**COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY**
**SUBCOMMITTEE ON ENERGY**
**U.S. HOUSE OF REPRESENTATIVES**
**HEARING CHARTER**

*Accelerating Discovery: The Future of Scientific Computing at the Department of Energy*
Wednesday, May 19, 2021
11:00AM EDT

## Purpose

The purpose of this hearing is to explore the unique scientific computing capabilities of the Department of Energy (DOE), including the forthcoming exascale systems, and to discuss the implications of these capabilities for other scientific disciplines and their relevance to pressing societal challenges. In addition, the Subcommittee will use the hearing to understand the role of DOE research and workforce development programs in driving innovation in scientific computing, especially in light of advancements in artificial intelligence, quantum science, neuromorphic computing, and other new and emerging capabilities and computing paradigms. The hearing will also examine ways in which Congress can contribute to DOE's scientific computing mission.

## Witnesses

- **Dr. J. Stephen Binkley**, Acting Director, Office of Science, Department of Energy
- **Dr. Georgia (Gina) Tourassi**, Director, National Center for Computational Sciences, Oak Ridge National Laboratory
- **Dr. Karen Willcox**, Director, Oden Institute for Computational Sciences and Associate Vice President for Research, University of Texas at Austin
- **Dr. Christopher Monroe**, Co-Founder and Chief Scientist, IonQ, Inc.
- **Dr. Seny Kamara**, Associate Professor of Computer Science, Brown University

## Overview of Scientific Computing at DOE

The Department of Energy (DOE) possesses some of the most powerful high-performance computing (HPC) systems in existence, including two of the top three fastest currently operating supercomputers in the world,[1] and is slated to deploy even more capable systems later this year. Housed at several of DOE's national laboratories, these systems enable researchers from both within and outside of the Department to accelerate discovery across numerous fields of scientific inquiry and technology development. Specifically, researchers use HPC systems to analyze huge data sets and test computational models that can help identify promising drug candidates, aid in the development of advanced manufacturing techniques, and derive new knowledge from the information generated by DOE's other experimental facilities. In so doing, DOE's HPC systems serve as critical resources for academic and industry users from the U.S. and around the world and are a significant component of U.S. economic competitiveness, scientific leadership, and national security.

---

[1] https://www.top500.org/lists/top500/2020/11/

An essential element of DOE's national and international leadership in HPC is co-design. Co-design refers to the synergistic development of a computer's physical components such as central processing units and data storage devices—the hardware—and the sets of instructions that enable the hardware components to work together—the software. This approach enables the design and development of HPC systems, such as the forthcoming exascale computers, that are dramatically more powerful, reliable, programmable, and energy efficient than their predecessors. DOE's focus on co-design will also be critical to the successful development of systems that rely on more nascent paradigms, such as quantum computing, and which are able to fully integrate artificial intelligence and machine learning capabilities.

DOE's HPC enterprise and the approaches that have enabled its successes, such as co-design, rest on a portfolio of research activities supported by the Office of Science. These include research that is foundational for future computing capabilities, as well as efforts to build tools that allow for the application of HPC to other disciplines. In addition, crosscutting initiatives in quantum, artificial intelligence, and microelectronics seek to leverage the relevant expertise contained within various Office of Science programs to advance the frontiers of scientific computing.

**Advanced Scientific Computing Research (ASCR) program**

The DOE Office of Science's Advanced Scientific Computing Research (ASCR) program is the Department's main sponsor of research in foundational areas such as applied mathematics and computer science as well as emerging computing technologies. In addition, ASCR supports four scientific user facilities that collectively offer a suite of HPC systems as well as research testbeds for advancing novel computing and networking technologies. ASCR received $1.015 billion in fiscal year (FY) 2021.

ASCR's agenda has been largely shaped by the DOE-wide Exascale Computing Initiative (ECI), which is aiming to deploy the nation's first exascale computer later this year at Oak Ridge National Laboratory. ECI has focused ASCR's research and construction activities on meeting the short-term needs associated with realizing that goal, but as the initiative comes to fruition, ASCR is reorienting toward longer term research on emerging technologies and computing paradigms that will enable the advancement of HPC capabilities after the end of Moore's law.[2] This shift was precipitated by strategic visioning efforts that produced a series of recommendations for reinvigorating ASCR's core research programs while enabling the sustainment of the exascale ecosystem.[3]

---

[2] In 1975, Gordon Moore, founder of Fairchild Semiconductor and Intel, observed that the number of transistors that could be included on a silicon chip was doubling every 18-24 months while the cost of computers was decreasing by a proportional amount at the same time. This trend has enabled consistent increases to computing speed and capability, but fundamental physical and economic limitations are predicted to usher in a post-Moore's law world in which new computing technologies must be used if productivity gains are to be maintained. This impending shift is driving exploration into new and emerging paradigms such as quantum and neuromorphic computing, the development of novel materials for use in microelectronics, and new mathematics for overcoming physical constraints of existing computing architectures.

[3] https://science.osti.gov/-/media/ascr/ascac/pdf/meetings/202004/Transition_Report_202004-ASCAC.pdf?la=en&hash=5164916FE5158EE8919C26804B4CF7F6DDA36E9D

Additional detail on ASCR programs, activities, and user facilities can be found below.

*Mathematical, Computational, and Computer Sciences Research*

ASCR's core research activities provide the foundation for strengthening the scientific computing ecosystem through the development of software, algorithms, methods, tools, and workflows that anticipate future hardware challenges and opportunities as well as science application needs. ASCR's research programs also support interdisciplinary activities through partnerships with other programs both within Office of Science and in DOE's applied energy offices. In so doing, ASCR leads the development of advanced scientific computing applications that help accelerate discovery in other disciplines and areas of strategic importance to DOE. Finally, ASCR's research account also supports workforce development activities aimed at equipping scientists with the knowledge and expertise required to successfully employ computational science to advance research in other disciplines.

Despite consistent growth to the ASCR topline, the research budget dropped steeply and then stagnated in recent years as resources were shifted toward ECI activities such as the Exascale Computing Project. However, funding for these activities has since rebounded as the office continues its shift toward more future-focused research activities as mentioned above. In FY 2021, ASCR's research budget was $250 million.

ASCR's research and workforce development subprograms include the following:

- **Applied Mathematics** – The Applied Mathematics subprogram supports research in scalable algorithms and libraries, multi-physics and multi-scale modeling, artificial intelligence and machine learning (AI/ML), uncertainty quantification, optimization, and efficient data analysis among many other areas which underpin DOE's computational and data-intensive science efforts. Among the largest activities funded under this subprogram are the Mathematical Multifaceted Integrated Capabilities Centers (MMICCs), which drive mathematical advances that address complex challenges in science and energy research, and that require new integrated processes across multiple computational and statistical disciplines.[4]
- **Computer Science** – The Computer Science subprogram supports research to enable computing and networking at extreme scales and the ability to understand complex data generated by both simulations and experiments. In particular, the activity supports the development of adaptive software tools, data infrastructure, and cybersecurity innovations that enhance HPC productivity and enable the transition to increasingly complex and heterogeneous computing paradigms.[5]
- **Computational Partnerships** – The Computational Partnerships activity primarily supports the Scientific Discovery through Advanced Computing (SciDAC) program. Established in 2001, SciDAC supports research collaboration between discipline scientists, applied mathematicians, and computer scientists with the goal of accelerating progress in areas of strategic importance through the development of novel modeling and

[4] https://www.energy.gov/sites/default/files/2020/03/f72/doe-fy2021-budget-volume-4_0.pdf
[5] Ibid.

simulation tools. ASCR also leverages SciDAC as one of its primary mechanisms for developing new algorithms and applications for AI and future computing paradigms such as quantum information systems.[6]

- **Research and Evaluation Prototypes (REP)** – The REP program serves as a mechanism for partnerships with industry and academia in the development of next-generation computing systems and emerging architectures. Specifically, it is through REP that ASCR is supporting quantum testbed activities.[7]

In addition to these research activities, ASCR's research account also funds the **Computational Sciences Graduate Fellowship (CSGF)**. Established in 1991, CSGF is designed to meet DOE's growing need for computational scientists across numerous core mission areas. Fellows enrolled in the program are required to undertake a curriculum that combines study in a scientific or engineering discipline along with substantive work in applied mathematics and computer science. The fellowship also includes a 12-week research experience at a DOE national naboratory.[8] In so doing, CSGF supplies DOE with an ongoing pipeline of scientists and engineers with computational research experience and close ties to DOE's national laboratories. CSGF has been flat funded at $10 million since FY 2016 even as ASCR's research mandate has expanded to include areas such as AI and quantum computing.

*User Facilities*

ASCR stewards four user facilities, including three HPC facilities and one network user facility. Collectively, these facilities were funded at $575 million in FY 2021, maintaining the high-water mark set during the previous year as ASCR prepares to deploy the first exascale systems. Additional detail on each facility can be found below:

- **Leadership Computing Facilities (LCFs)** – ASCR supports two LCFs, housed at Argonne and Oak Ridge National Laboratories. The LCFs enable open scientific applications, including industry applications, to harness the potential of leadership computing to advance science and engineering.[9] The capabilities offered by these facilities have been applied to a diverse array of areas, from simulations of cosmological phenomena to understanding the interactions between drug receptors and signaling proteins. Additional information on the LCFs is available below.
    - **Oak Ridge Leadership Computing Facility (OLCF)** – OLCF currently operates testbeds in support of ECI and the 200 pf IBM/NVIDIA system (Summit), which achieved the global number one ranking as the world's fastest system in June 2018, November 2018, June 2019, and November 2019. OLCF supported nearly 1,500 users in FY 2019 and is set to become the site of the nation's first exascale computer, Frontier, later this year.[10] [11]

---

[6] Ibid.
[7] https://science.osti.gov/ascr/Facilities/REP
[8] https://science.osti.gov/ascr/CSGF
[9] https://www.energy.gov/sites/default/files/2020/03/f72/doe-fy2021-budget-volume-4_0.pdf
[10] https://science.osti.gov/ascr/Facilities/User-Facilities/OLCF
[11] https://www.olcf.ornl.gov/

- **Argonne Leadership Computing Facility (ALCF)** – ALCF operates an 8.5 pf Intel/Cray system (Theta) and testbeds to prepare users, applications, and software technology for the nation's second exascale system, Aurora, which is slated for deployment in 2022. ALCF supported over 1,200 users in FY 2019.[12] [13]

- **National Energy Research Scientific Computing Center (NERSC)** – Housed at Lawrence Berkeley National Laboratory (LBNL), NERSC provides access to HPC resources to the entire Office of Science community, enabling and enhancing computational research across a variety of disciplines including astrophysics, chemistry, earth systems modeling, materials, high energy and nuclear physics, fusion energy, and biology. NERSC users come from nearly every state in the U.S., with about 49 percent based in universities, 46 percent in DOE laboratories, and 5 percent in other government laboratories and industry.[14] NERSC supported over 6,600 users in FY 2019.[15] [16]

- **Energy Sciences Network (ESnet)** – Operated by LBNL, ESnet is a high-speed network engineered and optimized to support DOE's large-scale scientific research. It connects the entire national laboratory complex, including its supercomputer centers and user facilities, and allows scientists to access and use these assets independent of time and location. ESnet also serves as a testbed for network design and operations systems.[17]

## Crosscutting Initiatives in Computing

DOE supports several crosscutting initiatives that primarily concern, or directly involve, computing. ASCR either leads or is a major partner in executing each of them.

### *Exascale Computing Initiative*

The Exascale Computing Initiative (ECI) is a partnership between the DOE Office of Science and DOE's National Nuclear Security Administration (NNSA) focused on developing and deploying exascale computing systems capable of serving the Department's needs in mission critical areas. The initiative will culminate in the delivery of three exascale systems, including Aurora at ALCF and Frontier at OLCF, as well as El Capitan at Lawrence Livermore National Laboratory. The latter will be designed to meet needs associated with NNSA's mission to maintain the nation's nuclear deterrent and will serve as a resource for all three NNSA laboratories (Los Alamos, Sandia, and Lawrence Livermore).

Within the Office of Science, ASCR is the lead organization for ECI. As part of its efforts, ASCR launched the **Exascale Computing Project (ECP)** in FY 2017. ECP represents the research aspects of ASCR's participation in ECI. Specifically, ECP supports research and development to ensure that relevant hardware and software, including applications, are available

---

[12] https://science.osti.gov/ascr/Facilities/User-Facilities/ALCF
[13] https://www.alcf.anl.gov/
[14] https://www.energy.gov/sites/default/files/2020/03/f72/doe-fy2021-budget-volume-4_0.pdf
[15] https://science.osti.gov/ascr/Facilities/User-Facilities/NERSC
[16] https://www.nersc.gov/
[17] https://science.osti.gov/ascr/Facilities/User-Facilities/ESnet

to equip researchers to use the exascale systems to meet scientific needs.[18] Outside of ASCR, the Basic Energy Sciences (BES) and Biological and Environmental Research (BER) programs support research efforts to deliver exascale-relevant software for materials and chemical sciences and Earth systems modeling, respectively.

*Quantum Information Science*

The DOE Office of Science has implemented a multi-program research effort in quantum information science (QIS) as part of the government-wide initiative authorized in the *National Quantum Initiative Act*. DOE's QIS efforts focus on fundamental research, developing instrumentation and other scientific tools and equipment, and providing community resources. Target applications of relevance to this hearing include the following:

- **Quantum Computing** – Quantum computers exploit the properties of quantum states, such as superposition and entanglement, to perform computations. While quantum computers are not a substitute for classical computers such as the exascale systems, they hold the potential for driving transformational progress in areas such as chemistry, drug development, financial modeling, and climate and weather prediction.

- **Quantum Communication** – Quantum communications relies on quantum mechanics to analyze, process, and transmit information. Such systems possess the potential for extremely secure encryption, rapid exchange of massive amounts of data, and a major leap in sensing technologies. [19]

In FY 2021, appropriators provided $245 million for the Office of Science to support QIS research across all six of its program offices. ASCR received the largest share of this funding—at least $86 million—to support the development of algorithms, applications, and data infrastructure for quantum computing, as well as research and testbeds for quantum computing hardware and networks.

*Artificial Intelligence*

As authorized under the *National Artificial Intelligence Initiative Act of 2020*, DOE has implemented a cross-agency activity in artificial intelligence and machine learning (AI/ML). DOE considers AI/ML to be critical to advancing science and enhancing the capabilities of its user facilities and other research infrastructure. These views were articulated in DOE's *AI for Science* report, which provides an assessment of the current state-of-the-art in AI/ML and outlines challenges and opportunities for leveraging AI/ML to accelerate progress in a variety of topic areas.[20]

In FY 2021, the Office of Science received $100 million to support AI/ML activities, with a mandate to accelerate scientific discovery by applying AI/ML research outcomes to analyzing data from DOE's user facilities and large experiments. As with QIS, ASCR received the largest

[18] https://www.exascaleproject.org/
[19] https://science.osti.gov/Initiatives/QIS
[20] https://www.anl.gov/ai-for-science-report

share of this funding and is responsible for supporting foundational research in AI/ML and developing AI/ML software with relevance to DOE's mission area applications and scientific data sets. Other programs within the Office of Science are supporting research activities aimed at leveraging AI/ML tools for managing and analyzing the experimental data generated by their user facilities and accelerating discovery in the disciplinary sciences.

*Microelectronics*

In FY 2021, the DOE Office of Science launched a $45 million initiative in Microelectronics Innovation involving ASCR, BES, the Office's Fusion Energy Sciences (FES) program, and its High Energy Physics (HEP) program. Through this initiative, DOE is seeking to advance the state-of-the-art in microelectronics through a system of co-design in which materials, chemistry, devices, systems, architectures, algorithms, and software are developed in an integrated fashion.[21] BES received the largest share of the funding for these activities in FY 2021, though spending through ASCR is likely to grow in the coming years.

[21] https://www.energy.gov/sites/default/files/2020/03/f72/doe-fy2021-budget-volume-4_0.pdf