ACCELERATING DISCOVERY: THE FUTURE OF SCIENTIFIC COMPUTING AT THE DEPARTMENT OF ENERGY

HEARING

BEFORE THE

SUBCOMMITTEE ON ENERGY

OF THE

COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY
HOUSE OF REPRESENTATIVES

ONE HUNDRED SEVENTEENTH CONGRESS

FIRST SESSION

MAY 19, 2021

Serial No. 117–16

Printed for the use of the Committee on Science, Space, and Technology

Available via the World Wide Web: http://science.house.gov

U.S. GOVERNMENT PUBLISHING OFFICE
44–540PDF WASHINGTON : 2021

COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

HON. EDDIE BERNICE JOHNSON, Texas, *Chairwoman*

ZOE LOFGREN, California
SUZANNE BONAMICI, Oregon
AMI BERA, California
HALEY STEVENS, Michigan, *Vice Chair*
MIKIE SHERRILL, New Jersey
JAMAAL BOWMAN, New York
BRAD SHERMAN, California
ED PERLMUTTER, Colorado
JERRY McNERNEY, California
PAUL TONKO, New York
BILL FOSTER, Illinois
DONALD NORCROSS, New Jersey
DON BEYER, Virginia
CHARLIE CRIST, Florida
SEAN CASTEN, Illinois
CONOR LAMB, Pennsylvania
DEBORAH ROSS, North Carolina
GWEN MOORE, Wisconsin
DAN KILDEE, Michigan
SUSAN WILD, Pennsylvania
LIZZIE FLETCHER, Texas
VACANCY

FRANK LUCAS, Oklahoma, *Ranking Member*
MO BROOKS, Alabama
BILL POSEY, Florida
RANDY WEBER, Texas
BRIAN BABIN, Texas
ANTHONY GONZALEZ, Ohio
MICHAEL WALTZ, Florida
JAMES R. BAIRD, Indiana
PETE SESSIONS, Texas
DANIEL WEBSTER, Florida
MIKE GARCIA, California
STEPHANIE I. BICE, Oklahoma
YOUNG KIM, California
RANDY FEENSTRA, Iowa
JAKE LaTURNER, Kansas
CARLOS A. GIMENEZ, Florida
JAY OBERNOLTE, California
PETER MEIJER, Michigan
VACANCY

SUBCOMMITTEE ON ENERGY

HON. JAMAAL BOWMAN, New York, *Chairman*

SUZANNE BONAMICI, Oregon
HALEY STEVENS, Michigan
JERRY McNERNEY, California
DONALD NORCROSS, New Jersey
SEAN CASTEN, Illinois
CONOR LAMB, Pennsylvania
DEBORAH ROSS, North Carolina

RANDY WEBER, Texas, *Ranking Member*
JIM BAIRD, Indiana
MIKE GARCIA, California
RANDY FEENSTRA, Iowa
CARLOS A. GIMENEZ, Florida
PETER MEIJER, Michigan

C O N T E N T S

May 19, 2021

Page

Appendix II: Additional Material for the Record

ACCELERATING DISCOVERY: THE FUTURE OF SCIENTIFIC COMPUTING AT THE DEPARTMENT OF ENERGY

WEDNESDAY, MAY 19, 2021

HOUSE OF REPRESENTATIVES,
SUBCOMMITTEE ON ENERGY,
COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY,
Washington, D.C.

The Subcommittee met, pursuant to notice, at 11:03 a.m., via Zoom, Hon. Jamaal Bowman [Chairman of the Subcommittee] presiding.

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

Chairman BOWMAN. This hearing will now come to order. Without objection, the Chairman is authorized to declare a recess at any time.

Before I deliver my opening remarks, I wanted to note that, today, the Committee is meeting virtually. I want to announce a couple of reminders to the Members about the conduct of this hearing. First, Members should keep their video feed on as long as they are present in the hearing. Members are responsible for their own microphones. Please also keep your microphones muted unless you are speaking. Finally, if Members have documents they want to submit for the record, please email them to the Committee Clerk, whose email address was circulated prior to the hearing.

Good morning, and thank you to all of our witnesses who are joining us virtually to discuss the importance of scientific computing at the Department of Energy (DOE). This hearing is one of a series on research and development (R&D) activities sponsored by the Department of Energy's Office of Science. Today, we will be examining the current status and needs of DOE's scientific computing programs, as well as the research, development, and workforce training required to ensure that DOE and the Nation maintains its leadership in this crucial area.

Stewardship of DOE's scientific computing ecosystem is led by the Office of Science's Advanced Scientific Computing Research program, or ASCR. ASCR is also DOE's main sponsor of research in foundational areas such as applied mathematics and computer science. This year, ASCR was funded at just over $1 billion, about 1/7 of the total Office of Science budget.

DOE possesses some of the most powerful supercomputers in existence. It will deploy the Nation's first exascale system this year, signaling an exciting new era in the field of scientific computing. Housed at several national laboratories, DOE's supercomputers help researchers analyze huge data sets and test complex computational models, greatly accelerating the pace of discovery in the design of life-saving medical treatments, advanced manufacturing, and the prediction of climate systems, among many other fields of research. DOE's supercomputing ecosystem serves as a critical resource for academic and industry users from the U.S. and around the world. I am looking forward to discussing with our witnesses the real-world applications of these incredible systems, and how Congress can ensure that they are continuously maintained and improved.

It is also critically important for DOE to support research that will lay the groundwork for future computing capabilities. We are fast approaching the point at which the computing architectures we have relied upon for decades will reach their physical and economic limitations. Therefore, ASCR must continue to invest in the applied mathematics, computer science, and the game-changing technology development activities that will enable powerful new paradigms like quantum computing.

As we craft a forward-looking Office of Science authorization bill, I will be looking to our witnesses for insights into how we in Congress can ensure that these activities are robustly supported. As we will explore today, scientific computing holds tremendous promise for accelerating scientific discovery. But we need to use these capabilities responsibly, ethically, and to advance the public good. For example, as computing and artificial intelligence become more powerful, we must ensure that algorithms are designed to protect people's privacy and eradicate bias. We must also stop these tools from fortifying the structures of systematic racism, as we have seen happen with things like predictive policing and facial recognition technology. This will only become more important as DOE's supercomputing capabilities are used to process, analyze, and store sensitive information, such as biomedical datasets.

Let's also discuss how to retain a strong role for the public sector here, to fully tap into computing's potential to help solve humanity's most pressing problems, from curing diseases to addressing the climate emergency. And let's involve the public, especially marginalized communities, in shaping the development and aims of new technologies like these so that all can share in the benefits equally. As you will hear from our witnesses today, we need to pursue an agenda of scientific computing for the people.

Finally, as I have said before, research and infrastructure funding represent just one piece of the puzzle. We need a skilled and diverse workforce to maintain the vitality of DOE's scientific computing ecosystem long into the future. I am particularly interested in leveraging programs such as the Computational Science Graduate Fellowship to forge closer connections between the Department and minority-serving institutions. We can all agree on the need for greater diversity, equity, and inclusion across our research enterprise.

I want to again thank our excellent panel of witnesses assembled today, and I look forward to hearing your testimony.

[The prepared statement of Chairman Bowman follows:]

Good morning, and thank you to all of our witnesses who are joining us virtually today to discuss the importance of scientific computing at the Department of Energy.

This hearing is one of a series on research and development activities sponsored by the DOE's Office of Science. Today, we will be examining the current status and needs of DOE's scientific computing programs as well as the research, development, and workforce training required to ensure that DOE, and the nation, maintains its leadership in this crucial area.

Stewardship of DOE's scientific computing ecosystem is led by the Office of Science's Advanced Scientific Computing Research program, or ASCR. ASCR is also DOE's main sponsor of research in foundational areas such as applied mathematics and computer science. This year, ASCR was funded at just over a billion dollars, about one-seventh of the total Office of Science budget.

DOE possesses some of the most powerful supercomputers in existence. It will deploy the nation's first exascale system this year, signaling an exciting new era in the field of scientific computing. Housed at several national laboratories, DOE's supercomputers help researchers analyze huge data sets and test complex computational models, greatly accelerating the pace of discovery in the design of life-saving medical treatments, advanced manufacturing, and the prediction of climate systems, among many other fields of research. DOE's supercomputing ecosystem serves as a critical resource for academic and industry users from the U.S. and around the world. I am looking forward to discussing with our witnesses the real-world applications of these incredible systems, and how Congress can ensure that they are continuously maintained and improved.

It is also critically important for DOE to support research that will lay the groundwork for future computing capabilities. We are fast approaching the point at which the computing architectures we have relied upon for decades will reach their physical and economic limitations. Therefore, ASCR must continue to invest in the applied mathematics, computer science, and the game-changing technology development activities that will enable powerful new paradigms like quantum computing. As we craft a forward-looking Office of Science authorization bill, I will be looking to our witnesses for insights into how we in Congress can ensure that these activities are robustly supported.

As we will explore today, scientific computing holds tremendous promise for accelerating scientific discovery. But we need to use these capabilities responsibly, ethically, and to advance the public good. For example, as computing and artificial intelligence become more powerful, we must ensure that algorithms are designed to protect people's privacy and eradicate bias. We must also stop these tools from fortifying the structures of systemic racism, as we have seen happen with things like predictive policing and facial recognition technology. This will only become more important as DOE's supercomputing capabilities are used to process, analyze, and store sensitive information, such as biomedical datasets.

Let's also discuss how to retain a strong role for the public sector here, to fully tap into computing's potential to help solve humanity's most pressing problems - from curing diseases to addressing the climate emergency. And let's involve the public, especially marginalized communities, in shaping the development and aims of new technologies like these - so that all can share in the benefits equally. As you will hear from one our witnesses today, we need to pursue an agenda of scientific computing for the people.

Finally, as I have said before, research and infrastructure funding represent just one piece of the puzzle. We need a skilled and diverse workforce to maintain the vitality of DOE's scientific computing ecosystem long into the future. I am particularly interested in leveraging programs such as the Computational Science Graduate Fellowship to forge closer connections between the Department and Minority-Serving Institutions. We can all agree on the need for greater diversity, equity, and inclusion across our research enterprise.

I want to again thank our excellent panel of witnesses assembled today, and I look forward to hearing your testimony. With that, I yield back.

Chairman BOWMAN. Finally, I want to note that it is a busy day on the Hill, and I may have to step out briefly to ask questions on another Committee.

With that, I now recognize Mr. Weber for an opening statement.

Mr. WEBER. Well, thank you, Mr. Chair, and I'll be glad to, you know, conduct the hearing while you're gone.

I do want to thank you for hosting this hearing and to our esteemed witness panel for being here this afternoon or technically, I guess, this morning. I'm excited to hear about the critical advanced scientific computing research and development activities being carried out through the DOE, Department of Energy's Office of Science.

The Advanced Scientific Computing Research program, ASCR as you referred to it, is one that enjoys bipartisan support as a priority within the Office of Science. For the past 30 years, researchers within this program have led advances in mathematics and computing that form the foundation for those complex models and simulations. These developments, in turn, have translated to increased knowledge and understanding of everything from bioenergy and climate change to Alzheimer's disease and health models.

Today, ASCR hosts some of the world's most powerful supercomputers and a high-speed network that moves enormous volumes of scientific data at light speed. In the rapidly evolving fields of quantum computing and artificial intelligence, ASCR is dedicated to maintaining U.S. competitiveness and leadership. The program also supports DOE's goal of completing the world's first exascale computing system this year and a second system within the next year. As our competitors race to develop exascale systems on their own, DOE's strong support of advanced computing research within ASCR is essential to maintaining U.S. leadership in this field.

And it's more than just hardware that needs additional focus, I might add. We need significant modifications to today's tools and techniques to deliver on the promise of high-performance computing. Researchers are in need of a new suite of software tools, programming models, and applications to enable effective use of exascale systems. Without software and application R&D, we will simply have high-powered machines collecting dust.

Additionally, in order to fully and effectively support innovation in next-generation science, DOE must also support and also encourage cross-cutting research initiatives within the Department, as well as other Federal agencies. Within the Office of Science alone, ASCR resources and capabilities can be used to drive innovation in computational chemistry and nanomaterials for energy applications, improve simulations of fusion energy reactors, and enhance our ability to predict changes in the global climate with next-generation Earth system models.

Other Federal agencies could also capitalize on these unique, world-leading resources. As authorized by the *Energy Act of 2020*, the Department of Veterans Affairs (VA) is partnering with DOE to use high-performance computing in analyzing massive amounts of health data. This data analysis will help the VA better understand diseases and improve veterans' overall quality of life.

We should seek to build upon and expand partnerships like this so that the entire Federal Government benefits from ASCR's tools, as well as its technologies. At the end of the day, we're all supporting one thing: U.S. leadership in science, technology, and innovation. There is no Federal entity in a better position to lead this change than DOE's Office of Science.

That's why I am pleased that we are very close to finalizing legislation that provides strong support and long-term guidance for the Office of Science. We're making sure that rubber meets the road, and that the U.S. research enterprise is well-equipped with all available resources to successfully overcome the generational challenges they face.

I want to again thank my colleagues for their bipartisan outreach and collaboration. And I want to thank the witnesses for offering and bringing their effort, their input.

Thank you, Mr. Chairman, and I yield back.

[The prepared statement of Mr. Weber follows:]

Thank you, Chairman Bowman, for hosting this hearing, and thank you to our esteemed witness panel for being here this afternoon. I am excited to hear about the critical advanced scientific computing research and development activities being carried out through the Department of Energy's (DOE) Office of Science.

The Advanced Scientific Computing Research Program, or ASCR program, is one that enjoys bipartisan support as a priority within the Office of Science. For the past thirty years, researchers within this program have led advances in mathematics and computing that form the foundation for complex models and simulations. These developments, in turn, have translated to increased knowledge and understanding of everything from bioenergy and climate change to Alzheimer's disease and health models.

Today, ASCR hosts some of the world's most powerful supercomputers and a high-speed network that moves enormous volumes of scientific data at light speed. In the rapidly evolving fields of quantum computing and artificial intelligence, ASCR is dedicated to maintaining U.S. competitiveness and leadership. The program also supports DOE's goal of completing the world's first exascale computing system this year and a second system within the next year.

As our competitors race to develop exascale systems of their own, DOE's strong support of advanced computing research within ASCR is essential to maintaining U.S. leadership in this field. And it's more than just hardware that needs additional focus.

We need significant modifications to today's tools and techniques to deliver on the promise of high-performance computing. Researchers are in need of a new suite of software tools, programming models, and applications to enable effective use of exascale systems. Without software and application R&D, we will simply have high-powered machines collecting dust.

Additionally, in order to fully and effectively support innovation in next- generation science, DOE must also encourage cross-cutting research initiatives within the Department and with other Federal agencies. Within the Office of Science alone, ASCR resources and capabilities can be used to drive innovation in computational chemistry and nanomaterials for energy applications, improve simulations of fusion energy reactors, and enhance our ability to predict changes in the global climate with next generation Earth System Models.

Other federal agencies could also capitalize on these unique, world-leading resources. As authorized by the *Energy Act of 2020*, the Department of Veterans Affairs is partnering with DOE to use high performance computing in analyzing massive amounts of health data. This data analysis will help the VA better understand diseases and improve veterans' overall quality of life.

We should seek to build upon and expand partnerships like this so that the entire federal government benefits from ASCR's tools and technologies. At the end of the day, we are all supporting one thing: U.S. leadership in science, technology, and innovation. There is no federal entity in a better position to lead this charge than DOE's Office of Science.

That is why I am pleased we are very close to finalizing legislation that provides strong support and long-term guidance for the Office of Science. We are making sure rubber meets the road, and that the U.S. research enterprise is equipped with all available resources to successfully overcome the generational challenges they face.

I want to again thank my colleagues for their bipartisan outreach and collaboration. And I want to thank the witnesses for offering their input on our efforts. Thank you, Mr. Chairman, and I yield back the balance of my time.

Chairman BOWMAN. Thank you so much, Mr. Weber.

The Chair now recognizes the Chairwoman of the Full Committee, Ms. Johnson, for an opening statement.

Chairwoman JOHNSON. Thank you very much, Mr. Bowman, for holding this hearing, and Ranking Member Weber. And I also want to thank our witnesses for your participation, and I have enjoyed reading your thoughtful written testimony.

The Department of Energy has long been a leader in advancing new energy technologies, as well as the fundamental and the foundational sciences of physics, chemistry, engineering and math and computational science that support energy innovation. High-performance computing, or supercomputing, is one area the Department has led for decades, and DOE shows no signs of slowing down. The Department currently stewards two of the top three fastest supercomputers in the world. And as we will learn more about from our witnesses here today, the United States is on track to finish building the first exascale computer in the world this year. These systems serve as critical resources for academic and industrial users and are a key component of our economic competitiveness, scientific leadership, and national security.

In the past, high-performance computers were needed almost solely for specialized scientific and engineering applications. Now, as we enter the world where thousands of devices all around us are generating millions of bytes of data every minute, high-performance computers can be used to fundamentally improve our quality of life. Public policies play a critical role in supporting the advancement of these capabilities and enabling our society and economy to directly benefit from them. Additional Federal investments in high-performance computing will enable the development of new industries, grow our technology economy, and advance our technological leadership internationally.

All that said, as we continue to support the development and use of these breakthrough technologies, we almost—we must also do everything we can to ensure that we are doing this in a responsible and ethical manner even in the face of competition from our adversaries.

I thank you again for being here, and I look forward to this important discussion today, and I yield back.

[The prepared statement of Chairwoman Johnson follows:]

Thank you, Chairman Bowman, for holding this hearing, and I also want to thank this excellent panel of witnesses for your participation and thoughtful written testimony.

The Department of Energy has long been a leader in advancing new energy technologies, as well as the foundational sciences of physics, chemistry, engineering, mathematics, and computational science that support energy innovation.

High performance computing, or supercomputing, is one area the Department has led in for decades, and DOE shows no signs of slowing down. The Department currently stewards two of the top three fastest supercomputers in the world. And as we will learn more about from our witnesses here today, the United States is on track to finish building the first exascale computer in the world this year.

These systems serve as critical resources for academic and industrial users, and are a key component of our economic competitiveness, scientific leadership, and national security.

In the past, high performance computers were needed almost solely for specialized scientific and engineering applications. Now, as we enter a world where thousands of devices all around us are generating millions of bytes of data every minute, high performance computing can be used to fundamentally improve our quality of life.

Public policies play a critical role in supporting the advancement of these capabilities, and in enabling our society and economy to directly benefit from them. Additional federal investments in high performance computing will enable the development of new industries, grow our technology economy, and advance our technological leadership internationally.

All that said, as we continue to support the development and use of these breakthrough technologies, we must also do everything we can to ensure that we are doing this in a responsible and ethical manner. Even in the face of competition from our adversaries.

Thank you all again for being here, and I look forward to this important discussion today. With that I yield back.

Chairman BOWMAN. Thank you so much for your opening statement, Madam Chairwoman.

The Chair now recognizes the Ranking Member of the Full Committee, Mr. Lucas, for an opening statement.

Mr. LUCAS. Thank you, Chairman Bowman, for hosting this hearing, and thank you to all our witnesses for being with us this afternoon.

Earlier this month, the Energy Subcommittee held a hearing on the Department of Energy's Office of Science, which emphasized the essential role of DOE in our Federal research enterprise and highlighted our shared support of these programs.

Today, we have an opportunity to examine the activities of another Office of Science program in Advanced Scientific Computing Research, or ASCR. Advanced computing research and infrastructure is the backbone of scientific discovery, not just at the Department of Energy but at U.S. research institutions nationwide. Through the ASCR program, DOE supports the development of tools and technologies in high-performance computing, applied mathematics, advanced networks, data analytics, and next-generation computing initiatives. It also hosts some of the most advanced computing resources in the world at its national laboratories.

There is a great potential for Federal agencies and U.S. industry partners to leverage ASCR's unique computing resources. With adequate support, DOE's program will revolutionize our relationship with advanced technology and our capacity for scientific progress. This work is vital to our clean energy economy, our national security, and our leadership in science and technology.

Yet we know that our international competitors like China are outpacing us in basic research investment and are closing the gap in key computing focus areas like artificial intelligence and quantum sciences. Expanding our capacities in these fields requires a strategic effort with strong Federal investment and active public-private partnerships.

That's why in this Congress I've introduced legislation to address those challenges. My bill, the *Securing American Leadership in Science and Technology Act, SALSTA*, roughly doubles funding for ASCR over the next 10 years.

Another bill I introduced, the *Quantum User Expansion for Science and Technology Act, QUEST*, establishes a program at the Department of Energy to expand public-private partnerships for quantum resource use and encourage greater participation in the development of quantum information sciences (QIS).

Mr. Chairman, at this time, I'd like to ask unanimous consent to submit for the record a letter from the Quantum Industry Coalition on the need to maximize the value of the U.S. quantum industry and the role that DOE and its national laboratories can play in this high-priority work.

Chairman BOWMAN. Without objection.

Mr. LUCAS. Thank you, Mr. Chairman.

I'm also proud to join my colleague and the Ranking Member of the Investigations and Oversight Subcommittee, Jay Obernolte, on a bill to strengthen the other high-priority computing research program carried out at the Department. This week, Representative Obernolte introduced the *Next Generation Computing Research and Development Act*, which authorizes various DOE advanced scientific computer programs. These will support beyond excellent energy computing, computing workforce development, and applied mathematics and software development activities. This bill, along with the *QUEST Act* and *SALSTA*, is an important step to move forward in improving our Nation's global standing in science and technology.

We know that maintaining U.S. leadership will require a shared commitment to prioritize DOE and its Office of Science, and nowhere is this clearer than in the advanced computing space. The United States relies on computing capacities that only the Department of Energy can provide. We know that the Nation who takes the lead in advanced computing will set the stage for the next generation of technologies and technology standards. We cannot afford to fall behind in this race.

Last week, I was encouraged by the progress made by our friends in the Senate to recognize the important role the Department of Energy plays in advancing U.S. innovation. But DOE and the national labs shouldn't be an afterthought when we consider the U.S.

research enterprise. They're integral to our scientific progress. That's why Chairman—woman Johnson and I have been working on bipartisan Office of Science legislation that will make a strong commitment to the Department of Energy and its work, including successful programs like ASCR.

This legislation to support research at the Department of Energy will go hand-in-glove with the *NSF (National Science Foundation) For the Future Act*, which supports basic research, STEM education, and technology transfer at the National Science Foundation. Together, these research bills will solidify the long-term stability of our international leadership in science.

I once again want to thank our witnesses for being here today. I look forward to a productive discussion. Thank you, Chairman Bowman, and I yield back the balance of my time.

[The prepared statement of Mr. Lucas follows:]

Thank you, Chairman Bowman for hosting this hearing, and thank you to all our witnesses for being with us this afternoon.

Earlier this month, the Energy Subcommittee held a hearing on the Department of Energy's Office of Science which emphasized the essential role of DOE in our federal research enterprise and highlighted our shared support of its programs.

Today, we have an opportunity to examine the activities of another Office of Science program in Advanced Scientific Computing Research, or ASCR. Advanced computing research and infrastructure is the backbone of scientific discovery, not just at the Department of Energy but at U.S. research institutions nationwide. Through the ASCR ("Oscar") program, DOE supports the development of tools and technologies in high performance computing, applied mathematics, advanced networks, data analytics, and next-generation computing initiatives. It also hosts some of the most advanced computing resources in the world at its national laboratories.

There is great potential for federal agencies and U.S. industry partners to leverage ASCR's unique computing resources. With adequate support, DOE's program will revolutionize our relationship with advanced technology and our capacity for scientific progress. This work is vital to our clean energy economy, our national security, and our leadership in science and technology.

Yet we know that our international competitors like China are outpacing us in basic research investment and are closing the gap in key computing focus areas like artificial intelligence and quantum sciences. Expanding our capacities in these fields requires a strategic effort with strong federal investment and active public-private partnerships.

That's why, this Congress, I've introduced legislation to address these challenges. My bill, the *Securing American Leadership in Science and Technology Act (SALSTA)*, roughly doubles funding for ASCR over ten years. Another bill I introduced, the *Quantum User Expansion for Science and Technology Act (QUEST) Act*, establishes a program at the Department of Energy to expand public-private partnerships for quantum resource use and encourage greater participation in the development of quantum information sciences.

Mr. Chairman, at this time I'd like to ask unanimous consent to submit for the record, a letter from the Quantum Industry Coalition, on the need to maximize the value of the

U.S. quantum industry, and the role that DOE and its national laboratories can play in this high-priority work.

I'm also proud to join my colleague and Ranking Member of the Investigations and Oversight Subcommittee, Jay Obernolte, on a bill to strengthen other high-priority computing research carried out by the Department.

This week, Representative Obernolte introduced the *Next Generation Computing Research and Development Act*, which authorizes various DOE advanced scientific computing programs. These will support beyond-exascale and energy efficient computing, computing workforce development, and applied mathematics and software development activities. This bill, along with the *QUEST Act* and *SALSTA*, is an important step forward in improving our nation's global standing in science and technology.

We know that maintaining U.S. leadership will require a shared commitment to prioritize DOE and its Office of Science. And nowhere is this clearer than in the advanced computing space.The U.S. relies on computing capabilities that only the Department of Energy can provide. We know that the nation who takes the lead in advanced computing will set the stage for the next generation of technologies and technology standards. We cannot afford to fall behind in this race.

Last week, I was encouraged by the progress made by my friends in the Senate to recognize the important role the Department of Energy plays in advancing U.S. innovation. But DOE and the National Labs shouldn't be an afterthought when we consider the U.S. research enterprise. They're integral to our scientific progress. That's why Chairwoman Johnson and I have been working on bipartisan Office of Science legislation that will make a strong commitment to the Department of Energy and its work-including successful programs like ASCR.

This legislation to support research at the Department of Energy will go hand-in-glove with the *NSF For the Future Act*, which supports basic research, STEM education, and technology transfer at the National Science Foundation. Together, these research bills will solidify the long-term stability of our international leadership in science.

I once again want to thank our witnesses for being here today. I look forward to a productive discussion. Thank you Chairman Bowman and I yield back the balance of my time.

Chairman BOWMAN. Thank you, Mr. Lucas, for your opening statement.

If there are Members who wish to submit additional opening statements, your statements will be added to the record at this point.

At this time, I would like to introduce our witnesses. First, Dr. J. Stephen Binkley is the Acting Director and Principal Deputy Director in the Office of Science at the U.S. Department of Energy. Prior to his experience in various leadership positions in DOE, Dr. Binkley has held senior positions at DOE's Sandia National Laboratories and the Department of Homeland Security. He has conducted research in theoretical chemistry, materials science, computer science, applied mathematics, and microelectronics.

Next, Dr. Georgia Tourassi is the Director of the National Center for Computational Sciences at the Oak Ridge National Laboratory. She also holds appointments as an Adjunct Professor of Radiology at Duke University and as a Professor of the Bredesen Center Data Science Program at the University of Tennessee at Knoxville.

Next, Dr. Karen Willcox is Director of the Oden Institute of Computational Engineering and Sciences, Associate Vice President for Research, and Professor of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin. She holds the W.A. "Tex" Moncrief, Jr. Chair in simulation-based engineering and sciences and the Peter O'Donnell, Jr. Centennial Chair in Computing Systems.

Dr. Christopher Monroe is Co-Founder and Chief Scientist at IonQ Inc. and the Gilhuly Family Distinguished Presidential Professor of Electrical and Computer Engineering and Physics at Duke University. He is an atomic physicist and quantum engineer specializing in the isolation of individual atoms as the core of a quantum computer.

Last but certainly not least, Dr. Seny Kamara is an Associate Professor of Computer Science at Brown University where he co-directs Brown's Computing for the People Project and the Encrypted Systems Lab. He is also affiliated with Brown's Center for Human Rights and Humanitarian Studies, the Data Science Initiative, and the Policy Lab. Kamara is a principal scientist at MongoDB, a company that provides one of the most widely used platforms to store and process data.

Thank you all for joining us today. As our witnesses should know, you will each have 5 minutes for your spoken testimony. Your written testimony will be included in the record for the hearing. When you all have completed your spoken testimony, we will begin with questions. Each Member will have 5 minutes to question the panel. We will start with Dr. Blinky—Blinkley, excuse me. Dr. Binkley, please begin.

TESTIMONY OF DR. J. STEPHEN BINKLEY, PRINCIPAL DEPUTY DIRECTOR, OFFICE OF SCIENCE AT THE DEPARTMENT OF ENERGY

Dr. BINKLEY. OK. Thank you, Chairman Bowman and Ranking Member Weber. I'm pleased to come here before you today to discuss the scientific computing capabilities of the Department of Energy, including the forthcoming exascale systems.

DOE computing traces its roots back to the Manhattan Project where extensive use was made of computers. During the 1950's, John von Neumann, the pioneer in computing, advocated for a program that would advance computer development. Over the years, ever more powerful computing capabilities were developed at the national laboratories beginning with the Lawrence Livermore and Los Alamos National Laboratories.

DOE and its predecessor agencies have supported applied mathematics and computer science, along with major investments in computer hardware and computational science that have been a major driver of progress in high-performance computing, spurring the U.S. computing industry forward. DOE computing applications have expanded from their original national defense focus to a broad portfolio of scientific research and significant use by industry beginning with the establishment of the leadership computing facilities at Argonne and Oak Ridge National Laboratories in 2004.

Today, DOE computing is a partnership between the National Nuclear Security Administration (NNSA) and the Department's Office of Science. Our two organizations are working hand-in-hand to advance high-performance computing, including the exascale computing project.

The strategic importance of high-performance computing has grown enormously. High-performance computing has become an essential pillar not just of America's national security but also of our leadership in science. DOE's supercomputing has brought major computational and driven advances in a wide range of fields such as climate science, fusion energy, and high energy and nuclear physics, materials science, chemistry, particle accelerator design, and biology, to name a few.

Over the last 6 years, we have been very focused on achieving exascale computing. The first exascale computing system is scheduled for delivery at the Oak Ridge National Laboratory to be complete by October of this year. The second system will go to Argonne National Laboratory in 2022, and a third system to Lawrence Livermore National Laboratory in 2023.

Exascale has the capability to deepen our understanding of climate change and hasten the development of clean energy. Partnerships between the Office of Science and NNSA with major computing and microelectronic vendors have been key in the development of exascale. A series of five partnership programs have brought DOE-supported researchers to work hand-in-glove with U.S. high-performance computing vendors, including AMD (Advanced Micro Devices), Cray, IBM (International Business Machines), Intel, Nvidia, and HPE (Hewlett Packard Enterprise) to overcome the key technical hurdles in exascale. In total, DOE has invested $460 million in this effort alone, matched by at least an additional $307 million contributed by industry.

Current and planned upgrades to Office of Science scientific user facilities, including light sources, neutron scattering sources, nanoscale, and genomic facilities will bring more sophisticated and precise observations and vastly larger data outputs. Artificial intelligence and machine learning will play a key role in this.

AI also holds the promise of more sophisticated and autonomous facility operations. It has the potential to monitor observations and adjust instrument operations in real time to further enhance efficiency and utilization of the facilities. DOE's ESnet provides ultra-high broadband connectivity across the DOE laboratories as connectivity will be increasingly vital as facility operations are controlled computationally.

We are looking forward and beyond exascale to new frontiers such as quantum information science. Leadership in science remains indispensable to our high-performance—to the country's prosperity, and high-performance computing is key. Continued stewardship and development of the skilled HPC workforce is essential. Our Computational Science Graduate Fellowship program is one such activity. Since its establishment in 1991, the program has sponsored over 450 fellows from more than 60 universities. DOE's response to the COVID–19 pandemic demonstrates the enormous value of DOE's high-performance computational research resources.

In summary, opportunities for accelerated scientific discovery will be enabled by current—the current era of high-performance computing marked by the advent of exascale systems and the rapid development of AI and machine learning.

And I'll end there.

[The prepared statement of Dr. Binkley follows:]

STATEMENT BY

J. STEPHEN BINKLEY

ACTING DIRECTOR, OFFICE OF SCIENCE,

U.S. DEPARTMENT OF ENERGY

BEFORE THE

HOUSE SCIENCE, SPACE AND TECHNOLOGY COMMITTEE,

SUBCOMMITTEE ON ENERGY,

ON

ACCELERATING DISCOVERY: THE

FUTURE OF SCIENTIFIC COMPUTING AT THE DEPARTMENT OF ENERGY

MAY 19, 2021

Thank you, Chairman Bowman and Ranking Member Weber. I am pleased to discuss the scientific computing capabilities of the Department of Energy (DOE), including the forthcoming exascale systems, the implications of these capabilities for other scientific disciplines, and their relevance to pressing societal challenges.

Introduction

DOE computing has deep historical roots, going back to the Manhattan Project, where computers were extensively used, including one of the earliest IBM punch card systems. During the 1950s, John von Neumann, a commissioner of the Atomic Energy Commission (AEC) and a pioneer in computing, advocated for a Mathematics program that would advance computer development. This started a succession of investments over the years into ever more powerful computing capabilities at the National Laboratories, beginning with what are now two National Nuclear Security Administration labs, Lawrence Livermore and Los Alamos National Laboratories.

With the eventual dissolution of the AEC and the establishment of the Department of Energy in 1977, DOE became the heir and steward of these capabilities. DOE and its predecessor agencies support for applied mathematics and computer science, along with major investments in computer hardware and computational science, have driven progress in high-performance computing (HPC) in the United States, a major force spurring U.S. computing industry.

The scope of DOE computing applications has expanded from their original national defense focus to encompass an increasingly broad portfolio of scientific research and significant use by industry, especially with the establishment of the Leadership Computing Facilities at Argonne and Oak Ridge National Laboratories in 2004.

Today, DOE computing is a partnership between the National Nuclear Security Administration and the Office of Advanced Scientific Computing Research within the Department's Office of Science. Our two organizations work hand-in-hand to advance high-performance computing, including our partnership in the Exascale Computing Project (ECP).

Strategic Importance

Over the decades, the strategic importance of high-performance computing has grown enormously, and it is fair to say that a nation's capabilities in high-performance computing are one of the most important measures of its overall competitive standing in the global economy.

Thus, high-performance computing has become an essential pillar not just of America's national security but also of our leadership in science and—increasingly—of our national economic competitiveness.

A key measure of this competitiveness is the power of hardware, the raw capabilities of our most capable systems, and this informs the deployment of exascale machines. Additionally, an entire research and development ecosystem both contributes to, and benefits from, these capabilities. The defining strength of DOE computing is the vitality of the total effort involving a whole of government interagency approach and the multiple partnerships among the computing and microelectronics industries, the national laboratories, universities, and industry at large.

A key term is "co-design," meaning a highly integrated and cooperative effort where end users and application requirements contribute to the design of systems, and the systems in turn enable scientific discovery through advances in hardware, software, algorithms, and applications. A recent report by the Advanced Scientific Computing Advisory Committee speaks of "the enviable culture of co-design teams consisting of scientific users, instrument providers, theoretical scientists, mathematicians and computer scientists," which is an apt characterization of the approach.

Leading with the Science

When DOE's Office of Science established the Office of Advanced Scientific Computing, or ASCR, in 1999, their first major initiative was to establish Scientific Discovery through Advanced Computing, or SciDAC, a program that thrives to this day. The aim was to make scientific research and discovery the driver of the computational effort, to lead with the science and thereby let the quest for scientific discovery drive progress in computing.

SciDAC supports multi-disciplinary partnerships of computer scientists and applied mathematicians with domain scientists supported by the Office of Science, which has brought major computational-driven advances in a range of fields, from climate science, fusion energy, and high energy and nuclear physics, to materials science, chemistry, particle accelerator design, and biology, to name just some areas. SciDAC has also spurred major progress in software development and algorithm design.

The program encompasses all major Office of Science program offices along with DOE's Office of Nuclear Energy.

To mention just a few notable achievements:

- Computational scientists at Lawrence Berkeley National Laboratory and climate scientists at the Los Alamos National Laboratory and the University of Bristol have teamed to create the most accurate model for understanding the retreating Greenland and Antarctic ice sheets. By providing sub-kilometer resolution of key areas and less computationally demanding resolution of less important regions, it provides a much more accurate picture of the dissipating mass, which threatens to contribute to significant sea-level rise over the coming century.
- Fusion holds the promise of abundant, carbon dioxide-free energy. Progress toward fusion will depend on experiments run on ITER, currently under construction in France. ITER success requires understanding the interactions between the burning plasma contained in the reactor and the reactor walls. These involve extraordinarily high temperatures—millions of degrees in the case of the plasma and up to a thousand degrees for reactor walls—that are impossible to reproduce in a laboratory and difficult to simulate accurately at less than exascale capabilities. Scientists as Oak Ridge National Laboratory and University of Tennessee have developed a new approach using machine learning that accurately simulates the interaction between the plasma and the tungsten coating of the reactor walls—using a fraction of computing power normally required.
- The nature of supernova explosions is one of the enduring mysteries of astrophysics. With support by SciDAC, scientists led by Princeton University used Argonne National Laboratory's Cray XC40 system, called Theta, created some of the first three-dimension simulations of supernova explosions. The development of fully three-dimensional simulations is opening the way to more detailed comparison between simulation and observed events, facilitating a new understanding of the mechanism by which supernovae explode and collapse into neutron stars.

SciDAC continues to be an important force for progress in both domain sciences and computational science and engineering.

The Advent of Exascale

Exascale computing will provide the capability to tackle even more complex problems. Exascale has the capacity to deepen our understanding of climate change and hasten the development of clean energy. It will aid in the development of advanced manufacturing and holds major promise in the area of cancer research. The potential applications are too numerous to name.

The first exascale system, a Cray system called Frontier, is scheduled for delivery at the Oak Ridge Leadership Computing Facility (OLCF) beginning in July, with deployment completion expected in October. The Argonne Leadership Computing Facility's (ALCF) Aurora system is expected to be deployed in 2022. Lawrence Livermore National Laboratory is expecting delivery of the Department's third exascale machine, El Capitan, in 2023.

The drive toward exascale has been a major co-design effort, on two levels.

First, there has been a co-design partnership of Office of Science and NNSA with major computing and microelectronic vendors, beginning in Fiscal Year 2012. Through a series of five "forward" programs—FastForward, DesignForward 1 and 2, Fastforward 2, and the current ECP PathForward program—DOE-supported researchers have worked with vendors including AMD, Cray, IBM, Intel, NVDIA, and HPE to overcome a series of key technical hurdles, including surmounting power and speed limitations through the deployment of Graphic Processing Units (GPUs), improving interconnect, and developing new approaches to memory. In total DOE, through ASCR in the Office of Science and NNSA combined, has invested $460 million in this effort, while industry contributed at least an additional $307 million toward the research. DOE invited multiple agencies to participate in the project reviews as observers along the way. This has significantly disseminated the knowledge gained through the ECP project development process to other agencies.

Many of the challenges to be overcome stem from the impending end of Moore's Law. Moore's Law is shorthand for the progressive process of microchip miniaturization over the past several decades that has led processor speeds to double approximately every two years. As this process of miniaturization reaches its limits, and feature sizes on chips narrow toward the width of atoms, significant creativity is required to address these limitations. Two successful recent examples from the ECP include the partnership with AMD to develop architectures based on AMD's chiplets, which are separate pieces of silicon within a single processor package; and a partnership with Intel to develop interconnect based on Silicon Photonics, which communicate via photons rather than electrons—with potential increases in speed and major power savings.

These research partnerships with DOE have been critical in enabling the U.S. computer industry to enter the exascale era.

A second, parallel and overlapping, co-design effort encompasses scientific end-users. A key goal of this effort is to ensure that the community is ready with scientific applications to run as soon as exascale systems are deployed and available. These partnerships echo those of SciDAC, where applied mathematicians, computer scientists, and software developers—and often hardware designers—team with domain scientists to harness computation for research.

<u>Challenges and Opportunities</u>

At the threshold of exascale deployments, the scientific environment continues to evolve and to pose new challenges and opportunities. Current and planned upgrades to Office of Science user facilities—including light sources, neutron scattering sources, nanoscale and genomic facilities, and the supercomputers themselves—will bring more sophisticated and precise observations but also vastly larger data outputs. Means must be found to tame these mountains of data, process them, and extract what is meaningful. Artificial intelligence (AI) and machine learning (ML) are likely to play a key role here.

Data do not just pose a challenge; they also represent an enormous opportunity. In the coming years, we expect a great deal of scientific discovery to be heavily data-driven, meaning discovery will come from new methods of extracting insights from massive data sets, many of which exist today as still largely untapped resources. That is one reason why the Office of Science recently designated several important existing data sets as Publicly Reusable Research (PuRE) Data Resources, which are to be carefully curated and made available to the public at large on a more systematic basis. Included are data sets from the Materials Project, the Atmospheric Radiation Measurement program, the DOE Joint Genome Institute, the DOE Systems Biology Knowledgebase, the Particle Data Group, and the National Nuclear Data Center, with more resources likely to be added. The Office of Science is committed to careful management of these resources, stewarding them for the benefit of the community. The expectation is that the community will take the opportunity to mine these resources for new insights. AI and ML are likely to be at the center of these efforts.

AI also holds the promise of more sophisticated and autonomous facility operations. Currently operations at facilities such as x-ray light sources are heavily dependent on human monitoring and control, which involves a measure of trial and error. AI has the potential to monitor observations and adjust instrument operations in real time to ensure that observations capture what is needed and beam time use is optimized.

Already important progress is being made in this direction. The Office of Science's Energy Sciences Network, or ESnet, now connects the Linac Coherent Light Source at SLAC National Accelerator Laboratory with the Cori supercomputer at Lawrence Berkeley National Laboratory's National Energy Research Scientific Computing Center, providing real-time analysis ported to users' laptops, enabling on-the-fly adjustments to experiments based on computational results.

Increasing integration of AI into this process will require not only increased computational power but also enhanced networking. The Office of Science ESnet user facility provides ultra-high broadband connectivity to the DOE laboratories and research institutions across the U.S. and internationally. This connectivity will be increasingly vital as facility operations are controlled computationally, potentially by physically distant computational resources.

<u>Beyond Exascale</u>

Even as we stand on the threshold of the deployment of exascale systems, we are looking beyond exascale to new frontiers. Quantum Information Science (QIS) is one of the most important such frontiers. We have seen early progress in quantum networking. Last year, Brookhaven National Laboratory and Stony Brook University demonstrated a three-node network prototype over approximately 87 miles. Scientists from Argonne National Laboratory and the University of Chicago achieved photon entanglement over a network stretching 52 miles through the Chicago suburbs. These early results bode well.

Through the National Quantum Initiative Act, we established five National QIS Research Centers, led respectively by Argonne National Laboratory, Brookhaven National Laboratory, Fermi National Accelerator Laboratory, Lawrence Berkeley National Laboratory, and Oak Ridge National Laboratory. They are pursuing research in quantum networking, quantum sensing, and quantum computing and will position us well as quantum applications, and quantum computing in particular, become more capable.

In addition, we are supporting early investigations into the potential of neuromorphic computing, or computing architectures modeled on the neural structure of the human brain.

However, quantum computing at any scale is still well over the horizon, and neuromorphic computing even further out.

The Imperative of Co-Design

That is why it is essential that we continue our co-design efforts with vendors and with the microelectronics industry in the context of current technologies. We must continue to promote and support the development of new approaches, new materials, and new architectures to ensure that our nation is invested in the technologies that our researchers will need to remain on the cutting edge. Leadership in science remains indispensable to our security and prosperity, and today there is no leadership in science without leadership in high-performance computing. International competition is fierce. We need to be investing today in the immediate next-generation systems to stay in the game and to lead.

Federal investment is essential, since our activities in high-performance computing serve the national interest but do not in any sense drive the market. Absent Federal investment, vendors have little incentive to pursue these systems.

A skilled HPC workforce is also imperative. The Computational Science Graduate Fellowship, which is available, with varying eligibility requirements, to undergraduate seniors through the first year of Ph.D. study. Since its establishment in 1991, the program has sponsored over 450 fellows from more than 60 universities. Currently there are 89 fellows using high-performance computing in their research over a wide range of fields. The Computational Science Graduate Fellowship plays a major role in building and sustaining the nation's HPC workforce and is a Federal investment that thereby pays dividends.

DOE Scientific Computing in the Pandemic

As a final point, crises have a way of revealing the true value of things. The COVID-19 pandemic has made abundantly clear the enormous value of DOE's high-performance computational resources—for molecular modeling, drug screening, epidemiology, and a host of other applications in the war against the disease. Early in the pandemic, OLCF's Summit supercomputer screened thousands of compounds for potential effectiveness against the SARS-CoV-2 virus. As requests for computer time multiplied and other public and private resources were made available, DOE joined with 15 partner institutions in March 2020 to establish the COVID-19 High Performance Computing Consortium. IBM established the consortium website linking to a central gateway, based on the NSF-funded XSEDE portal, for submission of proposals, which were reviewed by a panel of experts and assigned to appropriate resources first in the United States, then across the globe. Eventually, the consortium counted 43 members with over 600 petaflops of computing capacity. At last count, as many as 175 proposals were received from researchers in 15 countries.

An indication of the power of HPC to accelerate discovery in the fight against the pandemic is provided by a study titled, "AI-Driven Multiscale Simulations Illuminate Mechanisms of SARS-CoV-2 Spike Dynamics," which was awarded a 2020 Special Gordon Bell Prize for COVID-related applications. As part of this research, using AI, and building on multiple experiments supported by NIH, NSF, as well as DOE, a team involving Argonne scientists found a way to significantly speed the screening of molecules for effectiveness against the virus—increasing the speed of screening by a factor of 50,000. Extending the approach, the team managed to create an automated pipeline capable of screening literally a billion drug candidates per day. The work was supported through the National Virtual Biotechnology Laboratory, a consortium of DOE national laboratories focused on response to COVID-19, with funding provided by the Coronavirus CARES Act. (Like the HPC Consortium, NVBL employs a single gateway for efficient expert review of proposals and assignment of facilities.)

This last example of accelerated discovery hints at the enormous promise of the current era of high-performance computing, an era marked by the advent of exascale systems, the rapid development of AI and ML, and the establishment of computation as a powerful tool of discovery on equal footing with theory and experiment. The United States stands poised to lead in this new era of computing and discovery, and our future security and prosperity will depend on it.

J. STEPHEN BINKLEY

U. S. Department of Energy

J. Stephen (Steve) Binkley is the Acting Director and Principal Deputy Director in the Office of Science (SC) at the U.S. Department of Energy (DOE). In this capacity, Dr. Binkley is the senior career science official in the Office of Science, which is third largest Federal sponsor of basic research in the United States, the primary supporter of the physical sciences in the U.S., and one of the premier science organizations in the world.

As the Principal Deputy Director of SC, Dr. Binkley serves as the principal overall advisor to the Director on all aspects of the Office of Science. Dr. Binkley determines the financial and personnel resources needed to achieve mission objectives and support mission operations; oversees and directs the internal organization, staffing, policies, and personnel authorities required to carry out the responsibilities of the organization, including the recruitment of senior managers and technical experts necessary to ensure the success of the programs. He ensures that program activities are strategically conceived and executed to maximize the benefit to organization, the Department, and the United States. Dr. Binkley also serves as the champion for crosscutting issues that affect more than one program office and special research initiatives of priority to Director and the Department leadership.

Dr. Binkley has held senior positions at Sandia National Laboratories, the Department of Homeland Security (DHS), and the Department of Energy. He has conducted research in theoretical chemistry, materials science, computer science, applied mathematics, and microelectronics. At Sandia, Dr. Binkley managed computer science, fundamental chemistry, combustion science, and chemically reacting flow organizations. He also has served as the manager for the Office of Science's Combustion Research Facility, at Sandia's Livermore, California location. Dr. Binkley managed Sandia's Office of Science Program, comprising activities in materials science, chemistry, geoscience, magnetic fusion energy, atmospheric measurement technology, and scientific computing at Sandia's New Mexico and California locations. He also managed activities in Sandia's national security program, including distributed information systems technology.

At DOE, Dr. Binkley served as a technical advisor to the Assistant Secretary of Defense Programs (subsequently the Deputy Administrator for Defense Programs after the establishment of the National Nuclear Security Administration). At DHS, Dr. Binkley served as the deputy director for technology within the DHS Operations Directorate, where he led and managed the development of systems for monitoring and disseminating situational awareness to federal, state, and local law-enforcement organizations and for coordinating emergency response activities. Returning to DOE in 2006, Dr. Binkley served as a senior technical advisor to the Under Secretary for Science and the Director of the Office of Science.

As head of SC's Office of Advanced Scientific Computing Research, Dr. Binkley served as one of the Associate Directors for the Office of Science, and was responsible for the overall management of the ASCR program including: strategic planning; budget formulation and execution; project management; program integration with other Office of Science activities and with the DOE technology offices; and interagency integration.

Chairman BOWMAN. Thank you so much, Dr. Binkley.
Dr. Tourassi, you're now recognized.

TESTIMONY OF DR. GEORGIA TOURASSI, DIRECTOR, NATIONAL CENTER FOR COMPUTATIONAL SCIENCES AT OAK RIDGE NATIONAL LABORATORY

Dr. TOURASSI. Chairman Bowman, Ranking Member Weber, Chairwoman Johnson, Ranking Member Lucas, and distinguished Members of the Committee, thank you for the opportunity to appear before you today. My name is Georgia Tourassi. I lead the Department of Energy's Oak Ridge Leadership Computing Facility, OLCF, at the Oak Ridge National Laboratory in Tennessee. I'm a biomedical engineer and a computational scientist by education and training.

High-performance computing has been the cornerstone for the Nation's scientific advancement, technology innovation, competitive advantage, and economic prosperity. Its impact on global competitiveness has long been embodied by the saying, "You must out-compute to outcompete." With advances in data technologies, machine learning, and AI, this saying can be amended to, "You must learn faster to outcompete."

OLCF has been a global leader in high-performance computing for nearly 30 years. Currently, OLCF hosts Summit, the Nation's most powerful supercomputer for open science. Summit is in high demand for modeling and simulation, data analytics, and AI to better understand climate change, develop new ways to produce clean energy, design advanced materials, advance public health, and overall push the frontiers of science. The request for time on Summit is up to five times more than the hours available.

Our facility is both deliberate and responsive to national needs. In the past year, I have experienced firsthand our staff's Herculean efforts to be both fast and offer world-leading computing resources and computational and data expertise in the fight against the coronavirus. Through the COVID–19 High-Performance Computing Consortium, Summit and our competent staff, using world-leading AI, helped accelerate discovery, understand the virus, and inform management of the pandemic response. I would like to thank Congress for the *CARES Act* funds OLCF received to augment Summit and help support the COVID–19 research community.

Now, in 2021, OLCF is at the brink of delivering the first exascale computer in the United States called Frontier. This supercomputer will perform calculations up to eight times faster than Summit and will keep the United States at the forefront on high-performance computing. To prepare, the DOE's Exascale Computing Project is developing critical applications across many scientific and technical disciplines to run on Frontier on day 1.

In addition, exascale will offer training opportunities to grow a more high-tech and computationally savvy workforce in our Nation. It is imperative for the United States to expand and enhance the national research computing ecosystem.

DOE has asked us to deliver Frontier 1 year earlier than planned, and we're focusing our efforts on meeting that schedule. Once the system is delivered, we will need to properly fund the operation and applications to solve complex real-world problems in partnership with leading research institutions, industry, and other Federal agencies.

We need to continuously invest in new technologies such as AI and accelerated computing methods to maintain our competitive advantage and ensure our global leadership. We need to make investments in a national data infrastructure that makes the most of our high-performance computing and national data assets. The COVID–19 pandemic demonstrated the vital importance of having established interagency programs and data integration ahead of these anticipated crises and the utility of high-performance computing and AI for rapid, complex, real-world data analysis.

The DOE leadership computing facilities are uniquely positioned to support and integrate its research infrastructure, combining our leadership computing with national experimental facilities and Federal data assets to deliver unprecedented technological, scientific, and economic advantages to the Nation.

We know high-performance computing is high on other nations' priorities. We know that China, Japan, and the European Union are all investing heavily in exascale computing, and this has implications for both our national security and our overall global competitiveness. We cannot afford to be left behind.

Thank you again for the opportunity to testify. I welcome your questions on this important topic.

[The prepared statement of Dr. Tourassi follows:]

Testimony of Dr. Georgia Tourassi
Director, National Center for Computational Sciences
Oak Ridge National Laboratory

Before the
Subcommittee on Energy
Committee on Science, Space, and Technology
U.S. House of Representatives

"Accelerating Discovery: the Future of Scientific Computing at the Department of Energy"
May 19, 2021

Chairman Bowman, Ranking Member Weber, and distinguished members of the Committee: Thank you for the opportunity to appear before you today. My name is Georgia Tourassi. I am Director of the National Center for Computational Sciences of the U.S. Department of Energy's Oak Ridge National Laboratory (ORNL) in Oak Ridge, Tennessee. I am a biomedical engineer and computational scientist by education and training. My research work lies at the intersection of high-performance computing (HPC), artificial intelligence (AI), and biomedicine. I am privileged to direct one of the Department of Energy's User Facilities, the Oak Ridge Leadership Computing Facility (OLCF), which hosts the nation's most powerful supercomputer for open science, Summit. It is an honor to provide this testimony on the transformative role of the Department of Energy's Leadership Computing User Facilities in spearheading a national high-performance computing and data infrastructure for accelerating research and development (R&D) giving the nation a significant advantage in the ever-changing scientific and economic landscape.

The views I offer in my statement are my own, shaped by my scientific experiences, my experience with inter-agency partnerships to advance respective missions, and my experience as Director of a DOE Leadership Computing Facility dedicated to open science.

INTRODUCTION

High-performance computing (or supercomputing) has been the cornerstone for research advancement, scientific innovation, competitive advantage, and economic prosperity[1]. The Department of Energy's Leadership Computing Facilities have a long history of enabling researchers to accelerate scientific discovery and deliver practical breakthroughs for some of the most computationally challenging problems. These research and development advances happen across many disciplines such as materials science, earth science, nuclear science, astrophysics, biology, and engineering to name a few. In addition, supercomputing has proven to be an effective ally to our society by helping us address critical and pressing challenges, such as climate change, national emergencies, and public health. For instance, supercomputers help researchers develop personalized medical treatments as well as better predict and manage the effects of natural disasters such as floods and earthquakes through the use of advanced computer simulations.[2,3] These discoveries help shape our understanding of the universe, bolster US economic competitiveness, and contribute to a better future.[4]

I will begin with a brief overview of activities at the Oak Ridge Leadership Computing Facility, exemplifying its lasting impact on scientific discovery, including its impact on global emergencies, most recently the COVID-19 pandemic. Then, I will discuss the need for sustainable R&D investments in supercomputing and data management strategies to ensure the nation's continued economic competitiveness and to help expand our nation's high-tech workforce. Lastly, I will discuss a vision of how DOE's national laboratories' leadership computing user facilities and computational

---

[1] US DOE Office of Science. High-performance Computing. https://www.energy.gov/science/initiatives/high-performance-computing. Accessed May 13, 2021.
[2] K.R. Milner, et al, Toward Physics-Based Nonergodic PSHA: A Prototype Fully Deterministic Seismic Hazard Model for Southern California," Bulletin of the Seismological Society of America (2021). doi:10.1785/0120200216.
[3] C.W. Tessum, et al. "InMAP: A model for air pollution interventions." PloS ONE 12, no. 4 (2017): e0176131. https://doi.org/10.1371/journal.pone.0176131
[4] Report. *The State of the DOE National Laboratories.* https://www.energy.gov/sites/default/files/2021/01/f82/DOE%20National%20Labs%20Report%20FINAL.pdf

workforce can support a thriving research and development ecosystem while promoting AI innovation, energy-efficiency, stakeholder engagement, and national security.

THE OAK RIDGE LEADERSHIP COMPUTING FACILITY (OLCF)

Collectively, the DOE laboratory complex occupies a distinctive position in the national science and technology innovation ecosystem. By fostering strong partnerships with industry and academia, the DOE national labs bring together one-of-a-kind and often first-of-a kind scientific instruments, powerful computing infrastructure, and computational and data science experts in multiple disciplines to push the boundaries of computing and solve some of the world's biggest challenges in science and engineering.

ORNL has been a global leader in high-performance computing for nearly 30 years. The Oak Ridge Leadership Computing Facility (OLCF) provides access to supercomputers of unprecedented capability to the scientific community on behalf of the US Department of Energy (DOE). These versatile systems are 10 to 100 times more powerful than most supercomputers available for research. Our users run a variety of codes and analyses, from traditional modeling and simulation to new and emerging classes of data science and artificial intelligence. But in every instance, our users tackle the most complex and difficult problems that cannot be easily solved using conventional computing resources.

From the beginning, the OLCF has been at the forefront of the rapid evolution in scientific computing. OLCF provided the **first teraflop system** (IBM Power3 Eagle) for open science and the science community's **first petaflop system** (Cray XT5 Jaguar). **Three OLCF systems ranked as the fastest in the world by the TOP500 list,** including OLCF's current **pre-exascale system**, Summit, an IBM system. Summit debuted in June 2018 and has a theoretical peak performance of 200 petaflops. This is over a million times the capability of Eagle and provides scientists with incredible computing power to solve critical problems in areas as diverse as clean energy, human health, advanced materials, climate, nuclear physics, and other research areas.

OLCF is at the brink of delivering **the first exascale system** in the U.S., the 1.5 exaflops Frontier supercomputer, expected this fall. Frontier will perform calculations up to 50 times faster than today's top supercomputers. Frontier will support research in solving problems in energy, medicine, and materials that were impossible as recently as 5 years ago. Exascale is the next level of computing performance and is needed to continue U.S. leadership in high-performance computing. DOE's Exascale Computing Project is developing applications across two dozen scientific and technical disciplines to run on Frontier on day one. These applications will offer new insights in quantum materials, chemistry, physics, additive manufacturing, fuels, fusion, fission, and more. Equally important, Frontier will mark an impressive trajectory for our facility, for DOE, and for the nation, demonstrating that in the past decade we improved computational power by a factor of 500 while suppressing growth of the carbon footprint which increased only by a factor of 4.

Summit is in high demand. Researchers annually submit proposals for groundbreaking computational challenges and are granted time on DOE leadership computers based on scientific merit and their need for and ability to use the full compute power and capacity of a leadership-class system. OLCF typically receives requests for **three to five times** more node hours than what the programs are able to award**.** Still, the facility preserves the ability to shift resources to address national emergencies, as you will hear later in my testimony. Last but not least, to train users for success on our HPC systems, we host workshops, conference calls, and training events. These events can help prepare new users for large-scale computing systems, provide useful tools to veteran users, and facilitate engagement both with the public and the user community.

ADVANCING SCIENCE VIA SUPERCOMPUTING

Using supercomputers, scientists have expanded the scale and scope of their research, solved complex problems in less time, and filled critical gaps in scientific knowledge. The following is a short list of recent examples of how scientists leveraged Summit and OLCF expert support to advance science.

Reaching new heights in weather forecasting[5]: Using Summit, a team of researchers achieved a computational first: a global simulation of the Earth's atmosphere at a 1-square-kilometer average grid-spacing for a full 4-month season. The milestone marks a big improvement in resolution which today operates at 9-square-kilometer grid-spacing for routine weather forecast operations. This improvement is a critical step in the effort to create multi-season atmospheric simulations at high resolution, pointing toward the future of weather forecasting—one powered by the emerging exascale supercomputers.

Gaining new insights into quantum materials[6]: Quantum materials—those that have correlated order at the subatomic level—have potential applications in electronic devices, quantum computers, and superconductors. For the first time, a team at ORNL used supercomputing and artificial intelligence (AI) to find patterns in neutron scattering data that can guide future neutron scattering experiments and lead to deeper understanding of the physics inside quantum materials.

Designing efficient and durable gas turbine jet engines[7]: General Electric (GE) designs gas turbine jet engines that are world-leading in efficiency, emissions, and durability. A key requirement to successfully design these engines is the ability to accurately predict the temperature and motion of fluids in the high-pressure turbines that power them. Internal turbine temperatures become so high during operation that turbine components will melt if they are not cooled properly. A GE team used the Summit supercomputer at ORNL to complete first-of-their-kind 3D flow simulations that provided breakthrough insights about these fluid flows and revealed how to better cool engine parts for more efficient and durable jet engines. The team performed the computationally intensive large simulations with unprecedented speed and exceptional detail that more closely matched results of actual engine tests.

---

[5] Wedi, N. P., et al. A baseline for global weather and climate simulations at 1 km resolution. *Journal of Advances in Modeling Earth Systems*, 12 (2020), e2020MS002192. doi: 10.1029/2020MS002192
[6] Anjana M. Samarakoon et al., "Machine-Learning-Assisted Insight into Spin Ice $Dy_2Ti_2O_7$." *Nature Communications* 11 (2020): 892, doi:10.1038/s41467-020-14660-y.
[7] OLCF Website. *Sparks Fly in Marriage of GE's Genesis Code and the Summit Supercomputer*. https://www.olcf.ornl.gov/2020/08/10/sparks-fly-in-marriage-of-ges-genesis-code-and-the-summit-supercomputer/. Accessed May 13, 2021.

Deploying AI solutions for modernization of the national cancer surveillance program: Supporting efforts to reduce the cancer burden in the United States, the national Surveillance, Epidemiology, and End Results (SEER) Program manages the collection of curated data from population-based cancer registries on cancer incidence, prevalence, survival, and associated health statistics for the advancement of public health. As part of the National Cancer Institute (NCI), the SEER program is tasked with supporting cancer research to improve the understanding of patient care and outcomes in the "real world" beyond the clinical trial setting. By leveraging ORNL's pre-exascale computing technologies, a DOE-NCI team delivered state-of-the-art AI tools to help US state cancer registries improve the accuracy and efficiency of their operations in cancer incidence reporting. The AI tools make the dream of "real time" cancer incidence reporting a reality, while also enabling real-time eligibility assessment of cancer patients for clinical trials. This achievement to modernize the national cancer surveillance program exemplifies the benefits of a federally coordinated strategy to harness AI, high-performance computing, and sensitive health data assets for real-world application.

THE ROLE OF SUPERCOMPUTING IN FIGHTING THE COVID-19 PANDEMIC

In the past decade, supercomputing has emerged as a great ally not only to biology but also to biomedicine. This growing trend is even more apparent with life sciences being one of the leading application domains for artificial intelligence. Therefore, from the start of the COVID-19 pandemic, supercomputers, and by extension AI, have helped accelerate discovery, understand the virus, and inform management of the pandemic response.

In my role as director of the Oak Ridge Leadership Computing Facility, I have experienced firsthand our staff's herculean efforts to pivot fast and offer our world-leading computing resources and computational and data expertise in the fight against the coronavirus. Following are recent examples of how OLCF has leveraged its scientific tools and expertise to address challenges posed by the COVID-19 pandemic.

COVID-19 high-performance computing consortium: In March 2020, The OLCF and the Summit supercomputer joined forces with other U.S. federal agencies, industry, and academic leaders to provide access to the world's most powerful high-performance computing resources in support of COVID-19 research through the COVID-19 High-Performance Computing Consortium[8].

The Consortium is a unique private-public effort spearheaded by the White House Office of Science and Technology Policy, the U.S. Department of Energy and IBM to bring together federal government, industry, and academic leaders who are volunteering compute time and resources on their world-class machines. Since the beginning of the Consortium, OLCF has allocated 2,206,200 compute node hours across 21 most computationally demanding projects, supporting multi-disciplinary teams to improve our understanding of the virus's structure and biology in order to develop targeted therapies and vaccines, build better diagnostic and prognostic disease models, model how the disease may spread in communities as conditions change on the ground, as well as improve our understanding of how the virus may change as it spreads through the population.

Guiding vaccine development and therapeutic targeting: Research by a multi-disciplinary team leveraging OLCF's Summit led to a novel understanding of SARS-CoV-2 and a new method for studying disease on Summit. Since researchers first mapped the SARS-CoV-2 spike protein—the main infection machinery of the virus that causes the COVID-19 disease—scientists around the world embarked on a quest to understand the movement of the virus's spike protein, namely how it behaves and gains access to the human cell. A multi-institutional team built a first-of-its-kind AI-based workflow and ran it on the Summit supercomputer to simulate the spike protein in numerous environments, including within the SARS-CoV-2 viral envelope comprising 305 million atoms—the most comprehensive simulation of the virus performed to date. The team first optimized their molecular dynamics codes, which model the movements of atoms in time and space, on multiple smaller HPC systems. These optimizations prepared the team to run full-scale simulations leveraging the whole OLCF Summit supercomputer. The Summit runs have led to more detailed understanding of the virus

[8] The COVID-19 High-performance Computing Consortium Website. https://covid19-hpc-consortium.org/. Accessed May 13, 2021.

behavior but also its vulnerabilities, to guide vaccine development and better therapeutic targeting. Because one set of the Summit calculations generated a whopping 200 terabytes of data (equivalent to 62 million high resolution images), the team used AI to identify the intrinsic features from the simulations and break down the information to help them interpret what was happening. By layering the experimental data and the simulation data and combining it with their AI-based approach on Summit, the researchers were able to capture the virus and its mechanisms in unprecedented detail. The team received the prestigious ACM Gordon Bell Special Prize for HPC-Based COVID-19 Research at the 2020 International Conference for High Performance Computing, Networking, Storage and Analysis[9].

I would like to thank Congress for the CARES-Act funds OLCF received to augment Summit and help support the COVID-19 research community.

WHERE DO WE GO NEXT?

High-performance computing's impact on national competitiveness has long been embodied by the saying, "You must compute to compete." With advances in sensors and other data generation technologies, machine learning, and artificial intelligence, this saying can be amended to, "You must learn faster to compete." In today's world it is evident that the scientists, engineers, companies, and nations that learn the fastest are the most adaptive and most competitive. Those that can harness computation to analyze, identify patterns, and draw inferences from digital models and simulations, empirical measurements from large scientific instruments, or the wealth of data being collected from the connected sensors that have become ubiquitous will make discoveries and innovations at a faster pace than those that do not. It is imperative for the United States to expand and enhance the national research ecosystem and empower researchers with the computational tools so that they can unlock new insights hidden in data. We can accomplish this by taking the following steps.

---

[9] bioRXiv. "AI-Driven Multiscale Simulations Illuminate Mechanisms of SARS-CoV-2 Spike Dynamics." https://www.biorxiv.org/content/10.1101/2020.11.19.390187v1

First, we need to make the most of our Exascale systems: The upcoming exascale systems such as Frontier mark the beginning of a 5 – 7 year operations cycle. Researchers will be able to answer questions of national importance that simply can't be addressed with today's computing platforms. For example, users of Frontier will model the entire lifespan of a nuclear reactor to enhance efficiency and safety of operations. Biomedical scientists will use Frontier to improve health outcomes in new ways, from uncovering the underlying genetics of disease to extracting important information from complex patient data. Frontier will build on recent developments in science and technology to further integrate artificial intelligence with data analytics and modeling and simulation. These new capabilities will drastically reduce the time to discovery by automatically recognizing patterns in data and guiding simulations beyond the limits of traditional approaches. Overall, exascale systems will allow us to expand the application space and increase the impact of our investments. Furthermore, the exascale systems will offer training opportunities to grow a computationally-aware and computationally-savvy scientific workforce across academia, industry, national labs, and federal agencies.

Second, we need to continuously invest in new technologies and workforce development to maintain competitive advantage and ensure our global leadership: Ongoing investments in computing, data, and software technologies will provide the engine that drives an advanced research ecosystem that is energy-efficient, secure, and performant. Data infrastructure and software are needed to enable the flow of data between facilities. Scientific software must be maintained and extended to support these activities. Research in microelectronics will increase the power of computer processors to the limits of current transistor technology.

Furthermore, the energy demands of digital infrastructures and compute- and data-hungry AI algorithms pose additional challenges. AI and other computing activities are projected to use over half of the world's energy by 2040.[10,11] We cannot easily

[10] Greenpeace, "Clicking Clean: Who is winning the race to build a green internet?" Washington, DC, 2017.
[11] Semiconductor Industry Association and Semiconductor Research Corporation, "Rebooting the IT revolution: A call to action," September 2015.

predict how we can balance AI's energy demands with AI's demonstrated ability to guide management of our energy resources more efficiently. DOE's national laboratories are well positioned to lead hardware, software, and algorithmic innovations and deliver AI solutions that consume less energy, a challenge DOE is already working to address as part of the Exascale Computing Initiative. Energy-efficient AI is key to providing sustainable and affordable solutions that benefit our environment and our national and economic security.

Our societal advances in understanding the brain and its ability to process information in an energy-efficient manner will in turn influence how computing technologies will support society. While interest and advancements in the computing technologies of neural networks and neuromorphic computing were originally inspired by human neural capabilities, they have become broad fields of research in their own right. We anticipate that investments in AI methods and accelerated computing methods — from domain-specific accelerators and learning-and-inferencing at the network's edge to neuromorphic technologies and quantum accelerators — will directly enhance our national competitiveness.

Third, we need to make additional investments to support a national data infrastructure: The explosive growth of AI is driven by the convergence of big data, massive computational power, and novel algorithms. There is a pressing need to consider expanding federal investments in managing our national data assets in a secure environment that can provide the compute resources and scalable data analytics capabilities. The Department of Energy national laboratories are uniquely positioned to lead this effort for the nation, given the Department's long history of serving as the steward of large data infrastructures and of the nation's nuclear security enterprise. With their leading role in high-performance computing and their extensive data science and AI capabilities, the national laboratories could serve as a neutral entity, an honest broker for democratizing AI, while providing meaningful and responsible access to sensitive data assets and compute resources. This would further drive the need to educate and train a new workforce necessary to support these activities. Such an investment is critical to support the continuum of scientific discovery for effective domain-specific application of the nation's high performance computing and data assets. The COVID-19 pandemic demonstrated the value and importance of having established interagency programs and data integration ahead of this unanticipated crisis, and the utility of high-performance computing and artificial intelligence for rapid, complex, real-world data analyses.[12]

CONCLUDING REMARKS

The Department of Energy's national laboratories are a remarkable asset for the nation. Over the past 75 years, they have consistently provided the science and technology needed to address national problems. As a DOE national laboratory and Federally Funded Research and Development Center, ORNL is equipped with exceptional computational and experimental resources to support the Department's mission needs and extend to the nation's research community. Researchers at Oak Ridge and other national laboratories have a wealth of experience in delivering, operating, and using supercomputers to solve the most complex problems of our times. With the steadily increasing power of today's supercomputers and the massive data sets that are becoming available in a variety of areas, we are also in a position to build AI solutions that help us accumulate, analyze, and automate the delivery of functional knowledge in many application domains.

Our nation faces a formidable set of challenges: ensuring national security in a changing world; increasing the availability of clean, reliable, and affordable energy while protecting the environment and addressing climate change; improving human health; and enhancing U.S. competitiveness in the global economy by fostering scientific leadership and encouraging innovation. Leadership computing is expected to offer solutions to many of these challenges, but the importance of building a sustainable and integrated compute and data infrastructure cannot be underestimated.

A cohesive national plan for scientific computing, supporting emerging technologies as well as advancing AI, is imperative to secure the nation's economic

---

[12] Ramoni, Rachel, et al. "COVID-19 Insights Partnership: Leveraging Big Data from the Department of Veterans Affairs and Supercomputers at the Department of Energy under the Public Health Authority." Journal of the American Medical Informatics Association (2021).

competitiveness and well-being. At the same time, we as a nation have the unique opportunity to create a well-defined, federally coordinated roadmap for integrating our leadership computing facilities with national experimental facilities and federal data assets to deliver benefits across private and public sectors. The DOE Leadership Computing Facilities are uniquely equipped and positioned not only to make substantial contributions to address these opportunities and challenges, but also to support the execution of a national research computing ecosystem.

Thank you again for the opportunity to testify. I welcome your questions on this important topic.

Georgia Tourassi, Ph.D.
Director, National Center for Computational Sciences
Oak Ridge National Laboratory

tourassig@ornl.gov
865-576-4829

Dr. Georgia Tourassi is the Director of the National Center for Computational Sciences at the Oak Ridge National Laboratory (ORNL). Concurrently, she holds appointments as an Adjunct Professor of Radiology at Duke University and as a joint UT-ORNL Professor of the Bredesen Center Data Science Program at the University of Tennessee at Knoxville.

Her scholarly work includes 13 US patents and innovation disclosures and more than 260 peer-reviewed journal articles, conference proceedings articles, editorials, and book chapters. She is elected Fellow of the American Institute of Medical and Biological Engineering (AIMBE), the American Association of Medical Physicists (AAPM), the International Society for Optics and Photonics (SPIE), and the American Association for the Advancement of Science (AAAS).

Her research interests include high performance computing and artificial intelligence in biomedicine. For her leadership in the Joint Design of Advanced Computing Solutions for Cancer initiative, she received the DOE Secretary's Appreciation Award in 2016. In 2017, she received the ORNL Director's Award for Outstanding Individual Accomplishment in Science and Technology and the UT-Battelle Distinguished Researcher Award. In 2020, Dr. Tourassi received the DOE's Secretary Honors Award for her contributions to the COVID 19 Insights Partnership Team and to the COVID 19 HPC Resource Team.

Tourassi holds a B.S. in Physics from Aristotle University of Thessaloniki, Greece, and a Ph.D. in Biomedical Engineering from Duke University.

Chairman BOWMAN. Thank you so much, Dr. Tourassi.
Dr. Willcox, you are now recognized.

TESTIMONY OF DR. KAREN WILLCOX, DIRECTOR, ODEN INSTITUTE FOR COMPUTATIONAL ENGINEERING AND SCIENCES AT THE UNIVERSITY OF TEXAS AT AUSTIN

Dr. WILLCOX. Thank you, Chair Bowman, Ranking Member Weber, Chair Johnson, Ranking Member Lucas, and Members of the Subcommittee.

Today, I have three main points. First, the future of scientific computing must be interdisciplinary. Second, the DOE ecosystem that supports mission-driven basic research in scientific computing is a national scientific treasure. And third, the future of scientific computing hinges critically on the availability of a highly skilled workforce passionate about addressing the Nation's challenges in science, security, and sustainability.

So first, on the interdisciplinary future of scientific computing, the pace at which scientific computing can accelerate discovery and innovation will be limited by the rate at which we address foundational challenges that currently limit the complexity, scale, and trustworthiness of computational tools. This requires scientific computing research that draws on many fields, including computer science, computational science, the mathematical sciences, the domain sciences, and engineering.

Particularly important is the role of the field of computational science. Computational science differs from computer science because at its core, computational science involves developing mathematical models and simulations rooted in physical and mechanistic principles.

As we look to the future of scientific computing, the boundaries between computational science and computer science are becoming increasingly blurred. The future of scientific computing will involve new approaches that span the two fields such as AI and machine learning, and indeed the DOE has been at the forefront of defining notions such as AI for science and scientific machine learning.

However, when it comes to AI approaches in science and engineering, we must be careful not to chart our course based entirely on the successes of data science and machine learning in vastly different domains such as social media and online retail. We must instead recognize that energy, environmental, and nuclear challenges by their very nature require predictions that go well beyond the available data. There's a critical need to quantify uncertainty and to make informed decisions that account for risk. The future of scientific computing will only address these needs through a balanced investment in the foundational mathematical sciences and in computational science, along with data science and computer science. And we must also not underestimate the criticality of continuing to invest in experimental research and development since advancing discoveries through computational models really requires validation.

That brings me to my second point on the value of the DOE's mission-driven basic research ecosystem and its role in addressing these challenges. DOE supports for basic research at the national labs and at the Nation's universities has fostered interdisciplinary computing research in a way that community-driven basic research has struggled to achieve. And as one example, I highlight the Mathematical Multifaceted Integrated Capabilities Centers, or MMICCs, of the DOE applied math program. These centers focus on applied math basic research but strongly driven by application needs. For example, our AEOLUS (Advances in Experimental Design, Optimization and Learning for Uncertain Complex Systems) MMICC is addressing the basic mathematical research needs for advanced materials and additive manufacturing.

The MMICC program has been transformational in how it has shaped my own basic research portfolio, and one of the critical elements, first, the size of the center is large enough to bring together a diverse team that includes mathematicians, computer scientists, computational scientists, engineers, and domain experts spanning universities and national labs. This in turn enables a much-needed holistic approach for a complex system.

Second, the long funding horizon provides the stability to invest in challenging, high-payoff basic research ideas.

And third, the mission-driven nature challenges my mathematical research to target problems that are of high relevance to practitioners if the focus on basic research permits us to lay long-lasting foundations.

My final point is that achieving this future vision for scientific computing hinges critically on the availability of a highly skilled workforce. The challenges in front of us are twofold. First is training the workforce with the interdisciplinary skills that cut across the mathematical sciences, computing, and domain sciences, and second is ensuring a strong, diverse pipeline of highly trained professionals who remain committed to scientific and engineering domains rather than being lured away by more lucrative positions in commercial and business sectors. A critical part of this training is the immersive research experiences enabled by basic research grants such as the MMICC program I described earlier. Maintaining a strong investment in DOE basic research funding for universities while also continuing to support the collaborative and academic alliance programs at the national labs is absolutely critical to addressing the Nation's future workforce needs.

Thank you, and I look forward to your questions.

[The prepared statement of Dr. Willcox follows:]

Testimony for the *Subcommittee on Energy* of the
House Committee on Science, Space, and Technology (117th Congress)

hearing on

Accelerating Discovery: the Future of Scientific Computing at the Department of Energy
Wednesday, May 19, 2021 - 11:00am EDT

By

Dr. Karen E. Willcox
Director, Oden Institute for Computational Engineering and Sciences
Associate Vice President for Research
Professor of Aerospace Engineering and Engineering Mechanics
The University of Texas at Austin
Email: kwillcox@oden.utexas.edu   Website: kiwi.oden.utexas.edu

Chairperson Bowman and Ranking Member Weber: I appreciate this opportunity to submit testimony for this hearing on the future of scientific computing at the Department of Energy (DOE).

My name is Karen Willcox and I serve as the Director of the Oden Institute for Computational Engineering and Sciences at The University of Texas at Austin, where I am also Associate Vice President for Research and a Professor of Aerospace Engineering and Engineering Mechanics.

I offer my remarks today based on my expertise in scientific computing as well as my extensive interactions with the DOE in multiple capacities, which include serving as the current Co-director of the AEOLUS Mathematical Multifaceted Integrated Capabilities Center (MMICC) funded by the DOE Advanced Scientific Computing Research (ASCR) program, serving as a current member of the External Review Board for the Computing and Information Science Research Foundation at Sandia National Laboratories, and serving as a current member of the Advisory Board for the Advanced Simulation and Computing program at Los Alamos.

I wish to convey to you three main points:

1. **The future of scientific computing must be interdisciplinary**. Its very core must involve computer science, computational science, the mathematical sciences, the domain sciences, and engineering.
2. **The DOE ecosystem that supports mission-driven basic research in scientific computing across the National Laboratories and the Nation's universities is a national scientific treasure** that must be nurtured and bolstered to ensure a secure, sustainable, and competitive future for the Nation.
3. **The future of scientific computing hinges critically on the availability of a highly skilled workforce passionate about addressing the Nation's challenges in science, security, and sustainability.**

The interdisciplinary future of scientific computing

Scientific computing has an unquestionably central role to play in future scientific discovery and technological innovation to address societal grand challenges. The pace at which scientific computing can accelerate discovery and innovation will be limited by the rate at which we address foundational challenges that currently limit the complexity, scale, and trustworthiness of computational analysis, prediction, and decision tools. To address these challenges requires scientific computing research and development that draws on many fields, including computer science, computational science, and the mathematical sciences, and also includes close collaborations with domain scientists and engineers.

Particularly important is the critical role of the field of computational science (sometimes called computational science and engineering). Computational science is an interdisciplinary field that uses mathematical modeling and advanced computing to understand and solve complex problems. Computational science differs from computer science because at its core, computational science involves developing mathematical models and simulations rooted in physical and mechanistic principles, in order to understand, analyze, predict, design, and control natural and engineered systems. Quoting Rüde et al.[1]

> *"Advances in computational science have led to more efficient aircraft, safer cars, higher-density transistors, more compact electronic devices, more powerful chemical and biological process systems, cleaner power plants, higher-resolution medical imaging devices, and more accurate geophysical exploration technologies—to name just a few."*

We can draw insight as to the nature and impact of computational science through the historical example of the finite element method—the workhorse of modern computer-aided engineering analysis and design, and the foundation for the multi-billion-dollar computer-aided engineering (CAE) software industry, which has transformed the practice of engineering. What has it taken to make the finite element method a powerful broadly applicable analysis and design tool that is literally in the hands of every engineer? What has it taken to grow the impact of the finite element method from its origins in engineering structural analysis to its modern-day use in diverse applications from nuclear power plants to subsurface contaminants to polar ice sheets to materials processing to combustion processes and everything in between? A key part of the answer is the decades of investment in foundational mathematical and computational basic research, inspired by and connected to driving applications across engineering and the sciences. This research developed the foundational mathematical theory, such as the error analysis that established the finite element method's reliability,[2] the mathematical formulations that tackled the challenges of numerical stability,[3] and the theory and methods that extended the reach of the finite element method to nonlinear problems.[4]

As we look to the future of scientific computing, the boundaries between computational science and computer science are becoming increasingly blurred. The field of computational science is evolving in the face of increased data in its application domains, while computer science is beginning to impact domains in science, engineering and medicine. The future of scientific computing will involve promising new approaches that span the two fields, such as artificial intelligence (AI) and machine learning, enabled by the increasing amount of scientific data and by advances in scalable algorithms. Indeed, the DOE has been at the forefront of defining the notions of *AI for science*[5] and *scientific machine learning*,[6] with the goal of accelerating research and development breakthroughs in energy, basic science, engineering, medicine, and national security. However, as stated in my recent *Nature Computational Science* perspective piece,[7] when it comes to the development and adoption of AI approaches in scientific and engineering fields we must not lose sight of the need for a balanced investment that goes beyond computer science:

> *"For the last six decades these [scientific and engineering] fields have been advanced through the synergistic and principled use of theory, experiments, and physics-based simulations. Our increased ability to sense and acquire data is clearly a game-changer in these endeavors. Yet in our excitement to define a new generation of data-centric approaches, we must be careful not to chart our course based entirely on the successes of data science and machine learning in the vastly different domains of social media, online entertainment, online retail, image recognition, machine translation, and natural language processing—domains for which data are plentiful and physics-based models do not exist."*

We must recognize that energy, environmental, and nuclear challenges by their very nature require *predictions* that go well beyond the available data. There is a critical need to *quantify uncertainty* and our associated confidence in predictions; there is a critical need to make *informed decisions that account for risk*. The future of scientific computing will only address these needs through balanced investment in the foundational mathematical sciences and in computational science, along with data science and computer science. We must also not underestimate the criticality of continuing to invest in *experimental research and development*; advancing discovery through scientific computing requires validated computational models. Again, we can look to examples from our scientific past and appreciate that advances in time/space resolved experimental diagnostics have contributed significantly to establishing trust and credibility in physics-based computational models, because of the ability to do meaningful comparisons and validations at small scales.

The value of the DOE's mission-driven basic research ecosystem

The DOE research and development ecosystem is uniquely positioned to play a leading role in addressing these challenges and in crafting a strong interdisciplinary future for scientific computing.

The National Laboratories exemplify a culture of careful, measured, validated, and verified research that addresses vital scientific and technical application domains. Their leadership in scientific computing is complemented by their world-leading experimental and scientific user facilities. The growing efforts to support research that cuts across computational and experimental domains are essential to the future of scientific computing, and here DOE plays a unique role.

DOE support for basic research at National Laboratories and at the Nation's universities has fostered interdisciplinary computing research in a way that community-driven basic research has struggled to achieve.

As one example, I highlight the Mathematical Multifaceted Integrated Capabilities Center (MMICC) program of DOE ASCR. The Applied Mathematics Program invests $9M per year to fund three MMICC centers. The focus of these centers is on basic research in applied mathematics, but strongly driven by application needs. For example, our current AEOLUS MMICC, led by The University of Texas at Austin, is addressing basic mathematical research needs to enable predictive modeling, optimal process control, and optimal experimental design for applications in advanced materials and additive manufacturing. Over the past eight years, the MMICC program has been transformational in how it has shaped my basic research portfolio. What are the crucial elements? (1) The size of the center is large enough to bring together a diverse team that includes mathematicians, computer scientists, computational scientists, engineers, and domain experts, spanning universities and National Laboratories. This in turn enables a much-needed holistic research approach for increasingly complex systems. (2) The long funding horizon (4 or 5 years) provides the stability to invest in challenging high-payoff basic research ideas. It also provides the opportunity for PhD students to truly integrate with the team and the project. (3) The mission-driven nature of the center goals challenges my mathematical research to target problems that practitioners actually care about, yet the focus on basic research permits the research to lay long-lasting foundations that may ultimately impact a broad range of problems.

This notion of mission-driven cross-cutting mathematical research has been a mainstay of the DOE Applied Mathematics Program. It has provided, and will continue to provide, the rigorous mathematical and computational underpinnings that are essential to advancing scientific computing.

The criticality of the workforce

Achieving this future vision for scientific computing hinges critically on the availability of a highly skilled workforce passionate about addressing the Nation's challenges in science, security, and sustainability. The challenges in front of us include (1) training the workforce with the interdisciplinary skills that cut across the mathematical sciences, computing, and domain sciences, and (2) ensuring a strong, diverse pipeline of highly trained professionals who remain committed to scientific and engineering domains, rather than being lured away by more lucrative positions in commercial and business sectors.

The Oden Institute at The University of Texas at Austin has a globally recognized interdisciplinary graduate program in Computational Science, Engineering, and Mathematics (CSEM).[8] The CSEM program is unique in that we sit outside the usual departmental and school structure; our students are truly trained at the interfaces. A critical part of that training is the immersive research experience enabled by basic research grants, such as the MMICC program I described earlier, or the Predictive Science Academic Alliance Program (PSAAP).[9] Our graduate students work with collaborators from the National Laboratories and from industry partners. They engage in internships. They are immersed in the notion of basic research that targets societal grand challenges together with a culture of rigorous mathematically grounded approaches and a culture of high performance computing at scale. This prepares them to contribute to some of the Nation's most pressing scientific and technological challenges. For example, under our previous Diamond MMICC center, we trained scores of doctoral students and postdoctoral researchers, many of whom have gone on to careers in academia and the National Laboratories.[10]

Maintaining a strong investment in DOE basic research funding for universities, while also continuing to support the collaborative and academic alliance programs at the National Laboratories, is absolutely critical to addressing the Nation's future workforce needs.

Summary

Scientific computing will play a central role in future scientific discovery and technological innovation to address societal grand challenges. Scientific computing has and will thrive in an ecosystem that fosters interdisciplinary basic research and that provides the culture, environment, and resources needed to train a highly skilled workforce passionate about addressing the Nation's challenges in science, security, and sustainability. The DOE has been uniquely strong in providing this ecosystem in the past decades, and, with the proper support, is well positioned to do so in the future.

---

[1] Rüde, U., Willcox, K., Curfman McInnes, L. and de Sterck, H., 2018. Research and education in computational science and engineering. *SIAM Review*, 60(3), pp.707-754.
[2] Babuška, I., Strouboulis, T. and Whiteman, J.R., 2001. *The Finite Element Method and its Reliability*. Oxford University Press, Oxford, United Kingdom.
[3] Hughes, T., Franca, L. and Balestra, M., 1986. A new finite element formulation for computational fluid dynamics: V. Circumventing the Babuška-Brezzi condition: A stable Petrov-Galerkin formulation of the Stokes problem accommodating equal-order interpolations. *Computer Methods in Applied Mechanics and Engineering*, 59(1), pp.85-99.
[4] Oden, J.T., 1972. Finite Elements of Nonlinear Continua. McGraw Hill, New York, NY.
[5] Stevens, R., Taylor, V., Nichols, J., Maccabe, A.B., Yelick, K. and Brown, D., 2020. AI for Science (No. ANL-20/17). Argonne National Laboratory (ANL), Argonne, IL.
[6] Baker, N., Alexander, F., Bremer, T., Hagberg, A., Kevrekidis, Y., Najm, H., Parashar, M., Patra, A., Sethian, J., Wild, S. and Willcox, K., 2019. Workshop report on basic research needs for scientific machine learning: Core technologies for artificial intelligence. USDOE Office of Science (SC), Washington, DC.
[7] Willcox, K., Ghattas, O. and Heimbach, P., 2021. The imperative of physics-based modeling and inverse theory in computational science. *Nature Computational Science*, 1(3), pp.166-168.
[8] https://www.oden.utexas.edu/graduate-studies/
[9] https://psaap.llnl.gov/
[10] http://dmd.mit.edu/young-gems

Karen E. Willcox, MNZM, PhD
The University of Texas at Austin
email: kwillcox@oden.utexas.edu website: kiwi.oden.utexas.edu

Karen E. Willcox is Director of the Oden Institute for Computational Engineering and Sciences, Associate Vice President for Research, and Professor of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin. She is also External Professor at the Santa Fe Institute. She holds the W. A. "Tex" Moncrief, Jr. Chair in Simulation-Based Engineering and Sciences and the Peter O'Donnell, Jr. Centennial Chair in Computing Systems. Before joining the Oden Institute in 2018, she spent 17 years as a professor at the Massachusetts Institute of Technology, where she served as the founding Co-Director of the MIT Center for Computational Engineering and Associate Head of the MIT Department of Aeronautics and Astronautics. Prior to joining the MIT faculty, she worked at Boeing Phantom Works with the Blended-Wing-Body aircraft design group. Willcox has co-authored more than 120 papers in peer-reviewed journals and advised more than 60 graduate students. She is the recipient of multiple best paper awards and several awards for leadership and teaching. In 2017 she was appointed Member of the New Zealand Order of Merit (MNZM). She is a Fellow of the Society for Industrial and Applied Mathematics (SIAM) and a Fellow of the American Institute of Aeronautics and Astronautics (AIAA).

Education

1994 University of Auckland, Bachelor of Engineering, First Class Honours (Engineering Science)
1996 Massachusetts Institute of Technology, Master of Science (Aeronautics and Astronautics)
Thesis: *Aeroelastic Computations in the Time Domain using Unstructured Meshes*
2000 Massachusetts Institute of Technology, PhD (Aeronautics and Astronautics)
Thesis: *Reduced-Order Aerodynamic Models for Aeroelastic Control of Turbomachines*

Experience

University of Texas at Austin
2020-present Associate Vice President for Research
2018-present Director, Oden Institute for Computational Engineering and Sciences
2018-present Professor of Aerospace Engineering and Engineering Mechanics

Santa Fe Institute
2019-present External Professor

Massachusetts Institute of Technology
2001-2018 Assistant/Associate/Full Professor, Aeronautics and Astronautics
2011-2013 Associate Department Head, Aeronautics and Astronautics
2008-2018 Founding Co-Director, MIT Center for Computational Engineering

Singapore University of Technology and Design
2018 Visiting Professor (7-month stay)
2015 Visiting Professor (6-month stay)
2011 Visiting Associate Professor (6-month stay)

University of Auckland, New Zealand
2015 Visiting Professor, Department of Engineering Science (8-month stay)
2008-2009 Visiting Associate Professor, Department of Engineering Science (15-month stay)

Sandia National Laboratories
2005 Visiting Researcher, Computer Science Research Institute (5-month stay)

Stanford University
2005 Visiting Scholar (1-month stay)

Boeing Phantom Works
2000-2001 Visiting Researcher, Blended-Wing-Body Aircraft Design Group (1-year stay)

NASA Dryden Flight Research Center
1996 Aerospace Intern, Aerodynamics Branch

Karen E. Willcox, MNZM, PhD

Professional Interests

Research: Data to decisions in engineering systems. Computational models and methods for design, optimization, control and uncertainty quantification of engineering systems. Predictive data science and scientific machine learning. Reduced-order modeling and multi-fidelity methods. Future aircraft technologies, aircraft system optimization, aircraft environmental impact, multidisciplinary design, unmanned aerial vehicles, Digital Twin, Digital Thread.

Education: EdTech for data visualization, modeling and analytics (mapping.mit.edu). Fly-by-Wire intervention to enable scalable differentiated instruction in community colleges (fbw.mit.edu). Mapping learning outcomes across the undergraduate engineering curriculum (xoces.mit.edu); linking topics across the curriculum (crosslinks.mit.edu).

Teaching: Principles of Automatic Control (undergraduate), Computational Methods in Aerospace Engineering (undergraduate), Signals and Systems (undergraduate), Multidisciplinary System Design Optimization (graduate), Flight Vehicle Aerodynamics (graduate), Numerical Methods for Partial Differential Equations (graduate).

Diversity, Equity and Inclusion: Established new Diversity, Equity, Inclusion and Outreach Committee at the Oden Institute. Grew diversity of undergraduate and graduate aerospace engineering student body as Associate Department Head in MIT. Led Rising Stars events at MIT and UT Austin to foster gender diversity in aerospace engineering and computational sciences. Active in outreach activities to promote girls' interest in science, mathematics and engineering, including volunteer grade school science extension classes, many outreach visits to K-12 schools, and participation in the Advisory Board for Girls' Angle. First-generation student mentor at MIT.

Professional Memberships

Fellow, American Institute of Aeronautics and Astronautics (AIAA)

Fellow, Society for Industrial and Applied Mathematics (SIAM)

Member, American Society for Engineering Education (ASEE)

Member, American Mathematical Society (AMS)

Member, Design Society

External Boards and Committees (current)

Advanced Simulation and Computing (ASC) Advisory Board at the Los Alamos National Laboratory (2021 – present)

MATH+ Scientific Advisory Board, Germany (2021 – present)

SIAM Activity Group on Data Science (Inaugural Program Director, 2021 – 2022)

Co-Chair, SIAM 2022 Conference on Mathematics of Data Science

AIAA 2022 SciTech Forum Executive Steering Committee (2021 – 2022)

AIAA Board of Trustees (2020 – 2023)

Institute for Mathematical and Statistical Innovation (IMSI) Board of Trustees (2020 – 2024)

Advisory Board, Center of Excellence on Sustainable and Energy Efficient Aviation, TU Braunschweig, Germany (2020 – present)

NSF Advisory Committee for Cyberinfrastructure (2019 – present, Co-chair 2020 – 2022)

Science Board, Santa Fe Institute (2019 – present)

SIAM Journals Committee (2019 – present)

External Advisory Board, Michigan Institute for Computational Discovery and Engineering, University of Michigan (Member, 2017 – present)

National Academies Board on Mathematical Sciences and Analytics (BMSA) (2016 – present)

Advisory Board, Girls' Angle (2014 – present)

Karen E. Willcox, MNZM, PhD

External Boards and Committees (past)

National Academies Planning Committee on the Workshop on the Frontiers of Mechanistic Data-Driven Modeling for Additive Manufacturing (2019)

AIAA Fellows Selection Committee (2019 – 2021)

SIAM Fellows Selection Committee (2018 – 2020)

Department of Energy Working Group on Basic Research Needs for Scientific Machine Learning (2017 – 2019)

National Academies Committee to Assess the Risks of Unmanned Aircraft Systems (UAS) Integration (2017 – 2018)

SIAM Committee on Science Policy (2016 – 2018)

SIAM Activity Group on Computational Science and Engineering: Vice President (2013 – 2015), Program Director (2011 – 2013)

Co-Chair, SIAM 2013 Conference on Computational Science and Engineering

Co-Chair, Institute-wide Task Force on the Future of MIT Education (2013 – 2014)

MIT OpenCourseWare Faculty Advisory Committee (2011 – 2018; Chair 2015 – 2018)

Advisory Board, Department of Engineering Science, University of Auckland (Member, 2008 – 2018)

National Research Council, Committee to Conduct an Independent Assessment of the Nation's Wake Turbulence Research and Development Program (2007)

National Academies Decadal Survey of Civil Aeronautics, Aerodynamics and Acoustics Panel (2005 – 2006)

AIAA MDO Conference Technical Chair (2011 – 2012)

AIAA Multidisciplinary Design Optimization Technical Committee (2001 – 2021); Chair (2019 – 2021); Vice-Chair (2017 – 2019); Awards Subcommittee Chair (2003 – 2006); Publications Subcommittee Chair (2011 – 2018)

Visiting Committees and Review Boards

ExxonMobil Corporate Strategic Research, Capability Assessment External Review Panel (Physics and Mathematical Science and Scientific Computing) (2021)

Review Committee, Research Assessment of Aerospace Engineering, Delft University of Technology, Netherlands (2020 – 2021)

Committee to Visit Harvard University Information Technology (Member, 2019)

National Academies Panel on Review of the Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) (2018)

Review Committee, TU Braunschweig Universities of Excellence, German Excellence Initiative (2018)

External Review Board, Computing and Information Sciences Research Foundation, Sandia National Laboratories (Member, 2016 – present)

Visiting Committee, Applied Mathematics & Statistics Department, Colorado School of Mines (Member, 2017)

HarvardX Review Committee, Harvard University (Member, 2016)

Board of Visitors, Institute for Computational Engineering and Sciences, University of Texas at Austin (2012 – 2018; Chair 2015 – 2018)

Assessment Committee, Accreditation of Aerospace Engineering, Delft University of Technology, Netherlands (2013)

Committee of Visitors, Division of Mathematical Sciences, National Science Foundation (Member, 2010)

Karen E. Willcox, MNZM, PhD

Editorial Boards

Acta Numerica (Editorial Board Member, 2021 – present)

IEEE Computing in Science and Engineering (CiSE) (Associate Editor, 2021 – present)

AIAA Journal (Editorial Board Member, 2021 – present; Associate Editor, 2015 – 2020 and 2009 – 2011)

Journal on Data Centric Engineering (Advisory Board, 2019 – present)

SIAM Journal on Scientific Computing (Section Editor, 2013 – 2019; Associate Editor, 2008 – 2013)

ASA/SIAM Journal on Uncertainty Quantification (Associate Editor, 2012 – 2013)

SIAM Book Series on Computational Science and Engineering (Editorial Board Member, 2009 – present)

Leadership Activities

Academic: Given hundreds of invited lectures in the US and internationally, including multiple plenary/keynote talks at major international conferences. In 2021 delivered plenary talks at AIAA Scitech Forum (largest aerospace engineering conference) and SIAM Conference on Computational Science and Engineering (largest computational science and engineering conference). Published over 120 papers in refereed archival journals. Supervised theses for 60 graduate students (40 M.S., 20 PhD). Multiple graduate students and postdocs hold academic positions at prestigious universities and leadership positions in industry. Secured funding and managed multi-institutional research projects from many sources including the U.S. Air Force, Boeing, U.S. Department of Energy, Federal Aviation Administration, NASA, National Science Foundation, DARPA, and U.S. Department of Education.

Major multi-institution research grants as lead include: Co-lead PI and Co-Director, AEOLUS Multifaceted Mathematics Capability Center on Advances in Experimental Design, Optimal Control, and Learning for Uncertain Complex Systems (Department of Energy, $10M total budget over 4 years). Lead PI, Multidisciplinary University Research Initiative (MURI) project on Managing Multiple Information Sources of Multi-physics Systems (Air Force Office of Scientific Research, $7.2M total budget over 5 years). Lead PI, MURI project on Machine Learning for Physics-Based Systems (Air Force Office of Scientific Research, $2M total budget over 3 years). Lead PI, RISE of the Machines: Robust, Interpretable, Scalable, Efficient Decision Support (Department of Energy, $4.4M total budget over 3 years). Co-lead PI and Co-Director, DiaMonD Multifaceted Mathematics Capability Center on Mathematics at the Interfaces of Data, Models, and Decisions (Department of Energy, $16.7M total budget over 5 years). Lead PI, Dynamic Data Driven Methods for Self-aware Aerospace Vehicles (Air Force Office of Scientific Research, $2.5M total budget over 6 years). Lead PI, Towards Scalable Differentiated Instruction using Technology-Enabled Competency-Based Dynamic Scaffolding (Department of Education, $2.9M total budget over 4 years).

Administrative: Director of the Oden Institute for Computational Engineering and Sciences at UT Austin (2018-present). Oversees Oden Institute operations involving >350 people, >$80M in active research contracts/grants, and >$150M endowment funding. Served as the founding co-director of the MIT Center for Computational Engineering (2008-2018) and the Associate Head of the MIT Department of Aeronautics and Astronautics (2011-2013). In Associate Head role, led reforms in the undergraduate degree program and put in place initiatives that successfully increased undergraduate enrollment in aerospace engineering.

Professional: Active professional service and leadership through multiple conference organizing committees, conference chair positions, technical committee leadership, organizational review committees, advisory boards, and editorial positions.

Karen E. Willcox, MNZM, PhD

Selected Awards and Honors

Best Paper Award, “Toward predictive digital twins via component-based reduced-order models and interpretable machine learning”, AIAA Multidisciplinary Design Optimization Best Paper, 2020

SIAM Student Paper Prize (E. Qian), “Multifidelity Monte Carlo estimation of variance and sensitivity indices,” 2020

Southwest Research Institute Best Student Paper Award (M. Kapteyn), “Toward predictive digital twins via component-based reduced-order models and interpretable machine learning,” AIAA Non-Deterministic Approaches Conference, Scitech Forum, 2020

Paper “Variance-based sensitivity analysis to support simulation-based design under uncertainty” one of the top 10 most accessed articles in *Journal of Mechanical Design* in 2019.

AIAA Fellow, Class of 2019

SIAM Fellow, Class of 2018

Conference on Neural Information Processing Systems (NeurIPS) paper “Contour location via entropy reduction leveraging multiple information sources” selected for Spotlight Presentation (3% of submissions), 2018.

Best Paper Award, “Towards a Low-Order Model for Transonic Flutter Prediction,” AIAA Theoretical Fluid Mechanics Conference, AIAA Aviation Forum, 2017

Member of the New Zealand Order of Merit (MNZM), 2017

Distinguished Alumni Award, University of Auckland, 2016

Member, Harvard Higher Education Leaders Forum, 2016 – 2019

SIAM SIGEST Award for paper “Goal-oriented inference: Approach, linear theory, and application to advection-diffusion,” 2013

Sir Peter Blake Trust Emerging Leader Award, 2010

Selected for National Academies Frontiers of Engineering Education Symposium, 2010

AIAA MDO Technical Committee Service Award, 2008 and 2013

J. T. Oden Faculty Research Fellow, University of Texas at Austin, 2006

New Zealand Management Magazine, Young Leader, 2006

MIT Junior Bose Teaching Award, 2005

MIT Department of Aeronautics and Astronautics Teaching Award, 2004

Best Paper Award, “A Framework for Aircraft Conceptual Design and Environmental Performance Studies,” AIAA Multidisciplinary Analysis and Optimization Conference, 2004

Chairman BOWMAN. Thank you so much, Dr. Willcox.
Dr. Monroe, you are now recognized.

TESTIMONY OF DR. CHRISTOPHER MONROE, CO-FOUNDER AND CHIEF SCIENTIST, IONQ, INC.

Dr. MONROE. Good morning. Thank you, Mr. Chairman, Members of the Subcommittee, for this opportunity to testify before you today. I'm here on behalf of IonQ, a company that builds quantum computers. IonQ is headquartered in College Park, Maryland, and was spun out of the University of Maryland and Duke University about 5 years ago. I'm also a Professor of Electrical and Computer Engineering and Physics at Duke University. I have over 2 decades of experience in the field of quantum computing technology from both academic and industrial perspectives, and I'm here to talk about the future of computing in terms of quantum information.

Quantum computers, as you may have heard, are—they're as revolutionary as they are challenging to grasp and build. Their might, given these challenges, demand special attention. As you know, the 2018 National Quantum Initiative, or NQI, was initiated by the House Committee on Science, Space, and Technology to ensure that the United States remains at the forefront of this technology. The NQI endowed the Department of Energy, the National Science Foundation, and the National Institute of Standards and Technology (NIST) with coordination from the Departments of Defense (DOD) and the intelligence community to stimulate foundational research in quantum computing and other quantum technologies and translate this technology from laboratory to industry.

So how does a quantum computer work? It's essentially not hard. It's just that quantum computers follow laws of physics that have no analogy in everyday life, so it's confounding. It's not exactly that it's hard. Information in quantum computers can exist in superposition; that is, multiple values can be stored and processed simultaneously in a single memory device. But each time you expand a quantum computer by just a single bit—we call them quantum bits or qubits—its power essentially doubles. So with just 300 quantum bits, that's a pretty small chunk of matter, a quantum computer can process more possibilities than there are atoms in the entire universe. This massive parallelism in quantum computers allows certain computations to be performed that could never be accomplished using regular computers.

So a few far-reaching applications for this new mode of computing include optimization of complex problems dealing with huge amounts of data, including logistics and things like pattern recognition; secondly, molecular and materials design for energy, medical and defense applications; and finally, security, including secure communication, encryption, and decryption or code-breaking.

IonQ has collaborative projects in all of these areas, and one thing that's interesting in this field is it's such an early stage of this technology that it's——

Chairman BOWMAN. Mr. Monroe, time has run out. Can you just finish that last point?

Dr. MONROE. Yes. I think somehow the clock never started, I noticed.

Chairman BOWMAN. Yes, that's probably right. That's weird.

Dr. MONROE. I think I was talking for about 2 minutes but——

Chairman BOWMAN. Yes, let's put 2 minutes on the clock if we can.

Dr. MONROE. OK. I'll go quickly. So it's critical that quantum computer builders co-design applications with the systems they build. It's really not physics anymore, but all of the physical and chemical sciences, all of the engineering fields, computer science, algorithm design, economics, and even social sciences. So it's no surprise that one of the most important applications in quantum computers is energy and that the Department of Energy is an important player in advancing this field.

So IonQ machines and those built by others are still too small to beat regular computers in these types of problems, but we're just at the beginning of this commercial phase, and this situation will change very soon.

The core of a quantum computer is exotic, and its key attribute is isolation. It involves devices either cooled to nearly absolute zero temperature—that's negative 460 degrees—or in the case of our technology at IonQ, we use individual atoms suspended in a small vacuum chamber and poke with laser beams. I should also note that this technology was developed at the National Institute of Standards and Technology in the 1990's where I worked with David Wineland in developing the first quantum logic gate.

So, as exotic as this is, the core technology is not necessarily the main challenge. The real challenge as I see it is creating the workforce to understand how to deploy quantum systems. I like to say at the universities we typically don't build components for people to use, but industry so far has been a little bit slow to develop this technology because they don't have a basis in quantum. And so this is, I think, historically where government laboratories can play a role. At IonQ, our systems are now available on cloud servers, and at Duke University we're setting up a quantum—Duke Quantum Center that will be a scientific user facility that will serve the scientific use cases. And this is very important for the field.

So I want to conclude. I think I—it was a lot faster than I thought it would be. But now is a critical time for DOE, NSF, NIST, DOD, and the intelligence community to redouble and coordinate their efforts in translating quantum computers to the real world. One example is the QUEST program that Congressman Lucas mentioned that would subsidize access to industrial quantum computers. Another is endowing the NSF with a technologically driven division mandate. This is a particularly good way to ensure that this emerging technology gets used in the field.

I'm a member of many advisory boards in Europe and in Asia, and I'm well aware of the coordinated investments overseas. The United States must lead the race to build quantum computers and other quantum technologies, and I think that the programs and continued stewardship of the National Quantum Initiative by DOE and the other agencies I mentioned are critical to continued leadership.

I again think the Committee and Chairman for his leadership and for the opportunity to testify today even though I took, I think, 4 minutes. Thank you.

[The prepared statement of Dr. Monroe follows:]

TESTIMONY OF DR. CHRISTOPHER MONROE
Chief Scientist
IonQ, Inc., College Park, MD
Professor of Electrical and Computer Engineering and Physics
Duke University, Durham NC

"Accelerating Discovery: the Future of Scientific Computing at the Department of Energy"
Wednesday, May 19, 2021
Energy Subcommittee
House Committee on Science, Space, and Technology

Mr. Chairman and members of the Subcommittee, thank you for this opportunity to testify before you today.

My name is Chris Monroe, and I'm here on behalf of IonQ, a company that builds quantum computers. IonQ is headquartered in College Park, Maryland, and was spun out of the University of Maryland and Duke University in 2016. I'm also a Professor of Electrical & Computer Engineering and Physics at Duke University, and a member of the National Quantum Information Advisory Committee, coordinated by the White House Office of Science and Technology Policy.

I have over two decades of experience in the field of Quantum Computing Technology, from both academic and industrial perspectives, and I am here to talk about the future of computing in terms of Quantum Information.

Quantum computers are as revolutionary as they are challenging to grasp and build. Their might, given these challenges, demands special attention. As you know, the 2018 National Quantum Initiative (NQI) was initiated by the House Committee on Science, Space, and Technology, to ensure the US remains at the forefront of this technology. The NQI endowed the Department of Energy, National Science Foundation, and National Institute of Standards and Technology, to stimulate foundational research in quantum computing and translate this technology from laboratory to industry.

So how does a quantum computer work? It's not hard – it's just that quantum computers follow laws of physics that have no analogy in everyday life. Information in quantum computers can exist in superposition – multiple values stored and processed simultaneously in a single memory device. But each time you expand a quantum computer by just a single bit (we call them quantum bits), its power essentially doubles. With just 300 quantum bits, a quantum computer can process more possibilities than there are atoms in the entire universe! This massive parallelism in quantum computers allows certain computations to be performed that can never be accomplished using regular computers.

Here are some far-reaching applications from this new mode of computing:

1. Optimization of complex problems dealing with huge data, including logistics and pattern recognition
2. Molecular and material design for energy, medical and defense applications
3. Security, including secure communication and code-breaking.

IonQ has collaborative projects in all these areas. For example, we are just starting a collaboration with Breakthrough Energy Ventures, a fund backed by Bill Gates and others, to investigate how Quantum Computers can be applied to issues associated with climate change.

It's no surprise that one of the most important applications of Quantum Computers is Energy, and that DOE is an important player in advancing Quantum Computing.

IonQ machines, and those built by others, are still too small to beat regular computers on these types of problems, but we are just at the beginning of the commercial phase of quantum computers, and this situation will change soon.

The core of a quantum computer is exotic, and its key attribute is ISOLATION. It involves devices either cooled to nearly absolute zero temperature (-460 degrees F) or, in the case of our technology at IonQ, individual atoms suspended in a small vacuum chamber and poked with laser beams. (Incidentally, this so-called "ion trap" technology was originally developed at NIST in the 1990s, where David Wineland and I demonstrated the first quantum logic gate. Wineland was later awarded the Nobel Prize in Physics, partly based on this work.)

Exotic as it is, our core technology is not necessarily the main challenge -- atoms are perfectly identical and sufficiently isolated from the environment, so we have a good idea how to scale our systems at IonQ. Instead, the challenge is that the way quantum computers process information is totally different – programming them, controlling their "bits", interpreting results, is a radical departure from regular computation. The challenge is that we need to bring up a generation of builders and users together to "think quantum" as the technology matures.

I see this juxtaposition every day, at both IonQ and in my academic research laboratories. At IonQ, our systems can be accessed by anybody via Amazon and Microsoft cloud servers; we also have more intimate partnerships with users and companies wishing to get up to speed in quantum for their future commercial needs. At Duke University, we are standing up a Quantum Computer User Facility, following programs funded by the National Science Foundation, the Intelligence Advanced Research Projects Activity, and the Department of Defense. The Duke Quantum Center will host researchers from across the US to spend weeks or months at our facility in downtown Durham NC, to use the most advanced quantum computers in the academic community. Like IonQ, the Duke machines are based on atomic qubits. Unlike IonQ, the Duke machines will be used for scientific purposes – from understanding black holes to investigating models of exotic materials. This "co-design" process, integrating users and builders for both commercial and scientific applications, will be key to unlocking the power of quantum computers.

At both IonQ and Duke, we have vital partnerships with several DOE labs, including Oak Ridge National Laboratory, Pacific Northwest National Laboratory, and Sandia National Laboratories. They are all interested in using our devices and fielding their own, but these labs also supply key components for our machines.

As a member of advisory boards for Canadian, European, and Asian quantum centers, I am well aware of the outsized and coordinated investments overseas, especially from China. The US must lead the race to build quantum computers and other quantum technologies. We must train and place the next generation of quantum scientists and engineers. We already have the best higher education and attract the best minds in the world. Through the opportunities presented by our policies, we should seek to keep these scientists and engineers in our universities, laboratories and industry and thereby accelerate our progress in Quantum.

Now is a critical time for DOE, NSF, NIST, the DOD and the Intelligence Community to redouble and coordinate their efforts in translating quantum computers to real-world applications such as energy production and climate science. This will create vast opportunities for workforce development, and economic growth in Energy, Medicine, and Security.

The Endless Frontiers Act, which would endow the NSF with a technologically-driven division and mandate, is a particularly good way to ensure this emerging technology will be used for scientific applications and that scientific developments will continue to inform the fabrication of commercial and academic devices and systems. Another example is the Quantum User Expansion for Science and Technology or “QUEST” program (also under consideration as HR 1837), which would subsidize users to get access to commercially available quantum computers right now, to stimulate their future co-design.

These programs and continued stewardship of the National Quantum Initiative by the DOE, NSF, NIST, and the DOD and IC, are critical to continued American leadership in quantum computing and the future of advanced computing.

I again thank the Chairman and Committee for its leadership and for the opportunity to testify today. I look forward to answering your questions and working with you in the future.

Dr. Christopher Monroe

CEO and co-Founder at IonQ, Inc., a start-up company that is developing and commercializing the world's first fully-expressive, full-stack quantum computer, based on trapped atomic ions.

Lead a large experimental research group at Duke University, fabricating and using quantum computers and for fundamental studies of quantum physics and quantum entanglement. Investigate the storage and processing of quantum information using individual electromagnetically confined atoms and photons, the communication and teleportation of quantum information, and the use of quantum systems to simulate the complex behavior of magnetic materials. Tools include advanced laser sources, photonic technology, fast electronics, semiconductor structures, and stable microwave and radiofrequency sources.

Chairman BOWMAN. Thank you, Dr. Monroe. Apologies about the issues with the clock.

Dr. Kamara, you are now recognized.

TESTIMONY OF DR. SENY KAMARA, ASSOCIATE PROFESSOR, BROWN UNIVERSITY

Dr. KAMARA. Thank you. Chairman Bowman, Ranking Member Weber, Chairwoman Johnson, Ranking Member Lucas, and distinguished Members of the Committee, I appreciate the opportunity to testify at today's hearing on the future of scientific computing at the Department of Energy.

By the end of the year, the Oak Ridge National Lab will receive the world's first exascale supercomputer. This computer will be able to process 10 to the 18 or one quintillion operations per second. It is hard to overstate how difficult this is to achieve and what an accomplishment it is. This considerable leap in computing power will open the doors to new discoveries and significantly impact a multitude of fields, including medicine, meteorology, cosmology, and artificial intelligence.

It is clear that the world-class research and high-performance computing that has been conducted by U.S. universities, national labs, and industry in order to achieve exascale computing will affect our lives for the foreseeable future. But as we enter the era of exascale computing, I would like to provide a word of caution. I'm sure we can all agree that computing and the technologies it enables have had a tremendous impact on society. Because of this, it is easy to assume that technological progress always leads to positive outcomes and that new technologies benefit everyone equally. But this is not the case. Technology, like policy, can have disparate impact. It can enable positive outcomes for some and cause great harms to others.

Consider, for example, advances in facial recognition which allow us to log into our smartphones faster but also enables suspicionless mass surveillance with the progress in computer vision and robotics that enables new drones that can deliver medicine to hard-to-reach rural areas or missiles at the push of a button by somebody sitting in a room thousands of miles away. We must always remind ourselves that technology is not inherently good and does not always benefit everyone equally by default. In fact, we need to think hard about the harms technology can cause and work even harder to mitigate those harms.

One of the many important applications of exascale computing is artificial intelligence and machine learning, for example, to predict how a cancer patient might respond to a particular treatment. But as we know, thanks to the work of scholars like Cathy O'Neil, Joy Buolamwini, and Timnit Gebru and to outlets like Pro Publica, machine learning algorithms can be biased and can exhibit different behaviors on different populations. And as has been widely documented, these biases in machine learning most often harm people of color and those from marginalized communities.

So while we should appreciate that thousands of world-class scientists and engineers across the country are diligently working toward making exascale machine learning for cancer a reality, we also have to ask how many are working on ensuring that these cancer treatment prediction models work for people of all genders and of all races?

The investments we are making in exascale computing will improve national security, the U.S. economy, and industry, but will everyone benefit equally from this investment? Will the 13-year-old girl from Washington Heights, New York, benefit from this investment as much as the tech, energy, and pharmaceutical industries? Will there be as much effort to use these supercomputers in the fight against sickle-cell anemia as other diseases?

Exascale computing is not only an incredible achievement but it's an incredible resource with the power to shape our lives and those of future generations. As such, we must be careful and thoughtful about how we make use of it. In particular, it is incumbent upon us to make sure that we deploy and use this resource in a manner that is fair and inclusive that benefits not only the powerful but those who have historically been marginalized by society and by technology.

Thank you, and I look forward to answering your questions.

[The prepared statement of Dr. Kamara follows:]

Statement by Seny Kamara

Associate Professor
Department of Computer Science
Brown University

before the

U.S. House of Representatives
Committee on Space, Science and Technology
Sub-committee on Energy

May 19, 2021

Chairman Bowman, Ranking Member Weber and distinguished members of the committee. I appreciate the opportunity to testify at today's hearing on the future of scientific computing at the Department of Energy.

I am an Associate Professor of Computer Science at Brown University, where I conduct research in cryptography and privacy and direct the Encrypted Systems Lab and co-direct Brown's Computing for the People project. I am an affiliate of the Brown Data Science Initiative, the Brown Center for Human Rights and Humanitarian Studies and the Policy Lab. I am also a Principal Scientist at MongoDB, a company that provides one of the most widely used platforms to store and process "big data". Prior to this, I was a research scientist at Microsoft Research. I am testifying today in my capacity as an academic researcher.

By the end of the year, the Oak Ridge and Argonne National Laboratories will receive the world's first exascale supercomputers. These computers will be able to process 10^{18}—or a quintillion—operations per second. It is hard to overstate how difficult this is to achieve and what an accomplishment this is. This considerable leap in computing power will open the doors to new discoveries and significantly impact a multitude of fields including medicine, meteorology, cosmology and artificial intelligence. It is clear that the world-class research in high-performance computing that has been conducted by US Universities, National Labs and Industry in order to achieve exascale computing will affect our lives for the foreseeable future.

But as we enter the era of exascale computing, I would like to to provide a word of caution. I am sure we can all agree that computing and the technologies it enables have had a tremendous impact on Society. Because of this it is easy to assume that technological progress always leads to positive outcomes and that new technologies benefit everyone equally.

But this is not the case. Technology—like policy—can have disparate impact: it can enable positive outcomes for some and cause great harm to others. Consider, for example, advances in facial recognition which allows us to log into our smartphones faster but also enables suspicionless mass surveillance. Or the progress in computer vision and robotics that enables new drones that can deliver medicine to hard to reach rural areas or missiles at the push of a button by someone sitting in a room thousands of miles of away. We must always remind ourselves that technology is not inherently good and does not benefit everyone equally by default. In fact, we need to think hard about the harms technology can cause and work even harder to mitigate those harms.

One of the many important applications of exascale computing is artificial intelligence and machine learning; for example to predict how a cancer patient might respond to a particular treatment. But as we know thanks to the work of scholars like Cathy O'Neil, Joy Buolamwini and Timnit Gebru and to outlets like Pro Publica, machine learning algorithms can be biased and can exhibit different behaviors on different populations [3, 2, 1]. And, as has been widely documented, these biases in machine learning most often harm people of color and those from marginalized communities.

So while we should appreciate that thousands of world-class scientists and engineers across the country are diligently working towards making exascale machine learning for cancer a reality, we also have to ask: how many are working to ensure that these cancer treatment prediction models work for people of all genders and of all races? The investments we are making in exascale computing will improve National Security, the US Economy and Industry. But will everyone benefit equally from this investment? Will the 13 year old girl from Washington Heights, New York, benefit from this investment as much as the Tech, Energy and Pharmaceutical industries? Will there be as much effort to use these supercomputers in the fight against sickle cell anemia as other diseases?

Exascale computing is not only an incredible achievement but an incredible resource with the power to shape our lives and those of future generations. As such, we must be careful and thoughtful about how we make use of it. In particular, it is incumbent upon us to make sure that we deploy and use this resource in a manner that is fair and inclusive; that benefits not only the powerful but those who have historically been marginalized by society and technology.

Thank you. I look forward to answering your questions.

References

[1] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias, 2016.

[2] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR, 2018.

[3] Cathy O'neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown, 2016.

Seny Kamara is an Associate Professor of Computer Science at Brown University, where he co-directs Brown's Computing for the People project and the Encrypted Systems Lab. He is also affiliated with Brown's Center for Human Rights and Humanitarian Studies, the Data Science Initiative and The Policy Lab. Kamara is also a Principal Scientist at MongoDB, a company that provides one of the most widely used platforms to store and process data. Prior to this, he was a Research Scientist at Microsoft Research in the Cryptography Research group.

Professor Kamara conducts research in cryptography with a focus on problems motivated by social and policy issues; especially issues that impact marginalized groups. His pioneering work on encrypted search algorithms laid the foundation and pushed the state-of-the-art of end-to-end encrypted database technologies. His work has consistently appeared in the top peer-reviewed venues in data security and cryptography and has been featured in numerous outlets including Wired, Forbes and The Register.

At Brown, Professor Kamara teaches "Algorithms for the People", a course that surveys, critiques and tries to address the ways in which computer science and technology affect marginalized communities.

In 2016, he was appointed by the National Academies of Sciences to study the impact of end-to-end encryption on law enforcement and intelligence and in 2019 he testified to the Financial Services Committee of the U.S. House of Representatives about the privacy and fairness implications of Big Data. In 2020, he was appointed by the National Academies of Sciences to study the future of encryption. He has received a Google Faculty Award and was named a Leadership Fellow by the Boston Global Forum for his work and commitment to global peace.

Chairman BOWMAN. Thank you, Dr. Kamara.

At this point, we will begin our first round of questions. The Chairman now recognizes himself for 5 minutes.

Dr. Kamara, I'm going to start with you. Thank you so much for your testimony and for your attention to making sure that all people benefit equally from investments in scientific computing and other new technologies. And I understand that you are cofounding a new research institute at Brown University called Computing for the People. Can you talk more about why we should incorporate these kinds of questions into the R&D agenda from the very beginning? What can we do as Congress to design research programs that will prevent harmful applications down the road, keep the needs of marginalized communities and all people in mind, and prepare the computing workforce to engage with these issues as well?

Dr. KAMARA. Yes, so, as I said in my statement, it's clear that computer science and technology have had a huge impact, but the reality is that, as a field, we haven't really centered the problems of marginalized groups. It's just not something that comes naturally to the field. And there's many reasons for that. Some include the lack of diversity in computing, which is something that is well-documented. And so there's no natural way for—you know, for computer science research and technology research to really address those problems. And so this is why we're building this institute at Brown, and the motivation is to really make it a priority, right, understanding the problems that marginalized communities face and how technology can help and really focusing on that as our main motivation. So that's what we're doing.

The way that Congress can help is both in funding this kind of research and in asking questions, right, making sure that the needs of all people are addressed by technology and by computing research, just as we're doing in this hearing today.

Chairman BOWMAN. I was muted. Sorry about that. How much is a deep dive understanding and analysis of implicit bias in algorithms and computer science a part of this work that you're referring to?

Dr. KAMARA. Yes, it's crucial. So when we design—and it's—and I also want to highlight that it's not only at the level of research. It's also the level of education. So when we teach our students computer science and we teach them how to design algorithms, we're not teaching them how to think about bias and how to address it. We're not teaching them how to think about energy efficiency and how to design algorithms that are not only fast but also that minimize the amount of energy consumption. So these are all things that we want to do in our institute. And basically we want to integrate what we call responsible computing into computer science education and into computer science research as well.

Chairman BOWMAN. OK.

Dr. KAMARA. Learning how to find biases in data, learning how to find biases in algorithms is a crucial part of that.

Chairman BOWMAN. Thank you very much.

Dr. Tourassi, thanks for your testimony as well. I was intrigued by your suggestion for investing in a national data infrastructure that could be housed at DOE labs and that could play an important role in democratizing AI. Can you say more about what this would look like?

Dr. TOURASSI. Absolutely. We know that the explosive growth of AI is based on the three pillars, the supercomputing, the algorithms, and data. We talk a lot about investments on the supercomputing side and algorithms, but data is the fuel that will make the engine—the airplane fly. And, as I pointed out in my written and oral testimony, the past few years have taught us the importance of having a data infrastructure that makes the most of our Federal data assets, examples with the partnership with the National Cancer Institute, as well as with the Veterans Administration.

Building on that thread, because that infrastructure was in place and the interagency partnership was in place with the VA, we were able to pivot fast to address challenges with the COVID–19 pandemic because the infrastructure was in place to accept data, new data related to COVID–19 cases in the veteran population to start doing within 48 hours large-scale epidemiological studies and observational studies.

So this is what I meant in my statement that this was an important lesson learned that we need to be proactive and to put all resources necessary to support that infrastructure. DOE has a long history of building and sustaining successfully data infrastructures and enable—enabling the broader scientific community to make the most of it.

Chairman BOWMAN. Thank you so much, Dr. Tourassi.

I now recognize Mr. Weber for 5 minutes.

Mr. WEBER. Thank you, Chairman. I appreciate that. And I'm going to go back to you, Dr. Tourassi. We're going to keep you on the hot seat for a little longer.

First of all, I'd like to congratulate you on all your hard work at Oak Ridge as Director of its Leadership Computing Facility. I'm sure you just didn't wake up one day and decide you were going to get into something like that. You had probably been studying and working at that a long time. And I know I speak for all of us when I say I cannot wait to see Frontier in action this fall, so congratulations.

You've noted, Doctor, that Frontier will on day one of its operation no less, be ready to run applications across two dozen scientific and technical disciplines. Can you take a moment for us not-so-technical people maybe to explain why, when it comes to the race to the world's fastest supercomputer that the hardware development can't be our only focus? Talk to us, please.

Dr. TOURASSI. Absolutely. As you said, the tool by itself is not the enabler. We need to see the scientific impact. And scientific impact is measured in many different ways. First of all, by the breadth of application domains that benefit by the tool, how much faster we can do scientific discovery across that breadth of applications, and also how efficiently we use the tool.

So in partnership with the Exascale Computing Project, there is a portfolio of applications that have many critical mission areas, some of which I highlighted in my report. And we are working very closely with the exascale computing team to make sure that the necessary software tools will be in place to enable the scientific community that will need them, as I said, from day one. Some of these applications in climate change, in renewable energy, in advanced manufacturing, in biology and public health because we know that these are pressing application areas with great societal impact.

Mr. WEBER. Absolutely, thank you. How does the utility or usability of DOE's supercomputer like Summit, for example, we're talking about competitors. How does it compare to some of our competitors' supercomputers?

Dr. TOURASSI. So, clearly, these metrics are not monitored and publicly known, but we can say for sure for the United States and the leadership computing facilities not only we are oversubscribed, we operate in the high 90's efficiency. Anecdotally, some of our competitors are certainly behind, so they have the tool, not necessarily the most effective use of the tool. But, as I said, this is a race. We cannot just relax in that race. We need to keep moving forward and making sure that we are spearheading the specific domain.

Mr. WEBER. Well, thank you for that. I'm quite sure China has all the information on everybody's capabilities, so we know that.

So we do want to keep up with our competitors. We'd love to have that information. And point out to us why it really matters. I know we want to be first in the race, but what other examples would you use other than just time and being first? Why does it matter that we get ahead of our competitors? Give us some examples that tie to our country. Can you do that?

Dr. TOURASSI. Well, certainly the issues of national security are extremely important, and we know that supercomputers have played a very important role in addressing applications in the national security space, and we expect that space to grow even further.

As the other panelists and you all mentioned the issues of adversarial use of the technologies we develop, so we'd need to be very much aware that this is not only in our horizon, it is present.

Mr. WEBER. Well, perfect. That's kind of where I hoped you were going.

I'm going to jump over to you, Mr. Binkley—or Dr. Binkley, I'm sorry. As I mentioned in my opening statement, I believe that in order to fully realize the potential of the world-leading computing capabilities stewarded by our national laboratories, the Office of Science must prioritize expanding ASCR research partnership with other Federal agencies. Very quickly, Department of Energy's position to conduct scientific computing research to resolve diverse challenges, how so in your opinion? I'm going to be over time a little bit.

Dr. BINKLEY. So we have in fact over the last half a dozen years really put priority on growing the ASCR program. It was clearly in recognition of the need to develop and then deploy exascale. And, as has been elaborated on here in the last 5 or 6 minutes, concurrent with the development of the exascale computer systems has been a very concerted effort to develop applications that are ready—that will be ready on the day the machine is first turned on.

Mr. WEBER. Yes.

Dr. BINKLEY. The other thing that we've done is we have worked through the National Science and Technology Council processes to coordinate closely with other Federal agencies, including the National Science Foundation, NASA (National Aeronautics and Space Administration), NOAA (National Oceanic and Atmospheric Administration), et cetera, and have tried to encourage the other agencies to also make investments in the high-performance computing capabilities for their specific applications. And in the case of the National Institutes of Health, as Dr. Tourassi has pointed out, we—one of the exascale applications is focused on cancer problems stemming from the National Institutes of Health. And I'll stop there.

Mr. WEBER. All right. Well, thank you. I appreciate that. And then, Mr. Chairman, thank you for the indulgence.

STAFF. Ms. Bonamici is next.

Ms. BONAMICI. Thank you so much to Chair Bowman and Ranking Member Weber, and thank you to the witnesses.

In the district I represent in northwest Oregon, researchers at Intel are developing the foundation for exascale computers, commercially viable quantum systems, and also partnering with the Department of Energy to advance other high-performance computing technologies. And I know these efforts will help us transition to clean energy economy, better predict extreme weather events, strengthen preventative medicine, improve emergency response, and more.

So, Dr. Binkley, it's sometimes challenging to conceptualize the benefits of the Department of Energy's work on scientific computing. I am the Co-Chair of the House Oceans Caucus, so I especially appreciated your specific example of the work Lawrence Berkeley National Laboratory (LBNL) and Los Alamos and the University of Bristol are doing to provide more accurate pictures of retreating ice sheets contributing to sea level rise. How will future exascale capabilities strengthen our understanding and response to the climate crisis, and how can Congress better support this work expanding SciDAC (Scientific Discovery through Advanced Computing) partnerships across the DOE?

Dr. BINKLEY. So let me take your last question first. The—with the incoming Biden-Harris Administration, there has been a reorganization within the Department of Energy that brings the applied energy programs back into the Under Secretary for Science organization, and that gives us very close coupling with the applied energy programs. That is between the Office of Science and the applied energy programs. And I'm hopeful that that will set the stage where we can expand SciDAC to reach the other parts of the Department.

Going back to the question about climate simulations, our Earth systems model, which has been in development now since about 5 or 6 years and is slated to be, you know, up and running on the exascale computer in the fall, we are systematically increasing the resolution of that model by use of the—you know, the power of the exascale computers to have higher resolution, more predictive models of climate effects and then, you know, we will be able to predict areas that are going to be problematic in the future. You know, that will become, I think, a fairly standard tool for predictive earth systems modeling. I'll stop there.

VOICE. You are muted.

Ms. BONAMICI. Sorry about that. Dr. Binkley, thank you. I do want to get in another question.

Dr. Kamara, we know that technology is not developed or used in a vacuum. The growing body of evidence suggests that, left unchecked, digital tools can absorb and replicate systemic biases that are ingrained in the environment in which they are designed. For example, one hiring algorithm often referenced for its apparent biases identified high-performers as anyone named Derek who played lacrosse even though those features had no connection for the actual jobs which the firm was screening. Unfortunately, digital tools are opaque in their design and operation, and it's hard to hold it accountable. So what steps can Congress take to prevent harmful bias in the research, development, and commercialization of supercomputing technologies? And will advances in neuromorphic computing help to end/or minimize bias or simply make it harder to detect?

Dr. KAMARA. Yes, so there are a lot of ways that bias can creep into algorithms, and there's a lot of research going on on how to mitigate that, so a lot of people are thinking hard about that. And that's great, but what we also have to be careful of is that just because we can automate something doesn't mean that we should, right? So it's not just—it's not enough to just say, OK, well, let's design unbiased algorithms. We can make some effort, you know, toward that, but we also have to ask ourselves should we be using algorithms, you know, in these particular cases, right? So should we be using algorithms for—you know, for estimating the probability that someone is going to commit a crime? Is that the right use of this technology? OK. So there's a lot of different sort of aspects of this that we have to think about.

And the way Congress can help is in funding more research on fair algorithms and also on providing some structure for auditing algorithms. That is also an important component of this is that we need to be able to say, OK, well, you know, if an algorithm is going to be deployed, right, if we are going to make the decision to deploy some kind of algorithm in a particular use case, then we really should have rules in place for how we're going to audit those algorithms, how we're going to make sure that they don't discriminate and that they are not biased.

Ms. BONAMICI. Great. And your thoughts on neuromorphic computing and what that means to bias and minimizing bias?

Dr. KAMARA. Well, that's a little bit outside of my scope of expertise, but I would definitely say that, you know, these different types of sort of—as we—as we diversify the kinds of algorithms that we're using, it's going to be harder and harder to detect bias, right? So we also have to really be really, really mindful of this. It's easy to think that, well, just because an algorithm is running or some kind of algorithm—there's many different types of—that, you know, there's sort of this veneer of objectivity, but this really isn't the case, right, and so we have to be really vigilant against that.

Ms. BONAMICI. Thank you, Mr. Chairman. I yield back.

STAFF. Mr. Baird is next.

Mr. BAIRD. Thank you, Mr. Chairman, and thank you, Ranking Member Weber, for holding this session. It's very interesting and exciting to me. Of course, my background is agriculture, and so I'm going to relate agriculture to this quantum computing.

I recently introduced a bill, 2961, that really says that the *Department of Energy Biological Innovation Opportunities Act*, and that's making sure that DOE has sufficient infrastructure to be able to provide access to quantum computing to researchers and university people, as well as the industry. So I was trying to mention that because it'll be a part of Ranking Member Lucas's *SALSTA* bill and keeping us in the leadership around the world.

But, Dr. Monroe, you noted in your written testimony that translating quantum computing to practical applications will create opportunities for the workforce. I would also suggest that it will provide students and researchers the ability to analyze large-scale, complex, practical situations. And an example of this is the National Institute of Food and Agriculture's Genome to Phenome Initiative. It's designed to look at plant materials and give plant researchers the ability to analyze those plants and look for the visual characteristics and tie that to the genome. And that all is designed to improve our ability to raise food and feed a hungry world in the future.

So, Dr. Monroe, would you care to elaborate—you were talking about practical applications and the need for that. Would you care to elaborate on what quantum computing can do for us in agriculture and how you think that fits into the progress of our world?

Dr. MONROE. Yes, thank you for that question. So, yes, sorry for my garbled statement. I think I had to kind of jump around a little bit on the hot seat there.

So quantum computers appear to be good at solving generic optimization problems, that is, problems that have way too many inputs that we can't sample all of the configurations on a conventional computer. Those are called combinatorial optimization problems, and we take guesses with even high-performance computers. And as I hinted, there are problems out there that we'll never be able to solve because the number of configurations is just way too big.

And so you mentioned agriculture, but I would maybe broaden in it to pharmaceutical, energy, and, you know, gas and oil industry. The need to understand the structure of molecules, this is something computers are also very poor at because even a small molecule, if it has more than a few hundred electrons—and this is—you know, even a caffeine molecule has 100 electrons—we can't easily model that molecule to see how it interacts with others to form catalyzers, better fuels, to form better drugs even.

But at the same time there's a logistics problem even in big pharma if you have 10,000 compounds, which combination of 10 of them make a good drug for something. We may have models of that, but, once again, we can't optimize that.

And on materials science side—and this is closely related to what you had mentioned in plant science and agriculture—can we develop new materials that harvest sunlight much more efficiently than known materials today? It's—this—for the same reason. We can't compute or simulate the behavior of these very fundamental things.

So, you know, quantum computing is just starting, but that—those are the types of problems that are naturally attackable by a quantum computer. And I'll put it this way. Those problems will never be solved using conventional computers. If they are to be solved, they will demand a quantum computer. So it is a very researchy field now, but we're starting to launch into building devices. And the Department of Energy over the last half-dozen years has played a big role in their laboratories in trying to translate that research, basic research to product. And I think industry is starting to play a role as well, very big-named industries that you've heard of.

Mr. BAIRD. You're exactly right. You know, this genome to phenome, they take these spray rigs that are 60 foot wide with cameras down into look at each and every plant, and there's no way in my background in the past that we could have had the ability to analyze that much data, so thank you for that.

And I see I've got about 13 seconds left, and I have other questions I would like to ask the other witnesses, but I thank all of them for being here and appreciate this opportunity, and I yield back.

STAFF. Mr. McNerney is next.

Mr. MCNERNEY. Well, first of all, I want to thank the witnesses. It's great to hear your viewpoints and share the excitement about what's going on with these exascale computers.

Dr. Monroe, you mentioned some of the potential applications, including pattern recognition in huge data sets, molecular and material design, and secure communications. The DOE and several other agencies are engaged in the National Quantum Initiative, which was authorized by Congress in December 2018. This initiative includes basic research activities and establishes large-scale Quantum Information Science Research Centers. How have these activities and the DOE research ecosystem in general helped strengthen the Nation's quantum computing industry?

Dr. MONROE. Thanks for the question. Indeed, I think DOE has shown in force their establishment of five very highly scoped centers distributed throughout the country that each have a separate mission and aim in the general field of quantum information, not just quantum computing but quantum sensing, things like quantum simulation. And these—I think the DOE laboratories are a great place, think about this, because the DOE labs in a sense are that ideal combination of having seasoned engineers, device people, but also not being a stranger to the weird laws of quantum physics that underlie these devices. So things are just getting started.

I—as a disclaimer, I am a key contributor to a DOE quantum center headquartered at Lawrence Berkeley National Laboratories in your State, and it's a very exciting collaboration that will—our consortium will build systems, and we're interested in sort of the system engineering of—when you build a big system, it's a—sort of breathing thing that's not just a sum of the individual parts. And with the DOE laboratory at LBNL and also Sandia National Laboratory, New Mexico is involved, we hope to really further the field.

But I should also mention the National Science Foundation and NIST are big players. They have historically been at the field since the beginning in quantum, and they—you know, one beauty of our national system in funding science is that we have many agencies, all with different missions, and, you know, DOE obviously I think the centerpiece are the laboratories. NSF, they are more vertically organized in a certain sense. They can bring together physicists with computer scientists and everything in between to come to bear on this. So I think, you know, collaboration between those agencies is really what's going to keep the American in the lead in this field.

Mr. MCNERNEY. Excellent. Dr. Binkley, you first discussed the importance of strengthening the core research that feeds into technologies like quantum sensors and networks. Would you discuss that a little bit, please?

Dr. BINKLEY. Yes, the—one of the major premises of the DOE activities in quantum information science is to be really focused on the fundamental science first and then follow that through with technology developments to get practical applications. And so, you know, there—if you look across the Office of Science portfolio of activities, you know, it's chemical and material sciences—well, know, those are systems that are governed by the laws of quantum mechanics. It's nuclear physics and high-energy physics governed by the laws of the standard model. And, essentially, quantum mechanics permeates through all of the physical sciences activities that are supported by the Office of Science. And, you know, that is, I think, a major focus of how we're organizing our QIS activities.

Mr. MCNERNEY. Well, thank you. Moving on, Dr. Monroe, in your written testimony you note the proposed QUEST program. How much annual funding is needed for this program? Do you think it would make sense to ramp up the program from a relatively smaller amount over the first few years?

Dr. MONROE. I think there's some amount of sense there. I think quantum computing systems that are capable, that itself is ramping up right now. I think things started maybe 4, 5 years ago, very small systems, and these are provided by companies you've heard of, some maybe you haven't like my company IonQ but also, you know, Google, IBM, other companies are putting these devices out there. They're very expensive and, you know, I speak from experience that industry, the bet that quantum computing will have a commercial payoff, it's—you know, it's going to be a long-term bet, and it's not easy for companies to play that risk, and I think this will allow the Department of Energy and the U.S. Government to help connect these systems to users and help subsidize the use of those machines.

And I'll also say—and one thing that—the wildcard to me is that the killer application for quantum computing, I don't—I'm not sure when it happens. I'm not—I can't predict what it will be, but I'm pretty sure that it will happen out of left field from somebody that thought of a problem that I don't know about or none of us makers know about, and getting them connected to the system will just hasten progress in the field because as soon as we hit that killer application, I think it's just going to explode.

So, you know, your mentioning of ramping up the access program. I certainly think that's maybe not a bad play, but I think it needs to get going. We need more industrial involvement in this field, less risk. Thank you.

Mr. MCNERNEY. Well, I've run over my time, so I'm going to yield back at this point.

STAFF. Ranking Member Lucas is next.

Mr. LUCAS. Thank you. As I've made clear, I believe that ensuring strong support for the Department of Energy and its world-leading national lab system is essential for our global leadership in science and technology. And this is one of my highest priorities as Ranking Member of the Science Committee.

So I ask my first question to Dr. Monroe and then I'd like the rest of the panel's comments, too. And bear in mind that part of our responsibility as Ranking Member and my colleague as Chairman of this Committee is not only to try and create the right policy here in the Science Committee but also we have to be able to persuade our colleagues in the body as a whole and ultimately the American taxpayers that we are on the right track.

So I ask the following question in that vein. In your respective areas of expertise, what would it mean to U.S. leadership in advanced scientific computing if we fail to provide adequate support for DOE Office of Science? Start with you, Dr. Monroe, and whoever in the panel would care to touch on that.

Dr. MONROE. Sure, I'll——

Mr. LUCAS. And this is a message from my colleagues who may not be spending time listening to this Committee hearing, which flabbergasts me, I acknowledge.

Dr. MONROE. Sure, I'll answer very briefly. In my own field of quantum computing that, again, we're at the early stages of this field, and it is absolutely critical that industry eventually take it over. And I think the DOE labs in particular but also university laboratories supported by NSF and NIST under the NQI is essential that they are able to make the investments necessary to keep the United States ahead of the world in this field. We're ahead now. We have mighty industry just waiting in the wings to take it over. We have to translate it. It's absolutely critical given the coordinated investments from folks around the world.

Mr. LUCAS. Anyone else care to take a stab?

Dr. TOURASSI. So I could go next. Certainly if we——

Mr. LUCAS. Please.

Dr. TOURASSI [continuing]. Start with investments, essentially, we are going to stagnate scientific innovation. We will stop innovating not only across basic sciences but also across applied sciences. Since we are using quantum computing as an example, as Dr. Monroe said, this is the beginning of a promising and disruptive technology, but it will take more than a decade for that technology to become reality for all of us. What happens in between? And it is classical computing that will lead to those innovations of materials, new materials that are needed to advance quantum computing. So it is a relay, and we cannot just pause and wait for the next big thing to come out from another nation.

Dr. BINKLEY. I could go next. There's another dimension to this that I think is really important to keep in mind, and that is that to be first in innovation, to be first in economic security, and so on, national security, you know, it's really a race for getting the best and brightest people into the United States to do new research in areas like quantum information science that have not been a topic of heavy investment before. You know, it's essentially—the best tools attract the best people, and that works in our advantage in an international context. And so having, you know, the best workforce ultimately is the—something that we need to strive for.

Dr. WILLCOX. I echo Dr. Tourassi's comments and to say that under this the Office of Science would absolutely stagnate and in fact set back our ability to tackle the critical challenges in energy security and environment particularly, and the Nation, the world just can't afford that.

Mr. LUCAS. With that, I thank you. You make very compelling cases. My time is winding down. I would observe one other the thing. We are very sensitive in Congress these days about not only international competition but how we protect U.S. research from theft, at the same time encouraging transparency and a cooperative environment. Anyone in a few seconds who'd like to touch on that, I'm more than pleased to hear what you have to say. But that's the struggle we're also facing, how to give you the tools, create the environment, but at the same time preserve your good work and the work being supported by the American taxpayers.

Dr. BINKLEY. Well, I would add, Congressman Lucas, that's something that we are really focused on. In fact, the meeting I was in just prior to this hearing was a National Science and Technology Council meeting focused on research security. And the—I can attest to the fact that a number of—in fact, all of the major science funding agencies of the Federal Government are really focused on this problem these days, and it's a very difficult problem. You know, we're trying to develop policies that will protect the U.S. interest. It's a very thorny issue. And in the Department of Energy, we've already—because we have such a large research establishment in our national labs, you know, we have over the last 3 or 4 years begun implementation of policies to help protect the results of research from being, you know, taken illicitly.

Mr. LUCAS. Thank you, Doctor, and thank you to the entire panel. My time is expired. I yield back, Mr. Chairman.

STAFF. Ms. Stevens is next.

Ms. STEVENS. Great, thank you so much. Congresswoman Haley Stevens from southeastern Michigan and a major fan of supercomputers and supercomputer technology and today's hearing. I couldn't be more grateful.

In particular, I've had the privilege of working on a supercomputer program called the National Digital and Engineering Manufacturing Consortium working in partnership is a public-private partnership between Purdue, OSU (Ohio State University), and a handful of small businesses in the Great Lakes area funded through the Economic Development Administration.

Also, I think I'm getting some feedback, Weber. That might be you. I think I'm getting Weber's chit chat, which I love hearing, but, you know, I want——

Chairman BOWMAN. Mr. Weber, please mute your mic.

Ms. STEVENS. OK, great, thanks. Thank you, Mr. Chair. Thanks.

So we get this NDEMC (National Digital Engineering & Manufacturing Consortium) program. We saw it with Jaco Plastics was able to develop a new product line resulting in 100 new jobs.

And, Dr. Tourassi, first of all, it's always a privilege to have somebody from Oak Ridge at a science hearing, and we have a tremendous amount of respect for your capabilities and partnerships out of Oak Ridge and all that you have represented. And your hearing touched a little bit on this, but I was just wondering if you could shed some additional light—and this could also be open to anybody, but I'm just feeling like Oak Ridge might have some insight into this with supercomputer technologies, public-private partnerships, and also abilities to lead to the creation of jobs either at small enterprises or large and what else we could be doing in Congress to advance these opportunities.

Dr. TOURASSI. So, as you know, building the supercomputers represent the strong partnership with our vendors, and from that point on, it's what kind of science we enable, right? These are the scientific innovations that will lead to technology transfer, and they will lead to also job growth opportunities.

We talk a lot about artificial intelligence these days and about what artificial intelligence will do to our workforce, both positive and negative effects. And secondly, when I'm—I'm an optimist and a proponent of developing a computational savvy, aware, and ready workforce. And I see that all of the leadership computing facilities being pivotal in that space as well. It is not only how to use the computers, it is not only how to do the science, but how exactly we develop the workforce that can support the data infrastructure that we mentioned earlier. And that is a different flavor of data—of the workforce that is absolutely needed for our Nation. So I would encourage all of you to keep thinking along those lines as well. It takes a lot of effort to collect, curate, and manage data. These are technical expertise. And I know that we don't have enough in the Nation, so we should be building that workforce and create opportunities for smaller companies to play a role in that.

Ms. STEVENS. Yes, and I—you know, we could go on here, too, because part of why we had this NDEMC program, the National Digital Engineering Manufacturing Consortium, was particularly because of cost. And that was about 10 years ago now, so I don't know if any of our other panelists while I have about a minute left could shed some light on—in addition to Dr. Tourassi's really great statements around workforce and some of the barriers to entry because of workforce and human capital, but any barriers to entry that we need to be considering for small to midsize enterprises with supercomputer technology and costs in the year 2021? I don't know if Dr. Willcox or Dr. Monroe have any insights into that.

Go ahead, Dr. Monroe.

Dr. WILLCOX. Go ahead, Dr. Monroe.

Dr. MONROE. OK, yes. Yes.

Ms. STEVENS. You unmuted first.

Dr. MONROE. OK. Yes, thank you for the question. Indeed, so IonQ is a sub-100-person company notably about to go public in the next month, and of course we're still building machines and looking for use cases. And, as a small company, I think it—you know, while this hearing is predominantly about high-performance conventional computing, as it turns out, I think was mentioned that you do need high-performance computers to not only optimize and run but use quantum computers. They have to be developed in parallel. And this is why actually we have—we do have a relationship with Oak Ridge, a working relationship for—they can run small algorithms on our machines, and we in exchange can maybe use some of the advanced computing machinery located in Tennessee.

So, indeed, it's important. These are capital-intensive purchases to—you know, to be able to have your own array of GPUs (graphics processing units) on your own site. We don't have to do that all the time, but it's very important in my own field that we're able to get access to very high-performance computing machines, and we do rely on DOE for some of the——

Ms. STEVENS. Right, you're relying on DOE and our labs and maybe universities. Well, I'm out of time, but this was great. Thank you, Mr. Chair, great hearing.

STAFF. Mr. Feenstra——

Ms. STEVENS. I yield back.

STAFF. Oh, forgive me. Forgive me, Ms. Stevens. Mr. Feenstra is next.

Mr. FEENSTRA. Well, thank you. Thank you, Chairman Bowman and Ranking Member Weber. I want to thank you to the witnesses for their testimony and in sharing their extensive research and experience. It is truly outstanding, all the leaps that we're getting into on the cutting-edge of this technology.

I want to direct this question to Dr. Tourassi, and thank you for being here. I mean, it's just a pleasure to see you. Accurate weather forecasting is incredibly important in my district not only for the farmers in Iowa as they predict the growing seasons and the harvest but also for covering severe weather events for the public. Dr. Tourassi, in your testimony you highlighted the use of Summit to achieve a milestone in global weather forecasting simulations. Can you elaborate on the milestone and how it will help in the future of weather forecasting and how the advent of exascale supercomputing will also be able to build on this?

Dr. TOURASSI. Absolutely. And this is something that Dr. Binkley addressed earlier on. Simulating modeling climate and weather forecasting is one of the most complicated problems because we deal with many different physical systems interacting with each other. So—and this is something very similar that happens with the biomedical and the biological space. If you try to understand only one dimension of the problem, you just miss the big picture. This is what high-performance computing and supercomputing has done throughout the years because it enables modeling and simulation and weather forecasting by increasing the complexity of the components we add in the equation. We can take into account atmospheric patterns. We can take into account soil information or ocean information, and that's how these models, with the support of Summit and in the future with Frontier, will provide models that are more detailed, and they are more precise. Now, of course, it's an oxymoron in some ways to say we will do weather forecasting in a more precise way because the atmospheric system itself is a chaotic system. Still, though, we're decreasing the uncertainty according with which we do these predictions.

The way the community has been moving—and again, I'm speaking a little bit at a high level because I'm a biomedical scientist—we know that when you're dealing with a problem that is very difficult to predict, effectively, you try to create a number of models and just throw them in a pile and see what they do. Think of it as an ensemble approach. And then we aggregate the results of these models.

As the supercomputers are getting bigger and bigger, what are we able to do? First of all, run more complex models. Second, run more of them, therefore increasing the precision of our predictions. And this is what we are experiencing with Summit, and this is what we will see with Frontier even more.

Mr. FEENSTRA. Wow, thank you so much for that answer. I've got one more question quickly. Dr. Binkley, as part of Argonne's Early Science Program, a team from Iowa State and Ames Lab is rewriting a software program called NWChem for the exascale era. The new program could provide an increase in the size of chemical systems that can develop new methods for converting biomass into biofuels. This could be groundbreaking. Can you talk about how exascale computers can be used for breakthroughs such as in the many different areas including biofuels?

Dr. BINKLEY. Well, certainly, the example that you're citing, NWChem, which is one of the foremost quantum chemistry codes, it's been around for a number of years, and it was developed originally at the Pacific North—yeah, Pacific Northwest National Laboratory, and, you know, making an effort to bring that software set into operation on the exascale computers is a major step, and it's something that our Basic Energy Sciences Program has been supporting.

There are a number of other codes that I think are going to have transformational effects. There—we—there are fusion simulation models that are under development for use on the exascale computers. There's also software for understanding seismic events. You know, there is—if you look at the list of the 24 exascale applications that are slated for first use, a number of them I think are poised to have breakthroughs.

Mr. FEENSTRA. Well, thank you so much. And my time is up. Thanks, everyone, for being on the panel. This is great information. I yield back.

STAFF. Mr. Casten is next.

Mr. CASTEN. Thank you so much. I really appreciate our witnesses being here, and I have some questions for you, Mr. Monroe. My district is just north of Argonne National Lab and just to the east of Fermi, which means I can't claim any of it but spent a lot of time at both. And of course that's where a lot of the really groundbreaking deployment of quantum computing is happening.

I want to first just be a little bit nerdy with you. When you were listing some of the types of questions that you can uniquely answer with quantum computing, I think you mentioned optimization problems, encryption, decryption, some issues around protein chemistry. All of these sort of classes—sort of NP hard problems where the solution space is huge and you can just jam through a lot of data, as you look at sort of the opportunities, is it uniquely to those classes of problems that just you can churn through a lot more data fields a lot more quickly, or are there other classes of problems that you're also excited about?

Dr. MONROE. Maybe I'll answer it in two ways. Thanks for the question. I'm always glad to put the nerd hat on. So the one disclaimer I'll say is that it's very hard to prove that a quantum computer can find the absolute optimal configuration of variables in some complex model. However, it could be a heuristic, meaning that it could do better than any conventional computer. You don't really need proof. All you need to do is show, well, it gave me a better answer. It gave me a shorter path——

Mr. CASTEN. That's correct.

Dr. MONROE [continuing]. Of somewhere out subject to certain constraints. You don't need to prove it. And that's sort of why quantum computers have to be deployed. You need to build them and deploy them as soon as possible and sort of see how well they perform.

So the other way I might answer this question—this might be too universal—is that I think every application in a quantum computer can be cast in terms of an optimizer. Optimization is a very general thing. Even code-breaking application, which is actually factoring numbers into their primes, that's a very hard problem. It's easy to multiply, it's hard to do the inverse, to factor. You can cast that in terms of an optimization problem. It's optimizing the two numbers that actually multiply to give the original number, and that has, you know, revolutionary impacts on security.

But one thing I want to conclude here with is you shouldn't think—I hesitate to describe quantum computers as big data machines. They don't process big data in a way that you can get access to all the big data. You can't compute massive—you can't model every molecule in the atmosphere and predict exactly what the weather is going to be. You need problems that sort of have a very small—a very simple answer at the end like there's a model of climate I have in the upper atmosphere. I don't know how to solve it. I don't know what conditions will validate this model. It's those problems that quantum computers can do. They take lots of—they sample lots of data, but you only get very—there's like a winnowing down to getting only a very small amount of information at the end, so it's a little bit subtle. And so high-performance computers are one-to-one machines. They can compute brute force function evaluation with all these inputs and all these outputs. Quantum computers do something a little differently.

Mr. CASTEN. Fascinating. I'd love to follow up on that with more time than we have, but let me leave with a question to you but invite all of our witnesses to respond with your thoughts. The—when I've gone down and toured these facilities at our labs, I'm always struck by the fact that part of the excitement is these NP hard huge data set problems for all classes of supercomputing. And then the other conversation we always end up getting into is the ethical questions around what does it mean to have a computer that's capable of asking questions that our brains aren't capable of thinking of and what's the appropriate boundaries there? And I'd welcome any of your thoughts on what we should be doing as a legislative body to—you know, maybe we're doing enough already to put those ethical boundaries in place and really understand how do we make sure that—you know, that we don't—we don't hit whatever that moment was in *Terminator* where the computers are smarter than we are and we can't figure out why.

Dr. MONROE. Well, I'll answer very briefly there and leave it to the other witnesses. This is a little bit passive, but I think the only—I think there's an impetus for us to get there no matter what because if we don't understand some revolutionary form of computing and we don't get there, others will. And so I think that that's some type of ethics. At least we're at the forefront of that new technology. I agree with you it can be vexing when you get a new technology, how to use it ethically and so forth, but, you know, it's very important to get there and not just ignore it and decide not to get there. And if we do that, others will.

Mr. CASTEN. Well, thank you. I'm out of time but would welcome any thoughts that the rest of the panelists have in writing afterwards. Thank you, I yield back.

STAFF. Mr. Obernolte is next.

Mr. OBERNOLTE. Thank you very much, Mr. Chairman, and thank you to our witnesses for a fascinating hearing. Also thank you for allowing me to participate. I know I'm not a regular Member of the Subcommittee, but this is a subject that I find extremely interesting.

I think it's important to recognize just how much this concept of beyond exascale computing has the potential to change computer science and its contribution to humanity in general. I mean, we're talking about computers that can actually perform more calculations per second than the human brain can, we can approach these artificial intelligence problems in ways that we just hadn't been able to in the past, so it's a very exciting thing to be part of, certainly something that we as a Federal Government need to be stimulating investment and research into.

So a question for Dr. Binkley. When you were talking in your testimony about beyond exascale computing, you mentioned that quantum computing is probably going to be the lead in that. And my question to you is do you think that quantum computing is the only technology that's capable of beyond exascale power, or are there other technologies, traditional technologies that might also be capable of that?

Dr. BINKLEY. Well, I don't think that quantum computing is going to lead to the elimination of the type of computing that can be achieved on current supercomputers, including exascale. I mean, as Chris Monroe has pointed out, quantum computers are very specialized in the type of calculations that they can do. Conventional computers, including exascale, use numerical methods and follow principles of sort of orderly input and orderly output. And so I think that what is going to come next after exascale, if you look at the typical exascale computer, each node in the computer has both a conventional CPU (central processing unit) chip and some type of a graphical processor. And it's the combination of those two that allows one to attain 10 to the 18 floating-point operations per second.

Experiments are beginning now to look at other types of computing elements that can be included into a supercomputer, and so other types of GPUs, one could imagine neuromorphic chips that could be incorporated. I think there's a lot of research that can be done in those areas, and I think that the lineage that has led us up to exascale still has more steps to go through as we go forward.

Mr. OBERNOLTE. Well, thank you, Dr. Binkley. As a computer scientist myself, I find this extremely interesting and really inspirational.

Before I go, if I could just make everyone aware of a piece of legislation that I've introduced. It's H.R. 3284 that would direct the Department of Energy to establish a program for capabilities beyond exascale. And one thing that I think we need to do more talking about is the need to research energy-efficient computing because if you look at the current generation of exaflop-capable supercomputers, one thing that strikes me is the pure amount of power that they consume. The El Capitan supercomputer that's being installed at Lawrence Livermore Laboratories I believe is going to have a power consumption on the order of 40 megawatts, so the amount of energy they consume and the heat that they consume is going to quickly become a barrier to our ability to employ these kinds of technologies. So as part of this bill I'm suggesting that we also establish a program to stimulate research into energy-efficient computing.

And then of course I know the testimony was mentioned in support for the Computational Science Graduate Fellowship Program. That's a program that I am very supportive of. I think that at the same time that we are stimulating research into these technologies, we also need to make sure that our academic population and our workforce is prepared to employ these technologies. It's going to be no good to make the tools if we're not also stimulating the kind of workforce that's going to be able to help us take that to the next level. And this bill would also do that. So I would certainly invite everyone's support and participation on that bill.

And I'll yield back, Mr. Chair, but thank you very much for letting me be part of this discussion.

STAFF. Mr. Lamb is next.

Mr. LAMB. Thank you, Mr. Chairman. And I apologize up front if I have any internet connection problems here. It hasn't been the best day for me on that.

The first thing that I wanted to ask was for Mr. Monroe. There's a statement on the Department of Energy's website that says—that right now we're at the same point in quantum computing that scientists in the 1950's were with computers. Do you consider that an accurate statement as far as kind of placing in context where we might be in the development of quantum computing?

Dr. MONROE. Yes, thanks for the question. It's an interesting comparison, and I buy into that with the caveat that we're sort of on this experiential curve on progress and so, you know, 10 years back in the 1950's is maybe like 2 years now, so, you know, all of us students of history, it's wonderful to see the progression from vacuum tubes to germanium transistors to silicon and then silicon we built silicon out. It took 20 or 30 years for that to happen. We don't expect that to be the same with quantum. I think we are on the cusp of several different technologies, superconducting circuits. At IonQ we work with individual atoms. We know how to scale these things. We have to put the, you know, engineering effort into it.

So the comparison is a very good one because we are at very low levels now. We're not—we don't have high—we don't have Windows for quantum yet. You know, we're working with individual gates, very low-level stuff that they were doing in the 1950's with silicon transistors. But we also know that software exists for classic computers, conventional computers that we can also deploy and accelerate in quantum. So, yes, it's 1950's in quantum, but, you know, we're sort of moving much faster.

Mr. LAMB. Good, thank you. And you put in your testimony that you thought the *Endless Frontiers Act* would be helpful to our efforts to go even further in quantum computing. Could you maybe just say a little bit more concretely why that is? Because I'm a supporter of that bill. I think it's incredibly important. I do think there is some hesitation for those familiar with the NSF that it could somehow maybe detract from the core basic science mission of the NSF, so would you be able to address that and assuage those concerns?

Dr. MONROE. Yes, I—my understanding is there are many different approaches to re-tasking the NSF to have a technology edge to them. My understanding is also it's not a zero-sum game, that the NSF will still have in its core mission the ability to do research, blue-skies research, research for research sake.

You know, I want to link that to technology, especially quantum technology and why NSF is very well-suited to play a big role in the development of the tech, not just the science, and that is the future of quantum computing, we need applications, and right now they're coming in the name of science, mainly in universities. At Duke, in the University of Maryland, we're deploying our laboratory systems to do models of black holes and wormholes believe it or not, things that maybe a company would never do. And it's those scientific applications that are happening right now. They're also happening at Department of Energy laboratories across the country.

Companies won't pay for this. They're not going to build a device to do black holes. So having an agency like the NSF that has an eye on the blue-skies fundamental science research, building and using quantum computers, making a user facility, for instance, is a really good idea because it's going to tide us over until industry really controls the building of these devices. And then we can use those devices for new science just like we're using exascale and high-performance computers at Oak Ridge and other DOE labs for current science.

Mr. LAMB. I agree. I agree. Thank you.

I do want to just sneak in one more question. Dr. Tourassi, thank you for everything you're doing, very excited for the launch of Frontier this year. Is there in somewhat simple layman's terms a way to describe a problem that we cannot solve until we get Frontier? I understand that it'll probably solve other problems faster. Are there new actual types of problems that this platform will allow us to solve that we couldn't before?

Dr. TOURASSI. That's a good question. Actually, what we are hoping that the next machines will enable is that pure integration of modeling and simulation with large-scale AI, bringing together models and observational data that we can get from our different experimental facilities and from the different Federal agencies. That level of computing has not been—has not happened yet.

It goes back to the example that was given earlier about the phenome-genome association studies, which actually are a very challenging and attractive problem for exascale computing when we're looking at large populations of, let's say, humans when we want to do the—at the population level phenome-genome associations, that is which will drive new treatments, specialized treatments, precision medicine.

Mr. LAMB. Great, thank you. I'm out of time. Mr. Chairman, I yield back.

STAFF. Ms. Ross is next.

Ms. ROSS. Thank you, Mr. Chairman, and thank you for holding this hearing and to all the panelists for joining us today.

The Committee Members have heard this before, but to the panelists, I represent the Research Triangle area of North Carolina. I represent Wake County. And the Research Triangle area has a growing ecosystem of hundreds of innovative and collaborative companies, including science and technology firms, government agencies, academic institutions, including Duke, startups and nonprofits. The IBM Quantum Hub at NC State University is a cross-disciplinary center of quantum computing education, research development, and implementation. The IBM Quantum Hub at NC State works with researchers from Duke and UNC (University of North Carolina) Chapel Hill to help partners develop quantum teams and explore promising use cases and promotion of quantum computing in real-world applications.

And so, Dr. Monroe, congratulations on the quantum computing user facility at Duke. And I want to see how these partnerships work. But there's a real-world application that we've been dealing with these past couple weeks on cybersecurity and how we might be able to collaborate between research institutions and the private sector and use quantum computing to maybe prevent some of the kinds of things that we're seeing, particularly when we're dealing with 20th century technology that hasn't moved into the 21st century. And I don't know if the Committee has dealt with this earlier. I was in another Committee meeting. But I'd love to hear from you about how quantum computing might help us with cybersecurity.

Dr. MONROE. OK. Thank you for the question, Congresswoman. So indeed the known killer application of quantum computers is code-breaking, cracking the most popular data encryption schemes we have. It looks like that application is one of the hardest out there. We need much bigger machines to do that. That said, we also are—this has been said before. We're at such an early stage in the game. We know we can shrink the problem, we can make it more efficient through software by being more clever on how we structure not only the instructions that we run the quantum computer but the quantum computer itself, how we control it. And this is why at Duke, for instance, we're collaborating quite—we're starting a very large collaboration with the colleagues you mentioned at NC State, and they're experts in computer architecture. We make machines, and, you know, this is sort of the marriage that we really like the direction of that.

Now, making code-breaking problems easier is something that won't happen in isolation. It won't happen by me or, you know, it won't happen by somebody that doesn't have a machine, so I think they have to work in unison. And, you know, I've—at IonQ we've worked closely with IBM. IBM also writes software that supports our system, the expression of our particular hardware. And, you know, IBM I think has really, you know, taken a lead at educating the public on using systems. They were the first to put their system on the cloud, and we've sort of followed a few years later. So it's a wonderful ecosystem down here in North Carolina, but, you know, also at the company IonQ, which is a Duke-Maryland startup inside the beltway in Washington, they—they're very close to the National Security Agency and, you know, we—we're—we understand that community's needs in this field.

So it's—I don't know what else to say. It's a very exciting time for the field. We'll take problems anywhere they come from, but security is a big one, and I think one—one last comment, what's not mentioned so much is the energy grid. That's a huge logistics problem, and I was horrified to find out that apparently I think the way energy is coordinated across the country, it's not—maybe it's not done very smartly. It's not my field, I shouldn't comment on that, but we do know that's one of the most vexing optimization problems we have, and the news in the last few weeks, especially here in North Carolina, is that, boy, it would be nice if we made that a little smarter in the future. I'm not sure exactly how quantum will play a role there, but if it's an optimization problem, we want to take a crack at it.

Ms. ROSS. Well, you anticipated my second question, and that is how open are, you know, some of our big energy companies to your help?

Dr. MONROE. Well, you know, at IonQ, we are—you know, we're not a company that can afford to talk to 50 different companies in parallel, but we do have conversations with a few energy companies that have teams on the ground that are ready, they want to deploy quantum for their uses. These could be oil and gas or, you know, related to big pharma or, you know, developing new fuels. It is a researchy time in the field right now, but they're very open to working with companies like ours, like IBM, and so forth and going in the future. So in terms of the energy grid, that's definitely not my area of expertise, yes, so, thank you.

Ms. ROSS. Thank you, Mr. Chair. I've exceeded my time, and I yield back.

STAFF. Dr. Foster is next.

Mr. FOSTER. Thank you, and thanks to all of our witnesses here.

I—well, first off, for those of you that don't know me, I'm Congress's Ph.D. physicist and also the—Congress's AI programmer, although I only have really dabbled in TensorFlow recently.

The question I have was a more general one about where you see this whole field is going of advanced computing. It seems to me that it's likely to fragment, that traditionally we've had a whole big line of, you know, megaflops, big floating point pipelines to—for local partial differential equations, and that will continue for all the usual reasons from, you know, weather maps to nuclear weapons and everything in between. And so that will continue perhaps with an overlay of AI to do the grid scale optimization on the fly but largely just, you know, a similar thing.

And then if you look at what's happening in AI, the state-of-the-art AI engines are not in the Federal realm any more. They are—you know, if you're just high-speed execution of GPT–3 and other AI interesting algorithms, you know, the best engines are from places like Google for that. And I don't believe that the Federal Government is likely to be able to keep up with that, that industry will continue to pull ahead. In that case, the Federal investment will largely be in paying for university groups to have access to these very advanced AI engines, and that's probably as ambitious as we should be on that front.

And then, you know, farther down the chain of, I guess, narrower calculation of widths, you have things like neuromorphic engines, which may end up as not floating-point or even low-precision integer but single-bit computing engines for neuromorphic thing—and so those will be experimental for a while and may or may not become, you know, really important.

And so from a Federal point of view we have to decide where to put—you know, there's a bipartisan agreement that we should something like double the Federal budget, but we have to understand, you know, whether to write, you know, a big new—a check for a big new initiative wholly in National Science Foundation, whether we should sort of scale up all of the existing enterprises and then see what they're able to produce in a few years. And a big part of that decision is whether we think we have to—whether we're going to—what the mixture of small research and big facilities is going to be because that—anyway, so I'd be—I guess, Steve, we can start with you and what you see as DOE's role in that.

Dr. BINKLEY. Well, clearly, Bill, the—our programs are really intended to advance the state-of-the-art in scientific computing writ large, and so that essentially encompasses all of the technology paths that you just outlined. I mentioned earlier that the—sort of the next obvious move post-exascale is to look at heterogeneous computing systems where, in addition to GPUs, one could have neuromorphic chips and other computational technologies, and those types of machines would provide avenues to solve essentially new problems.

But also I think continued investment in quantum computing is going to be necessary. I don't—I do see a separation between sort of the classical Turing-type computation and quantum computation, and so there may—you know, that may eventually divide—you know, there may be a parting of the roads there. But at least at this stage in time I think we need to really look at the diverse portfolio of technologies and not foreclose any options.

Mr. FOSTER. Yes, well, are there areas where you sort of believe we'll not be able to compete with industry? You know, I'm thinking, as I mentioned, specifically of AI. Are there areas there where it's clear already that the government is not going to be leading?

Dr. BINKLEY. Well, when it—so there are certainly AI applications that are going to be dominated by the high-tech companies, Googles and so on. The one area, though, that I don't see Google moving into is there is—with the advent of the exascale machines and the type of GPUs they have, they are actually the most powerful AI engines that exist. And the—you know, this very, very high end of AI is something that I think is going to still be the domain of the leadership computing class facilities. You may disagree, but, I mean, I think——

Mr. FOSTER. No, I—you know, I'd be interested because there's a different narrative coming out of Silicon Valley——

Dr. BINKLEY. Yes.

Mr. FOSTER [continuing]. And a lot of worry that small startups in AI simply can't compete because they do not have access to Googles and so on, Amazon's big AI engines so——

Dr. BINKLEY. Yes.

Mr. FOSTER. Now, does anyone else have a—Georgia or——

Dr. TOURASSI. Yes, I would like to add a couple of data points to make the case of why the supercomputers and exascale computers will actually be the powerful AI engines, and I would like to draw these experiences from the exascale computing project and particularly the one that was mentioned earlier to support cancer. When we started that effort, the specific application we were focusing was not a traditional application that anybody would have thought about supercomputers, but it was a challenge application for us to start building frameworks to prepare for the scientific community.

Well, in the process, we actually faced the philosophical debate, do you fit hardware to algorithms or do you build algorithms, modify algorithms to fit hardware? Typically, we think of I have a specific application, I will prepare a certain type of hardware to fit that particular application. Well, in the case of a supercomputer in transit, we had to do the opposite. And what we discovered in the end is that for this particular application, by modifying the algorithm to make the most of the supercomputer, we exceeded performance, state-of-the-art performance. So that opened a completely new way of thinking about using supercomputers for applications that traditionally were not thought suitable to supercomputing.

The second example I would like to give was the very recent Gordon Bell award for the COVID–19 pandemic that brought together lots of modeling and simulation of the complete envelope, viral envelope of 305, you know, million atoms, but there was an AI workflow to accelerate that modeling and simulation, so that coupling of AI and modeling and simulation that was only possible on Summit. So I do believe that we see plenty of examples out there to support that these general-purpose machines will be able to deliver a lot even in the AI space.

Mr. FOSTER. All right, fascinating discussion. I've exceeded my time, so I guess I have to yield back.

Chairman BOWMAN. Thank you, everyone. Before we bring the hearing to a close, I want to thank our witnesses for testifying before the Committee today. The record will remain open for 2 weeks for additional statements from the Members and for any additional questions the Committee may ask of the witnesses.

The witnesses are excused, and the hearing is now adjourned. Thank you all. Have a good day.

[Whereupon, at 1:06 p.m., the Subcommittee was adjourned.]

Appendix I

Answers to Post-Hearing Questions

ANSWERS TO POST-HEARING QUESTIONS

Responses by Dr. J. Stephen Binkley

QUESTIONS FOR THE RECORD
U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

Hearing: "Accelerating Discovery: the Future of Scientific Computing at the Department of Energy"
Wednesday, May 19, 2021

QUESTIONS FROM CHAIRWOMAN EDDIE BERNICE JOHNSON

Q1. The Scientific Discovery through Advanced Computing program (SciDAC), within the Advanced Scientific Computing Research program (ASCR), supports research to develop computational solutions to challenging problems in other disciplines. SciDAC partnerships are currently limited to collaborations between ASCR and the other Office of Science programs, with the one exception being the Office of Nuclear Energy.

Q1a. Given the importance of modeling and simulation tools for addressing challenges in areas such as energy technology development and infrastructure resilience, what plans, if any, does ASCR have for expanding SciDAC to include partnerships with other applied energy offices?

A1a. For the past two years, the Advanced Scientific Computing Research (ASCR) program has had several productive conversations with applied energy offices such as the Office of Electricity (OE) and the Office of Energy Efficiency and Renewable Energy (EERE) regarding areas of possible collaborations through the Scientific Discovery through Advanced Computing (SciDAC) program. Recognizing that the Exascale Computing Project (ECP) is scheduled to end by 2024, ASCR will work closely with the other applied energy offices to ensure that the partnerships established via ECP will continue.

Q1b. Are there other mechanisms for building connections between ASCR's research programs and needs in other parts of the Department that should be explored?

A1b. In the past year, ASCR established across the Department of Energy (DOE), an internal DOE working group to redefine the ASCR Leadership Computing Challenge allocation program that provides ASCR computing resources to support critical DOE mission areas as well as support national emergencies. We now have broad participation in the working group from across the Department, which has been fruitful in strengthening connections between ASCR and other parts of the Department. ASCR will continue to explore other mechanisms to connect our research programs to other parts of the Department.

Q2. As scientific discovery relies increasingly on computing and computational tools, it becomes increasingly important to ensure that the Department has the skilled computational workforce it requires to ensure long-term U.S. leadership in this area. One program that exists to meet this need is the Computational Science Graduate Fellowship (CSGF), which supports

QUESTIONS FOR THE RECORD
U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

Hearing: "Accelerating Discovery: the Future of Scientific Computing at the Department of Energy"
Wednesday, May 19, 2021

graduate students at the intersection of disciplinary science and applied math and computer science. Fellows also participate in a 12-week research traineeship at a national laboratory, helping to ensure the availability of computationally trained scientists with close ties to DOE.

Q2a. CSGF has been flat funded at $10 million for several years now. Is this level of funding adequate given both the level of demand on the program as well as ASCR's expanding research mandates in areas such as artificial intelligence (AI) and quantum information science (QIS)?

A2a. Because of the continued high quality of the applicants to the Computational Science Graduate Fellowship (CSGF) and the very high demand for workers with the skills fostered by CSGF across the DOE complex and the Nation, the FY 2022 President's Budget Request includes a $5,000,000 increase in funding for the CSGF program.

Q2b. How can the Department use this fellowship to diversify the computational workforce and foster deeper connections with minority-serving institutions?

A2b. We have been working to increase the diversity of the fellows and the institutions engaged in CSGF. The 2021 CSGF applicant pool and selected fellows are the most diverse in the 31-year history of the program. Through the proposed Reaching a New Energy Sciences Workforce (RENEW) program, the Department will build connections with minority-serving institutions to broaden the pool of diverse CSGF applicants and institutions where the fellows receive their degrees.

Q3. The Department is preparing to deploy the nation's first exascale computer at Oak Ridge National Laboratory this year, bringing to fruition an initiative that began well over a decade ago. While this is an exciting milestone worth celebrating, we must not lose site of the fact that these forthcoming systems—and the software that enables them—must be continuously maintained and improved in order for their full potential to be realized.

Q3a. What sort of initiative should the Office of Science be taking to ensure that the ecosystem the Department stewarded to establish these systems is sustained even as ASCR shifts its attention toward long-term research needs?

QUESTIONS FOR THE RECORD
U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

Hearing: "Accelerating Discovery: the Future of Scientific Computing at the Department of Energy"
Wednesday, May 19, 2021

A3a. The Office of Science (SC) will continue to evaluate recommendations from the Advanced Scientific Computing Research's (ASCR) Federal Advisory Committee Act (FACA) advisory committees to identify initiatives and programs to ensure that the exascale ecosystem is sustained and nurtured in the future. ASCR has also established a task force comprised of research and facilities program managers to identify possible sustainability mechanisms for the exascale software ecosystem. ASCR will work closely with the other SC and Applied Energy Program Offices, through the Scientific Discovery through Advanced Computing (SciDAC) program, to ensure that a broad set of applications are capable of fully utilizing the exascale ecosystem.

Q4. As the Exascale Computing Initiative comes to fruition, ASCR is reorienting toward longer term research on emerging technologies and computing paradigms that will enable the advancement of high-performance computing capabilities after the end of Moore's law.

Q4a. Please outline the importance of investing in fields such as applied math and computational science for enabling the computing capabilities of the future.

A4a. Evolving applied mathematics and computer science algorithms and methods, as well as machine learning and data analytics, continue to improve the effectiveness of scientific application codes on emerging technologies. As high performance computers become more heterogeneous, continued investments in applied math and computational science methods and tools that hide the complexity of the systems from the user are crucial. Further, applied math provides the foundation for scientific machine learning and artificial intelligence (AI/ML). As noted in the 2019 Basic Research Needs workshop report on Scientific Machine Learning, research is needed to ensure that AI/ML techniques are repeatable, include domain awareness, and explainable. AI/ML techniques also rely on high quality data that is findable, accessible, interoperable and reusable.

Q4b. How can Congress better enable that shift back toward forefront research?

QUESTIONS FOR THE RECORD
U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

Hearing: "Accelerating Discovery: the Future of Scientific Computing at the Department of Energy"
Wednesday, May 19, 2021

A4b. As represented in current legislative actions, Congress continues to recognize the importance of fundamental research for economic and national security.

Q5. The Office of Science has launched a new initiative to advance the state-of-the-art in microelectronics through an approach that integrates materials science, chemistry, devices, and algorithms among other areas. The Basic Energy Sciences program seems to be taking the lead on this topic given its focus on materials and chemical sciences.

Q5a. How do you see this Office of Science-wide effort evolving in the years ahead, particularly as the federal government acts to address the current semiconductor shortage and its impacts on the broader economy?

A5a. In FY 2020 the Office of Basic Energy Sciences (BES) initiated two new Energy Frontier Research Centers (EFRCs) focused on microelectronics research. In FY 2021, five Office of Science (SC) programs – Advanced Scientific Computing Research (ASCR), BES, Fusion Energy Sciences (FES), High Energy Physics (HEP), and Nuclear Physics (NP) – have come together to support multi-disciplinary microelectronics research. The aim is to accelerate the advancement of microelectronic technologies in a co-design innovation ecosystem, in which materials, chemistries, devices, systems, architectures, algorithms, and software are developed in a closely integrated fashion. A Department of Energy (DOE) national laboratory announcement solicited proposals in this area and award announcements are expected in August 2021. Subject to FY 2022 Congressional Appropriations, DOE plans to expand investments in this research, across the five SC programs noted above, that will support both microelectronics and advanced manufacturing initiatives. The fundamental knowledge and discoveries that result from this research will help lay a foundation for continued U.S. innovation and leadership in the semiconductor sector. DOE is coordinating these efforts with other federal agencies through the newly established National Science and Technology Council (NSTC) Subcommittee for Microelectronics Leadership.

QUESTIONS FOR THE RECORD
U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

Hearing: "Accelerating Discovery: the Future of Scientific Computing at the Department of Energy"
Wednesday, May 19, 2021

Q5b. How can Congress help enable the nascent microelectronics initiative within the Office of Science?

A5b. Based on recent legislative actions, Congress recognizes the importance of both microelectronics and fundamental research for economic and national security.

QUESTIONS FOR THE RECORD
U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

Hearing: "Accelerating Discovery: the Future of Scientific Computing at the Department of Energy"
Wednesday, May 19, 2021

QUESTION FROM REPRESENTAIVE SUZANNE BONAMICI

Q1. Dr. Binkley and Dr. Tourassi, will advances in neuromorphic computing help to end or limit bias in high performance computing and artificial intelligence? If it does not end bias, will it make detecting bias more difficult?

A1. Neuromorphic computing, which is one of the National Academy of Engineering grand challenges for the 21st century, seeks to reverse engineer the computing architecture of the brain. The primary reason for the Department of Energy to explore neuromorphic computing is for its potential to foster great advances in energy efficiency computing since the brain is approximately a million times more energy efficient than our supercomputers. As noted, artificial intelligence (AI) bias is a source of errors in machine learning algorithms; therefore, it represents a major obstacle and limiting factor in advancing overall use of AI. Once we are able to understand and harness the brain's computing architecture and develop effective neuromorphic systems, it may be possible to measure and characterize bias similarly to how we evaluate bias in people. However, it is important to note that there is still much more basic research and development needed to understand, simulate and emulate the brain and to be able to engineer a neuromorphic computing system on such principles.

Responses by Dr. Georgia Tourassi

QFR Response by Dr. Georgia Tourassi

Question submitted by Representative Suzanne Bonamici

Will advances in neuromorphic computing help to end or limit bias in high performance computing and artificial intelligence? If it does not end bias, will it make detecting bias more difficult?

Response: Bias in AI is due to bias inherent in the data used to train the AI model. The bias is introduced either during the data collection process or during the data preprocessing phase (when the data is "manipulated" in preparation to train the AI model). There is nothing inherent in neuromorphic (or any other) computing that can detect or mitigate such bias. For this reason, advances in mathematical and statistical sciences will be important. AI scientists should promote a rigorous statistical framework to monitor for potential biases in the collected data and for measuring AI reproducibility and repeatability based on statistically and application-specific appropriate standards.

Responses by Dr. Karen Willcox

Questions for the Record

Subcommittee on Energy of the
House Committee on Science, Space, and Technology (117th Congress)

hearing on

Accelerating Discovery: the Future of Scientific Computing at the Department of Energy

Submitted July 6, 2021

By

Dr. Karen E. Willcox
Director, Oden Institute for Computational Engineering and Sciences
Associate Vice President for Research
Professor of Aerospace Engineering and Engineering Mechanics
The University of Texas at Austin
Email: kwillcox@oden.utexas.edu Website: kiwi.oden.utexas.edu

1. In your testimony, you noted that you lead a large center funded by the Advanced Scientific Computing Research (ASCR) Applied Mathematics subprogram known as a Mathematical Multifaceted Integrated Capabilities Center (MMIC). These centers seek to address complex challenges in energy research that require integrated processes across multiple computational and statistical disciplines.

a. Can you give us your perspective on why this type of effort is so impactful for the Department to meet its mission needs?

b. Why is it important for ASCR to balance its investments in world-leading computing systems with foundational research activities like the MMICCs?

c. Many of these activities provide direct funding to academic institutions. How does university-based research provide unique contributions to ASCR's research agenda?

The Mathematical Multifaceted Integrated Capabilities Center (MMICC) program of DOE Advanced Scientific Computing Research (ASCR) funds centers to conduct focused basic research in applied mathematics that is strongly driven by mission needs. The centers bring together diverse teams that include mathematicians, computer scientists, computational scientists, engineers, and domain experts, spanning universities and National Laboratories.

Why is this type of effort so impactful for the Department to meet its mission needs? Scientific computing will continue to accelerate discovery and innovation only if we address foundational challenges that limit the complexity, scale, and trustworthiness of computational analysis, prediction, and decision tools. The MMICC program is so impactful because of the unique way that it fosters **mission-driven interdisciplinary basic research** to address these foundational challenges. These challenges cannot be addressed in isolation by mathematicians or computer scientists alone, but rather need deep integration across mathematics, computing, and application domain expertise. Mechanisms for funding interdisciplinary basic research teams exist in other agencies (e.g., the DOD Multidisciplinary University Research Initiative (MURI) and NSF crosscutting programs) but **the MMICC program is unique in the way it brings together DOE mission needs with a mandate to conduct foundational mathematical research**. My experience is that the driving mission pull plays a critical role in ensuring that interdisciplinary research flourishes. In the absence of such a pull as a lens by which to evaluate research ideas, **foundational scientific computing research runs the risk of falling through the disciplinary cracks**, as has historically been the case at a number of institutions and agencies. One should also not overlook the transformational positive impact that is achieved through the immersive interdisciplinary research experiences the MMICC program provides for graduate students and early-career researchers.

Why is it important for ASCR to balance its investments in world-leading computing systems with foundational research activities like the MMICCs? Investments in supercomputing systems are necessary but not sufficient. These advanced machines present fundamental challenges to the mainstream algorithms of computational science that have been developed in an era of much simpler architectures. Foundational research in numerical methods and algorithms is required to unlock the great potential of these new systems, which include heterogeneous nodes with advanced accelerators. The leadership computing facilities of today are having impact because they are employing the scalable algorithms built on the mathematical foundations laid over the past five decades. Simply put, **we will be unable to take advantage of future computing systems unless we also invest in future mathematical methods and computational algorithms**.

How does university-based research provide unique contributions to ASCR's research agenda? University-based research has played a critical role in building the mathematical and algorithmic foundations of scientific computing. Whether it is the foundational theory of the finite element method, practical methods for computing singular value decompositions, or the generalized minimal residual method for solving linear equations, **university-based research underpins the most powerful modeling, simulation and data-analysis tools used across science and engineering today**. The Nation's universities attract the top talent from around the world, and are best positioned to address the pressing basic research challenges discussed above. Further, **DOE direct funding to academic institutions is the key driver in creating the future workforce**. It is absolutely essential that graduate students and postdoctoral fellows continue to have immersive research experiences that hone their specific technical skills, cultivate their broader problem-solving skills, and expose them to the wider scientific and technological community. These experiences can only be provided if basic research funding to universities remains strong.

2. Much attention has been focused on the role of ASCR's research programs in advancing artificial intelligence, quantum, and other future computing technologies, but it is important to recognize that the applied mathematics and computational science activities funded by ASCR are also vital to meeting needs across the Department's mission space.

a. In your testimony, you mention the importance of investing in this foundational research for addressing energy, environmental, and nuclear challenges. Can you clarify that connection?

b. Please provide your perspective on whether the Department should expand existing mechanisms, such as the Scientific Discovery through Advanced Computing (SciDAC) program, and/or create new ones for fostering deeper connections between ASCR's foundational research programs and other offices within the Department.

Why is it imperative to look beyond artificial intelligence, quantum, and other future computing technologies and also invest in applied mathematics and computational science foundational research for addressing energy, environmental, and nuclear challenges? It is important to recognize that **quantum computing, if it comes to practical fruition, will be applicable to a narrow class of problems**. The vast mainstream of problems in computational science – including astrophysics, climate modeling, structural mechanics, subsurface flow, and turbulence – do not lend themselves to quantum formulations. To quote Scott Aaronson, world-renowned quantum computing expert

> "The trouble is that quantum computers will not revolutionize everything. Yes, they might someday solve a few specific problems in minutes that (we think) would take longer than the age of the universe on classical computers. But there are many other important problems for which most experts think quantum computers will help only modestly, if at all."[1]

Regardless of the form that future computing technologies take, advances in hardware by themselves are not sufficient to address the great energy, environmental, and nuclear challenges facing us. Overcoming these challenges requires fundamental advances in modeling, prediction, control, and optimization of increasingly complex systems. In particular, the increasing (1) **complexity of models** (driven by the need to incorporate more physics into these models), (2) **complexity of computer architectures** (driven by advances in hardware), and (3) demands to move beyond simulation to optimization and uncertainty quantification (driven by the **increasing role of modeling in decision-making for critical systems**) require fundamental new advances in applied mathematics and computational science.

Should the Department should expand existing mechanisms, such as the Scientific Discovery through Advanced Computing (SciDAC) program, and/or create new ones for fostering deeper connections between ASCR's foundational research programs and other offices within the Department? Now is an **ideal time to forge deeper connections between ASCR's foundational research programs and other DOE offices**. Such connections will permit DOE to capitalize fully on past investments, as well as position it to embrace the increasing footprint of computational science across all areas of science and technology. Outside of the Office of Science and NNSA (which already have strong connections with ASCR), several high-value connections could be made, with the Office of Energy Efficiency and Renewable Energy (EERE) being one obvious target. However, even within the Office of Science and

[1] Aaronson, S. "What Makes Quantum Computing So Hard to Explain?" *Quanta Magazine*, June 2021. https://www.quantamagazine.org/why-is-quantum-computing-so-hard-to-explain-20210608/

NNSA, there are untapped opportunities to strengthen connections between ASCR's foundational research programs and other program offices. For example, there are clear opportunities for integrated multifaceted mathematics research to have high impact across Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Sciences, High Energy Physics, and Nuclear Physics; yet, with only three funded MMICC centers, the MMICC program can touch only a small fraction of these scientific application areas. Expansions of the SciDAC program, which has primarily focused on scientific computing software, and the MMICC program are mechanisms to address these opportunities, as is a growth in the base Applied Mathematics research program, but new mechanisms are also needed. For more than three decades, the DOD has used the **highly successful Multidisciplinary University Research Initiative (MURI) program to support DOD mission-relevant basic research that intersects more than one traditional science and engineering discipline**. MURI research is crafted to be highly aligned with DOD priorities, through the specific topics that are issued each year. A similarly structured program, led by DOE, could play a pivotal role in catalyzing interdisciplinary basic research that addresses gaps and key mission needs, especially when it comes to forging new connections between the mathematics, computational science, and scientific application domain communities, and to readying scientists and engineers to take advantage of future computing systems. A second new mechanism could be through an expansion of the **DOE Early Career Research Program**. The current Early Career Research Program solicitation states that "Research topics are required to fall within one of the Department's Office of Science's eight major program offices." Restructuring this program to explicitly encourage Early Career research that cuts across program offices could again catalyze interdisciplinary basic research that addresses key gaps and mission needs.

3. As scientific discovery relies increasingly on computing and computational tools, it becomes increasingly important to ensure that the Department has the skilled computational workforce it requires to ensure long-term U.S. leadership in this area. One program that exists to meet this need is the Computational Science Graduate Fellowship (CSGF), which supports graduate students at the intersection of disciplinary science and applied math and computer science. Fellows also participate in a 12-week research traineeship at a national laboratory, helping to ensure the availability of computationally trained scientists with close ties to DOE. The Committee recognizes that CSGF represents a laudable effort to strengthen the pipeline, but notes that the program has been flat funded at $10 million for several years now.

a. From your perspective as the director of both a computational science-focused institute and an ASCR-funded center, do you think the Department's current investment in CSGF is adequate given both the level of demand on the program as well as ASCR's expanding research mandates in areas such as artificial intelligence (AI) and quantum information science (QIS)?

b. Do you have any suggestions as to ways in which this program might be expanded or modified to focus on workforce challenges that are currently unaddressed?

Computational Science Graduate Fellowship (CSGF) program Fellows exemplify the type of students required to address the workforce needs I highlighted in my testimony: CSGF Fellows are inspired by interdisciplinary topics at the interfaces of mathematics, computing and engineering/scientific applications, and they are deeply committed to addressing the Nation's scientific and technological challenges. CSGF Fellows are more likely than other students to pursue careers in scientific and engineering domains (e.g., the National Laboratories or in academia)[2]. Thus, the CSGF program plays a critical role in achieving the workforce goal I stated in my testimony of "ensuring a strong, diverse pipeline of highly trained professionals who remain committed to scientific and engineering domains, rather than being lured away by more lucrative positions in commercial and business sectors."

Furthermore, the CSGF program is unique among other prestigious national fellowships in the way it proactively shapes its trainees. This is achieved through requirements on graduate courses (and associated stringent oversight), through the required 12-week research traineeship at a national laboratory, and through concerted attention to mentoring and networking. A consequence of the graduate course requirement is that CSGF has played an important role in incentivizing universities to create new interdisciplinary computational science graduate curricula that address the Nation's workforce needs. As such, **an expansion of the CSGF program would not only increase the supply of talented students trained at the interfaces of mathematics, computing, and science/engineering – greatly valued at the national laboratories and in industry – but also provide DOE with a key lever to encourage new university programs that respond to pressing scientific computing workforce needs**.

Is the Department's current investment in CSGF adequate? There is no doubt that demand for those trained through the CSGF program already outstrips supply, and this demand will only increase in the coming years. Further, the pool of highly qualified applications far outstrips the availability of CGSF awards. **DOE could double the size of the CSGF program and both the demand for and quality of**

---

[2] Wells et al., "U.S. Department of Energy Computational Science Graduate Fellowship: 1991–2016. A follow-up study of recipients and programmatic outcomes." February 2017. https://www.krellinst.org/doecsgf/docs/2017_DOE_CSGF_Longitudinal_%20Study-Web.pdf

the Fellows would remain extremely high. Given the criticality and urgency of the Nation's scientific computing workforce needs, especially with regard to having highly trained US citizens, the benefits of increasing the investment in CSGF appear to far outweigh the costs.

How might the CSGF program might be expanded or modified to focus on workforce challenges that are currently unaddressed? As noted above, the CSGF program has played a key role in incentivizing universities to create new interdisciplinary computational science graduate curricula. **An expansion of the CSGF program, such as through a series of tracks each with a carefully crafted set of requirements, could play a similar role in catalyzing universities to expand class offerings and degree programs in areas such as exascale computing, AI for science and engineering, and quantum information systems**. This in turn will have ensuing positive impacts on workforce development well beyond just the CSGF Fellows. A further expansion would be to **introduce a Computational Science Postdoctoral Fellowship program** – an equivalent to the CSGF program targeted towards postdoctoral fellows. The National Laboratories all currently offer prestigious postdoctoral fellowships, but these are for postdoctoral fellows employed at the laboratories. There would be tremendous value in a postdoctoral program paralleling CSGF where the postdoctoral fellow's primary appointment is at a university, with a required research traineeship at a national laboratory. Such a program would enable early career researchers to broaden their skillsets in areas that are highly aligned with DOE workforce needs, while still providing the flexibility to build their academic research portfolios before committing to a particular career path. Just as with CSGF, such a program would also create a long-lasting network of young professionals and would contribute towards continued alignment of university basic research with DOE mission needs.

Appendix II

Additional Material for the Record

LETTER SUBMITTED BY REPRESENTATIVE FRANK LUCAS

QUANTUM INDUSTRY COALITION

May 18, 2021

The Honorable Eddie Bernice Johnson
Chairwoman
House Science, Space and Technology Committee

The Honorable Frank Lucas
Ranking Member
House Science, Space and Technology Committee

The Honorable Jamaal Bowman
Chairman
Energy Subcommittee
House Science, Space and Technology Committee

The Honorable Randy Weber
Ranking Member
Energy Subcommittee
House Science, Space and Technology Committee

Dear Chairwoman Johnson, Ranking Member Lucas, Chairman Bowman and Ranking Member Weber:

I am writing to you on behalf of the Quantum Industry Coalition regarding this week's hearing on "Accelerating Discovery: the Future of Scientific Computing at the Department of Energy." The Quantum Industry Coalition is a group of companies dedicated to US leadership in the quantum fields. Its members include Accenture, AgilePQ, Amazon Web Services, Anametric, ColdQuanta, D-Wave Government Inc., EeroQ, Entanglement Institute, Founders Fund, Google, Honeywell, IBM, MagiQ Technologies, Nanotronics, Quantum Xchange, Qubitekk, Rigetti Computing, Safe Quantum, Inc., Seeqc, Strangeworks, Xofia, and Zapata Computing.

The Department of Energy (DOE) has been a leader in high-performance computing for decades. In the last few years, thanks in large part to research conducted through DOE, other federal agencies, academia, and the private sector, quantum computing has begun to emerge as a valuable aspect of high-performance computing. While quantum computing will not supplant classical computing, it will provide critically important new capabilities that will make it possible to solve computing problems that classical computers could not solve for the foreseeable future. The combination of quantum computing with classical computing, also called hybrid, will help drive the next generation of advances and breakthroughs across a broad spectrum of scientific and technological fields. Today, nearly all quantum computing hardware is now made available securely via the cloud, which provides the added benefit of breaking down any barriers to access the power of quantum computing. The technology is now at the point when accelerated innovation can occur.

The scientific, economic, and national security benefits of leadership in quantum information science have the potential to be extraordinary, which is why so many nations are making major quantum investments. China in particular is investing billions of dollars in quantum R&D, quantum infrastructure, and quantum workforce development because it sees quantum technologies as opportunities to leapfrog the U.S. economically and militarily.

In 2018, the House Science, Space, and Technology Committee helped lead a bipartisan, bicameral effort to improve our quantum standing by passing the National Quantum Initiative Act. It was an excellent start, and the Quantum Industry Coalition strongly supported it. But still the U.S. Government is not investing as much in quantum R&D as China is. The early lead we enjoyed is shrinking, and has disappeared in certain areas of quantum technology.

This should be alarming. Fortunately, answering our competitors' spending with more money of our own is not the only solution. The federal research enterprise - including the National Labs, NSF, NIST, and federally-supported research at universities, FFRDCs, and elsewhere - works best when it can leverage the full power of U.S. companies.

The level of quantum innovation taking place at U.S. companies today is unmatched anywhere in the world. From emerging companies to global powerhouses, the private sector is developing the quantum hardware, software, and applications that will help secure U.S. leadership.

The key question is how the federal government can maximize the value of the U.S. private sector's contribution. We offer a few thoughts, focused on the role of the Department of Energy and the National Laboratories:

- **Encourage translational research with near-term applications.** Too often, federal research priorities are limited to fundamental, basic research - excluding applied research that leads to near-term applications. This was a weak point of the otherwise largely successful National Nanotechnology Initiative; basic research conducted in the United States was often commercialized overseas, to our competitors' benefit. It is a challenge we see across federal research. Near-term applications should be a priority, and separate funding should be made available for applications with a development timeline of one, two, and three years.

- **Host user facilities with a variety of hardware, software, and applications.** DOE and the country would benefit greatly from a network of user facilities with a wide variety of commercially viable quantum computers and related hardware, available via cloud access, and supporting a wide variety of commercially viable software and applications. These facilities would help reduce the cost of access to quantum computing capabilities. They would also become hubs that would begin generating research and economic value based on the size and quality of the user group. The user facilities should be public-private partnerships involving small and large companies.

- **Provide on-ramps for small and medium-sized businesses.** Beyond the SBIR and STTR programs, DOE should look for ways to make it easier for small and medium-sized businesses to participate. This includes funding opportunities, but also being a good customer: operating at the speed of small business; avoiding unnecessary contracting complexity; and being willing to work with companies of all sizes.

- **Promote quantum workforce development.** A highly-qualified, large, and diverse domestic quantum workforce is key to long-term quantum leadership. Throughout its activities, DOE should seek to expand the domestic quantum workforce. This begins with education, including support for quantum-track instruction and research opportunities at the high school, secondary, and post-secondary levels. It also includes providing access to quantum computing capabilities for a diverse and growing group of students, researchers, and professionals, along with public-private partnerships to provide opportunities to enter the quantum fields.

Thank you for your continued leadership on these important matters. We would be happy to answer any questions you may have.

Sincerely,

Paul Stimers
Executive Director
Quantum Industry Coalition
www.quantumindustrycoalition.com