Draft Minutes
Advanced Scientific Computing Advisory Committee Meeting
April 4–5, 2016
American Geophysical Union, Washington, D.C.

ASCAC Members Present

Martin Berzins
Barbara Chapman
Jacqueline Chen
Silvia Crivelli
John Dolbow
Jack Dongarra (via telephone)
Thom Dunning
Sharon Glotzer (via telephone)
Susan Gregurick
Anthony Hey
Gwendolyn Huntoon (via telephone)
David Levermore
Juan Meza (Monday only)
John Negele (via telephone)
Linda Petzold (via telephone)
Daniel Reed (Chairman)
Dean Williams

ASCAC Members Absent

Keren Bergman
Vinton Cerf
Vivek Sarkar

Also Participating

John Steve Binkley, Associate Director, Office of Advanced Scientific Computing, Office of Science, USDOE
David Brown, Director, Computational Research Division, Lawrence Berkeley National Laboratory
Richard Carlson, Computer Scientist, Office of Advanced Scientific Computing, Office of Science, USDOE
Christine Chalk, ASCAC Designated Federal Officer
Jody Crisp, Oak Ridge Institute for Science and Energy
Barbara Helland, Director, Facilities Division, Office of Advanced Scientific Computing, Office of Science, USDOE
Alexander Larzelere, Program Manager, Modeling and Simulation Energy Innovation Hub, Office of Nuclear Energy, USDOE (retired)
Michael Martin, Fellow, Office of Advanced Scientific Computing, Office of Science, USDOE
Patrick McCormick, Programming Models Team Lead, Los Alamos National Laboratory
Paul Messina, Senior Strategic Advisor, Argonne Leadership Computing Facility, Argonne National Laboratory
Cherry Murray, Director, Office of Science, USDOE
Frederick O'Hara, ASCAC Recording Secretary
Franklin Orr, Under Secretary for Science and Energy, USDOE
Robinson Pino, Program Manager, Office of Advanced Scientific Computing, Office of Science, USDOE
Brian Plesser, Attorney, Office of the General Counsel, USDOE
Rick Stevens, Associate Laboratory Director, Argonne National Laboratory

About 55 others attended in the course of the two-day meeting.

Monday, April 4, 2016
Morning Session

Before the meeting, **Brian Plesser** of the DOE's Office of the General Counsel conducted the annual ethics briefing for the non-federal members of the Committee.

The meeting was called to order by the chair, **Daniel Reed**, at 8:30 a.m., and he set out the themes of the meeting: post-exascale computing and the time scale of the research agenda for the transition to the exascale.

Reed introduced **Cherry Murray** to present an overview of the Office of Science (SC).

DOE is a mission agency, covering energy, science, nuclear safety and security, and environmental cleanup. About 40% of the Department's budget goes to the National Nuclear Security Administration (NNSA), 40% goes to science and energy, and 20% goes to everything else (mostly environmental cleanup). High-performance computing underpins all of these areas. SC funds high-risk, high-payback scientific research. The Advanced Research Projects Agency–Energy (ARPA–E) deals with high-risk science and technology. The applied technology offices deal with environmental cleanup and with energy technologies that are close to development and deployment.

Federal agencies are always working on three budget cycles. They are currently executing the FY-16 cycle, having hearings before Congress on the FY-17 cycle, and internally preparing for the FY-18 cycle.

SC is the largest supporter of the physical sciences in the United States. Its FY-17 budget request is for $5.67 billion, a 6.1% increase over the FY-16 appropriation. Under the terms of the 21st Annual Conference of Parties (COP21) agreement, the United States is pledged to double its clean-energy research budget. To that end, $1.8 billion has been requested for Mission Innovation.

The President's FY-17 proposed budget for DOE requests $13.1 billion for nuclear security, $7.2 billion for energy technology, $5.8 billion for science research, and $6.1 billion for environmental management. Of the $7.2 billion requested for energy, $1.3 billion is for investments in the 21st-Century Clean Transportation Plan, which is highly unlikely to get fully funded.

The FY-16 budgets and requested increases for the SC programs are

- Advanced Scientific Computing Research (ASCR): $621 million; +6.8%
- Basic Energy Sciences (BES): $1.849 billion; +4.7%
- Biological and Environmental Research (BER): $609 million; +8.7%
- High Energy Physics (HEP): $795 million; +2.9%
- Fusion Energy Sciences (FES): $438 million; -9.1% [it had a large bump-up last year]
- Nuclear Physics (NP): $617 million; +3.0%

The proposed budget for FY-17 includes $100 million of funding for Mission Innovation. Of the $100 million, $10 million is for ASCR, $51 million is for BES, $35 million is for BER, and $4 million is for FES. There is computation embedded in everything. In FY-17, investments are made in all of the SC programs. SC also proposed $100 million of new funding for academic research.

In ASCR, the FY-17 budget will support the Exascale Computing Initiative (ECI) and Exascale Computing Project (ECP), optimally operating facilities at more than 90% availability, Scientific Discovery through Advanced Computing (SciDAC) partnerships, applied mathematics research addressing challenges of increasing complexity, and computer science research addressing exploration of beyond-Moore's Law architectures, and the Computational Sciences Graduate Fellowship.

The ECI was initiated in FY-16 to support research, development, and computer-system procurements to deliver an exascale computing capability by the mid-2020s. It is a partnership between SC and NNSA. The Oak Ridge and Argonne leadership computing facilities are being upgraded, and it is then on to the exascale! One thing that has become clear is that extreme-scale computing cannot be achieved by a business-as-usual, evolutionary approach.

The Energy Sciences Network (ESnet) will increase its bandwidth to 400 Gb/s on high traffic links.

In BES, the FY-17 budget will support

- Increased funding for Energy Frontier Research Centers (EFRCs),
- A new activity in computational chemical sciences,
- Advancement of the mission-innovation agenda that targets materials and chemistry for energy efficiency and for use in extreme environments,
- Both energy innovation hubs (a Joint Center for Energy Storage Research and the Joint Center for Artificial Photosynthesis),

- Maintaining international competitiveness and discovery science, and
- Five upgrades to facilities, including the Linac Coherent Light Source-II and the Advanced Photon Source upgrade.

In BER, the funding will support genomic sciences, mesoscale-to-molecules research, climate-system modeling, atmospheric-system research, environmental-systems science, climate and environmental data analysis and visualization, and user facilities. That Office is developing Accelerated Climate Modeling for Energy (ACME), a high-performance climate code, as a step toward the next community model.

In Fusion Energy Sciences, support is provided for the DIII-D National Fusion Facility (DIII-D) and the National Spherical Torus Experiment Upgrade (NSTX-U) programs; the U.S. research involvement in international machines; general plasma-science activities; and the U.S. contributions to the internationally funded experimental fusion reactor ITER.

In High-Energy Physics, the plan developed by the Particle Physics Project Prioritization Panel (P5) is being followed, focusing on construction and project-support increases to implement the P5 strategy and accelerator stewardship. That strategy calls for continuing active engagement in the highly successful Large Hadron Collider (LHC) Program, solidifying international partnerships for a U.S.-hosted Long-Baseline Neutrino Facility/Deep Underground Neutrino Experiment (LBNF/DUNE), and advancing the understanding of dark matter and dark energy.

In Nuclear Physics, an upgrade for the Relativistic Heavy-Ion Collider, the 12-GeV upgrade to the Continuous Electron Beam Accelerator Facility (CEBAF), the continued construction of the Facility for Rare Isotope Beams (FRIB), the Gamma-Ray Energy Tracking Array (GRETA), and fabrication of a Stable Isotope Production Facility (SIPF) are being supported.

In FY-17, SC is expected to contribute to five DOE crosscuts: advanced materials, exascale computing, subsurface technology and engineering, the Energy–Water Nexus (EWN) Program, and cybersecurity.

In developing the FY-17 budget, SC is grappling with a number of issues and priorities. It seeks to balance discovery research versus science for clean energy and Departmental crosscuts. It seeks to balance research funding versus scientific user facility construction versus operations. The ECI is important for national and economic security and will continue into the out-year budgets. A defining moment in fusion sciences is occurring this year. Upgrades of the CMS [Compact Muon Solenoid] and ATLAS [A Toroidal LHC Apparatus] detectors at the LHC are being requested at the same time as construction of the LBNF/DUNE. For that facility, the world's most intense neutrino beam will be produced at Fermilab and directed 800 miles through the earth to Lead, South Dakota, where a 40-kiloton liquid-argon neutrino detector will be placed in the Homestake Mine. How to fit all these large expenditures into the budget is a huge challenge. This challenge will be addressed through enhanced communications with Congress and research universities and through instituting best practices in national-laboratory management.

Exascale computing is a grand challenge. DOE is partnering with industry to codesign machines and codes for exascale computing.

May 2, 2016, is a make-or-break moment for the fusion community. The program is 10 years into the ITER project, and the question is whether the United States is in for the long haul or not. The nation committed to $1 billion of support for ITER, and it has now contributed $4 billion to $6 billion. Congress is not happy. How to fit it into the budget is a major question. At the same time, there is also the LBNF/DUNE. How to get that project built is another major question; international support is being sought.

In addition to SC, ARPA–E is also focused on transformative science. Within SC, 70% of the applied technology and 32% of the scientific programs are scored as Mission Innovation by the Office of Management and Budget (OMB). In all, a third of SC's budget is underpinning clean-energy research. Secretary Moniz wants the new $100 million directed toward the interests of Mission Innovation. SC got a 2% increase in discovery science funding.

In FY-14, SC user facilities had 33,671 total users; 65% of those came from universities, 23% came from national laboratories, 8% came from other research institutions, and 4% came from the private sector. Among the construction projects proposed in the FY-17 budget are five upgrades to user facilities. In the FY-16 budget, 30% of construction funding went to universities (mainly Michigan State University). Much of the ECI is research, increasing the percentage of research in the overall budget.

SC partially funds ten science laboratories. Each of these national laboratories gets funds from many programs.

Reed asked if there were anything specific that Murray wanted the Committee to look at. Murray replied, the theme you mentioned at the beginning: beyond the exascale.

Hey mentioned that an ASCAC Committee of Visitors (COV) had looked at the Office of Scientific and Technical Information (OSTI). He asked how that organization fitted into the budget. Murray answered that it will be supported. There is a struggle going on to get principal investigators (PIs) to submit their publications to OSTI. Such submission will become mandated. A reasonable budget is being proposed to make this mandate happen.

Levermore noted that "innovation" has different meanings to different people. In the engineering community it is not the bedrock but what one does with the bedrock. He asked if she were happy with the scoring process and the 70/30 ratio. Murray implied, yes. However, how the Secretary and the White House are scoring innovation is by assessing how much gets into the market and making money. Some of SciDAC would be included in the scoring. She considered the 30% as being quite reasonable and noted that the cost of white light-emitting diodes (LEDs) had gone down 90%, making energy-efficient lighting more widely deployed.

Chapman asked if there were any other activities that would aid education in energy. Murray answered: workforce-development technician training, graduate students at the national laboratories, and undergraduate internships. SC is trying to develop the workforce that the nation will need in energy.

Franklin Orr was asked to address high-performance computing and the *Quadrennial Technology Review*.

The first *Quadrennial Energy Review* looked at 17 energy research areas that might be beneficial to pursue. The review was to be repeated every 4 years. This time, there are 65 technology assessments in the renamed *Quadrennial Technology Review* (QTR).

There have been interesting changes in the energy landscape with a large increase in wind technology and a 90% reduction in LED costs. How to get to scale in other technologies is a great problem. DOE's portfolio includes

1. Economic security: drive down the costs of clean-energy technologies
2. Energy security: develop a portfolio of diversified sources
3. Environment: improve the quality of air and water (great progress has been made in this area since the 1980s)

In an input-output analysis of energy use, there are primary energy sources and energy services. The grid or power generation system pulls them all together.

The future grid will differ radically from the present one. Today, the response time of a grid is much longer than the time it takes for a problem to develop. The Recovery Act of 2009 devoted a lot of resources to this problem in the form of communications systems, diverse types of sources, and control systems and centers. This challenge includes great computational problems. There is also the problem of intermittency to deal with as more-diverse sources come online. One research opportunity is energy storage as a way to deal with intermittency. Pumped hydro is now used for energy storage; more-efficient ways to store and retrieve energy are needed (e.g., through electrochemistry). One needs to understand the fundamental scientific underpinnings to base this technology on.

Clean electric power is another challenge. One approach is carbon capture and storage, although there is a question of cost. Improved compression technology has been modeled and simulated on the Titan at the Oak Ridge Leadership Computing Facility to better understand how compressors work.

Today, 19% of the United States' power comes from nuclear energy. Four new nuclear power stations are under construction. The Consortium for Advanced Simulation of Light Water Reactors (CASL) is also addressing life extension of current reactors and the certification of renewed reactors through modeling and simulation on the Titan.

Going from the Rankine cycle to the Brayton cycle and from steam to supercritical CO_2, one can increase efficiency from 30% to more than 40%. Wind turbines have a lot of opportunities for improvement, also. Computational analysis can further the understanding, effectiveness, and efficiency of these turbine systems.

In transportation, because of the high energy density of liquid carbon fuels, combustion will not go away. However, fuel cells, plug-ins, and light-weighting are making alternative-energy sources more useful and attractive.

In advanced manufacturing, researchers went from drawings to a 3-D printed (additive manufacturing) vehicle in six weeks. Reducing the design and manufacturing time makes the manufacturing cycle more efficient and cheaper.

Pretty much all of technology space is underpinned by an understanding of fundamental science. Characterization of materials at all scales is what is being done at all the DOE user facilities. Computing is an essential feature of all that is done at DOE. People will look back on this time as one in which there was an amazingness acceleration of the development of energy technology. A portfolio approach is required, and enabling science and computation are essential to the future success of the energy enterprise.

Hey asked about the status of the clean carbon capture power station project. Orr replied that several coal-fired pilot plants are under construction and will be in operation this calendar year. Much carbon dioxide is separated from natural gas. It is more difficult to separate carbon dioxide in coal-plant emissions than in natural gas. New separation technologies are being piloted. The costs of wind and solar power have come down so much that India and China may deploy them instead of coal. The use of coal in China decreased last year, according to the *New York Times*.

Dunning stated that the new data-driven way to manage cities calls for collaboration among offices of SC and asked what the status of this collaboration was. Orr answered that a lot of funding has been put into cross-cutting initiatives, especially the grid and the ability to heavily instrument the environment with sensors. Both of these cross-cut activities require high-performance computing. OMB and Congress support this approach.

Levermore said that modeling is a huge challenge. With the exascale, as the resolution of modeling is increased, new science is discovered and new algorithms need to be written. Orr said that there was no disagreement there. Magnetohydrodynamic models of the Sun are proof of that. The low-resolution models did not work, and high-resolution models were needed.

Chen said that it would be great to leverage the available codes and models to overlapping scientific areas. Orr said that links are being built, and they have been useful. It is expected that this linking effort will be continued. The EFRCs provide a rich space for operating in that way.

A break was declared at 9:57 a.m. The meeting was called back into session at 10:18 a.m.

Steve Binkley was asked to review the activities of ASCR.

The FY-17 President's request for ASCR is $663 million, which is up 6.8% or $42 million over the FY-16 enacted appropriation for the Office. Of this amount, 50% is for facilities and 50% for research. The SciDAC partnerships will be re-competed in FY-17. These partnerships have been very successful, and the Office does not want to alter that. It is likely that there will be partnerships with other parts of DOE. A modest effort in R&D for post-Moore's Law computing is included in the request as is a modest effort in support of the White House's BRAIN Initiative. The Computational Sciences Graduate Fellowship Program has been funded at $10 million, which might support 25 fellows.

In the funding summary, the ECI portions have been pulled out and reported under an Exascale line, producing a lot of negative numbers for various line items. Those negative numbers simply reflect the transfer of funding to the ECI budget line. The ECI was initiated in FY-16; and monies have been projectized according to the principles of DOE Order 413.3b and reflect the work breakdown structure (WBS). The first 4 years focus on research in software and hardware, followed by acquisition of systems to be operational about 2024. BES and BER are contributing to the development of exascale applications (climate modeling in the case of BER). The FY-15 enacted, FY-16 enacted, and FY-17 requested exascale funds for ASCR are $91 million, $157.9 million, and $154 million, respectively; for BER, they are $0,

$18.7 million, and $10 million; for BES, they are $8 million, $12 million, and $26 million. Historically, exascale funding has grown quite dramatically. The nature of R&D will change significantly in the coming years, and that will have to be managed carefully.

A new ASCR investment priority is to begin R&D for the post-Moore era. Going from 14-nm gaps to 10-nm gaps is challenging; going to 7-nm gaps will be even more difficult. New techniques will be added to supplement floating-point computing.

In the Office staffing, Claire Cramer has joined the staff with a focus on future computing technologies.

The Office is on track for upgrades at the National Energy Research Scientific Computing Center (NERSC) and the Leadership Computing Facilities (LCFs). The new Summit computer at Oak Ridge National Laboratory (ORNL) and Aurora at Argonne National Laboratory (ANL) will have system peaks at 200 petaflops and 180 petaflops, respectively, quite a step up from today's machines.

SC is investing $100 million in Mission Innovation, of which $10 million will go to ASCR. The Secretary is pushing hard for greater cooperation between DOE and the National Institutes of Health (NIH), specifically on the BRAIN Initiative. A DOE/NIH team on collaboration has been at work for 2 years. New funding for the collaborative effort is in the FY-17 request. In addition, changes in lithographic etching resolution from 35 to 14 nm have been based on technological advances; however, complementary metal-oxide semiconductor (CMOS) lithographic-feature sizes are now approaching the fundamental limits. One needs to think about what computing will look like 10 years from now. Three approaches under consideration are quantum computing, neuromorphic computing, and probabilistic computing. The technology that is being used now will last another 8 to 9 years. Planning needs to be started for going beyond that technology. The Office is setting up test beds and workshops on the new technologies to look at what algorithms can be run on these new systems. Cybersecurity is another area for which ASCR will initiate a modest research effort in FY-17 with an emphasis on the unique challenges of the Department's high-performance computing facilities.

Williams asked about the upgrade of ESnet to 400 Gb/sec. Binkley replied that there is a 100-Gb/sec backbone and links at almost all DOE research sites. Some areas can use 400 Gb/sec (the Bay Area and Chicagoland). The transatlantic connections have been increased from 30 Gb/sec to 100 Gb/sec, mainly to CERN and the LHC. This occurred in June 2015 in coordination with the LHC Run II. For Asia, ESnet partners with the National Science Foundation (NSF) and the Internet 2 Consortium. For South America, the details are still being worked out.

Chapman asked what collaboration is being conducted with the NIH on the BRAIN Initiative. Binkley answered that these are not DOE missions. "Touch places" are being looked for where DOE can further its mission and NIH's, also. In the National Cancer Institute (NCI), both agencies have multiscale modeling interests, for example, from the molecular scale to the whole healthcare system.

Paul Messina was asked to give an update on the Exascale Computing Project.

On the path to the exascale, the transition is being made from an initiative to a DOE project. DOE is the lead agency in the National Strategic Computing Initiative (NSCI). In DOE, SC and the NNSA will execute a joint program focused on advanced simulation through a capable exascale computing program emphasizing sustained performance on relevant applications.

In 2015, ASCAC recommended that the effort develop a detailed management and execution plan; use codesign to meet application needs; distinguish essential system attributes from aspirational ones; mitigate software risks by developing evolutionary alternatives to complement more innovative, but untested alternatives; support both data-intensive and computation-intensive workloads; conduct periodic external reviews; and work with other federal research agencies, as appropriate.

DOE has been funding research related to the exascale for more than 5 years. The Exascale Computing Project (ECP) is being launched as a joint SC/NNSA partnership. The ECP is managed by DOE national laboratories following DOE Order 413.3b.

It should be emphasized that the ECP does not procure exascale systems. Procurements will be separate from this project. [Dunning pointed out that the output from this project will influence the specifications and purchases of exascale equipment.]

The goals of the project include developing a broad set of modeling and simulation applications (that is where the payoff is); developing a productive exascale capability in the United States by 2023, including the required software and hardware technologies ready to support science; preparing two or more SC and NNSA facilities to house this capability; and maximizing the benefits of high-performance computing (both hardware and software) for U.S. economic competitiveness and scientific discovery including benefits across scientific computing, for example, affordable terascale laptops.

The project will fund and manage work at the national laboratories, industry, and universities. In some cases, the project will not be funding anything new but will build on existing activities and provide incremental funding. There is a formal solicitation and selection process. There are major deliverables and various reviews of major milestones and deliverables.

The project will pursue a 10-year plan structured into four focus areas: applications; software (enhancing the software stack that SC and NNSA applications rely on to meet the needs of exascale applications and evolving it to use exascale systems efficiently); hardware (continuing to fund supercomputer vendors to do the research and development of hardware-architecture designs needed to build and support the exascale systems envisioned by DOE now rather than their current business plans); and systems estimates (contributing funds to continue advanced-system engineering development by the vendors, incremental site preparation, and cost of system expansion needed to acquire capable exascale systems. Each system is projected to be $150 million; actual cost for the first of a kind (FOAK) may be $250 million.

Currently, the project office has been established at ORNL. Project team leadership has been selected. The project scope has been determined, and a detailed work-breakdown structure (WBS) has been created. An independent cost review was successfully completed in January, and an independent design review was conducted in March. Requests for information (RFIs) have been issued for applications, codesign centers, and software technology for the exascale; this effort started with 133 white papers, leading to 57 pre-proposals, of which 33 will be invited to submit full proposals; it is expected that 4 years of funding will be provided, with some extensions. A vendor information meeting for hardware technology was held in April with an RFI that went out in mid-March. The project is awaiting CD-0 (critical decision zero, approval of mission need) approval. A CD-1/3 review has been scheduled for Summer of 2016.

The project scope that has been worked out is based on mission needs and requirements: the breadth of the mission-critical SC and NNSA applications; historical and current software requirements; input on future needs; reports of workshops on application and software needs for the exascale; reports of workshops and analyses of hardware requirements; analyses of computing technology trends; the identification of gaps in vendor product plans; and experiences from the NNSA Accelerated Strategic Computing Initiative (ASCI) program.

Major issues include collaboration with other government agencies, training and productivity, codesign integration, and workflows. People have been selected to staff all of these WBS elements. Resilience needs to be integrated among all elements, but that has not been done, yet. Integration requires a holistic structure of application development, software technology, hardware technology, and exascale systems. A detailed, integrated project timeline has been established, and a proposed project-management structure has been developed. The project-management structure will include a board of directors that will include six laboratory directors, a Science Council to give advice, and an Industry Council to give advice. A memorandum of agreement (MOA) is being drafted to guide the Board of Directors and other project management.

The project has three phases: Phase 1 encompasses R&D conducted before the DOE facilities issue the exascale-systems RFP in 2019 (covering the first 4 years). In Phase 2, CD-2 is received, the exascale architectures and necessary engineering are known, and targeted development is conducted. In Phase 3, the exascale systems are delivered, gaps are discovered and addressed, and software is validated.

The solicitation and selection process will include an RFI, a review against the published criteria, down-selection to a reduced number eligible to respond to the request for proposals (RFP), a review against published criteria, and selection. Selection criteria include quality and makeup of the team, relevance to and expertise in the exascale, a match to mission needs, and technical feasibility. Process steps will be tuned, as appropriate, for the technical project under consideration and the target audience. And ECP focus-area leaders will select teams of subject-matter expert reviewers.

The ECP has developed a detailed set of about 75 technical risks. Among them are the inability of applications to achieve the science goals; vendor issues that could result in project delays, technical challenges, or financial losses; the inability to inject needed software technologies into production; and the unavailability of math libraries and software frameworks or the insufficient application integrity for adequate throughput and productivity.

In summary, ECP is off to a strong start and has vigorously responded to the seven ASCAC recommendations.

Meza asked about application integrity in terms of compute reproducibility and validation of applications. Messina agreed that validation and integrity are very important. The intention is to work very closely with software developers and to look at the effects of hardware faults on applications.

Levermore asked how links will be added horizontally to break silos in the WBS. Messina answered that there are a lot of face-to-face meetings. It is a contact sport. All the people have to know each other and work together. There are cross-cutting committees. Levermore asked how, in codesign, one avoids surprises. The vendors and users need to work together closely. The vendors will be driven by computer gamers. Specifications need to be drawn up early. Messina pointed out that codesign can be software as well as hardware. In hardware, one cannot change much, but there are things that can be managed and effected. What are desired are products that are real, marketable products for the vendors. One needs to be aware of what is coming down the pipeline.

Berzins liked the evolutionary approach, but it limits agility. He asked if there were any formal process for bringing in new approaches. Messina responded that the applications will drive innovation; targets are being set, and people can suggest changes. Berzins asked what the timeline after delivery might be. Messina replied that scientific results will be obtained and will provide information on what the next systems will look like. Facilities will take up the development process and continue, say, library development and software-environment development. Berzins said that that was not very clear and that a more-specific outlook would be good.

Chen asked if preliminary models could be seen and whether, in codesign, the WBS would allow cross-cutting discussion. Messina answered that a lot of codesign is spread out. The WBS can be modified. Alternative programming languages can be adopted.

Dunning noted that Messina had said that there would be few new applications, but that he expected that all applications would be new because they will have to deal with new environments and uses. The teams need to be realistic on what they can accomplish. Well-rounded teams are needed to consider all the issues. Messina pointed out that there is little funding for applications and that they were quibbling about the word “new.” The team has to have experience to re-implement codes. All applications are new in that they do not solve today’s problems but solve new problems. Some will just require higher resolution, but most will need new functionality. Dunning said that there is a great challenge ahead. Applications will not just be moved over from one system to another. Messina contended that it is better to fund a few applications and do them right rather than just sprinkling the funds around.

Gregurick asked how Messina saw this project moving forward with other agencies. Messina said that that will probably be done by the DOE offices and the interagency groups that they participate in. Cooperation with the NIH and NSF has been formalized in white papers last summer. Binkley added that, in principle, the Interagency Coordination Committee will allow cross-linkages with other agencies. In the next 4 to 6 weeks, the NSCI will decide on how to deal with questions like this.

Crivelli said that it is hard to express new information and asked how that problem will be dealt with. Messina answered that experts in the various areas will deal with those complex subjects. The applications will be relied upon to drive the needs.

Negele presumed that some of the preproposals would be asked for full proposals. Messina said, yes. For applications, invitations for full proposals would go out the week of this meeting and would be due on

May 6. For software technology, invitations for full proposals would go out in mid- to late April and would be due 5 weeks later.

The meeting was adjourned for lunch at 12:29 p.m.

Monday, April 4, 2016
Afternoon Session

The meeting was called back into session at 1:57 p.m.

Richard Carlson was asked for the ASCR response to the Next-Generation Networks for Science (NGNS) COV.

A COV was conducted in October 2016 to assess the processes used to solicit, review, recommend, and document proposal actions on the NGNS Program. The COV made 13 recommendations, and ASCR agreed with all of them. Specifically, the Program's management agreed to:

- Continue to evaluate meeting locations to encourage participation by the members of this community while maintaining control over the financial aspects of hosting a workshop
- Encourage cross-disciplinary research in network modeling by holding a workshop in examining the network protocol challenges that must be overcome to ensure that science communities can effectively use future network infrastructures
- Implement and deploy the Portfolio Analysis and Management System (PAMS) to support and document the complete research-funding process of SC
- Maintain a database of reviewers in PAMS
- Continue to use PAMS to track and manage funded projects
- Work with ASCR management to ensure that an appropriate level of travel funding is available for community-outreach activities
- Actively monitor and engage with the research community to keep abreast of current research activities and explore future opportunities
- Work with ASCR management to ensure that an appropriate level of travel funding is available
- Work with peers in other federal agencies to track and collaborate in multi-agency research programs through the Networking and Information Technology Research and Development (NITRD) Program
- Specifically call out in future funding opportunity announcements (FOAs) ESnet's 100-Gb/s network testbed, as appropriate
- Actively engage with the research community to advertise workshops and FOAs
- Interact with peers in other federal agencies to define and implement a coordinated set of strategic plans
- Continue to support PIs working in international standards groups [Internet Engineering Task Force (IETF), Institute of Electrical and Electronics Engineers (IEEE), and Open Grid Forum (OGF)] to increase the level of international collaborative efforts

There were no disagreements with the COV.

Williams asked if the COV had said exactly how to engage with the research community. Carlson replied, no. Williams asked if he had any ideas. Carlson answered, through PI meetings, NITRD community, and workshops. The Office is doing what it needs to do through its information infrastructure program.

Hey asked if there were any agency collaborations. Carlson answered, yes: DOE's perfSONAR [Performance focused Service Oriented Network monitoring Architecture] and NSF's DMZ [Demilitarized Zone]. The Office has dealt with the Department of Defense (DOD) and National Oceanic and Atmospheric Administration (NOAA) periodically. Hey noted that perfSONAR has been a major success for DOE.

Williams asked when the 400-Gb/sec network will be up and running. Carlson said that the vendor is putting together the infrastructure, and DOE is implementing and testing the system as funding allows. Software is being looked at to exert more control over the network.

Reed asked Chalk and Binkley to comment about the suggestion to combine COVs. Chalk answered that, in the NGNS COV, the question about the periodicity of COVs arose. Cerf suggested combining math, computer science, and NGNS into one COV conducted every 3 years rather than three COVs conducted one per year. Combining them would be in agreement with DOE rules and more consistent with other Office of Science programs. Williams commented that every 3 years would give a broader view, but coordination of personnel would be more difficult. Chalk noted that large COVs break up into subject areas and come back together to write the report. Normally, a COV operates over a summer and reports in November. It would put a large load on the COV chair. Levermore stated that, in a large COV, a final report is put together in about a month, with the writers staying an extra day after the meeting. Hey commented that it would be logistically tricky to get a large COV together. Meza said that, with a large COV, one cannot dig down into any project very deeply because there is such a broad subject area. It would be difficult on the DOE staff, also. Chalk pointed out that the PAMS would help organize and retrieve the required program information. Dunning suggested that one could focus on one program in depth and look at how all the other programs interact.

Williams asked when DOE needed to have a decision about this issue. Chalk replied, before the Summer meeting of ASCAC.

Levermore noted that NSF selects a date and then polls a large pool to find people who can meet on that selected date. They split into small groups; the chair then surveys the small groups and summarizes the threads of the discussions. Chalk said that another challenge is that one would have a difficult time getting a large group of reviewers who did not have a conflict of interest.

Reed said that he would confer with Chalk, and they would make a decision before the Summer meeting.

Patrick McCormick was asked to review NNSA's investigation of advanced programming models and runtime systems.

The Advanced Technology Development and Mitigation Subprogram was established to investigate advanced programming models and runtime systems for the exascale. The programming model seeks to achieve a balance among performance, portability, productivity/programmability, interoperability, and correctness in reproducibility. It was established to explore new models, shape and influence standards, and provide a tight coupling between code development and applications on the one hand and architecture and software development on the other.

Programming abstractions are important at the exascale. Today's programming models focus on control flow, explicit parallelism, and low-level data abstraction. Two node-level models were considered: Kokkos and RAJA. Both use C++ template metaprogramming for performance portability and functional-like programming techniques. Kokkos focuses on multicore central processing units (CPUs), Intel Xenon Phi, NVIDIA GPUs, IBM Power 8, and AMD Fusion. It relies on patterns, policies, and spaces to determine how the computations will be executed, where the computations will execute, where user data reside, and how user data are laid out in memory. In RAJA, loops are the main conceptual abstraction. It is lightweight and can be adopted incrementally. Its key abstractions are traversals and execution policies, IndexSets, and reduction types. The idea is to minimize traversals and allow programmers to color the data in partitioning it.

In looking at performance metrics, both Kokkos and RAJA increase the figure of merit. C++ metaprogramming produces impact on productivity resulting in long compile times, like executable sizes, and code-optimization challenges. This is the trade-off for performance and portability.

The challenge is to look at "model-awareness" inside the compiler. Functional programming is being considered, but today they still go through compilation. One needs to see whether to do things in a library or in the compiler. Kokkos uses semantics-aware code generation.

LLVM [the Low-Level Virtual Machine compiler-infrastructure project] continues to make strong inroads in high-performance computing. The ATDM initiated a project in FY-15 to fund NVIDIA/PGI to release their production front-end Fortran compiler to the LLVM community. It is expected to be available for community feedback, input, and contributions in late 2016. It is currently being tested with a small set of alpha testers.

Reed asked what will make this different from earlier attempts. McCormick replied, standards bodies' changing C++. Having LLVM available is better than black-box rules. Hey asked what their business model is. McCormick replied that they want to be responsible for the back end (with help).

One hears that data movement is expensive and compute is free, but idle processors are not free. With Trinity if one dumps data from memory to disk, one spends 10 times more power waiting for the data to move than to move the data. The strategy is to keep the processors busy.

The Legion programming model targets heterogeneous, distributed-memory machines. It is a tasking model and is a unit of parallel execution. Tasks must specify how they are going to use their regions, specifying privileges and coherence. Task launches follow sequential semantics with relaxed execution order. Next up is to have the task run the mapper, which picks up the needed pieces and hands them off to the processor(s). Mapping allows tasks to be written in different programming models and stays independent of tunings of applications. A lot of white space is eliminated by the mapper.

It is desirable to integrate external resources into the programming model; however, one cannot ignore the full workflow. Amdahl's Law sneaks in, and overhead increases 15 to 76% from the 2 to 12% for the original Fortran code. If one introduces new semantics for operating with external resources, these resources get incorporated into the deferred-execution model. Consistency is maintained between different copies of the same data, and the underlying parallel I/O is handled by Hadoop File System (HDFS) but scheduled by runtime. Thus, overhead is not affected, and applications are allowed to adjust the snapshot interval based on available storage and system-fault concerns instead of overheads.

Currently, there is a move on to refactor legacy Fortran code to use well-defined, pure functions/subroutines. The remaining aspects of the Legion–Fortran interface are a work in progress.

Sandia National Laboratories (SNL) is assessing leading asynchronous many-task runtime systems to inform the Advanced Technology Development and Mitigation (ATDM) subprogram's technical roadmap. Empirical studies show an AMT runtime can mitigate performance heterogeneity inherent to the machine itself. Message-passing-interface (MPI) and AMT runtimes perform comparably under balanced conditions. Previous experiments showed strengths of AMT runtimes for dynamic applications. Legion runtime needs hardening and lacks an application-facing application-programming interface (API). Uintah is targeted at Cartesian-structured mesh applications. Charm++ requires new abstractions and improved component implementations to realize its full potential.

DHARMA [Distributed asyncHronous Adaptive Resilient Models for Applications] is a portability layer for AMT runtimes that addresses key gaps. It insulates applications from the runtime system and machine architecture, improves application programmability, and synthesizes application requirements for high-performance-computing runtime-system developers. The idea is to make DHARMA an embedded domain specific language (DSL) with three layers: an application-facing front-end API, a translation layer, and a back-end API. There are active collaborations with industry, vendors, and researchers to help ensure success. Compiler teams and programming-model developers are improving support for C++ based encapsulation. Tool support needs to be developed for C++ templates and new models/tasking.

In conclusion, ATDM has used a holistic approach to tie together vendor research; Advanced Technology System (ATS) procurements; Predictive Science Academic Alliance Program II (PSAAP II); ATDM; ASCR; and research in algorithms, software, and hardware.

Berzins said that legacy Fortran is better than legacy C++. Scale and applications change and are not completely communicated to the proxy application. McCormick agreed that all the challenges out there have not been uncovered.

Chapman asked how much experience had been gained. McCormick replied that a lot of lessons have been learned with S3D code and 20 applications. A lot of errors were found that required one to go back and correct. A boot camp was held in December on the Legion Programming System, where it was learned that runtime optimization and compiler optimization could be tightly coordinated. Partitioning is hard to do. These are early growing pains.

A break was declared at 3:16 p.m. The meeting was called back into session at 3:45 p.m.

Robinson Pino was asked to review neuromorphic computing.

Recently, a grand challenge was issued to create a new type of computer that can proactively interpret and learn from data, solve unfamiliar problems using what it has learned, and operate with the energy efficiency of the human brain. SC convened a roundtable in collaboration with ASCR and BES. The roundtable convened a workshop and produced a report, *Neuromorphic Computing: From Materials to Systems Architecture*. The goal of the workshop was to support development of a new paradigm for extreme and self-reconfigurable computing architectures that go beyond Moore's Law and mimic neural biological computing architectures.

Neuromorphic computing has been defined as biological information-processing systems that operate on completely different principles from those with which most engineers are familiar. When input data are ill-conditioned, biological solutions are many orders of magnitude more effective than digital methods are because they use elementary physical phenomena as computational primitives and represent information by the relative values of analog signals rather than by the absolute values of digital signals. Adaptive techniques are used to mitigate the effects of component variation, leading naturally to systems that learn about their environment. Large-scale adaptive analog systems are more robust to component degradation and failure than are more-conventional systems, and they use far less power. Recently, emerging technologies like memristive systems and their potential integration with CMOS has ignited renewed interest in this field.

Memristor systems incorporate a memory-resistor in two-terminal passive devices. A memristor, the functional equivalent of a synapse, could revolutionize circuit design. The technology could have significance in large-scale parallel computing, high-density nonvolatile memory, inherent radiation hardness, and optimum power/size/weight ratios. Parallel architectures can perform many computations simultaneously moving from serial to parallel computation. Traditionally, software has been written for serial computation, which has the disadvantage of excessive communication between computing nodes, degrading performance.

Currently, about 5 to 15% of the world's energy is spent in some form of data manipulation, transmission, or processing. Conventional computing devices fail in some of the most basic tasks that biological systems have mastered, such as language and vision understanding. It has been predicted that conventional approaches to computation will hit a wall and become critical in the next 10 years. Fundamental (atomic) limits exist beyond which devices cannot be miniaturized. Local energy dissipation limits device packing density. The increase in overall energy consumption is becoming prohibitive. Thoughtful approaches and new concepts are needed to achieve the goals of developing increasingly capable computers that consume decreasing amounts of power. There is a need for enhanced computing.

DOE has charted a path to exascale computing by early in the next decade. Even though they will be incredibly powerful, these machines will still consume between 20 and 30 MW of power and will not have intrinsic capabilities to learn or deal with complex and unstructured data. The mission areas of DOE in national security, energy sciences, and fundamental science will need even more computing capabilities than what can be delivered by exascale-class systems. Neuromorphic computing systems are aimed at addressing these needs.

Data science supports research involving complex and unstructured data. Machine learning has been applied in many areas in DOE where traditional methods are inadequate (e.g., climate models, recognizing features in large-scale cosmology data, and predicting maintenance needs for accelerator magnets). Traction is being gained in designing materials and predicting faults in computer systems. Machine learning has affected nearly every major DOE research area. Neuromorphic computing may even play a role in replacing existing numerical methods, where lower-power functional approximations are used, and could directly augment planned exascale architectures. What areas of science are most likely to be impacted, and what might the requirements be? This roundtable did not focus on applications.

Among state-of-the-art systems, IBM's True North chip has research groups using it experimentally; common architectures arrange synapses in crossbars. A challenge for such systems is information in/out ratios. Biology can have up to 2×10^4 synaptic connections per neuron; the current state in CMOS is 256.

The envisioned research path calls for developing potential architectures inspired by biology, enabling accelerated platforms or prototypes, developing software for extreme-scale simulation and verification, and scaling promising architectures into end-to-end systems.

The main conclusions of the roundtable were

- Creating a new computational system will require developing new system architectures to accommodate all needed functionalities.
- One or more reference architectures should be used to enable comparisons of alternative devices and materials.
- The devices to be used in these new computational systems require the development of novel nano- and meso-structured materials. This last would require (1) unlocking the properties of quantum materials based on new materials physics and (2) designing systems and materials that exhibit self- and external-healing, three-dimensional reconstruction, distributed power delivery, and colocation of memory and processors.
- The development of a new brain-like computational system will not evolve in a single step.
- Successfully addressing these challenges will lead to a new class of computers and system architectures.
- Future computing systems with these capabilities will offer considerable scientific, economic, and social benefits.

In summary, this initiative is high-payoff R&D and a foundational technology for systems able to perform autonomous operations, sense-making for scientific discovery, and accelerated and reconfigurable in situ analysis. There is plenty of opportunity to make substantial contributions in the area of information processing and understanding. Brain-inspired architectures in software or hardware offer a path for much-needed technological evolution. It is a truly multi- and cross-disciplinary effort in materials, physics, chemistry, biology, mathematics, engineering, computer science, neuroscience, etc. to ensure success in this field.

Crivelli asked if people from multiple disciplines were going to submit ideas and listen to neuroscientists. Pino replied, yes. There should be more workshops than just this one.

Dunning noted that this technique could solve different types of problems, but there are classes of problems that it could not address. Entanglement is where a lot of work needs to be done, and it holds great promise. He asked what DOE will be exploring in addition to this technique to get beyond the exascale. Pino responded that this is not a universal computer. It addresses the problems of science. Today, a lot of data does not get analyzed because there are not enough computing resources. Going from megawatts to watts comes from low energy consumption for operation, and one can reconfigure it. It will never replace digital logic. However, it can make a large contribution in some situations. It is needed. A human will still be needed to analyze the data. Having said that, there are lots of places where it could help DOE, such as in monitoring and analyzing the operations of the electrical grid and looking for patterns of incipient failure.

Reed commented that a criticism of deep-learning computing is that it is not known how it works; however, it is not known how human learning works, either. Pino said that one can write equations and algorithms, but not everyone will understand how they work.

Hey asked if DOE were interested in funding something like SpiNNaker. Pino responded that making hardware is very expensive. Making hardware with new types of materials is even more expensive. DOE is funding a lot of materials science, and BES talks about single-atom control. There is a gap between that and ASCR's research portfolio. ASCR is not funding work on circuits. The question arises: Can ASCR spend money on advanced materials and circuits for studying their material properties? How would DOE benefit from making that advance? DOE does not build hardware. True North came out of the Defense Advanced Research Projects Agency, but when one looks at True North software, one has to sign a nondisclosure agreement. The problem space is very new, and getting the research community to spend time understanding it is difficult. Basic science must be conducted in a way that takes on risk.

Michael Martin was asked to discuss future directions for DOE's Advanced Computing Technology Team.

Tech teams and crosscuts exist to address the same challenge: How does one get different parts of the Department that face similar challenges to work well together? Tech teams are unfunded vehicles for coordination. Crosscuts pool together funding and have congressional reporting requirements. The Tech Team has made efforts for five years to move high-performance computing to a cross-cut status. The decision was made that the jump to a crosscut could not be made. If there were no need for a crosscut, would there be any need for a tech team? The QTR pointed out existing partnerships, and the need for more. Networking with the Federal Information Technology Acquisition Reform Act (FITARA) showed that a mechanism is needed for all programs to talk to each other, and the team has showed that the interest is there. A 2014 survey of needs showed there was an advantage to keeping the tech team going. In some cases, applied groups faced exactly the same challenges. In other cases, problems are analogous, with the potential to learn from shared approaches. While the physical nature of the problems varied, applied groups pointed to four basic sets of needs: high-performance computing for discovery and fundamental physics, using high-performance computing for model reduction, system/plant-level simulation, and post-design simulation.

There is an interest in using high-performance computing at the innovation/discovery stage on a range of topics: materials design, fluid mechanics, heat transfer, and combustion physics. It needs to be remembered that the terms "discovery" and "innovation" are often understood differently by SC and the applied programs, which can lead to gaps in both the physical understanding of the problem and the computational tools available to attack the problem.

Moving from discovery to design requires a different approach. The goal here is the highest-fidelity simulation that allows the engineer to explore the parameter or design space. Often, the length and timescales of the problem may require moving from first-principle simulations to models (knowing that all models are wrong, but some models are useful). The question is: How does one use approaches that require leadership-class simulation in an integrated manner with modeling?

Many technologies have problems that require multiple subsystems to be simulated to capture interactions. In post-design simulation, for example, engineering simulation is used to compare options and pick the most effective one. Accurate simulations may contribute to deployment of new technologies by improving the assessment of risk in three areas: performance, insurance, and regulation. A high-fidelity, validated simulation capability can produce improved performance prediction and more-accurate risk assessment. These two improvements lead to less financial risk, a lower cost of capital, lower insurance costs, and quicker regulatory approval. All of these lead to cheaper, more rapidly deployed technology. The situation, however, requires simulations that financiers, insurers, and regulators trust.

The Tech Team's primary goal is to allow communication across the Department. The goal is for the scheme to highlight approaches that have the potential to jump between programs. The Team is moving toward regular, monthly meetings, with 10- to 15-min updates and 30- to 40-min deep dives. Adding more national-laboratory members will broaden participation. The goal is now to enable collaboration and partnership across programs and national laboratories rather than *to be* the partnership.

Reed asked what the major challenge was. Martin replied, like any collaboration, the need is to get a standard language and to set common priorities so one can find gaps.

The floor was opened to public comment. There being none, the meeting was adjourned for the day at 4:52 p.m.

Tuesday, April 5, 2016

The meeting was called to order at 8:36 a.m.

Steve Binkley was asked to present a summary of quantum computing.

Quantum computing is a subset of quantum information science. A decade's worth of research is needed to get beyond the CMOS computing regime. Quantum computing will not replace but will augment CMOS (big iron) computing.

The origins of quantum computing can be traced to remarks and papers by Richard Feynman, to wit, if one has something solvable by the Schroedinger Equation, one will need a quantum computer. Peter Shor said that, with a quantum computer (assuming that one can be built), one can factor an N-digit number in about N^3 steps. In 1994, Cirac and Zoller proposed the first controlled-NOT (two-qubit) gate for trapped ions. In 1995, Haroche suggested that quantum computing will not work because one cannot correct quantum errors, but Shor and Steane immediately proposed quantum-error correction. In that same year, Monroe and Wineland realized the first 2-qubit quantum logic gate and got the Nobel Prize in Physics for 2012.

It is difficult to make qubits. Quantum devices have to be or have a scalable physical system with well-characterized qubits (i.e., there has to be reproducible manufacturing). In addition, there must be (1) an ability to initialize the state of the qubits, (2) long coherence times relative to gate-operation time, (3) a universal set of quantum gates, and (4) a qubit-specific measurement capability.

There have been many forms of qubits proposed: ion traps, superconductors, nitrogen vacancy centers in diamonds, quantum dots (phosphorus atoms in a silicon substrate), and materials with topological properties that are self-protected against coherence problems (which makes it difficult to get information in and out).

With quantum computing, software looks nothing like today's programming languages. Progress in this field spans multiple disciplines (computer science, mathematics, material sciences, quantum engineering, and high-energy physics).

Quantum computers could break all present-day public key encryption systems, but quantum encryption would not be susceptible to computational attack. Quantum computing will allow more-direct solutions for materials design, pharmaceutical design, chemical processes, etc.

Classical bits are on or off. Quantum qubits have state superposition with on and off and space (a unit vector). As a result, classical bits store a single value between 0 and 2n – 1, while qubits can hold 2^n values.

A notional quantum computer would more closely resemble a massive computer-controlled quantum-physics experiment than a classical computing engine. It would have a large, external, classical control computer that would drive a quantum core. How far in the future is this? Five years is not likely, but ten years is possible. A small-scale version could be built to explore quantum computing. Quantum computers could solve linear equations or run search algorithms that could be extended to other types of optimization problems. Problems that are not tractable today, such as calculating the ground state of a complex molecule like ferredoxin, could be solved with only 200 qubits.

Several workshops on quantum computing have been held, including a DOE/ASCR 2015 workshop that explored mission relevance, the impact on computing, and challenges of quantum computing. These workshops have concluded that quantum-computing testbeds could be established to explore (1) algorithms and computational approaches and (2) scientific applications that are important to DOE missions that can be attacked with quantum algorithms. Quantum computing would be applicable to linear algebra; integration and summation; optimization; and graph theory.

The end of the Moore era will end the exponential increase in data-gathering capabilities seen over the past four decades. The development of a new generation of detectors and sensors based on novel quantum-entanglement approaches has the potential to sustain needed increases in detector and sensor performance as well as to open new frontiers in sensor and detector technologies. The FY17 objective is to initiate R&D in quantum measurement to develop technologies that could impact next-generation scientific facilities across numerous SC offices. The proposed activity is consistent with and coordinated with the broader, government-wide, interagency quantum-information science activity, specifically that at the National Institute of Standards and Technology (NIST) and NSF.

Hey pointed out that there have not been any algorithm breakthroughs since 1996. A lot needs to be done. When one makes a measurement, one gets one value; researchers have not been able to find any other general-purpose algorithm than Shor's. Entanglement is the magic. Binkley replied that quantum computing gives one superposition and entanglement. Einstein asserted that quantum mechanics cannot be correct. Entanglement is where a lot of work needs to be done, and it holds great promise. Heisenberg's uncertainty principle assumes that uncertainty is distributed homogeneously. However, that is not necessarily correct. Hey said that John Bell said that Einstein could be right if there were a hidden variable with certain conditions. Bell objected to this kind of thing. Quantum states evolve with the Schroedinger Equation, but actually, when the observer makes a measurement, it jumps into an eigenstate, and that is not described by quantum mechanics. The conventional argument is that interactions with the environment will interfere with the measurement. Bell was not happy with that. He called it Heisenberg's "shifty split" in which kinematics is given a wavefunction for the quantum part, classical variables (which have real values that can be measured) for the classical part, and "collapse recipes" for the regions of interaction. But the system is made up of atoms and electrons, so one can push the boundary further back. But when one *does* push it back, one finds that it does not collapse in the expected place but in another place. There is no mechanism for actually doing the collapse. It can be in the variables that one is looking at; in the classical system it might appear that there is no collapse. There is no satisfactory opinion about this situation. As a result, quantum computing could push quantum mechanics into an interesting area.

Berzins said that it is important to have a quantum-computing program. With big data, there would be a lot of applications. Binkley agreed and said that there was a lot of interest in large data across SC (e.g., climate science and cosmology). A lot of researchers outside of DOE are going in that direction, also.

Rick Stevens was asked to describe the BRAIN [Brain Research through Advancing Innovative Neurotechnologies] Initiative.

There are two partnerships between DOE and NIH: Cancer and the brain.

The NIH could use assistance in high-performance computing from DOE in two areas: cancer research and the BRAIN Initiative.

DOE currently has a partnership with the National Cancer Institute (NCI) to advance cancer research and high-performance computing in the United States. The problems being addressed fall into three categories:

- Cancer biology, focusing on the explication of Ras oncogene pathways with unsupervised learning and mechanistic models to develop an understanding of the mechanism and to identify drug targets. Some 30% of cancers have mutated Ras, producing about 1 million deaths per year. The use of simulation could lead to understanding the variations in the Ras molecule. The heart of this effort is machine-learning-guided dynamic validation.
- Preclinical cell culture/animal models to speed the development of drugs and the production of drug responses. A new investigative model is patient-derived xenograft models, in which reprogrammed cell lines are grafted into another species (the mouse) to get 100 mice that are proxies of the human patient, allowing hundreds of experiments to be conducted on the proxy mice. These experiments produce masses of data that can be subjected to machine-learning-based predictive models, uncertainty and optimal experiment design, and hypothesis formation and mixed modeling.
- The Cancer Surveillance Program, which has data on 20 million patients, needs those data to be interpreted and abstracted into agent-based modeling of cancer-patient trajectories. The database will be mined to produce models that can be used to track treatment application and outcomes.

A National Strategic Computing Initiative (NSCI) public–private partnership of biopharma, national laboratories, universities, technology, and government is emerging for the development of computing precision medicine.

The NCI–DOE partnership will extend the frontiers of DOE computing capabilities in simulation, data analytics, and new computing architectures. This is a large-scale codesigned effort. A partnership with NCI could become a model for other work within NIH. The partnership will allow DOE to integrate simulation, data analytics, and machine learning in addressing a complex-features problem. Supercomputers and exascale systems will address deep learning and scalable data analytics while traditional high-performance-computing systems will address large-scale numerical simulation. (It is expected that one could get convergence with neuromorphic constructs.) A recently published paper indicated that deep learning outperformed all other methods with respect to the area under the curve of accurate measure of certainty and was significantly better than all commercial products.

The ultimate goal is to use mechanistic biological models and machine learning in concert to produce hybrid models that will, in turn, produce predictive performance and insight. How the mechanistic models and machine learning are to be integrated is what needs to be solved.

The White House BRAIN Initiative is another area of possible cooperation between DOE and NIH. The challenge here is to (1) map the circuits of the brain, (2) measure the fluctuating patterns of electrical and chemical activity flowing within those circuits, and (3) understand how their interplay creates humans' unique cognitive and behavioral capabilities. This goal should be pursued simultaneously in humans and in simpler nervous systems that we can learn important lessons from more quickly. The BRAIN Initiative aims to help researchers uncover the mysteries of brain disorders, such as Alzheimer's and Parkinson's diseases, depression, posttraumatic stress disorder (PTSD), and traumatic brain injury (TBI). In some ways, this is similar to the Human Genome Project, which produced $800 billion in economic return (so far).

The goals of this Initiative are to discover diversity in cell types, produce maps at multiple scales, observe the brain in action, demonstrate causality of behavior by brain activities, identify the fundamental principles, advance human neuroscience, and integrate the results with and translate them to medical practice.

An earlier effort, the NIH Blueprint for Brain, started in 2004 but was only modestly successful. A reboot of the NIH effort has subsequently held six workshops with about 100 participants each, but they included no computer scientists or mathematicians despite the fact that they had discussions on big data. In FY16, NIH contributed $135 million, DARPA contributed $95 million, NSF contributed $72 million, and other funds came from the Intelligence Advanced Research Projects Activity (IARPA) and the Food and Drug Administration (FDA).

The brain has 100 billion neurons, 100 trillion synapses, and one zettabyte in "Google brain map" (about the annual global Internet traffic). The scales of the structures of interest in the brain run from 10 cm to 1 μm.

To investigate the connectome, one must prepare the tissue, section and prepare the wafer, acquire images, perform registration of the images, detect segmentation and synapses, and visualize and analyze the data. High-throughput imaging is needed. Currently, a 61-beam parallel electron microscope produces about 12 TB of data per day or 2 petabytes in about six months.

It is now believed that there may be 1000 types of brain cells that need to be identified and investigated. Instead of slicing and photographing the brain, one can use the Advanced Photon Source and X-rays to image and characterize the connectome. The visualization technique NeuroLines is used to perform neuronal connectivity analysis to determine a wiring diagram, and BigNeuron [a community effort to define and to advance state-of-the-art single neuron reconstruction algorithms] is used for automated image reconstruction for large-scale phenotyping of neuron morphologies.

But critical disease processes take place at scales that simply cannot be seen, even by today's best tools. Optogenetics allows visualization of neural activity. The functional connectome represents the structure of brain functions and a weighted, direct graph describes the dynamic, casual interactions among neurons in the functioning brain. Thus, visualization is being furthered, and brain-mapping tools are being developed.

ASCR can uniquely contribute to the brain initiative through advances in applied mathematics and computer science together with high-performance-computing facilities for probing and developing function, theory and models, and structure. Top-level computing opportunities include large-scale data analysis, large-scale predictive modeling, and large-scale explanatory modeling. These efforts could be performed at the national laboratories to build flexible teams that would cross disciplines and laboratories; support sustained technology development to reach a goal that involves partners; provide large-scale project management; provide production-quality software development and software engineering; build user communities around new scientific capabilities; and integrate across multiple domains and facilities.

NIH already taps into national laboratories via university grants and contracts. NIH-supported PIs are users of laboratory user facilities through direct funding/hosting arrangements and agency-to-agency arrangements. The future will call for a flexible framework for larger-scale partnerships that tap into the capabilities and culture of the national laboratories in a more-direct fashion.

To truly understand what something means, one can build models that predict future states or outcomes and can give insights or explanations of the behavior of the system. To do this often requires integration of knowledge from many sources into a coherent, computable representation that can be used to test hypotheses and conjectures. In this sense, computing and modeling collaborations often play the role of grand scientific integrators. This is true in DOE mission space and also is often true in NIH mission space. Computing collaborations with DOE could improve how NIH integrates science across domains, institutes, and projects.

Levermore said that the interplay between large data and modeling is appropriate for DOE to be engaged in and needs to be transferred to NIH. He asked if weather forecasting is a parallel experience. Stevens replied that the NIH is inspired by the climate-change community. DOE has found a legitimate reason to work on climate change and collaborate. Congress raises questions about why DOE is working on the brain.

Crivelli said that bioimages in breast cancer show beautiful networks of filaments indicating that something is going on. It is important to gain the context of the tumors. Stevens replied that one of these strategies/problems is integrating images and modeling. If imaging can replace bioassays, much more information can be gleaned and comparisons allowed across cell types and times of disease progression.

Dolbow noted that mechanical models are not predictive. He asked if they are too slow to be useful. Stevens said that it is desirable to bring uncertainty quantification and rational design to modeling. There are stochastic pathway models that can lend to understanding the required time samples needed to effectively model.

Chapman said that the European Human Brain Project is more similar to the U.S. program than generally recognized. Europe had medical and computer people in the same institutions, which was beneficial. European computing science learned what future computer architectures were needed to be responsive to the neuroscience researchers' needs. Stevens answered that the United States is building an environment that makes bioscientists more productive.

Berzins pointed out that some researchers in Western Europe have crowd sourced data collection and need access to DOE computing resources. Someone should reach out to those researchers. Stevens replied that there are programs like the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program that make allocations available to such researchers. There are also LCF allocations that are available to DOE researchers.

Chen said that a codesign center on data analytics is similar to other centers on deep learning at the national laboratories. Stevens acknowledged that there is a lot of conversation going on and asked whether codesign would address new architectures, new machines, convergence, or what. The exascale is being launched, and its roadmap could accommodate such codesign efforts, but one does not want to disenfranchise the traditional users of these systems. A model of the entire system itself is needed. That would show what is missing, where work needs to be done, etc. This is becoming the age of constructionism after the end of the age of reductionism.

Gregurick said that this pilot is very ambitious. She hoped that Stevens would come back to give an updated progress report. NIH needs advanced computing to gain insight into their problems. Stevens replied that they are using computers to do common tasks. DOE can help them on that level, but the gain needs to be turned up on what is being done even now.

Hey said that the functional stuff is what needs to be done. What can be learned from small systems needs to be identified. Stevens said that a shrew has a brain that is only 1 cm^3 in volume. Almost all of the shrew's sensory system is in 10,000 whiskers that are connected to that brain. The structure of the connectome is being investigated. It is hoped to get the basic tools for this investigation built and to show how they can be applied.

Hey noted that some are skeptical about the ability of deep neural networks. Stevens said that he was skeptical, too. It has limitations that require complementary techniques. Hey said that Bill Gates said that Microsoft was the only place where deep learning is possible because universities do not have the resources needed to process the big data. This is where DOE could enable the universities.

Levermore said that, sometimes, physical or biological systems have structures that deep learning cannot make use of. Stevens replied that soon there will be deep-learning applications and basic science research.

Reed noted that a lot of these things are at the boundaries of fields and require collaboration. Stevens responded that one has to be fearless and learn the languages of all the disciplines. One must be deeply humbled to make headway. There have been people at the national laboratories doing optimization, which is applicable to deep learning. Half of the people at those laboratories working on the BRAIN Initiative are physicists. More science-savvy and medicine-savvy computer scientists are needed. Maybe ASCR needs to support a sabbatical program so computer scientists can learn about biology and medicine.

A break was declared at 10:40 a.m. The meeting was resumed at 10:54 a.m.

Barbara Helland was asked to describe the laboratory planning and laboratory appraisal processes.

DOE is the steward of 17 national laboratories across the United States. DOE has been asked for a management plan for each national laboratory. SC had already recognized the importance of the planning process as an element of its stewardship responsibility for these laboratories. Laboratory leadership teams are asked to define a long-range (10-year) vision for the institution. This derives a starting point for discussion between SC leadership and each laboratory about the laboratory's future direction. The process starts with the Director of SC setting laboratory planning expectations. This document is sent to the laboratories to review their core capabilities and assignments. Draft guidance is issued, which is reviewed and upgraded before final guidance is issued. The laboratories submit their draft plans, and each program office reviews those laboratory plans. Annual laboratory planning briefings are held, leading to the publication of the final laboratory plans.

In FY15, ASCR was responsible for reviewing the following core competencies of the national laboratories: advanced computer science, visualization, and data; applied mathematics; and computational science. The laboratory plans present mission/overview, the laboratory at a glance, current laboratory core capabilities, science strategy for the future and for major initiatives that they would like to do, work for others, infrastructure/mission readiness, human resources, and cost of doing business.

In 2016 there are some changes to this process. Plans will include a new section on each laboratory's R&D computing and commodity information technology (IT). R&D computing is defined as computation resources, ranging from individual data-collection computers to supercomputers used in scientific discovery or in other roles that directly support the laboratory mission. Commodity IT is defined as conventional IT resources similar to those that would be found in any operational business. For each of these categories, the national laboratories are being asked to submit an overview that includes (1) existing systems, including networking capabilities; (2) a table of planned major procurements; and (3) a gap analysis pointing toward future needs. In the future, this requirement will be expanded to include the energy laboratories [Idaho National Laboratory (INL), National Energy Technology Laboratory (NETL), and National Renewable Energy Laboratory (NREL)]. The laboratories will be briefed on the planning process in June and July. These meetings will include the chief information officers and will cover any work for others of more than $1 million.

Since 2006, SC has conducted a yearly evaluation of the scientific, technological, managerial, and operational performance of the contractors who manage and operate its 10 national laboratories. This appraisal process is designed to improve the transparency of the process, raise the level of involvement by SC leadership, increase consistency in the way that the national laboratories are evaluated, and more effectively incentivize contractor performance. There are eight performance goals: mission accomplishment; design, construction, and operation of research facilities; science and technology project/program management; leadership and stewardship of the laboratory; integrated environment, safety, and health protection; business systems; facilities maintenance and infrastructure; and security and emergency management. The Office gives feedback to the laboratories for the first three of these performance goals. The site offices give feedback to the laboratories for the remaining five performance goals. For 2016, the established goal/objective language includes assessment of public-access support and, specifically, submission of accepted manuscripts as a result of ASCR's recommendations for OSTI.

The Performance Evaluation and Measurement Plan offers opportunities for SC programs and site offices to identify a small number of notable outcomes that illustrate or amplify important features of a laboratory's performance for the coming year. Performance goals, objectives, and notable outcomes are documented and appended to the respective laboratory contracts at the beginning of each year. The SC appraisal process uses a 0 to 4.3 scoring system with corresponding grades for each performance goal and objective. As an example, for 2015, ASCR produced a notable outcome for NERSC and the LCFs to define a program-review process to collect exascale application requirements from the Office of Science programs by the end of FY16 and to present the plan to ASCR by February 27, 2015. The first exascale application program review was held in September 2015.

When the Office observes a problem at a national laboratory, this planning process is seen as a way to help that laboratory.

Reed asked why the limit was 4.3. Chalk said that it was a compromise. Helland pointed out that the scores are normalized SC-wide.

Williams asked what happens to the items not released to the public. Helland answered that DOE uses these items to give feedback to the national laboratories.

Chen asked who defined the competencies of the national laboratories. Helland replied that the national laboratories propose their competencies, and the associate laboratory directors review these requests and validate them.

Levermore noted that many agencies use the National Research Council to review large-expenditure programs and asked whether DOE does that. Chalk said that some programs do that, but most do not because it is a long and expensive process. Levermore pointed out that it lends to the long-term credibility before Congress of the agency's stewardship.

David Brown was asked to present an update on the Computational Research Leadership Council (CRLC).

The directors of computing at the eight multipurpose national laboratories have gotten together to form a leadership council. Each of these national laboratories has a computer science division headed by an associate laboratory director and containing a computational research division. About 1000 employees are engaged in computer science at the eight national laboratories. Today, math and computing play an essential role in all areas of science in high-fidelity modeling and simulation and in the experimental and observational data spaces. The computational research divisions have significant areas of common interest: They develop and maintain research programs and a skilled workforce to support DOE's computing and applied math needs. They are significant stakeholders in the ASCR research program and the Exascale Computing Project (ECP).

Now, management and funding-model changes associated with the startup of the ECP represent a challenge for all of these divisions, and it is helpful to know what the issues are at all the national laboratories. The CRLC was formed in January 2016. It is made up of representatives from the computer science divisions of the eight national laboratories. The purpose of the CRLC is to (1) identify and address issues of common interest to the computational research divisions (workforce development, research roadmaps, major initiatives, workshops, and collaborations); (2) represent these interests to major stakeholders; and (3) work in partnership with stakeholders on major initiatives, such as the ASCR research roadmap and the ECP laboratory leadership.

Under the CRLC charter, members are appointed and annually reconfirmed by the associate laboratory directors for computing annually (one per national laboratory). The chairmanship rotates annually through the national laboratories alphabetically. The initial meeting was January 11, 2016. The plan is for the CRLC to have two face-to-face meetings per year. Council members have weekly teleconferences that include other ASCR points of contact at the national laboratories. The Council may establish subcommittees. It is meeting with ASCR research leadership to contribute to developing a long-range plan for the research program. To that end, it has had several meetings with ASCR personnel, conducted a deep dive on convergence of high-performance computing and data-intensive science, and holds regular bi-weekly teleconferences. The Council has also been meeting regularly with the ECP laboratory leadership. Topics include keeping communication lines open between the project management and the laboratory research divisions, identifying laboratory capabilities relevant to ECP software technology, and reviewing ECP software technology white papers.

Williams asked if there were counterparts to the Council in BES or BER and, if so, how ASCR cooperates with them. Binkley replied that there is something like this in FES, and there has been something like this in BES, but they were not laboratory-specific like this one is. Williams asked if the Council worked out the research plan with ASCR. Binkley said that a distinction is made between long-term and strategic planning. Chalk added that long-range planning is community driven.

Berzins asked if there were things that should be done differently to ensure the needed workforce for the exascale. Brown replied, yes; the workforce needs to be mapped onto the expected activities, and DOE will need to hire new personnel with specific, needed skills. Binkley added that everyone is deeply engaged in workforce development. The offers being made to new hires put a strain on the national laboratories. DOE is always looking for ways to ease that pressure, which is not going to decrease. Brown said that it used to be that people worked at a national laboratory for 40 years. Today, the national laboratories have to recruit all the time as the workforce turns over.

Chapman noted that there was difficulty getting data from the national laboratories in drawing up the workforce-development report. Brown said that that issue could be addressed by the Council.

Gregurick asked whether the Council's reports were for ASCR or for the public. Brown said that these are mostly internal discussions.

Crivelli asked if a discussion could be added on why people do not stay a long time. Brown agreed that this issue should be understood better. Only anecdotal information is available. The University of California has a retirement system that incentivizes people to stay there. The national laboratories do not have such a system.

Levermore pointed out that ASCI was also designed as a workforce-development program. The employment environment has changed, and employment strategies need to be reconsidered. Brown agreed. There are diversity issues that need to be addressed at the undergraduate level. Internship programs have also been successful.

Hey asked if the program got any credit for or equity in any spinoffs. Binkley answered, not at Lawrence Berkeley National Laboratory. It can vary from laboratory contract to laboratory contract.

Chen said that many new hires are looking to work at national laboratories so they can work on cool projects. Fellowships to go back to school and the use of Laboratory Directed Research and Development (LDRD) funds to allow people to follow their research interests can aid retention, also. Brown said that there are programs for sending employees back to school at Lawrence Livermore National Laboratory, Lawrence Berkeley National Laboratory, and others.

Levermore said that bright students are attracted to data science. DOE's leadership role in data-driven research, advanced computing, etc. should be advertised. Brown said that DOE has never been great at letting people know what it does. One needs to get people in the door so they can see what unique things are going on.

Chapman said that the ASCAC study found that little was known at universities about DOE. Reed said that that is a worry. There is a marketing issue here. Brown said that going out to colleges and talking is a good way to inform students.

Reed opened the floor to final comments from the Committee. Negele pointed out that named fellowships have an attractiveness to them. Invoking the National Research Council in independent reviews strengthens one's case.

Reed announced that Frederick O'Hara would be retiring as the ASCAC Recording Secretary after this meeting and thanked him for his long service to ASCAC, dating back to the Committee's very first meeting in November 2000. [Round of applause.]

Larzelere noted that the Committee had heard about beyond-CMOS, which is devoted mostly to the Department of Defense and the intelligence community. He asked if ASCR were getting any pushback from those agencies for its efforts in this area. Binkley replied, no. All of these communities are very collaborative about that topic.

Chalk asked for suggested dates for the summer meeting of ASCAC.

There being no further discussion, the meeting was adjourned at 11:59 a.m.

Respectfully submitted,
Frederick M. O'Hara, Jr.
ASCAC recording secretary
April 25, 2016