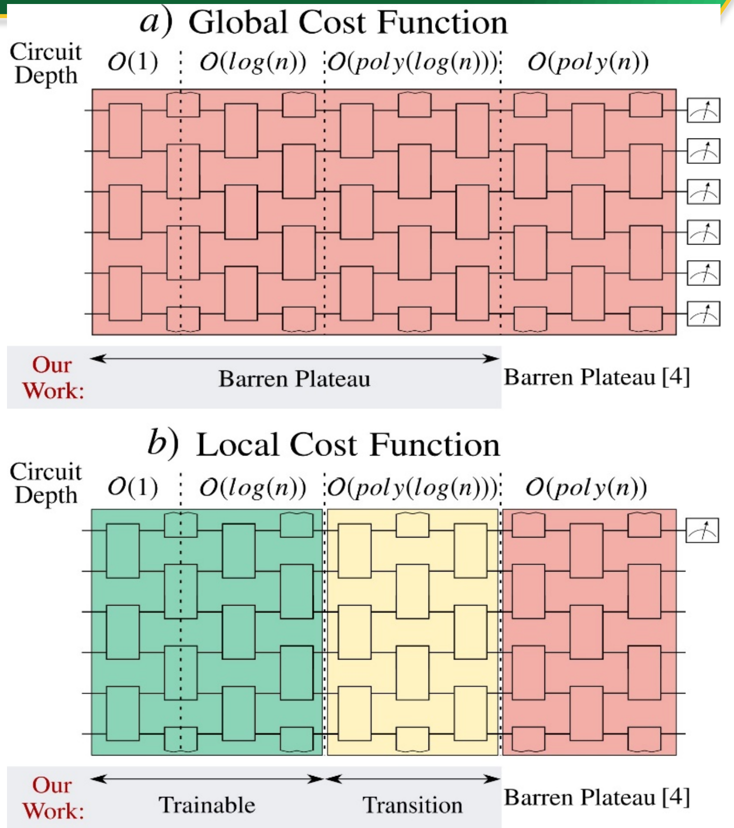

Caption: We find that barren plateaus are cost-function dependent. a) For global cost functions the landscape will exhibit a barren plateau for all depths D. b) For local cost functions, the gradient vanishes polynomially and hence is trainable when D = O(log(n)), while barren plateaus occur for D = O(poly(n)), and between these regions the gradient transitions from polynomial to exponential decay.

FlexFlow & Legion: Accelerating Deep Learning at Scale

Legion Enables Improved Deep Learning Performance

The days to weeks required to train Deep Neural Networks (DNNs) from large data sets is a bottleneck for productive science. FlexFlow is a DNN framework that leverages the Legion parallel and distributed runtime system to provide a significant performance improvement over TensorFlow.

Improved Productivity, System Utilization, and Broader Impacts

Reduced training times by ~15x for the CANcer Distributed Learning Environment (CANDLE).
Faster deployment of machine learning technologies to help address mission-focused efforts.
Provides significant reductions in time-to-solution and scaling compared to existing methods.
Enabling insights from extreme scale data sets and ML approaches to better match the computing capabilities at the Office of Science Leadership Computing Facilities.

FlexFlow now supports “entry points” for industry-standard API’s: Keras, PyTorch, ONNX.

Significant intertest and engagement across industry and academia (e.g., Facebook, NVIDIA, MIT, and more).

FlexFlow: https://flexflow.ai Legion: https://legion.stanford.edu

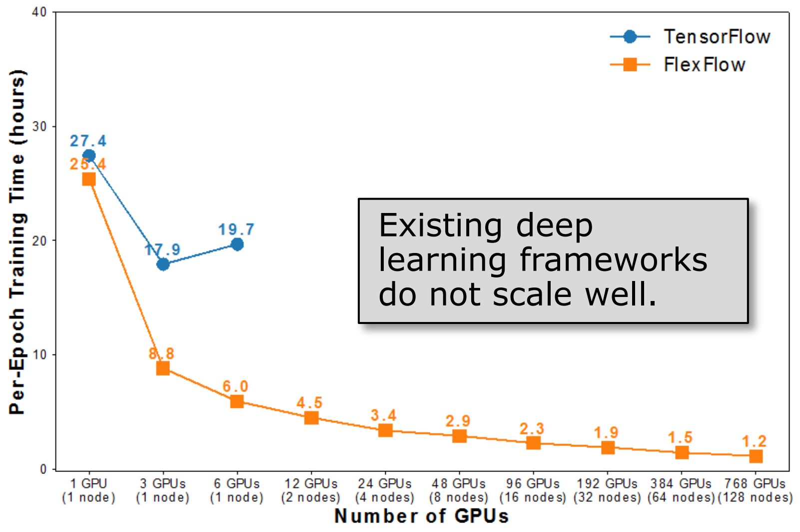

Improved Productivity & System Utilization: The FlexFlow system, built on top of Legion, provides improved training times in comparison to Google’s TensorFlow. Scale to 768 GPUs and reduces the per-epoch training time from 18 hours to 1.2 hours.

Scaling Implicit Parallelism via Dynamic Control Replication: M. Bauer, W. Lee, E. Slaughter, Z. Jia, M. Di Renzo, M. Papadakis, G. Shipman, P. McCormick, M. Garland, and A. Aiken, Principles and Practices of Parallel Programming 2021, Feb. 2021.
Exploring Hidden Dimensions in Parallelizing Convolutional Neural Networks, Z. Jia, S. Lin, C. R. Qi, A. Aiken, in Proceedings of the International Conference on Machine Learning (ICML), Stockholm, Sweden, July 2018.

The Mochi Project: Improving science productivity through specialized data services

The Science

Exascale supercomputers will generate and consume massive amounts of data and storing and exploring that data to extract scientific insights are challenges that are exacerbated by the increasing diversity in the way data are used in scientific endeavors. The Mochi project is addressing this challenge by providing a flexible framework for developing new data services that closely match specific DOE mission needs. This framework makes it much easier for teams to develop these services than if they were built independently.

Mochi, developed by a collaboration between Argonne National Laboratory, Carnegie Mellon University, Los Alamos National Laboratory, and The HDF Group, adopts the microservices paradigm common in industry to compose multiple data services in a flexible framework. These microservices handle everything from traditional high-performance data management to specialized data services to feed machine-learning workflows.

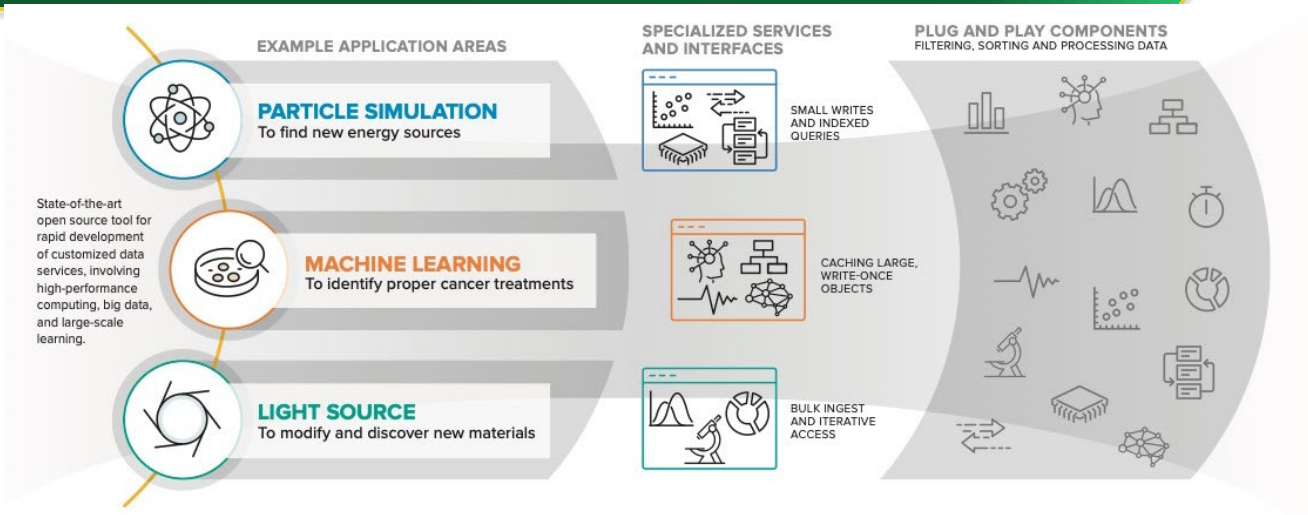

The Mochi framework allows for co-design of data services with science teams. Science teams identify requirements and work with computer scientists to define a productive interface and data model, then computer scientists compose this service using a set of plug-and-play components.

The Impact

Mochi technology is now used by industry, academia, and research organizations, including by Kitware Inc. in their Phase 2 SBIR work. The high-performance core of Mochi, called Mercury, is a key part of Intel's DAOS storage solution that will be featured on DOE's upcoming Aurora exascale supercomputer. **The Mochi technology is a finalist for the 2021 R&D100 award for its impact on service development, and Rob Ross was a recipient of the 2020 E. O. Lawrence award in part for his contributions to this project**.

Supported by the Exascale Computing Project (ECP) and ASCR's research program, Mochi will be available on upcoming exascale supercomputers and many other computing resources across the country.

2021 R&D 100 FINALIST

Robert B. Ross et al. "Mochi: Composing Data Services for High-Performance Computing Environments", *Journal of Computer Science and Technology*. 35, 121–144 (2020).

Argonne NATIONAL LABORATORY | Carnegie Mellon University | HDF | Los Alamos NATIONAL LABORATORY EST. 1943

Recognition of Our Community Members

Ernest Orlando Lawrence Award

Rob Ross, ANL

RAPIDS-2 Director

2020 AAAS Fellow

Bert de Jong, LBNL

AIDE-QC Director

2020 AAAS Fellow

Esmond Ng, LBNL

FASTMath Director

2021 SIAM Fellow

Rob Falgout, LLNL

FASTMath PI

2021 SIAM Fellow

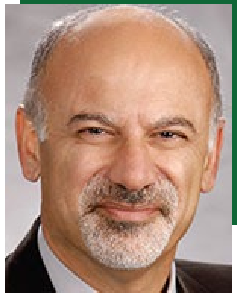

Habib Najm, SNL

FASTMath PI

2021 AWM Fellow

Carol Woodward, LLNL

FASTMath PI

Barry Smith "retires" at Argonne

Barry's work transformed how large-scale software libraries are developed, supported, and used – not only within DOE, but also across the scientific and engineering community

DOE highlights

- 1990 joined Argonne as the Wilkinson Postdoctoral Fellow
- 1994 rejoined Argonne, starting work on PETSc 2.0
- 1997 paper *"Efficient management of parallelism in object-oriented numerical software libraries"* w/ Satish Balay, Bill Gropp, & Lois McInnes
    - Named DOE SC 40 major papers that changed the face of science
- 1999 Gordon Bell special prize with Kaushik, Keyes, & Gropp
- 2009 R&D100 award to PETSc
- 2011 DOE E.O. Lawrence Award w/ Lois McInnes
- 2015 core PETSc dev team awarded SIAM/ACM CS&E Prize
- 2015 named Argonne Distinguished Fellow

2020-2021: Barry is a part-time Argonne Associate and Senior Research Scientist at the Flatiron Institute, focusing research on PETSc and the changing future of large-scale numerical simulation

Tammy Kolda Retires (from Sandia)

Tammy made foundational contributions to algorithms and software for tensor decompositions, graph analysis, and data science.

DOE highlights

- 1997 joined ORNL as the Householder Postdoctoral Fellow
- 1999 joined Sandia in Livermore, CA
- 2003 Presidential Early Career Award for Scientists and Engineers (PECASE)
- 2004 R&D100 Award for Trilinos Software (Team Award)
- 2005 Release of Tensor Toolbox for MATLAB
- 2015 Elected SIAM Fellow
- 2018 Founds SIAM Journal on the Mathematics of Data Science (SIMODS)
- 2018 Member, Board on Mathematics Sciences and Analytics (BMSA) National Academies
- 2019 Elected Fellow of the Association for Computing Machinery (ACM)
- 2020 Elected Member of the National Academy of Engineering (NAE)

2021: Tammy is an independent consultant and researcher under the auspices of her company MathSci.ai.

Phil Colella retires from Berkeley Lab June 2019 with 40 years service to University of California

- PhD 1979 UC Berkeley, Applied Math
- *Recognized for fundamental advances in high-resolution schemes and adaptive mesh refinement methods for solution of PDEs*
- Long time ASCR Applied Math PI
- Founding Lead of Applied Numerical Algorithms Group at Berkeley Lab; led development of Chombo
- 1998 IEEE Computer Society Sidney Fernbach Award
- 2003 inaugural SIAM/ACM Prize for CSE (with John Bell)
- 2004 Elected to National Academy of Sciences
- 2009 first class of SIAM Fellows
- 2018 Berkeley Lab Director's Lifetime Achievement Award

1980

2018

FY2022: A Preview

- **Data Visualization Beyond 4D**
- **Storage Systems and I/O**
- **Parallel Discrete Event Simulation (PDES)**
- **Explainable AI**
- **High-Productivity Environments for Scientific Computing**
- **Randomized Algorithms for Scientific Computing**
- **Federated Scientific Machine Learning**
- **MMICCs**
- **SciDAC Partnerships**