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

MEETING MINUTES

April 23-24, 2020

Teleconference

ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE

The U.S. Department of Energy (DOE) Advanced Scientific Computing Advisory Committee (ASCAC) convened virtually on Thursday and Friday, April 23-24, 2020 via Zoom. 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

Daniel Reed (Chairperson)
Keren Bergman
Martin Berzins
Vinton Cerf
Barbara Chapman
Jacqueline Chen
Silvia Crivelli
John Dolbow
Jack Dongarra
Thom Dunning
Tim Germann
Susan Gregurick
Anthony Hey
Gwendolyn Huntoon
Sandy Landsberg
Richard Lethin
David Levermore
John Negele
Vivek Sarkar
Krysta Svore

ASCAC Members Absent

Satoshi Matsouka

Also Participating

Steve Binkley, Deputy Director for Science Programs, Office of Science (SC), Department of Energy (DOE)
Barbara Helland, Associate Director, Advanced Scientific Computing Research (ASCR), DOE
Christine Chalk, ASCAC Designated Federal Officer, Program Manager, Oak Ridge Leadership Computing (OLCF), ASCR, DOE
Roscoe Giles, Boston University
Bruce Hendrickson, Lawrence Livermore National Laboratory (LLNL)
Kerstin Kleese Van Dam, Brookhaven National Laboratory (BNL)
Doug Kothe, Director, ECP, Oak Ridge National Laboratory (ORNL)
Andrew Siegel, ORNL
Tzanio Kolev, ORNL
Rick Stevens, Argonne Leadership Computing Facility (ALCF)
Julia Phillips, National Science Foundation (NSF)
Valerie Taylor, Argonne National Laboratory (ANL)
Gina Tourassi, OLCF

Attending

There were approximately 131 attendees in the teleconference.

Thursday, April 23, 2020

Reed, ASCAC Chair, called the ASCAC meeting to order and share opening remarks.

VIEW FROM WASHINGTON, Steve Binkley, Principal Deputy Director of the Office of Science

Binkley provided an update on the FY20 enacted appropriation, the President's FY21 Budget Request, and the Office of Science (SC) reorganization.

SC's FY21 request is $5.838B. The FY20 enacted budget is $7B – an increase from FY19 ($6.585B). Budget allocations are research (42%), facility operations (40%), and projects/ infrastructure (18%). Of the research portion, 38% goes to universities (single investigators).

SC will continue investments in exascale computing, artificial intelligence/machine learning (AI/ML), quantum information science (QIS), microelectronics, DOE's Isotope Initiative, biosecurity, and the U.S. Fusion Program acceleration.

FY21's new research initiatives include integrated computational and data infrastructure for scientific discovery, next generation biology, rare earth/separation science, polymer upcycling, accelerator technology, and data and computational collaboration with the National Institutes of Health (NIH).

The Office of Science reorganization effects the top levels of the organization. Dr. Binkley will assume the new position of Principal Deputy Director (SC-2). Dr. Harriet Kung has been appointed as the Deputy Director for Science Programs (SC-3). Juston Fontaine will be the Deputy Director for Field Operations (SC-4).

Discussion

Dongarra asked about the SBIR/STTR budgets for FY20 and FY21. **Binkley** explained the budgets for these programs are based on a percentage of the research budget. SBIR is ~3.2% and STTR is ~0.45%. This is a requirement in the statute (15 U.S. Code §638). The budget numbers have not yet been analyzed for FY20 and FY21.

Lethin requested more detail on the status and form of the fusion research commercialization efforts. **Binkley** said not all the details for the "SpaceX-like" portion have been finalized. INFUSE (Innovation Network for Fusion Energy) is funded in the FY20 budget and is based on a similar program in nuclear energy. INFUSE allows commercial companies to request assistance from DOE for resources and capabilities to further develop their fusion ideas. Thus far, the program has received favorable responses from the Fusion Industry Association.

Landsberg asked how the new initiatives in FY21 will factor into a reduced budget. **Binkley** said the FY21 budget submission contains a table that cross-walks those initiatives. If the budgets are lower, SC will have to make difficult decisions.

Levermore noted that both the "SpaceX-like" program and the nuclear energy program demonstrated basic technology proof of principle before they attracted private company investment and asked about the rationale related to the Fusion Energy program. **Binkley** explained that the fusion landscape has changed dramatically in last few years. For example, tokamak size is a function of magnet strength and there have been significant advancements in high field magnets. MIT plans to explore that part of the phase space and expects to be able to demonstrate that technology within 7-8 years.

Hey expressed concern about the budget for research being dramatically reduced due to the present extraordinary times. **Binkley** said the appropriations' response to the COVID issues and the stimulus for small business do not reflect budget reductions. However, the amassing U.S. debt could lead to a reduction in the SC budget. **Reed** added that this is a global framing issue about the U.S. position relative to the rest of the world.

VIEW FROM GERMANTOWN, Barbara Helland, Associate Director of the Office of Science for Advanced Scientific Computing Research

Helland discussed COVID-19 efforts, ESnet, ASCR highlights, the FY20 budget, announcements and workshops. The COVID-19 High Performance Computing (HPC) Consortium is a public/private effort that brings together government, industry, and academic leaders who are volunteering free compute time and resources of world-class high performance computing in support of COVID-19 research.

ESnet's iperf3 tool has measured the change in internet traffic during the COVID-19 pandemic. The iperf3 tool is available as part of perfSONAR or a stand-alone tool; the largest known iperf3 deployment is Comcast.

Two new FY21 initiatives are the Integrated Computational and Data Infrastructure for Scientific Discovery and the Data and Computational Collaboration with NIH. ASCR contributes to SC's priorities in AI/ML ($56M), QIS ($86.2M), and Strategic Computing ($428.9M).

Funding opportunities discussed include SciDAC Institutes, Scientific ML for Model and Simulation, AI and Decision Support for Complex Systems, Fair Data and Models for AI/ML, and the National QIS Research Centers. Accomplishments highlighted were the hybrid Quantum Classical Approach to Chemistry Calculations, FASTmath ML for Neutron Science, and the DOE Applied Math Codes for the Nuclear Regulatory Commission licensing process for nuclear technologies.

The 5G-enabled energy innovation workshop was held in March 2020. Draft priority research directions include wireless communication in extreme environments, scientific instrumentation and critical national infrastructure, the digital continuum, AI-enabled edge computing, and innovation through community testbeds. Upcoming workshops are the Future Scientific Methodologies workshop (August 4-6, 2020), a curated "unconference", and the Data Reduction workshop (October 2020).

The final configuration for Perlmutter will include the AMD Milan chip, GPU cabinets (late 2020), and CPU cabinets (2021) connected with the Cray slingshot high performance network. Helland closed with ASCR staff changes (retirements, job changes, and new positions). Laura Biven was recognized as a graduate of the first cohort from the DOE Leadership Development Program.

Discussion

Cerf inquired if there is a voluntary agreement to share information among the COVID-19 Collaboration parties. **Helland** stated that everything has to be available in a public venue. A small effort to match partnerships with existing research exists with the European Union as part of the NSF Exceed program. **Cerf** described the quantum internet as a collection of distributed quantum computers that optically moves entangled photons to other locations making use of fact that entanglement is distance independent. Delivery of an entangled photon over an optical channel or free space laser is a challenge with respect to maintaining the quantum state. **Helland** indicated there will be a number of years of research to develop the repeaters and routers that can go a further distance. Ideally, this will begin with connecting the QIS centers with a computational internet. In partnership with NASA (National Aeronautic and Space Administration) it may have to begin with satellite communications. **Cerf** asked if NITRD (Networking and Information Technology Research and Development Program) is being used to coordinate quantum research among the various labs and agency programs. **Helland** said NITRD is not being used as much in quantum. A

national coordinating committee, much like in NITRD, was established by the National Quantum Initiative. The National Science and Technology Committee has a quantum subcommittee with all agencies. And DOE is establishing the National Quantum Advisory Committee that will allow broader input from the community.

REPORT FROM EXASCALE TRANSITION SUBCOMMITTEE, Roscoe Giles, Boston University

Giles provided updates to the Exascale Transition report. The changes since the January 2020 update include additional examples, elaborations, and clarifications of the text. There are four report sections: A – advance and build on the Exascale Computing Project (ECP), B – advance ASCR resources, C – current and future workforce, and D – national/ international leadership. For each section Giles discussed selected recommendations and sources of input.

Discussion

Berzins suggested two things be brought out more clearly in the report – the math research opportunities to use the abilities of modern hybrid architectures and consideration of more power efficient and less expensive architectures which includes hardware research to develop a next generation of focused architectures for quantum. **Giles** explained that the report is not trying to say that the exascale machines will solve all problems or are the end of the discussion for algorithm design and scientific workloads. **Levermore** noted that the new physics of the new architectures, as well as new AI ideas, will yield much larger wins for modeling than what will be seen from just hardware and algorithm speed-up. Algorithms, hardware, and modeling are likely to be important in varying degrees for any application. **Berzins** pointed out the report states it is essential to grow and move forward with the current code-base. But that code-base is at the leading edge and it has to be built on and improved. It is also necessary to hold on to other areas. A balance must be struck; new architecture's give new opportunities. State-of-the-art algorithms and state-of-the-art codes have a lower complexity but there is a contradiction between the proposed architectures and what the codes actually do. **Levermore** added that the committee viewed lessons-learned from the ECP ecosystem as helpful to address that contradiction.

Dunning complimented the committee on addressing the longevity of software. He mentioned that in some ways the software is more serious than the hardware because the software can change on a year to year, month to month basis. It is imperative to maintain investment in software which means ongoing expenses. Dunning encouraged ASCR and other parts of SC to think seriously about the legacy of code. The codes capture the intellectual understanding of the phenomena of interest; it is extremely valuable, it needs to be maintained, and it needs investments like ECP that adapt to a truly new type of computing architecture. **Giles** said the report includes a recommendation about distributing application or experimental software. There is a point where productization or hardening of software for use with a broader audience is desirable. The report focused on making it usable by others which requires ongoing investment and a method to accomplish that.

Helland shared that there is significant effort within SC to address many of the issues brought up in report. There are SBIR topics that allow companies to build a business model around SciDAC for example. SC's open source software is available for people to develop support mechanisms around it or to expand it.

Landsberg commented that the increasing math and computer science research budget and the solicitation aimed at the high risk high payoff aligns with a recommendation in the report.

Reed thanked everyone involved with the report.

Reed dismissed ASCAC for a break from 1:00 p.m. – 1:15 p.m.

REPORT FROM THE SUBCOMMITTEE ON 40th ANNIVERSARY ACCOMPLISHMENTS, Bruce Hendrickson, Lawrence Livermore National Laboratory

Hendrickson reminded ASCAC of the charge and provided an update on the progress of the two documents (historical and accessible). In January 2020, ASCAC members provided three elements of feedback – references, climate science, and foundational mathematics. The subcommittee created an "additional reading" appendix with pointers to key non-technical documents; OSTI (Office of Scientific and Technology Information) will archive and make these available. Jim Hack was engaged to assist with a sidebar on ASCR contributions to climate science. For foundational mathematics a vignette was added on compatible discretization. Also more clarifying text was developed observing that the time to impact for foundational work in mathematics is long. Recognition of recent work in this area will require time.

The accessible, impact-centric document, is structured around a few exemplar impact stories that are 3-4 pages each. The articles are written by professional technical writers and shepherded by a member of the committee. All of the articles have been written and the document is in the layout process with the final version printed by June 1. The historical document is in the final edit and layout stages.

Discussion

Reed thanked Hendrickson and complimented the response to earlier discussions with ASCAC. **Berzins** stated the report is wonderful and the committee has gone beyond expectations. **Hendrickson** said that this was a team effort by individuals who wanted to convey what we all believe is an important part of the scientific timeline and what we want to share with the world. **Landsberg** commended the subcommittee for its work and said the glossy document is exceptional. **Levermore** appreciated the way the reference problem was addressed saying it was brilliant and serves the purpose that was needed.

Helland echoed the congratulations. ASCR has needed such a report for a long time. Helland thanked the subcommittee for its hard work.

Reed stated that the histories highlight a dynamic community of people who have contributed over a very long time. Capturing that history reminds people that our current position relies on similar people.

Hey appreciated all that has gone into the report. **Levermore** indicated the glossy report will be more impactful initially. **Hendrickson** said the hope is the glossy version will be useful for the short-term but the history document will have a long tail. Capturing this for posterity is a service to the community that hopefully will still be of use years from now. **Reed** said capturing the important developments that occurred in the past will be important for the younger generation of scientists. **Chalk** reminded ASCAC that they voted to accept the report in January 2020 and that this is the final version with suggestions addressed.

COVID-19 EFFORTS AT OLCF, Gina Tourassi, Oak Ridge National Laboratory

Supercomputing use at OLCF for COVID-19 research includes bioinformatics, drug discovery, diagnosis, prognosis, and epidemiology, the majority of which has been in informatics. Supercomputers are being used in non-traditional ways such as in diagnosis and prognosis (using

medical imaging and clinical records' datasets) and in epidemiology (for virus spread, conditions on the ground that impact the spread, and virus evolution in the population).

The COVID-19 HPC Consortium consists of members from industry, national labs, and academic and federal agencies. Proposals submitted to OLCF for computer time come through the Director's Discretionary (DD) program or through the Consortium. The typical award is 50K – 100K Summit node hours. To date, OLCF has allocated 1.2M node hours for eight projects.

Tourassi highlighted six projects – Rapid Antiviral Drug Discovery for SARS-CoV-2 (ChemRxiv), Discovering Molecular Mechanisms of the Human Coronavirus (BioARxiv), and Systems Biology of COVID-19 on Summit (BioARxiv). Both Using MD and QM/MM to Improve Drug Candidates for nCoV-19 Targets and AI-driven Integrative Biology for Acceleration of Therapeutic Discovery against SARS-CoV-2 are in progress. AI Text Mining of Coronavirus Scientific Publications is in development.

Summit is adding 54 "larger memory" nodes allowing COVID-19 jobs that need larger on-node memory to run on up to 54 nodes. OLCF is preparing to handle sensitive data according to the national data protection standards. The data is being used to predict adverse events and effectiveness and safety of treatments, and to model efficient and effective use of medications, mitigation of supply shortages, and optimization of resource allocation.

Discussion

Crivelli asked if the codes and data results will be publicly available, and when Summit will be accessible for running models for sensitive data. **Tourassi** said dissemination of the output is a requirement from the Consortium. Publications will be openly available, but it is unclear when the results will be available. Information on OLCF supported efforts can be found on the Consortium and OLCF websites. The protocol for Summit availability was reviewed and given the green light several months ago. The hope is this will go live in two months.

Gregurick inquired how researchers will access the Veterans Administration data assets and about plans to extend the HPC Consortium beyond September. **Tourassi** clarified that OLCF is providing compute cycles rather than access to data. In terms of the long-run, OLCF is providing cycles through the DD. **Helland** added that the executive board will review the Consortium every 90 days. ASCR is committed to the long-haul.

Cerf asked if documentation for the encryption/decryption is available. **Tourassi** explained that documentation needs to go through review and are not yet publicly available.

Madduri asked about use cases and protocol. **Tourassi** said the first use case will be with the National Cancer Institute (NCI) data and we will be ensuring all computations are secure. The desire is to have applications that enable comparison of the computational results. The first use case is enabled by the DOE/NCI partnership and range of models that involve AI. But it is not the only use case.

SCIENCE AND ENGINEERING INDICATORS, Julia Phillips, National Science Foundation

Phillips reviewed the 2020 Science and Engineering Indicators. Since World War II, advancements in science and technology have driven over 85% of the U.S. economic growth. U.S. preeminence in science and engineering (S&E) has been sustained through bipartisan investment in fundamental research. Future U.S. preeminence is not assured but the U.S. is well-positioned to compete, collaborate, and thrive.

Phillips focused on investment, determining output (quantity, quality), collaboration, and workforce and students (foreign, domestic).

Since 2000, global investment in R&D has tripled. While the U.S. remains a leading player other countries are following our example. In 2017, the U.S. spent $548B on R&D and continues to spend more on fundamental research than other country. Since 2000, China has accounted for 1/3 of the total global growth while the U.S. global share of R&D has declined from 37% to 25% The U.S. is the largest producer of output in high R&D intensive industries.

The production of new knowledge from R&D can be measured by peer-reviewed publications. The EU, China, U.S., India, Japan, and South Korea, together produce more than 70% of the worldwide refereed S&E publications. China's S&E publication output has risen 10-fold since 2000. The U.S. still has largest share of highly cited publications.

The U.S. has seen dramatic growth in international collaboration, from 19% (2000) to 39% (2018). U.S. researchers collaborate most with Chinese researchers. International student enrollment in U.S. higher education has declined since 2016. The U.S. remains the destination for internationally mobile students worldwide (19%). Asian countries (China, India, and South Korea) are the largest source countries and accounted for >50% of all international recipients of U.S. S&E research doctoral degrees since 2000. In critical fields, nearly 60% of the U.S. workforce are foreign born. Since 2000, the U.S. share of S&E bachelor's degrees awarded annually to Hispanic students nearly doubled while the share awarded to black students remained flat. Numbers of women and minorities in the S&E workforce has increased. Since 1993, the number of under-represented minorities has increased nearly four-fold. However, the changes of women and minorities varies significantly by field.

There is more competition, collaboration, and knowledge production across the global S&E environment than ever before. Scientific advances come with opportunities and risks. To mitigate those risks in an increasingly competitive world, the U.S. must stay at the forefront of science and cutting-edge research. Public funding of fundamental research is a sustained commitment over a long period of time and it has been a competitive advantage for the U.S. for decades. To remain competitive the U.S. must adapt more quickly through partnerships and collaborations, reaffirm our values, give Americans the STEM knowledge and skills they need to thrive, and ensure we have the infrastructure and resources to provide a welcoming home for the world's best talent and ideas.

Discussion

Cerf asked if the R&D expenditures shown were in current or constant dollars. **Phillips** believed the chart was showing constant dollars.

Sarkar asked who is fostering strategic discussions between government, academia, and industry. **Phillips** stated that the President's Council of Advisors on Science and Technology (PCAST) is working on ideas to encourage interactions to spur the translation of discoveries into innovations and into the private sector. The national labs have technology transfer activities. NSF also has various programs to encourage those discussions and the National Science Board is working on a vision for 2030.

Antypas inquired if the publications and collaborations with Chinese authors at U.S. institutions was distinguished from those at Chinese institutions. **Phillips** indicated that those data were looking at the institutional affiliations listed on the publication.

Bergman asked about acceleration of the trends beyond 2018. **Phillips** said while the National Center for Science and Engineering Statistics data is exceptionally trustworthy, it is lagging. She projected that as more data become available there might be interim publications, but those are not available at this time.

PUBLIC COMMENT

None.

Reed dismissed ASCAC for the day at 3:00 p.m.

FRIDAY, April 24

Reed, ASCAC, called the meeting to order and introduced Rick Stevens.

COVID-19 RESEARCH, Rick Stevens, Argonne National Laboratory

Stevens explained what is known about SARS-CoV-2, how it replicates, and ALCF efforts in COVID-19 research. SARS-CoV-2 is 50% similar to the common cold and 80% similar to SARS-1. The ALCF research effort focuses on three computational research problems in antiviral drug screening, epidemiology, and evolution.

The goal of the antiviral effort is to discover small molecules that will inhibit the viral replication cycle. Drug targets focus on slowing entry into the cell (interrupt recognition process of the virus and the host receptor) and disrupting the protease activity (interrupt the replication cycle). ALCF is investigating drugs for the host protein rather than the viral protein. Host proteins are not evolving as opposed to a virus which could evolve a defense against any drug. ALCF has built a workflow using physics-based and AI-base models that can search ~ 4B molecules overnight by running drug docking programs on a library of ~ 1M molecules. ALCF also has AI-based models that can generate molecules around a chemical structure space and then test them in the docking. ALCF is trying to use computation as much as possible to reduce the space of molecules for wet labs to look at.

ALCF has been gathering molecules across 21 databases. The current database that drives a finger-printed based or an image-based search is ~80TB of computed features. ALCF has been mining the literature for drug discovery and repurposing. There are 40K papers in the COVID-19 Open Research Dataset (CORD-19) that contains curated and tagged information from articles to use as training data for the AI systems to assess the larger collection of articles and pull out drugs and relationships that are interesting for research.

For epidemiology ALCF has an agent-based model that can handle millions of agents. This is being used to model Chicago (ANL's CityCOVID). It has a representation of 2.7M people (agents) and 1M different locations in the city. A comprehensive set of interactions and intervention strategies can be programmed into CityCOVID. A core element is the incubation process in the model that replicates what is seen in the clinic.

In terms of evolution research there are currently >10K viral sequences, and ALCF is producing phylogenetic trees daily to give to agencies such as BARDA (Biomedical Advanced Research and Development Authority) and FEMA (Federal Emergency Management Agency). Each of these sequences includes the place and date of collection. The trees are built using the whole genome sequence and snips enabling mutations to be tracked. The clusters show the effects of travel, for example, on mutation and transmission.

Discussion

Gregurick asked about optimization of AI algorithms and in the synthesis route. **Stevens** said the SAVVY database has ~250M molecules and their synthesis pathways. There are also ML models that can estimate the difficulty of synthesis. The priority is repositioning existing drugs as the primary vehicle because that is the fastest way to clinic. Off-the-shelf compounds (~800M of these) have also been prioritized.

Brown inquired if CityCOVID is being compared with other epidemiological models. **Stevens** explained ALCF is collaborating with three epidemiological groups – the 4 lab group, the Illinois Governor's Task Force, and the MITRE group. There are ~12 models being compared although there are very few agent-based models. The groups working on epidemiology models are actively comparing outputs and calibrating with each other.

Crivelli asked about docking algorithms that consider flexibility of the proteins. **Stevens** said both flexible and rigid models are being tried. There are several different docking programs (OpenEye suite, Dock6, Gold, Glide) and molecular dynamics are utilized to validate any docking high hits with fully flexible models. There is also a separate protocol for the covalent docking work.

Hey mentioned timescales in the development or repurposing of drugs. **Stevens** explained in the broader pharmaceutical pipeline this work is at the beginning. Repurposing existing drugs is faster. Clinicians in the group are willing to do off-label use because they have already been tested for toxicity, safety, etc. For the new molecule leads, this is longer term game (years).

EXASCALE UPDATE, Doug Kothe, Andrew Siegel, Tzanio Kolev, Oak Ridge National Laboratory

Andrew Siegel provided a six month update on ECP application development (AD). All of the AD projects have seen significant progress on multi-GPU nodes, co-design centers have surpassed the original vision, and there have been fundamental changes to data structures, movement, and algorithms. AD projects are exercising and providing feedback for performance portable programming models. Pre-exascale systems shifted from all CPU systems to CPUs + NVIDIA GPUs to having three types of accelerators (AMD, Intel, and NVIDIA). Use of early access hardware includes Tulip from Frontier and Iris from Aurora.

The AD groups have coordinated publication efforts through special issue journal themes. The ECP Industry Council Deep Dive was held in March 2020. The 2019 Annual report provides background on the key performance parameter (KPP) definitions, and a 5-10 page summary from each project on their experiences and lessons learned. Common themes are flat performance profiles of codes, strong scaling on GPUs, accelerator performance, programming models, new mathematical models, and software dependencies.

The next steps for the AD projects is to focus on the performance envelope, timelines and requirements for software dependencies, exascale programming models, new GPU-resident physics models, and KPP issues for initial target architectures.

Discussion

Berzins asked what proportion of the portfolio is constrained by the limited compute intensity. **Siegel** said for sparse linear algebra and Monte Carlo you do not ride the curve of peak flop performance between one machine and the other. Over 50% of the AD projects have enough computational intensity to take advantage of the increased flops.

Dunning expressed concern about people leaving the project before the ECP end date. **Siegel** noted that ECP is incentivizing teams to demonstrate their KPPs early and to define challenge problems that go beyond project completion. **Kothe** welcomed recommendations and indicated ECP would work closely with the Exascale Transition subcommittee.

Chapman commented on the importance of the training program in AD efforts and asked if there are any additional needs for the coding efforts. **Siegel** complimented the phenomenal job Ashley Barker has done running the training program. Themes for future pursuit include performance on GPUs, deeper understanding of HPC hardware to benefit the principal investigators, and ML that is different than in industry.

Hey asked about a repository for the coordinated publication and the annual publication rate on ECP. **Siegel** indicated Julia White has created the repository. **Kothe** shared that tracking will occur moving forward. **Hey** suggested archiving be done by OSTI. **Kothe** said ECP is talking to OSTI regularly.

Tzanio Kolev discussed the Center for Efficient Exascale Discretizations (CEED), one of six co-design centers in ECP. CEED is working on computational algorithms that are common motifs in many applications – Partial Differential Equations (PDE)-based simulations on unstructured grids. CEED is focused on high-order methods and spectral finite elements. High-order methods are a better match for hardware and expected to be a better match for HPC. It is known that high-order methods yield benefits on computational results and quality of simulations.

CEED is targeting ECP applications in MARBL, ExaSMR, E3SM, ExaWind, Urban, ExaAM, GEOS, and Nek5000. CEED also works with non-ECP applications. CEED uses discretization libraries to make the tools easily accessible to applications. Nek5000 and MFEM are the two large projects, both of which have been improved and are able to work on exascale systems. CEED helps applications by developing mini-apps which are used to prototype improvements, to work with vendors, and help software technology projects. Kolev shared information about three CEED products: libCEED (low-level mapping library), MFEM 4.0 (GPU support in linear algebra and finite elements), and NekRS (GPU-oriented version of Nek5000). CEED makes all of their products freely available on their website.

Discussion

Cerf asked if there is a fundamental reason why the GPU process is so successful. **Kolev** described high-order methods as introducing a dense kernel in the middle of the computation. GPUs have been optimized for simulation to dense linear algebra.

Berzins inquired about the challenges faced with the application as a whole, the levels of performance, and the challenges beyond the kernels. **Kolev** noted there are issues with the matrix (algebraic issues), with meshing, and with preconditioning. Currently, remeshing components are still performed on a CPU while error estimation and assembly are performed on the GPU. The issue is not just porting but also upgrading to high-order.

Levermore asked about the spectral element approach. **Kolev** said spectral element is a particular case of high-order methods.

Reed asked for additional comments for the Exascale Transition report and the potential for ASCAC to vote on tentative approval.

Berzins requested the math section be revised to reflect the current status of algorithms and challenges. **Lethin** agreed to connect with Berzins to address his comments.

ASCAC unanimously accepted the draft report.

Reed called for a break from 12:30 p.m. to 12:45 p.m.

REPORT FROM AI TOWN HALLS, Valerie Taylor, Argonne National Laboratory

Taylor provided information on the AI Town Halls' report which consists of nine chapters that consider the use of AI in Science in the domains as well as AI for computer science. Seven of the chapters focus on research to advance AI. Anticipated AI for Science impacts include accelerating discovery rates, semi-autonomous "self-driving" labs with active learning loops, simulations and AI hybrids, accessible and integrated knowledge bases, and comprehensive transformation of science support and operations.

The vision for AI at DOE over ten years focuses on learned models, refactored experimental discoveries, moving from semi-autonomous to fully autonomous, simulation and AI approaches mergers, theory as data for next generation AI, and AI integration into science, engineering, and operations.

AI for Science Applications are looking at AI enabled design workflows (what to make), experimental workflows (how to make it), and scientific comprehension (what it means). AI is a fundamental shift in the economic and military landscape. The Office of Artificial Intelligence and Technology includes industry, academia, and federal agencies. DOE's unique role is mission-driven development and application of AI/ML in science, energy, and national security.

Discussion

Lethin asked about DOE's vision for AI and realizing this report. **Taylor** noted that each lab and the AI subcommittee are considering these aspects. **Helland** added that integrated computational data infrastructure for scientific discovery is the first step for connecting AI and what ASCR does with the experimental user facilities. Moving forward SciDAC institutes and SciDAC calls are being discussed for 2022 to address grand challenges in the document.

Levermore inquired how SC will effectively integrate with the existing AI communities. **Taylor** said the collaborations will be important and are currently taking place between labs and universities. Labs focus on the science, while the universities focus more broadly. **Levermore** stated that pharmaceuticals have a lot of applications involved in AI. AI for Science cannot be treated like exascale where DOE has the prime view because there is a huge community exploring these ideas in ways that are relevant to DOE. **Hey** added presentations to the AI for Science subcommittee shed light on industry's different perspectives on AI/ML use. What is unique about the DOE is the facilities which produce large scale data sets.

Jim Ang (Pacific Northwest National Laboratory) recommended investigating how AI/ML is used to analyze test data and experimental measurements, as well as control to help experimental scientists work remotely. **Dunning** commented that a supplement to scientists would be AI research assistants with domain knowledge embedded in them to help computational scientists be more effective. AI research assistants can help run computational experiments and organize the output data in an understandable way. **Taylor** agreed, stating that automation will be very important.

UPDATE ON NEW CROSS-CUTTING AI SUBCOMMITTEE, Tony Hey, ASCAC

The AI for Science subcommittee initially focused on information gathering through the AI Town Halls, presentations and submissions from the Office of Science programs, Agency

presentations and submissions from NSF, NIH, as well as industrial presentations from the IT Software, IT Hardware, and non-IT companies.

Four meetings have occurred thus far, in February, March, and April. Two additional online meetings are planned for April 30 – May 1, with follow-up meetings to be determined. A large number of reports on AI have been collected from the following sources: Office of Science, National Reports, and Research Community Reports. A preliminary report is anticipated in May 2020, with the final report planned for August 2020.

Discussion

Reed agreed to circulate draft reports to ASCAC members for preliminary feedback.

Berzins asked where DOE's emphasis should be given the input received from other communities, what is distinctive in the DOE setting? **Hey's** perspective is that DOE is distinctive because it produces huge amounts of data. Domain scientists need help because their data is hundreds of terabytes. Scientists need help in finding signals, doing analysis, and such things.

WORKSHOP ON QUANTUM NETWORKING, Kerstin Kleese van Dam, Brookhaven National Laboratory

The goal of the Quantum Networking workshop was to define a roadmap leading from lab-based experiments to a nation-wide DOE quantum internet. The workshop was held in February 2020 with ~80 attendees from DOE labs, agencies, universities, and industry. Input was solicited before the workshop to address the report, identify gaps, and focus on the roadmap.

An internal DOE assessment of the state of the art in quantum networking has been drafted. The workshop brochure and the report are in draft phases with anticipated delivery in May/June 2020. The report includes four priority research opportunities (PRO) and five key roadmap milestones.

The four PRO's are to: (1) Provide foundational building blocks for a quantum internet, (2) Integrate multiple quantum networking devices, (3) Create repeating, switching, and routing for quantum entanglement, and (4) Enable error correction of quantum networking functions.

The five milestones identified are: (1) A multi-institutional ecosystem, (2) Secure quantum protocols over fiber networks, (3) Inter-campus and intra-city entanglement distribution networks, (4) Intercity quantum communication using entanglement swapping (quantum memory networks), and (5) Interstate quantum entanglement distribution using cascaded quantum repeaters (network connectivity).

Discussion

Cerf pointed out the fundamental problem of running a quantum internet keeping the quantum machines in sync. **Kleese van Dam** explained the quantum internet idea is not about connecting quantum computers rather it is about Quantum Information Science.

Hey asked if Oxford's connection of two ion-trap computers using optics is a prototype for small quantum internet device. **Kleese van Dam** indicated it was to some extent (moving information from one to the other machine). There have been more experiments like that. **Hey** added in the case of the ion-traps, it does not scale like supercomputers and qubits, rather they must be connected coherently. **Cerf** stated that taking a result from a quantum computation to conventional form and moving it to other machines can be done with conventional transmission mechanisms. The next step is to figure out how to force the state of the quantum machine to take advantage of that partial result. The quantum state itself, which is not readable out of the machine,

is what is valuable to transfer. **Kleese van Dam** explained that one idea behind distributed quantum computing via quantum networking is that if there are no large-scale quantum computers by that time it is possible to build bigger quantum computers. The question is how to do the message passing. **Cerf** added that getting scaling of quantum computation is an important target and doing that in a local architecture is the best thing to focus on right now.

Svore suggested pairing the milestones with a set of applications, workloads, and specifications. The requirements of validation and verification, replication, capacity, and the like must be clearly outlined. Additionally, those requirements and milestones on a quantum internet should be aligned to the roadmaps across quantum computing. **Kleese van Dam** said that such discussions were had but there was no conclusion. When it comes to metrics a large community effort beyond a 1.5 day workshop is needed; the metrics must be built before the devices.

Hey asked about the international situation. **Kleese van Dam** noted that Europe and China have led the way. China developed the first long-distance network with trusted nodes, and have completed the first satellite link. Europe has been setting standards and defining the first blueprint of what a quantum network could look like. The U.S., with the help of DOE/NSF funding, has made great strides in last few years and is on the cusp to take leadership in some areas. ANL has a 50-mile loop testbed which will soon become an 80 mile loop. BNL has an 80 mile loop now and is working on entangling quantum memories over the normal communication fiber – those are things no one else has done at that scale.

PUBLIC COMMENT

None.

Reed adjourned the meeting at 2:50 p.m. ET.

Respectfully submitted,
Tiffani R. Conner, PhD, PMP, AHIP
ORISE/ORAU