DRAFT

Advanced Scientific Computing Advisory Committee Meeting
American Geophysical Union, Washington, D.C.
August 24-25, 2010

ASCAC members present:

F. Ronald Bailey
Marsha J. Berger
Jacqueline H. Chen
Jack Dongarra
Roscoe C. Giles, Chair
Susan L. Graham
James J. Hack
Anthony J. G. Hey
John W. Negele (by telephone)
Linda Petzold
Larry L. Smarr (by telephone)
Vivek Sarkar (by telephone)
William Tang
Robert G. Voigt
Victoria A. White

ASCAC members absent:

Thomas A. Manteuffel

Also participating:

Melea F. Baker, Office of Advanced Scientific Computing Research, Office of Science, USDOE
William Brinkman, Director, Office of Science, USDOE
Christine Chalk, ASCAC Designated Federal Officer, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Dona Crawford, Associate Director for Computation, Lawrence Livermore National Security
Ping Ge, Office of Workforce Development for Teachers and Scientists, Office of Science, USDOE
Barbara Helland, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Daniel Hitchcock, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Thuc Hoang, Advanced Simulation and Computing Division, National Nuclear Security Administration, USDOE
Douglas Kothe, Director of Science, National Center for Computational Sciences, Oak Ridge National Laboratory
Alexandra Landsberg, Computational Science Research and Partnership (SciDAC) Division, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Lucy Nowell, ASCAC Committee Manager, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Frederick O'Hara, ASCAC Recording Secretary
Walter Polansky, Office of Advanced Scientific Computing Research, Office of Science, USDOE
James Roberto, Deputy Director for Science and Technology, Oak Ridge National Laboratory
Rachel Smith, Oak Ridge Institute for Science and Education
Rick Stevens, Associate Laboratory Director for Computing, Environment, and Life Sciences, Argonne National Laboratory
Michael R. Strayer, Associate Director, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Alvin Trivelpiece, Chairman, Exascale Workshop Panel
Andrew White, Deputy Associate Laboratory Director; Theory, Simulation, and Computation Directorate; Los Alamos National Laboratory

About 70 others were in attendance in the course of the two-day meeting.

Tuesday, August 24, 2010
Morning Session

The meeting was called to order at 9:00 a.m. by Chairman **Roscoe Giles**. The Designated Federal Officer (DFO), **Christine Chalk**, took the new Committee member, Susan Graham, to the Forestall Building to be sworn in. **Lucy Nowell** served as the DFO pro tem until Chalk returned. Giles reviewed the agenda. **Rachel Smith** made convenience and safety announcements. Three members attended by a telephone.

William Brinkman was asked to give an overview of activities in the DOE Office of Science (SC). President Obama wants the country to increase science funding and has been delivering on that promise. "When we fail to invest in research, we fail to invest in the future. Yet, since the peak of the space race in the 1960s, our national commitment to research and development has steadily fallen as a share of our national income. That's why I set a goal of putting a full 3 percent of our Gross Domestic Product, our national income, into research and development, surpassing the commitment we made when President Kennedy challenged this nation to send a man to the moon."

He reviewed the organization of SC and of the offices under the Under Secretary for Energy. SC has six research offices and the Office of Workforce Development for Teachers and Scientists (WDTS). SC's priorities are in scientific computing and climate science. Its requested budget for FY11 is $5.1 billion, a 6.1% increase. The 2011 budget has now come out of the congressional committees. SC took a heavy hit in the House markup; the Senate took $109 million away from the request. The Office of Basic Energy Sciences (BES) took a heavy hit, a loss of $95.9 million in the Senate markup. There were some earmarks (e.g., for the artificial retina), and nuclear medicine was moved to the Office of Nuclear Physics. There will likely be a continuing resolution for the beginning of FY11.

The SC Graduate Fellowship Program is very important. $10 million is needed in FY11 to fund about 170 additional fellows. About $16 million would be available in FY11 to fund about 60 additional Early Career Research Program awards at universities and DOE national laboratories. These budget changes reflect the reversion to a more normal budget after the American Recovery and Reinvestment Act (ARRA) infusion. An effort will be made to try to have these cuts restored.

The ultimate goal for the Fuels from Sunlight Hub is to imitate photosynthesis with a factor-of-10 increase in productivity. The winning team was California Institute of Technology and Lawrence Berkeley National Laboratory (LBNL); it will be led by Nate Lewis and will partner with six other institutions. The Hub will push to integrate processes to form a complete, useful system. The Department is pushing for a Battery and Energy Storage Hub that will deal with the problems produced by intermittent energy sources (wind, solar, etc.). It is in the Senate version of the FY11 budget.

The Office of Advanced Scientific Computing Research (ASCR) has the number 1, number 3, and number 17 most powerful computers for open science and an exascale initiative for investigating oxycombustion, carbon sequestration, and other scientific topics. The major components of the Exascale Initiative include platform R&D on power, integration, and risk mitigation; critical technologies; software and environments; co-design and integration with vendors; and platforms that ensure component integration and usefulness.

The Linac Coherent Light Source (LCLS) at the Stanford Linear Accelerator Center (SLAC) produces an amplified X-ray beam and has already produced several science results, such as refracting nanocrystals in water.

The International Thermonuclear Experimental Reactor (ITER) has as its goal a power of 500 MW and a 300- to 500-second burn time. In the past year, it has developed a real schedule, estimated a realistic cost (that translates to a billion dollar annual contribution from the United States), and installed a new director. It has been established as an independent international legal entity with about 400 personnel from all of the member nations. The United States has a 9% share in this enterprise. Roughly 80% of the contribution will be in-kind components manufactured largely by U.S. industry. In addition, the United States will contribute 13% of the cost for operation, deactivation, and decommissioning. The U.S. share of construction was estimated to be $1.45 billion to $2.2 billion. ITER is located in Cadarache, France, and the site has been prepared.

The Inertial Fusion Energy project at the National Ignition Facility recently began full operations and is on track for demonstration ignition.

In particle physics, SC is supporting work at the Large Hadron Collider (LHC) and at the Long-Baseline Neutrino Experiment (LBNE). The LHC is still a factor of 100 off the luminosity at 3.5 GeV; it is to go to 7.0 GeV before shutting down for upgrades. There is a big push to keep the Tevatron running. The LHC is being slow in starting up. It is not clear how the ramp-up of the LHC is going to proceed. The Tevatron was to have been shut down in 2011, but areas of exclusion for the Higgs boson have been constantly expanded, leading to valuable prospective insights from the Tevatron’s continued operation for another 2 or 3 years. But extending the operation of the Tevatron would cost a small fortune. An advisory committee will consider the question of its continued operation later this summer.

In accelerator technology, the questions are:

- Can accelerators be built with about 50 MW of power in the beam?
- Can associated targets be constructed?
- Can accelerators be built to burn the actinides that dominate nuclear-waste-storage issues?

A workshop on this topic was held in July. It is hypothesized that a few such accelerators could destroy all the long-lived actinides produced by nuclear power reactors.

Steps are being taken to strengthen the Small Business Innovative Research (SBIR) program, which amounts to $150 million to $200 million per year in the DOE budget, Those steps include its being moved up to report to the Deputy Director of SC. The Office is being enhanced to deal with SBIR better.

Giles asked if there were ASCR-related activities in the SBIR program. Brinkman replied, yes; the program is distributed across all SC offices.

Hey noted that, in addition to supporting ITER, materials research is needed for that program. Brinkman agreed and pointed out that such research is being conducted in the offices of BES and Fusion Energy Sciences (FES) to build stable, high-temperature materials. Strayer noted that there will be a report on a workshop on that topic later in this meeting.

Berger stated that the Graduate Fellowship Program is great and asked if there were a mechanism to get the graduate students’ advisers more involved in the laboratory experience. Brinkman replied that the only reason that the Department can have a program is to have the students involved in its science. Involving the advisers is an interesting idea. Bell Laboratories always provided a mentor, who played an important role.

Giles noted that the budget picture is shady and asked how a community maintains its enthusiasm in such a situation. Brinkman reassured him that the community will be well-funded and advised that he stay enthused.

The floor was opened for public comment. There was none.

Michael Strayer was asked for an update on the activities of ASCR.

Going forward, computation will play an important role in U.S. science. The ASCR FY11 budget request to Congress was a substantial increase over FY10. The budget has doubled during the past 5 years. In its markups, the Senate cut $109.4 million from the SC request and $8 million from the ASCR request. The House markup cut $221.0 million from the SC request. In the FY10 budget request, increases in Applied Mathematics and Computer Science are maintained; exascale research programs are continued; a focus on commutation of partnership teams for running at scale on multicore computers is continued; and obligations at computing facilities and to partners are fulfilled. For example, the Argonne Leadership Computing Facility (ALCF) upgrade, NERSC-6 operations at the National Energy Research Scientific Computing Center (NERSC), and the deployment a 100-Gbps ESnet are supported.

Exascale-related proposals have been processed in uncertainty quantification, advanced architectures, and software (the X-Stack). Six projects have been funded in advanced architectures at $5 million per year. Eleven projects on the X-Stack were selected at $8.5 million per year; these represent new research directions for the future. Six projects in uncertainty quantification were funded at $3 million per year; these projects will turn modeling and simulation into science and technology at greater fidelity; all of the proposals were formidable.

Meetings are being held with the Department of Defense (DoD) and the Defense Advanced Research Projects Agency (DARPA) to identify critical exascale technology issues.. Ongoing meetings are being held with the National Nuclear Security Administration (NNSA) to create a DOE-wide exascale program. .

In Applied Mathematics, the Hybrid Optimization Parallel Search Package (HOPSPACK) was released as a mature software package for tuning codes with multi-threading parallelism. It is applicable to all known systems. In computer science, SMARTMAP significantly improves communication performance on multicore processors by an order of magnitude. The activities of the Joint Math/Computer Science Institute would be discussed on the following day of the meeting. In Scientific Discovery Through Advanced Computing (SciDAC), the Visualization and Analytics Center for Enabling Technologies (VACET) has developed an automatic feature-detection method for visualizing, recognizing, and tracking topological and graphic features. Also in SciDAC, the Center for Enabling Distributed Petascale Science (CEDPS) has developed and now provides secure, high-bandwidth "fire-and-forget" file-transfer services that have been applied to the STAR [Solenoidal Tracker at RHIC] collaboration at Brookhaven National Laboratory (BNL) and to TeraGrid resources sent to the Open Science Grid (OSG).

In the Facilities Division, NERSC is completing the installation of the Hopper machine, a 1.25-petaflops Cray XE6, which will provide more than 1 billion core-hours to science each year. Resources at NERSC has been upgraded by a factor of 100 during the past 5 years. NERSC is a workhorse of the Department. NERSC has been collaborating with NNSA because Los Alamos National Laboratory (LANL) is standing up a similar machine. The NERSC/Cray is a Center of Excellence, and its programming models ensure effective use of the new 24-core nodes. NERSC has made scientific accomplishments in climate, energy resources, fusion energy, combustion, materials, and nanoscience.

The Argonne Leadership Computing Facility (ALCF) is being upgraded to a 10 petaflop IBM Blue Gene /Q. The new machine will have more, faster CPUs per node, more storage and more memory per CPU making it 20 times more powerful and using only 20% more space than the current IBM Blue Gene/P. Water cooling will enable this technology. This facility has made scientific accomplishments in climate, gas turbines, nuclear energy, nanocatalysts, fusion energy, and materials science.

The Mission Need for Leadership Computing Facilities at Argonne and Oak Ridge was signed in January, 2009. A Lehman review in July 2009 of Oak Ridge Leadership Computing Facility's (OLCF) upgrade plan approved the acquisition strategy for a heterogeneous processor machine. A follow-on application readiness review, held this past month, recommended acceptance of the application-readiness plan. Another Lehman review is scheduled for December 2010. The OLCF has contributed to advances in fusion efficiency for the ITER; simulations of an organic polymer widely used in light-emitting diodes and televisions; understanding the origins of "cosmic lighthouses"; exploring the carbon-water union; and exacting fuel efficiency from smart trucks.

Four of the six Association for Computing Machinery (ACM) 2010 Gordon Bell Prize finalists work in ASCR facilities.

ESnet has garnered a number of awards for excellence. It easily handled an almost 300% increase in network traffic from July 2009 to June 2010. Its 12-month reliability is now 99.985%. and is now in the process of adopting 100-Gbps optical switching. The On Demand Secure Circuits and Reservation System (OSCARS) virtual circuits now deliver cloud service between NERSC and the Joint Genome Institute, a feat that is repeatable for all the laboratories. An interagency proof-of-concept virtual circuit was set up between NASA [National Aeronautics and Space Administration] Ames and the U.S. Geological Survey for the Towards Optimal PDE Simulations (TOPS) project.

Two positions will be posted, one for the Director of the Computational Science and Research Partnerships Division and a computer scientist. Three vacancy announcements just closed, two for a mathematician/physical scientist and one for a computer scientist.

During the presentation, Chalk returned and resumed the duties of DFO.

Graham stated that the Argonne National Laboratory (ANL) system's power consumption should be stated. Stevens answered that Q is about 6 MW for that machine; computing power increases by a factor of 20 as power increases by a factor of 3.

Graham asked about DOE–DARPA collaboration. Hitchcock responded that ASCR has some collaborations, especially with mission agencies, to establish a path forward. It does not talk too much about those collaborations because of mission sensitivity. Strayer added that the Office has ongoing planning activities with DARPA in very collegial relationships governed by a memorandum of understanding (MOU). The collaboration is very vigorous.

Dongarra asked about interactions with the National Science Foundation (NSF). Strayer replied that the NSF cofunds some SciDAC projects. Discussions are being held about how NSF might contribute to the exascale initiative. ASCR reviews their facilities for them.

Chen asked if there could be a connection with NNSA on codesign. Strayer replied that the Office is still trying to see what the exascale will look like and is engaged in plans and discussions with NNSA on the science and codesign. The two organizations meet every other week.

Giles asked if the Lehman reviews were useful. Strayer emphatically replied, yes. They improve project management and allow the Department to build large, complicated facilities on time and on budget. They are a very effective tool.

V. White asked if there were a vision for SciDAC and how it would evolve. Strayer said that all current SciDAC funding will run out in 2012. Science never stands still. The Office will try to meet science's needs and be more focused in its solicitations. There will be separate solicitations in different fields (high-energy physics, biological and environmental research, etc.). They will be accompanied by workshops for the affected communities.

Negele asked what the plan was for adding more codesign centers. Strayer answered that outstanding proposals were received in codesign, and the Office is still deciding what to do with them. There are a lot of options. One possibility would be to issue additional solicitations in FY11 or FY12.

The floor was opened to public comment. There was none. A break was declared at 10:33 a.m.

The meeting was called back into session at 10:55 AM to hear **Linda Petzold** report on the Committee of Visitors (COV) to the Division of Applied Mathematics.

The charge to the COV was to assess the operations of the Applied Mathematics program during the fiscal years 2007, 2008 and 2009, evaluating the processes used to solicit, review, recommend, and document proposal actions along with those used to monitor active projects and programs. It was also to assess how the award process has affected the breadth and depth of portfolio elements, and the national and international standing of the program.

The COV met on May 12-13, 2010. A COV web page was available a month in advance of the visit, presentations were made during the visit, files were well-organized and available, and information was made available upon request.

The solicitation and review processes appear to be effective and fairly administered. The program is to be commended for their work in streamlining the proposal review process. The documentation seems to be done very well, but summary statistics were not provided. Delays in processing approved grants, which are outside of the control of the program, affect the principal investigators' (PIs) ability to recruit students and postdocs and also affect tenure decisions for junior faculty.

The COV recommends that further consideration be given to improving the level of outreach on new funding opportunities and to exploring a more flexible approach with a broader proposal-acceptance window. Proposal project descriptions should be limited to 15 pages. The merit review criteria for large multi-investigator proposals should ensure that the elements of the proposed research are appropriately integrated, coordinated and synergistic. Actions should be taken to accelerate the processing of approved grants.

The Applied Mathematics research program managers use generally effective mechanisms, including site visits, PI meetings, and progress reports, to monitor ongoing projects and collect information about major awards and accomplishments and to maintain the high quality of the research.

Explicit guidelines should be instituted for progress reports, including a prescribed length and a clear description of the information that should be in the report, such as a list of publications, presentations, awards, and patents attributable to the project. The metrics for impact should also be clearly stated and explained.

The COV found the portfolio to be exceptionally strong in both depth and breadth. The balance of awards with respect to innovation, risk, and interdisciplinary research appears to be appropriate. The COV was very impressed with the long-term perspective of the DOE Applied Mathematics program and its simultaneous agility at funding new program areas.

The DOE Applied Mathematics program has been, and continues to be, of extremely high quality and standing, both nationally and internationally. A great strength of the program is the willingness it has demonstrated to invest in projects with a longer-term perspective than is possible at most U.S. agencies, enabling the support of breakthrough research and ensuring its success and eventual adoption.

Hey noted that, in regard to the national and international standing of the program, the DOE Applied Mathematics Program has a longer-term perspective than does any NSF program. Petzold replied that DOE supports researchers rather than short-term projects, allowing research, algorithm development, and code writing. Berger asked if that carried over to the national-laboratory environment. Can they do the long-term code development, or are they suffering from the three-year limits? Petzold said that this issue did not come up in the review. Landsberg added that it goes back to renewals. National laboratories can do decades of research. They are successful in developing mature software. National laboratories and universities are evaluated equally.

Voigt asked if the decision to award were passed on to the recipient before the announcement was made. Landsberg replied that the Office can let the awardees know so that they can work “at risk.” Voigt stated that the 10-year process should not be affected, then. Petzold said that there is anecdotal evidence about hindrance of work by universities that had not yet gotten their funds. Voigt suggested that the balance issues should be mentioned in the report. Giles offered that the Subcommittee could amend the report to include that comment.

Chen, noting the current balance between funding for national laboratories and universities, asked if the balance should shift as the big stack is developed for the exascale. Landsberg pointed out that every exascale-related award has been to a joint national-laboratory–university proposal. How the dollars are split between the two is determined within the collaboration. Giles asked if the funding level was as severe as the award level indicates. Landsberg replied, yes; $3 million is insufficient to fund uncertainty quantification.

Negele noted that all researchers spend time writing proposals and reading them. Large proposals with many PIs need more space than more modest proposals do. Petzold said that the COV discussed this and believe that the PIs should think out the important aspects of the proposed research. The Office might want to extend the limit a little bit for really big collaborations. Giles asked if the COV had a strong feeling for the 15-page limit. Petzold replied, yes; that is not as drastic as the recent National Institutes of Health (NIH) recommendation. Giles suggested making the recommendation and seeing what the ASCR response was. Chalk noted that the Office is required to respond to the Committee’s recommendations within 30 days.

Hey suggested that the DOE’s strengths in its national laboratories could be mentioned.

Giles called for a motion to accept the report as amended. Voigt made such a motion; Dongarra seconded it; the motion was unanimously adopted.

Alvin Trivelpiece was asked to present the results of the Exascale Blue Ribbon Committee’s workshop. He asked the Committee how many were at the Chattanooga SciDAC meeting; a few had attended. He had been asked by Michael Strayer to chair the Exascale Blue Ribbon Committee. James Decker and Trivelpiece wrote the final report. The results were presented at the SciDAC meeting.

The Superconducting Super Collider (SSC) started out as a simple idea coming out of a workshop. The discussion of it was very closed. The design was done by Maury Tigner, and it was estimated to cost $4 billion. With contingencies, the cost was set at $4.4 billion. At a White House meeting, President Reagan used a football analogy to justify “throwing deep,” and the SSC was launched. It was not built for

$4.4 billion because of politics that called for international participation. The Italians had pledged $1 billion, and the Japanese had pledged $400 million. The price tag for the Italian part was that the international director was to be an Italian for one year; however his political foes scuttled the Italian cooperation. A similar process played out in Japan. There were not a lot of people involved in the decision making. The exascale initiative has a lot more people and institutions involved; therefore, it has a much higher risk of failure.

One needs to put together a lot of things to accomplish the exascale. In the 1980s, the fusion energy research effort was centered at the Mirror Fusion Test Facility-B (MFTF-B) at Lawrence Livermore National Laboratory (LLNL) and was connected by a 50-kilobit data-communications link to sites around the country. A single transport code on one computer was to be accessed and used remotely. It did not work well. The Controlled Thermonuclear Research Computer Center (CTRCC) that was set up at LLNL eventually turned into NERSC.

The key finding of the Exascale Blue Ribbon Committee Workshop was that

> " there are compelling needs for exascale computing capability to support the DOE's missions in energy, national security, fundamental sciences, and the environment. The DOE has the necessary assets to initiate a program that would accelerate the development of such capability to meet its own needs and by so doing benefit other national interests. Failure to initiate an exascale program could lead to a loss of U.S. competitiveness in several critical technologies."

Moving forward is a good idea. The workshop concluded that the exascale initiative is a good way to move forward.

Graham noted that there are scientific challenges in deploying the exascale and asked how risk should be managed. Trivelpiece replied that there is the science that can be done with computers and there is science that has to be done to make the machine work. The two are related, and both need to go forward. The United States occasionally goes off and does something and then loses it, usually to overseas institutions. The country needs to get the science of computers right. Graham said that the Committee understands these issues but needs to explain the exascale to the broader community, stating the challenges without making the effort seem futile. Trivelpiece stated that the scientists made the case in a compelling way. This is not the SSC. There are a lot of stakeholders and a lot of pitfalls, but the exascale is worth examining to see how it would succeed before the Office of Management and Budget, the White House, and Congress.

Bailey stated that programs are successful if they meet a focused national need. However, this report appeals to a broad community. The high-energy-physics communications program succeeded with a similar constituency. He asked if there were a compelling reason for the United States to do this. Trivelpiece said that some programs were pushed forward under one person's authority. Sometimes dumb luck favors a project's going forward (e.g., Rosen's support in a National Academy of Sciences report for light sources). It is not easy to determine a path forward. Sometimes the support produces a backlash that kills the program (and vice versa). Strayer and the DOE need ammunition to push this cause in a way that does not produce internecine warfare. There are a great number of variables. Bailey said that everyone seems to agree that there is an energy problem. That might be a hook to hang this program on. It seems that this program has too many such hooks. Trivelpiece said that the computing program probably owes its existence to the 1973 oil crisis. There is a compelling case to proceed on the exascale. However, 20% of the federal budget is discretionary, and only a small part of that is devoted to science, and its priority is not great. The workshop said there was a compelling case.

Giles said that Committee has to think how this report will be incorporated into the Committee's report to DOE. A break for lunch was declared at 12:09 PM.

Tuesday, August 24, 2010
Afternoon Session

The meeting was called back into session at 1:38 p.m., and **Andrew White** and **Dona Crawford** were asked to report on the status of the ASCAC Subcommittee on Exascale Modeling and Simulation. Crawford introduced other members of the Subcommittee who were in attendance. A. White pointed out that an oral report was transmitted to SC on March 31, 2010. The draft written report is complete. Structuring the report is problematical. The initiative is made up of three parts: a shared piece, a community piece, and a proprietary piece. Codesign, an iterative process of simultaneous code and hardware development, is what binds all the pieces together. Uncertainty quantification is applicable across all activities in each of the three pieces.

The strategy for the report is: (1) do not go discipline by discipline, but use the workshop reports as input; (2) focus on formational calculations; (3) identify the barriers to all applications; and (4) think broadly, including NNSA, engineering, and industry.

The key observations from the March 31 transmittal letter are:

1. Exascale computing offers transformational opportunities for key science and engineering areas.
2. The application workshops identify some of these opportunities and their impact.
3. The cross-cutting workshops identify some of the challenges with a goal of identifying an R&D agenda.
4. There is a fundamental transformation of computing taking place with new families of processor chips.
5. The time to start is now.
6. The magnitude of the changes produces large challenges, including many in education and training.

What was expected was for the report to include the remaining areas; to elaborate on case studies of successes in moving from a megascale to the gigascale to the petascale; and to identify ASCR strategies, such as using the lessons learned from the NNSA ASC, employing codesign, and engaging industry. However, the report does not talk about the role or place of experiments, the benefits of a more holistic approach, or technology transfer strategies for interacting with the private sector.

Crawford said that an outline of the report was put together during Subcommittee telephone calls. After an Introduction, the outline covers (1) Why the Exascale? (including this being a transformational, not an evolutionary, enterprise; the complexity challenge; and the need for connecting to industry); (2) What History Can Teach Us; (3) What Applications May Be Transformed by Going to the Exascale (where the community is making progress and what progress needs to be made); and Challenges in Going to the Exascale. The outline also calls for a Summary.

Giles noted that pieces of this report are more complete than others are and pointed out that the Committee does not have a copy to discuss.

Chen commented that the outline looks reasonable but could be more inclusive of transformational applications. A. White pointed out that the applications are astrophysics, climate, combustion, biology and medicine, airframes, jet turbines, and materials. Chen asked if fusion should be in there. Bailey stated that he would be skeptical of a DOE report on airframes. A. White stated that that was an example of successes in the past, backed up by reports. Crawford added that each application could be talked about in terms of its past history, its readiness for transformational change, and how to go forward.

Bailey suggested identifying a strategy that ASCR can use; he did not see any strategies in the outline. A. White said that they are embedded in the technical discussions. The Subcommittee has emphasized applications in challenges and should balance them with strategy.

Tang suggested pointing to events in the past that have been transformational changes and using them as a point of departure for discussing the basis for the exascale. Fusion energy is a good application domain that should be included in the report.

Voigt asked why the exascale should be pursued *now* on such a short timeframe. He questioned whether one gains credibility from these multiple applications rather than stressing national security and independence. A. White said that the Subcommittee focused on the DOE mission in science. The information technology industry’s competitiveness is mentioned, but not as strongly.

Graham commented that the conversation seemed to vacillate between the transformational and the extant suite of applications. A. White said that the Subcommittee picked the applications by their impact on DOE's mission. Crawford added that the report's length had never been set and the level of summarization should be thought about. A. White offered that the report should stand alone. One has to be convinced to do the exascale from this report.

Graham asked who the audience was. A. White responded, Koonin and Brinkman. Graham noted that it will be a public document and suggested considering who else will read it.

Bailey said that something should be said about ASCAC's opinion of the validity of the reports of the workshops. A. White emphasized that the Subcommittee *is* basing a lot of this on the workshops. Bailey asked if there were a strategy for success and suggested involving the computer industry to support this initiative. He asked if the challenges were in design or process, and how one can overcome these challenges. A. White responded that industry is involved in the path forward, in national laboratory–industry partnerships, and in codesign. Bailey suggested going into the details of the path, the partnerships, and codesign to be convincing. Giles added that the report has to capture the depth and meaning of the workshops and not underestimate the amount of work that has been done. He suggested bringing across the work that has been done, the fact that the time is now, and that there is a window for assuming leadership. He suggested a findings and recommendations section be inserted at the very beginning. That section should be followed by the stories behind the transformation and where there may be pitfalls and challenges for the exascale to come into its own. These reports should be readable to everyone.

Hey said that the Committee is making the case for the exascale as a national priority. Many of the applications cited could be related to energy independence. A. White said that the Subcommittee was trying to strike a balance between science and energy independence. Chen suggested grouping the sciences together.

A. White admitted that he was concerned about the process of how the transformational applications were selected. To leave out the work that was done in the workshops is a mistake. All the people in the community should be involved, even though that will produce friction. The effort should be taken to a level higher.

Hack stated that a compelling story has been told for the need of this enhancement. Everyone will benefit from this type of investment. A hands-off evolutionary process will not make this go forward. It should not be a hard case to make. The workshops should be rolled up; they point out the effects on national security, energy independence, U.S. competitiveness, etc.

Petzold summed up by saying that the Committee seems to be saying that the report will be more effective if it is crisp and concise and makes its case by telling the effects of the exascale on the economy, military, etc. The scope should be inclusive.

Negele added that one should not be shy. He urged making a strong case. There is an opportunity in the Executive Summary to make a sharply focused statement. He liked crisp examples. There was some good grist in the workshops. A. White pointed out that the report says that computer science is needed for the exascale and that more workers need to be drawn into this field.

V. White asked if the report addresses the balance between the intellectual effort needed on hardware and software. A. White answered that he believed that the report was balanced and that the balance was stated implicitly. There is a perception that the exascale is about hardware; perhaps that should be corrected.

Berger stated that no one would argue about the exascale's being beneficial to a broad base of people. The question is, why now? Giles said that it is essential to move forward in DOE's missions. Crawford added that there are other questions, such as does the United States want to be subservient to someone else's programming models.

Voigt stated that everyone agrees that many disciplines will benefit from the exascale; the temporal urgency is the question to be addressed.

Tang said the United States is facilities poor in this area. This initiative could get the United States ahead of the curve. All the application areas should be subjected to the litmus tests of urgency, competitiveness, etc. Some time could be spent to bring more people to the table and make the report more compelling.

The floor was opened to public comment. Stevens said that this is a case where not taking action has a huge downside. The importance of DOE's making progress is an argument for pushing this forward in tough times. There are some application chapters missing. Everyone wants their own thing in the exascale. One could imagine many more disciplines that would benefit. One has to constantly ask, is this enough?

Trivelpiece pointed out that the report of the Exascale Blue Ribbon Committee had a biography for Lee Hood but there was nothing in the report from him. He gave a talk about how a drop of blood would, in the future, be used to diagnose an individual's health. It does not fall in DOE's bucket, but it was a very compelling argument for the exascale.

Hack noted that Jack Worlton, in writing about patterns of technological change in high-performance computers, had said that the transition to a new technology is a discontinuity and that to manage a discontinuity improperly is a disaster. Hack noted that the exascale is a discontinuity in both hardware and software. It must be managed carefully.

Giles said that ASCR needs to move forward on the exascale initiative, and the report will reflect that. The Committee should have a teleconference to consider and vote on this report in the next 6 to 8 weeks. The Chairman's letter to SC will reflect the sentiments of the Committee.

A break was declared at 2:52 p.m. The meeting was called back into session at 3:24 p.m. **James Roberto** was asked to report on the ASCR–BES Workshop on Computational Materials Science and Chemistry for Innovation.

Advances in materials and chemistry have shaped history. However, energy technologies are limited by the availability of advanced materials and chemical processes. The solar cell, electric cars, and rechargeable batteries have not realized their potentials because of a lack of materials; and fossil-fueled plants operate at two-thirds of their optimum efficiency because of a lack of materials to deal with pollutants. Materials that operate at extreme temperatures and environments; materials and chemical processes that efficiently separate greenhouse gases from effluent streams; high-strength, lightweight materials for transportation; and new catalysts for efficient chemical processes are needed. Materials science and chemistry also underpin industrial competitiveness; the company or nation with the best environment for discovering and deploying new materials and chemical processes will be more competitive. Transformative advances in materials and chemistry *will* be achieved; the question is how quickly and by whom.

Achieving performance gains requires exploiting many degrees of freedom in composition and structure. For example, the parameter space for advanced steels has increased 1-million-fold compared to early steels. New superconductors and high-field magnets are much more complex than their predecessors. This means that intuitive, trial-and-error discovery is impractical for future discoveries because there is not the time or resource base to explore all the options experimentally. The discovery process must be transformed, and the United States is in an excellent position to do that. It has developed and deployed the world's most powerful collection of tools for the synthesis, processing, characterization, and simulation and modeling of materials and chemical systems at the nanoscale. For the first time in history, these materials can be synthesized, characterized, and modeled at the length scale where this behavior is controlled. The leap of nanoscience holds great competitive advantages.

During the past decade, computing power has increased by a factor of 1000, and the United States leads the world in computing power. At the same time, software advances have added another factor of 1000 for many applications, producing a 1-million-fold increase in effective capability. This has profound implications for the pace of discovery and the creation of new technologies, and it impacts the entire innovation cycle. In airfoil design, engine development, tire design, and aluminum casting design, drastic decreases in testing and development time have been achieved by simulation-based engineering and science.

A Workshop on Computational Science and Chemistry for innovation was held in Bethesda, Maryland, July 26-28, 2010. Its premises were that (1) advances in computing and computational science offer the potential for predictive capability in many areas of science and engineering and (2) experimentally validated simulations will accelerate discovery and innovation, providing a competitive advantage for both science and technology.

160 experts from 69 organizations were assembled to build on the existing basic research needs, grand scientific challenges, exascale, and FTAC [Fast Track Action Committee on Computational Modeling and Simulation] reports and to assess the potential of experimentally validated simulations to accelerate discovery and innovation. There were plenary sessions on basic energy sciences, computational sciences, industrial context, computational design of materials, and accelerating the innovation cycle. Breakout sessions were formed around seven themes: materials for extreme conditions, chemical reactions; thin films, services, and interfaces; self-assembly and soft matter; strongly correlated electron systems; electron dynamics, excited states, and light-harvesting materials and processes; and separations and fluid processes. Each breakout session explored the respective state-of-the-art, the most significant research opportunities and barriers, the acceleration of discovery and innovation by computational models and simulations; and the computational and experimental challenges that must be overcome to enable this acceleration.

This is an opportune time to pursue materials and chemistry by design because the experimental and computational facilities are in place; new materials and chemistry are outpacing the development of many new technologies; computational capabilities are outpacing materials research; predictive design is key to accelerating discovery and innovation; and advances are urgently needed for economic competitiveness and for developing energy technologies. The scale and quality of U.S. scientific infrastructure currently convey a significant competitive advantage.

A study at the Massachusetts Institute of Technology in 1995 showed that it took an average of two decades to go from discovery to commercialization. One example of a possible accelerant of the design of new materials is the Materials Genome Project that is computing the materials properties for all inorganic compounds, which is leading to the discovery of new structures and chemical classes and the synthesis and characterization of the most promising materials. Already, 200 new ternary compounds that are candidates for battery materials had been identified. This process needs to be integrated and made robust. It would allow sifting through all the possible material structures much faster than can be done now. In such a process, software equals infrastructure

The workshop came up with seven foundational challenges:

1. Predicting and optimizing structure, which are essential in materials for extreme conditions
2. Understanding and controlling self-assembly, which are essential in designing and engineering materials at the nanoscale
3. Light harvesting, which is essential in photons to energy
4. Controlling chemical reactions, which is essential in combustion and designer catalysts
5. Separations and carbon capture, which are essential in chemical engineering by design
6. Designer thin films and interfaces, which are essential in interfacial materials and advanced batteries
7. Predicting and controlling electronic structure, which are essential in understanding spins and superconductivity for high-performance magnets

The outcome will be the creation of an innovation ecosystem that integrates synthesis, processing, characterization, and simulation and modeling. This ecosystem requires that a predictive capability be achieved and strengthened in foundational challenge areas; that approaches be developed that span vast differences in time and length scales; that validation of models and quantification of uncertainty be developed in simulation and modeling; that robust and sustainable computational infrastructure be created; and that simulation-based engineering and science be efficiently transferred to industry.

A first draft of the workshop report was expected in a week or two after this meeting.

Voigt asked what the states of the prediction and validation capabilities were. Roberto said that that is a work in progress. To a degree, it could be built into the machines and algorithms.

Tang said that it is clear that there will be an immediate social impact. Validation is important to convince the public that progress is being made. He asked if there were an industry developing to do this.

Roberto responded that the line of sight from fundamental discussions to the ultimate application is very clear. Industry will be able to justify taking scientific advantage. The benefits are enormous. Industry is way ahead of materials development; these advances will be highly leveraged; these are multi-year efforts; they need time, resources, and the right people.

Strayer said that the benchmark of uncertainty quantification is couple-cluster theory with errors that are smaller than what can be measured.

The floor was opened to public comment. There was none.

Giles asked when the final report would be ready. Roberto said that they expected a draft in September with the final version ready soon after that.

Douglas Kothe was asked to speak about the Consortium for Advanced Simulation of Light Water Reactors (CASL), a new DOE simulation hub that brings together national laboratories, universities, nuclear-power trade groups, owner-operators of nuclear plants, fuel vendors, and many individual contributors.

The United States is a leader in nuclear power; of the 24 combined operating licenses (COLs) for planned sites, 14 are in the South. The U.S. nuclear-energy industry has gone from poor capacity factors to capacity factors of better than 90%, a model for industrial production. Since 1977, the current fleet has extracted 5.7 GW of additional power without building a new plant; and with the current modeling and simulation software, an additional 3.4 GW will be achieved from the current fleet by extracting more heat from the fuel.

There are numerous safety, operating, and design issues to consider for nuclear reactors; CASL selected 10 operational criteria to model and deal with. All products from CASL should hit the manufacturing floor and garner acceptance by the user community, acceptance by the regulatory authority, and acceptance of outcomes by the public. This is a tough community to introduce new ideas into.

There is a sense of urgency about life extension that is driven by the economic decision on whether or not to continue to operate a plant. The oldest plant is 41 years old, and it is licensed to a life of 60 years. Can it go to 80 years? A tool, such as an advanced "virtual reactor," that could predict the answer to that question would drive decisions. The high-level goals are (1) to reduce capital and operating costs per unit of energy by power upgrades and lifetime extensions, (2) to reduce nuclear waste by enabling higher fuel burnups, and (3) to enhance nuclear safety.

CASL is focusing on pressurized water reactor (PWR) fuel failures, which are largely caused by grid-to-rod fretting. CRUD [corrosion residual unidentified deposits] and corrosion preclude a lot of upgrades. CRUD will build up in the upper reaches of the core and cause corrosion and an axial offset on the fuel tubes. As a result, power in those areas goes down to unacceptable levels. The industry needs to know how CRUD operates and how to stop it. To do that, a multi-physics, chemistry, flow, and neutronics model is needed to predict CRUD growth. Interestingly, when one thermally cycles the reactor, CRUD goes away. When CRUD remains, it produces hot spots on the fuel, leading to localized boiling, CRUD dryout, accelerated corrosion, and a leak in the fuel rod. A high-fidelity, high-resolution capability to predict hot spots, localized crud thickness, and corrosion is needed.

With thermal cycling, the fuel and the springs holding the fuel rods in place shrink and expand. Rattling of the fuel rod produces fatigue wear and failure (fretting). An interactive tool is needed to predict turbulent-flow excitation, rod vibration, and wear.

There are other processes that require modeling and simulation of the reactor:

- As the fuel is a radiated, fuel rods expand and take on an S-shaped bowing.
- The maximum heat transfer rate is accomplished just before boiling nucleation. Power uprates require improved quantification of the margins for the departure from nucleate boiling and for dryout limits.
- The structural integrity of a reactor vessel and its internals depends on its temperatures, stresses, and material performance (fatigue and cracking) over long-term operation.

New materials and fuel concepts are needed for any transformational performance improvement. Silicon carbide cladding is of interest because it could produce enrichment savings (because of its low cross-section to neutrons), it could lead to an uprate capability, it is insensitive to dryout, and it is immune from fretting failure. Uranium nitride fuel is also of interest because of its much higher thermal conductivity and increased thermal output capability, leading to a cooler fuel and lower fission-gas release.

The development of a virtual reactor and its successful application to these problems would benefit the nuclear industry, so CASL has selected these key phenomena as "challenge problems." Such a virtual reactor would leverage current state-of-the-art neutronics, thermal-fluid, structural, and fuel-performance applications; leverage existing systems and safety analysis simulation tools; develop new requirements-driven physical models; develop an uncertainty-quantification framework; deliver a validation basis against 60% of the existing fleet by using data from Tennessee Valley Authority (TVA) reactors; and deliver a base modeling and simulation capability for light water reactors.

During the next 5 years, CASL's priorities are to deliver an improved predictive simulation of the PWR core, internals, and vessel and to couple the virtual reactor to an evolving out-of-vessel simulation capability; to work in five technical focus areas; and to establish a focused effort on boiling water reactors and small modular reactors. The five technical focus areas are

- materials performance and optimization,
- models and numerical methods,
- virtual reactor integration,
- validation and uncertainty quantification, and
- advanced modeling applications.

It will use the LIME multi-physics integrator. Level-1 milestones and capabilities to address the six challenge problems mentioned above (CRUD, grid-to-rod fretting, etc.) have been developed. The virtual reactor development cycle is planned around major releases for Level-1 simulations by the public, and there is a plan for developing and releasing these products. The CASL virtual reactor has a mature starting point, building on existing capabilities to deliver versatile tools. One of these existing capabilities is the Denovo high-performance-computing transport code. The code is performing very well in parallel and has been scaled up. The neutron-energy spectrum in nuclear reactors needs to be analyzed; this code allows a 44-energy-group resolution of that spectrum, a capability that will allow better prediction of phenomena.

CASL will try to identify and resolve uncertainties in the various challenge problems. The Predictive-Capability Maturity Model (PCMM) will be used to measure the progress of the virtual-reactor development. PCMM measures the maturity of the process by objectively assessing seven technical elements. The CASL virtual reactor will be assessed annually against the challenge problems.

In terms of computational requirements, one neutron state of a core has about 10^{23} uncertainty states because of the number of degrees of freedom under steady-state and transient conditions. With a system like Jaguar and better algorithms, models, and software, one can make a good impact on the technology. CASL will make sure that its capabilities in technology and computer science are world-leading. The CASL partnership also possesses the key elements required for success (e.g., physical reactors; Nuclear Regulatory Commission engagement; an education, training, and outreach program; validation; and virtual office, community, and computing).

Berger asked if they were using a lot of meshes that do not scale well. Kothe replied, yes. The core is modeled on a structured mesh, and the pins are homogenized. Another mesh is used for the flow. One need is to be able to change geometry decisions quickly. That is a large driver for fidelity.

Tang asked how ready LIME was and whether new versions would be released with new models. Kothe replied that they are going to try to do that on a 6-month cycle. It will be a challenge; the difficulty of doing this should be known in the next 3 months. Westinghouse will help at the beginning.

Bailey asked whether some codes will be parameterized for design when the effort is all done. Kothe replied, yes; it is expected that the code will run on the desktop as well as push the exascale. What is needed to be known is whether or not one needs 10,000 energy groups to follow a neutron from birth to death. The big systems will be needed to guide such decisions.

Hey asked whether Westinghouse used parallelized HPC codes in their existing simulations. Kothe replied that they use some moderate-scale machines with a few hundred processers (cores) but not petaflop or teraflop machines. Commercial codes will be brought in to run on the virtual reactor. Hey asked whether the advanced reactor designs will be looked at. Kothe replied, yes, including pebble bed, molten salt, etc., as time allows.

The floor was opened to public comment. There being none, the meeting was adjourned for the day at 4:47 p.m.

Wednesday, August 25, 2010

The meeting was called to order at 8:34 a.m. **Jack Dongarra** was asked to provide an update on the Extreme-Scale Algorithms and Software Institute (EASI), one of four math/computer science institutes, and on the International Exascale Software Project (IESP).

EASI's goal is to produce architecture-aware algorithms for scalable performance and resilience on hybrid architectures by studying and characterizing application-architecture performance gaps; identifying features that future systems may need; developing multi-precision and architecture-aware implementations of software for heterogeneous multi-core architectures; exploring new methods of algorithm resilience; looking at run-time support for adaptable algorithms for resilience, scalability, and performance; demonstrating architecture-aware algorithms in full applications; distributing the new algorithms and run-time support; and establishing a strong outreach program to disseminate results, interact with colleagues, and train students. It is an integration of mathematics, computer science, and application experts working together to create new architecture-aware algorithms and associated run-time support, applications, and numerical libraries. It is studying a number of applications (e.g.,HOMME, MADNESS, and Charon).

EASI started in the fall of 2009 for national laboratories and in the spring or summer 2010 for universities. The locally self-consistent multiple-scattering (LSMS) code is a first-principles computer model that simulates the interactions between electrons and atoms in magnetic materials. It was the first application to hit a teraflop and a petaflop. A parallel implementation and scaling of the LSMS method is perfectly scalable at high performance and has a sustained performance similar to that of Linpack.

Complex multiplication usually involves four multiplications and two additions. One can do three multiplications and five additions, simplifying matrix computation. But there is no free lunch. One needs extra storage, and the imaginary part may be contaminated by relative errors much larger than those of conventional multiplication. However, with certain ways of measuring errors, they are just as small as for conventional multiplication. It was used on an Nvidia board, and performance increased about 33%.

EASI also wants to develop robust multi-precision algorithms, exploiting a floating-point architecture as much as possible and then enhancing the precision to get back the 64-bit precision. With mixed-precision iterative refinement, the accuracy attained is the same as that of the 64-bit arithmetic solution. However, one needs additional memory, and, if the matrix is too ill-conditioned, the matrix will converge. But 32-bit arithmetic is twice as fast as 64-bit arithmetic. Single precision is faster than double precision because there is higher parallelism within floating-point units, there is reduced data motion, and there is higher locality in cache. Running an example on a graphics processing unit (GPU) board showed a factor of 2 improvement for single precision over double precision with mixed precision performing at a slightly lower level than single precision. The same strategy can be used in iterative methods, with outer iterations using 64-bit floating-point and inner iterations using 32-bit floating-point. Generally speaking, using mixed precision gets applications to run faster (often by a factor of 2), although it does not work for everything. A possible technique would be to automatically switch between single precision and double precision to match the desired accuracy. This would apply to sparse direct and iterative linear systems and eigenvalue, optimization problems, where Newton's method is used.

EASI is also trying to devise communication-avoiding algorithms and has written such algorithms for QR factorization (orthogonal matrix triangularization). That is not a new algorithm, but this is a new way of expressing it. Benefit is obtained on a shared-memory machine with a significant bump-up in speed resulting from communication avoidance. In experiments on different matrices, communication avoidance sped up the run time per kernel on an eight-core Clovertown.

EASI research is developing the supporting architecture-aware run-time support required by a variety of algorithms. It is also extending Message Passing Interface (MPI) (which is the only "defined" communicator but which assumes a flat architecture) for hierarchical architectures. EASI has extended the MPI interface and run-time support

to enable existing MPI algorithms to discover and take advantage of the hardware hierarchy and multi-core shared memory. EASI is looking at what processes are running on a node, processor, or cache. These features are going into Open MPI and other widely used software packages and are being adopted into the standards.

EASI startup at the University of Illinois was delayed by delays in funding. Graduate students have now joined the project.

The report of the International Exascale Software Project (IESP) has been produced. It contains a roadmap for software development for exascale computing.

A plot of the total performance of all U.S. top-500 computers showed an increase from about 50 Tflop/sec to about 20,000 Tflop/sec between 2000 and 2010. A similar plot for Europe is remarkably parallel but uniformly slightly lower in total performance. Japan kept pace with Europe during the first half of the decade but fell back a bit in the latter half. China has made remarkable progress; in 2000 it had no computers in the top 500, yet it surpassed Japan last year and will surpass Europe in November.

In looking at the exaflop machine projected for 2018, one has to be terrified or elated by the projected system memory (32 to 64 PB), node performance (2 to 4 TB/sec), total concurrency (10 to 100 billion), and mean time between interrupts (1 day). The attributes of this projected machine that necessitate redesign are the extreme parallelism and hybrid design, the tightening of the memory/bandwidth bottleneck, and the necessary fault tolerance. In addition, the software infrastructure does not exist today. Hardware has changed dramatically while software ecosystem has remained stagnant. Previous approaches have not looked at the codesign of multiple levels in the system software stack. There is no global evaluation of key missing components. New hardware trends that cannot be handled by the existing software stack need to be exploited. Emerging software technologies exist, but have not been fully integrated with system software. The community codes are not prepared for the coming sea change in architectures.

The goal of the IESP is to improve the world's simulation and modeling capability by improving the coordination and development of the high-performance-computing software environment. It requires an international community effort because of the scale of investment, the need for international input on requirements, the international scope of software development, the lack of global evaluation, and the existence of hardware features that are coordinated with software development.

The effort started before 2008; a series of meetings in 2008 and 2009 produced the IESP roadmap in 2010. That roadmap is a planning instrument to enable the international high-performance-computing community to improve, coordinate, and leverage their collective investments and development efforts. After the needs are determined, organizational structures will be constructed. The roadmap covers operating systems, programming models, applications, and cross-cutting dimensions (i.e., resilience, power management, performance optimization, and programmability). The programming models include the development of the numerical libraries that will be needed to consider technology drivers, alternative R&D strategies, a research agenda, and cross-cutting considerations.

This roadmap will need to be refined; a prioritized list of software components for exascale computing will need to be set; the software and algorithm needs of applications will need to be assessed; a structure for the IESP will need to be developed; the participation of funding agencies will need to be coordinated; how laboratories, universities, and vendors can work together will need to be explored; and a codesign work plan will need to be created. Execution of this project will involve a strategic plan, a realistic timeline, community development techniques and risk plans, and intellectual property agreements.

The Europeans are partners with the United States but have their own organization [European Exascale Software Institute (EESI)]. They have 18 months of funding from the European Union. They have set up eight groups to establish their strategies. The G8 has a call out for an Interdisciplinary Program on Application Software Towards Exascale Computing for Global Scale Issues. It commits €10 million over 3 years. It has garnered 78 proposals and expects to select 8 to 10 projects. It has a computing, software, and simulation component with €2.5 million that will fund two or three projects.

Bailey asked about a roadmap for platforms to demonstrate this software. Dongarra replied that the program was working with the vendors. Funding agencies are also involved. Applications people are also guiding the program.

Bailey asked when an integrated milestone chart (a unified roadmap) would be available. Dongarra replied that one gets some roadmaps from vendors, but not a crisp view. Bailey noted that it is difficult to do codesign if only part of the community can see the target. Strayer said that ASCR talks with vendors about R&D for building a roadmap, but there is no funding or program yet. All there is are projections. There will be an integrated roadmap after 2012. This is just the planning stage.

Berger pointed out that the multi-precision algebra for solving partial differential equations on finer grids builds up a lot of errors. Dongarra said that it works for these problems but not for all applications, yet. It is a research problem.

Barbara Helland was asked to speak on the ASCR Computer Science Graduate Research Fellowship Program.

DEIXIS is a magazine that tells what the fellows are doing. The fellowship program started in 1992 to fill a growing need for computational scientists. Sixteen classes have completed the program, and four are in progress. The program requires fellows to have broad training through an approved program of study and to participate in a 12-week practicum at a national laboratory. There were 530 applications in 2010.

The steering and review committees have an uncanny ability to predict the future application areas.

There are 69 fellows as of September 1, 2010. The program is open to U.S. residents or permanent residents in the first or second year of graduate school and to exceptional senior undergraduates. It provides 4 years of support at $36,000 per year plus a $5000 academic allowance plus tuition and fees. The obligations of the fellows are to maintain full-time enrollment in a program of study that will provide a solid background in science/engineering, computer science, *and* applied mathematics. This program must be completed by the third year of the fellowship. The practicum provides the fellows with experience outside the main thesis path and makes them better aware of the areas that define computer science. The practicum must be completed by the second year of the fellowship and carries an additional financial allowance. There is also an annual fellowship meeting.

Of the former fellows, 26% are now at DOE facilities, 30% are in industry, 33% are in academia, 3% are in graduate school, and 8% are in other positions (e.g., at other government agencies). Reviews of the program cite it as an “incredibly effective program.”

Graham was struck by the low success rate of applicants (5%) and asked if anyone studies what happens to the unsuccessful candidates. Helland replied, no.

Petzold asked if the staff had thought to leverage this program to require the university programs and departments to be part of a computer science program network. Helland replied, no, but the advisors had been invited to come to the annual meeting.

Hack asked what the relationship of this program was to the SC Graduate Fellowship Program. Helland answered that there had been no discussion. This program is funded by ASCR. Assurance has been received that funding will be continued. The SC program is funded out of WDTS.

Hoang noting that this program is also supported by NNSA said that other NNSA activities will be modeled after it.

A break was declared at 9:57 a.m. The meeting was called back into session at 10:09 a.m., and **Walter Polansky** was asked to provide an update on the ARRA investments in ASCR.

The purpose of the ARRA is to create new jobs and save existing ones, spur economic activity and invest in long-term growth, and foster unprecedented levels of accountability and transparency in government spending. ASCR had to develop project operating plans, identifying public benefit and impact, cost and schedule, milestones, performance measures, and management. All funds had to be obligated by September 30, 2010, and spent by September 30, 2015.

The $154.9 million of funds went to

- the advanced networking initiative ($66.8 million);
- leadership computing facility upgrades ($19.9 million), which are now complete;
- advanced computer architectures ($5.2 million), which will be completed in the next month;
- Magellan (cloud computing) research ($33 million); and
- SciDAC ($30 million), which is not meeting its timetable.

These investments are designed to

- determine whether a cloud can serve the DOE's mid-range computing needs,
- accelerate the commercialization of 100-Gbps networking technologies by deploying a national-scale prototype network with a test bed (which is now operational), and
- stimulate applied mathematics and computer science research efforts to advance the DOE mission across a wide range (applied-mathematics awards were made in support of DOE electricity-grid efforts).

There were problems in making the SciDAC awards because they were to live proposals from another call and transparency issues arose. Another call had to be made to current SciDAC awardees, and again transparency issues had to be dealt with, delaying the program by 6 months. Fourteen awards were made to EFRC–SciDAC [Energy Frontier Research Centers; SciDAC-e] collaborations, an $11 million engagement of 20 collaborators at 15 EFRCs.

In addition, approximately 30 new 2-year postdoctoral appointments were made at ASCR facilities to enhance user support. At the ALCF, six of the eleven postdocs are already on board, and five are starting in the next 6 months. They are working on such topics as negative-index-material calculations and an extended Hamiltonian replica-exchange molecular-dynamics simulation method. At the OLCF, six postdocs are already working, and two more are starting this year. They are working on a new chemical method for tailoring the band gap of semiconducting graphene nanoribbons and on developing the LAMMPS [Large-scale Atomic/Molecular Massively Parallel Simulator] kernels and the local second-order Møller–Plesset (LMP2) chemistry code to use the MADNESS run-time environment. At NERSC, eight positions were filled with postdocs. Two of the postdocs are co-authors of the best paper at the Cray User Group 2010 meeting in Edinburg, Scotland, for their poster presentation. Two of the postdocs served as teaching assistants for the NERSC-sponsored user course on GPUs.

All Recovery Act projects are under way. SciDAC-e is providing exciting early results despite being delayed in startup.

Giles asked if any ARRA money will be left on the table. Polansky replied, no.

Giles returned to the previous day's consideration of the exascale report. A letter to Brinkman had been crafted overnight. The letter stated that the Committee concurred with the main assertion of the Trivelpiece report that the case to pursue the exascale is compelling. The key findings of the report include:

1. Eleven workshops and three panels have explored the impact of exascale computing across DOE mission areas and key areas of national need, finding that exascale computing will be transformative for some applications and will benefit all science and engineering.
2. The United States and DOE have been leaders in high-performance computing since its beginning, and the exascale is the only path forward to maintain that leadership; exascale modeling and simulation will be key to competitiveness, and will benefit many other technologies and industries.
3. Codesign (linking developments in computer hardware and software, mathematics, and applications) will mitigate risk by allowing adaptation along the path to success.

Work continues on the full report, which will be delivered in early November.

The Committee concurred with the phrasing of the letter. Example applications were not to be cited so as not to tie the hands of the Subcommittee in drafting the final report.

Ping Ge was asked to describe the new SC Graduate Fellowship Program.

The program is housed in WDTS, which supports students at all levels from middle school to graduate school through internships and graduate research awards. It also supports educators from kindergarten to college. Three of six research offices (Advanced Scientific Computing Research, Biological and Environmental Research, and Fusion Energy Sciences) already had graduate research fellowships. A coordinated and prestigious SC Graduate Fellowship Program was seen to be needed for all six research offices. It will not replace the extant fellowship programs. The goal is to identify and support outstanding students pursuing advanced degrees in areas relevant to the SC and DOE missions, encourage the development of the next generation of the U.S. scientific and technical workforce, and encourage graduate students to pursue careers in basic research at DOE national laboratories and in academia.

$12.5 million of ARRA funding provided a jump start for the program, allowing 80 fellowships. The WDTS FY10 budget allocation of $5 million provides continuity. The WDTS FY11 budget request of $15 million, if appropriated, will cover a second year and a new cohort. Within SC there is a program management working group that identifies, recruits, and moderates reviewers with the help of the Oak Ridge Institute for Science and Education at Oak Ridge Associated Universities. That contractor lends logistic and software support in dealing with reviewers, applicants, and fellows.

The eligibility requirements are U.S. citizenship, being an undergraduate senior or first- or second-year graduate student, and pursuing an advanced degree in areas of basic research important to SC with a letter of acceptance into a PhD program. About 10% of the awardees are undergraduates, and 90% are graduate students. The application consists of an application form, proposed plan of research, transcripts, and three letters of recommendation.

The fellowship is a 3-year award, totaling $50,500 per year and includes a living stipend, tuition and fees, and a research stipend. It also includes attendance at a Fellowship Program research conference, which was held this year at ANL, where fellows share their research with other fellows and invited researchers from universities and national laboratories. Guest lectures, tours of the host laboratory, professional-development seminars, and workshops on how to access the DOE user facilities supplement an orientation for new fellows.

The first merit review of the 3216 applicants was of academic performance; the second was of the scientific and/or technical merit of the proposed plan of research; and the third was of scientific and technical contributions by the applicant outside the classroom. These reviews are carried out in (1) an eligibility and compliance review, (2) an online review with three external reviews, and (3) an on-site review of 400 to 500 finalists to select 150 awardees.

The fellows will go to 51 institutions in 25 states. They are 33% female and 67% male; 22 are undergraduates, 56 first-year graduate students, and 72 are second-year graduate students. 147 of the fellows met at ANL; they toured the ALCF, heard science talks, saw posted presentations, and learned about user facilities.

For FY11, the program will be announced in September 2010, applications will be due in October 2010, the online review will be conducted in December 2010 and January 2011, the on-site review will be held in March 2011, and notification will occur in March or April 2011.

The program had a COV from the Basic Energy Sciences Advisory Committee (BESAC) in May 2010. Lessons have been learned from the first-year experience, the application was improved, and the review process was improved. Significant improvements are being made each year.

Giles noted that most people in the Computer Science Graduate Research Fellowship Program are doing multidisciplinary work and asked if interdisciplinarianism could be stressed in the selection process. Ge said that that interest is equally important across offices, and some offices “shared” a fellow. The Office of Biological and Environmental Research (BER), ASCR, and FES fellowship programs stress a strong focus in their respective areas. The SC fellows are encouraged to be very creative. She did not see any emphasis on a strong focus or on interdisciplinarianism.

Giles asked how the program would compare to the NSF fellowship program. Ge replied that the two programs cannot be compared. The NSF program is 59 years old and has supported 33,000 fellows and produced 33 Nobel laureates. This program is in its infancy.

Voigt asked if there were any plans for supporting fellows beyond the current term limit. Ge answered, no; they can be supported by research grants.

Hack asked what the steady-state goal was. Ge replied that the goal was to establish three cohorts with a total of 350 to 400 fellows. Hack asked if there were any issues of interactions between the applications and national laboratory support of the research proposed. Ge said, no. Also, the SC program does not have a practicum, as does the Computer Science Graduate Research Fellowship Program.

Berger noted that the SC Graduate Fellowship Program is not open to permanent residents and asked if that would change in the future. Ge responded that she did not know; it may change.

Tang said that it would be good to include interdisciplinarianism in the selection process. Shared experiences are good. The practicums might be included in the SC program; they are broadening experiences. Ge stated that, currently, the practicum is optional; it will be further considered.

Giles stated that ASCR has carefully crafted a Computer Science Graduate Research Fellowship Program and wished to support the SC program in a way that does not undermine the Computer Science Graduate Research Fellowship Program in any way. It has taken a while to tune the program to get it to where it is today.

The floor was opened to public comment. There being none, the meeting was adjourned at 11:37 a.m.

Respectfully submitted,
Frederick M. O’Hara, Jr.
Recording Secretary