Meeting Minutes
ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE (ASCAC)
April 17-18, 2018
Sheraton Hotel Pentagon City – 900 South Orme St., Arlington, VA 22204

ASCAC Members Present

Keren Bergman (online)
Martin Berzins (online)
Jacqueline Chen
Silvia Crivelli
John Dolbow (online)
Jack Dongarra (online)
Thom Dunning
Tim Germann
Anthony Hey (online)
Gwendolyn Huntoon
Richard Lethin (online)
David Levermore
John Negele (online)
Daniel Reed (Chairperson)
Vivek Sarkar (online)

ASCAC Members Absent

Vinton Cerf
Barbara Chapman
Susan Gregurick
Satoshi Matsouka
Linda Petzold
Krysta Svore
Dean Williams

Also Participating

J. Stephen Binkley, Principle Deputy Director for Programs, Office of Science (SC), Department of Energy (DOE)
Rich Carlson, Advanced Scientific Computing Research (ASCR)
Christine Chalk, ASCAC Designated Federal Officer, Program Manager, Oak Ridge Leadership Computing Facility (OLCF), ASCR, SC, DOE
T. Reneau Conner, Oak Ridge Institute for Science and Energy (ORISE), Oak Ridge Associated Universities
Paul Dabbar, Under Secretary of Energy for Science
Cory Hauck, Oak Ridge National Laboratory (ORNL)
Barbara Helland, Associate Director, ASCR, SC, DOE
David Kahaner, Asian Technology Information Program (ATIP)
Steve Lee, ASCR
Joe Lykken, Fermi National Accelerator Laboratory (Fermilab)
Paul Messina, Argonne National Laboratory (ANL)
Lucy Nowell, ASCR
John Sarrao, Los Alamos National Laboratory (LANL)

Attending

James Ang, Pacific Northwest National Laboratory (PNNL)
Nathan Baker, LBNL
Paul Bayer, DOE, Biological and Environmental Research (BER)
Jim Belak, Lawrence Livermore National Laboratory (LLNL)
Laura Biven, SC, DOE
Buddy Bland, ORNL
Bill Brantley
Ben Brown, SC, DOE
David Brown, LBNL
Mark Byrd, Southern Illinois University
Julie Carruthers, SC, DOE
Chris Chang
Phillip Collela, Lawrence Berkeley National Laboratory (LBNL)
Claire Cramer, SC, DOE
T.L. Cubbage, DOE
Lori Diachen, LLNL
Steve Elbert, PNNL
Kristen Ellis, DOE
Miriam Elsayed, SC, DOE

Farida Fahrod, Air Force Office of Scientific Research
Al Geist, ORNL
Shaun Gleason, ORNL
Mark Guston, Cray
Jim Hack, ORNL
Bruce Harmon
Bruce Hendrickson, LLNL
Graham Heyes, Jefferson Lab (JLab)
Jeff Hittinger, LLNL
Mike Ignatowski, LLNL
Ben Kalber, Lewis-Burke
Miriam Kiran, LBNL
Kerstin Kleese van Dam, Brookhaven National Laboratory (BNL)
Doug Kothe, ORNL
Carolyn Lauzon, SC, DOE
Ted Lavine, SC, DOE
Alexander Lazelere, US Council on Competitiveness (USCOC)
Ji Lee
Arthur Maccabe, ORNL
Michael Martin, LBNL
Sonia McCarthy, SC, DOE
Sandra McLean, SC, DOE
Tameka Morgan
Jeff Nichols, ORNL
Michael Papka, ANL
D. Parks
Abani Patra, SC, DOE
Robinson Pino SC, DOE
Brian Plesser, DOE, General Counsel
Robert Roser, Fermilab
Arjun Shankar, ORNL
Andrew Sherman
Michelle Sneed, DOE
Julie Stambaugh, SC, DOE
James Stewart, Sandia National Laboratory (SNL)
Ceren Suset, SC, DOE
Peter Swart, LANL
Angie Thevenot, SC, DOE
Pieter Tseronis, DOE Retired
Jason Turner, Entanglement Institute
Jack Wells, ORNL
Justin Whitt, ORNL
Stefan Wild, ANL
David Womble, ORNL
Kathy Yelick, LBNL

Tuesday, April 17, 2018

OPENING REMARKS FROM THE COMMITTEE CHAIR

Dan Reed, ASCAC, called the meeting to order at 9:00 a.m. Eastern Time (ET) and introduced the first speaker.

VIEW FROM WASHINGTON

Paul Dabbar, Under Secretary of Energy for Science, stated his excitement for ASCR, shared his activities over the past few months, and explained his primary focus for SC and his background in nuclear and the renewable energy sector. Dabbar is currently visiting all 17 labs in the National Lab complex and focusing on environmental management (nuclear cleanup), commercialization of R&D work in all offices including SC, applied energy, and NNSA's commercial component.

The budget for ASCR in FY17 was $647M, in FY18 $722M was requested and Congress appropriated $810M, and the FY19 request is $899M. DOE leadership has been highlighting the exascale build out and quantum computing indicating that DOE leadership team is focused on the Exascale Computing Project (ECP) and accelerating funds for exascale and quantum information systems (QIS). Last week a Request for Proposal (RFP) was announced for the second and third machines at ORNL and LLNL, and an upgrade to AURORA.

Dabbar is concerned about the lack of public exposure to SC and has an objective of establishing more press coverage about accomplishments in hardware and applications across all of SC, and about the people working on new cutting edge applications. A recent example of press coverage applicable to ASCR efforts and research areas is an article on the oil and gas industry using supercomputers for modeling production.

It is important to DOE leadership to help drive industry. The budget has been increased for quantum computing and DOE leadership is focused on ways to support growth in quantum computing. Other areas of interest are semi-conductors and microelectronics, there has been and continues to be discussions on moving forward with these two areas in conjunction with other agencies and with industry. DOE is considering hosting a large summit on machine learning (ML) and artificial intelligence (AI) in Washington, DC to explore how these can be applied in science and for broader use. Dabbar shared three examples of ML for energy applications. LBNL has built a ML algorithm to process all energy data to optimize building energy management. PNNL's ML algorithm can provide predictions and guidance on electric grid management. And the Department of Defense (DOD) has a ML algorithm that can provide incoming field officers with predictions and guidance on carrying out a successful mission in an unfamiliar war zone.

Discussion

Reed asked about the most important elements necessary to facilitate quantum advancement in the US. **Dabbar** said that DOE is still considering this topic. Two potential options are to fund centers like the Energy Innovation Hubs and Energy Frontier Research Centers and another is to help focus on a testbed. Dabbar is reaching out to major universities to discuss collaborations in quantum and working with other agencies to ensure there is coordination on the focus and importance of quantum.

VIEW FROM GERMANTOWN

Barbara Helland, Associate Director, ASCR, reviewed budget information and provided updates on ASCR activities and projects. The FY19 budget request for ASCR was $899M, the FY18 enacted appropriation was $810M, 11% higher than the President's Budget Request. In the FY19 budget ASCR's direction provides $205M for ECP, $110M for Argonne Leadership Computing Facility (ALCF), $162M for Oak Ridge Leadership Computing Facility (OLCF), $94M for the National Energy Research Scientific Computing Center (NERSC) at LBNL, $10M for the Computational Science Graduate Fellowship (CSGF), and $79M for ESnet.

The ECP Independent Project Review (IPR) was held January 2018 at ORNL. Charge questions focused on addressing recommendations, documentation, milestone-based plan, baseline change, management structure, actionable engagement, and tracking to meet milestones. The review committee had several recommendations for ECP including: annual gap analysis revisions, initial level 4 milestone revisions, earliest-possible testbeds, break-out sessions for future IPRs, non-disclosure agreements, and key performance parameters. ASCR held a two-day meeting to address the IPR recommendations and establish a process for developing joint milestones.

Helland provided updates on the machines and facilities. ORNL released the CORAL-2 RFP in early April 2018 with proposals due May 24, 2018. Installation of hardware on Summit was completed in March 2018. Innovative and Novel Computational Impact on Theory and Experiment (INCITE) opened an RFP on April 16 with proposals due June 22. Three facility project reviews were completed November 2017 through January 4, 2018; four other reviews are anticipated between June 2018 and 2019, with an annual review for ECP occurring in late October 2018.

Lab reverse site visits, which inform budget and strategic plans, will take place in May 2018. The Early Career Research Program funding opportunity announcement (FOA) was released in December 2017 with final proposals due April 4, 2018. ASCR received 105 proposals; 67 from universities and 38 from labs. A decision on awards will be made by the end of May 2018. The FOA for 2018 Mathematical Multifaceted Integrated Capability Centers (MMICCs) was released April 4, 2018 and full proposals are due May 23, 2018. The expectation is to award 2-3 MMICC projects. Quantum Algorithm Teams (QAT) projects were awarded to LBNL, ORNL, and SNL. Quantum Computing Application Teams (QCAT) lab program announcement was released April 13, 2018 and proposals are due June 29, 2018. The FY18 Quantum Testbeds Pathfinder lab announcement was issued March 19 and full proposals are due May 14, 2018 with anticipated funding of $2M per year. Quantum Testbeds for Science lab announcement was issued April 6, 2018 and full proposals are due June 8, the anticipated funding is $9M per year. Some

ASCR accomplishments include the simulation of an atomic nucleus and tensor networks. And Scientific Discovery through Advanced Computing (SciDAC) accomplishments includes Global Adjoint Tomography at Scale, and Unstructured Mesh Technologies for Fusion Applications.

Discussion

Sarkar asked about the extent to which ASCR is, or should be, engaged with DARPA semi-conductor and micro-electronics initiatives. **Helland** said Robinson Pino is familiar with DARPA activities, and areas in SC have a vested interest in this. Both Basic Energy Sciences (BES) and High Energy Physics (HEP) are interested in the semi-conductor industry. SC is planning a small roundtable on micro-electronics in May with a follow-on Basic Research Needs workshop in the fall. Results are expected to be available by the December ASCAC meeting. Going forward SC has to determine its niche.

PERSPECTIVE ON QUANTUM INFORMATION SCIENCE

Joe Lykken, Fermilab, discussed quantum digital computers, quantum circuits, speed-up, and progress in building quantum computers. The power of quantum computing grows exponentially with each qubit; however analog computers are more powerful for certain applications.

Using quantum hardware to solve quantum problems is very different than building gate-based quantum computers. Superconducting Josephson Junction circuits are being studied by multiple groups. A group at Yale showed the advantages to putting Josephson Junction qubits inside superconducting microwave cavities. Fermilab is using superconducting cavities for particle accelerators made from Niobium. Challenges to using Niobium cavities for quantum computers include the number of photons (quantum computers use a single microwave photon), and temperature (accelerators operate at 2 Kelvin while quantum computers are ~20milli-Kelvin).

Fermilab is devoted to research in HEP which use computing and is interested in supercomputing qubits. Fermilab is identifying pieces of problems that can be addressed with near-term quantum technologies. Lyyken shared several examples including quantum machine learning (QML), QML with quantum annealers, QML to identify Higgs bosons, quantum simulations for gauge theories, dark matter detection (quantum sensors), cold atom interferometry, quantum entanglement, teleportation, and quantum teleportation through a wormhole.

Fermilab is engaging the SC QIS initiative in ways appropriate to their HEP mission using collaborations with partners, exploiting Fermilab expertise, leveraging resources with funding streams, keeping Fermilab aligned with HEP needs, and producing high impact science in the near-term while building capacity for HEP needs in the long-term.

Discussion

Reed asked for Lykken's perspective on accelerating collaborations across groups. **Lykken** mentioned two things that are necessary, 1) labs talking to one another and 2) avoiding duplication by coordinating activities. **Reed** stated that some companies are investing heavily in quantum and others are taking a wait and see position. **Lykken** explained that this indicates the appropriate role for labs in terms of basic research. Fermilab is working on supercomputing technology with the existing infrastructure that cost >$200M and took 10 years to build. This infrastructure allows Fermilab to learn about superconducting materials in quantum. Reticent companies will take over when practicality and acceptance are demonstrated.

Jack Nichols, ORNL, asked about companies that are moving towards logical qubits and if needing a large number of error correcting qubits arise with other quantum devices, quantum communication, and quantum sensing. **Lykken** said the need for error correcting qubits is known. By running real algorithms to solve a particular HEP problem applications that are more fault tolerant are uncovered. Benchmarking will illuminate error effects and the variances in different technologies.

Levermore asked about the connection between the need for error correction and improved coherence. **Lykken** said counting photons with enough reliability is a form of error correction. It allows one to determine if the input matches the output. That is tied to the improvements in cavity performance.

EXASCALE UPDATE

John Sarrao, Los Alamos National Laboratory, provided a status report on the ECP, an element of the Exascale Computing Initiative (ECI) landscape. ECP will deliver exascale-ready applications, create and deploy the software stack, and leverage high-performance computing (HPC) vendors' research and development activities and products. Sarrao reminded ASCAC of ECP's history from FY16 to the present. Governance is an important consideration in ECP and includes a Board of Directors and a Laboratory Operations Task Force (LOTF); advice is provided by an Industry Council and a Science Council.

The ECP IPR held in January 2018 yielded six recommendations outlined in Helland's remarks (above). ECP funding in FY18 was slightly increased; in FY19 there may be a large drop in ASCR support. ECP today is 50% larger than in FY16.

Doug Kothe's objectives for the year include execution, engagement, evaluation, and Critical Decision 2 (CD2). Sarrao shared highlights of each project area, the five clusters, and co-design. ECP risks and challenges include cost/ funding issues, meeting requirements, unclear messaging, insufficient standard programming models, PathForward funding, and the relationship between ECP and facilities. ECP is moving forward in project development, is actively managing risks, and is putting facility engagement plans into action. ECP still has to track progress, maintain agility for the highest return-on-investment on research design and development, and adjust the plan to ever-changing budget and cost profiles.

Discussion

Germann asked about the timeline and purpose of the Science Council. **Sarrao** said there is not a specific timeline. Having a Science Council will help illustrate and understand where the opportunities are, initially in applications.

Levermore outlined potential weaknesses that occur when clustering and asked if there has been any thought about how to defend against such weaknesses. **Sarrao** said that is the reason the evaluated additional strategic scope is important. **Levermore** mentioned that two things need to occur from the applications and software side 1) adaptation to the new hardware environment, and 2) create genuinely new things. **Sarrao** said that is an important and risky element of the ECP. What the project delivers the facilities have to own, run, and stand behind.

Reed asked about messaging to stakeholders given the developments in China. **Sarrao** stated sea changes must be made in national security and in science through ECI. The overall ECI needs lead the change for deep and specific mission needs.

Dunning expressed concern about the overhead costs associated with managing the teams in ECP. **Sarrao** acknowledged his concern and said that in addition one needs to strive to make the oversight as lean as possible, good project management is good science delivery.

Huntoon asked if the right people are in the pipeline to support the ECI going forward. **Sarrao** said ensuring the future workforce is appropriate is one reason for the LOTF. Many people outside of the DOE labs see ECP as a great training opportunity.

Chen amplified Dunning's point stating there is a need for more money for management to interact with pieces of ECP, additional staff to interact with the code, applications developers to interact with software teams, and the new data science. **Sarrao** agreed.

WORKSHOP ON EXTREME HETEROGENEITY

Lucy Nowell, ASCR, shared information on the Extreme Heterogeneity workshop held in January 2018. The primary focus was on the software stack and programming models, environments, and tools. The workshop attracted over 150 participants and received over 100 white papers to contribute to the Factual Status Document (FSD). The charge was to define basic research needs and opportunities in computer science research to develop smart and trainable operating and runtime systems, execution

models, and programming environments that will make future systems easier to tailor to scientists' computing needs and for facilities to securely deploy.

Nowell explained the meaning of extreme heterogeneity and shared the five PRDs developed from the workshop: programmer productivity, resource management, performance, reproducible science, and data management. For each of the five PRDs, Nowell provided additional, in-depth information.

The draft report is under review, the PRD brochure has been drafted, and lab and ASCR-supported facility visits are occurring. What is emerging from conversations with the facilities is a much richer mode of interaction which needs to continue and create a closer relationship to ensure ASCR is addressing the appropriate challenges.

The difficulty, especially from the extreme heterogeneity workshop, is getting people to think far enough out. The workshop was a grand experiment for two reasons 1) being modeled after BES Basic Research Needs workshops, and 2) being forced to run a workshop in the virtual mode on short notice. Some participants found the virtual format to be better than a face-to-face meeting, but the majority definitely said the virtual workshop was okay but face-to-face would be better. The workshop was meant to bring representatives from multiple communities together to learn how to work in concert and to share information. The personal relationship-building did not happen as well in the virtual setting.

Discussion

Reed asked how Nowell ensured thinking about radical opportunities when attendees were bound by a connection with the past. **Nowell** thought the FSD was both a friend and foe. Nowell continually referred the report writers back to the charge because they were tempted to arrange their report around the structure of the FSD. Many of the people who submitted white papers had been involved in the FSD process which may have reinforced that tendency. Nowell mentioned that asking people to review white papers over the holiday break lead to lower quality and in some cases signs of bias.

UPDATE ON HPC IN CHINA

David Kahaner, Asian Technology Information Program (ATIP) provided updates on HPC, quantum technology, and artificial intelligence (AI) activities in China. China has a top-level push for independent and controllable technology, they want to move from "made in China" to "created in China".

China has four exascale prototypes from National University of Defense Technology (NUDT), JiangNan, and Sugon all of which must use domestically developed processors and achieve greater than 512 nodes; they are funded at ~$5M from 2016 to 2018. China will begin full exascale system development in 2018, with a machine for scientific computing and potentially a machine for AI. Kahaner discussed exascale development for science and locations of the new machines including Qingdao National Lab for Marine Sciences and Technology and the China National Genome Bank. He provided information on leading Chinese system vendors and HPC manufacturing. Factors influencing Chinese software development include generous funding from the central government, growing demand, and more talented people. Barriers for software development include an immature business environment, protection of intellectual property, and lack of a reward system. The Chinese wish to dominate in HPC consumption and manufacturing, however application and system software development is weak.

Quantum technology in China is funded in three areas 1) quantum communications, 2) quantum computing, and 3) quantum sensing. There has been an increased focus in China on quantum computing since 2016. Quantum sensing utilizes overlapping technologies. The new National Lab for Quantum Information Science (NLQIS) was established in 2017. There is a fiber-based quantum communications network running from Beijing to Shanghai. China will launch the world's first satellite dedicated to quantum communications and experiments in 2016. China is realizing quantum computing in superconductors, semiconductors, ions, photonics, and more.

The Chinese government announced a national strategy for AI in July 2017. AI research is increasing and focuses on theory, algorithms, and model development. AI start-ups include SenseTime for computer vision and deep learning technologies and DJI who is a world-leader in the civilian drone market. China's AI strengths include government elevation of AI, rapid progress in HPC development, a supercomputing

center, cloud computing centers, enormous amounts of data, and easy access to people and applications. The challenges include a shortage of talent, lacking original AI research, and domestic companies pursuing near-term gains. Finally, Kahaner reviewed the China Brain Project. The technology roadmap is to repair relevant brain disorders, simulate brain-inspired computing and brain-machine intelligence technologies, and understand the brain's cognitive function.

Discussion

Messina asked about the quality of universities in China. **Kahaner** provided anecdotal information, that the good universities are able to attract good students who are willing to stay in China because both faculty and facilities are much better.

Nichols noted that Kahaner mentioned Shenwei was upgraded by a "couple of clicks" and asked for clarification. **Kahaner** stated that the chip in the Taihulight is 28nm, the exascale prototype (next-gen) will be 14nm; however, what will be in the full exascale system is unknown. SMIC is the premier fabricator in China and the government is putting money there.

Kerstin Kleese van Dam asked about the price point of Chinese commercial systems. **Kahaner** was unsure but said of the Top 500 systems 40% are coming from China.

Crivelli mentioned the China Brain Project and asked about other major developments at the intersection of AI and biology or medicine as well as efforts to engage in multi-disciplinary work. **Kahaner** said the China Brain Project is comparable to the US in terms of scope and funding size. While companies can take advantage of common resources, for example the China National Genome Bank, both Huawei and BGI are companies with their own proprietary issues making collaboration untenable.

Chen asked if applications China is investing in are open science or military applications, and if any are open source. **Kahaner** said from those who have attempted to use the codes the common agreements are that documentation is messy and work must be done directly with the framework developers in their location. It is very far from open science.

Bland commented that although 40% of the Top 500 list computers were Chinese they tend to be cloud vendors, which is not seen in the US. **Kahaner** said the Chinese break out their top 100 into two parts, cloud oriented and not.

Reed asked Kahaner for his sense of instruction-set design and the Chinese perspective on risk 5. **Kahaner** did not know but agreed to send information to Reed.

WORKSHOP ON SCIENTIFIC MACHINE LEARNING (SciML)

Steve Lee, ASCR discussed the Basic Research Needs workshop on SciML held January 30-February 1, 2018. The charge focused on status and recent trends, opportunities and barriers, grand challenges and PRDs, and basic research needs and opportunities. Lee shared three working definitions of machine learning: a field of study, a set of rules, and a process of automated learning. Deliverables from the workshop included articulation of grand challenges and PRDs.

The six PRDs include domain-aware, interpretable, and robust (foundational and crosscutting), data intensive, inner loop, and outer loop (capabilities). Key foundations for SciML PRDs include DOE Applied Math Base Program (foundational themes), DOE Applied Math Research Initiatives and Scientific Inference & Data Analysis (data-intensive SciML), Multiscale Models & Algorithms (inner-loop SciML), and Integrated Capabilities for Complex Systems (outer-loop SciML).

In summary SciML is a powerful scientific enabling technology. DOE leadership, with roots from the previous decade of Applied Math basic research, has encouraged growth and action. The future of science and energy research includes advanced technologies, greater automation and adaptivity, and cross-cutting initiatives.

Discussion

Hey asked how SciML compares with similar NSF initiatives. **Lee** stated he has to find the unique DOE purview. For example, the theme of interpretability, DARPA has its own xAI research and there are things to learn there.

Dunning stated that complex systems can be computational or experimental. He illustrated experimental system with two examples on complex light sources where 1) ML could determine when part of the system will go bad and allow for preventive maintenance, and 2) ML could provide better understanding of what controls the quality of the beam, the quality of the light source itself.

Chen mentioned the ML must be orchestrated and must work effectively with the rest of the workflow. **Lee** agreed and stated that some great science use cases came out of the workshop.

Levermore asked if there was any discussion about the interplay between SciML and exascale. **Lee** said there had not been any such discussions yet but that the capabilities connect with the applications and robustness with the vendors and computer scientists.

Crivelli commented that it is important to create teams in the context of big data. She mentioned that the Million Veteran Program (MVP) is working with heterogeneous data including genomic data, health records, and images, all of which are excellent applications of ML.

UPDATE ON SUBCOMMITTEE DOCUMENTING ASCR IMPACTS

Paul Messina, Argonne National Laboratory provided an update on activities and plans for the Subcommittee on ASCR 40th Report. The first in-person meeting was held April 16, 2018. High level objectives for the report include highlighting examples of major scientific accomplishments shaping the fields of ASCR research and lessons learned to guide future strategies.

Messina discussed the elements of accomplishments the subcommittee is seeking. He showed ASCAC the draft outline of the report, which included accomplishments divided into different research areas. He concluded by asking ASCAC to share their stories or offer suggestions to any subcommittee member.

Discussion

None.

PUBLIC COMMENT

None.

Reed adjourned the meeting at 3:25 p.m. ET.

Wednesday, April 18, 2018

OPENING REMARKS FROM THE COMMITTEE CHAIR

Dan Reed, ASCAC, called the meeting to order at 11:05 a.m. ET.

FUTURE LAB COMPUTING WORKING GROUP REPORT (FLC-WG)

Rich Carlson, ASCR discussed the FLC-WG, explained the drivers to form the group, and provided background demands on resources and emerging trends in lab computing across the 10 DOE SC labs. The FLC-WG was formed to encourage labs to utilize decades of research and experience and assist labs in developing common policies and practices, understanding the current state of the art in Distributed Computing and Data Ecosystems (DCDE), and developing a strategy for adopting a Federal DCDE model. The FLC-WG merged with the National Lab Research Computing Group (NLRCG) and defined four focus areas: economies of scale, barriers to collaboration, scientific data management, and future lab computing.

Carlson discussed the FLC-WG working group charter, goals, membership, observers, and program managers. The FLC-WG holds regular audio-visual meetings and held a face-to-face meeting at SuperComputing 2017. Carlson shared the observations of the science use cases, DOE SC facilities, and research activities. The FLC-WG draft report findings include: seamless user access, coordinated resource access and cross-facility workflows, scientific data management (movement, dissemination, and archival

storage), functional variety and portability, and organizational concerns and governance issues. FLC-WG's PathForward is to establish a pilot project that implements four items/ topics in a coherent and progressive manner. Conclusions include the SC DCDE is considered a reasonable response to the demands of scientists and limited budgets, that technical and operational expertise exists in the labs, that components and services needed already exist, and that policy and governance issues need to be resolved.

Discussion

Levermore feared there is a blind spot in the committee concerning scaling up a network and economic changes. He stated that when the economics of a network are changed, essentially making what is there more valuable and more vulnerable, off-the-shelf solutions no longer work. There is greater incentive espionage or mischievousness, evidence of this exists in the finance industry, in DOD, in the health industry, and in Google, Amazon, etc., when they went to cloud computing. He asked if there was any discussion of scale up and economic changes. **Carlson** said Vaughn Welch from Indiana, an expert in security issues, was brought in. While the FLC-WG currently focuses on identity access rather than the protection of the devices, the DCDE need to be a secure environment.

Chen commented on the need for connections to the university infrastructure. **Carlson** assured Chen that the FLC-WG recognizes that 80% of ESnet traffic goes outside the DOE complex and that a major element of the science community is at universities. These scientists must be able to access the data as well, but the FLC-WG has not yet specifically addressed universities.

Reed stated the DCDE will permanently require labs to cede control to an organization that sits above them and may actually require a lab to buy a mid-range resource to support the collaborative good as opposed to its local needs. This is a social engineering issue which is much harder than technical issues. **Carlson** said that ASCR has experience in facilitating such things and understanding how ESnet benefitted the entire lab complex. FLC-WG built a management structure and governance opportunity to make DCDE work; there is an understanding of what needs to be done and what can be tried. It is a starting point that can be pushed forward. **Helland** asked if the national labs' Chief Information Officers also started the same task. **Carlson** said yes, that their charge was issued simultaneously with ASCR's.

Dunning stressed that the way the pilot is designed and expanded is going to be critical. From the beginning the pilot must identify which labs and science communities will participate, the focus areas, and how the pilot project will expand over time. **Carlson** indicated he was encouraged because there have been volunteers for the pilot project already. He is comfortable that there will be enough motivation to get the labs engaged and that it will be completed in 2-3 years demonstrating the probability of success.

Huntoon stated that developing trust will be at the heart of this endeavor and agreed with Chen that the user community is on the campuses, which have a different environment than the labs. Some of the assumptions made in terms of how a lab-based distributed environment works are different than observations from campus-based users. In an NSF environment more data is used in smaller colleges and universities, including the community college level. **Carlson** agreed and said that he thought a DCDE could enable lone researchers, who do not have support staff, to do very productive science using DOE resources. He speculated there may be a reemergence of single investigator discoveries in the US.

Chen suggested interacting with one of the NSF-funded integration support activities to encourage the NSF single-investigator communities to adopt one method thus establishing the linkage between the DOE and NSF infrastructures.

Reed noted a difficult problem will be the forced subsidization of a heavily utilized capability for a national community. The broader that community gets the more funding models must be examined. **Carlson** conjectured that the pilot project will evolve and the people who own the facilities will develop appropriate metrics for the funding models and other policies.

Michael Martin, National Renewable Energy Laboratory, asked how a DCDE will affect the energy labs. **Carlson** indicated the energy labs have not been engaged yet, but he assumed that once something is demonstrated that engagement, with energy labs and NNSA, will begin.

EARLY CAREER: HYBRID METHODS FOR COLLISIONAL PARTICLE SYSTEMS

Cory Hauck, ORNL, discussed his project on creating hybrid, hierarchical, and multi-level algorithms for the simulation of complex many-particle systems. The challenge in simulating particle systems is one of computational complexity, a consequence of the huge number of unknowns in the system and the large variations in temporal and spatial scales over which they evolve. Hauck's project had two goals: to connect the fluid and kinetic descriptions in a single efficient method for attacking multiscale problems, and to improve the efficiency of molecular dynamics solvers using the solution of the kinetic model as a preconditioner. The effort thus far has focused on the first goal.

Hauck shared publications that demonstrate the progress being made on the project. He provided examples in simulation, modeling, and analysis. The technical outlook is that development of the hybrid kinetic strategy is progressing; three major endeavors to complete include non-linear problems, error estimation, and adaptivity. The practical challenges of impact and stability are: the workflow to reach and "end product", and Hauck's role in that workflow, and the best way to create a sustainable effort.

Discussion

Levermore asked if Hauck had thought about workflow directions. **Hauck** said that it is easy to design algorithms to address classes of problems, but it is difficult to ask someone who has a big project with hard deadlines to try an idea; there is a jump to be managed.

Chen stated that there is a lot of interest in multi-scale methods, particularly in BES and Hauck's work is at the bleeding edge. She encouraged him to connect with other ECP projects that can help implement his work at scale.

PROGRAM RESPONSE TO COV REPORT ON ASCR RESEARCH

Steve Lee, ASCR, shared the program response to the Committee of Visitors (COV) report on the ASCR Research Division and pointed out variations from previous methods for program reviews. The COV charge was to assess the efficacy and quality of the FY13-FY15 award processes, and to comment on how the award process affects portfolio elements, emerging challenges, and national and international standing. There were 22 COV recommendations grouped by ASCR into processes and awards, and portfolio and impact.

For processes and awards there were two program elements: 1) solicit, review, and document actions, and 2) monitor active project and programs. For portfolio and impact there were three program elements: 1) portfolio breadth and depth, 2) anticipate and address emerging challenges, and 3) stature in HPC and Data. Lee discussed ASCR's response for each of the five elements listed.

Discussion

Crivelli asked about the composition of the SC-wide working group to promote diversity and inclusion. **Helland** stated that Binkley's office is running the working group. All Associate Directors have volunteered one person to serve and ASCR has appointed Claire Cramer. SC is looking at diversity and inclusion to determine if there are other guidelines SC should give to review committees.

PUBLIC COMMENT

None.

Reed adjourned ASCAC at 11:21 a.m.

Respectfully submitted,
T. Reneau Conner, PhD, PMP, AHIP
ORISE
May 9, 2018