Meeting Minutes
ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE (ASCAC)
April 18-19, 2017
Crystal City Marriott - 1999 Jefferson Davis Highway, Arlington, VA 22202

ASCAC Members Present

Martin Berzins
Keren Bergman
Barbara Chapman
Silvia Crivelli
John Dolbow
Jack Dongarra
Thom Dunning
Tim Germann
Susan Gregurick
Anthony Hey
Gwendolyn Huntoon
Richard Lethin
David Levermore
Satoshi Matsouka
John Negele (telephone)
Daniel Reed (Chairperson)
Vivek Sarkar
Krysta Svore
Dean Williams

ASCAC Members Absent

Vinton Cerf
Jacqueline Chen
Juan Meza
Linda Petzold

Also Participating

Steve Binkley, Acting Director, Office of Science (SC), Department of Energy (DOE)
Ben Brown, Office of Advanced Scientific Computing Research (ASCR)
Christine Chalk, ASCAC Designated Federal Officer, Program Manager, Oak Ridge Leadership Computing, ASCR, SC, DOE
Tiffani R. Conner, Oak Ridge Institute for Science and Energy (ORISE)/ Oak Ridge Associated Universities (ORAU)
Claire Cramer, ASCR
Jack Deslippe, Lawrence Berkeley National Laboratory (LBNL)
Tara Dunderdale, Westat
Jeff Hammond, Intel
Barbara Helland, Associate Director, ASCR, SC, DOE
Paul Messina, Argonne National Laboratory (ANL)
Abani Patra, ASCR
John Shalf, LBNL
Deniese Terry, ORISE/ ORAU
John Wells, Westat
Wolfgang Windl, Ohio State University (OSU)
Shinjae Yoo, Brookhaven National Laboratory (BNL)

Attending

Mitch Ambrose, American Institute of Physics
Nathan Baker, LBNL
Laura Biven, DOE, SC
David Brown, LBNL
Lori Deachen, Lawrence Livermore National Laboratory (LLNL)
James Hack, Oak Ridge National Laboratory (ORNL)
Jeff Hittinger, LLNL
Paul Hovland, ANL
Beth Kaspar, Los Alamos National Laboratory (LANL)
Ben Kallen, Lewis-Burke Association, LLC
Alexander Lazalere, US Council on Competitiveness (USCOC)
Stephen Lee, LANL
Anthony Lewis, National Nuclear Security Administration (NNSA)
Sandra McLean, DOE, SC
Esmond Ng, LBNL
Jeff Nichols, ORNL
Lucy Nowell, DOE, SC
Karin Remington, DOE
David Riegner, OSU
Carly Robinson, DOE, SC
Michelle Schwalbe, National Academies of Science (NAS)
Jim Stewart, Sandia National Laboratory (SNL)
Louis Terminello, Pacific Northwest National Laboratory (PNNL)
Kathy Yelick, LBNL

Tuesday, April 18, 2017
Afternoon Session

OPENING REMARKS FROM THE COMMITTEE CHAIR Dan Reed, ASCAC

ASCAC Chairperson, Dan Reed called the meeting to order at 12:07 p.m. Reed reviewed the agenda with the ASCAC Committee and introduced Steve Binkley.

VIEW FROM WASHINGTON, Steve Binkley, Acting Director of the Office of Science

Binkley expressed appreciation for ASCAC members and the audience. Several DOE political appointments remain unfilled. Secretary of Energy, Rick Perry, has made clear that High Performance Computing (HPC) is important to him and the pursuit of exascale will remain a priority.

Binkley announced Dan Brouillette will be nominated as Deputy Secretary of Energy. However, there have been no specific names mentioned for the Under Secretary of Energy, Director of the Office of Science, or other political appointments.

A continuing resolution (CR) for the FY17 budget has been in place since October 2016. Congress is poised to pass an omnibus appropriation. If the omnibus does not pass, Binkley projected a CR will be in effect for the remainder of FY17.

A blueprint budget for FY18 is available, the central feature of which emphasized an increase in defense spending, requiring a decrease in other areas, including science budgets. The SC budget is slated to decrease 17% ($920M) in FY18. This decrease resets the SC budgets to 2008 levels. The details of the proposed budget are expected to be released by May 15.

There is a lot of motion in relation to the administration's statements about decreasing the workforce. SC has been operating at a lean level and Binkley is hoping SC will not experience any major impact.

Finally, Binkley stated that it is imperative to focus on exascale, quantum technologies, artificial intelligence (AI), and machine learning (ML) in the future.

Discussion

Alexander Lazalere asked about the status of the National Strategic Computing Initiative (NSCI). **Binkley** responded that the Executive Order -- Creating a National Strategic Computing Initiative (2015) is still in effect. Unless rescinded, an Executive Order remains in effect and NSCI has not been rescinded. **Lazalere** asked about the future for NSCI. **Binkley** said discussions are occurring now about the path forward for very high-end computing.

Reed thanked all in SC for their work in an uncertain time.

VIEW FROM GERMANTOWN, Barbara Helland, Associate Director of the Office of Science for ASCR

Helland covered changes in ASCR staffing. Helland is now full-time Associate Director of ASCR. William (Bill) Harrod is working on strategic computing across the federal government. Steven Lee serves as Acting Director for the Computational Science Research and Partnerships Division. Carolyn Lauzon, who is in charge of ASCR Leadership Computing Challenge (ALCC), is the Program Manager for the Cancer Initiative and is working on the Million Veterans Program. Betsy Riley will be helping Lauzon with ALCC. Sonia Sachs will lead the Argonne Leadership Computing Facilities leaving a vacancy in Computer Science (CS). Lucy Nowell was asked to step in and look at the overarching future directions for CS. Ben Brown will run ESnet, and Laura Biven will run data and visualization for ASCR. Interviews have been completed for the Director of Advanced Computing Technologies Division. Helland now must justify the position.

DOE's Scientific Discovery through Advanced Computing (SciDAC) partnerships are closed, some reviews have been completed, and reviews should be finalized in June. Status of the Basic Energy Sciences (BES) Funding Opportunity Announcement (FOA) and the SciDAC FOA is still in the Department. Helland focused on current highlights from four SciDAC Institutes: Frameworks, Algorithms, and Scalable Technologies for Mathematics (FASTMath), Quantification of Uncertainty in Extreme Scale Computations (QUEST), Scalable Data Management, Analysis, and Visualization (SDAV), and the Institute for Sustained Performance, Energy, and Resilience (SUPER).

Helland provided an update on the facilities and reminded ASCAC of the computing upgrades. The National Energy Research Scientific Computing Center's (NERSC) newest supercomputer, Cori, is on the floor and available for early science runs until June. ORNL's Summit supercomputer will start delivery of cabinets this summer and the machine should be up and running around June of 2018; Aurora will start delivery after Summit. NERSC has a request for proposal (RFP) out for machines in the 2019-2020 timeframe, and Helland has the mission needs statement for upgrades at ORNL and ANL.

The FY18 call for the Innovative and Novel Computation on Impact on Theory and Experiment (INCITE) program was opened April 17, 2017 and will close on July 23, 2017. Helland shared the INCITE award statistics for FY17. Fifty-five projects were awarded, 17 of which were renewals. The acceptance rate is 45% for non-renewals and 85% for renewals.

Helland reminded ASCAC of the exascale requirements review goals and schedule. All six reviews were completed between September 2016 and March 2017 when a cross-cut review was held. There were specific requests for a session on models, methods and algorithms. Issues with training and the workforce were expressed, especially for software developers and technicians.

Helland described a workshop report from the September 2016 National Security Agency (NSA)-DOE meeting on HPC to analyze foreign HPC and provide recommendations.

Discussion

Sarkar asked for details on the limitations in finding candidates for the Program Manager positions. **Helland** stated that currently the main limit is on approvals but it is also very difficult to find candidates. **Sarkar** offered information about channels that could be used for potential applicants. **Helland** stated they use USAJobs, HPCwire, and SIAM (for mathematics candidates) and requested information on any other, better channels to use. **Nowell** commented that they provided potential listing sites but were told they could not do anything that required a fee.

Williams asked if the DOE Computational Science Graduate Fellowship (CSGF) would be impacted with the current funding status. **Helland** said ASCR is committed to going forward at the current funding level.

UPDATE ON COMMITTEE OF VISITORS, David Levermore, ASCAC

Membership was completed for the 2017 ASCR Committee of Visitors (COV) on April 10, 2017. There are 10 people composing the broad group, all of whom have different backgrounds. They are scheduling a face-to-face meeting in July 2017 with virtual meetings for subcommittees after. The COV has a target to provide a report at the Fall 2017 ASCAC meeting.

Discussion

Hey asked about the charge for the COV. **Levermore** indicated the charge was to produce a report on the standard COV process and determine if it results in a high-quality research portfolio and to do an international comparison.

Williams asked when the COV would ask for feedback from ASCAC. **Levermore** indicated that three of the COV members are ASCAC members, but the COV will also ask for feedback from ASCAC as a whole.

REPORT FROM THE LDRD SUBCOMMITTEE, Martin Berzins, ASCAC

The Lab-Directed Research and Development (LDRD) subcommittee performed a review of the program on behalf of all the Advisory Committees in SC.

There are 1,700 projects per year, approximately 4.5% of the lab base costs go towards LDRD projects, and each project receives around $300K. Thirty percent of all post-docs across the labs are supported by LDRD projects; the bulk of the post-docs are at NNSA labs.

Berzins reviewed the process for the committee and lab visits to LBNL, LLNL, ORNL and National Renewable Energy Laboratory (NREL). The labs use a multi-layered process with feedback loops to select specific projects. The amount and type of rigor in the review is based on funding levels, but at all levels rigor is as expected for any funding agency. One challenge expressed is locking post-docs in when funding does not start until October. The schedule is off from the academic calendar.

Evaluation of the LDRD projects occurs during and after funding, at quarterly, mid-year and annual times for multi-year projects. Reporting outcomes require annual metrics (publications, presentations, intellectual property), and the LDRD long-term impact is evaluated as part of the lab strategic activities; fine scale impacts of LDRD projects are reviewed as part of employee performance reviews.

The committee reported nine observations, recommendations, and best practices for LDRD. Overall the committee found that the LDRD program provides a unique combination of high-level laboratory-driven strategic and fundamental research, it appears to be very well run and

monitored, the processes are both appropriate and necessary for the labs, and a more systematic approach to monitoring the long-term impact was suggested.

DISCUSSION OF LDRD REPORT, ASCAC

Levermore suggested having statistics on LDRD conversion to permanent staff at different DOE labs.

Crivelli asked if all labs keep track of post-docs. **Berzins** indicated this information is in the annual reports from the labs, and **Hey** stated that the number is in the LDRD report.

Matsouka asked if LDRD funds can be used to sponsor internships at the labs. **Berzins** stated he did not recall hearing about internships explicitly. **Levermore** added that graduate students go to the labs in the summer and are often mentored by LDRD scientists and that internships are funded from a different source.

Gregurick referred to unsuccessful LDRD applications asking if applicants received written feedback and the ability to reapply. **Berzins** stated there is feedback and there were no stated restrictions on resubmitting.

Sarkar referred to the percentage variability across the labs and asked for insights into what drives the money and who is in charge of making the decision. **Berzins** indicated that each lab decides the percentage relative to their strategic plans and they have a discussion with DOE concerning the appropriateness of the funding percentage. NNSA labs need to be at the congressionally allowed limits because of the type of projects they do, much of which is classified work. **Sarkar** commented that despite only three data points, the funding trend appears to be up and then down again. **Berzins** clarified that numerous legislative changes in the program have taken place over the past years that changed what was used to calculate the percentage.

Williams asked if there was any coordination between the labs at the end of LDRD projects. **Berzins** stated that the COV heard that such a thing takes place at the NNSA labs, but the COV did not hear of any equivalent procedure across SC labs. **Williams** asked why Sandia National Laboratory (SNL) has a 7% LDRD limit when others are 6%. **Berzins** reiterated that LDRD limits are lab decisions and are approved by DOE.

Dolbow asked about laboratory staff collaborations outside the labs. **Berzins** stated that while outside collaborations were mentioned, it was not the focus of the COV's discussions with the labs. **Karin Remington** mentioned that while the committee did not specifically ask about external collaborations there were numerous examples in the projects. **Berzins** reminded ASCAC that external collaborations, while implicit, are not one of the aims of LDRD.

Lethin requested any insight into the factors that led to higher research productivity of the LDRD projects. **Berzins** indicated there is not a clear understanding of why this works as it does. He offered some observations: projects of a shorter duration are targeted in a different way from a long running program. In other words, short-term science programs answer a couple of question quickly, and the influx of new energy from new people is helpful to creativity. **Dunning** suggested that the higher productivity is also due to the cutting-edge nature of the LDRD.

Levermore referred to the difference between NNSA and SC labs and asked if the committee felt LDRD should be developed more broadly or if cultural differences at some of the SC labs make LDRD unnecessary. **Berzins** said that LDRD clearly has high value and high impact and is helpful in assisting the labs look towards the future.

Crivelli commented on the exit strategy and asked if any mentoring occurs after the project is completed. **Berzins** said the COV did not hear about subsequent mentoring aspects. The LLNL exit strategy was a good way to formalize the process; writing it down is significant.

Remington added that mentoring occurred for rejected proposals and informal mentoring happened throughout the process.

Matsouka asked how the success of the LDRD is being communicated to outside peers. **Berzins** noted publications. **Hey** interjected that the question focused on how the brand name of the LDRD program was communicated. **Berzins** stated that the brand name is not necessarily coming through despite high publication rates. **Remington** indicated that the LDRD programs have been producing highlight reports about these programs and results. **Reed** said that visibility can be good or bad because this is discretionary investment. **Hey** added there may be some value for SC in knowing the program's successes because of the congressional oversight.

Dongarra stated that although a compelling case was made about the importance of LDRD, he was surprised that the report says the funding level is appropriate. **Berzins** noted there is a tricky balancing act because of how LDRD is funded. **Hey** added that an earlier draft of the COV report says "at least" at the current funding level but it is not in the current draft.

Chapman indicated she had expected the number of post-docs supported by LDRD to be higher than 30%. **Berzins** noted the percentage variability is in the yearly reports, however a comment was made that the number is going down and it is becoming more difficult to attract people. **Remington** mentioned the committee had a small sampling bias because two labs were in the San Francisco Bay Area (LBNL, LLNL) and those labs were particularly challenged. **Berzins** pointed out that the same issue existed at NREL because it was around Denver. The importance of LDRD is that it gives new people an opportunity to define a research direction and increases their awareness of what is possible.

Levermore echoed Matsouka, that more publicity is needed, and recommended that DOE put the report on the website and bring it to the attention of professional societies that advocate for ASCR.

Reed asked if the committee felt the four labs constituted a representative sample and if there were best practices that might be codified. **Berzins** said he looked at every LDRD report available from all the different labs and got a sense that these four labs were representative of the overall LDRD approach. **Williams** asked how the list of labs was selected. **Reed** indicated that the charge asked for a cross-section of the different kinds of labs, and there was a 10-year requirement of the LDRD program. **Chalk** added that labs under recompetition for Management and Operations (M&O) contracts were specifically avoided. **Reed** reemphasized that ASCAC was asked to complete the report on behalf of all.

Alan Stone, Office of High Energy Physics, asked if there was an unfair competitive advantage for the labs that invest in LDRD at a higher percentage. **Berzins** said it depends on how you define competition, but the program is adaptive and the labs can change the percentage or how they structure their programs. LDRD makes it possible to address any perceived imbalance or advantage that other labs may have by changing what they do themselves. The key strength of LDRD is that it is an adaptive program; adaptive to circumstances, changing priorities, difficulties in recruitment, and so on.

Mitch Ambrose, American Institute of Physics, asked if the COV heard about any potential downsides of having large amounts of post-docs relying on LDRD. **Berzins** stated that the COV did not hear that, rather the COV's sense was that the processes were aimed at bringing in high-flying candidates.

Reed asked ASCAC members to vote on the LDRD draft report; ASCAC was unanimously in favor of accepting the committee's report.

Reed adjourned ASCAC for a break at 3:15 p.m. and reconvened the meeting at 3:30 p.m.

UPDATE ON HPC IN ASIA, John Shalf, Lawrence Berkeley National Laboratory

From the Top 500 list, China has as many machines on the list as the U.S. Chinese colleagues said their interest lies in ancillary benefits to the aeronautic industry shifting to the basic sciences. There has been aggressive growth in domestic fabrication capability. Currently China has 28nm capacity in Shenzhen and Nanjing regions and just broke ground on 14nm fabrication to be ready by 2018 near Shanghai with another one planned in Shenzhen.

China is planning three prototype exascale systems in 2017. Shenway had a 40nm prototype in 2015 and delivered 28nm technology in 2016. Phytium has a completely domestic microarchitecture and the Instruction Set Architecture (ISA) is compatible with Advanced RISC Machine (ARM) allowing the Chinese developers to take advantage of the software ecosystem. Sugon Silicon Cube, a meteorological supercomputer for weather and climate modeling, uses a 12 core Xeon processor with DDR4.

Overall, the Chinese are moving at a fast pace and have a wide-range portfolio; their risk profile and choice of architectures is extremely broad. China has little incentive to play it safe because they have no alternative. Their focus is on creating a domestic supply chain to serve domestic industry.

In Japan there is a strategic delay of the Post-K machine because there is a process technology coming online that would dramatically improve the efficiency of the system.

Cray will deliver a 10,000+ core ARM system called Isambard to the United Kingdom (UK) Bristol. The European Union (EU) recently released a revised strategy for Horizon 2020, to "refocus in domestic technologies". The EU is interested in creating a completely, end to end, domestic solution.

Discussion

Dunning commented that China's big push right now is what they have to do to develop applications for new machines. **Shalf** mentioned that everyone he spoke to wanted to know how the U.S. develops applications.

Berzins asked if the ascension of ARM-based architectures is something that should be taken notice of. **Shalf** indicated that the cost of development is in the design and verification of circuits. There is nothing special about ARM; what is special is the economic model being used. **Berzins** continued, is it the case that there are too few swim lanes in the U.S. **Shalf** noted that the U.S. owns all incumbent leaders in HPC, while China has nothing to lose.

Dongarra asked about the air cooled version of the Sunway machine. **Shalf** stated with the air cooled version is there are four areas of the die which are completely independent meaning it is possible to break the chip up into multiple dies. Like the Japanese strategy, the hope is to build a big system and to sell smaller copies to other places. The problem is that university class institutions cannot accept a machine that has onerous liquid cooling requirements, but they do have closet space which makes the air cooled machine desirable.

Bergman asked if Shalf thought the U.S. was doing the right thing and investing in the right areas. **Shalf** cautioned that the U.S. should focused on the long haul; that the attention should be having a supply chain that delivers machines that solve problems for big sciences. The Chinese seem steady handed at this because they recognize the metric for success is that industry depends on HPC and not the machine itself.

Sarkar asked if the Chinese were running into workforce and skills challenges. **Shalf** did not ask that specific question, but workforce development is certainly on their minds. **Sarkar** asked about discussions on the strategy for Post-Moore's Law computing. **Shalf** said there were a lot of wide-ranging discussions about beyond Moore's Law. Fujitsu and NEC want to become more like services companies, hence the AI interest. In the short-term, there was a lot of discussion about Tensor Processing Units (TPU) and neuromorphic chip.

Levermore asked Shalf to comment on his perception of the Google TPU changing the equation. **Shalf** said that the lesson of Google TPU is that a lot can be gained from specialization. Shalf thought ways to apply it will be discovered. A family of tools and methodologies for reducing the cost of developing things that are specialized are imperative. However, Shalf said it is obvious from the long author lists that industry is not leading.

Matsouka noted that the K computer is 5-6 years old and on any metric other than the Top 500, K is still #1. However, K is a bandwidth monster. He also noted that NEC is coming out with Aurora; NEC has gone from more dedicated vertical integration to an accelerated base of design, suggesting this is where industry is going. **Shalf** mentioned that the SX-ACE vector supercomputer is an anomaly in terms of architectural efficiency. **Matsouka** offered a third point that as much as Fujitsu and NEC are building HPC chips, there is very heavy activity in terms of AI, and it is being funded by private-sector business. While China is building things aggressively, there are Japanese companies who are more aggressively pursuing chip design in terms of HPC and AI.

APPLIED MATH CENTER INVESTMENTS-SUCCESS AT SCALE, Abani Patra, ASCR

Patra provided a bit of history concerning the creation of Mathematical Multifaceted Integrated Capabilities Center (MMICCs), which arose from a realization that the traditional ways of supporting mathematics through individual awards, did not address everything that needs to be done. Two things drove the creation of MMICC centers: one was a longer term horizon, 5-10+ years' impact, and two was cross-cutting, where more than one particular mathematical topic is needed.

The program review for the three MMICC centers was held in October – November 2016. In four years' time, there have been 329 peer reviewed publications, more than 30 faculty and lab researchers were trained, and multiple awards were received and fellowships bestowed. The study group of researchers, Principal Investigators (PIs), and other agency program managers stated they would never have embarked on this successful line of research without the MMICCs program.

The study group discussed 10 questions concerning the scope, uniqueness, integration, project management, best practices, sustained support, exploration, response design, incentives, and new structures of MMICC centers.

The scope was well-defined but needed a mechanism to adapt to changing priorities, research directions, or poor performance.

The MMICC projects were considered "right sized" but the group indicated flexibility was needed to form cross-institutional teams.

Laboratory and university skills were integrated among those with existing DOE connections, early career investigators, and junior researchers.

The study group suggested that graduate students and post-docs should be engaged directly through laboratory internships.

In terms of project management practices, activities that occurred early in the process were considered important to encouraging open dialog among team members and to provide focus.

Retention was a bigger challenge than recruitment. Collaborative opportunities, guaranteed funding, DOE laboratory time, and high-profile PIs attracted top talent.

The study group suggested that the centers remain organized around mission-driven problems; it is necessary to maintain several diverse centers, and finding ways to disseminate successful results will help sustain support in Applied Mathematics and Statistics.

Encouraging exploration of new and risky research paradigms was considered. The group indicated that flexibility is the key; there should be opportunities for teams to self-assemble, thus allowing the centers to evolve. Additionally, PIs should have the freedom to issue new subcontracts and add new people.

Responding to simultaneous support of long standing research topics and fast changing demands was discussed. The group suggested providing long-term commitments with the latitude to rebalance and reconfigure as necessary; funding parallel work; and that including personnel from labs and academics was essential.

To incentivize the integration of researchers, the centers' funding cycle needs to match that of academia and funding needed to be provided for a fixed term.

The study group offered four structural ideas to improve MMICCs: add a fellows program, allow supplemental proposals for collaboration, create a major milestone at three years for serious feedback, and develop a mechanism for renewals that is appropriately customized.

Discussion

Gregurick stated that at National Institutes of Health (NIH) - National Institute of General Medical Sciences (NIGMS) is the idea to have the option of pilot projects in the bigger centers. NIGMS has had great success bringing in junior, untenured PIs to have a 2-year project that is a bit more high-risk, they then spin off and become their own independent investigators. The idea of adding that flexibility in NIH centers with junior PIs has been fairly successful.

Levermore stated that an intermediate, incubator program might be good so that the creation of such centers does not introduce the risk of creating concentrations to the detriment of other areas. **Patra** mentioned that deep partnerships among applied math people in the centers probably would not have happened otherwise. The virtual handoff to concurrent development speeds everything up to 1-2 years rather than 1-2 decades.

Crivelli mentioned bringing the work of MMICCs to the next level and asked if MMICCs support other areas. **Patra** stated that one of the things discussed among the group was how these centers become repositories of findings that can be accessed by the entire community seamlessly. Currently they hold workshops or researchers call them, but there is no easy way to partner.

David Brown, LBNL, mentioned things that came out of the MMICCs workshop in 2011. The basic model for MMICCs was based on a very successful model at the labs within the applied math base problems. Two features of the program that were identified was one, having a larger effort with a single PI and concentrating on a single scientific area, and two, focusing on developing mathematics that would address challenges that came up in those scientific areas. He asked Patra to compare and contrast the MMICC centers with similar activities at labs. **Patra** stated that while he had good knowledge of MMICCs he was not familiar with the lab programs. However, he thought there was a fair degree of complementarity. MMICCs have done a significant amount of cross-institutional collaborations, whereas the traditional lab model would not allow that to be easily done. While there are still issues about the flexibility of moving

resources around, most of the scientific areas have a multiple lab base and most of the work has partners across institutions. The notion of intellectual partnerships formed by a common challenge is translated nicely here. The idea of having a PI who is strong enough and intellectually engaged enough to influence the outcome of particular sub-projects is necessary and borrowed completely from the labs. **Levermore** offered an example that George Karniadakis became an employee of PNNL. In the course of going to PNNL, he took some young people with him. It is too early to say, but if this kind of thing is replicated, then it becomes a big workforce plus.

DOE PROJECT LEADERSHIP INSTITUTE, Ben Brown, ASCR

The DOE Project Leadership Institute (PLI) is a cohort-based program created to connect and develop project leaders across the DOE complex. Started in CY17, the PLI is a year-long commitment that results in a formal certificate.

The focus of PLI is on the project delivery side of project management. Project Management Institute and Construction Industry Institute have a larger zoom than the DOE project management concept. DOE owes it to itself to have a project management program targeted to what DOE does because since 2002 SC has completed 77 projects each with costs over $5M (total cost ~$6B).

The former Secretary of Energy, Ernest Moniz signed a memorandum in 2014 that mandated the creation of a project leadership institute. The PLI mission is to cultivate a diverse network of successful DOE project delivery practitioners. There are four learning objectives, including strategic thinking and analysis, organization and general management skills, team building, and communications.

There are 22 participants representing 13 organizations in the 2017 cohort. The program content is focused on five in-person themed events, one online course, and a year-long capstone project. The PLI program is a partnership with an academic institution; it is a hybridization of Stanford University's Project Management content plus DOE content. The capstone project is designed to provide a team-oriented and integrative experience to culminate the course of study. The 2017 capstone project assignment is to write a case study of the Superconducting Super Collider project.

Discussion

Dunning noted that he thought the PLI was going to be a very valuable educational opportunity. He stated that there is a big difference between managing projects and managing big computing projects, and coming up with best project management practices for big computing projects would be extremely valuable. **Brown** said he thought comparative analysis was exactly what the PLI was built to generate within cohorts. The Stanford course module steps through several different paradigms of project management including agile. Surveying conventional project management approaches is part of what the PLI is built to do. Over time it is essential that the PLI attracts people with experience in different types of projects that are run in different ways so that the shared experience can be offered to the whole group. **Hey** sought clarification from Dunning, on whether he was referring to computing hardware or software projects. **Dunning** indicated that sometimes they are coupled, like the ECP, it is actually both hardware and software components. Dunning also suggested that the PLI create short courses for individuals who were responsible for some aspects of the project, but who did not want to learn

all of the details of project management. **Reed** added that managing projects successfully is both an art and science and many in the research community are sadly lacking.

PUBLIC COMMENT

None.

Reed adjourned the meeting at 5:46 p.m.

Wednesday, April 19, 2017
Morning Session

Reed called the meeting to order at 8:40 a.m. and introduced the first presenter.

QUANTUM TESTBED STAKEHOLDER WORKSHOP, Claire Cramer, ASCR

Quantum computing is a completely new area for DOE and despite the amount of expertise within DOE the question remains how to build up a community in this area. Cramer held a Quantum Testbed Stakeholder Workshop to share capabilities and interests, discuss how to set up a quantum testbed (who can do what where and what does that mean), determine how to achieve goals, and identify relevant technological challenges and opportunities.

In February 2017, a 2.5 day workshop was held. The workshop used a modular agenda enabling people to pick and choose the session of interest to them. The first half of the first day was inspirational, with a series of plenary talks surveying the most mature quantum computing devices, the state of the art of the devices, and basic research and development needs to continue to advance those devices into maturity. During the first day, the weapons and SC multi-purpose labs and FermiLab gave overviews of their interest and capabilities in quantum computing.

The second day consisted of break-out discussions covering a wide range of topics, including the look of a quantum testbed, similarities to a leadership computing facility, sharing the capabilities of the testbed, identifying user communities, qubit technologies, challenges in simulation and emulation, and exploring different architectures.

On day 3, industries came in and shared their needs and interests in quantum computing. One message that came through loud and clear was the need for the interested community to explore the application space. The participants also discussed what role a DOE testbed could play in advancing this technology, such as standards and benchmarking.

Discussion

Dongarra asked about global quantum computing. **Cramer** stated while that was not discussed at this workshop, there were some non-U.S. attendees. However, within the quantum information science interagency group there has been a lot of focus on developments in quantum computing overseas.

Svore asked if Cramer learned more about the balance or focus on hardware versus software and how that will play out in a testbed program. **Cramer** indicated that along the spectrum of user groups, there are those who want full transparency and full control and to figure out the failure mechanisms and how to make it better. On the other end of the spectrum are folks who would like to do software and application development and algorithm testing and they want a stable system. Since one testbed cannot meet all of these elements, there are two options: to have two testbeds or to have one testbed that evolves over time.

Hey mentioned that one notable absence in the presentation was Weinmann's group at National Institute of Standards and Technology (NIST) because outside the U.S. they would be regarded as the leaders in U.S. quantum computing. **Cramer** said she reached out to NIST and asked how they would like to participate in this workshop. While Weinmann is absolutely a leader in the field, he is very much on the basic research side; he wants to build laboratory prototypes.

Svore asked if NIST can play a role in benchmarking and standards. **Cramer** noted that DOE would proceed collaboratively with NIST in that role. NIST develops community-based standards based on consensus. DOE has also played a role in standards development. While there was not a lot of discussion of standards in the workshop, there was a general consensus that this needs to be considered now rather than later. A government testbed was considered a neutral meeting ground.

Reed asked about next steps Cramer would like to see. **Cramer** mentioned a series of intelligent investments to allow genuine and honest exploration of the potential of this technology, broader access to early stage device technologies, and ways to do pre-competitive research and development to the basic building blocks of a quantum computer.

Helland noted that Binkley is taking a holistic approach and has established a working group between all but the Fusion Energy Program Office within SC to think about the strategic focus.

MACHINE LEARNING, Shinjae Yoo, Brookhaven National Laboratory

Yoo discussed why ML is important, explained ML concepts, Big Data & ML, and Research Areas. Tom Mitchell (2006) asked "how can we build a computer system that automatically improves with experience, and what are the fundamental laws that govern all learning processes?" ML joins architectures and algorithms that effectively handle data and orchestrates multiple learning subtasks in large systems. In other words, ML sits at the center of an algorithm-infrastructure-data triangle.

Yoo described different types of learning used in ML, including generative, discriminative, active, multi-task, kernel, metric, dimensionality reduction, and feature. ML algorithms include Bayesian, instance-based, regularization, decision trees, association rule mining, and ensemble learning. Within big data ML concepts include volume, velocity, veracity, variety. Yoo described MapReduce, Power-Iteration-based Method, and Streaming approximations.

The potential research areas in ML include unsupervised, active learning, in-situ and streaming analysis, new architectures, programming models, compiler technologies, workflows to leverage HPC more effectively, new mathematical solutions, solvers, and libraries for HPC, foundational theories, automation for simulation or experiments, fusing theory, simulation, experimentation and ML, interactive analysis in petabyte scale data, and text mining.

Discussion

Reed asked about the concept of AI technology as a black box as a research direction. **Yoo** indicated that while this area is important, interpretation is similarly important.

Crivelli asked how Yoo is getting ready for AI. **Yoo** said that AI is a much bigger model; that deep learning (DL) is one branch of ML, and ML is one branch of the approach to enabling AI; it is an inclusive relationship utilizing ML algorithms to solve challenges.

Crivelli stated that in her field, different results are obtained depending on the metric used. She asked if Yoo had the same observation and if there is a need to develop metrics to compare apples to apples. **Yoo** mentioned that there are two fundamentals to be able to compare apples to

apples: a shared benchmark data set, and a benchmark of a variation metric to design better or different algorithms.

Sarkar inquired about any tension between science and ML. **Yoo** thought that the science community is becoming more accepting of ML due to excitement about applying ML to a research problem. He felt this change in attitude was based in part on a Journal of the American Chemical Society article. **Sarkar** asked if using simple models for large volumes of data applies to ML in the context of science or if applied math could play a larger role to develop advanced models. **Yoo** expressed that the model should fit the problem if the problem can be solved with a simple model. For more challenging problems, it makes sense to create and use more complex models.

Lethin asked about the use of generative models where the neural network is run backwards. **Yoo** indicated it is great for data simulation and also one of the popular features of the ML effort. Generative models are a great direction to explore and are included under the big umbrella of the unsupervised and active learning aspects.

Matsouka asked about DOE's competitive advantage in ML. **Yoo** said that industry is moving faster than imagined and it is not easy to compete with them. HPC technologies are still contributing to the DL or ML community because they are using the technology developed by the ASCR and SC communities. There is also a lot of room for improvement in model and data parallelism. **Matsouka** added that from the chip and system aspects, the improvements are not driven by HPC market economics but by industry needs; although industry exploits a lot of the technologies from HPC, they are utilizing those to satisfy the market needs. **Helland** added that the real benefit in ASCR and other program offices will be to think about what these processes are going to look like and how we can help development of detectors that have smart learning embedded in them. **Matsouka** noted that rather than just focusing on the processes we need to think about the machines being designed. **Helland** agreed noting that as the exascale program has strong partnerships with vendors, the facilities put on the floor invest in non-recurring engineering costs to make sure DOE needs are getting met. The vendors are building for industry that wants ML, so DOE has to make sure to take advantage of all those chips and work with industry to design machines that our scientists can use. **Hey** suggested that labs have an advantage in possessing large, complex data sets that are publicly accessible.

Bergman stated there is no strategic, cohesive plan and asked what is being done holistically. **Reed** stated that historically, in HPC space, governments have long been the dog not the tail, but in ML DOE is definitely the tail not the dog. DOE is trying to carve out that space.

Reed adjourned the meeting at 10:15 a.m. for the morning break and reconvened at 10:31 a.m.

EXASCALE PROJECT UPDATE – SOFTWARE STACK TECHNOLOGIES, Paul Messina, Argonne National Laboratory

Messina provided a top-level view of the Exascale Computing Project (ECP) giving a brief review of the project and the leadership team. He discussed software, hardware and co-design examples of activities. The next step for software (SW) is to complete the gap analysis that is currently underway. The gap analysis includes responses from the vendors to the ECP request for information (RFI).

Messina gave an update on ECP activities which consists of six strategic pillars: national security, energy security, economic security, scientific discovery, earth system, and health care. There are currently five ECP co-design centers and 27 application development milestones have

been delivered. There are two new activities since December 2016 including eXproxy applications and IDEAS-ECP.

In terms of hardware (HW) technology activities there are four contracts fully signed, one being routed for signatures, and one at DOE/NNSA for approval.

The ECP teams are beginning work on the SC system via early access, ALCC awards, and testbeds. In terms of communication and outreach, ECP has a website, held its 1st annual meeting, and launched a newsletter "The ECP Update." ECP also has developed the Industry Council Charter. The Industry Council consists of representatives from 14 Forbes Category industries. The first meeting of the Industry Council was held in March 2017 and the next will occur in October 2017.

Discussion

Berzins asked about industrial applications that run on exascale machines, the SW stack coming together at scale on pre-exascale machines, and when it will be confidently working. **Messina** stated industrial applications on machines will not necessarily happen except on a case-by-case basis. At the first ECP meeting, the 15 companies represented expressed interest in giving input on what they see as their needs. The next step is to have phone calls with each representative which will explore, in depth, what they see as their applications and their current understanding of the requirements to help them get to exascale. Many of the companies are already using the DOE facilities. There has been no promise made that their applications would be part of DOE's application portfolio, but to the extent that they would like to be users down the line, once there is enough of the SW stack available as a testbed, ECP would try to make time available to industry. An RFP for industry applications in our focus area is still desired. There is anticipation of having an integrated SW stack large enough to be evaluated in 1.5 years. Regarding running at scale, it will be running at whatever scale is available; by the end of 2018, Summit and Sierra at Livermore will be open and ECP expects to have access to that. Some applications are also early science applications for labs and work with facilities on how to install exascale SW stack is occurring.

Chapman asked if there are any bigger plans for training. **Messina** indicated that ECP plans to have internal training with participants in which discussions on future architectures will be held.

Bergman asked about shared milestone plans for HW. **Messina** mentioned there were similar plans to SW but not quite the same because vendors are doing HW research and development. There is a formal review of vendor products every 6 months. **Bergman** asked if there were any workforce challenges because of the prominent researchers involved in ECP. **Messina** noted that part of the researcher's time is being paid by ECP and that they find the project compelling and complimentary to their own work, therefore, they have dedicated themselves to ECP. Messina had not heard much concern about finding people for the projects. Short-term staffing is okay, but there is concern for the long-run for two major reasons: people involved will be very sought after due to their accomplishments on the project; and funding uncertainty, some may leave due to fear of lack or loss of funding. Sustaining, not attracting, is the workforce issue.

Sarkar asked about mitigating possible risks concerning performance portability. **Messina** clarified that Sarkar was talking about the conflicting requirements and said the hope is to have enough tools with a high enough level of extraction that a large fraction of the code does not have to be changed. One of the reasons for investing heavily in libraries and tools is the hope that

they can be modules that are integrated into the application codes and those modules have different versions for the different architectures. Measurement of performance portability is on the horizon in a couple of years, but a rule of thumb would be a factor of two.

Matsouka noted that the usual SW development scheme is to have a testbed and a procedure to test the SW and that the added difficulty is these will have to work at scale with only a few machines that are at scale. He asked how ECP will solve the difficult problem of aggregating all of the SW and making sure they cohesively work together in a production environment. **Messina** said that ECP staff is aware of the problem but they do not have a solution yet. All the projects are using a formal process, but it is not the same one. **Matsouka** asked if the timeline of accelerated delivery is the same. **Messina** confirmed that the timeline has not changed at this time.

Huntoon asked if the effects of access to external resources on the network have been taken into account. **Messina** stated that the Energy Sciences Network (ESnet) is evolving and they will play a role in the ecosystem. ESnet has been invited to the Annual Meeting and the first PI meeting to make a presentation on their plans. Currently, ECP is ensuring that we have the right communication. **Huntoon** asked for examples of collaborations with NSF. **Messina** said there is nothing formal; however, ECP is ensuring the National Science Foundation (NSF) is aware of the plans and vice versa.

Lethin asked, in the context of the PLI, what is the degree to which ECP is adopting the agile method for software development. **Messina** indicated ECP is considering using more tools from Atlassian and the agile development method. **Lethin** asked if the Confluence database would be open. **Messina** said no, that would be a problem for participants.

CSGF LONGITUDINAL STUDY, John Wells and Tara Dunderdale, Westat

Wells and Dunderdale described a follow-up study of the U.S. DOE CSGF: 1991–2016A recipients. This was an update to the 2012 study and included 414 of 436 participants. The methods used include a survey of alumni and fellows (67% response), telephone interviews with targeted alumni, and curriculum vitae (CV) data analysis (75% provided updated CV).

The themes that came out of the interviews were skills, content expertise, interdisciplinary skills, and flexibility to pursue topics within computer science and engineering. In surveys and interviews, the majority of respondents felt the CSGF fellowship was useful for obtaining their first professional position. Virtually all alumni are employed in academia, industry, or DOE labs; the most common current employment area was industry followed by academia. From the CV data, 165 out of 243 had received at least one award ($\bar{x} = 3.5$); 51 grants, ($\bar{x} = 6.5$); and 35 patents ($\bar{x} = 3.5$). In terms of publications, these increased every year and were overwhelmingly published in highly influential journals (2,100 articles were mostly in the top 10% of journals).

DOE CSGF is a highly sought after fellowship. The alumni gave the program high praise for its unique benefits and stated that CSGF provided a broad scope of training and development. Alumni remain involved in the program through recruitment, guidance and professional development, and mentoring. A long-term career benefit included direct translation to permanent employment where alumni have achieved leadership positions.

Discussion

Dolbow asked to what extent fellows, who have gone to academia, maintained active ties back to the DOE labs. **Dunderdale** said that was not something they captured. **Wells** added that they looked at connections and the ways alumni stayed with the program.

Svore asked if Dunderdale and Wells spoke with anyone outside the program, how many people recognize or know about the program, and how the fellows learned about the program. **Wells** indicated they did not ask if the program was prestigious; they asked if it was a factor in obtaining the respondents' first professional position. **Svore** asked for any data on how the program is perceived by outside people, and how the winners of the fellowship learned about the program. **Dunderdale** indicated that 44% said their academic advisor was most important to them in learning about the program. **Svore** continued by asking how many industrial companies know about program in order to assign prestige to the CSGF. **Wells** reminded ASCAC that this data was from the recipients' perspective rather than industry.

Reed adjourned the ASCAC meeting for lunch at 12:18 p.m.

Wednesday, April 19, 2017
Afternoon Session

UPDATE FROM THE SUBCOMMITTEE ON FUTURE COMPUTING TECHNOLOGIES, ASCAC

Sarkar reviewed the subcommittee charge to examine the opportunities and challenges for HPC and potential research areas with a report due in December 2017. The subcommittee interpreted the charge to focus on the timeframe of different technologies, especially post-exascale, to seek community input, which is considered critical, and to provide prioritization with respect to the HPC requirements for advancing science.

The plan is to finalize the subcommittee by the last week in April and then begin bi-weekly meetings. Sarkar is keen on identifying a broad set of resources from which the subcommittee can draw, and they want to talk to subject matter experts in government labs. The subcommittee will discuss these opportunities and challenges in the context of advancing science. Sarkar told ASCAC to be prepared for an update at the next ASCAC meeting.

Discussion

Matsouka asked how the subcommittee plans to coordinate with the International Roadmap for Devices and Systems (IRDS) activities. **Sarkar** said that within the subcommittee they can discuss this with Tom Conte, but Sarkar's sense is that there is a lot of activity going on and the subcommittee can draw from those resources. They are trying to make judgement calls about prioritization from the viewpoint of relevance to advancing science applications.

CORI EARLY SCIENCE AND APPLICATION PERFORMANCE, Jack Deslippe, NERSC

The latest news about the Cori system is scientists are currently stressing the system, running at full scale, and doing exceptional science. Deslippe highlighted two groups, one from cosmology, and the other from materials science.

NERSC, being the mission HPC facility at DOE SC, has an astounding number of users, projects, codes, locations, and institutions. NERSC users are coming from traditional Central Processing Unit (CPU) systems and NERSC must support the entire SC community as well as transition to more energy efficient systems. To meet this challenge, NERSC Exascale Scientific Application Program (NESAP) committed two tiers of resources, including staff time, to 20 areas based on the workload. To meet user needs, NESAP chose different science areas,

algorithms, and methods to ensure they had representative codes from each area. NESAP is covering approximately 50% of the workload, and with a proxy or third tier NESAP is covering approximately 60% of the workload.

Deslippe described the post-doc program and the activities being performed by the post-docs in the Application Performance Group. The post-doc program helps to generate the kind of workforce the community needs.

NERSC decided to utilize the Roofline Model to communicate the 700 codes to users.

Facilities provide a venue for successful interaction between applications and vendors. The role of the facility is to be a bridge, an engagement point, and it can be a location for hosting training sessions and providing a central point for documentation, lessons learned, and information about how to use tools.

Deslippe shared preliminary NESAP Code Performance on Knight's Landing (KNL) and average code speedups. Optimized Haswell code is 2.3x faster than the baseline code; KNL is 3.5x faster when optimized. In terms of performance ratio, baseline code is slower on KNL than Haswell, but optimized-speedups are higher on KNL. Comparing the performance of optimized codes on KNL to baseline codes on Haswell indicated where DOE would be if "business as usual" was chosen and a Haswell machine purchased. The same comparison was performed on KNL versus Ivy-Bridge (Edison). Both the baseline and optimized codes performed better on KNL than on Edison. There is a high demand for Cori KNL, utilization is very high at around 90%, and there is a backlog of the queue. The top 10 projects on KNL include projects from each of the six SC offices.

Deslippe shared three projects that are currently utilizing the Cori system. One is a 45 qubit simulation that is the largest quantum computer simulation to date. Another is Quantum ESPRESSO which includes MD simulations, MG^{2+} diffusion, and Nudged Elastic Band (NEB) calculations; and a full system DL science to classify climate data.

Discussion

Hey asked how KNL compares with NVIDIA. **Deslippe** talked about comparisons at peak performance level that are within the same 50% of peak performance. Although there are a couple of applications running on both, Deslippe had no example to show Cori versus Titan.

Lethin asked to what extent the network made a difference in performance versus the distributed ML training being done now in clusters. **Deslippe** stated he is not an expert in DL. The researchers were using convolutional neural networks that were built into IntelCaffe. He speculated that they tweaked the depth and stride enough to maximize the performance of the system. **Lethin** asked what was simulated in the quantum simulation, and is there a way to compress representation to save memory. **Deslippe** said they simulated a number of gates and in these particular runs it was not compressed. **Svore** added that on the application they are looking at quantum supremacy experimentation.

Berzins asked about the difference in speed between the MCDRAM and Haswell. **Deslippe** said the total bandwidth is about a factor of 5. **Berzins** asked about the comparative energy efficiencies on the different processors. **Deslippe** said that for KNL versus Haswell, KNL is a bit lower in terms of total watts; peak performance is 2x higher on KNL; however, if you measure energy efficiency in terms of real applications it depends on performance δ for the different applications. **Berzins** asked if Linpac runs hot on KNL. **Deslippe** said yes, but Linpac increases the power draw on all architectures. **Berzins** commented that it comes back to peak and realizable peak performance; his final question was how much should be measured in imagined

flops versus petaflops, and how close to the real peak performance for some applications can be achieved. **Deslippe** said, on KNL, the frequency for AVX instructions is lower than non-AVX instruments. Second was to use the empirical Roofline model developed at Berkeley where they have defined the limits on the Roofline through running real applications rather than computing theoretical peak performance.

Matsouka brought up a DL question and asked about conversion properties. **Deslippe** stated he is not an expert on DL algorithms but he could provide a connection. **Matsouka** asked how people designate the usage of Haswell versus KNL. **Deslippe** said while there is currently one scheduler for the whole system, which makes it possible, in principle, to run a job on both systems, the user decides which queue.

Crivelli asked for more information about the post-doc program and its continuation. **Deslippe** said they plan to continue the post-doc program with a steady state of 10. **Crivelli** asked if there was a plan to have fat nodes on the system or terabyte memory nodes. **Deslippe** said there is no plan to have that level of memory per node on Cori right now.

Reed called a break at 2:50 p.m. and reconvened ASCAC at 3:05 p.m.

DOWN THE RABBIT HOLE: FROM B3LYP TO X86, Jeff Hammond, DOE-CSGF Alumni

Hammond provided a biographical look at his development from pre-computers to his current career path. Hammond, an alumnus of the CSGF program, described his practicum at PNNL as the highlight of his fellowship. He utilizes the skills he learned while a post-doc to move industry to be more like DOE in terms of working in teams, interdisciplinary skills, and communication skills. In summary, Hammond praised the CSGF program.

Discussion

Dolbow asked if Hammond had considered how CSGF could reach out to a broader set of students from non-traditional channels. **Hammond** noted that the applicant pool is 800 and was then cut down to 20. There is currently a lot of outreach success and the only thing he would do differently is to make the program bigger.

Svore asked if Intel recognized the CSGF Fellowship. **Hammond** said that within DOE, CSGF is immediately known, but he is unsure about the knowledge at Intel. The beneficial outcome of going through the CSGF program is that your work is evident on your CV; he encourages CSGF fellows to complete a DOE post-doc before going to industry.

SciDAC – EFRC COLLABORATION, Wolfgang Windl, Ohio State University

Windl and Wigner described their program called WastePD, which is about performance and design of encapsulation of waste materials from spent nuclear fuel. At Hanford and Savannah River sites there is still a lot of liquid waste which is highly corrosive and radioactive. The U.S. is attempting to get rid of the liquid waste by turning it into glass, ceramic, or concrete which has to last 100K years; however some of the containers for the liquid waste are corroding and leaking. The common method for testing corrosion is atomic-scaled modeling which includes the tasks of structural stability and energetics, surface processes (e.g., oxidation and desorption), and externally-controlled processes (e.g., deformation and field evaporation).

Windl described WastePD's collaboration with SciDAC and discussed the wish list WastePD submitted to SciDAC. The wish list included high computational demands in many calculations,

empirical potential fitting, visualization and analysis of complex atomistic simulation results, and uncertainty quantification. WastePD was interested in expanding their Molecular Dynamics (MD) capabilities. SciDAC QUEST and SUPER focused on generation and evaluation of new potentials. SDAV focused on acquisition and analysis of relevant data from large systems.

Wigner has been managing the SciDAC collaborations. WastePD had automated code that they knew worked but they were unsure what it could or could not do. The SciDAC collaborations gave them a chance to investigate the code more thoroughly. Wigner described the work accomplished with each SciDAC center in detail.

Discussion

Lethin asked how long it will be until the collaboration results transition to the practical problem. **Windl** said that all of the groups (glass, ceramics, and metals) need to describe the interaction potential fitting capability. They will hold a workshop in a month where the other teams will be taught how to use the new potentials capability; basically it is entering the process flow now. All the groups are aware of VisIt software and are using it. **Wigner** added that to design a waste container you must know what makes a good waste container, then design the material that does that. People do not know why things are corrosion resistant; most of the research in this area has been entirely trial and error, or just experience and knowledge. With faster codes we can cast a really wide net and test 10,000 alloys all at once and pick the best one. By making the codes faster and the work smarter it expedites the process of figuring out what makes something corrosion resistant.

PUBLIC COMMENT

None.

Reed adjourned the April 2017 ASCAC meeting at 4:11 p.m. (ET).

Respectfully Submitted,
Tiffani R. Conner, PhD, PMP, AHIP
Science Writer
ORISE