Draft Minutes
Advanced Scientific Computing Advisory Committee
December 9–10, 2015
American Geophysical Union
Washington, D.C.

ASCAC Members Present

Keren Bergman
Martin Berzins
Vinton Cerf
Barbara Chapman (via telephone)
Jacqueline Chen (via telephone, Wed. only)
Silvia Crivelli
John Dolbow
Jack Dongarra (Wed. only)
Thom Dunning
Susan Gregurick
Anthony Hey
Gwendolyn Huntoon
David Levermore (Wed. only)
John Negele (via telephone)
Linda Petzold
Daniel Reed (Chair)
Vivek Sarkar

ASCAC Members Absent

Sharon Glotzer
Juan Meza
Dean Williams

Also Participating

John Steve Binkley, Associate Director, Office of Advanced Scientific Computing, Office of Science, USDOE
Christine Chalk, ASCAC Designated Federal Officer, Office of Advanced Scientific Computing, Office of Science, USDOE
Susan Coghlan, Deputy Division Director, Argonne Leadership Computing Facility, Argonne National Laboratory
Kory Contreras, Office of Human Resources, USDOE
Richard Gerber, Senior Science Advisor, National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory
Roscoe Giles, Department of Electrical and Computer Engineering, Boston University
Barbara Helland, Director, Facilities Division, Office of Advanced Scientific Computing, Office of Science, USDOE
Brian Hitson, Director, Office of Scientific and Technical Information, Office of Science, USDOE
Alexander Larzelere, Director, Modeling and Simulation Energy Innovation Hub, Office of Nuclear Energy, USDOE
Paul Messina, Senior Strategic Advisor, Argonne Leadership Computing Facility, Argonne National Laboratory
Lucy Nowell, Program Manager, Data and Visualization, Office of Advanced Scientific Computing, Office of Science, USDOE
Frederick O’Hara, ASCAC Recording Secretary
Kalyan Perumalla, Senior Researcher, Computational Sciences and Engineering Division, Oak Ridge National Laboratory

John Shalf, Chief Technology Officer, National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory
Edmund Synakowski, Associate Director, Office of Fusion Energy Sciences, Office of Science, USDOE
Erin Szulman, Policy Advisor to the Chief of Staff, Office of Science and Technology Policy, EOP
Deneise Terry, Oak Ridge Institute for Science and Energy

Wednesday, December 9, 2015
Morning Session

Before the meeting started, the new Committee members were sworn in by **Kory Contreras** of the Office of Human Resources of the U.S. Department of Energy (DOE).

Deneise Terry made safety and convenience announcements and announced that the meeting was being recorded.

The meeting was called to order by **Chairman Daniel Reed** at 8:34 a.m. He had each member introduce himself or herself, and he introduced the four new members of the Committee.

Steve Binkley was asked to report on the activities of the Office of Advanced Scientific Computing Research (ASCR).

ASCR is located in DOE's Office of Science (SC), which supports 47% of the U.S. federal research in the physical sciences. ASCR focuses on mathematics research, computer science research, Scientific Discovery through Advanced Computing (SciDAC) partnerships with other offices and programs in SC, the push from petascale to exascale computing, facilities to supply computer capabilities to the user community, and the 23-year-old Computational Science Graduate Fellowship postdoctoral program. The management and exploitation of large scientific data will be very important in the next 6 years.

At the time of this meeting, ASCR's federal budget for 2016 to 2017 was before Congress and therefore embargoed. ASCR's FY15 enacted appropriation was $541 million, and its FY16 president's request was $621 million, a 14.8% increase. In comparison, the SC budget increased 5.4% from its FY15 appropriation to its FY16 request; however, its $5 billion annual budget still shows good support for science by Congress. Funding for the ASCR exascale effort shows continued support, as does that for the National Nuclear Security Administration (NNSA) effort, for which an additional $14 million was requested for FY16.

The DOE *Quadrennial Technology Review* (QTR) was published in September 2015. It examines the status of the science and technology that are the foundation of the nation's energy system, together with the research, development, demonstration, and deployment (RDD&D) opportunities to advance them. It was called for by the President to analyze government-wide energy policy. It frames trade-offs that all energy technologies must balance and guides investment in RDD&D. In the past 4 years, changes in the energy landscape include fracking, increased vehicle gas mileage, and increased deployment of wind and solar energies. An input-output analysis of U.S. energy use showed a growth from 88 quads in 1980 to 98.3 quads in 2014. Significantly, the percentage of energy wasted has been reduced. SC and ASCR are featured in Chapter 9 of the QTR: Enabling Capabilities for Science and Energy. The chapter highlights a lot of opportunities for the DOE facilities to enable science. The same can be said for ASCR in computing facilities, networking, modeling and simulation, etc.

In DOE, Cherry Murray has been nominated to lead SC. She has had confirmation hearings; a full Senate vote is still needed.

In ASCR, Ceren Susut-Bennett returned in July from a four-month detail to the National Science Foundation (NSF). Karen Pao, a program manager in applied math, left federal service in August. The Office is currently recruiting to fill two vacant applied-mathematics program-manager positions. Barbara Helland has taken the lead for the Exascale Computing Initiative to get it executed on budget and on schedule. William Harrod is leading the effort to develop a long-range plan for ASCR research; this planning activity will be reviewed by ASCAC at its March meeting.

In facilities, the Lawrence Berkeley National Laboratory (LBNL) Computational Research and Theory Building was dedicated on November 12, 2015. It was funded by the University of California at Berkeley. The European extension of the Energy Sciences Network (ESnet) has been supporting the Large Hadron Collider (LHC) Run 2 since commencement in June 2015.

DOE's Exascale Computing Initiative is a partnership between ASCR and NNSA. It includes hardware, software, applications, large data, and the underpinning applied mathematics and computer science. It supports the DOE's missions in national security and science. Extreme-scale computing cannot be achieved by a business-as-usual evolutionary approach. The key performance goals are a performance of a sustained 1 to 10 ExaOPS, a power of 20 MW, and a system memory of 128 to 256 PB.

During the past year, a broadening of the reach of exascale science applications has occurred. The initial call for input from all 17 national laboratories was released on May 31, 2015, to identify potential applications that could deliver new science capabilities on exascale systems; 126 application white papers were received. On September 15, 2015, the National Institutes of Health (NIH), NSF, and DOE issued a request for information (RFI) to identify scientific research topics that need high-performance computing capabilities that extend 100 times beyond today's performance on scientific applications. By the Nov. 13 deadline, 114 responses were received, and they are now being reviewed. An ASCAC report identified the top 10 technical challenges for the exascale. These challenges are now being addressed in planned research.

The exascale initiative will follow established DOE review and decision protocols for its execution. A project office has been established at Oak Ridge National Laboratory (ORNL) with representation from the major participating laboratories. An Integrated Project Team (IPT) has been established and is refining the work breakdown structure (WBS) and is preparing required project documentation. A top-level WBS activity has been established to develop and implement exascale applications on the basis of an enterprise-wide request for information. Lead people for the top levels have been named. The Excascale Computing Initiative timeline calls for procurement between 2017 and 2020. Critical-decision (CD) steps have been set for establishment of the need, spelling out the conceptual design, identification of needed procurements, conduct of research and development, and operation of the facility.

In August 2015, ASCAC conducted an Exascale Computing Initiative (ECI) review. It found that, although the proposed ECI involves some risks, the benefits of the initiative to scientific discovery, national security, and U.S. economic competitiveness are clear and compelling. ASCAC made some recommendations, and ASCR is addressing and responding to each of those recommendations. ASCR is developing a detailed management and execution plan (the Exascale Computing Plan or ECP); implementing an advisory board; scheduling independent project reviews and periodic reviews; incorporating appropriate activities in the ECP; pursuing collaborations with other federal agencies via the National Strategic Computing Initiative

(NSCI); implementing evolutionary as well as more innovative but untested alternatives in application development; defining essential system attributes in the ECP conceptual design; implementing application performance and stability in the ECP conceptual design; and remaining cognizant of the need for the ECI to support both data-intensive and computation-intensive workloads.

Industry, academia, and DOE are looking at advanced concepts in computing, such as quantum computing and neuromorphic computing. In 2015, DOE conducted a workshop on neuro-inspired computational elements. Also in 2015, DOE issued a report on machine learning and understanding for intelligent extreme-scale scientific computing and discovery. A nanotechnology-inspired grand challenge was announced for future computing to create a new type of computer that can proactively interpret and learn from data, solve unfamiliar problems with what it has learned, and operate with the energy efficiency of the human brain. A joint ASCR–Basic Energy Sciences (BES) Neuromorphic Computing Study Group held a roundtable discussion in October 2015 in Gaithersburg, Maryland. Its purpose was to evaluate both advanced materials and scientific computing research opportunities to support development of a new paradigm for extreme and self-reconfigurable computing architectures that go beyond Moore’s Law and mimic neuro-biological architectures. Approximately 15 experts in emerging computer architectures and related materials science attended along with observers from across the national laboratories. A summary report is being developed.

Dongarra asked what part of the budget goes into the ECI, and what is left for the Office. Binkley said that those numbers are still being debated by Congress, and the budget details are currently embargoed. The answer to the question will be given at the March meeting. The Office has been careful to plan funding for the research activities needed in the field across SC. Once the transition is made to an established exascale computing program, planning funding for R&D will get serious.

Edmund Synakowski was asked to review the role of high-performance computing in the Office of Fusion Energy Sciences (FES).

This is an unprecedented era of transformation with the benefit, burden, and promise for making life better for the globe. For the first time in history, people have an understanding of the linkage between quality of life (i.e., life span) and energy availability, the impacts on the globe of their efforts to improve life quality, and the range of possibilities for the path forward. The scientific community has also come to understand that fusion energy, in step with high-performance computing, can be a transformative, clean energy source enabled by frontier science.

The aspiration of fusion science is to create a star on Earth with a hot plasma in which isotopes of hydrogen collide hard and often enough to fuse into a helium nucleus and a neutron with a large net energy gain and self-heating plasma. Experiments in the U.S. to date have used 40 MW of power to heat the fusion fuel to generate 10 MW of fusion power. Fusion is ready for a major step in demonstrating net fusion power: ITER [formerly the International Thermonuclear Reactor Experiment].

Fusion science is exciting and broad. A great intellectual framework has been developed, with rich and complex interplay. Energy-directed research spawns great science on a scale running from feet to lightyears.

All fusion reactors require a burning plasma. The key challenge is to confine the hot and dense plasma while it burns. One needs a high-density plasma, an extended confinement time, and a certain temperature range (about 250,000,000 K). ITER is designed to get into the regime of $n\tau_E T \geq 5.10^{21}$ m^{-3} s keV, near the fusion-reactor regime. Fusion's progress rivals that of computer chips. However, it still has not achieved the critical triple point. There have been many problems, like turbulence. These challenges are of the highest order.

ITER is being built in the south of France by an international partnership representing more than one-half of the world's population. It can be a truly transformative scientific instrument. Such a reactor is a candidate for baseload power.

The airline industry achieved an order-of-magnitude cut in the prototyping of wing designs with modeling and simulation, a tremendous savings in time and cost. Fusion science must do a similar thing.

FES research is carried out at 53 universities, 12 businesses, and 10 national laboratories. Foundational burning-plasma science includes several SciDAC endeavors and the employment of high-performance computing.

The aspect ratio is central to much of toroidal-confinement science. It is also a driver in the ultimate cost and viability of a fusion power plant. The DIII-D [Doublet III tokamak, version D] and NSTX-U [National Spherical Torus Experiment Upgrade] machines have a significant impact on the ITER design and operations, and both operate in a highly collaborative environment. The physics of sustaining and driving the current in a plasma is the focus of FES's R&D.

The Tokamak Fusion Test Reactor held the record of the highest man-made temperature. However, in the data produced, theory did not capture the reality of the situation. The change of confinement from L-mode to supershot challenged theorists to develop a predictive model. Their results indicated that the T_i turbulence data were wrong, and they were right. When turbulence was correctly measured, the experimental data came right into line with the theoretical predictions.

The challenge is that the physical processes in a tokamak discharge span multiple time and spatial scales. The typical time scales in a next-step experiment cover a wide range from 10^{-10} s to multiple seconds. Partnership with ASCR has increased computational performance by a factor of 4. Now, details of the turbulence responsible for cross-field transport of heat and fuel are being computed and measured, and boundary-heat-flux simulation is enabling assessments of conditions at the wall of the tokamak, where a heat transfer of 10 to 20 MW m^{-2} must be managed for months at a time.

There has been significant progress in understanding how the fusing plasma interacts with the material boundary. Helium bubbles are detrimental to plasma-facing materials, such as tungsten. Understanding how helium bubbles form and grow is important for predicting the large-scale material response to the extreme fusion environment. All relevant kinetic processes must be accounted for as well as how these processes enhance the interaction of the hot helium bubble with the local microstructure of the boundary-wall material. These processes need to be understood and modelled, and the models need to be validated and coupled with all other processes.

The FES Strategic Plan places a strong emphasis on high-performance computing. The major themes of the FES Strategic Plan are:

- Massively parallel computing is an emphasis.
- Materials science will provide the scientific foundations for greatly improved plasma confinement and heat exhaust.
- Research in the prediction and control of transient events will provide greater confidence in machine designs and operation with stable plasmas.

- Discovery in plasma science will address frontier science issues underpinning great mysteries of the visible universe and will help attract and retain a new generation of plasma/fusion science leaders.
- FES user facilities will be kept world-leading.

The Strategic Plan has just been delivered to Congress

Recent efforts continue a strong tradition of planning and working together. The challenges are significant. Joint program review meetings and workshops shape future directions and provide unique opportunities to exchange information with all stakeholders.

A vision for integrated extreme-scale simulations has been developed. Computational and enabling technologies in integrated fusion simulations require high software integration and performance in multiphysics coupling, uncertainty quantification, data management, multiscale coupling, numerical optimization, and data analysis and assimilation.

Superconducting magnets are essential for fusion reactors. A proposal to develop a capability in developing and producing superconducting magnets in the United States was not pursued. Republic of Korea took the lead in the field. Now there is a collaborative operation that exploits remote operation through communication networks.

The ITER construction is progressing. Many components are built onsite and installed in the facility. Things will go nowhere without burning-plasma science, and ITER is crucial for understanding and furthering burning-plasma science into the long-pulse and high-power regimes.

The ITER is not necessarily the best design for a burning plasma reactor. It will produce scientific knowledge from which the processes can be modeled and simulated to determine that optimum reactor design.

Berzins asked about the role of computing in scale bridging. Synakowski replied that the boundary conditions were very important and they control the edge temperature. Lithium was used as a getter on the wall, hugely increasing the temperature there and increasing the bulk plasma temperature to a new world record. Computing was critical in recognizing that this was possible and how it occurred.

Hey asked about the role of commercial startups. Synakowski answered that, if the production costs can be shrunk by lowering temperature and pressure and if size can be shrunk by identifying and applying control physics, fusion can become competitive.

Levermore noted that a breakthrough in modeling for fusion energy has furthered the science at Lawrence Livermore National Laboratory. A major lesson was that the hardware, software, and algorithms have to work together.

Gregurick asked where the biggest payoff would be in altering the ITER design. Synakowski responded that the design of ITER is set. Materials can have an impact on its later life. A reduced model for ITER experimental conditions would be very important. Data management will be important, as well.

Sarkar pointed out that cooperation among DOE, industry, and academia are important and asked what types of joint programs would be helpful. Synakowski said that subcontracts to universities have been used. There will be more competition for those subcontracts coming from the national laboratories. Professors and deans need to understand that national laboratories are the places that academia should use to keep faculty and students engaged in the cutting edge of science. There are many aspects of fusion and material sciences where universities can do a better job than the national laboratories can. People must work internationally, also. One must go where the research opportunities are. The Optimized Accelerator, which is just coming online in

Germany, is developing highly complex magnetic coils, and U.S. students there play a major role in determining the direction of the project.

A break was declared at 11:06 a.m. The meeting was called back into session at 11:19 a.m.

Erin Szulman was asked to give an update on the National Strategic Computing Initiative (NSCI).

The NSCI was created by Executive Order 13702 on July 29, 2015. The NSCI is a whole-government effort designed to create a cohesive, multi-agency strategic vision and federal investment strategy, executed in collaboration with industry and academia, to maximize the benefits of high-performance computing for the United States.

In 2010, the President's Council of Advisors on Science and Technology (PCAST) recommended that NSF, the Defense Advanced Research Projects Agency (DARPA), and DOE invest in a coordinated program of basic research on architectures, algorithms, and software for next-generation high-performance computing systems, including efficiently analyzing vast quantities of both numerical and non-numerical data.

The lead agencies in the NSCI are the Department of Defense (DoD), DOE, and NSF. The foundational R&D agencies are the Intelligence Advanced Research Projects Activity (IARPA) and National Institute of Standards and Technology (NIST); and the deployment (application) agencies are the National Aeronautics and Space Administration (NASA), National Oceanic and Atmospheric Administration (NOAA), Federal Bureau of Investigation (FBI), Department of Homeland Security (DHS), and NIH.

The strategic objectives listed in the Executive Order are a capable exascale; coherence between the technology base used for modeling and simulation and that used for data-analytic computing; a post-Moore's Law era; an enduring national high-performance computing ecosystem (networking technology, workflow, software, and workforce development); and an enduring public–private collaboration.

To date, the Executive Order has been issued; an NSCI private roundtable has been held at the White House to gather perspectives; the NSCI Executive Council has held an inaugural meeting; a request for information (RFI) on science drivers for a capable exascale has been issued; a 2-day White House NSCI workshop has been held; and an implementation plan has been delivered. Next steps include the development of the FY17 budget, further private-sector engagement opportunities, and an annual plan.

Related White House initiatives include the Materials Genome Initiative, Advanced Manufacturing Initiatives, National Nanotechnology Initiative, Brain Research Through Advancing Innovative Neurotechnologies Initiative (BRAIN Initiative), Precision Medicine Initiative, National Big Data R&D Initiative, and National Photonics Initiative.

The October NSCI Workshop developed several ideas that could potentially inform the NSCI implementation:

- The evolutionary path for high-performance computing is more uncertain than during the previous decades.
- A diversity of approach is key in the short term.
- NSCI must accommodate a breadth of choice.
- Clouds could be a viable model for NSCI broad deployment.
- Deeper engagement with the industrial (non-computing) sector will be key.

NSCI should strive for the convergence of numerically intensive and data-intensive computing, keep the United States at the forefront of high-performance computing capabilities, streamline high-performance computing application development, make high-performance computing readily usable and accessible, and establish hardware technology for future high-performance computing systems.

Berzins asked if there were any attempt to involve workforce development and education in the architectural development effort. Szulman said that an effort was being made to address that issue. Feedback is coming in from academics about where the gaps are. An attempt will be made to fill those gaps.

Dunning noted that industry does not develop applications anymore. They rely on national laboratories and universities. He asked if there were an effort to foster development of the new applications that will be needed. Szulman answered that the initiative is trying to (1) get agencies to identify the breadth of applications needed; (2) get industry to do the same thing from its perspective; and (3) establish lines of communication among industry, government agencies, and academia. Dunning said that many universities have partnerships with industrial sectors that could be exploited. Industry needs to be told where high-performance computing can be employed in their organizations. Szulman responded that the initiative is trying to recruit people who can do that. Dunning pointed out that national laboratories and NSF centers have a lot of experience in this area.

Levermore said that the bottleneck is the people on board. The best way to get people on board is to define education and training very broadly. Szulman agreed that they should be thinking along those lines.

Dunning said that academia has to learn how to integrate high-performance computing into its curricula.

Reed pointed out that new ideas take off when old scientists die.

Messina suggested rephrasing the comment on the need for software development. Today’s systems are not being re-implemented. Szulman said that what she had meant to say was that software development will become easier but will be part of the infrastructure. Dunning stated that one needs good front ends for the more complicated software. Hey added that the complexity must be hidden behind the user interface.

The meeting was adjourned for lunch at 11:47 a.m.

Wednesday, December 9, 2015
Afternoon Session

The meeting was called back into session at 1:44 p.m.

Brian Hitson was asked to present the response of the Office of Scientific and Technical Information (OSTI) to the report of ASCAC’s Committee of Visitors (COV) to OSTI.

He expressed gratitude to the COV, ASCAC, and DOE for their support and aid to OSTI.

In Patricia Dehmer’s 2014 charge letter, ASCAC was asked to determine if OSTI products and services are best in class, whether they met customers’ current and future needs, what OSTI’s national and international standing was, and where it must be a clear leader. Additional questions posed were:

Is the mission statement sensible in light of the statutory authorities?

Is OSTI organized and staffed to accomplish today’s mission?

Are the current and planned products and services the correct ones?

What suggestions would the COV make for the next steps?

In March 2015, Hitson provided an overview briefing to ASCAC. In May, the COV, chaired by Tony Hey, performed an on-site review at OSTI in Oak Ridge. In July, Hey presented the

COV report summary to ASCAC. And in September, the ASCAC Chair transmitted the formal report to Dehmer.

The COV observed that progress had been made in addressing the 2009 COV recommendation to focus resources on DOE R&D results and in fixing "leaky pipes" and a lack of comprehensiveness in scientific and technical information (STI) submissions. It pointed out that OSTI has exhibited leadership in implementing new public-access requirements (DOE PAGES Beta) in partnership with NSF and DoD to minimize the submission burden. Several of OSTI's products were seen to be best in class, but the Energy Science and Technology Software Center (ESTSC) was not considered to be best in class. OSTI's primary customer groups were seen to be the public, librarians, and researchers, with OSTI effectively reaching the first two but needing to better understand the needs of researchers. To researchers, OSTI services seem cumbersome compared with existing domain-specific solutions. Researchers also see a need for more integration of OSTI products and improved user interfaces. The COV commended the OSTI team's overall enthusiasm, competency, innovation, and adaptation to evolving technological trends but cautioned that a change in the mix of technical expertise and skills will be needed.

The COV recommended that OSTI

- Initiate a vigorous outreach program with the national laboratory researchers;
- Work with the DOE research community to reinvent the ESTSC software service;
- Work with the national laboratories to identify "researcher champions" who can work with the scientific and technical information program (STIP) community to strengthen the link to researchers;
- Continue toward a unified user environment with a limited number of clearly delineated, non-redundant tools and develop a master plan for future development and areas of expansion through community input; and
- Address publication content gaps.

The COV recommended that SC

- Promote a successful implementation of the public-access requirement issued by the Office of Scientific and Technical Policy by expressing a measurable expectation in the national laboratories' annual Performance Evaluation and Measurement Plans (PEMPs) and
- Define a useful role for OSTI and the STIP management team in managing DOE data.

OSTI has adopted four strategies to address these recommendations:

1. Strengthening ties to DOE researchers by holding a series of results-oriented workshops/listening sessions at national laboratories to address such questions as where and how do researchers use STI in their workflow and how can OSTI's STI products or content be more useful to researchers and by holding another series of workshops focused on specific STI types (data, software, etc.).
2. Enhancing product cohesiveness and comprehensiveness by reinventing the software service and integrating it with other STI types in tune with researchers' workflow needs; by developing enhanced focus-group processes and more granular metrics to understand user behavior within products; by defining and implementing a unified user environment as a content environment where diverse but linked forms of STI are seamlessly available; by applying a numerator/denominator (i.e., number of articles received/all DOE publications in the *Web of Science*) comprehensiveness model to public access; and by other secondary actions.

3. Implement public access by establishing goal/objective language related to public access support and submission of accepted manuscripts into the laboratories' FY16 PEMPs; by establishing a notable outcome addressing progress in public-access implementation in SC laboratories' FY16 PEMPs; and following up to provide specific guidance and examples of successful implementation.
4. Through national-laboratory and community-specific workshops, assess and characterize DOE needs related to the COV's six suggested roles for OSTI in DOE's data landscape and, with SC approval, integrate the resulting new goals and strategies into OSTI's strategic plan and budgeting/staffing; explore providing institutional and operational data-management support to SC's Working Group on Digital Data (SCWGDD) and the DOE WGDD; leverage the Data ID service and E-Link supplemental-material metadata to enable linking of publications, software, and data; and identify and obtain new data and software-management skills/expertise.

The COV's six suggested roles for OSTI in data management are

- Collecting digital versions of tables, graphs, and images from papers;
- Developing better solutions for linking data and software to publications;
- Helping identify data needs by discipline to identify explicit commonalities and differences between disciplines;
- Participating in collaborative pilots that establish the open data and open science end-to-end infrastructures;
- Assessing how well the data management plan (DMP) and OSTI services support the community; and
- Developing cost models for manageable and cost-effective data solutions.

In conclusion, OSTI's goal is to be best in class; the COV's work helps immensely. DOE-researcher needs will shape the "unified user environment." Public access is an area in which OSTI must be a clear leader to fulfill its mandated responsibilities. OSTI appreciates ASCAC's and the COV's efforts and looks forward to continuing to work with the COV Subcommittee in helping to shape the future of OSTI.

Dongarra asked if OSTI had any interaction with the Research Data Alliance (RDA). Hitson replied, yes and applauded RDA's focus on specific results-oriented projects. Hitson also noted OSTI's interaction with a number of other data-related groups, including CODATA and DataCite. Hey noted that academic librarians on the COV had called attention to the need to get publishers to make data available within 12 months. He asked what mechanisms were available for monitoring openness. Hitson said that the obligation for fulfilling public access is on DOE labs, grantees, and their authors, and that OSTI is leveraging the existing infrastructure of DOE's Scientific and Technical Information program, which incorporates 40,000 items into OSTI per year. DOE-funded authors must now submit metadata or full publications to OSTI or to a laboratory depository library. OSTI will use these submissions to fulfill public access. The publishing community is serving to complement this system, but DOE isn't relying on publishers to fulfill public access requirements. There are automated systems to track DOE publications. OSTI also does random manual checks to make sure the automated system is not missing anything.

Crivelli encouraged Hitson to visit all the national laboratories and the sooner the better. She asked what constituted the "low-hanging fruit" in integration. Hitson replied that OSTI's management will visit as many laboratories as permitted by resources. There are a lot of interdependencies among the strategies, and the workshop feedback will serve to shape actions across all of the strategies. We've pursued and implemented some low-hanging fruit in the areas of software search and public access performance metrics, and we expect the workshops to help shape longer-term changes in OSTI's product mix and features. The whole process will probably take 18 to 36 months.

Gregurick noted that the discussion was focused on national laboratories' products and asked whether this effort will be expanded to all DOE-funded research. Hitson answered that the national laboratories have 60% of funding and produce around 50% of journal articles. The national laboratories have administrative methods that can be leveraged to increase the comprehensiveness of submissions. DOE has also modified grant terms and conditions to specify accepted manuscript submission requirements.

Petzold asked Hitson if he had any thoughts of collaborating with other federal funding agencies so submission is not so time-consuming. Hitson pointed out that OSTP had held a meeting a week before this meeting to address that very issue. The goal is one submission, many agencies. One issue to resolve is that each agency has unique metadata needs, such as funding identifiers (e.g., contract and grant numbers).

Reed expected that universities will adopt their own depositories, so they will be held harmless. Two groups that represent university interests are the Association of Public and Land-Grant Universities and the Association of American Universities. It might be good to confer with these institutions about this issue. Hitson replied that that was an excellent suggestion and pointed out that there are two cooperative organizations, CHORUS [Clearinghouse for the Open Research of the United States] on the commercial side and SHARE [Shared Access Research Ecosystem] on the academic side. APLU and AAU were key forces behind the establishment of SHARE. SHARE's progress is modest to date, but their goals are worthy, and OSTI can foresee a role where grantees can make one deposit through the SHARE service, and agencies will be able to harvest that information. Reed said that he understood that the problem for universities is more on the data side of public access than with the publications side. Hitson said that he had heard similar feedback because data is the newer challenge.

Barbara Helland was asked to speak to ASCR's requirements reviews and related activities.

ASCR has four major facilities: the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory; the two Leadership Computing Centers at Argonne National Laboratory (ALCF) and Oak Ridge National Laboratory (OLCF); and the Energy Sciences Network (ESnet)

In scientific grand challenges, previous requirements-gathering efforts indicated that one should lead with the science. In this approach, review meetings establish consensus on requirements, capabilities, and services; scientists, program offices, and facilities have the same conversation; a solid, fact-based foundation for service and capability investments is provided; and DOE mission goals are addressed by ensuring DOE science is effectively supported.

The challenges that can be tackled with the proposed LCF upgrades include energy storage, nuclear energy, combustion, fusion, electric grid, accelerated design, catalysis design, biomass to biofuels, high-resolution climate modeling, and rapid climate and Earth-system change.

Computing upgrades pursued by ASCR reflect major changes in memory, shifting toward persistent and nonvolatile memory.

The first objective of a new exascale requirements review is to capture the whole picture by identifying the continuum of computing needs for the program office from institution clusters to leadership computing and by including modeling and simulation, scientific user facilities' and large experiments' needs, data needs, and near-real-time needs. The information gathered will inform the requirements for ecosystems for planned upgrades in 2020 to 2023, including the pre-exascale and exascale systems, network needs, data infrastructure, software tools and environments, and user services. Another objective is to communicate to SC scientists the known/fixed characteristics of upcoming computing systems in the 2020 to 2025 timeframe and ask the DOE scientists for feedback on proposed architectures. The process will start with SC to identify what science would push the development of exascale computing.

These workshops will be held in the Washington, D.C., area so program managers from all offices can participate. They will be led by a program committee made up of community representatives and be attended by about 50 individuals, including DOE program managers, community representatives, and ASCR-supported applied mathematicians and computer scientists. There will be a preplanned agenda and pre-meeting homework for the attendees. A summary workshop report will be written for each workshop. A case-study template has been developed. The HEP requirements review was held in June 2015 with 71 total attendees, 9 white papers, and 9 case studies; breakout sessions were held on computing-intensive modeling and simulation and on data-focused analysis and workflows.

The BES Requirements Review will probably be the largest one. It was held in Rockville in November 2015 with 99 total attendees; 50 of 60 attendees submitted white papers. Breakout sessions were held on quantum materials, core challenges in heavy-element chemistry, exotic states, emergence; catalysis, photosynthesis, light harvesting, and combustion; materials and chemical discovery; computing and data challenges at BES facilities; next-generation programming; advances in quantum algorithms; math and computer science; and soft matter, biochemistry, and bioinspired materials. The participants concluded that user facilities will be the drivers.

In a related needs-gathering activity, a call for input from all 17 national laboratories was released on May 31, 2015, to identify potential applications that could deliver new science capabilities on exascale systems. The 133 responses received will be used by ASCR/ASC to identify additional key scientific areas for exascale discovery and specific opportunities for new and existing scientific applications. They will also provide broad input on the kinds of partnerships and investments required to address the technical challenges of exascale applications.

An NIH–NSF–DOE Request for Information was issued in September 2015 to identify scientific research topics that need high-performance computing capabilities that extend 100 times beyond today’s performance on scientific applications. Responses identified a range of science areas, with biology, physics, materials, geophysics, and health leading the pack. A common taxonomy was developed to understand the science drivers and the portfolio across ASCR facilities; it works across science disciplines and also highlights DOE science areas. These science categories provide weight to high-use application areas and are inclusive of possible future applications. The top-level categories are mathematics, computer and information science, physics, chemical sciences, biological sciences, geosciences, materials sciences, engineering, applied energy technologies, social, behavioral, and economic sciences, and health sciences.

These responses provide compelling science drivers for the exascale in climate challenge problems, nuclear-energy challenge problems, and health-science challenge problems. They also identified common barriers and challenges in applications (scaling, using accelerators, lack of high-performance-computer-ready codes, and lack of software containers); in workforce development (producing scientific professionals who can use highly parallel systems); in data workflows [determining the memory and input/output (I/O) footprints on future architectures, data protection, data access from multiple sources, and data sharing]; and in ease of resource access. Of the 248 barriers and challenges identified, 52 emphasized data as an important component.

Software technologies cited in the responses included math libraries, components, and functions; compilers and languages; accelerator programming; shared and distributed memory parallelism; development environments and runtime systems; data science, data management, and visualization; and domain-specific frameworks, environments, runtime systems, libraries, and workflows.

A workshop was held in September 2015 on scientific software architecture for portability and performance to overcome the problem of running applications on machines with different architectures. The primary findings of the workshop were that true performance portability is very hard, and the community still needs to determine what is possible; software engineering for applications and libraries is key to tackling the application challenges of the future; and communication paths are needed to share the huge amount of work being done today.

The workshop concluded that, in application architecture, portable functionality is easier than portable performance; the best practice today is separate node and core-level parallelism; concerns should be separated, including data structures; and libraries are prohibitive for early functionality and performance in early science periods. In libraries and tools, users and developers need confidence in funding and persistence; transparent roadmaps are needed; a homogenous software stack helps adoption; and better communication with science/application developers is needed. In software engineering, scientific software needs test-driven development; good practices impact productivity, portability, performance, and personnel retention; and not enough training or incentives are provided for scientists to adopt the best practices.

Dunning noted that BES had asked for software best practices for materials science and asked if there would be such forward-looking actions by the other SC offices. Helland replied, yes. One needs to build bridges with the program offices. Dunning suggested that using wikis is a good way to provide information across the community. Helland said that the Office *was* using wikis.

Crivelli noted that when Helland discussed integrated imaging under the exascale, it was exciting news. But no software technologies in integrated imaging were mentioned. Helland replied that it was found that there are not a lot of tools for managing data, and some new tools will have to be developed.

Berzins found it interesting that there is no sense of how many of the best-practices efforts are of a preliminary nature. Helland answered that more information along that line is being asked for. Berzins said that there was not a lot of pure engineering. Helland said that that was not surprising because the request was sent to national laboratories, who do little engineering.

Bergman asked how it would work for machines that hit the floor in 2016 and thereafter. Helland replied that the national laboratories are working on libraries. They are supposed to be working together, but there has not been a report from them on their cooperative actions.

Sarkar asked if there were any milestones for the portability of applications. Helland said no, but wished that there were some.

Levermore asked if enough were known about the machines to judge portability. Helland responded that because of the initiative's partnerships with vendors, it is able to build emulators.

Chen asked if there will be tight linkages between the development of performance portability of the applications and the development of the software stack, with opportunities for the application teams to iterate with the software-stack development team. Helland replied, yes; that is already being worked on.

A break was declared at 3:14 p.m. The meeting was called back into session at 3:29 p.m.

Alexander Larzelere was asked to present an update on the Consortium for Advanced Simulation of Light-Water Reactors (CASL).

A Brookings Institution report published in February 2009 advocated the establishment of several dozen Energy Discovery-Innovation Institutes (e-DII) to foster partnerships; develop and rapidly transfer highly innovative technologies; build the knowledge base and human capital necessary to address the nation's energy challenges; and encourage regional economic development. Eight hubs were proposed; three were funded, one in the Office of Nuclear Energy (NE).

The question was finding the right problem for the NE hub. Building a virtual reactor might show how to get more from existing reactors that were known well and had a lot of performance data. The virtual reactor was based on TVA's Watts Bar 1, a Westinghouse pressurized water reactor (PWR). Its core is 12 feet tall and 10 feet across. It has more than 51,000 pins and more than 16 million fuel pellets. Modeling such a complex system is a complicated problem. Industry was interested in this project. These reactors have been operating for decades.

"Virtual" means that one has to model three dimensions (four dimensions with time) and coupled multi-physics, multiscale (by coupling the simulation processes) at high resolution with high fidelity, sufficient simulation time, sufficient simulation space, enabled by high-performance parallel-processing computing with rigorous verification, validation, and uncertainty quantification.

CASL built a team from four national laboratories, three universities, three industrial partners, and numerous associate members that used an innovative mix of collaboration technologies and geographic co-location. CASL developed and delivered the Virtual Environment for Reactor Analysis (VERA), which includes interoperability, chemistry, thermal mechanics, thermal hydraulics, and neutronics. One needs one input/output for all code modules. There must be interoperability internally and with commercial codes.

In Phase 1 of the program, CASL was built and used to create insights into industry-defined challenge problems, such as crud, pellet-cladding interaction, and cladding integrity during reactivity insertion accidents. VERA was implemented with a user environment that is appropriate for an industry setting. A simulation that calculates in all the characteristics of the 51,000 pins and 16 million fuel pellets presented a new view of what is happening in the whole reactor.

CASL's business model called for it to focus on use-inspired research and technology deployment as enabled by a light federal touch with an educational program imbedded in it.

Not every technical approach worked out, and not everybody worked out. CASL was given sufficient time and funding to figure that out and make corrections. NE acted as a partner in the success of CASL. It implemented the General Groves Model of Lab Oversight: bet on success, respect the special relationship instilled in the Federally Funded Research and Development Center (FFRDC) business model, leave the technical program planning and management to the national laboratories, and minimize the number and direct involvement of federal staff. The national laboratories are a good place to put hubs; they have the needed facilities, infrastructure, management, etc.

The CASL team developed a management plan, and DOE determined that it would satisfy the Office of Management and Budget (OMB) oversight requirement. As the program developed, the emphasis was on describing demonstrated taxpayer return on investment, certification by industry and science councils, and reporting to stakeholders. At the end of CASL, the taxpayers will have invested about a quarter of a billion dollars.

Lessons learned during CASL first 5 years (Phase 1) included the needs to use the competitive-proposal process; to focus on solving specific game-changing problems; and to maintain a light federal touch; and to maintain a fierce sense of urgency.

After CASL started, that led to the question, who else is out there? What other DOE programs are building and/or using advanced computing? The Advanced Computing Tech Team (ACTT) was set up to promote and facilitate the improved use of advanced computing technologies by scientists and energy innovators (in government, academia and industry) in support of the applied energy technology missions of DOE. The ACTT will work to move DOE advanced computing activities towards harder implementation challenges and greater potential impact. It was officially designated a Tech Team by former Secretary Chu in June 2012. It meets bimonthly and focuses on information exchange, common issues and challenges, and opportunities for technology reuse and collaboration.

CASL is now in Phase 2, the application process, which will focus on end-user nuclear energy modeling and simulation, science and engineering, workforce training and education, and sustaining taxpayer investments. The program will shut down upon successful completion. R&D will continue.

Key accomplishments of CASL in FY15 were VERA simulation of Watts Bar Nuclear Unit 1 operational history; Integral Pressurized Water Reactor (iPWR) small-modular-reactor (SMR) modeling demonstration; simulation of product-induced power shift challenge problem; 3-D modeling of pellet–cladding interaction and comparison with the Braidwood experience; and a second major externally available release of VERA.

The addition of advanced computing provides the opportunity to obtain insight into details about how complex physical processes work in ways that are not available through theory or observation.

Hey asked what analytical tools were available when these reactors were built. Larzelere replied that one-dimensional simulations were used; they worked, but not very well.

Dolbow asked how predictive the simulations will be in the next generation of reactors. Larzelere said that there are simulations of the AP1000. The first AP1000 will go online next year in China, and one will be able to see how good the predictions are. Westinghouse is very happy with what has been done.

John Shalf was asked to discuss memory errors, presenting the results of a 4-year field study of memory error on DOE supercomputers, with projections for the future.

Reliability is crucial for large-scale systems. One must confirm that reliability models are accurate, and one must use data from real systems to correlate to models. This study used more than 500 million CPU socket-hours and 40 billion dynamic random-access memory (DRAM) device-hours from Cielo at Los Alamos National Laboratory and Hopper at NERSC.

Failures were found at solder joints and other connectors/mechanical devices, and ephemeral big upsets are tied to energetic particle strikes. In the latter case, the probability is proportional to the surface area exposure and is very localized. However, software is the predominant mode of failure. Most of the computer is memory with a lot of discrete components. Static random-access memory (SRAM) structures consume a lot of area on modern CPUs. DRAM has a lot of components, but these components are resilient to these upsets. DRAM reliability is important today in everything from laptops to petascale supercomputers, and DRAM reliability will be critical in the future.

There are architectural and micro-architectural approaches to reliability. To get it right, one must know the faults to expect. The study looked at faults collected in production systems in the field.

A FIT is one failure per billion hours of operation or a failure every 115,000 years for a single device. That corresponds to a failure every 12.8 years for a machine with 8944 nodes, every 1.6 years for 71,552 dual in-line memory modules (DIMMs), and every 36 days for 1,144,832 DRAM chips. A target node FIT rate of 1000 is a failure every 4.6 days. A target DRAM chip FIT rate of 35 is one failure every day.

A fault is the underlying cause of an error, such as a stuck-at bit or high-energy-particle strike. A transient fault will return incorrect data until overwritten; it is random and not indicative of device damage. A hard fault (like a solder-joint failure) consistently returns an incorrect value. It must be repaired by disabling or replacing the faulty device. Intermittent faults sometimes return an incorrect value under specific conditions, such as elevated temperatures, and are indicative of device damage or malfunction. These are so rare that the probability of one component failing twice is infinitesimal.

An error is an incorrect state resulting from an active fault, such as an incorrect value in memory. Some get corrected, some are not correctable, some are not detected but do matter, and some are benign and are never detected and do not matter.

Data on failures are collected, and a scrubber is used to coalesce errors into faults. A lot of stuff on chips is parity protected, which just detects errors. Error-correcting code (ECC) protection can do single-error correction and double-error detection. Data are interleaved for protection. When data are stored in contiguous locations, a failure is uncorrectable because of a four-bit error. When data are stored in noncontiguous locations, a failure is correctable because of a one-bit error.

Transient faults are pretty much constant over time. The rate of permanent faults declines with time. More than 50% of faults are permanent. In double-data-rate (DDR) machines, there is ECC plus address parity, a valuable addition to the DDR specification.

An SRAM case study was run on Hopper and showed that the uncorrected errors were dominated by the L1DTag. The L2Tag has ECC but still has a significant number of failures because only one bit per entry is covered by parity. Details matter; seemingly small decisions can have a large impact on system reliability. For SRAM faults, accelerated testing with neutron beams correctly predicts error rates in the field, demonstrating that SRAM faults and mitigation techniques are well-understood. The fault model has been validated by accelerated testing and field data. Chip architects must pay attention to reliable design (and they do).

It was expected that altitude would produce higher fault rates because of cosmic ray and gamma-ray air showers. Cielo at 9000 feet and Hopper at sea level were compared. Actually, the failure rate was correlated with the number of boxes. It turned out that the chips were manufactured by three vendors. Once chip quality was corrected, one could see the altitude effect. One could also see differences in failure structures from vendor to vendor.

Chipkill ECC is the ability to correct any error from a single DRAM device. It requires more overhead than SEC-DED ECC. It was found that 30% of multi-bit errors were detectable by SEC-DED, but 70% were not, Chipkill corrects 42 times more errors than does SEC-DED. It can detect more errors and correct more errors. Therefore, single error correcting and double error detecting (SEC-DED) code is poorly suited to modern DRAM technology.

However, this conclusion assumes that all DRAM is being used all of the time, which is not the case. Counting logged errors overemphasizes the impact of permanent faults. Error events are not independent. A single fault produces some arbitrary number of errors. Permanent faults tend to cause more errors than transient faults. Logged error rates are of no help in evaluating a system; uncorrected error rates are helpful.

Hopper has a memory error rate four times that of Cielo, but a memory fault rate that is 0.625 times that of Cielo because of the lower altitude. Error counts are confounded by other factors, such as workload behavior; they are not an accurate measure of system health.

In projecting to the exascale, it must be borne in mind that, at scale, even small structures see faults. Error-rate targets can be met by adding layers of protection until the target is met. Vendors must pay attention to reliable design.

In projecting the SRAM uncorrected error rate from Cielo to the exascale, large systems (100,000 nodes) have double the uncorrected error rate of small systems (50,000 nodes). This is the same fault rate as 45 nm (the sky is falling). If faults are scaled according to the current trend, the sky falls more slowly, and a switch to fin-shaped 3D field-effect transistors (FinFETs) may make the situation even better. If some engineering effort is added, the sky stops falling. At the exascale, SRAM faults are unlikely to be a significantly larger problem than they are today.

In projecting DRAM error rates to the exascale, it depends on which chipset is selected. Error rates are 10 to 70 times greater than those of current systems. This is not just a problem for the exascale; it is a cost problem for data centers and the cloud. Solutions are out there, and lots of people are working on this problem. DRAM subsystems need higher reliability than those in use today.

The conclusion is that large systems require reliable design and reliability modeling. Field-data analysis is crucial to correlate reliability models and to guide DOE investments. Component suppliers must be tracked to draw proper conclusions. Collaboration among DOE researchers, vendors, integrators, and facilities is critical to achieving this goal.

A number of risk factors must be faced in the future. For SRAM structures, the likelihood of faults is very high, but the risk is low because the model for particle strikes on CMOS SRAM remains solid. Addressing this issue is just a matter of engineering and cost. For JEDEC DIMM structures, the likelihood of faults is medium, and the risk is medium because vendor effects have more effect than the conventional fault model projects. For stacked DRAM, the likelihood of faults is medium and the risk is high because there are no field-test data to confirm very-well-thought-out models.

The error-tolerance requirements for self-driving vehicles are approaching those required for the exascale. The same microarchitecture error-tolerance techniques will be employed in both places.

Sarkar asked if the choice between SRAM and scratchpad memory had any implication. Shalf said that the tag contributes more errors than SRAM does. Sarkar asked if there were any discussions of the potential impact of hardware bugs. Shalf replied that there are verification challenges for billions of transistors, but complexity is the biggest problem rather than cosmic rays.

Levermore asked how the conclusions apply to other, smaller machines. Shalf responded that, in this study, error rates scaled down proportionately to the size of the machine.

Berzins asked how the transition to a new consumer market will affect reliability for the exascale. Shalf answered that the consumer market does not require the same reliability as the exascale, except for smart cars. This initiative will progress ahead of the smart cars. Vendors will be incentivized to work with the DOE to deal with error rates. The technology will translate 100%.

Dongarra asked how software manufacturers should react to mitigate errors. Shalf responded that what should be expected in the exascale is what we see in current machines. Software does not trap errors soon enough. The calculating moves forward too quickly. Dongarra noted that nonvolatile random-access memory (NVRAM) was not addressed. Shalf said that it was not studied.

Hey asked if many jobs were stopped by these errors. Shalf said, yes, about once an hour for jobs that run the whole machine. Some sort of recovery mechanism will be needed.

The floor was opened to public comment. There being none, the meeting was adjourned for the day at 5:20 p.m.

Thursday, December 10, 2015

The meeting was called back into session at 8:36 a.m.

Gwendolyn Huntoon was asked to report on the COV to the Next-Generation Network for Science (NGNS).

To help the research communities make efficient and effective use of current and future computing capabilities, ASCR supports a basic research program in networking. To ensure the integrity of this research program, a COV reviewed the management processes for the NGNS elements of the ASCR program. The COV met in October 2015, drafted and reviewed its report in November, and presented it to ASCAC for comment and acceptance in December.

On the basis of presentations by and discussions with the NGNS Office, the COV considers the NGNS to be a good program with quality execution, including in-depth reviews resulting in an effective and well-managed program. Overall, the COV was impressed with the quality and amount of work that gets done with a limited number of staff and funds.

The COV looked at the solicitation and review process. The basic finding was that the process is effective and well administered. The COV recommended that the program continue to broaden the breadth and diversity of workshop participants without diluting the focus of the workshop topics on the SC mission. Other cities beside Washington, D.C., might be considered. At least one workshop on network modeling should be hosted to better understand the issues associated with stimulating research in the area. The full implementation of the Portfolio Analysis and Management System (PAMS) into the program was encouraged, including linking the portfolio brief for each funding opportunity announcement (FOA) into the program.

The COV monitored active projects and programs and found that program officers effectively use a number of mechanisms to both interact with and track the progress of funded projects, including PAMS. It also found that program officers are limited in their ability for on-site visits to monitor projects and to interact with the project principal investigators (PIs) and associated staff. It recommended that the program continue to integrate PAMS into the tracking and management of the funded-project portfolio during the post-award. (This is being done for all programs.) This will make it easier for future COV's to do their work. It recommended that support, including appropriate travel funding, be provided to the program offices so they can perform project-site visits. Program officers should continue to participate in a broad set of community events where funded projects are presented or discussed as well as participating in strategic meetings where future network requirements and technologies in the support of scientific applications are discussed so they can get a good sense of what is going on in the field.

The COV examined the breadth and depth of the research portfolio and found that the portfolio, developed through a coherent set of FOAs, is of high quality and addresses challenges that are distinctly relevant to the mission of SC and DOE. Because of funding limitations, a number of quality proposals were not funded, particularly for the larger FOAs. The COV encouraged cross-agency collaboration, particularly in identifying overlapping high-performance-networking issues and technologies that each agency is working to address. Overlaps and possible synergies with NSF and NOAA were discussed. The program should develop a tighter relationship with ESnet where use of ESnet resources are specifically written into the FOAs as a target platform. The program should identify institutions from which they should be getting proposals but are not and should broaden workshop participation as well as the target groups for the FOA announcements accordingly.

The COV considered the program's anticipation of and addressing of emerging challenges. It found that the NGNS has the ability to influence the arc that high-performance network technologies and infrastructure developers take in support of and as part of the overall DOE and SC programs. It encouraged strategic planning among NOAA, NSF, DOE, and other agencies working in this realm to leverage total dollars spent on these efforts most efficiently and effectively.

The COV assessed the international standing of the program and found that it contributes to DOE's leadership role in the development and deployment of high-performance networking tools, technology, and middleware in support of science, producing prominent research both nationally and internationally. It recommended that the program continue to encourage this participation, including through international collaborations, when appropriate. In addition, the program office was encouraged to track activities that underscore the program's national and international standing, possibly through the PAMS system or through the annual reporting process.

Sarkar asked whether the laboratories and universities understood the larger context when they were doing research. Huntoon said, yes; they were aware of the motivations and extent of the program.

Reed asked the Committee if it had any comments on the draft report itself.

Cerf noted that the process was rigorous and asked a lot of hard questions. It reflects a genuine effort to make the program more relevant. The technology is being pushed, and cooperation with the Advanced Research Projects Agency (ARPA) and others is important.

Berzins said that the scope of the program is limited, and the field holds much promise.

Chalk asked for a vote for accepting the report with minor editorial changes. The vote was unanimous in favor of accepting the report.

Richard Gerber was asked to discuss the systems roadmap and plans for supporting extreme-data science.

NERSC provides mission-supportive high-performance computing and data resources for SC R&D, where theories, simulations, and experimental data are coming together. NERSC has a strong focus on science with 1,808 refereed publications in 2014. It deploys first-of-a-kind systems. It has more than 6000 users computing at scale and at high volume with a diversity of algorithms. SC offices allocate 80% of the computing and storage resources at NERSC, the ASCR leadership computing centers allocate 10%, and the NERSC director allocates 10%.

The NERSC-8 System, Cori, supports the broad SC research community and will begin to transition the workload to more energy-efficient architectures. It is a great system for supporting data-intensive science. It uses an Intel Knights Landing processor, a next-generation Xeon-Phi with a greater than 3-TF peak. Up to 72 cores per processor run at 1.2 GHz with support for four hardware threads each; this is more cores than the current-generation Intel Xeon Phi. It has 512b vector units with three times the single-thread performance of the current generation Xeon Phi co-processor. It has high-bandwidth on-package memory (up to 16-GB capacity) with a bandwidth projected to be 5 times that of DDR4 DRAM memory and higher performance per watt.

Cori will be installed in the Computational Research and Theory (CRT) Facility, which has exceptional energy efficiency and natural air and water cooling. The NERSC staff is moving into the building in late December 2015. Everything went well in the Cori Phase 1 installation, and it is up and running in the CRT now, with all NERSC users in pre-production mode. 83 million NERSC MPP [Message Posting Protocol] hours have been delivered to science, and it is running at high concurrency.

NERSC's current big system is Edison, the first Cray petascale system with Intel processors, an Aries interconnect, and dragonfly topology. It has a very high memory bandwidth. Hopper is also running for another week. Since Cori Phase 1 came online, Edison has been better serving the demand for running large jobs. More than 16-K-core jobs are using 80% of the time now. Hopper will be retired; Edison is being moved to the new facility from downtown Oakland; Cori's installation will be completed next summer; and NERSC-9, -10, and -11 are upcoming.

The CD-0 for NERSC-9 was signed in August 2015. The RFP draft technical specifications were released in November, and vendor feedback is due in December. The design review is scheduled for January 2016, and an independent project review will be conducted in the second quarter of CY16. The request for proposals (RFP) will be released in the late spring or early summer of 2016.

The NERSC Exascale Science Application Program (NESAP) is designed to prepare the SC user community for the Cori many-core architecture. NERSC is partnering closely with about 20 application teams drawn from across SC and will apply lessons learned to the broad SC user community.

The NESAP code teams get early access to hardware, technical deep dives with Cray and Intel staff on-site, multi-day deep dives (dungeon sessions) with Intel staff in Oregon to examine specific optimization issues, and user-training sessions. These activities have been extremely productive, so far. NERSC staff provides hands-on assistance, and NERSC is hiring eight postdocs with application-performance expertise.

To run programs effectively on Cori, users will have to manage domain parallelism, increase thread parallelism, exploit data parallelism, and improve data locality through cache blocking and the use of on-package memory.

The effort is at or ahead of schedule at this time. Astounding progress has been made, and many codes are ready to run. All of the NESAP codes have demonstrated speedups (some running up to three times faster).

What have gone well are the setting of requirements for dungeon sessions, engagement with the Intel Xeon Phi Users Group (IXPUG) and user communities, and a large number of the NERSC and vendor training sessions. A massive amount has been learned about tools and architecture. The NERSC-9 workload analysis shows a large fraction of jobs using less than 16 GB of memory per node, producing higher concurrency.

NESAP plans to increase the excitement and effort in 2016 with extra training events, on-site hackathons, and more dungeon sessions; to continue the successful Cray + Intel pipelining approach; to continue the application-readiness program as an ongoing center effort through 2025; and to maintain a community database of lessons learned.

Cerf asked if there were any difference between planned allocation and the actual usage. Gerber replied that there has been mid-year shuffling among programs within an office; all allocations are fully used. Cerf asked if resources of the machine failed to be used because of the allocation process. Gerber replied, no; NERSC is always resource constrained; it is always at 99% utilization. Cerf asked who will be competing for the use of NERSC-9. Gerber said that it is open to all vendors now. Cerf asked if there were any instrumentation to track how codes run. Gerber replied, yes; there are tools at all levels that users can use to optimize their codes. There is not a holistic performance analysis capability, however. Cerf referenced one of the slides and asked how one interprets a speed-up factor of 0.5. Gerber said that it means that the code runs half as fast and pointed out that the next slide refers to a 50% faster running, a speed-up factor of 1.5. Dunning pointed out that Intel says that the processor runs at “up to” 16 GB, but the slides refer to “16 GB.” Gerber answered that that processor has been upgraded to 16 GB now. Dunning pointed out that Pacific Northwest National Laboratory has a Knight’s Corner machine and asked if NERSC were making use of it. Gerber replied that they were not because they had their own Knight’s Corner machine.

Bergman said that this level of success is very commendable and asked what went wrong, if anything, and what should be changed going forward? Gerber replied that the things that did not go well were really trivial.

Berzins asked if there were a simulator that would predict performance. Gerber answered that that would be a good idea. Berzins asked if the final delivery will spill over into 2017. Gerber responded, no. Some activity will be going on, but the machine will be in place. Berzins asked what the socket power use will be. Gerber said that it would be about the same as current machines, but the efficiency will be much greater (a factor of 6 to 8). Determining that value is not entirely straightforward.

Sarkar asked whether anything was said about portability in the dungeon sessions. Gerber said that the three national laboratories have been having discussions about portability for some time. There was a best-practices workshop, and the program is looking at more libraries.

In regard to programming practices, Cerf asked what range of programming languages were being used. Gerber replied that Fortran and C5 are heavily used, mostly Fortran. There is also some use of Python. Cerf asked if compilers go directly to instruction sets. Gerber said that they do in most cases. Cerf asked whether one can use simulators to optimize codes before the machine is available. Gerber answered that emulators have to come from the vendors for particular chipsets, and the vendors are not very forthcoming. Dunning noted that emulators are slow, applicable to only small segments of a code, and not very useful. Coghlan said that a lot of code-level collaboration is performed at Argonne National Laboratory (ANL). Emulators were slow and only good for small codes. With the KNL and KNH processors, good numbers are being produced.

Gerber resumed the presentation, noting that extreme data science is playing a key role in scientific discovery, with NERSC providing the high-performance computational needs of research efforts that have led to five Nobel prizes in theoretical chemistry, high-energy astrophysics, cosmology, neutrino physics, and climate change. NERSC has been supporting data-intensive science for a long time involving the scientific investigations of supernovae, the cosmic microwave background radiation, the Large Hadron Collider, neutrinos, the light sources, and bioinformatics.

In addition, NERSC archives an enormous amount of data for the scientific community: 60 PB of data are stored in NERSC's High-Performance Storage System archive.

NERSC's goal is to increase the productivity, usability, and impact of DOE's experimental user facilities and other data-intensive science by providing comprehensive data systems and services to store, analyze, manage, and share data.

Through 2015, NERSC deployed separate computing-intensive and data-intensive systems. But really, how different are those platforms? In policies, they all have fast turnaround times and can run large numbers of throughput jobs. In software/configuration, they all have support for complex workflows, communication and streaming data from external databases and data sources, and easy-to-customize user environments. In hardware, they all have local disks for fast input/output, some systems have larger memory nodes, and they all provide support for advanced workflows.

NERSC is making significant investments in Cori to support data-intensive science. It is implementing new queue policies, high-bandwidth external connectivity to databases from compute nodes, more login nodes (23) for managing advanced workflows, virtualization capabilities, and a NVRAM/flash burst buffer as an I/O accelerator.

The burst buffer motivation is that flash storage is significantly more cost effective at providing bandwidth than is disk; that flash storage has better random access characteristics than does disk; and that users' biggest request (after wanting more cycles) is for better I/O performance. There are different ways to implement a burst buffer. NERSC is putting in a layer between the calculations going on and the back-end file system. It can be used like a fast file system or as cache. Burst buffers accelerate I/O by employing checkpoint/restart or other high-bandwidth reads/writes, applications with a high number of I/O operations per second, out-of-core applications, and fast reads for image analysis.

NERSC put out a Burst Buffer Early User Program call for proposals for the award of exclusive early use of the burst buffer on Cori Phase 1; 30 proposals were received. Successful proposals came from ASCR, Biological and Environmental Research, Basic Energy Sciences, Fusion Energy Sciences, High Energy Physics, and Nuclear Physics.

NERSC has heard from a number of users that the lack of real-time access to the system is a barrier to scientific productivity. With NERSC's new batch scheduler, the Simple Linux Utility for Resource Management (SLURM), it has the capability to offer immediate or real-time access on Cori Phase 1 for projects and users with requirements for fast turnaround. It added a question to the Energy Research Computing Allocations Process (ERCAP) form about real-time needs to assess demand and size real-time resources and got 19 yeses, a very small fraction of the workload.

Shifter is a container for high-performance computing that allows users to bring a customized operating-system environment and software stack to a high-performance computing system.

Hey asked if it destroyed the parallelism. Gerber replied that that issue is being looked at; communication is the hard part.

Cori's external connectivity is being upgraded to enable the connection of a 100-GB+ instrument to Cori.

There is now a superfacility concept. Data volumes are increasing faster than Moore's Law. Facility data exceeds local computing and networking capabilities. It is unfeasible to put a supercomputing center at every experimental facility. Fortunately, there is ESnet to connect all of DOE's facilities. Based on NERSC's experience supporting science from experimental facilities, a common design pattern can be perceived: data production by sensors; local data processing/filtering; predictable data movement; on-the-fly calibration; analysis and modeling; real-time access and visualization; and storage, archiving, and sharing. As a response to this phenomenon, NERSC has engagements with a number of experimental facilities, such as the Advanced Light Source, Linac Coherent Light Source (LCLS), Large Hadron Collider, and Joint Genome Institute.

Hey asked how networking with the LCLS at the Stanford Linear Accelerator Center (SLAC) was coming along. Gerber replied that proof of principle has been demonstrated in getting the data to NERSC.

Crivelli asked if there were a plan to hire more people for the Data Analytics Group. Gerber explained that the facility has a Data Analytics Group to help people in machine learning. There were two new hires during the week of this meeting. This group is getting more resources.

Gregurick asked if the facility's vision was to host and distribute climate data. Gerber replied that that was being looked at on a case-by-case basis, looking at levels of support, longevity of data maintenance, etc. No data have ever been deleted.

A break was declared at 10:35 a.m. The meeting was called back into session at 10:48 a.m.

Kalyan Perumalla was asked to review his Early Career Award work on reversible software-execution systems.

Reversible software provides a new path to exploiting inherent model-level reversibility. It provides an efficient alternative to checkpoint/restart approaches. It addresses fundamental computational-science questions with respect to limits of energy and computation time.

Normally, as one starts a program, one lets it run to completion, and gets results. In reversible software, there are points where the program can be directed to go back to an earlier point and start running again from that point. The information needed to provide this flexibility is stored in memory. Execution can be switched between irreversible and reversible programs and components. Reversible computing can make subsets of irreversible languages, cross-compile an irreversible compiler and a reversible compiler, cross-compile an irreversible instruction set with a reversible instruction set, simulate a reversible computer with an irreversible instruction set, or emulate an irreversible computer with a reversible instruction set.

Issues to be considered include (1) reversibility (the ability to design an inverse circuit for every forward circuit, where inverse circuits recover input signals from output signals that may be built from the same or different gates as the forward circuit); (2) universality (the ability to realize any desired logic via composition of gates); (3) conservation; and (4) adequacy. There has been a lot of work on synthesizing circuits. Manifestations of reversible computing include energy-optimal computing hardware. New uses are relevant for high-performance computing: synchronization in parallel computing, processor architectures, efficient debugging, fault detection, etc.

The objectives of this study were to enable and optimize reversible computing to overcome the formidable challenges in the exascale and beyond in terms of the memory wall, concurrency, resilience, and emerging architectures.

The approach is to employ automation (reverse compilers and reversible libraries), runtime (reversible execution supervisor and reversible extensions to standards), theory (unified reversible execution complexities, memory limits, and reversible physicals system modeling), and experimentation (prototypes, benchmarks, and scaled studies).

There are several relaxations of forward-only computing to reversible computing. One is compute-copy-uncompute (CCU), which determines the minimum amount of energy needed for any computation. This is a basic algorithmic building block to avoid the erasures in arbitrary programs. Another is forward-reverse-commit (FRC). This is a basic operation in optimistic, parallel, discrete-event simulations, such as the time-warp algorithm. A third is undo-redo-do (URD, which is used in graphical user interfaces. And a fourth is begin-rollback-commit (BRC), which is used in databases, nested-tree computation scheduling, and high-performance computing languages.

Cerf noted that this is like recalling an email. Reed added that, here, one gets rollback avalanches.

Interest was aroused in reversible computing because of its relationship to energy consumption by computers. People asked, what is the minimum energy needed or dissipated during computational operations? It was assumed that every bit operation dissipated a unit of energy. However, it was seen that not every *bit operation* but every *bit erasure* dissipates a unit of energy. Indeed, other bit operations can be implemented without energy dissipation. The question then arose, what is the minimum number of bit erasures needed to run an algorithm? It was expected that there would be a nonzero computation-specific number, but Charles Bennett's surprising solution showed that zero bit erasures were possible for any arbitrary program. Further refinements in algorithmic complexity and trade-offs led to partial reversibility.

Indeed, reversible-computing-based recovery is significantly more efficient than memory-based recovery, as shown by comparing the runtimes for an algorithm on an irreversible computer and on a simulated reversible computer. This large runtime efficiency was seen both with CPUs and GPUs. The performance increases because of better memory behavior. A fault-tolerance scheme that builds on reversible-computing software relieves file-system congestion, relaxes the need for a global snapshot, enables the node-level freedom of checkpoint frequency, and avoids message replay.

The directions to be taken in programming languages (Janus, R, SRL, ESRL, Reversible C, etc.) were studied.

In Janus, reversible conditional execution can store a state in memory, and the program can jump back to that state from a point later in the program. Reversible looping is possible as is reversible subroutine invocation. Other reversible constructs include swap, arithmetic, and input/output. Because of this symmetry, jumpfrom and jumpto commands can simply drop their tags to become a single instruction type: jump.

Upper-bounded rejection sampling generates samples from any complicated distribution $p(x)$ without the need for any saved (checkpointed) memory to enable repeatable and reversible (bidirectional) sampling.

A new framework has been proposed for reversible integer arithmetic (reversible math) as well as reversible basic linear algebra subprograms (RBLAS). There is evidence and motivation for fully optimized reversible software at scale that will be a community exercise. Compiler libraries have to be brought together with virtual testbeds, implementation scaling, and proof of concept *inter alia*.

The evolution from irreversible to reversible computing is envisioned to entail the use of irreversible programs on irreversible machines (the current situation), reversible programs on irreversible machines (in the short term), reversible programs on reversible machines (in the medium term), and irreversible programs running on reversible machines (for the long-term). This is a natural extension of computing science.

Petzold asked whether one could avoid running out of bits if one allows a certain number of bits for this purpose. Perumalla replied that, yes, one can use up the allocation and then go on to another, higher allocation.

Sarkar said that it seems that reversible computing can touch on many of the challenges found in the exascale and asked which applications of reversible computing might have the greatest impact on exascale computing. Perumalla said resilience would show the greatest impact, then memory, then energy, and finally debugging.

The floor was opened to public comment. There being none, the meeting was adjourned at 11:51 a.m.

Respectfully submitted,
Frederick M. O'Hara, Jr.
Recording Secretary
December 28, 2015