ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE
to the
U.S. DEPARTMENT OF ENERGY

MEETING MINUTES

March 26-27, 2019

Cambria Suites Rockville
1 Helen Heneghan Way, Rockville, MD 20850

Meeting Minutes
ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE (ASCAC)

The U.S. Department of Energy (DOE) Advanced Scientific Computing Advisory Committee (ASCAC) convened on Tuesday and Wednesday, March 26-27, 2019 at the Cambria Suites Rockville 1 Heneghan Way, Rockville, MD. The meeting was open to the public and conducted in accordance with the requirements of the Federal Advisory Committee Act. Information about ASCAC and this meeting can be found at https://science.osti.gov/ascr/ascac

ASCAC Members Present

Keren Bergman (online)
Martin Berzins
Jacqueline Chen (online)
Silvia Crivelli
John Dolbow
Thom Dunning (online)
Tim Germann
Susan Gregurick
Anthony Hey
Gwendolyn Huntoon (online)
Richard Lethin
David Levermore
John Negele (online)
Linda Petzold
Daniel Reed (Chairperson)
Vivek Sarkar
Krysta Svore (online)

ASCAC Members Absent

Jack Dongarra
Vinton Cerf
Barbara Chapman
Satoshi Matsuoka
Dean Williams

Also Participating

Debbie Bard, Lawrence Berkeley National Laboratory (LBNL)
Frazier Benya, The National Academies of Science, Engineering, and Medicine (NASEM)
Steve Binkley, Deputy Director for Science Programs, Office of Science (SC), Department of Energy (DOE)
Christine Chalk, ASCAC Designated Federal Officer, Program Manager, Oak Ridge Leadership Computing (OLCF), Advanced Scientific Computing Research (ASCR), DOE
Paul Dabbar, Under Secretary for Science, DOE
Lori Diachin, Deputy Director, Exascale Computing Project (ECP)
Roscoe Giles, Boston University
Barbara Helland, Associate Director, ASCR, DOE
Elise Jennings, Argonne National Laboratory (ANL)
Bill Johnston, LBNL
Harriet Kung, Associate Director, Basic Energy Sciences (BES), DOE
Steve Lee, ASCR, DOE
Bronson Messer, Oak Ridge National Laboratory (ORNL)
Lynne Parker, Office of Science and Technology Policy (OSTP)
Nicholas Peters, ORNL
Terri Quinn, Lawrence Livermore National Laboratory (LLNL)
Rob Ross, ANL

Attending

Francis Alexander, Brookhaven National Laboratory (BNL)
James Ang, Pacific Northwest National Laboratory (PNNL)
Laura Biven, DOE
David Brown, LBNL
Richard Carlson, DOE
Phil Colella, LBNL
Leland Cogliani, Lewis-Burke
T. Reneau Conner, Oak Ridge Institute for Science and Energy (ORISE)
Claire Cramer, DOE
Kristin Ellis, DOE
Katherine Evans, ORNL
Richard Gerber, LBNL
Bruce Hendrickson, LLNL
Jeff Hittinger, LLNL
Paul Hovland, ANL
Mike Ignatowski, LLNL
Fred Johnson, Retired, DOE
Kerstin Klease van Dam, BNL
Alexander Larzelere, US Council on Competitiveness (USCOC)
Randall Laviolette, DOE
Michael Martin, NERL
Sonia McCarthy, DOE
Sandra McLean, DOE
Chris Miller, DOE
Esmond Ng, LBNL
Nicholas Peters, ORNL
Sonia Sachs, ASCR, DOE
John Sarrao, LANL
Michelle Schwalbe, NASEM
Mike Schulte, AMD Research
Bill Spotz, DOE
James Stewart, Sandia National Laboratory (SNL)
Ceren Suset, DOE
Hikmet Terzic, DOE
William Vanderlinde, ASCR, DOE
Stefan Wild, ANL
Justin Whitt, OLCF
Carol Woodward, LLNL

Tuesday, March 26, 2019

Daniel Reed, ASCAC Chair, called the meeting to order at 8:30 a.m.

View from Germantown, Barbara Helland, Associate Director, ASCR, DOE

In fiscal year (FY) 2019 ASCR started quantum application and computing application teams, funded two new Mathematical Multifaceted Integrated Capabilities Centers (MMICCs), four new scientific data management, analysis, and visualization projects, as well as Artificial Intelligence (AI) and Machine Learning (ML) projects. NERSC-9, Perlmutter, was baselined in September 2018 (delivery expected October 2020), Aurora at Argonne Leadership Computing Facility (ALCF), and Frontier at Oak Ridge Leadership Computing Facility (OLCF) have been baselined and expect delivery in 2021 and 2022 timeframe. Summit was completed and accepted at OLCF in February 2019. Annual reviews are planned for Aurora, CD-2 baseline review for ESnet, and baseline review for Exascale Computing Project's (ECP). A Funding Opportunity Announcements (FOA) on quantum algorithms and other crosscutting technologies is forthcoming. Facilities will continue to operate at greater than 95%; Edison, Titan, and Mira will cease operation in 2019. Facilities funds in FY19 enable a heterogeneous upgrade at NERSC, long lead-time procurements for ESnet-6, Exascale upgrades to ALCF and OLCF, and funds for PathForward milestones on ECP with vendors.

In the President's FY20 budget request, ASCR has funding for ML, Quantum Information Science (QIS) and Exascale Computing, ESnet-6 upgrade, and facility operations.

ASCR will complete the final site preparations for NERSC-9, continue SciDAC partnerships, and have partnerships in Beyond Moore's law, QIS, AI, and data-intensive science. The applied math program will reorganize investments in fundamental research in AI, ML, and other activities. Funding is available for centers to support the National Quantum Initiative (NQI) in partnership with Basic Energy Sciences (BES) and High Energy Physics (HEP). The Computer Science Graduate Fellowship (CSGF) is funded at $10M. AI money was moved back into the applied math ($41.5M) and Computer Science (CS) ($38.7M) programs. SciDAC funding is $60.9M, and Small Business Innovation Research/ Small Business Technology Transfer (SBIR/STTR) is increased to $5.3M. FY19 funding allows NERSC to complete a power upgrade; OLCF and ALCF receive an extra $21M between the two; and ESnet is down slightly. The Exascale Computing Project is down $40M in accordance with the project baselines.

Hey asked about the funding increase for applied mathematics. **Helland** explained the increase is primarily for AI, which is now in both CS and applied math. **Berzins** interjected that AI is described as being part of the applied math program. **Helland** said ASCR received $36M for AI. Funding supports foundational research in applied math and co-design of the distributed computing ecosystem, including vendor partnerships and collaboration tools. AI requires data; CS is focused on data management and data curation. SciDAC and SC programs have funding for partnerships to broaden the applicability of AI for big data solutions to SC challenges.

William Spotz is the new ASCR applied math program manager. The NERSC program manager position was posted March 22, 2019 and a CS program manager position has been approved. ASCR will soon release a Request for Information (RFI) based on questions from the in-situ management workshop (January 2019); an FOA is anticipated late summer 2019.

Discussion

Levermore inquired about the differences in AI and QIS impacts. **Helland** said industry is ahead in AI. Opportunities in AI include applying AI to scientific computing, improving reproducibility, and understanding how results were obtained. ASCR must be ready when SC programs need to take advantage of QIS capabilities.

Sarkar asked about plans to add additional CS program managers. **Helland** explained ASCR has approval to hire one more CS person; currently there are three CS program managers.

View from Washington, Paul Dabbar, Undersecretary for Science, DOE

Kelvin Droegemeier was confirmed as Director of the Office of Science and Technology Policy (OSTP) at the White House. Since the DOE Leadership Team has been in office the SC budget is up by 25%; ASCR's budget is up 45%. Interest in innovation, science, and economic competitiveness has driven stakeholders to fund research. The Secretary of Energy routinely expresses the importance of building tools that move research forward. Increased support to SC has translated into funding major user facilities, and continuing work on Facility for Rare Isotope Beams (FRIB), Linac Coherent Light Source-II (LCLS-II), and Long-Baseline Neutrino Facility/Deep Underground Neutrino Experiment (LBNF/DUNE). In FY20, Electron Ion Collider (EIC), stable isotope production, and extreme condition upgrades move forward.

DOE, National Science Foundation (NSF), and National Institute of Standards and Technology (NIST) are the three major federal agencies in the National Quantum Initiative (NQI). NSF focuses on education centers for QIS-based programs, NIST will work on commercialization, and DOE will focus on up to five quantum research centers. DOE requested the quantum centers be fully funded.

Dabbar expressed that the uniqueness of the national lab complex is not the amazing user facilities but the people and culture. AI applications will bring together top hardware, applied math and algorithm developers, and subject matter experts from a wide range of science, applied, and defense applications.

What is important about an organization is the core values. The core value of DOE is the passionate and ingenious people that move the body of knowledge forward for humanity. People with passion can discover the undiscovered universe.

Discussion

Sarkar offered to provide a briefing on the exploratory research conducted by members of ASCAC. **Dabbar** stated keeping 40% of all dollars for research has always been very important; it is one way SC balances building user facilities and supporting research.

Levermore expressed appreciation that Dabbar differentiated between technology and product. Moving technology development under product, as part of the exascale funding, raised concern. **Dabbar** said as exascale moves forward it is dropping a bit; balance is very important. The DOE national lab complex exists because even universities with huge endowments cannot build what the DOE user facilities provide; a point of pride in discovery science.

Update on 40th Anniversary Subcommittee, Bruce Hendrickson, LLNL

Hendrickson reminded ASCAC of the charge for the ASCR 40th report. Outcomes from the March 25, 2019 meeting included redefining document sections, developing a consolidated plan, and finalizing responsibility for content production. The three-element plan is to continue with a detailed history document, contract with Krell to write and produce a glossy, impact-centric document, and continue to collect raw materials for future use.

The impact-centric document will be structured around ~10 exemplar impact stories. Possible stories are being finalized; these will have a clear impact, cut across ASCR's investments portfolio, and will touch upon many different ASCR investments over the decades. The impact-centric document may be 40-60 pages, and is anticipated in July-August 2019; the history document by September-October 2019.

Discussion

Hey suggested involving Office of Science and Technology Information (OSTI) in archiving. **Hendrickson** stated the subcommittee's intention is to involve OSTI.

Levermore explained the NRC study, *Math 2025*, had a huge impact, and suggested the subcommittee explore archives for potential declassification stories.

Sarkar asked if the exploratory funding contributions to accomplishments will be discussed. **Hendrickson** said the report will include lessons learned from, and recommendations for, different modes of funding.

Crivelli asked why the structure of the history document is discipline-centric. **Hendrickson** explained multi-disciplinary achievements will be included, but the story flows better structured as historical ideas and they tend to be discipline-centric.

Germann suggested that the Krell interviews be captured and archived. **Hendrickson** agreed and indicated that Messina has collected several interviews.

Gregurick asked if other federal advisory committees have provided historical perspectives on impacts of ASCR investments. **Hendrickson** stated the subcommittee has collected SC office leaders' thoughts and appreciated the suggestion of the advisory committees.

Martin Berzins, standing in for Dan Reed, called for a break at 9:58 a.m. and ASCAC was reconvened at 10:20 a.m.

Quantum, AI and Exascale: SC Collaboration, Harriet Kung, Associate Director of BES

Kung stated that almost all of BES's research priorities touch on ASCR resources or priorities. ASCR, BES, and HEP will initiate new quantum centers in FY20. The Energy Frontier Research Centers (EFRC) have benefitted from co-funding from ASCR. The Energy Innovation Hubs utilize ASCR resources, especially in high performance computing (HPC). Computational Materials and Chemical Sciences (CMCS) is part of the Exascale Computing Initiative (ECI) and is preparing the BES community to take full advantage of the exascale computing platforms. BES is contributing to the QIS initiative via next generation quantum systems, quantum computing, and user capabilities. BES is also partnering with ASCR in quantum testbeds and algorithms.

The EFRCs were launched via a program in 2009 with Recovery Act funding and with some co-funding from ASCR. This modality combines grand science challenges and use-inspired challenges identified by Basic Research Needs (BRN) workshops. In FY20, BES plans to issue an EFRC FOA at $40M presenting the potential for a joint funding opportunity with ASCR. In anticipation of the ECI, the materials science community began a collaboration with ASCR to develop high accuracy software. Over 16 new or enhanced open-source software packages have added capabilities for predictive design of materials and simulation of atomic interactions.

Kung shared three examples of collaborations with ASCR, the Joint Center for Artificial Photosynthesis (JCAP), the BRN for Microelectronics workshop, and user facilities. Kung appreciated the positive reactions to the BES@40 report; ASCR has direct or indirect credit in every one of the stories. The High Stakes Race in High Performance Computing story highlighted BES' direct connections with ASCR.

Discussion

Levermore mentioned a partnership between BES and ASCR in training AI with heterogeneous data sets. **Kung** stated infusing physics, applied math, and CS into the ML and AI methodologies is critical; the validation component is important to BES.

Lethin asked about BES programs to accelerate the transition to market, entrepreneurship, SBIR, and the relationship of CO2 to carbon sequestration. **Kung** explained BES is the largest contributor to SBIR/STTR programs. BES does not have any dedicated program for entrepreneurship. EFRC Principal Investigators (PI) took it upon themselves to branch out secure funding from other sources. BES funds basic research and must consider the best use of federally appropriated dollars. Transitioning that knowledge into the applications side is a challenge, especially in renewable energy. BES has a life cycle view towards CO2, from capture to utilization to conversion and to some sequestration. Both the CMCS and the Geosciences portfolio's have sequestration aspects.

Storage Systems and Input/Output (SSI/O) Workshop, Rob Ross, ANL

Ross discussed the SSI/O workshop (September 2018) that focused on the need for new research in SSI/O. Workshop participants approached opportunities for research from a market

pull/ technology push perspective. There are four drivers for changes in application and technology opportunities. Five key questions were asked by attendees concerning productivity, AI/ML, streaming, workflow, and trust. Four research opportunities were discussed including rich data formats, metadata, and provenance; in situ and streaming data analysis; monitoring, production, and automation; and architecture of systems and services. An NSF workshop held in summer 2018 came to similar conclusions. Opportunity exists for ASCR to bring multi-disciplinary teams together and lead in this space.

Discussion

Levermore asked if integrating digital with analog data was considered. **Ross** explained discussions in the 5-7 year timeframe focused on traditional data sources. Workshop attendees talked about bio-inspired techniques for storing data, and analog data. The SSI/O community is still exploring what will happen in the quantum space.

Lethin inquired about recommendations for a research program aligned with ASCR's networking research. **Ross** said there is a need for greater connectivity between the communities and technologies such as scheduling the system, workflow, and resource management. The workshop focused more on connections between different classes of system software and network management software.

Hey mentioned DNA storage. **Ross** said its archival role was discussed. Making good use of DNA storage is an open question.

Gregurick asked if there was discussion on interconnects to avoid data movement. **Ross** stated addressing interconnects requires better decisions about placement and locality. The question raised was how to help tools fundamentally reduce the total volume of data.

Sarkar asked about the ASCR community sustaining system software and moving to new interfaces for new technologies. **Ross** explained there is balance to transitioning successful software research into production. In the context of the workshop, people recognized the potential to have impact by virtue of containerizing.

Reed requested lessons learned that distinguish what ASCR does in the scientific space from big data activity in the commercial space. **Ross** shared examples from his research team. His teams developed a tool for I/O characterization of HPC applications with facilities in mind. The composition project has been broken into pieces that are palatable to vendors. The team can explores known problems; only projects that can yield a production tool are pursued.

Reed dismissed ASCAC for lunch at 11:41 a.m. and reconvened at 1:05 p.m.

ECP Update and Partnering with Office of Science Computing Facilities, Lori Diachin, ECP Deputy Director, and Terri Quinn, Director, Hardware and Integration, ECP

The ECP mission is to develop applications, software stack, and research in hardware architectures for the U.S. vendor community and provide the connection between the application and software back to DOE facilities. ECP was launched in 2016; there are six core labs, three technical focus areas, and the program management office. ECP is funding ~100 R&D teams and 1,000 researchers. Three technical areas in ECP are application development (AD), software technology (ST), and hardware and integration (HI).

AD focuses on developing mission ready applications to enable new kinds of science. The mission areas support 10 DOE offices plus the National Institutes of Health. AD technical

areas include chemistry and materials science, earth science, energy, data analytics and optimization, national security, and six co-design centers.

ST funds 80+ products in programming models and runtimes, visualization tools, and math libraries. PathForward supports six vendors and ends in 2020. The ST ecosystem includes programming models, development tools, math libraries, and data and visualization. Key themes in ST are exploration and development of new algorithms, high-concurrency node architectures and advanced memory and storage techniques, and enabling and using APIs.

Over the past 6 months, ECP has held major meetings, conducted external reviews of ST and AD projects, and held the 3rd annual meeting in January 2019. Future activities include engaging with the facilities, reviewing application teams, and final design; CD2 review will occur in December 2019.

Berzins asked about the gap that exists between arithmetic intensity of the applications and new architectures. **Diachin** shared an example. The small modular reactors team reorganized their entire algorithm by grouping particles together that were looking at a similar kind of event; the new algorithm vecrotizes better on the accelerated architectures.

Quinn discussed the execution of ECP and covered HI. HI is primarily responsible for the partnership with facilities. Three example projects were hardware innovations with vendors, application preparation in porting and optimization, and software deployment and testing. Quinn shared accomplishments, an 18-month plan, and the state of deployments and integration at the HPC facilities. Facility engagement plans, signed over 1 year ago, are mapped to project plans. The plans were updated in a March 2019 meeting, developing a shared vision, with the facilities, of how ECP will proceed is a new goal; outcomes of the meeting were drafted, appendices updated, increased communications planned, and the shared vision noted.

Discussion

Reed asked for the next step for ECP. **Quinn** stated all the labs are interested in software deployment; some formalism or structure needs to bind this together. When PathForward ends all DOE hardware investment ends. Quinn stated DOE should invest in industry; those the companies have excellent people who are thinking about this community's problems.

Levermore asked if ECP had been thinking beyond porting and improving algorithms to opportunities for new codes, new modeling, and new simulations. **Diachin** said the selection of the science grand challenges integrated that kind of thinking. Many of the projects are considering different kinds of physics combinations. **Quinn** added that software has a role such as new algorithms targeted towards the anticipated architectures.

Petzold was impressed how soundly ECP brought the labs and companies together towards a common objective that will extend well past Exascale.

Hey was reminded of the Advanced Simulation and Computing program, stating that once the funding was gone there was intense competition between the labs.

Sexual Harassment, Frazier Benya, National Academies of Science, Engineering, and Medicine

Benya shared the findings and recommendations of the NASEM study on sexual harassment of women. Sexual harassment is a form of discrimination consisting of three types of behavior: sexual coercion, unwanted sexual attention, and gender harassment. There were five key findings: extensive sexual harassment in the fields of science, engineering, and medicine; gender harassment is the most common; sexual harassment undermines research integrity, reduces the talent pool, and harms targets and bystanders; legal compliance is necessary but not

sufficient; and changing the climate and culture can prevent and effectively address sexual harassment.

Severe, pervasive gender harassment can do the same professional and psychological damage as an isolated instance of sexual coercion. Fifty percent of faculty and staff have experienced sexual harassment in the last 1-2 years; 20%-50% of students experienced sexual harassment at the hands of faculty and staff members.

Sexual harassment has adverse effects on bystanders, co-workers, workgroups, and the entire organization. Dozens of studies have made clear that sexual harassment derails women's work lives. Research shows that those who simply witness sexual harassment targeted at others in the same setting will report negative outcomes that parallel those of direct victims; these effects emerge for both male and female employees.

In research environments, sexual harassment violates the standards and values of research integrity. The committee concluded that the cumulative effect of sexual harassment is significant damage to research integrity and a costly loss of talent.

The least common response to sexual harassment is formal reporting. The committee concluded that effectively addressing sexual harassment through the law, institutional policies or procedures, or through culture change requires taking into account that targets of sexual harassment are unlikely to report the harassment they experience and are often going to face retaliation. The committee recommended moving away from a culture of compliance and toward a culture of respect.

A male-dominated environment and organizational tolerance of sexual harassment are key predictors of a high-risk environment. Organizational climate was found to be the greatest predictor of sexual harassment. Four recommendations from the report included creation of diverse, inclusive, and respectful environments; diffuse hierarchical and dependent relationships between trainees and faculty; provide support for targets; and improve transparency and accountability. The committee also made four recommendations to federal agencies focusing on evaluating the effectiveness of policies; providing attention and resources; instituting rewards and incentives; and requiring violations be reported.

NASEM formed an Action Collaborative on Preventing Sexual Harassment in Higher Education. The action collaborative objective is to spur evidence-based policies and practices at the individual and system levels for addressing and preventing all forms of sexual harassment in higher education and promoting a culture of civility and respect.

Discussion

Lethin asked what actions were necessary for funding agencies and the labs. **Benya** suggested looking at the recommendations in the report. Federal agencies need to address issues similar to universities, and support academic institutions' actions to address the issues.

Sarkar was curious if professional organizations are taking similar actions to the Association for Computing Machinery (ACM) policies against harassment and unacceptable behavior. **Benya** explained the ACM policies are consistent with the committee's recommendations for professional societies. The American Geological Union (AGU) has been a leader on this issue. Policies that set expectations of behavior, especially ones that include a range of increasing consequences, are consistent with the committee's recommendations. American Association for the Advancement of Science and AGU also consider professional misconduct in their award decisions. AGU will provide free legal services to any trainee who requests it for addressing illegal discrimination or harassment.

Hey shared that in a previous position he addressed gender discrimination by forming a Women's Engineering Group, led by the first female engineer officer in the Royal Navy. She had experienced discrimination first hand, having to be helicoptered off the ship each night and helicoptered on the next morning. **Benya** said new research, by one of the committee members, reports alarmingly high rates of peer-on-peer harassment in a specific discipline of science. Student-on-faculty harassment, called contrapower harassment, is a problem and understudied. Benya shared a contrapower example of a female professor who addressed a male student's gender harassment every semester, noting it is an illustration of a hostile work environment and the academic institution is responsible.

Benya shared an anecdotal example about organizational tolerance. In a social science experiment, researchers created a lab to simulate tolerance for sexual harassment (images conveying and actors portraying gender harassment). Participants were identified as holding or not holding sexist beliefs. When exposed to the environment both groups started sexually harassing. The researchers hypothesize creating an environment that conveys zero tolerance for sexual harassment can get even those who hold sexist beliefs to stop sexually harassing.

Petzold shared her experience of discrimination stating it has to do with a small number of male colleagues who are threatened by women's success; the more successful a woman is, the more she will be threatened.

Crivelli mentioned that often when women complain informally they are discouraged. The men are too important, they are doing great science, they bring in money, or they are very useful. **Benya** said years of sexual harassment research shows that institutions fail to see all of the costs of keeping the harasser. The cost of setting up a lab only to have a harassing researcher push out a new faculty member is an expense institutions are not counting yet. Institutions were fearful of the litigiousness of perpetrators, but with the increase in legal and social action, institutions may become equally concerned that victims of sexual harassment may sue if policies and procedures are not adequately followed.

Quantum Networks for Open Science (QNOS) Workshop Report, Nicholas Peters, ORNL

DOE has a distributed science environment enabling more opportunities to use science resources that are networked and accessible. The QNOS workshop, held September 2018, was charged to identify opportunities and challenges of developing scalable quantum networks through optical fiber. The questions guiding the workshop focused on ASCRs involvement, technical challenges, and near- and long-term goals. The included quantum enabled science applications, the subsystems, architectures, protocols and components, and modeling and management. Motivating quantum science applications, quantum network architecture, devices and subsystems, and operations and controls were the four breakout topics discussed.

Attendees identified big advantages in hybridizing continuous and discrete variable approaches. Quantum networking is a new field. Most of the community lacks classical networking knowledge. A long-term collaboration between people with diverse skills working to solve all these problems is expected.

Discussion

Gregurick sought clarification on the network being leveraged by the University of Chicago, University of Illinois, University of Wisconsin, and ANL. **Reed** explained roughly two decades ago the state of Illinois in partnerships with ANL, Fermi, University of Illinois and federal partners acquired a Indefeasible Right of Use (IRU) for fiber that became part of the

trans-continental anchor for the NSF Teragrid. The IRU still exists and that is the testbed fiber connection being used for the quantum network among these institutions.

Hey asked if the papers from Delft University of Technology (TU Delft) have been considered. **Peters** explained TU Delft is following along the same lines as the rest of the quantum community. Their concept of operations in based on 2-qubit entanglement distribution with the idea of that being useful for quantum key distribution. While they have demonstrated some of the underlying functionality, a lot of work still needs to be done.

Svore asked what key applications of a quantum network would be targeted in advance of quantum computing nodes. **Peters** responded that quantum key distribution applications can be deployed now. Networking quantum sensors follows the security application and precedes computing applications.

Update from the Subcommittee on Future High Performance Computing (HPC) Capabilities (Technologies beyond Moore's Law), Vivek Sarkar, ASCAC

Sarkar reminded ASCAC of the charge and the committee's interpretation to contain two timeframes, post-Exascale (2020's) and post-Moore (2030's+). Post-Exascale technologies comprise reconfigurable logic, memory-centric processing, and silicon photonics; Post-Moore technologies include neuromorphic computing, quantum computing, and analog computing.

The six findings in the report concentrated on preparing for a period of uncertainty, extreme heterogeneity, preparing applications and system software for extreme heterogeneity, developing early testbeds, open hardware and open interfaces, and synergies between HPC and mainstream computing. Six recommendations focused on SC's leadership role in future HPC roadmaps, application readiness and exploratory research in applied math and algorithms, exploratory research in open hardware interfaces and components, and system software, preparing users for post-Moore computing, and recruiting, growing, and retaining talent.

Discussion

Berzins expressed the need to characterize the scale of arithmetic intensity and move the applications forward for future architectures. **Sarkar** said the subcommittee saw an analogy with Exascale readiness: first gaining a good sense of the characteristics of exascale computers, then holding workshops in different application areas to discuss possible implications for the applications.

Reed asked about the community's preparedness for future architectures. **Sarkar** stated there is a lot of uncertainty in future HPC directions. More exploration must be done first.

Hey asked if NSF has a complementary program. **Sarkar** said NSF is investing heavily in quantum computing such as the NSF Expeditions on quantum, and the Exploiting Parallelism and Scalability (XPS) program for parallelism. The science applications drive DOE to have a different scale and focus.

Crivelli asked about synergies with private investment and raising public awareness. **Sarkar** indicated where there are synergies, the open hardware trend can be helpful because there are domain-specific accelerators deployed everywhere. For example, Google is using the tensor processing unit; numerous start-ups are using hardware accelerators for deep learning. DOE has an opportunity to leverage components for different purposes, both for simulation and for data analytics of experimental data. DOE contributed to awareness for exascale; DOE can lead in this larger dialog as well.

ASCAC unanimously accepted the report.

International Collaborations in SC Programs, Steve Binkley, DOE SC

The science programs in DOE have strong international collaborations. Currently there are 62 SC international agreements (IA) with 16 countries, many involving ASCR facilities, research, and program activities. Seventeen additional IAs are under development. In addition to formal government-to-government agreements, there are informal agreements including Memoranda of Agreement (MOA), Cooperative Research and Development Agreements (CRADA), and Strategic Partnership Programs (SPP) through DOE/SC.

Development of IAs are governed by seven principles: quid pro quo–mutual benefit; Mutual respect; Maintain openness, transparency, respect for individuals; Research for peaceful purposes; Community engagement and buy in; Use of rigorous project management, where appropriate; and Governed by formal, government-to-government agreements.

The Deputy Secretary of Energy issued a memorandum in December 2018 concerning security issues related to sensitive technologies such as AI and QIS. Chief Research Officers (CRO) at labs are working on a list of sensitive technologies; it is anticipated by the end of 2019. There are four named sensitive countries, China, Iran, Russia, and North Korea.

A second memorandum, issued in January 2019, focuses on foreign talent programs. An individual in a research institution, lab, or university, will not receive U.S. funding and foreign talent funding at the same time. CRO's are working through details to implement this policy at national labs. Outreach activities are occurring with specific universities and professional organizations to construct policies and controls. DOE has also entered into discussions with the Association of American Universities and the Association of Public and Land-grant Universities, to gather input and feedback, with substantive information anticipated by summer 2019.

Discussion

Lethin asked how DOE is coordinating with the Department of Commerce (DOC), Department of State (DOS), and International Traffic in Arms Regulations (ITAR) on sensitive technologies. **Binkley** said there are active discussions with DOC in export control, and with DOS, and with ITAR, as well as interagency coordination. Classification guidance exists for quantum technologies, but not in biotechnology or AI. DOE will apply its own judgement, and work with interagency venues.

Reed asked about explaining, to the academic community, some judicious balance and assessment of open science and collaboration versus competitiveness and national security issues must occur. **Binkley** stated boundaries must be worked out. DOE is trying to identify technologies that are dual use and pose a risk to the U.S. Part of the solution is to clearly communicate and carefully articulate the risks.

Update from Subcommittee on Transitioning from Exascale Project, Roscoe Giles (online)

The charge for the subcommittee is to provide guidance, strategies, and approaches that are key to ensuring future U.S. leadership; and to identify key elements of ECP that need to be transitioned to ASCR research program or other new SC/ ASCR initiatives. The subcommittee will provide recommendations for capturing the lessons learned from ECP, support technologies and applications development, and inform ASCR's future investment strategy. The draft report will be ready by the September ASCAC meeting with the final report by December 2019.

The subcommittee has been formed and have gathered information from ECP via an all hands meeting (January 2019) and leadership team meeting (March 2019). Future meetings will

be held with stakeholders and there will be conversations with the ECP Industry Advisory Committee. The process for gathering information and advice will be a transparent and open.

Themes heard so far include maintaining the base research program in ASCR, support for exascale computing systems, sustaining development support and co-design teams, impact on the workforce, and management lessons for large R&D projects, maintaining awareness of transition effects in ASP and NNSA, and learning from interactions with industry.

Discussion

Berzins noted there is a challenge in moving all the applications forward after ECP. **Giles** said the subcommittee would address challenges and mechanisms for addressing them.

Levermore volunteered to help with this committee.

Public Comment

None.

Wednesday, March 27, 2019

Overview of the American AI Initiative, Lynne Parker, OSTP

The President signed an executive order (EO), Maintaining American Leadership in Artificial Intelligence, on February 11, 2019. AI.gov is a portal for federal AI activities for innovation, industry, American worker, and American values. The six areas of AI governance and activities are prioritize AI R&D, remove barriers to innovation, ensure AI-ready workforce, defend national security, leverage AI for government services, and lead international AI deliberations. The AI Interagency Working Group (AI-IAWG) is a community of practice, responsible for many actions related to the federal government and AI.

Sarkar asked if robotics, natural language processing, and visualization are within the scope. **Parker** indicated all of that is within the scope. **Levermore** pointed out the scope does not include AI for materials development and other engineering and science applications. **Parker** said the scope is broad and does not call out specific applications, but pattern recognition, robotics, and applications would be included.

The EO calls out the importance of investing in fundamental algorithmic advances as well as the ecosystem, the infrastructure, and the data. The EO does not say the data will be made publicly available but there will be increased access to the data for AI R&D research.

Crivelli asked about synthetic data sets. **Parker** said the European Union's General Data Protection Regulation (GDPR) is creating synthetic data. Often the outliers cause the most challenges for AI research and it cannot be properly modeled.

Lethin questioned if data on the internet might be enough. **Parker** explained it is not broad enough. Traffic and flow data, held by the Department of Transportation, may not be available, broadly speaking. Data.gov exists but the data are not discoverable. The Federal Data Strategy is looking for solutions to maintain data as well as metadata; AI is leveraging that for discoverability.

The EO directs agencies, who have the computational infrastructure, to prioritize resource access to AI researchers consistent with their mission.

Reed mentioned that the biggest impediment to federal data was often federal IT staff. **Parker** said there is a distinction; the AI initiative calls for access for federally funded AI research, not federal use itself.

The process is meant to help agencies consider and determine governance in the types of AI applications consistent with their mission. The EO calls for a two-step process. Office of Management and Budget issues a guidance memorandum, and then agencies develop best practices or recommendations for how AI should be used.

The EO calls for agencies to prioritize opportunities for people to learn about AI. The focus is not about building tools but about being comfortable using the tools.

Crivelli suggested internships target staff members at community colleges and small liberal arts colleges to gain teachable skills. **Parker** said the EO calls out opportunities such as curriculum development that are designed to engage teachers.

Gregurick asked about augmenting the skills of the highly trained workforce. **Parker** said there are a number of conversations about changing how to do higher education funding, such as Pell grants, to allow for short-term training.

The Department of Defense (DOD) established the Joint AI Center (JAIC) and the Defense Innovation Board is creating a set of principles for the appropriate use of AI in defense. Defense Advanced Research Projects Agency (DARPA) is investing in "3rd wave" AI with cognitive, predictive systems; common sense reasoning systems.

A National Security Presidential Memorandum on protecting investments in AI against competitors and adversaries was issued February 11, 2019. General Services Administration (GSA) has pilot projects to provide services to the federal government on how to use AI on simple, public-facing tools.

The EO is promoting the international AI industry. Conversations focus on setting norms for the international use of AI, recognizing liberal democracies do not want to become surveillance states.

Berzins inquired about U.S. investments in AI compared to China. **Parker** said China has civil-military fusion making it difficult to compare dollar-for-dollar. Regardless of the level of U.S. investment, impact is important. **Berzins** said cuts to NSF intelligence systems and reduced funding in R&D will challenge the U.S. leadership position in AI. **Parker** explained that AI is an incredibly bipartisan area; Congress is proposing many bills for funding AI R&D and are appropriating the necessary funds.

Hey asked about mid-career professionals and mentioned studies from Price Waterhouse and McKenzie in 2017 that estimated by 2030, 25% of U.S. employment was at risk of automation. **Parker** referred to a study indicating only tasks within jobs would be affected by automation. She argued that computational thinking would help mid-career individuals be able to use the new tools effectively. **Hey** expressed concern about unemployment in Spain, Italy, and Greece among people between 16-25 years of age, stating if they do not have a stake in society they cannot be trained to use the tools. **Parker** said that the overall policy environment creates opportunities.

Levermore asked about the effect of monetarizing data on the structure of the economy; it incentivizes companies to merge thus reducing competition. **Parker** mentioned two ways to address such incentives includes developing new methods that do not require large amounts of data, and providing access to data for the everyday researcher or the everyday industry.

Lethin asked how restructuring funding models would influence the AI plan, and how the national labs and DOE can participate. **Parker** said AI is inherently multi-disciplinary;

bringing people together in the interdisciplinary enclaves helps accelerate the advances by providing perspectives from different angles.

Overview of Data science Efforts at the ASCR Facilities, Elise Jennings (ANL), Bronson Messer (ORNL), Debbie Bard (LBNL), and Bill Johnston (ESNet)

Panelists shared information about the data science programs at their respective facilities. Jennings discussed ALCF's Data Science Program (ADSP), which began in 2016. The ADSP's goal is to support big science programs and proposals that require the scale and performance of the leadership computing facilities. Successful projects have high potential impact, data scale readiness, diversity of science domains and algorithms, and can fully exploit the architectural features of Theta. Two main targets for development in ADSP are science applications and tools. Jennings shared information on two projects currently running, emerging trends being seen at ALCF, and the emerging needs. ALCF is developing data science software including the Balsam workflow manager to simplify running workflows for large science campaigns, Deep Hyper for hyperparameter optimization, and PETREL for data movement and management.

Messer discussed strategy and tactics at OLCF. Their three-pronged strategy is engagement with applications, creating leadership class analytics, and enabling infrastructure for analytics, AI, and data-intensive facilities. Messer highlighted examples of these three approaches including the Directors Discretionary projects, Programming with Big Data in R (pdbR), and using MENNDL and Titan to generate deep neural networks. Plans are to create policies that enable a larger set of data services for storage, access, or end station programs.

Bard said 35% of the projects at NERSC are focused on data analysis; the scale of required resources has also increased. Cori was designed to support data science and experimental science; Cori experiences are being used to design NERSC-9, Perlmutter, to meet the needs of large-scale simulation and data analysis, and facilities. NERSC Data Science is providing a seamless experience for users, including large-scale computing and storage resources, policies, reusable building blocks, scalable infrastructure, and expertise. The LBNL Computer Science strategic plan is supporting the Superfacility model and emphasizing engagement with users, facilities, and all of SC.

Johnston stated ESnet traffic increases 10x every four years and shares a fiber optic network with Internet 2. ESnet-6 is a next generation, different network, which will own and manage the entire capacity of fiber around the country and will have aggregate capacity of 100 exabytes per year. ESnet's goal is to provide transparent access from instruments to facilities that handle data. When LHC data increased, ESnet detected several transatlantic network paths. LHCONE is an overlay that allows network engineers to steer traffic in a transparent way.

Discussion

Sarkar inquired if any projects are analyzing features in software using statistical analysis learning. **Messer** said time is spent examining characteristics of the software. There are ways to extract data in an automated fashion. Enough data to consider software analysis is just now being seen. **Bard** explained that NERSC is scouring their systems data and system logs to understand fault tolerance and fault prediction. **Jennings** mentioned that ADSP is analyzing software as part of Balsam; hot spots, bottlenecks, improvements, interactions, output errors, and segmentation faults. ADSP is using advanced statistical techniques to deploy the software and utilize the resources successfully.

Lethin asked Johnston how the global I/O across ESnet is changing relative to deep learning. **Johnston** said ESnet is using ML techniques for automated error detection analysis, with some success detecting soft errors and characterizing data flows. LHC analysis systems are the most likely to start changing the way people use ESnet. There is potential to optimize data flows using centralized control of data parceling and analysis.

Levermore asked about augmenting simulation data with experimental data to create heterogeneous training data. **Jennings** indicated some groups at ADSP are combining experimental and real data to create large databases for training. **Bard** explained NERSC uses ML to train algorithms on simulation data, and then uses generative adversarial networks to augment the simulation set. There are examples of groups using real data to validate ML algorithms as well. A different paradigm is developing new statistical techniques to do ML on a small amount of data. **Jennings** added the real data does not look like the simulated data; the challenge is making simulated data messy and adding noise. **Messer** noted that non-Gaussian transients, such as a microwave, are harder to simulate than noise.

Huntoon asked about ESnet's activities at the service edge for non-ESnet connected users (i.e., researchers at the University). **Johnston** noted a site may use ESnet so long as it is consistent with the policy of SC. ESnet has been a leader in software-defined networking (SDN). Network service interfaces are available through the SDN, a software suite is accessible and used for setting up virtual paths and pseudo wires between two end points.

Crivelli sought clarification about plans to share data with the public and information on hyperparameter optimization tools. **Jennings** explained ADSP prioritizes those projects that advance the community and make their data open source. Hyperparameter optimization software packages are numerous. Choosing one is based on the architecture being run, how heterogeneous architecture is making use of memory, and the specific problem. **Messer** reminded BERAC all of the data programs are able to assign digital object identifiers to data to aid in curation and discoverability. ORNL can admit any use case; the generators of the data control what happens with the data. **Bard** pointed out there is a difference between making data available and making data easily available. NERSC can support data being available but cannot support thousands of scientists requesting that data every day without a dedicated infrastructure.

Hey asked if the National Cancer Institute/ DOE collaborators are using ALCF facilities or benchmarks. **Jennings** said CANcer Distributed Learning Environment (CANDLE) is using ALCF facilities for real science.

Update from Applied Math Program, Steven Lee, DOE SC

The applied math program portfolio in FY19 is $30M and supports ~50 projects. Lee congratulated Jack Dongarra who was awarded the SIAM/ ACM Prize in Computational Science and Engineering. The ASCR Applied Math PI meeting was held January 2019. Eight Early Career Awards (ECA) were given in FY17 and FY18. There is a new SIAM Journal on Mathematics for Data Science (SIMODS). Lee touched on three math centers Advances in Experimental Design, Optimization and Learning for Uncertain Complex Systems (AEOLUS) at University of Texas at Austin, Physics-Informed Learning Machines for Multiscale and Multiphysics Systems (PhILMs) at PNNL, and Multifaceted Mathematics for Rare, High Impact Events in Complex Energy and Environment Systems (MACSER) at ANL. Lee also mentioned the Center for Advanced Mathematics for Energy Research Applications (CAMERA) project, the Julia (MIT) program, and Jupyter, a sharing mechanism.

The applied math program has laid the groundwork for harnessing ML and AI for scientific purposes. Three areas of interest are considering scientific ML and AI foundations and use cases, post-Moore algorithms and programming, and enabling technologies for complex systems research.

Discussion

Berzins commented on the long-term themes and highlighted goings-on in other programs noting funding is not available to carry out and take advantage of these activities. He suggested a joint math and CS activity on the robustness of ML. He was concerned about the difficulties for the traditional DOE community in exploiting the new heterogeneous architectures; the aspirations does not match available funding.

Levermore suggested looking for synergisms and leveraging ASCR investment. **Lee** mentioned coordination within OSTP is vital to understand community activities and identify gaps.

Dolbow asked about transitioning ECA winners into the core program. **Lee** replied that most have continued funding but many former awardees are approaching the transition.

Public Comment

None.

Reed adjourned ASCAC at 11:54 a.m.

Respectfully submitted,
T. Reneau Conner, PhD, PMP, AHIP
Science Writer
ORISE/ ORAU