Revolutionizing Neuromorphic Computing for Science

2024 Neuromorphic Computing for Science
Workshop Report

September 12-13, 2024
Bethesda, MD

Sponsored by the U.S. Department of Energy
Office of Advanced Scientific Computing Research

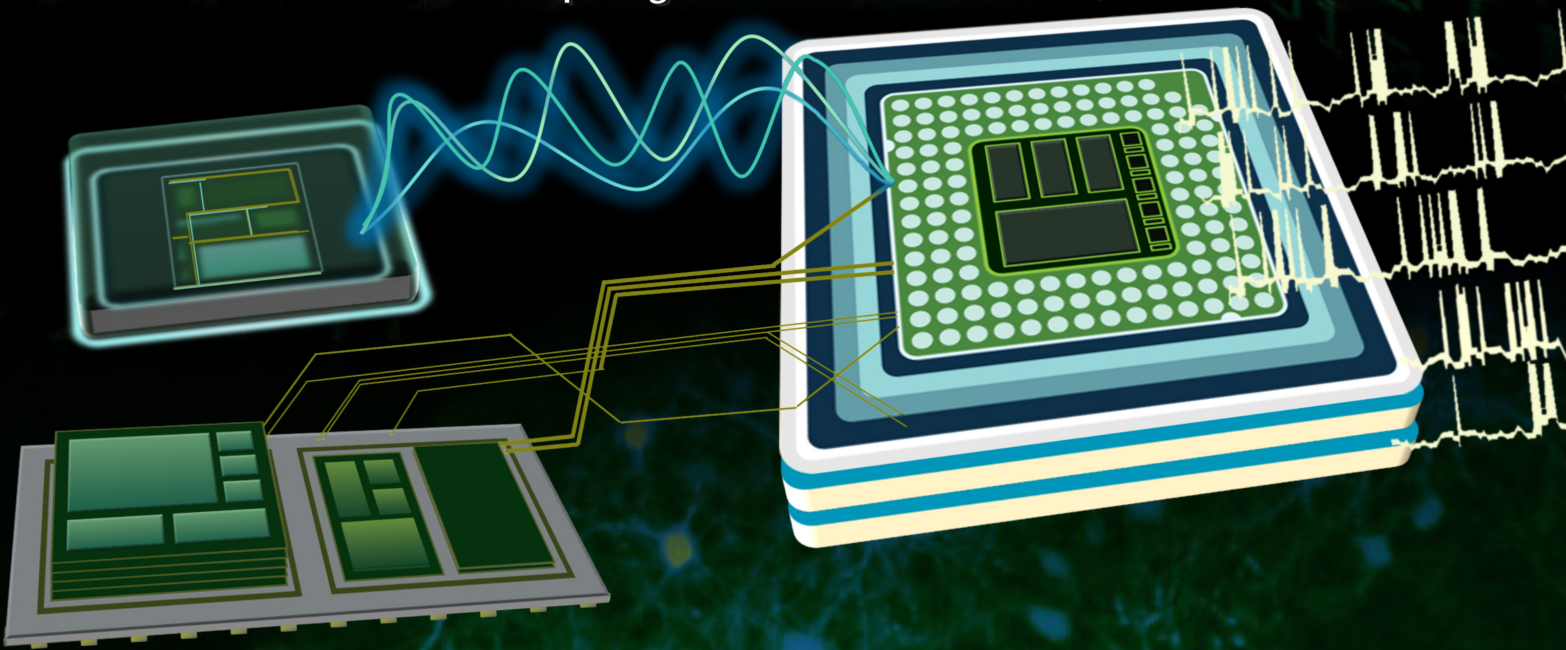

Co-chairs:
Gina Adam, George Washington University
Garrett Kenyon, Los Alamos National Laboratory
Thomas Potok, Oak Ridge National Laboratory

DOE Lead Program Manager:
Robinson Pino, Advanced Scientific Computing Research

U.S. DEPARTMENT OF ENERGY | Office of Science

Revolutionizing neuromorphic computing for science

Engineering novel neuromorphic computing systems with functionalities, capabilities, and energy efficiency similar to biological brains is one of the most exciting and challenging scientific endeavors of our time. In September 2024, the U.S. Department of Energy's Office of Science (DOE-SC) convened a workshop to identify potential opportunities and explore proof of principle neuromorphic circuits applicable for high performance computer (HPC) acceleration for large-scale AI and for scientific discovery. This workshop was motivated by the DOE-SC 's long-term investment in energy-efficient supercomputing and related technologies.

This workshop discussed the key research needs, challenges, and next steps necessary to develop biologically realistic neuromorphic circuit primitives (e.g. complex neurons, neuro-inspired connectivity motifs, etc.) that capture the functionality of neural systems found in nature. The need for innovations in the modeling and simulation of neuromorphic computing primitives is essential for understanding intelligent behavior at large scales meaningful for solving DOE-relevant scientific problems. Effective use and insights of neuroscience datasets and metrics are needed to vet proposed neuromorphic circuits and simulations.

Interdisciplinary collaboration is key to the success of revolutionizing neuromorphic computing. This workshop brought together a diverse range of experts from DOE national labs and academia to adequately address the multi-faceted challenges across the three main areas in:

- neuroscience,
- microelectronics, and
- large-scale simulations.

Workshop contributors submitted position papers that were discussed in breakout sessions cross-cutting neuroscience, microelectronics and large-scale simulation. Plenary sessions were used for keynote presentations to frame scientific opportunities and present roadmaps for future technologies. Interdisciplinary discussion groups identified four priority research directions (PRDs), detailed below.

The workshop report is available from https://www.osti.gov/ at DOI: https://doi.org/10.2172/2476969

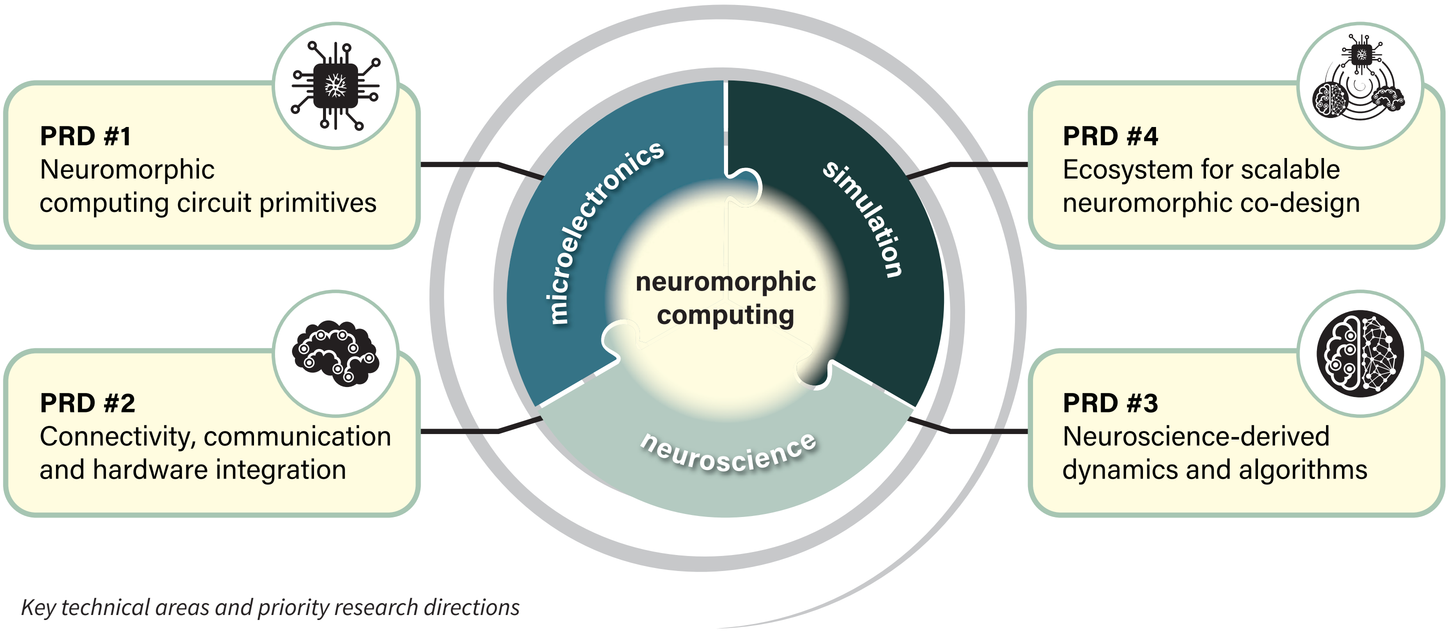

Key technical areas and priority research directions

Interdisciplinary innovation in neuromorphic computing Priority Research Directions

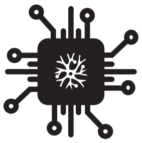

PRD #1: Accelerate the design and prototyping of neuromorphic computing circuit primitives

Key questions: What are the key neuromorphic circuit primitives that capture the functionality of critical biological computing mechanisms? How can appropriate models of the computational primitives and dynamics of diverse neural components be derived? How can these components be translated to neuromorphic circuits with the expressivity and complexity of biological neurons?

Revolutionizing neuromorphic computing requires understanding the computational components and principles, e.g., dendritic trees and local plasticity, that underpin the brain's functionality and robustness. These biological primitives must be translated into functionally equivalent circuits that could be implemented and validated with a computer system. Novel neuromorphic circuits based on current and emerging technologies guided by neuroscience-inspired functionality need to be engineered to realize energy efficient high-performing computational systems by six orders of magnitude.

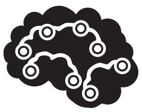

PRD #2: Transform neuromorphic connectivity, communication, and hardware integration

Key questions: What neural network connectivity from biology should be abstracted to enable the design of neuromorphic circuits and systems? How can brain connectivity motifs be translated into scalable communication mechanisms? How can neuromorphic primitives be organized for optimal energy efficiency and information movement?

Significant connectivity and communication challenges exist in tackling practical integration of neuromorphic computing hardware at scale and with instrumentation for scientific computing. Multi-scale connectivity, relevant circuit motifs from connectome data and efficient encoding schemes should be leveraged and appropriately translated to neuromorphic circuits. Advances in this critical domain will enable the design of high-bandwidth and massively parallel connectivity across neuromorphic processing units.

PRD #3: Leverage neuro-inspired dynamics and functionality to design neuromorphic algorithms

Key questions: How can scalable and high-performance neuromorphic algorithms be designed? How can multiscale neural dynamics be harnessed to develop adaptive, efficient, and robust neuromorphic systems with stable functionality? How can relevant emergent neural behavior and theories inform neuromorphic computing at scale?

To achieve the greatest impact, it is essential to develop strategies for creating neuromorphic algorithms that inherently leverage the dynamics and emergent behaviors of biological neural circuits at various scales, while also harnessing the efficiencies and parallelism offered by neuromorphic hardware. By taking advantage of the brain's heterogeneity, configurability, and dynamics, neuromorphic computing can revolutionize the next generation of advanced computing.

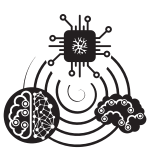

PRD #4: Pioneer an integrated ecosystem for scalable neuromorphic co-design

Key questions: How can the large multi-dimensional co-design space for neuromorphic circuits and systems be effectively explored? What are the capabilities needed for accelerating simulation and prototyping of neuromorphic circuits and systems? What infrastructure is needed for benchmarking and validation across the stack using neuroscience-derived data and metrics?

Understanding the functionality of neuromorphic systems of relevance to DOE applications requires the integration and scale-up of hardware-aware simulations up to hundreds of millions of neurons and synapses. Novel co-design methodologies and high-performance computing are therefore required. Capabilities for simulation and prototyping as well as data-driven benchmarks should support a strong integrated ecosystem in neuromorphic innovation relevant to the Department of Energy's mission.

Summary

A grand challenge for our time is to rethink computing in an efficient, intelligent, and extremely scalable way beyond the existing digital AI paradigm by using inspiration from the brain. Neuromorphic computing offers the promise of enhancing energy efficiency and enabling new algorithms and capabilities both for high performance computing and at the edge. Therefore, interdisciplinary research is critical to identify the key neuromorphic primitives and their scalable modeling and implementation. Breakthroughs in understanding, designing, and prototyping the circuitry and simulation capabilities for a truly neuromorphic computer are essential to enable progress in the field. By taking advantage of the brain's heterogeneity, configurability, and dynamics, neuromorphic computing has the potential to advance large-scale scientific simulations, revolutionize intelligent agents and deliver energy efficiency beyond the exascale. The impact of neuromorphic computing and the fundamental basic research needed in this area is essential to the success of scientific computing relevant to the Department of Energy.

Workshop Co-Chairs: Gina Adam, George Washington University, **Garrett Kenyon**, Los Alamos National Laboratory, **Thomas Potok**, Oak Ridge National Laboratory
DOE Lead Program Manager: Robinson Pino, Office of Science, Advanced Scientific Computing Research
Bethesda, MD, September 12–13, 2024, DOI: https://doi.org/10.2172/2476969.

Organizing committee: Giorgio Ascoli, George Mason University, **Frances Chance**, Sandia National Laboratories, **Yiran Chen**, Duke University, J**oseph Friedman**, University of Texas at Dallas, **Cory Merkel**, Rochester Institute of Technology, **Maryam Parsa**, George Mason University, **Midya Parto**, University of Central Florida, **Catherine Schuman**, University of Tennessee Knoxville, **Shinjae Yoo**, Brookhaven National Laboratory, **Yuping Zeng**, University of Delaware

Scientific contributors: Yohannes Abate, University of Georgia, **Tosiron Adegbija**, University of Arizona, **Brad Aimone**, Sandia National Laboratories, **Alireza Alemi**, UC Davis, **Harbir Antil**, George Mason University, **Kazi Asifuzzaman**, Oak Ridge National Laboratory, **Ahmedullah Aziz**, University of Tennessee Knoxville, **Adarsha Balaji**, Argonne National Laboratory, **Srutarshi Banerjee**, Argonne National Laboratory, **Frank Barrows**, Los Alamos National Laboratory, **Kristofer Bouchard**, LBNL/UC Berkeley, **Suma Cardwell**, Sandia National Laboratories, **William Chapman**, Sandia National Laboratories, **Holger Dannenberg**, George Mason University, **Hritom Das**, Oklahoma State University, **Prasanna Date**, Oak Ridge National Laboratory, **Rogene Eichler West**, Pacific Northwest National Laboratory, **Ashish Gautam**, Oak Ridge National Laboratory, **Andreas Gerstlauer**, The University of Texas at Austin, **Luisa Gonzalez**, Lawrence Berkeley National Laboratory, **Qing Gu**, North Carolina State University, **Chathika Gunaratne**, Oak Ridge National Laboratory, **Akhilesh Jaiswal**, University of Wisconsin Madison, **Sumit Jha**, Florida International University, Miami, Florida, **Erik Johnson**, Johns Hopkins University Applied Physics Laboratory, **Shubha Kharel**, Brookhaven National Laboratory, **Dhireesha Kudithipudi**, University of Texas at San Antonio, **Shruti Kulkarni**, Oak Ridge National Laboratory, **Suhas Kumar**, Sandia National Labs, **Hai "Helen" Li**, Duke University, **Siddharth Mansingh**, Los Alamos National Laboratory, **Andres Marquez**, Pacific Northwest National Laboratory, **Prashansa Mukim**, Brookhaven National Laboratory, **Alex Noy**, Lawrence Livermore National Laboratory, **Robert Patton**, Oak Ridge National Laboratory, **Praful Purohit**, Brookhaven National Laboratory, **Xiaoning Qian**, Brookhaven National Laboratory, **Qinru Qiu**, Syracuse University, **Brian Robinson**, Johns Hopkins University Applied Physics Laboratory, **Matt Roos**, Johns Hopkins University Applied Physics Laboratory, **Fred Rothganger**, Sandia National Labs, **Vishal Saxena**, University of Delaware, **Jeffrey Shainline**, National Institute of Standards and Technology, **Anshumali Shrivastava**, Rice University, **Andrew Sornborger**, Los Alamos National Laboratory, **Mircea Stan**, University of Virginia, **Ishan Thakkar**, University of Kentucky, **Dmitrii Torbunov**, Brookhaven National Laboratory, **Ing Tsang**, IMEC, **Lav Varshney**, University of Illinois Urbana-Champaign, **Felix Wang**, Sandia National Laboratories, **Cheng Wang**, Iowa State University, **Yingxue Wang**, Max Planck Florida Institute for Neuroscience, **Chris Yakopcic**, University of Dayton, **Angel Yanguas-Gil**, Argonne National Laboratory, **S. J. Ben Yoo**, University of California, Davis, **Aaron Young**, Oak Ridge National Laboratory, **Ramtin Zand**, University of South Carolina

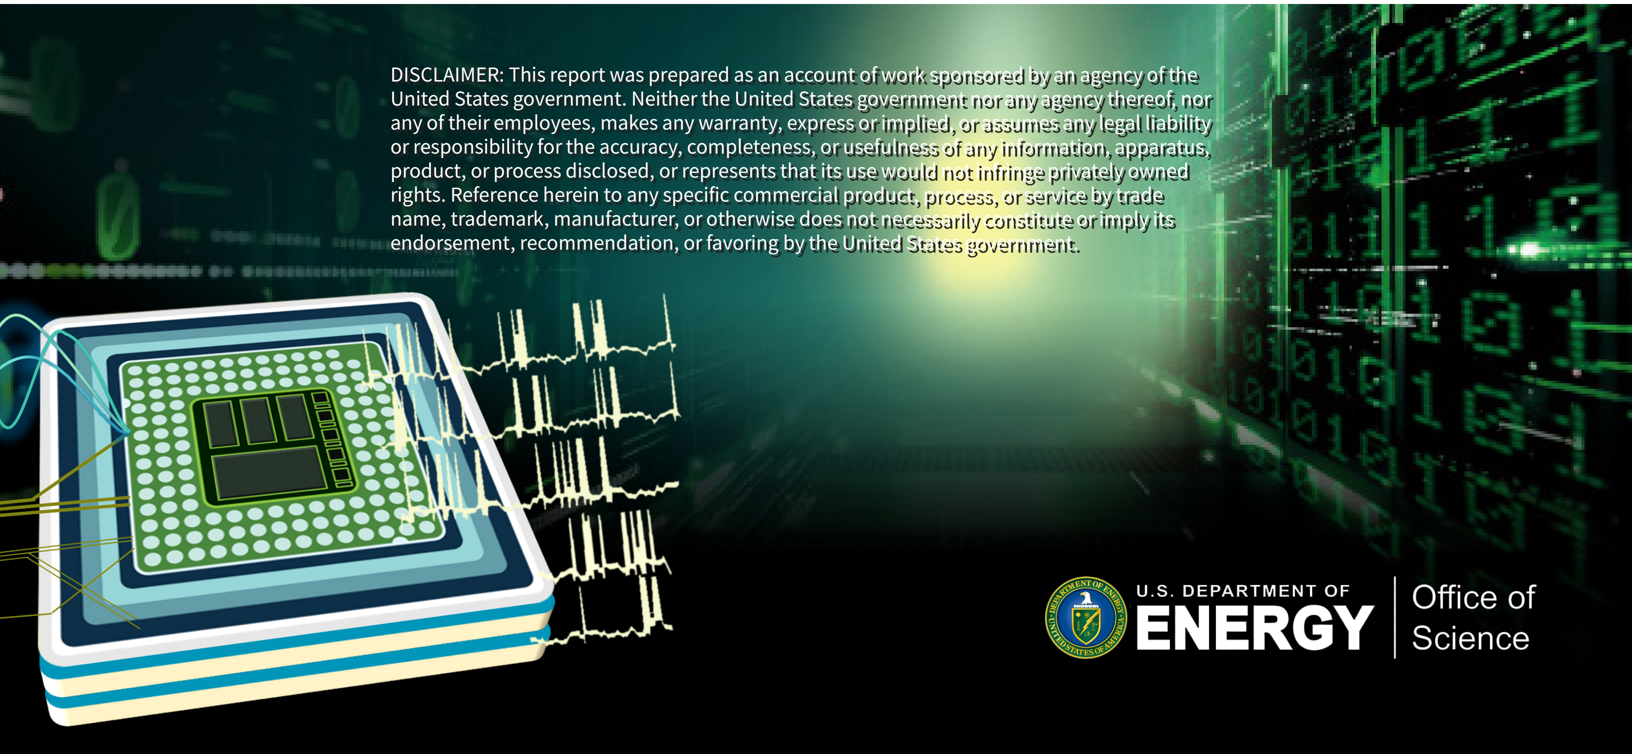

U.S. DEPARTMENT OF ENERGY | Office of Science