Draft Minutes
Advanced Scientific Computing Advisory Committee Meeting
July 27, 2015
Marriott Gateway Crystal City
Crystal City, Virginia

ASCAC Members Present

Martin Berzins
Barbara Chapman
Silvia Crivelli
John Dolbow (departed at noon)
Jack Dongarra (via telephone)
Roscoe Giles (Chair)
Sharon Glotzer (via telephone)
Anthony Hey
John Negele (via telephone)
Linda Petzold
Daniel Reed
Vivek Sarkar
Dean Williams

ASCAC Members Absent

Vinton Cerf
Jacqueline Chen
Gwendolyn Huntoon
Juan Meza

Also Participating

John Steve Binkley, Associate Director, Office of Advanced Scientific Computing, Office of Science, USDOE
Laura Biven, Senior Science and Technology Advisor, Office of the Deputy Director for Science Programs, USDOE
Arthur Buddy Bland, Director, Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory
Rick Borchelt, Director, Office for Communications and Public Affairs, USDOE
Ronald Brightwell, R&D Manager, Scalable System Software Department, Sandia National Laboratories
Christine Chalk, ASCAC Designated Federal Officer
Susan Coghlan, Deputy Division Director, Argonne Leadership Computing Facility, Argonne National Laboratory
Leland Cogliani, Associate, Lewis-Burke Associates
Jody Crisp, Oak Ridge Institute for Science and Energy
Patricia Dehmer, Acting Director, Office of Science, USDOE
Frederick O'Hara, ASCAC Recording Secretary
Robert Ross, Deputy Director, Scientific Data Management, Analysis, and Visualization Institute, Argonne National Laboratory
James Sethian, Head, Mathematics Group, Lawrence Berkeley National Laboratory
Robert Voigt, Consultant, Krell Institute
Kathy Yelick, Associate Laboratory Director for Computing Sciences, Lawrence Berkeley National Laboratory

About 55 others attended the one-day meeting.

Morning Session

The meeting was called to order by the Chairman, Roscoe Giles, at 8:30 a.m.

Christine Chalk, ASCAC Designated Federal Officer, made safety, legal, and convenience announcements.

Giles welcomed everyone to the 40th ASCAC meeting. He noted that the Computational Science Graduate Fellowship annual review was being conducted in the same building and suggested that Committee members observe as much of that review as possible.

Patricia Dehmer was asked to present the program-planning perspectives of the Office of Science (SC) of DOE.

Program planning in SC is conducted through both top-down and bottom-up processes, considering mission needs and scientific opportunities, respectively. Mission need is set by the Executive Branch priorities, within which

1. Administration priorities are set by deliberations at the National Science and Technology Council (NSTC) and its working groups, Office of Science and Technology Policy (OSTP) and its working groups, other administration-convened ad hoc working groups, and interagency coordination and
2. Departmental priorities are determined through DOE and program strategic plans and the *Quadrennial Technology Review* (QTR) and *Quadrennial Energy Review*.

Congressional-branch priorities are reflected in legislative authorities and annual appropriations. Within the myriad scientific opportunities, program priorities are set via the engagement of community experts and stakeholders through federal advisory committees' reports and recommendations along with DOE-sponsored scientific and technical workshops/reports and non-DOE scientific and technical workshops/reports.

During the three weeks prior to this meeting, the accomplishments of all of the advisory committees were reviewed by SC leadership. It was a very instructive exercise. Often, the advisory committees have changed the courses of their programs. Their studies and recommendations have been incredibly important to SC and DOE. Exemplary activities include the Basic Energy Sciences Advisory Committee (BESAC) charge to provide advice on the future of photon sources in science; the Nuclear Science Advisory Committee (NSAC) charge to annually study the National Nuclear Security Administration's (NNSA's) molybdenum-99 program; and the Advanced Scientific Computing Advisory Committee (ASCAC) Committee of Visitors (COV) to the Office of Scientific and Technical Information (OSTI), which is in progress. These studies are taken very seriously.

The Office of the Under Secretary for Science and Energy has been heavily involved in the QTR for the past year. The first draft was issued on the Friday prior to this meeting. Its purpose is to inform/brief the budget discussions. SC and its activities are represented in the 2015 QTR in a chapter entitled "Enabling Capabilities for Science and Energy RDD&D." In the broad view, SC provides tools for scientific discovery and technology development. The SC chapter in the QTR deals with the theme of understanding and controlling matter at the atomic scale. This means unique, cutting-edge experimental tools (light sources, neutron sources, and research centers) for characterization, discovery, and synthesis of novel materials and energy systems. A study cited was the traversing of a catalytic reaction pathway in femtosecond steps. In this study, ultrabright femtosecond X-ray pulses from the Linac Coherent Light Source (LCLS) allowed researchers to directly characterize catalytic reaction intermediates. Another study cited was modeling and simulation of complex phenomena, requiring leadership-class computers, production-class computers, and the Energy Sciences Network (ESnet). DOE computers have an enormous impact across the engineering and manufacturing space. Virtually all technology offices rely on high-performance computing and on materials by design.

An analysis of the list of annual ratios of funding over 35 years for the six offices within SC that have advisory committees [Nuclear Physics (NP), High Energy Physics (HEP), Fusion Energy Sciences (FES), Biological and Environmental Research (BER), Basic Energy Sciences (BES), and Advanced Scientific Computing Research (ASCR)] revealed significant changes over the years. Those changes were driven by the growth of the light sources, the importance of the materials-by-design era, and high-performance computing. BES has the biggest dollar growth by virtue of the growth in light sources. ASCR has the biggest percentage growth because of the broad need for and adoption of high-performance computing. FES started quite large but went down in priority. HEP expanded rapidly during the period of Superconducting Supercollider construction and then declined significantly with program closures (i.e., the B Factory and Tevatron) and with the off-shoring of high-energy-physics research. And NP has grown and will continue to grow. Many of these changes were rooted in the advisory-committee reports of the past.

All of the charges to and reports from ASCAC since 2013 have had to do with exascale computing. The first was to determine the potential synergies between the challenges of data-intensive science and exascale computing. The second was to determine the 10 principal research challenges and the technical approaches required to develop a practical exascale computing system. And the third was to review the Department's draft preliminary conceptual design for the Exascale Computing Initiative. This Committee has been devoted to the advancement of high-performance computing and has made a huge impact on the budget and the direction of scientific research.

BESAC has responded to two main charges. The first was to provide advice on the future of photon sources and science, considering both new science opportunities and new photon-source technologies in parallel. It responded with the report of the BESAC Subcommittee on Future X-Ray Light Sources, which was in type script (no color) and very short. It made clear that the upgrades planned were not good enough to maintain U.S. leadership in science. The day that report was delivered, the light-source directors were called in to Washington to revamp their upgrades, changing billions of dollars' worth of projects. The second huge charge was to revisit the BESAC 2007 "challenges" report, considering progress achieved, the impact of the challenges on energy sciences, funding modalities, and new areas of basic research not described in the original report.

The Biological and Environmental Research Advisory Committee (BERAC) has been looking at long-range planning for many years through many charges and recommended initiatives for field-based research that capture a multidisciplinary approach and build on observations and modeling. Its reports define the criteria for selecting sites for future BER field-based research and prioritize the sites identified or described.

The Fusion Energy Sciences Advisory Committee (FESAC) was charged to assess priorities among and within the elements of the Magnetic Fusion Energy Science Program, to develop a strategic plan for the Program, and to assess connections between research supported by the Program and other scientific disciplines and technological applications.

The High Energy Physics Advisory Panel [HEPAP, a joint advisory committee with the National Science Foundation (NSF)] was charged in September 2013 to develop an updated strategic plan for U.S. high-energy physics that can be executed over a 10-year timescale in the context of a 20-year global vision for the field. The panel used its standing planning subpanel, the Particle Physics Project Prioritization Panel (P5), to conduct this study. HEP had just closed its two largest facilities at the Stanford Linear Accelerator Center and the Fermi National Accelerator Laboratory (Fermilab). The International Linear Collider (ILC) was budget-prohibited at $12 billion. The culture of collider physics was changing. P5 provided a report, *Building for Discovery: Strategic Plan for U.S. Particle Physics in the Global Context*. SC immediately got on the telephone with congressional committees that were considering the budget. Members of the high-energy-physics community got Congress and DOE to re-think plans for the Long-Baseline Neutrino Experiment (LBNE) and enlarged it to an international facility.

The Nuclear Sciences Advisory Committee (NSAC, a joint advisory committee with the NSF) was asked to conduct a new study of the opportunities and priorities for U.S. nuclear physics research and recommended a long-range plan that will provide a framework for coordinated advancement for the nation's nuclear science research programs during the next decade. NSAC knows well what happened in BESAC and HEPAP and is expected to provide a report with similar impact.

The take-home message is that the work that advisory committees do has enormous, transformational impacts on research and facilities and on the field of science.

Giles said that the Committee had noted the references to ASCAC reports in Congressional budget requests. The Committee has been called on to think about programs that were more for DOE than just SC or ASCR. Are there corresponding advisory committees outside SC that we should work with? Dehmer responded that SC has a long history of using advisory committees. The Basic Energy Sciences Advisory Committee predates the emergence of the DOE technology offices. DOE leadership takes advisory committees very seriously. Other advisory committees in the Department are younger than ASCAC, less mature, and have less impact. Advisory committees struggle to figure out who they are and what their goals are. The charges given to them have challenged them and made them grow and evolve.

Williams asked how all of the advisory committees were interacting and integrating. Dehmer replied that one very successful result has been joint subcommittees. There have been cases where the advisory committees have found similar activities and have worked cooperatively.

Giles noted that the *Quadrennial Technology Review* and the Exascale Initiative have overlapped. He asked if there had been a correlation. Dehmer answered, not overtly, but high-performance computing has been a driving force in the development of scientific programs.

Steve Binkley was asked to present a review of the activities of ASCR.

ASCR conducts mathematics research, computer science research, Scientific Discovery Through Advanced Computing (SciDAC) partnerships, exascale computing, facility operations, and a postdoctoral program. ASCR's investment priorities are the exascale, facilities, and large scientific data.

ASCR's FY15 appropriation was $523.4 billion. Its FY16 President's Request was $621 billion. The House mark was $537.5 billion (making ASCR budget constrained). The FY15 budget had language that allowed moving forward with the exascale. The Senate mark was $621 billion, close to the President's request. The budget is now in conference. The exascale is a major theme in the ASCR budget, appearing in the Applied Mathematics, Computer Science, Computational Partnerships, and Research and Evaluation Prototypes budget lines.

The Office is still awaiting word on the appointment of a permanent Director of SC.

Pavel Bochev, an ASCR computational mathematician at Sandia National Laboratories received the DOE Lawrence Award.

The awards in the DOE SC Early Career Research Program are substantial. For university awards, the minimum award is $750,000; for national-laboratory awards, the minimum award is $2.5 million. To be eligible, no more than 10 years can have passed between the year that the principal investigator (PI) received his or her PhD and the year of the deadline for the proposal. The 2015 ASCR Early Career Research Program awards went to Cory Hauck at Oak Ridge National Laboratory (ORNL), Jonathan Weare at the University of Chicago, Henry Hoffmann at the University of Chicago, Christian Engelmann at ORNL, and Daniela Ushizima at Lawrence Berkeley National Laboratory (LBNL).

Operating and maintaining large-scale computing facilities is a major mission for ASCR. Seven machines are currently being installed or upgraded: Edison, Titan, Mira, Cori, Summit, Theta, and Aurora. After the upgrades, Summit will have a system peak of 150 PF, multiple IBM Power9 CPUs and multiple Nvidia Voltas GPUs, and about 3500 nodes. Aurora will have a system peak of 180 PF, Knights Hill Xeon Phi many-core CPUs, and more than 50,000 nodes. Aurora is one of two systems now planned for Argonne National Laboratory (ANL); the other is Theta. The delivery schedule for chips is uncertain. The situation is being followed very closely.

ESnet has been extended to serve Europe. The CERN link is operational. The Large Hadron Collider (LHC) has resumed data production. Tier-1 data laboratories at Brookhaven National Laboratory (BNL) and Fermilab are receiving those data.

There has been a lot of progress during the past few months on the Exascale Computing Initiative. The effort is led by the Integrated Project Team. A project plan and a management plan have been set up. The Exascale Computing Initiative will follow established DOE review and decision protocols for its execution. A project office has been established at ORNL with heavy representation from the major participating laboratories, exploiting earlier cooperative efforts. An integrated project team has been established; it is refining the Work Breakdown Structure (WBS) and is preparing required project documentation. A top-level WBS activity has been established. A start has been made on a draft of the management structure. Preliminary plans have been received by ASCAC, and another review of the final plans is to be done at the end of the calendar year.

ASCAC polls other SC offices to see what their program plans are and what high-performance-computing needs will need to be met. Review meetings establish consensus on requirements, capabilities, and services. This process addresses DOE mission goals by ensuring DOE science is effectively supported. A series of workshops is being held, one per SC office. These workshops started in June 2015 with HEP and will be completed in September 2016 with ASCR. A summary report is being written for each workshop. These reviews help the Office to define facility and research needs in computer science.

The trend is for more systematic use of modeling and simulation in R&D. The leadership computing facility (LCF) upgrades will allow researchers to address the mission needs and science challenges that can be tackled with the proposed upgrades, such as energy storage. The type of computing that will be needed has inspired the goal of ensuring the ability of ASCR facilities to support the SC mission science in the exascale regime in the 2020–2025 timeframe. The LCFs will serve the needs of a lot of research communities. The huge amount of HEP and LHC data is an order of magnitude increase. How these large amounts of data are handled will have to change in the next few years.

Giles asked how the outcome of the exascale report being presented at this meeting would be integrated into the program. Binkley replied that the Office has worked closely with the Subcommittee. A major recommendation is the development of a management plan. Such a plan is already being worked on. Reed noted that the Subcommittee had tracked what was going on in the initiative. It did not comment on things that were not yet public, but it was aware of them.

Williams asked if the Office would want a study of large-scale data handling. Binkley answered that that will be an area to deal with. The Office will need to know where the data will be, what the software data stack will look like, what application developers will need, the need for real-time processing, the changes needed to be made to the computational facilities of today, and trade-offs with other computational operations. Sarkar noted that a previous ASCAC report was focused on the synergies between data-intensive science and exascale computing. It would make sense to revisit that topic after the Exascale Computing Initiative is under way.

The floor was opened for public comment. Bland asked where the "virtual data facility" would go. Binkley said that he did not like the term "virtual." He has been calling it the science data facility. The Office needs to understand what the software, middleware, and computational science needs are, and then it can plan out a facility.

Cogliani stated that it appeared that the vast majority of exascale funding will be distributed to research and evaluation and a lesser amount to applied mathematics and computational science. He asked if ASCR were going to continue that ratio of funding in the future. Binkley responded that, because of how vendors' development efforts find their way into mature technology that can actually be purchased, it takes about 4 years, so the 2018–2020 range is a target for technological development in order that there will be a machine to buy in the 2023 timeframe. Most of the current funding is going for hardware development. Subsequently, the funding focus will shift to software. A lot of the R&D is done in DOE national laboratories and the vendors' laboratories. In the long view, the budget is pretty well balanced between hardware and software.

A break was declared at 9:47 a.m. The meeting was called back into session at 10:02 a.m.

Wendy Huntoon's presentation about the ASCAC Committee of Visitors (COV) on Next-Generation Networks for Science was given by **Christine Chalk**.

The COV has been charged to assess the operations of the networking programs of ASCR during the fiscal years 2011–14 and to examine both national laboratory projects and university projects. For both of those, the COV is to assess the efficacy and quality of the processes used to solicit, review, recommend, and document proposal actions and to monitor active projects and programs. Within the boundaries defined by DOE missions and available funding, the COV's report is to comment on how the award process has affected (1) the breadth and depth of portfolio elements, (2) the degree to which the program is anticipating and addressing emerging challenges from large-scale scientific facilities and collaborations in support of the DOE missions, and (3) the national and international standing of the program with regard to other computer-science research programs that are also focused on high-performance networking tools and middleware for science. The COV will meet in September 2015 to review program material, in October 2015 to review materials, and in November 2015 to present its report to ASCAC for comment.

Tony Hey was asked to present a report from the Standing Subcommittee on the Office of Scientific and Technical Information (OSTI).

The charge letter requested that ASCAC establish a standing subcommittee for 2 years to advise SC on OSTI and for that subcommittee to examine four questions. Additional guidance from SC added four supplemental questions.

The subcommittee visited OSTI in the spring and early summer of 2015. During the Subcommittee, all the presenters were knowledgeable and enthusiastic about their subject area. From the Director down, the staff in OSTI was very motivated and committed to delivering excellence. It was particularly impressive to see how OSTI had stepped up to the challenge of increased public access to research journal and conference papers. The staff had also thought innovatively about future developments and challenges.

The first charge question was: Are the current OSTI products and services best in class and are they the most critical for the OSTI mission given the present constrained budget environment? OSTI's products and services are professional and generally well done. OSTI is a charter member of CENDI [the Commerce, Energy, NASA, Defense Information Managers Group] and operates Science.gov, CENDI's flagship, cross-agency scientific and technical information (STI) product for searching U.S. Government scientific and technical information. OSTI was the first U.S. federal agency to be a member of the DataCite organization, which now provides researchers with digital object identifiers (DOIs) for their data sets and makes those data sets available to users via Data Explorer. Many of OSTI's services do incorporate leading-edge technologies and, in this sense, can clearly be regarded as "best in class." However, justification for selection of products (e.g., video) is not present; there is overlap among the products; and the software service is dated and out of touch with customers.

In terms of OSTI's public access role, the National Institutes of Health (NIH), with its National Library of Medicine and its PubMed Central repository, must be regarded as best in class among the federal agencies. That said, OSTI's rapid development of PAGES [Public Access Gateway for Energy and Science] and the associated further development of their E-Link service for submission have been impressive.

The Clearinghouse for Open Research of the United States (CHORUS) provides access to journal articles resulting from government funding through article metadata and links to publishers' websites. PAGES can operate successfully independent of CHORUS. There is still much skepticism from the academic library community about DOE's collaborating with CHORUS, and institutional repositories are concerned about the confusion this is causing.

According to the DOE backup plan, during an "administrative interval" of up to 12 months, PAGES will not provide access to the full-text manuscripts. During this time, metadata (including links to the publishers' version of record) will be discoverable through PAGES. In cases where the version of record is not accessible, PAGES will display a link to the accepted manuscript. In all cases, OSTI will maintain a dark archive of manuscripts to be used in the event links become broken or full-text access is otherwise interrupted or discontinued.

The Subcommittee found that, at present, the DOE PAGES service must be regarded as a promising newcomer to public access. It already seems clear that PAGES could prove to be an attractive alternative solution to that offered by the NIH and emerge as best in class in a few years.

The second charge question was: Do OSTI products and services fulfill customer needs now? From the usage statistics, there appears to be significant take-up of OSTI services by the public and by commercial services. Unfortunately, OSTI's services do not appear to be widely used or even recognized by the DOE researcher community. In respect to outreach to the DOE researchers, very little seems to have changed since the Subcommittee of 2009. The services are more targeted towards librarians than researchers, and the OSTI services seem cumbersome.

However, the Subcommittee noted that certain classes of customers (such as major DOE data-program managers in the Atmospheric Radiation Measurement Program and the Carbon Dioxide Information Analysis Center) are satisfied with existing OSTI products and services and see OSTI as one of the few solutions for their needs. Researchers see the need for more integration of the different services and an improved user interface.

The third charge question was: Are the OSTI products and services positioned to evolve to fulfill customer needs in the future? In their presentations, the OSTI staff showed a good awareness of the likely evolution of the services with respect to linking publications to data. The Data ID Service is a useful start, and OSTI's involvement with similar organizations will be valuable. The emerging challenge of collecting electronic versions of graphs, tables, and images in papers does not seem to be currently addressed in detail in OSTI's plans. The DOE scientific and technical information (STI) community must work closely with the DOE research community to develop new useful services for today's and tomorrow's researchers.

The fourth charge question was: What is the national and international standing of OSTI with respect to similar organizations? Nationally, OSTI has a leadership role with the CENDI interagency group in operating the Science.gov gateway to government science information. The recent development of PAGES for access to research journal articles has won the respect of the NSF and the Department of Defense, who are considering using the OSTI system for delivering their public-access plans. This is an area in which OSTI must be a clear leader to fulfill its mandated responsibilities. The OSTI services employ a range of innovative technologies not uniformly available from their peer international scientific information organizations. In terms of international leadership and recognition, OSTI is a founding member of the WorldWideScience Alliance and is responsible for providing novel real-time searching and translation services over globally dispersed multilingual scientific literature. OSTI products and services compare well with those delivered by similar organizations in Canada, France, and the German National Library of Science and Technology. OSTI was invited to chair the Technical Activities Coordinating Committee of the International Council for Scientific and Technical Information (ICSTI).

The first supplementary question was: Is the mission statement sensible in the light of the statutory authorities? The OSTI mission statement is entirely appropriate in targeting DOE researchers and the public.

The second supplementary question was: Is OSTI organized and staffed to accomplish today's mission? The recent re-organization of OSTI in terms of its three core functions has given OSTI clearer focus on DOE research results. It is likely that some changes to the mix of technical expertise at OSTI will be required to design and develop services suitable for modern science environments. If OSTI is to take on a larger role with respect to data, it needs to expand its expertise in this area.

The third supplementary question was: Are the current and planned OSTI products and services the correct ones? The products and services need to be targeted for at least three different communities: the traditional library and information management community, the DOE research community, and the general public. The automated collection of publications and provision of public-access versions should remain a top priority for OSTI. Optimization of the publication-collection method could significantly reduce the burden on the data submitter and significantly increase collection completeness. The start on collecting multimedia content is valuable, but improvements in metadata and consideration of the priorities of different types of multimedia should be undertaken. The Data ID Service is a critical first step towards making data sets citable and linking data to publications.

The fourth supplementary question was: What suggestions would the subcommittee make for the next steps?

- OSTI should initiate some serious two-way outreach and dialog with the national laboratories' research communities to better understand what services they would like and use.
- A more detailed analysis of Google/Bing search results on DOE R&D could help determine on which areas OSTI should focus to deliver complementary functionality.
- OSTI should discuss tool usability issues with the DOE research community with a view to developing an integrated "one stop shop" approach to STI services
- The STI management should be enlarged with researcher champions from each national laboratory.
- OSTI needs to develop the necessary skills to advise researchers about the required data management plans.
- OSTI should discuss approaches to partner with the national laboratories and researchers to improve content completeness and help reach the DOE goal.

The Subcommittee had six specific recommendations:

1. OSTI should initiate a vigorous outreach program with the national laboratory researchers.
2. OSTI needs top-down support from DOE in clearly communicating that providing public access is not a requirement/burden imposed by OSTI but rather a government-wide and DOE-wide requirement meant to share federal research results and accelerate scientific progress.
3. OSTI should work with the DOE research community to re-invent the Energy Science and Technology Software Center (ESTSC) software service with the Linear Algebra Package (LAPACK) team and the Message-Passing Interface (MPI) team. In addition, releasing software in support of a research publication needs to be supported.
4. OSTI should work with the national laboratories to identify researcher champions who can work with the STI community to strengthen the link to researchers.
5. OSTI should elicit community input from a unified user environment and develop a master plan for future development and areas of expansion.
6. Through partnership with the national laboratory librarians and researchers, OSTI should identify and address publication content gaps and develop clear instructions and guidelines regarding content-submission requirements.

In terms of developing a data-management framework for DOE, the Subcommittee suggested that OSTI

- Work with the SC program offices and the research communities in the national laboratories to develop better solutions for linking data and software to publications.
- Participate in reviews of the data needs by discipline and identify explicit commonalities and differences between disciplines.
- Participate in collaborative pilots that establish the open-data and open-science end-to-end infrastructures.
- Act as a brokering clearinghouse instead of the core infrastructure manager for some of these services.

- Become a clear leader in helping formulate the data policy.
- Assist in the development of an evaluation plan to assess how well the data management plans and OSTI's services support the community.
- Develop cost models for manageable and cost-effective data solutions.
- Pursue a full integration of software and data together with a link to the published article.
- Play a similar role in the United States as the Data Curation Centre does in the United Kingdom.

OSTI would need to develop some significant additional expertise to play the coordinating and advising roles suggested above.

Giles expressed the hope that the Committee could get to agreement on this report soon. If significant changes are not called for, the report could be accepted today with amendments. He asked the Committee members for their responses to this draft.

Hey noted that many people in the audience had used OSTI services.

Sarkar suggested setting some metric about OSTI's success and that that metric be added to the recommendations.

Berzins noted that there seems to be a gap in computer sciences in the publication of scientific activities (e.g., codesign activities). He asked if there should be a mandate for outreach recommended in this report that also addressed the shifting of publishers' interests and policies. Hey responded that there is a shift in culture occurring, and the report encourages that. He had introduced such openness to Microsoft R&D. However, copyright restrictions are in force, and they complicate compliance. When NIH *required* public access, compliance went from 30% to 70%.

Sarkar noted that a dominant portal for computer sciences is the Association for Computing Machinery (ACM). Software is hosted in ACM's digital library. They provide *links* to copyrighted material. OSTI could adopt such a process. Hey pointed out that, for 12 months, PAGES would keep the version of record of a publication in its dark archive and make it available after 12 months if the publisher does not make it publicly available. This practice started in August 2014.

Petzold asked about multi-agency-funded researchers' linking their publications in the NIH archive to the DOE archive. Hey replied that the dialogue between the two agencies is a good start. They are trying to fulfill the legal requirement placed on DOE.

Reed complimented the Subcommittee on its hard work.

Giles went back to the findings, recommendations, and suggestions and asked for a review of them. Hey reiterated the Subcommittee's responses to the charge questions, including the assessment of PAGES. He brought up the question of how to position OSTI in regard to competitive and cooperating services. OSTI is talking with a lot of similar organizations about relevant topics.

Giles asked the Committee if it was ready to vote on acceptance of the report. Sensing agreement, he moved that the Committee accept the report. The vote was unanimously in favor; there were no abstentions.

The floor was opened for public comment. Biven noted that the pilot was established in July 2014, and the requirements went into effect last October.

Daniel Reed was asked to present a report from the Exascale Computing Initiative Subcommittee.

The Subcommittee had a kickoff planning teleconference (1) to formulate questions to be posed to DOE, national laboratory staffers, and computer science community members and (2) to develop a work stream.

DOE's exascale plans have been reviewed extensively. There have been community workshops, technical studies, strategic assessments, congressional hearings, and secretary of energy advisory board assessments for more than 7 years. The Subcommittee focused primarily on organization and management, and the technical issues and previous studies informed this assessment.

The Subcommittee strongly endorses the DOE plan for exascale computing development and deployment. Like any ambitious undertaking, DOE's proposed Exascale Computing Initiative (ECI) involves some risks. Despite the risks, the benefits of the initiative to scientific discovery, national security and U.S. economic competitiveness are clear and compelling. The Subcommittee believes the ECI is a well-crafted plan designed to meet DOE mission needs while also advancing broader national security and competitiveness goals.

The Subcommittee identified seven detailed recommendations:

1. Develop a detailed management and execution plan that defines clear responsibilities and decision-making authority to manage resources, risks, and dependencies appropriately across vendors, DOE national laboratories, and other participants.
2. As part of the execution plan, clearly distinguish essential system attributes (e.g., sustained performance levels) from aspirational ones (e.g., specific energy-consumption goals) and focus effort accordingly.
3. Given the scope, complexity, and potential impact of the ECI, conduct periodic reviews by a carefully constituted external advisory board.
4. Mitigate software risks by developing evolutionary alternatives to more innovative, but risky, alternatives.
5. Unlike other elements of the hardware/software ecosystem, application performance and stability are mission critical, necessitating continued focus on hardware/software co-design to meet application needs.
6. Remain cognizant of the need for the ECI to support data-intensive and computation intensive workloads.
7. Where appropriate, work collaboratively with other federal research agencies and international partners on workforce development and long-term research needs, while not creating dependences that could delay or imperil the execution plan.

In terms of the management and execution plan, DOE should

- Develop a detailed management and execution plan that defines clear responsibilities and decision-making authority to manage resources, risks, and dependencies appropriately across vendors, DOE laboratories, and other participants.
- Establish a leadership structure that operates below and in concert with the present, high-level leadership at DOE headquarters. This leadership structure's sole focus should be the exascale program.
- Develop a formal risk-assessment plan in concert with detailed execution planning.

The key performance goals are to attain a sustained performance of 1 to 10 ExaOPS at a power up 20 MW with a system memory of 128 PB.

At a point when many new technologies' components are still maturing, the ECI must not commit prematurely. If target numbers are publicized and shape activities prematurely, there is a danger that the ECI could be perceived as a failure for not reaching initial objectives. Innovation is needed to make this work. Failure to create a broad ecosystem will very likely disincentivize both the users and the vendors, and, as a result, will fail to leverage their mainstream research and development efforts, ultimately resulting in fewer technological advances and lower overall performance.

External reviews should be conducted by an advisory board. There must be a well-defined process to monitor technology developments, potential risks and benefits; careful co-ordination across stakeholders; and rigorous assessment of project priorities and directions. The primary rationale for the ECI is the new scientific discoveries and technical capabilities it will enable. DOE must quantify what that means, ensuring there are *credible application and discovery measures* for the success, or failure, of the ECI.

Software risks should be mitigated by developing evolutionary alternatives to innovative, risky alternatives. Exascale software development has two distinct goals: The first is allowing applications to execute at scale as quickly as possible, with minimal change. The second is shifting the software base to post-petascale architectures and ensuring broader uptake and use of exascale systems. Applications should have both an evolutionary and a revolutionary path to exascale execution.

A key issue regarding applications is that today's parallelism is "weak scaling." In the future, one must be able to exploit "strong scaling."

Unlike other elements of the hardware/software ecosystem, application performance and stability are mission critical, necessitating continued focus on hardware/software co-design to meet application needs. The ECI should identify a set of these mission-critical applications from its target domains (computational materials science, next generation climate models, stockpile stewardship) and make them yardsticks against which exascale systems are evaluated.

On the data-analytics side, extreme-scale science is driven by exponential technology advances producing massive data sources, big data, and big compute. It is important that DOE remain cognizant of the need for the ECI to support data-intensive and computation-intensive workloads. Modeling and data analysis are inextricably intertwined enablers of innovation and discovery; both draw on the same ecosystem of hardware and software technologies; and both are crucial elements of DOE's ECI. A new generation of data analytic tools and libraries are needed to aid in the interpretation and validation of the data generated from exascale applications.

In terms of interagency and international collaborations, DOE should, where appropriate, work with other federal research agencies and international partners on workforce development and long-term research needs, while not creating dependences that could delay or imperil the execution plan. The February 2015 ASCAC report identified the need for long-term partnerships among the stakeholders, including government agencies, academia, and vendors, to address these fundamental requirements.

More experienced application developers are needed. DOE should develop plans for interagency research collaborations and mechanisms to incorporate salient research results. More broadly, interagency research collaborations would expand and accelerate development of a highly trained and flexible workforce that is aware of, contributing to, and utilizing exascale systems. There are additional opportunities for bilateral and multilateral international collaborations to ensure development of consistent and interoperable software ecosystems and applications.

In summary, advanced computing's benefits are broad and deep in scientific discovery, national security, and the U.S. economy. Exascale computing is not a destination, but the next milestone in a journey. The Subcommittee strongly endorses the DOE plan for exascale computing development and deployment.

Chapman thanked the Subcommittee for its pragmatic approach. She asked if there were anything in the report about workforce development. Reid said that there was more in the report about expanding and accelerating the development of the workforce. There are interagency discussions going on about this issue. Also, one should look for bringing more people into this effort.

Williams asked how engaged this process was with the applications domain space. Berzins replied that the Subcommittee spent about 4 hours of its meeting on applications and had follow-up discussions with the national laboratories about challenges and opportunities in this area.

Crivelli asked for an expansion on technical tools. Reed responded that the report does not focus on one R&D path or another. It focuses on the plan to move the Initiative forward. The algorithmically driven data are not dealt with in the report nor is machine-learning technology. The Subcommittee notes that these aspects of data analytics must be explored and developed.

Giles asked if the pointing out of the need for new and novel applications was a new observation. Reed replied, no. It is a concatenation of prior discussions. One must have backward compatibility. Co-creation and other such topics must be part of the Management Plan. One must have applications to run.

Giles asked for a vote to accept the report subject to editorial revisions. The motion was accepted unanimously with no abstentions.

The floor was opened for public comment. Yelick said she appreciated the focus on science in the report. In regard to the revolution/evolution comment, she pointed out the need for rewriting CUDA code. Reed replied that the comment was meant to be general in nature. Evolution meant a complete application rewrite would not be needed; rather, the existing applications would be evolving. Sarkar added that “revolution” might not be a good term. Perhaps, “final.”

Cogliani noted that the use of external reviewers for the recommendations gives Congress a better sense of whether the project is on track and that there are no lower-cost alternatives.

A break for lunch was declared at 12:02 p.m.

Afternoon Session

The meeting was called back into session at 1:37 p.m.

Susan Coghlan was asked to provide an update on the Argonne Leadership Computing Facility (ALCF).

Current ALCF resources include the Mira, a 10-PF IBM BlueGene/Q with 48,000 nodes, 786,000 cores, 786 TB of memory, a 5D Torus interconnect, and a 26-PB general parallel file system. It is supported by Cetus for application design and development, Cooley for data analytics, and Vesta for system software. 60% of Mira’s time is devoted to INCITE [Innovative and Novel Computational Impact on Theory and Experiment] users, 30% to ASCR Leadership Computing Challenge users, and 10% to director’s discretionary allocations to incubate projects.

The future is Aurora, which is coming from the Collaboration of Oak Ridge, Argonne, and Livermore (CORAL), which was founded in 2012 to provide the next generation of DOE computer resources with architectural diversity. The resulting set of computers best serves the DOE missions. Argonne ended up with two new LCF systems: Theta and Aurora. Both have Intel Xeon Phi compute architecture and deep-memory architecture. This was the right path for the ALCF users because they already had many-core evolution, easy-to-port codes, well-balanced memory, and a robust and well-known user environment combined with some new Intel innovations.

Theta is a stepping stone to Aurora. It will be moved to data analytics when Aurora comes into production. It will be delivered in 2016 and have a system performance of more than 8.5 PF, more than 60 cores per node, up to 16 GB of bandwidth per compute node, 192 GB of DDR4 memory, 128 GB of SSD, and a small file system. It will use the Cray Aries Dragonfly topology interconnect [already used at the National Energy Research Scientific Computing Center (NERSC)]. On Theta, Mira users will have the same MPI+, a different Dragonfly but with advantages, increased vectorization opportunities, faster memory, more memory per node, and better single-thread performance.

Aurora will have a peak system performance of 180 to 450 PF, an application figure of merit that is 13 times that of Mira, a third-generation Intel Xeon Phi processor, more than 50,000 nodes, more than 7 PB of local memory, more than 30 PB/s of bandwidth, a second-generation Intel Omni-Path with silicon photonics system interconnect, more than 500 TB/s of bisection bandwidth, a burst buffer, an Intel Lustre file system with more than 150 PB capacity, and more than 1 TB/s of file-system throughput.

On Aurora, Theta users will have next-generation Xeon Phi, a similar number of nodes as on Mira, similar tiered memory and input/output, the new Intel Omni-Path interconnect with the same topology, a similar Cray software stack, and the same programming models.

A major goal is having *science* on day one (even prior to acceptance). The early-science program (ESP) is made up of a science team, ALCF stack expertise and dedicated time on the computers. The vendor also provides expertise and early access to software. The Theta ESP was selected in 2015 and starts runs in 2017 with six projects. The Aurora ESP project will be selected in 2016 and starts runs in 2019 with ten projects. The Theta ESP will have a kickoff workshop in August 2015 and a hands-on workshop in 2016. The early science dedicated access begins in January 2017 and ends in March 2017. The teams are coordinated to provide tools and libraries, such as performance tools. Many of these codes run on a variety of machines and architectures, so a team has been formed to improve data locality and thread parallelism, promote the use of portable libraries, use MPI+OpenMP 4.0 as a common programming model, and encourage portable and flexible software development.

Sarkar pointed out that there is a difference in functionality among the different versions of MPI+ and asked what version was being used. Coghlan replied that she did not know the exact version numbers. Maybe Version 3. Sarkar asked what migrating effort might be needed. Coghlan answered that no formal discussions of that topic had been held, but a lot of discussion had been held on the path. Best practices are being selected.

Williams asked about the resiliency of the system. He asked whether programs will be able to run at both leadership-computing facilities. Coghlan replied that it is not that they will be able to be run on both systems but that it will be easy to adapt code from one system to run on another. She did not know the details of resiliency.

Berzins asked how one deals with performance targets for systems that do not exist yet. Coghlan replied that targets and a shared-risk target are established, and dates are set for the targets. If one cannot make those targets, it is a no go. In performance, expectations are built on projections. Scaling and vectorization will help performance.

Rick Borchelt was asked to speak on science communication among DOE, stakeholders, and others.

A reporter will ask, what is the big picture? A scientist hears, what are the details of your specific project?

SC has adopted an approach to communications with the view that a communication should be planned as a life-cycle process, not a one-off activity, and should be integrated with the science. It should include ways to reach various stakeholder audiences where they already are. It should be aimed more at curation than creation; it should *focus attention*, not *throw more information out there*. Thus, communications should be viewed as a strategy, not a tactic. And it should be recognized that the general public is not SC's audience.

People used to surf on the web. There was so little information there, they were likely to run across whatever an agency had put out there. Then the content on the web became huge. In the future, communications cannot rely on consumers' finding what one wants them to be aware of simply through happenstance. Rather, communications must produce awareness, promote understanding, and increase support. This is not necessarily a stepwise process, and the presence of one element does not necessarily indicate the presence of another. These repeating cycles can also be thought of as

- Recruitment to build visibility and generate buzz,
- Maintenance to build connections and generate dialogue, and
- Retention to build consensus and generate support.

There is a spectrum of sophistication among Web users. Different people with different interests and information needs will come to a website. 90% will be skimmers who will never go deep into the subject matter; some will be dabblers who slightly break the surface; and some will be divers who go deep into the subject matter. SC wants to provide content for all three types of users.

There is more content than can be used. Instead of content being created, it must be curated. Stakeholders' content must be put where it can be seen and used. DOE has banners, features, news, and other categories of information that it curates on its own webpages. It also has partnerships with Eurekaworks (a product of the American Association for the Advancement of Science) and with Newswise. People pay to put their research on Eurekaworks. Newswise has about 5000 reporters.

There's no such thing as the general public. What matters to DOE are the Congress, the President, and the President's proxies. There are about 9000 people in the United States who actually touch on science policy. One does not need the *New York Times* or NBC to reach these 9000 people. They can be targeted as a group.

Giles asked how well ASCR communicates. Borchelt replied that ASCR does a good job of reaching its target audience. It has a way to go to reach Congress and the President. It is not done in the right context. Efforts need to be coordinated across DOE and focused.

Kathy Yelick was asked to report on programming models and environments.

A previous workshop had discussed hardware challenges, application challenges, and ecosystem issues. Summit and workshop reports gave an overview of the programming model stack, including applications, languages, and machine-specific abstractions.

Future generic node architectures will have lightweight cores; dynamic random-access memory (DRAM); memory stacks on the package; nonvolatile random–access memory (NVRAM) for a burst buffers and rack-local storage; "bulky cores" for standard processes; and integrated network interface controllers (NICs), which offer new opportunities.

Lightweight cores will have all or most of the system performance. They cannot be ignored. They need fine-grained parallelism but may not be powerful enough. The on-chip interconnect offers opportunities for performance, and new models of communication may be essential. The hardware is heterogeneous; portability and performance portability are challenging. There will be new levels of memory hierarchy, and they may be software controlled; locality and communication-avoidance are paramount. Performance variability may increase because of software or hardware control-clock speeds. Resilience will be paramount at such a scale; failures grow with the number of components and connections.

OpenMP is popular for its convenient loop parallelism, but it is too coarse (there is implicit synchronization between loops, which limits of parallelism and adds overhead) and too fine (larger chunks of serial work need to be created to minimize data movement). There is a lot of unnecessary synchronization today, producing serialization. The locality in OpenMP4 is, at best, computation-centric. Locality needs to be annotated for every loop. Architecture may converge, but Titan, Mira, and Edison represent three distinct architectures within SC. There is not performance portability across systems.

Major programming-model research is being conducted in ASCR on performance portability through compilers and autotuning, data locality in languages and libraries, less-synchronous directed-acyclic-graph (DAG) execution models, correctness, and resilience models and technology.

Three approaches are being taken to performance portability. The first approach is compiler-directed autotuning, which must analyze the code to determine legal transformations and then select the best optimized version. The second approach is domain-specific languages, such as Halide, which was developed for image processing and uses a schedule that is either auto-generated by autotuning with opentuner or is hand created by an optimization expert. Domain-specific languages (DSLs) are used to produce code for hierarchical memory, and Rose/PolyOpt is used to apply DSLs to large applications and collaboration on adaptive mesh refinement (AMR). The third approach is selected, embedded, just-in-time specialization, which uses a general optimization framework. It produces within 50% of the performance portability of hand-optimized code and is composed of 1400 lines of DSL-specific code.

Locality control can be conducted by tiling, in which data layouts can be used to improve locality and find parallelism. OpenMP allows a user to specify a number of data layouts; however, the code is different for GPUs versus CPUs, and several approaches are pursued.

Providing support for applications without locality, such as random access to large memory, enables a new class of applications. The Meraculous Assembly Pipeline reduces a human task that takes 44 hours to 20 seconds. Data fusion in UPC++ “fuses” observational data into a simulation.

One communication technique is to lower the overhead for smaller messages. Today one has a bottleneck at the node. MPI+X can communicate on one lightweight core. One can do a reverse offload to a heavyweight core, but one wants to allow all cores to communicate. Lightweight communication is more important with the lightweight cores of the future.

Several projects are looking at avoiding synchronization. One project, HPX Asynchronous Runtime performs on manycore. Another project, Legion Programming Model and Runtime, is running on the Titan and overlaps communication of CPUs and GPUs. A big question is how one specifies this in a higher-level program.

A lot of research is going on regarding application performance with tuned remote memory access programming in MPI-3. It deals with distributed hash tables, dynamic sparse data exchange, three-dimensional fast Fourier transforms, and the MILC [MIMD lattice computation] code.

A lot of time at the workshop was spent on technology-transfer paths. In languages, one technique is adoption into popular programming models, putting one-sided into MPI and locality control into OpenMP, and another technique is adoption by a compiler community. In compilers, one can leverage mainstream compilers, leverage another existing domain-specific language, or use small compilers for small languages.

The next phase of the program is to focus on application partnerships.

Berzins stated that there will be better science with these codes. Yelick agreed. There is no limit of opportunities in science or in the number of techniques to improve performance, opening new possibilities to researchers.

Crivelli asked if there were a plan to reach out to students. Yelick replied that many of these projects have been conducted by students. However, they usually do not know what performance computing is when they start.

Reed asked about resiliency and energy consumption. Yelick answered that performance management may be affected by energy management. Resilience mechanisms will have an effect on performance. That is why dynamic programming is necessary. Lighter-weight resilience is being looked at.

Chapman said that she did not remember a lot of discussion of tools. Yelick replied that there was a breakout session on tools like debuggers. These tools come with a performance cost.

Sarkar noted that people are relying on embedded DSLs. Yelick agreed. There is a fuzzy line on what is a DSL and what is not.

A break was declared at 3:31 p.m. The meeting resumed at 3:45 p.m.

James Sethian was asked to review work on applied mathematics for experimental science.

In 2025, DOE will be dealing with more data, more users, and more discovery. There will be a need for computational tools for analysis, data reduction, and feature extraction in situ and a need for post-processing for reconstruction, intercomparison, simulation, and visualization. These activities are interesting because the problems have not yet been mathematisized, there are no equations of motion, and there are deep connections between the science and math. These problems require new mathematics that bridge mathematical disciplines. Fortunately, applied mathematics is undergoing a profound transformation, breaking down the walls between continuous math, discrete math, analysis, probability and statistics, topology, algebra, and geometry.

The Center for Applied Mathematics for Energy Research Applications (CAMERA) was established at LBNL to build the applied mathematics that can accelerate scientific discovery at DOE experimental facilities. It is a joint ASCR–BES pilot project with the goal of building advanced mathematics to extract information from murky data and to help interpret experimental results; to provide on-demand analysis as results are being produced; to steer experiments and suggest optimal solutions; to decrease turnaround time and save money; and to extend the capabilities of existing and future experimental facilities. To do so, it needs to have experimental scientists and applied mathematicians work together to develop a common language, to build new mathematical models, to invent algorithms, and to build prototype codes. The goal is to test those codes on the "shop floor" interatively until those codes are solid and useful, embedding advanced mathematics in useable software tools.

CAMERA's original projects have been on pytchography, automatic image analysis, new methods for density functional theory, design and materials, and X-ray nanocrystallography. Software that the Center has released includes codes for flexible grazing-incidence small-angle X-ray scattering; a fast method for electronic-structure calculations; multi-GPU accelerated pytchography; image enhancement, filtering, segmentation, and feature extraction; and analysis/assembly of crystalline porous materials. The mathematics group does modeling and simulation of vertical-axis wind turbines, cell-cluster growth, virtual colonoscopy, industrial inkjet printing, draining in coal hoppers, and industrial foams. Its algorithms are being used widely in industry.

A major project has been Scalable Heterogeneous Adaptive Robust Pytchography (SHARP), which combines a high-precision scanning microscope with high-resolution diffraction measurements. It replaces a single detector with a 2-dimensional charge-coupled-device (CCD) array to measure intensity distribution at many scattering angles. The phase retrieval comes from recording multiple diffraction patterns from the same region of the object. Pytchography uses a small step size relative to illumination geometry to scan the sample. Diffraction measurements from neighboring regions are related through this geometry. Thus, phase-less information is replaced with a redundant set of measurements. Lots of pytchographic equipment and codes are used throughout DOE and worldwide. However, there is no convergence proof yet available for the method. There are problems with efficiency, scaling, and initial guesses, ultimately becoming an over-determined problem in high-dimensional space. A better starting guess is produced by viewing every pixel of every frame as a dimension, building a relationship network, constructing a graph Laplacian of the relationship network, and determining the largest eigenvector of the connection graph. SHARP is an open-source, downloadable package that provides an 80-times speedup with algorithms, a 30-times speedup with GPUs, and more than a 16-times speedup with distributed GPUs. Currently, the user interface starts processing at the end of a full scan; SHARP reads the data directly off the CCD.

Another application is QuantCT, a set of automatic image-analysis tools for micro-computed-tomography (micro-CT) for the analysis of the porosity, pathways, interior voids, etc. of physical structures. QuantCT uses the Mumford-Shah functional for image segmentation of two phases to find the interface that minimizes energy. It becomes a partial differential equation (PDE) transport method with level-set method. It extends the Mumford-Shah energy functional to a multiphase multi-interface with the Voronoi implicit interface method to allow simultaneous extraction of multiple structures in three dimensions. It uses our augmented topological descriptors to determine conductivity and channel pathways. It is currently available with a browser/computer at the Advanced Light Source (ALS), it is available as a FiJi plug-in, and the prototype source is downloadable.

A third application is fluctuation X-ray scattering for structure reconstruction. It is an extension of small- and wide-angle X-ray scattering. X-ray snapshots are taken below rotational diffusion times. Significantly more experimental information is produced than available with traditional techniques. It is a powerful technique for modern synchrotron and free-electron lasers. Going from a real space structure to fluctuation scattering data is straightforward. However, the reverse is difficult. A CAMERA collaboration has produced a new technique: M-TIP, which exploits multi-tiered iterative phasing and solves this inverse problem. The new method figures out the structures of objects that cannot be crystallized at a far higher resolution than previously available.

CAMERA’s products are receiving a positive response from the community and many software requests. New joint projects are starting up. Knowing what to build, how to build, and how to use it requires close-knit, coordinated teams with many different skills. With careful attention to mathematics and algorithms, one can build codes and software tools that can transform data into the information that users really want.

Hey noted the inclusion of students and graduate students and asked what career paths were open to them. Sethian said that there was a crying need for such people at the national laboratories. Academia uses them in mathematics and computer sciences departments. However, academia does not want one's life to be "the code," but the math, the science, or the material.

Robert Ross was asked to present an update on the DOE Storage Systems and Input/Output (I/O; SSIO) workshops.

The workshop was held December 8–11, 2014.

SSIO is everything from the low-level parallel file system and archival storage up to libraries that service the interfaces to applications. They bring with them new challenges like deeper storage hierarchies, increasing scales, complex topologies, demand for greater resilience, and new science workflows mandated by new R&D on data management and scalable system software. The goals of the workshop were to review SSIO requirements for simulation-driven activities, assess the state of the art, and identify priority research directions in SSIO. Actually, three workshops were held, one each on requirements gathering, cross-cutting computer science, and the SSIO meeting itself.

The goal of the first workshop was to define a set of use cases for future systems to inform the SSIO R&D community. Application representatives presented future requirements and answered a detailed set of questions. It was apparent that application developers already see data services as part of their applications.

The goal of the second was to identify critical SSIO requirements and points for coordination between SSIO and other computer science areas. Experts in related computer science areas presented their views on how their area intersects with SSIO. An analysis team commented on in situ data analysis and visualization I/O.

The third identified potential research direction in SSIO is for extreme-scale DOE science. Initial talks summarize findings of prior workshops and other recent activities. An open discussion was organized around five areas: hardware/software architectures for SSIO; metadata, name spaces, and provenance; supporting science data; integration with external services; and understanding SSIO systems.

In regard to hardware/software architectures, discussions were held on network technologies and topologies (where should the stack be placed?), solid-state storage in and near the high-performance-computing system (as opposed to the distributed storage of today), computing and storage, system noise and reliability, and autonomics. The findings that emerged were that storage hierarchy is increasing in complexity and that scientists need an integrated view of storage resources.

Metadata, name spaces, and provenance are popular topics these days and include metadata and alternative data stores, automating provenance capture, and connection to other services. The main finding was that new requirements for validation of results will change the role of metadata in DOE applications. New methods for capturing provenance and exploring data sets will be needed.

In supporting science data, topics included programming model integration, SSIO services in support of workflow, self-tuning libraries, and data abstractions. The findings were that scientists' productivity is tied to the ability to represent and interact with complex and specialized data (writing to bare metal). Also, alternative programming languages and an increased need for workflow support drive new SSIO research. Time needs to be spent on allowing applications to determine the proper data structure. Today's software is poorly configured. There is a big gap between abstractions and the actual data structures.

In regard to integration with external services, discussions covered the required scheduling and resource management, system monitoring (which should be left to a specific agent, not placed in applications), workflow and orchestration, and archival storage. Current SSIO designs are hindered by isolation from system-level resource management, monitoring, and workflow systems.

Discussions concerning understanding the behavior of storage systems and I/O in a production environment covered workload characterization (how applications are behaving), modeling and simulation, and designing for understandability. A major research priority is the need to improve the ability to characterize storage activities to model and predict the behavior of SSIO activities on future systems.

The workshop report came out in May 2015. A solicitation was issued on storage and I/O for extreme-scale science. The solicitation identified three themes: measurement and understanding, scalable storage software infrastructure, and new paradigms in SSIO. Proposals were due July 13, 2015.

An ongoing discussion on burst buffers is occurring. Applications use checkpointing, which produces bursts of activity. A buffer for those bursts opens up new ways to use the resource. The debate is where to put such buffers, what services should use them, and what interfaces should be provided.

Experimental and observational data bring new challenges. High-performance computing has a role, so SSIO also has a role. However, it is not known what the new applications will be or what services they will need.

Glotzer thanked Ross and the panel for their hard work.

Sarkar asked what the impacts would be from the SSIO perspective. Ross replied that the current generation of high-performance computers has problems that will be solved without significant changes. There is a big gap between peak and sustained performance. Applications are not given the power to address that gap.

Ronald Brightwell was asked to report on the Exascale Computing Initiative Runtime Systems Workshop.

The purpose and goals of the workshop were to review the state-of-the-art in runtime systems; to identify the challenges being addressed by current runtime-system R&D; to identify research questions that need to be resolved; to devise metrics, measures, benchmarks, and means for testing and evaluating runtime system prototypes; and to discuss an R&D roadmap that would result in one or more high-quality runtime-system prototypes.

The workshop was held March 11–13, 2015, in Rockville, Md., with 45 domain experts in high-performance-computing runtime systems. It included invited talks and breakout sessions on the architecture for future runtime system software, runtime system design, outstanding research questions, and a roadmap for the future.

In system architecture, discussions were held on governing principles of the strategy of computation, semantics and control strategy in the presence of asynchrony, the scope of the runtime system from system to node level, responsibilities and interfaces of the operating system and runtime system, basic requirements for runtime systems, information that compilers can provide to the runtime system, and metrics for testing and evaluating a runtime system.

The design of runtime systems has questions about how the runtime system manages memory resources, how the runtime system can use dynamic adaptive techniques, the role of the runtime system in correctness analysis, the role of the runtime system in managing parallelism, specific capabilities of the runtime system for resilience, the role of the runtime system in minimizing energy costs, and how to test and evaluate runtime-system design.

This workshop builds on the runtime-system summit activity in April 2014, where 11 attendees from X-Stack projects held a day-long meeting to brainstorm about requirements for an exascale runtime system. The goal was to develop high-level requirements, roles, and responsibilities for runtime systems to provide some context for generating a roadmap for future investments in runtime systems. Services a runtime system needs to provide include (1) interfaces between the runtime system and node- and system-level hardware abstraction layers, the operating system, and the programming interface and (2) mapping these interfaces to the existing runtime system.

A significant number of projects are going on in runtime-system R&D. The X-Stack program has played a key role in supporting runtime-system R&D for the extreme-scale. The X-Stack renewal enables engagement across projects in runtime-system prototypes, interfaces, and evaluation strategies. There are also runtime-system projects in NNSA and academia, as well.

The first problem is nomenclature. The definition of a runtime system is incomplete. There is a strong desire to understand responsibilities of the runtime system. Certain characteristics are defined: It is non-privileged; it runs in the applications space. It is ephemeral; it does not live beyond the application. It can manage hardware directly as long as isolation and protection mechanisms are provided. It interfaces to the node-level operating system (OS). And it may interface to the system OS and the enclave OS. However, any definition may be platform specific.

There was also a struggle with nomenclature with the execution model. It depends on what runtime services being provided. Runtime services should be able to be bypassed. With asynchrony, there is performance variability; some programming models embrace asynchrony; and everything needs to be lightweight. There is a relationship between the runtime system and the OS. The issue is services used by an application versus across applications. The OS should still get out of the way but enable the runtime system. In the relationship to the program model, the issues are what gets exposed and what gets hidden (transparency); the connection to services like data management, security, and performance monitoring; and the flow of information between an application and the runtime system. The questions in evaluation are, what are the metrics, and what about runtime system portability?

In runtime-system architecture, there are blurry lines between the runtime system above (the programing model) and below (the operating system). There is a lack of clear taxonomy. Requirements are needed from the top (application driven); there is a loss of semantic information all the way down the stack. Uniform quality of service requirements across programming models would be helpful. There are also issues about managing shared resources, dealing with elasticity, and resilience (a cross-cutting problem).

In runtime-system design, there are memory-system questions in translation; the need to support static, semi-static, and dynamic use of memory; how to differentiate between memory and storage; how memory is virtualized; introspection; reliability; and energy/power management.

Scheduling and resource management should also be able to do load balancing, latency hiding, etc. The toolchain needs to be co-designed with the runtime system. It needs attribution of performance bottlenecks, interoperability of different programming systems and runtime systems, and understood detailed decisions by the runtime system. Adoption is a good metric of success; more specific metrics are scalability, flexibility, portability, completeness, and ease of use.

In articulating the runtime-system ecosystem, an ecosystem model needs to be developed for runtime-system components; which runtime-system services are to be stand-alone and which are to be embedded into larger components needs to be determined; interfaces that are ready for a standardization process need to be identified; and it needs a process for transitioning runtime system software from research to production.

Just performance metrics alone are not wanted. Relative metrics are needed to evaluate research progress such as time to solution, time to solution with failures, time to solution with system variability, and time to solution under power/energy constraints. Metrics are also needed for runtime overhead, CPU overhead, and memory overhead. Runtime-system portability is also needed.

There are a number of questions about dynamic control: What does each runtime system layer or component control? How do layers coordinate toward goal-oriented optimizations? What managed resources need to be identified? How should ordination and optimization across layers be carried out? How will the backplane be used for communication between layers?

Runtime systems need to support resilience and *be* resilient. Runtime-system-based strategies include (1) task replication and migration and (2) fine-grain checkpointing. This is a critical challenge for extreme-scale computing.

New runtime-system layers must be done with application developers and system-software developers. DOE needs to partner with application teams and to disseminate the runtime-system R&D impact. Co-design should include system software, applications, and platforms.

The key takeaways from the workshop:

- A process to work through several issues needs to be defined.
- The workshop only scratched the surface.
- Crisp definitions are needed for the basic terms.
- A set of services needs to be agreed upon to organize discussions.
- There is a tension between the monolithic approach and interoperable components.
- Everyone wants control of the layers below them (including applications).
- A bi-directional flow of information is needed between layers.
- Better agreement is needed on what are “operating systems” and “runtime systems.”
- Interoperability is needed between different runtime systems.
- Are dynamic runtime systems capable or necessary for the exascale?
- An awareness is emerging of the ties between the runtime system and SSIO.
- The runtime system itself will need to be resilient.

A draft report is being produced.

Sarkar noted that, as technology has changed, the software stack has changed along with it. Brightwell replied that the current effort is to deal with how to get things out of the way that have been used in the past. Sarkar noted that adoption should not be an “all or nothing” thing. Brightwell agreed. There are different products on different timelines. Some products are long-term, and some are near-term. One needs to figure out how to characterize those and understand the ramifications for transitioning, which is difficult.

Reed asked how one gets convergence on some of these topics. Brightwell said that there had been lots of discussions on different methods. The incremental method of transition was looked at for the exascale. In the MPI Forum, it was a relatively simple thing; this is a much more complex thing. A process of evaluation and adoption needs to be decided on. There is an incremental path forward. One needs to figure out what needs to be solved and how to solve it.

Berzins said that the central point is that what is being looked for are simple ways that will allow decomposition to work at the exascale without producing delays. Brightwell agreed. Over-decomposition is an approach that people are thinking about. Heterogeneity is another factor to be dealt with.

Robert Voigt was asked to give an update on the Computational Science Graduate Fellowship (CSGF) Program.

The CSGF program provides stipends of $36,000 per year for 4 years, full tuition and fees, professional-development support, practicum support (living expenses and travel), and an annual program review. Fellows must complete their program of study and practicum within the required time, make satisfactory progress on their thesis research, stay in good standing with their graduate program, attend annual program review, submit a renewal request each year, and provide status updates. The program of study must demonstrate breadth and include graduate work in science or engineering, computer science, and applied mathematics. All fellows must participate in a 12-week research practicum at an approved DOE site. The practicum must be completed within the first 2 years of the program and be approved in advance.

After staying stable or increasing slightly from 2004 to 2011, the NNSA funding was zeroed out in 2011. As a result, SC money had to be used to cover out-year expenses of participants, so the number of new fellows had to be cut. In the meantime, annual tuition costs increased. In April 2015, an additional $1.5 million was contributed by SC. However, both SC and NNSA funds are now at historic lows. The FY16 appropriation bill for SC has $8 million for the CSGF program in the House version and $10 million in the Senate version. NNSA plans to continue support at $1.5 million.

The program is managed by ASCR and NNSA under contract with the Krell Institute. There is a steering committee of nine people from academia, national laboratories, industry, and program alums.

In the 2015 application process, a program of study must be proposed that includes two courses in each of science/engineering, computer sciences, and mathematics; a course in parallel computing; and the statement about programming languages. Three essays are involved: field of interest, program of study, and role of high-performance computing in the research area. Three references must be submitted along with a list of technical and extracurricular accomplishments. Undergraduate seniors, first-year doctoral, and master's degree students are eligible along with employed individuals.

In 2014 and 2015, the application opening dates were delayed because of delays in funding. This resulted in delayed start dates, as well. In 2013, 2014, and 2015, there were 124, 103, and 82 undergraduate applications, respectively. The same years, the numbers of fellowships awarded were 10, 23, and 11.

This year, 12 teams of two people each screened the 164 initial applications. Each team read about 15 applications and selected about 8 for further consideration. This process was completed in February. In March, two groups of six reviewers each cut the number of applications be considered to 33. The full committee of 12 spent a full day discussing the remaining applications, resulting in nine finalists. Supplemental funding allowed an additional two applicants to be selected.

In the early years of the CSGF, biologists predominated among the fellows. Today, the fellows are fairly evenly distributed among the disciplines of biology, physical sciences, engineering, and computer sciences and mathematics. Since inception of the program, Massachusetts Institute of Technology, Harvard, Stanford, Princeton, and California Institute of Technology have been the leaders in producing fellows; and LBNL, ANL, Lawrence Livermore National Laboratory (LLNL), and ORNL have been the leaders in hosting fellows for their practica. Many of the fellows end up being employed at a national laboratory.

The annual program review was being held at the same time and in the same building as this ASCAC meeting.

Giles asked what the continuing resolution might mean for the CSGF. Voigt answered that the continuing resolution will let the program go forward at $4.5 million and allow more planning and better scheduling.

Berzins asked what level would be optimum to affect the workforce. Voigt responded that the program could handle a doubling of funding to cover 40 students per year. There are a lot of students interested in high-performance computing.

The floor was opened to public comment. There was no further comment.

Giles announced that this was his 45th and last ASCAC meeting; he is retiring on September 30, 2015.

Binkley said that, from the DOE side, it has been a total pleasure to work with Roscoe Giles, who provided great leadership and behind-the-scenes work, such as congressional testimony, while being a great teacher. [Round of applause.]

The meeting was adjourned at 5:48 p.m.

Respectfully submitted,
Frederick M. O’Hara, Jr.
ASCAC recording secretary
Aug. 9, 2015