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

MEETING MINUTES

September 23-24, 2019

Holiday Inn Capitol
550 C Street, Washington, DC

Meeting Minutes
ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE

The U.S. Department of Energy (DOE) Advanced Scientific Computing Advisory Committee (ASCAC) convened on Monday and Tuesday, September 23-24, 2019 at the Holiday Inn Capitol, 550 C Street, Washington, DC. 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 http://science.osti.gov/ascr/ascac

ASCAC Members Present

Martin Berzins
Jacqueline Chen (online)
Silvia Crivelli (online)
John Dolbow
Thom Dunning
Tim Germann
Susan Gregurick
Anthony Hey (online)
Richard Lethin
David Levermore
John Negele (online)
Daniel Reed (Chairperson)
Krysta Svore (online)

ASCAC Members Absent

Keren Bergman
Jack Dongarra
Gwendolyn Huntoon
Satoshi Matsouka
Linda Petzold
Vivek Sarkar
Dean Williams

Also Participating

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
Barbara Helland, Associate Director, ASCR, DOE
Laura Biven, ASCR
Susan Coghlan, Argonne National Laboratory (ANL)
Julie Carruthers, DOE
Lori Diachin, Deputy Director, Exascale Computing Project (ECP)
Roscoe Giles, Boston University
Bruce Hendrickson, Lawrence Livermore National Laboratory (LLNL)
Doug Kothe, Director, ECP
Jeff Nichols, Oak Ridge National Laboratory (ORNL)
Thomas Peterka, ANL
Bill Spotz, ASCR
Justin Whitt, ORNL

Attending

James Ang, Pacific Northwest National Laboratory (PNNL)
Srinivasan Arunajatesan, Sandia National Laboratory (SNL)
Lisa Arafune, Coalition for Academic Scientific Computing (CASC)
Jay Bardhan, GlaxoSmithKline
Sam Barish, DOE
Arthur Bland, ORNL
David Brown, LBNL
Suren Byna, LBNL
Richard Carlson, DOE

Matthew Carnavos, DOE
Scott Collis, ANL
Leland Cogliani, Lewis-Burke
T. Reneau Conner, Oak Ridge Institute for Science and Energy (ORISE)
Tricia Crumley, DOE
Al Geist, ORNL
Raul Grout, National Renewable Energy Laboratory (NREL)
Mark Guiton, Cray, Inc.
Steve Hammond, NREL
Jeff Hittinger, LLNL
Thuc Hoang, DOE/National Nuclear Science Administration (NNSA)
Holly Holt, ORISE
Paul Hovland, ANL
Fred Johnson, Retired, DOE
Ben Kallen, Lewis-Burke
Randall Laviolette, DOE
Steven Lee, ASCR
Arthur Maccabe, ORNL
Michael Martin, NERL
Scott McKee, U.S. House Committee on Appropriations, Subcommittee on Energy & Water
Sandra McLean, DOE
Esmond Ng, LBNL
Mary Ann Picone, DOE
Ashley Predith, DOE
James Ricci, ASCR, DOE
Katherine Riley, ANL
Sonia Sachs, ASCR, DOE
John Shalf, LBNL
Arjun Shankar, ORNL
Julie Stambaugh, DOE
James Stewart, SNL
Ceren Suset, DOE
Pieter Swart, Los Alamos National Laboratory (LANL)
Valerie Taylor, ANL
Hikmet Terzic, DOE
Angie Thevenot, DOE
Jeff Vetter, ORNL
David Womble, ORNL
Carol Woodward, LLNL

Monday, September 23, 2019

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

View from Washington, Steve Binkley, Deputy Director, Office of Science, DOE

Dr. Chris Fall was confirmed on May 23, 2019, as Director of Office of Science. The fiscal year (FY) 2020 President's budget request provides $5.5B for SC, $500M is budgeted for Exascale Computing, $169M for Quantum Information Science, $71M for Artificial Intelligence (AI) and Machine Learning (ML), and $25M for microelectronics research. The ASCR President's budget request for FY20 was $920K, the House mark was $956K and the Senate Mark was $1.03M.

Shared threats to research and development (R&D) include intellectual property theft, forced technology transfer, undisclosed affiliations and funding sources of U.S. researchers, and breaches in the peer review process. SC recognizes that international cooperation is important for the advancement of science. DOE Order 486.1, June 2019, addresses foreign talent programs. National lab directors are defining sensitive countries and technologies. Addressing these threats consists of DOE national labs, university grants, collaborations with other federal agencies, and establishment of the Office of Science Technology Policy (OSTP), Joint Committee on Research Environments (JCORE).

Discussion

Dolbow asked for more information about peer review breaches. **Binkley** stated that proposals that come for peer review have reappeared by those who had access to proposals during the peer review process, especially at the National Institutes of Health (NIH) and the National Science Foundation (NSF).

Lethin asked about the DOE's prioritization of the data infrastructure. **Binkley** said a National Science Technology Committee (NSTC) subcommittee is looking at R&D infrastructure across the nation. With rise of the importance of data and ML and the rapid growth in data volume, multiple agencies have acknowledged the need for state of the art infrastructures to systematically manage data. For example, the High Luminosity (HL) upgrade at the Large Hadron Collider (LHC) at CERN will increase the data volume by 2 to 5 times.

Reed inquired how to raise awareness of intellectual property infringements without impinging on open collaborations. **Binkley** suggested a systematic approach to sharing information. He noted that although specifics on IP thefts and issues are frequently classified, the information is compelling. This is not a new problem, it has been an ongoing since 2009.

Svore queried if there will be an equal emphasis across data analytics, data science, and data infrastructure. **Helland** stated Laura Biven will be providing an update on the Data Roundtable.

Berzins asked how R&D protections impact open source software given its reliance on symmetry. **Helland** indicated while there are no immediate impacts, going forward some software may be removed from open source. **Binkley** said some of the new and existing collaborations do not have that symmetry and reminded ASCAC these arrangements need to be quid pro quo.

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

ML/AI, biosecurity, quantum information science, exascale computing, microelectronics innovation, the national isotopes strategy and the U.S. fusion program acceleration remain DOE priorities. The House and Senate marks for ASCR funding are $956M and $1.029B, respectively. Senate marks notably recommend no less than $160M for research in high performance computing (HPC). House and Senate language encouraged the DOE to expand relationships with NIH, NSF, and National Institute for Standards and Technology (NIST).

DOE established the Artificial Intelligence and Technology Office (AITO), reporting directly to the Under Secretary for Science. In May 2019, a request for information (RFI) on a Quantum Center received 38 comments which are published in Federal Register. DOE anticipates release of a funding opportunity announcement (FOA) shortly after appropriations for FY20 with a decision by early summer 2020. The National Quantum Initiative Advisory Committee (NQIAC) was established by Executive Order on August 30, 2019. DOE was given the authority to organize the NQIAC. An RFI has been released and nominees for NQIAC are being accepted.

Helland delivered updates on quantum testbeds, characterizing tensor flow in graphics processing units, and using AI to predict nuclear binding energy and radius. The DOE and National Cancer Institute (NCI) are using HPC to improve understanding of cancer biology and applications for more effective therapies; the tools and methodologies developed have broad applicability to ML problems.

DOE exascale systems will be available in a limited capacity in 2020, with Frontier and Aurora activity scheduled for 2021-2023. Though Titan has been retired, there are several

notable accomplishments from Titan's 7 years of activity. The ASCR Leadership Computing Challenge (ALCC) now requires a letter of interest.

The new National Energy Research Scientific Computing Center (NERSC) Perlmutter Cray Shasta machine will be available at the end of 2020. ESnet6, scheduled for completion in 2023, will go through the Director's review for critical decision (CD)-2/3 in October 2019. During ESnet6's completion, ESnet5 will be available. ESnet6 is a key partner in a newly awarded NSF project, **F**ABRIC is an **A**daptive Programma**B**le **R**esearch **I**nfrastructure for **C**omputer Science and Science Applications (FABRIC).

ASCR presented eight Early Career Research Awards, six in computer science and two in applied mathematics; three awards were presented to women. Four researchers in ASCR areas were nominated for the Presidential Early Career Award for Scientists and Engineers (PECASE). The Society of Industrial and Applied Mathematicians (SIAM) recognized three ASCR leaders, and ASCR sponsored two American Association for the Advancement of Science (AAAS) fellows.

Discussion

Lethin requested details on the funding relationship between ESnet and FABRIC. **Ben Brown**, NSF, explained that funding for FABRIC originates from the NSF's 10 Big Ideas initiative. FABRIC was awarded to a group of academic principal investigators and ESnet's partnership greatly strengthened the proposal.

Levermore inquired if Brexit will impact Esnet6. **Helland** replied that Brexit is not expected to affect Esnet6 because there are several paths into Europe.

Svore asked about the number of National Quantum Initiative (NQI) Centers anticipated and if collaboration among award centers will be encouraged. **Helland** said, pending funding, the DOE anticipates conferring awards to 2-5 NQI Centers; she emphasized that the DOE will encourage collaboration with NSF centers.

Reed commented that Susan Gregurick has been named NIH Associate Director for Data Science which presents additional collaboration opportunities. **Gregurick** stated she is happy with the existing collaborations with NCI and CANcer Distributed Learning Environment (CANDLE) and looks forward to exploring future collaborations with DOE, ASCR, and others.

Jeff Vetter, ORNL, asked about the Basic Research Needs for Microelectronics RFI. **Helland** replied that the RFI was handled through Basic Energy Sciences (BES) and updates are pending.

Svore noted the diversity of Early Career Award winners this year and asked what was different this year. **Helland** said because application pipeline issues were addressed the applicant pool was stronger this year**.**

Jim Ang, PNNL commented that the Semiconductor Research Corporation, with sponsorship from DOE, will hold a series of workshops for semiconductors focused on the industry perspective. The first workshop will be in October 2019 in Sandia, CA.

Report on Data for AI Roundtable, Laura Biven, ASCR

The SC Working Group on Digital Data (SCWGDD) supported the workshop. Participants included members of all six SC program offices, Office of Science and Technology Innovation (OSTI), DOE national labs, NIH, and NSF. Discussion centered on enhancing and enabling access to high-quality and fully traceable research data, models, and computing resources to increase the value of such resources for AI R&D. The workshop aimed to identify

key challenges, opportunities and potential next steps for the SC with an additional goal of synergizing efforts with ongoing, relevant AI Town Halls and other AI-focused workshops.

Conversations addressed several challenges in using AI for science, including the generation of FAIR (Findable, Accessible, Interoperable and Reusable) data. Participants also catalogued opportunities specifically available to the SC to address AI challenges and broadly identified additional capabilities needed to enable data science.

Discussion

Berzins asked if discussion addressed National Cancer Institute data and facilitating study reproducibility. **Biven** said that participants from relevant collaborations shared lessons learned. Conversations highlighted the need for frameworks relating data, models, and tasks.

Lethin asked to what degree this SC roundtable, industry, and other relevant organizations are synchronizing data management efforts. **Biven** asserted that interactions between parties should be strengthened. Participants discussed how the DOE could engage with groups that are developing standards for data sharing and modeling.

Levermore observed that data generators seldom make raw data public; rather, data are frequently processed. Thus, there are not only technical challenges to data management, but also communication challenges between data generators and users. **Biven** agreed that for data generators, there is a tension between anticipating the use cases for a data set and not knowing what those cases might be.

Gregurick referenced ongoing conversations about how schema can be interfaced across different data communities. **Biven** replied that roundtable discussions addressed how to make data AI ready. Concluding that interactions between the AI community and data holders would be necessary to identify key metadata needed for AI use cases.

Berzins stated that at a previous NCI-DOE meeting, Rick Stevens asserted that a new theory of data is needed. **Biven** affirmed that this is still an open challenge in the AI domain.

Crivelli asked if multimodal data integration was discussed. **Biven** confirmed that this was one of the key challenges deliberated.

Hey prefaced his question by noting that it related to a previous comment addressing OWL (ontology web language) and other semantic web technologies; various communities (e.g., biosciences) and industry (Google and Microsoft) have adopted these technologies, but most scientists do not use Resource Description Framework. He inquired if there was discussion of these approaches at the roundtable. **Biven** replied that those building schemas were invited to the roundtable, but bringing these communities together is a future challenge.

Reed asked about incentives for AI researchers to work on data issues. **Biven** responded that data is a window into the SC's mission. Making data broadly available will attract people to participate in these far-reaching missions.

Reed called for a break at 10:15 a.m. and reconvened ASCAC at 10:31 a.m.

Report on AI Town Hall Meetings, Jeff Nichols, ORNL

Three Town Hall meetings addressing the SC's AI mission were hosted in Chicago, Oak Ridge, and Berkeley from July-September of 2019. A fourth Town Hall is scheduled for October, 2019 in Washington, D.C. The AI Town Halls captured ideas, problems, requirements, and challenges in using AI for science. The rationale for these meetings included 1) AI is disruptive; 2) approximately 35 countries have AI strategies; 3) AI will become increasingly

integrated into the business landscape; 4) the White House issued an executive order regarding AI; and 5) the DOE recently formed AITO.

Town Hall discussions centered on transformative ideas for AI uses 10+ years in the future. Several themes emerged: 1) learned models will begin to replace data; 2) models will replace experiments and experiments will improve models; 3) questions will be pursued semi-autonomously at scale; 4) simulation and AI approaches will merge; 5) AI will contribute to advancing theory; and 6) AI will become a common part of scientific laboratory activities. Discussions evaluated transformative ideas based on their ability to cut across DOE expertise, facilities, and applications and to address issues of broad importance. Example topics included cosmology and astrophysics; numerical aspects of learning; ML and AI approaches to increase trust; advanced deep network architecture design; characterizing the loss landscape of science-informed ML; and exploiting variable precision for performance.

Discussion

Dunning asked if pilot studies have been considered as the next step following AI Town Halls. **Nichols** responded that pilot studies are already ongoing through the ASCR office. BES, Biological and Environmental Research (BER), and the High Energy Physics (HEP) groups are collaborating with experimental colleagues.

Gregurick requested additional information regarding bias in trustworthy AI. **Nichols** affirmed that crosscut breakout groupss will address bias in AI.

Levermore commented that the DOE is behind commercial industries, other countries, the NSF, and most universities in AI. He asserted that ASCR needs to advance AI initiatives so that the DOE is not left behind.

Berzins noted that while the DOE is a world leader in modeling and simulation, continuous improvement in simulation science is a must to ensure that applications work well with AI. **Nichols** stated that the Town Halls, especially the Berkeley Town Hall, made significant progress in addressing this topic; discussion encompassed how AI will fundamentally change modeling and simulation procedures and how AI can be used to write applications or compilers.

Lethin inquired about funding to support AI research. **Nichols** noted that the Exascale Computing Initiative was funded at approximately $4.5-4.7B. If the AI initiative is on a similar scale of funding, many issues will be addressed.

Reed observed that new research communities may be funded in the future because the AI community is broad and disjointed. He added that data holders are frequently advancing AI. **Nichols** concurred, noting the importance of industry and academia community members.

Levermore asserted that the ECP is a poor model for AI because the exascale model was DOE centric while AI is broadly distributed. Because DOE researchers are learning AI, he suggested the DOE ensure researchers stay within programmatic lines and utilize new AI technology. **Nichols** encouraged the audience to think of the Exascale Computing Initiative (ECI) rather than the ECP as a model. **Hey** affirmed that an AI for Science initiative is an appropriate SC focus because the DOE laboratories lead the world in supercomputing and generate enormous amounts of scientific data. The application of AI and deep learning technologies to these datasets represents a unique SC opportunity.

Dunning stated the DOE should be pursuing AI because it is one of the largest holders of experimental data in the world and has a strong computational science program; pilot projects are an effective path forward towards unveiling challenges. **Nichols** agreed but asserted moving

forward rapidly is important so as to not fall further behind industry and other agencies. **Dunning** answered that he is less concerned about falling behind than conducting projects well.

Berzins offered an example of how the DOE has the potential to lead in AI, remarking on the knowledge of uncertainty quantification at LANL. **Reed** cautioned that academics are transitioning to industry to conduct AI at scale.

Bland emphasized the incorporation of ethics and legal frameworks into AI advances to protect people and their privacy.

Report on In-Situ Data Management Workshop, Thomas Peterka, ANL

Peterka summarized the outcomes from an ASCR workshop held on in situ data management (ISDM) in January 2019. ISDM was defined as the practices, capabilities and procedures to control the organization of data and enable the coordination and communication among heterogeneous tasks, executing simultaneously in a HPC system, and cooperating toward a common objective. ISDM has the potential to make critical contributions to the management and reduction of large volumes of data from computations and experiments. This methodology enables scientific discovery from a broad range of data sources and on a wide scale of computing platforms. Workshop priority research directions comprised pervasive and co-designed ISDM, in situ algorithms, controllable, composable, and transparent ISDM. The workshop produced a brochure and final report, and a journal article has been submitted.

Discussion

Berzins asked how important ISDM is to DOE grand challenges. **Peterka** replied that ISDM and its core capabilities underlie and enable ML, AI, and other future technologies.

Reed asked about cultural challenges in shifting from legacy approaches to ISDM. **Peterka** stated there is growing recognition of the importance of ISDM and the need for data infrastructures that can accommodate a variety of external tools.

Reed adjourned ASCAC for lunch at 12:10 p.m. and reconvened at 1:30 p.m.

40th Anniversary Accomplishments Subcommittee Report, Bruce Hendrickson, LLNL

Hendrickson highlighted progress on the ASCAC charge to assess and document historical accomplishments of the ASCR program and its predecessors over the past four decades. Two documents will be produced, a detailed historical and technical document, and a shorter, impact-centered, and accessible manuscript. Five high-level lessons emerged from the documents: 1) a compelling and consistent vision can drive scientific revolutions; 2) different funding models are required for diverse and impactful outcomes; 3) workforce investments have been critical; 4) partnerships are essential; and 5) testbeds and platform access funding models are important. Future challenges include technology disruptions, funding balance, a software support model, broader partnerships, and workforce. Drafts of the accessible document are available and writing for the detailed document will be completed by the end of 2019.

Discussion

Berzins suggested that the detailed document incorporate an additional story to illustrate the impact of software on supercomputing. **Hendrickson** agreed with this suggestion.

Reed requested cautionary and encouraging examples of technological and political factors that have affected ASCR progress. **Hendrickson** agreed to consider this suggestion.

Levermore commented that he found the future challenge of technology disruptions particularly apt because ASCR itself was formed in response to the disruptive rise of computing. **Hendrickson** concurred, noting that science has entered a computing era where improving performance is a challenge and changes to algorithms, workflows and architecture will alter the landscape of computing.

Giles asked if the reports will relay the contributions of women and minorities behind ASCR advances. **Hendrickson** stated the detailed report largely omits names and the shorter report highlights the contributions of three individuals. He will ask the committee to provide feedback on how to gather relevant data to share the story arc of women and minorities.

Update from Subcommittee on Exascale Transition, Roscoe Giles, Boston University

Giles reviewed ASCR's charges to the Subcommittee: 1) examine ECP lessons learned for managing large collaborations; 2) evaluate ASCR's historic fundamental research investments in applied mathematics, computer science, and computational partnerships with national laboratories and 3) identify new research and development priorities in AI, quantum information systems and strategic computing.

ASCR budgets increased steadily during the exascale era. ASCR's research budget for non-ECP activities remained relatively constant with increases attributable to ECP funds. Upturns to the facilities budget can be ascribed to the ECI activity. Post ECP, there is an opportunity to deemphasize development and reinvest in research.

Dunning asked what will happen to remaining ECP funds at the project conclusion. **Helland** replied that unspent funds will be returned to the treasury.

Subcommittee preliminary findings highlighted lessons learned for managing large ECP collaborations built from DOE, SC, and ASCR communities. Notably, ECP partnerships created a productive working ecosystem supported by agile project management. However, the current workforce pipeline has reduced direct university engagement, shifted the hiring focus from research to development and limited opportunities for early career investigators. Workforce retention is a pressing matter.

The Subcommittee determined that ASCR's historic fundamental research investments enabled ECP success and credited the ECP with substantial progress in application development, software technology, hardware and integration, and project management. Importantly, developing workable nondisclosure agreement (NDA) processes has been critical for protecting vendor intellectual property. Looking forward, there are opportunities to reinvigorate funding in applied mathematics and computer science as well as invest in novel cross-disciplinary research areas including quantum computing and ML/AI for science. However, large industry investments are changing the working landscape requiring ASCR to be more nimble in its operations.

The future vision is for the national laboratories to capitalize on the overall R&D ecosystem in an agile way, by integrating industry and university talent at all levels. Weekly Subcommittee meetings continue; a workshop is scheduled for October and a community discussion in November 2019. A final report will be presented at the next ASCAC meeting.

Discussion

Dunning asked how to sustain ECP software and applications; this need might constitute a shared responsibility for ASCAC and other offices. He also emphasized that the strong rapport between team leaders and members, due to leaders' expertise, contributed to ECP success. **Giles** agreed that a mechanism enabling all National Laboratory applications to become operational

with exascale systems is needed; forthcoming recommendations may address how to maintain requisite ECP layers for a lower cost, identify the parties that should invest, and determine how to scale investments to ensure broader participation.

Berzins stressed the importance of ECP software maintenance with an analogy to building maintenance. Codes must be portable and structured to be performant. Algorithms must be adjusted for future architectures. **Giles** noted the subcommittee report will address these concerns.

Levermore said ultimately scientists will judge ECP success by making use of existing software to fully exercise hardware capabilities.

Reed inquired how to balance ECP software sustainment against future investments. **Giles** replied this situation entails tradeoffs between short- and long-term activities. Long-term, future activities will require attention because of Moore's law, Dennard scaling, and scalability. He suggested a holistic view of HPC, bringing in those who are approaching HPC from the bottom, communities outside of ASCR. People are being forced to think more broadly.

Dunning commented enabling new science and varying software technology utility may determine which ECP components are supported in the future. **Giles** concurred that exascale applications are intended for discovery; applications should facilitate transition from an idea to an exascale computation in a well-defined manner.

Dolbow said academic engagement is important for discovery using exascale applications and requested details of academic ECP involvement. **Giles** responded that academic involvement is variable, and attributable to financial management rules and other factors.

Berzins observed that ECP software tools and libraries are critical to the future success of a diversity of architectures. Existing applications are broadly applicable to the DOE's core interests.

Ang commented the ECP is large enough to have created its own de facto non-disclosure agreement standards. ECP's true impact will be assessed when advanced architectures allow the entire HPC community to leverage their applications with ECP software at scale.

Reed dismissed ASCAC for a break 3:13 p.m. and reconvened the meeting at 3:32 p.m.

DOE Exascale Systems: Aurora Argonne Leadership Computing Facility, Susan Coghlan, ALCF-3 Project Director, ANL

Coghlan summarized ALCF historical supercomputing technologies and accomplishments, segueing to the latest supercomputer, Aurora, with scheduled delivery in 2021. Preparations for the site, the computer build, and tests and applications are ongoing. Coghlan highlighted Aurora's hardware specifics. Aurora will sustain performance at >1 exaflop with double precision. Aurora's programming environment will support enhanced graphics processing unit (GPU) programming tools designed for performance and portability (including SYCL, OneAPI and DPC++) along with many of the same compilers, programming models, libraries, and tools (e.g., OpenMP 5.x) available on most supercomputers as well as most data and learning frameworks. Aurora has assembled restricted secret NDAs for hardware and corporate NDAs for software details among vendors, institutions and other involved organizations.

Argonne selected 15 Early Science Program (ESP) projects emphasizing simulation, data and learning, and covering a broad set of science domains and codes. These ESPs, in addition to 22 ECP application development projects, are being used to prepare Aurora workflow technologies, optimize libraries, frameworks and tools, and harden the software stack. The

Argonne Aurora Programming Workshop (September 2019) addressed application readiness for Aurora as well as hardware, software, and tools. Notably, the applications use a variety of programming models that are not equally supported on exascale supercomputers, presenting portability and performance challenges. There is interest in a number of programming models (e.g., Kokkos, Raja, OpenACC).

The next Performance, Portability and Productivity in HPC Forum is scheduled for 2020 in Kansas City.

Discussion

Berzins commented that Kokkos and Raja tie in well with the ECP. He asked if the Aurora risk register accounts for the project's many moving parts. **Coghlan** replied the risk register stresses scheduling; facility scheduling is relatively straightforward, but application and software scheduling must account for more unknowns.

Lethin asked about NDA export controls. **Coghlan** confirmed that NDAs adhere to export control requirements. Vendors have provided export control classification numbers (ECCN) and Export Administration Regulations (EAR) information. Export Controls is consulted before new entities are added to the multiparty agreement. NDAs are issued on a need-to-know basis.

DOE Exascale Systems: Frontier Oak Ridge Leadership Computing Facility, Justin Whitt, OLCF-5 Project Director, ORNL

Whitt reviewed the OLCF's historic heterogeneous GPU/ computer processing unit (CPU) exascale systems noting that Frontier nodes will pair each CPU with four GPUs. Frontier will be composed of >100 Cray Shasta cabinets, incorporate Slingshot interconnect, include Dragonfly topology, and perform at >1.5 exaflops with a storage capacity 2-4 times that of Summit's performance. Preparations for the site, the computer build, and tests and applications are ongoing; delivery is scheduled for 2021.

Frontier's programming environment will support many of the same compilers, programming models, and tools available on Summit. ORNL, Cray, and Advanced Micro Devices (AMD) are co-designing GPU programming tools to aid in application transition from Summit to Frontier. The Cray Shasta software stack will introduce new capabilities. Furthermore, AMD has produced Heterogeneous-compute Interface for Portability (HIP) to allow users to transition from Compute Unified Device Architecture (CUDA) GPUs to Frontier's NVIDIA or AMD GPUs. Frontier will closely integrate AI with data analytics; modeling and simulation advances will drastically reduce discovery time.

The OLCF and the ECP are preparing eight and twelve applications, respectively, for use on Frontier. Applications cover a range of research domains. An application readiness workshop is scheduled for October 2019.

Discussion

Lethin requested facility power details for Frontier. **Whitt** said though the system is expected to use less than 30 megawatts (MWs) of power, a total of 40 MWs (including cooling power) will be supplied with 14 MWs repurposed from existing power.

Berzins asked about portability between Aurora and Frontier and if Kokkos, Raja, and OpenMP 5.x will be supported. **Whitt** replied that Aurora is working with HIP while Frontier is learning about SYCL. Many ECP applications are not machine specific; general portability

strategies are under development but additional work is needed. Frontier will support Kokkos, Raja, OpenMP 5.x, and OpenACC in addition to HIP for CUDA conversions.

Lethin inquired about resilience, a potential barrier to exascale computing noted in original reports. **Whitt** said resilience issues have not yet been encountered on the path to exascale. Early prototyping examined resiliency; with each generation, nodes are more instrumented so more fine-grain data is available for system management and early identification of problems.

Public Comment

None.

Reed adjourned ASCAC for the day at 4:20 p.m.

Tuesday, September 24, 2019

Exascale Update, Doug Kothe, ORNL, Lori Diachin, LLNL

The ECP has the necessary technical components (application development, AD; software technology, ST; hardware and integration, HI) to meet national goals. The ECP has generated external (Japan, UK) and internal (National Oceanic and Atmospheric Administration, National Aeronautics and Space Administration, Department of Defense) partnerships.

The ECP application portfolio comprises 24 projects and 50+ separate codes. Despite challenges most projects are on schedule. The ECP's six co-design centers address computational motifs common to multiple applications. For example, the Co-design center for Particle Applications (CoPA) develops software and algorithms for particle methods and has been paired with Cabana, a software library, to ensure portable application performance. Exascale ML has been applied to materials inverse problems and light source workflows.

ECP ST builds a comprehensive software stack that enables application developers to write highly parallel applications for diverse architectures. Approximately 70 software products have been grouped into various software development kits (SDKs) and stacked to generate Extreme-Scale Scientific Software Stacks (E4S) to enhance product interoperability.

ECP HI unites applications, software, and hardware within DOE facilities. Six level 3 (L3) technical HI projects (PathForward; Hardware Evaluation; Application Integration; Software Deployment; Facility Resource Utilization; and Training and Productivity) are ongoing.

ECP completion is scheduled for 2023. The ECP final design, consisting of three components (project structure; technical plans; management processes) was presented at the final design review in June 2019 at ANL. Based on feedback, three (application performance; application capability; software technology integration) of the four key performance parameters (KPPs) addressing AD, ST co-design centers, and HI have been revised to include stretch goals. Early access projects will help ensure that ECP subprojects have sufficient resources to meet their KPP and figure of merit targets. The ECP is actively managing several dependencies in real time within projects and the DOE facilities with support from an agile project management database and tool.

Discussion

Dolbow asked if the final design report is available. **Diachin** replied that a final design report draft will be accessible in two weeks following completion of AD and ST reviews.

Levermore stated that articulating ECP success in delivering exascale-capable applications and software that meet KPPs to funders and broader audiences is important. **Kothe** said that ECP success will be measured by the ability of applications to solve DOE problems. **Dunning** commented that application reports, in the final design report, offer a basis for connecting problems addressed by applications with the DOE's scientific mission.

Berzins said the KPPs demonstrate that the expected application performance has been achieved. Regarding future KPPs, he asked if performance issues and optimal architectures can be characterized for current codes. There are opportunities to investigate algorithms and architectures for mission-oriented application performance from a post-exascale perspective. **Kothe** replied the ECP utilizes hardware-driven algorithm design but co-design centers and proxy applications will allow a shift towards algorithm-driven hardware design. **Diachin** agreed that these concerns are important for the long-term sustainability of operations and research.

Reed asked how to shape hardware that is well-matched to computational science algorithms given that DOE influence with hardware vendors is waning at the ECP scale. **Kothe** said completion of the ECP project will help determine which hardware is best-suited for algorithms. Future convergence of AI and science will aid in algorithm-driven hardware design.

Ang asked about vendor engagement recommendations because PathForward is ending. AI and ML might attract vendors and create room for innovation. Developing algorithm-driven hardware at large scale may be impractical because of the potential cost of failure; hardware co-design at a smaller scale offers solutions. Defense Advanced Research Projects Agency (DARPA), for example, is investing in electronic computer-aided design (ECAD) tools that will accelerate hardware design. **Diachin** replied the Recovery Time Objective document recommends continuing similar investments for PathForward. AI and ML can serve as inner loop processes. Increased flexibility in network and chip design offers an avenue for smaller-scale explorations as a precursor to DOE-scale mission investigations. **Reed** commented that a system for generating small numbers of many different kinds of chips is lacking.

Diversity and Inclusion in the Office of Science, Julie Carruthers, Office of Science

The Office of the Deputy Director for Science Programs is leading SC efforts in diversity, equity and inclusion (DEI). The SC implemented a new annual process for National Laboratory DEI oversight in 2016. Laboratories are required to communicate their DEI strategies to the SC and the SC to review and provide feedback on these strategies. Laboratories are also required to publically post and annually update their workforce demographic data on their websites. Laboratory demographics show that 30% of employees are women, 19% under-represented minorities, and 10% other people of color. All 17 of the DOE National Laboratory Directors are actively engaged in advancing DEI and the National Laboratory Director's Council (NLDC) has held annual workshops for the past four years to share promising practices. Notably, laboratory-wide culture and climate surveys, followed by periodic pulse surveys are an effective way to identify and assess progress on DEI challenges. The SC will evaluate the efficacy of this oversight process through an external peer review in the fall of 2019.

The SC established an internal working group to evaluate DEI in SC business practices concerning processes and procedures for research awards to universities and the DOE labs. The working group is compiling recommendations and a report is due to the Deputy Director and

Associate Directors this fall. Beginning in December 2019, selected recommendations will be implemented using a transparent communication strategy.

The SC has published a website consolidating DOE policies and procedures as they apply to recipients of financial assistance, as well as the SC's statement of commitment to DEI.

The SC is coordinating DEI efforts with the DOE's Office of Civil Rights and Diversity (OCRD). OCRD administers DOE policies, practices, and procedures related to internal (federal employees) and external (DOE financial assistance recipients) civil rights. For external parties, ORCD conducts pre-award assurance reviews to insure award recipients meet Title IX institutional requirements as well as post-award compliance reviews. The DOE and ORCD have conducted more compliance reviews than any other science agency. ORCD also coordinates Title IX oversight with the Department of Education and the Department of Justice. Further interagency committees and working groups addressing DEI topics have been established with the Office of Science and Technology Policy (OSTP) and the National Science Talent Contest (NSTC) Committee on STEM Education.

Discussion

Berzins asked if the national laboratories post demographic histories online. **Carruthers** said cumulative records are not posted and agreed gauging trends would be informative.

Lethin commented that most SC DEI programs appear to focus on gender diversity and asked if there is a program centered on racial diversity. He also asked if there are active programs such as mentorship to prevent discrimination. **Carruthers** replied the recruitment process ensures that women and minorities are represented in applicant pools. The current review phase has not yet addressed how the SC can better communicate research opportunities to underrepresented minorities.

Dolbow inquired about the data tracking diversity hires' likelihood of accepting a position and their retention. **Carruthers** replied that the national laboratories are tracking hiring information, but data are not currently available. Laboratory exit surveys may not capture cultural reasons for employees' departures. The upcoming peer review will provide additional oversight guidance.

Crivelli asked about the level of national laboratory technical staff diversity. **Carruthers** responded that the NLDC website tracks technical staff demographics; diversity is lowest in this area across the national laboratories. **Dunning** stated more fine-grained demographics are needed.

Valerie Taylor, ANL asked about demographic data addressing progression rates, such as promotion and leadership opportunities. **Carruthers** answered that few laboratories track this information rigorously.

Reed asked about DOE policies requiring reporting of sexual harassment. **Carruthers** replied the DOE and other agencies met with the NSF before the NSF posted a Term and Condition addressing sexual harassment. A decision to enact a similar DOE Term and Condition has not been reached; implementation would require a federal rule-making process. Currently, the DOE requires notification if any employee is placed on administrative leave for any reasons, if it exceeds a certain number of months.

Hendrickson asked about the DEI status of the research community beyond the SC laboratories. **Carruthers** answered that the Office of Nuclear Energy and the Office of Energy Efficiency and Renewable Energy have asked their laboratories to develop DEI strategies using

SC guidance as a template. NNSA laboratories are also examining this topic because of a recent incident at the Nevada test site.

Negele shared that four women have recently accepted university positions in the field of computational science. These individuals require access to computational resources to achieve tenure. However, changes in available exascale supercomputers and funding may make supporting these women difficult. Providing resources to these role models is important.

Giles asked if climate survey data are aggregated across the national laboratories. **Carruthers** replied aggregate results are not available because laboratories voluntarily share survey results. The Office of Basic Energy Sciences is funding Geraldine Richmond to interview scientists at leadership levels across the DOE complex about concerns regarding DEI to identify common issues. Results are not yet available.

Lethin noted a previous National Academy of Sciences review recommended conducting climate surveys. **Carruthers** commented that though the review was for academia, the laboratories recognize the value of culture climate surveys. However, experts are needed to interpret survey data and provide recommendations.

Reed called for a break at 10:37 a.m. and reconvened the meeting at 10:55 a.m.

Update on Mathematical Multifaceted Integrated Capability Centers (MMICC), Bill Spotz, ASCR

Following an ASCR workshop and report in 2011 and an Applied Math Summit in 2012, a new program, Mathematical Multifaceted Integrated Capability Centers (MMICCs) was created; a solicitation was released 2012. Among other criteria, successful MMICCs applicants were evaluated based on their ability to apply cross-cutting math to DOE mission areas and identify research challenges that represent abstractions of DOE grand challenges. Centers will ideally transition from MMICCs to other funding sources such as Scientific Discovery through Advanced Computing (SciDAC) partnerships or institutes, co-design centers, or other DOE sources.

The first MMICCs solicitation awarded a total of $9M per year to three centers: Multifaceted Mathematics for Complex Energy Systems (M2ACS); Collaboratory of Mathematics for Mesoscopic Modeling of Materials (CM4); and An Integrated Multifacted Approach to Mathematics at the Interfaces of Data, Models and Decisions (DiaMonD). The current MMICCs portfolio is funded at $8M annually and comprises Multifacted Mathematics for Rare, High Impact Events in Complex Energy and Environment Systems (MACSER); Advances in Experimental Design, Optimization & Learning for Uncertain Complex Systems (AEOLUS); and Collaboratory on Mathematics and Physics-Informed Learning Machines for Multiscale and Multiphysics Problems (PhILMs).

MACSER quantifies the occurrence of rare high-impact events and designs and optimizes energy systems to withstand and recover from them. MACSER also develops novel algorithms and software and participates in extensive outreach to the mathematical and domain science communities (e.g. Institute for Computational and Experimental Research in Mathematics).

AEOLUS is developing a unified optimization-under-uncertainty approach to learning predictive models from data and optimizing design and control under uncertainty. This approach is applied to complex multiscale systems in advanced materials and manufacturing. AEOLUS exploits problem structure (geometry, sparsity, low-dimensionality) to achieve principled, rigorous, and scalable exploration of parameter and decision spaces.

PhiLMs develops physics-based and data-driven tools and approaches including non-local operators, multifidelity data and information fusion, deep neural networks, meshless methods, uncertainty propagation, and stochasticity. Ultimately, PhILMs aims to be established as a new DOE center at the interface of mathematics, physics, data science, and deep learning.

Spotz summarized accomplishments of and future work for MACSER, AEOLUS and PhILMs. MMICCs reviews are scheduled in October and November of 2019.

Discussion

Levermore commented that AEOLUS and PhiLMs have integrated ML and AI into prior programs and are leveraging NSF partnerships to bring new ideas into mainstream research. **Spotz** replied he has been involved in scientific machine learning reviews; these required branching into new areas to find reviewers representing this new research dynamic.

Berzins asked how MMICCs are meeting the aim of connecting with SciDAC partnerships and other opportunities. **Spotz** answered that future review processes will require MMICCs to report how they will maintain funding after MMICC support ends.

Lethin commented that mathematics are vital to performance, perhaps even more so than hardware in terms of achieving new scales. **Spotz** agreed, noting AEOLUS and PhILMs are focused on the outer loop around the forward solve; large computations require finding ways to reduce the flock count.

Reed asked about funding MMICCs in a post-ECP era. **Spotz** anticipates program reviews with ECP completion. He will bring an open mind to the process and the end result should be fostering the best applied mathematics program possible.

Public Comment

None.

Reed adjourned ASCAC at 11:25 a.m.

Respectfully submitted,
T. Reneau Conner, PhD, PMP, AHIP & Holly Holt, PhD
ORISE/ ORAU