Minutes
Advanced Scientific Computing Advisory Committee Meeting
Aug. 11–12, 2009, American Geophysical Union, Washington, D.C.

ASCAC members present:

Marsha J. Berger
Jacqueline H. Chen
Roscoe C. Giles, Chair
James J. Hack
Anthony J. G. Hey
Thomas A. Manteuffel
Linda Petzold
Vivek Sarkar
Horst D. Simon
Rick L. Stevens
William Tang
Robert G. Voigt
Victoria A. White
Thomas Zacharia (Tuesday only)

ASCAC members absent:

F. Ronald Bailey
David J. Galas
John W. Negele
Larry L. Smarr

Also participating:

Roch Archambault, Senior Technical Staff Member, IBM Toronto Lab
Melea F. Baker, Office of Advanced Scientific Computing Research, Office of Science, USDOE
William Brinkman, Director, Office of Science, USDOE
Christine A. Chalk, ASCAC Designated Federal Officer
Phillip Colella, Computational Research Division, Lawrence Berkeley National Laboratory
Dona Crawford, Associate Director of Computation, Lawrence Livermore National Laboratory
Vincent Dattoria, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Steven Hammond, Director, Materials Sciences Center, National Renewable Energy Laboratory
Barbara Helland, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Daniel Hitchcock, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Fred Johnson, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Steven Koonin, Under Secretary for Science, USDOE
Douglas Kothe, Director of Science, National Center for Computational Sciences, Oak Ridge National Laboratory
Paul Messina, Retired Director, Center for Advanced Computing Research, California Institute of Technology
Walter Polansky, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Frederick O'Hara, ASCAC Recording Secretary
Edward Seidel, Director, Office of Cyberinfrastructure, National Science Foundation

Michael R. Strayer, Associate Director, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Warren Washington, Head, Climate Change Research Section, Climate and Global Dynamics Division, National Center for Atmospheric Research

About 50 others were in attendance in the course of the two-day meeting.

Tuesday, August 11, 2009
Morning Session

Before the meeting began, new members of the Committee were sworn in by a member of the DOE Office of Human Resources. Chairman **Roscoe Giles** called the meeting to order at 9:00 a.m., welcomed those new members to the Committee, and asked each member introduce himself or herself. He reviewed the agenda and asked Michael Strayer to introduce the first speaker, **William Brinkman**, the new Director of the Office of Science.

The Office of Science (SC) has three missions: to perform science for discovery, to conduct science for national needs, and to construct and operate national scientific user facilities. SC has six major activities: Advanced Scientific Computing Research (ASCR), Basic Energy Sciences (BES, which has very large facilities), Biological and Environmental Research (BER, whose facilities are growing), Fusion Energy Sciences (FES), High Energy Physics [HEP, which is the mainstay of national scientific activities and which strongly supports the Large Hadron Collider (LHC)], and Nuclear Physics [NP, which has two main facilities, Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL) and the Continuous Electron Beam Accelerator Facility (CEBAF) at the Thomas Jefferson National Accelerator Facility (JLab)]. The total funding requested for SC for FY10 is $4.94 billion.

The major user facilities include advanced computational resources, five nanoscale-science research centers, the Joint Genome Institute, the Environmental Molecular Science Laboratory at Pacific Northwest National Laboratory (PNNL), atmospheric and environmental facilities, and three new biofuel centers.

The light sources dominate in the number of users served. Protein crystallography has become a major use of these facilities, and two Nobel prizes have resulted. SC facilities have served about 25,000 users in FY09, about half from universities; about one-third from national laboratories; and the remainder from industry, other agencies, and international entities.

The government had 3 years when budgets could not be passed by Congress. This failure wreaked havoc on the funding of science. SC’s funding has grown in FY09 and hopes to grow again in FY10. One caveat is that the Obama administration will have to bring the deficit down in FY11. The promised doubling of science funding is being based on the FY06 appropriation.

DOE has a long history of training scientists, mathematicians, and engineers through research grants, the national laboratories, and targeted education programs. SC will support more than 4400 graduate students and 2700 postdocs in FY09. The DOE National Science Bowl engages 22,000 high school and middle school students every year. With American Recovery and Reinvestment Act (ARRA) funds and the FY10 request, SC initiated the DOE SC Graduate Fellowship Program, supporting more than 160 graduate students in fields important to SC missions. SC proposes to increase the Graduate Fellowship Program to support approximately 400 graduate students in the out years.

DOE also has an Early Career Research Program (which was just started) to begin providing 5 years of support to about 65 people this year. That program has received about 2200 letters of intent.

For FY09, SC received an enacted appropriation of $4.76 billion plus $1.6 billion from the ARRA. Its FY10 request to Congress is $4.94 billion, about a $200 million increase. The House passed a $4.94 billion appropriation, and the Senate passed a $4.90 billion appropriation. Resolution of this difference by the House and Senate committees will occur after the August recess.

$409 million was requested for ASCR in the FY10 budget request. The House agreed with that amount, and the Senate set the appropriation at $399 million.

One controversial thing that DOE did was to propose Energy Innovation Hubs on solar electricity, fuels from sunlight, batteries and energy storage, carbon capture and storage, electrical grid systems, etc. This proposal went through Congress before Under Secretary Koonin was confirmed, and Congress was confused about the proposal. Three such hubs might get funded. Each hub will comprise a world-class, multidisciplinary, and highly collaborative R&D team working largely under one roof. These teams will focus on solving critical technology challenges that prevent large-scale commercialization and deployment of the energy systems needed to address our nation's greenhouse-gas-emission, energy security, and workforce-creation goals.

The goals of the ARRA are to preserve and create jobs and to spur technological advances in science and health. SC ARRA projects were selected if they were shovel-ready, enhanced research infrastructure and supported high-priority R&D, were low risk, and created no out-year mortgages [with two exceptions: the Energy Frontier Research Centers (EFRCs) and the Graduate Fellowship/Early Career Awards].

SC had 51 ARRA projects, totaling $1.6 billion. The National Synchrotron Light Source II (NSLS-II) at BNL got $150 million to take x-ray microphotographs. JLab got $65 million for its 12-GeV upgrade. Science laboratory infrastructure construction got $108.5 million. The NuMI Off-Axis ν_e Appearance (NOvA) major items of equipment (MIE) got $55 million. Advanced networking got $66.8 million. And EFRCs were created with $277 million in forward funding for 5 years.

The total ASCR ARRA funding was $152,715,000 in the advanced networking initiative, cloud computing for science, SciDAC-e, a leadership-computing upgrade, and advanced computing architectures.

Beyond FY10, ASCR's science for discovery will pursue developing mathematical descriptions, models, methods, and algorithms to further understand complex systems. These computers have changed the way science is done. Now at the petascale, they are moving toward the exascale. The hope is for upward scaling in architecture.

Computing is central to addressing complex issues in science, like climate. If the observed changes in global average surface temperature and global average sea level continue, the world will have to change how it lives. It is estimated that it will cost $80 to 100 per ton to capture carbon dioxide. There is a major challenge to bring the cost of sequestration down. Although methane emissions from agriculture, waste, and energy have remained relatively unchanged since 1970, slight increases in N_2O (from agriculture and other sources) and in CO_2 (from deforestation, decay, and peat) and large changes in CO_2 (from fossil fuel burning and other sources) are going to cause deep trouble in 20 years.

Stevens asked Brinkman what he would like to achieve in computing. Brinkman replied that computing will make a huge difference in understanding climate. People need to make huge changes in the way they live.

Giles noted that the purpose of ASCR is to ensure the success of the other offices in SC. To do so, workshops have been held. He asked how the information produced by those workshops can be integrated with the rest of SC. Brinkman replied that one worry was about integrating the computational efforts of SC and the National Nuclear Security Administration (NNSA). Cooperation is needed, not competition. Another worry is not having enough money.

White asked about his view of DOE’s relationship with the National Science Foundation (NSF). Brinkman answered that the members of SC and other DOE offices talk with NSF members continuously about cooperation. He has not worried about computing but has about the Deep Underground Science and Engineering Laboratory (DUSEL), where an effort is being made to set up neutrino experiments. That will likely go forward under NSF leadership.

Zacharia noted that Brinkman had mentioned climate as a major user of computational resources and asked whether, as the exascale is approached, DOE and SC should support the infrastructure of the science that could use these resources. Brinkman replied, yes; but clearly, the effort must be driven by the applications. Certainly, one must ensure that the couplings are there. One must balance facilities with research.

Stevens noted that, as the enterprise turns to software to enhance performance, the level of engagement with the user community will be greatly heightened. The engagement model’s structure is a great concern. Brinkman replied that there was a similar transition upon the introduction of parallel and vector machines. Stevens pointed out that this transition will be a multilevel challenge. Brinkman stated that he understood that.

Manteuffel asked Brinkman where he saw the algorithmic development etc. coming from. Brinkman replied that he did not know the answer to that question.

Sarkar stated that, in enhancing performance, one will have to pay attention to all the different levels of the software stack. Brinkman agreed that it is frightening.

Tang noted that there needs to be an encouragement of the younger members of the computational community and asked how the educational support will be sustained. Brinkman answered that it is under control. DOE needs to watch out for treading on the turfs of the Office of Education and NSF.

Simon asked about materials research. Brinkman replied that the Office is meeting with these facilities. The people in charge of DOE today are more oriented toward research than in the past. If carbon sequestration is not advantageous, one needs to turn to nuclear energy. Right now, the experts are being sought to understand what science needs to be done to develop the materials and fuels for that technology

Hey noted that the challenge of data (especially with the LHC, ocean observatories, and climate satellites) and their interpretation is huge. Brinkman agreed.

Zacharia said that Congress may fund fewer than eight hubs and asked whether there might be one within the computing sciences. Brinkman observed that in the computer sciences there are already some things that look like hubs and act like hubs.

Chen asked about engagement of computer sciences in the Outstanding Junior Investigator (OJI) and Early Career Award (ECA) programs. Brinkman answered that he did not know how to respond.

The floor was opened to public comment. Crawford asked what Brinkman would do to encourage SC and NNSA to work together. Brinkman replied that common goals and interests are identified, and a way to work together to accomplish those common goals and interests is sought.

A break was declared at 10:02 a.m. The meeting was called back into session at 10:18 a.m. **Michael Strayer** was asked to present an update on the activities of ASCR.

ASCR has been understaffed for several years, largely because of attrition. As a result, the staffs of the divisions work in teams. Walt Polansky is the Acting Director of the Research Division, which has teams in Applied Mathematics, Computer Science, Computational Partnerships (which are going to be expanded into the NNSA and agency-wide as well as into the NSF and other agencies), Next-Generation Networking (including collaboratories), and the Education Program (which will have an aggressive build-out). Vince Dattoria leads the Facilities Division team. Larry James leads the Small Business Innovative Research (SBIR) team of the Small Business Division. New positions will soon be posted for research program managers in computer science and SciDAC and a facilities program manager.

For the FY10 budget, the request for Applied Mathematics is an increase of $4.686 million over the FY09 appropriation, the request for Computer Science is up $13.182 million, the request for Computational Partnerships (including SciDAC) is up $1.171 million, the request for the National Energy Research Scientific Computing Center (NERSC) is up $210 thousand, the request for Leadership Computing Facilities is up $15 million, and the request for High Performance Network Facilities and Testbeds (ESnet) is up $5 million. ARRA funds are also being used to build out the networking capabilities.

The FY10 highlights a new effort in advanced computer architecture design for science, increases support lease payments and site preparation at Argonne National Laboratory (ANL) for a proposed upgrade, and allows ESnet to deliver 100 to 400 Gbps to SC laboratories.

The House Energy and Water Development Appropriations Subcommittee reported out $409 million for ASCR, the same as the request. The Senate came close to that, reporting out $399 million for ASCR. The Committee expects SC to continue to support joint research with the NNSA laboratories.

Solicitations just closed for

- Mathematics for Analysis of Petascale Data, which is funded at $4 million per year and had 81 proposals from universities and national laboratories; the topic areas include anomaly detection, machine learning, streaming data, dimensionality reduction, and visualization; 11 awards were made, with 50% of the funds going to national laboratories and 50% to universities.
- Mathematics for Complex, Distributed, Interconnected Systems, which was funded at $3.5 million per year and received 38 proposals, all national-laboratory-led; 5 to 7 awards are anticipated with some university subcontracts.

- A Joint Math/Computer Science Institute to bridge the gap between large complex scientific applications software and next-generation hardware, which is funded at $4 million per year and received 29 superb applications; three awards have been made so far, two laboratory-led and one university-led.
- Ice-Sheet Modeling to develop fully dynamic ice-sheet models and ocean/ice shelf interactions and to assess the rate and magnitude of sea-level rise caused by rapid ice-sheet melting, which is funded at $3 million per year for 3 years and received eight proposals of which six were awarded, two university-led and four laboratory-led.

Three ASCR-funded projects won *R&D 100* awards: PETSc, a suite of data structures and routines for solving partial differential equations at scale; ROSE, a compiler infrastructure; and Catamount, the lightweight-kernel operating system.

Presidential Early Career Award for Scientists and Engineers (PECASE) awards were won by Cecilia Aragon for her groundbreaking research in data-intensive scientific workflow management, and pioneering development of innovative methods for visualization, analysis, and organization of massive scientific data sets; Alexandre Tartakovsky for his research on subsurface flow that addresses past and future energy needs like storing carbon dioxide from fossil fuels underground; and Oliver Fringer, who was a fellow in the ASCR Computational Science Graduate Fellowship Program from 1997-2001.

Of the 191 Society for Industrial and Applied Mathematics (SIAM) fellows named in 2009, more than 40 have been or are currently funded by ASCR.

Expert panels are conducting decadal reviews and rankings of accomplishments in Applied Mathematics, Computer Science, and Computational Science.

NERSC is in the process of a quad-core upgrade to Franklin, which was accepted on June 17, 2009, a third of the petaflops. The NERSC-6 contract was awarded to Cray for an at-least-1-petaflop Cray XT5. At the leadership-class facilities (LCFs), the mission need was approved in January 2009. The upgrade of the LCFs to tens of petaflops by 2013 is driven by several important international programs, including climate science fusion energy research, and the Nuclear Energy Advanced Modeling and Simulation program. Follow-on Lehman reviews were held at the Oak Ridge LCF on July 7-8, 2009, and at the Argonne LCF on July 28-29, 2009. An operational assessment review of the Oak Ridge LCF will start in August 2009.

For the ARRA, ASCR identified several potential projects. SC reviewed the SC-Programs list and developed an SC-wide list. The SC priority list was reviewed by Secretary Chu to develop DOE's proposal to the Office of Management and Budget (OMB). ASCR's projects were the Advanced Networking Initiative, an Oak Ridge LCF upgrade that will take the Jaguar to 2 petaflops, advanced computer architectures, Magellan (cloud computing for science), and SciDAC-e that enhances SciDAC for energy applications.

Several workshops have been completed: the BER/Climate Workshop, HEP/High Energy Physics Workshop, NP/Nuclear Physics Workshop, FES/Fusion Workshop, and NE/Nuclear Energy Workshop. Upcoming workshops will be the Materials Workshop with BES, the Biology Workshop with BER, and the ASCR Workshop with NNSA. Other workshops are also expected.

The NNSA-ASCR Exascale Partnership has established an Executive Oversight Committee and a Steering Committee. A university consortium or partnership will be part of this Steering Committee. The initial tasks of the Steering Committee are to develop the science case for the Exascale Initiative that focuses on the Department's energy, environmental, security, and societal grand challenges that will require exascale computing to solve and to develop a high-level roadmap for the Exascale Initiative that identifies start dates, approximate duration, and high-level dependencies.

Berger asked what role ASCAC would play in achieving the objectives cited. Strayer replied that DOE is slow in getting a charge to ASCAC. It has formulated one and will get it to the Committee as soon as possible. Berger asked if it would be for a specific topic. Strayer answered that it would be for the exascale issues that need to be addressed. Berger asked him how he saw this Committee helping his Office in general. Strayer replied that ASCAC is very helpful in socializing the responses to issues that the Office faces. It is hoped that ASCAC is proactive and takes strong stands on issues. Giles pointed out that the Committee will have to figure out how to form the subcommittee for that charge, given that the charge likely will be received between Committee meetings. Stevens stated that the Steering Committee would like to have as much interaction as possible with ASCAC.

Sarkar asked if there were a feeling by the Office that hardware and software need to be codesigned. Strayer replied, yes. They have to be coherently together. That is the model. It was used with the Blue Gene LPQ and was very successful.

White asked if he could say more about what the charge will be. Strayer answered, no; it will come from upper DOE management.

Strayer introduced **Steven Koonin**, the new Under Secretary of Energy for Science. He is the chief scientist of DOE. His role in DOE is to enable cross-disciplinary research to flourish and to raise the quality of science in basic science, energy technology, and the NNSA. He also tries to ensure that the Department's policies have sound technical foundations.

Koonin stated that America's energy challenges include (1) the security of supply and (2) greenhouse-gas emissions. Domestic production has decreased, while imports have increased. The goal is to increase domestic production by 3.5 million barrels per day while reducing greenhouse-gas emissions.

Energy has many aspects, as the nation's energy-flow diagram shows. The United States uses about 100 EJ per year. This situation offers many places where science can make a difference in the energy used and how it is used. Energy sources and uses change over time and have long intrinsic timescales, which are driven by scale, ubiquity, incumbency, and interoperability.

Computation has claimed a role in energy-related science. Its influence and contributions should expand (e.g., into nuclear energy). If those problems are going to be tackled, science, tools, and algorithms have to be melded together, which will require laboratory work, validating the data, and simulation. We need to involve theorists, computational scientists, experimentalists, applied mathematicians, and computer scientists.

The tools have changed rapidly. There have been great increases in speed and power consumption. Power management will be a big deal. New computer architectures will need to be devised. The algorithms and models are important in improving performance.

DOE has been successful in going from the gigascale to the petascale. Going to the exascale will take a decade. It will require scientific justification, breakthroughs in hardware and software, and building interdisciplinary communities (which will require integrating the knowledge of theorists, experimentalists, computational and computer scientists, and applied mathematicians and will require collaborating with NNSA).

Giles commented that he appreciated scientifically based leadership.

Stevens asked how Koonin would envision SC laboratories' transforming their methods to those like NNSA. Koonin said that there is a discipline of scientific computing, but it is different from using simulation to optimize technological processes. What is being talked about is engaging talent from all three stovepipes in the needed simulations (e.g., for nuclear energy, carbon sequestration, and combustion engines). Climate models are being thought about seriously; they need to produce better results.

Zacharia asked whether, as computing matures, DOE should consider a different engagement model, looking at the offices as opportunities to engage the scientific community. Koonin stated that the Department should be doing the best science it can. Let the scientific community decide where that should take the Department. On the other hand, the scientific community has to be fully informed. Zacharia agreed; ASCR has an opportunity to create that group of scientists needed to tackle these problems. Koonin said that he would not stop at science; the Department should be doing simulations of, say, chemical engineering processes.

White stated that DOE seems to have become risk-averse and asked how risk should be balanced with certainty of results. Koonin replied that, in research, one is supposed to fail, otherwise it would not be research. Tolerance for failure goes down as one advances along the technological path.

Tang pointed out that SC invests more money in facilities and computers and asked how the need to be predictive could be balanced between these scientists and the facilities. Koonin replied that the theorists need to tell the observational scientists what the critical experiments are that should be performed and what variables are important and must be gotten right.

Hey pointed out that, for biological and environmental scientists to get their models right, they must have lots of good data. There are three parts to environmental science: discipline-based, Earth-system, and environmental applications. Koonin observed that the scientific community members are not the only ones interested in these issues.

Manteuffel asked if there were any ideas in fostering the bridging between science and computation. Koonin said that, in biology, before 1990, they did not think about computers as tools. Now, bioinformatics is a mainstay because it has demonstrated its utility. Scientists will move toward computational methods in materials science and other disciplines when they see the utility of computing. The coupling of high-throughput experimentation with computing will be a major player in new science.

Sarkar asked what was feasible in energy consumption for data storage and manipulation. Koonin said that he was not an expert in that topic but recognized that the issue had to be addressed. Giles said that the computing done now in energy consumption is in non-equilibrium. Also, the next generation of computer scientists needs to be trained.

Stevens stated that a mega-sea-change would occur if information technology were central to the research process.

The floor was opened to public comment. Colella said that the NNSA data point is scary because it is not just a management problem. It is spread over many disciplines, and the question of architectures arises. It will be expensive. Koonin replied that a couple of problems will need to be picked to see if the economic, physical, and sociological barriers can be overcome.

A break for lunch was declared at 11:48 a.m.

Tuesday, Aug. 11, 2009
Afternoon Session

The meeting was called back into session at 1:06 p.m., and **Steven Hammond** was asked to discuss high-performance computing for renewable energy.

The Workshop on Research Needs for Alternative and Renewable Energy was held about two years ago, preceding the Intergovernmental Panel on Climate Change (IPCC) fourth assessment report on climate change. Since then, budget of the Office of Energy Efficiency and Renewable Energy (EERE) went from $1.5 billion per year to $16 billion under the ARRA. The workshop had five parallel breakout sessions on renewable fuels: H_2, renewable electricity: photovoltaics, renewable fuels: bioenergy conversion, renewable electricity: wind, and grid reliability.

Molecular hydrogen has tremendous promise as a potential future energy carrier that is free of greenhouse gases, but it needs to employ onboard vehicle storage within the current form factor. Platinum-free fuel cells are essential for mass market adoption of hydrogen as a fuel.

The focuses in the hydrogen breakout sessions were (1) rate processes in hydrogen production, storage, and use (which need new methods for treating kinetics and reactions, new methods for determining reaction surfaces and predicting reaction pathways coupled to kinetics, and multiscale approaches to electronic-structure-based dynamics and atomistic approaches to large systems); (2) inverse materials and systems by design (which need novel optimization techniques tightly coupled with existing simulation tools to control catalysis, charge separation, spatial component arrangement); and synthesis of targeted material (which needs novel, scalable tools to calculate pathways and predict reactions in solution necessary for the fabrication of targeted materials). Models are needed to link scales ranging from atoms to systems because it is cost-prohibitive to simulate large systems with uniformly high levels of detail.

Photovoltaics seek to exploit the fact that sunlight is by far the largest of all potential carbon-neutral energy sources. The energy from sunlight striking the Earth in one hour exceeds annual global energy consumption. However, their efficiency needs to be improved, and their costs need to be lowered. There is a lot of interest in modeling materials for property prediction and design. What is needed are

- Improved scaling in density functional theory to model systems with more than 10^3 atoms,
- Better methods for carrier transport and charge separation in organic photovoltaics,
- New multiscale codes for predicting material properties that provide sufficient accuracy at variable scales, and
- Materials by design.

On the cost side, there is no good tool. For example, chemical vapor deposition (CVD) is used to produce thin-film photovoltaic materials. Numerical simulation of the CVD process would provide unique insights in the complex coupled chemistry, energy, multispecies transport, and flow patterns that determine the film characteristics that would improve film uniformity, reduce costs, and improve reliability.

In biomass conversion, most ethanol comes from corn starch. Cellulosic biomass offers the greatest source of renewable transportation fuel. A DOE and Department of Agriculture study projected that it would be possible to displace 30% of current U.S. fuel use by 2030. Fundamental advances in the understanding of the structure/function relationships governing the activity of soluble enzymes on insoluble polymeric substrates are essential to make these fuels cost-competitive. This task is being done at the three DOE Bioenergy Research Centers. Priority research directions include modeling the plant cell wall and modeling cell-wall-degrading enzymes. Feature-rich codes like CHARMM do not scale to large processor counts, and scalable codes (AMBER, LAMMPS, and NAMD) do not have the desired features.

For thermochemical conversion, what is needed is to develop integrated multiscale models to accurately simulate kinetics and thermodynamics covering pretreatment, pyrolysis, and gasification and to simulate and optimize design of biomass conversion plants, replacing the need to build numerous pilot-scale facilities and the traditional trial-and-error experimentation needed to bridge the gap between laboratory scale and production.

In wind, there is a National Academy of Sciences (NAS) study that states that, in the Midwest prairie states alone, there are sufficient wind resources to supply 16x current U.S. electrical demand using existing technology. It currently supplies 3% of U.S. demand. In 2008, 48% of new electrical supply came from wind. The DOE goal is to satisfy 20% of U.S. demand from wind by 2030. At one point, the DOE/EERE Wind Program sponsored efforts to develop codes to explore the complicated, unsteady, 3-D flow field that are critical to cost-effective design and efficient operation. Modeling wind turbine aeroacoustics as well as integrated gearbox reliability is an important research topical area for further efficiency and total-cost-of-ownership improvements.

Today’s electricity comes from a centralized architecture in which electricity generation is geared to peak demand. For tomorrow, a distributed-generation model is needed but has great needs for load management. A national grid simulation capability could greatly improve current simulation methods and would include advanced contingency analysis, multiple-physics effects, system interdependencies, advanced control algorithms, improved model granularity, and multiple spatial and temporal scales. The amount of data used to run such a system is going to be overwhelming. A flight simulator for the new grid, enhanced data streaming and analysis, predictive modeling with mixed digital and analog inputs, scalable nonlinear dynamics algorithms to describe coupled states, fast state-estimation algorithms for real-time control and operation, and large-scale application of game theory are going to be needed.

The common themes from the workshop were

- Rational design of materials, devices, and systems,
- Multiscale models, from atoms to systems,
- Multidisciplinary collaboration, and
- Additional computational resources.

In some instances, simulation codes exist but they do not exploit multicore architectures or scale to large processor counts. Over the years, ASCR has funded a significant body of work in advanced numerical libraries and in adaptive and multiscale methods that can be applied to challenges of interest to EERE.

From the DOE, there should be seamless progression from discovery research to use-inspired basic research to applied research to technology maturation and deployment. But traditionally, there is the valley of death between basic research and applied research, manifested between SC/ASCR and the technology offices like EERE. To bridge this valley of death, the offices could form some focused working groups for each area with representation from both SC/ASCR and EERE to review the broad array of solicitations [Advanced Research Projects Agency-Energy (ARPA-e), ARRA, EFRCs, Scientific Discovery Through Advanced Computing (SciDAC), etc.]; identify gaps and prioritize areas of mutual interest to the two offices; and then develop a SciDAC-like program to link applied mathematics, computational science, advanced algorithms, multiscale methods, and petascale computing to EERE priority areas.

Manteuffel asked where the big turbines are going in. Hammond replied, at the National Wind Technology Center (near Rocky Flats), part of the National Renewable Energy Laboratory in Colorado.

Berger asked where the research funding was coming from now. Hammond answered that General Electric does some research, but not the fundamental research. They are not interested in funding research; their order books are full. But there is not much research, just the use of available information.

Giles asked if there were some follow-up that the Committee needs to supply. Hammond said that inquiries placed today would be responded to differently than they would have been a few months ago.

Simon noted that this community is ready to run on a few parallel systems. There is a basic lack of understanding of what exascale simulations could do. Hammond said that there are a few select areas where a SciDAC project would be effective, like looking at fundamental functions of enzymes, aerodynamics, and wind-farm layout.

Manteuffel pointed out that there is a difference between a SciDAC project and an exascale project. This area could benefit from a SciDAC team.

Zacharia noted that engineers tend to use the tools available. They can use other tools. Making tools available to design engineers has a huge impact. SciDAC can help, and several examples exist that could use the peta- and exascales. High-performance computing would lend a unique capability in making transformational changes.

The floor was opened to public comment, but there was none.

Phillip Colella was asked to discuss the potential impact of high-end-capability computing.

The National Research Council (NRC) produced a report entitled *The Potential Impact of High-End Computing on Four Illustrative Fields of Science and Engineering*. This report was intended to make a case for computational science that would be useful for policy development. The purpose of the study was to identify the subset of those important questions and problems for which an extraordinary advancement in understanding is difficult or impossible without high-end capability computing; to identify the numerical and algorithmic characteristics of the high-end capability computing requirements needed to address the identified scientific questions and technological problems; and to categorize the numerical and algorithmic characteristics, specifically noting those categories that cut across disciplines.

Four fields of interest to the federal government were selected: astrophysics, atmospheric science, chemical separations, and evolutionary biology. The study included four meetings, small disciplinary workshops, and an external review. The report was released on August 26, 2008, with a related symposium at the NAS on September 22, 2008. It was then used to brief OMB and the Office of Science and Technology Policy (OSTP).

The study reached three general conclusions: (1) High-end capability computing is advanced computing that pushes the bounds of what is computationally feasible. It requires innovation and poses technology risks. (2) Advanced computational science and engineering is a complex enterprise that requires models, algorithms, software, hardware, facilities, education and training, and a community of researchers attuned to its special needs. (3) Decisions about when, and how, to invest in HECC should be driven by the potential for those investments to enable or accelerate progress on the major challenges in one or more fields of science and engineering. This should be a science-driven enterprise.

In astrophysics the major challenges that require high-end computing lie in answering the questions:

- How did galaxies, quasars, and black holes form?
- How do stars and planets form? And evolve?
- What are the mechanisms for supernovae and gamma ray bursts?
- What will the universe look like observed in gravitational waves?

In atmospheric sciences the major challenges that require high-end computing lie in

- Extending the range, accuracy, and utility of weather prediction;
- Improving the understanding and prediction of severe weather, pollution, and climate events;
- Determining the effect of seasonal, decadal, and century-scale climate variation on global, regional, and local scales;
- Understanding the effects of moisture and chemical exchange at Earth’s surface; and
- Creating the ability to predict global change for the next 100 years.

In evolutionary biology the major challenges that require high-end computing lie in understanding the history of life, how species originate, and the origin and evolution of the phenotype.

And in chemical separations the major challenges that require high-end computing lie in predicting physical properties for phase equilibria (for difficult separations) and in designing and producing mass separating agents.

These requirements were analyzed for each scientific discipline/technology and crosscuts were discussed. The potential rate-limiting issues were found to be the availability of (1) models that are amenable to computation and stable, (2) algorithms to represent such mathematical models; (3) software infrastructure (middleware, communication, and high-level libraries), (4) facilities, (5) data analysis and management, and (6) people. The specific manifestations of these crosscutting issues in the selected fields of interest were spelled out.

In astrophysics/star formation, there is a need to transition from well-characterized “first-principle” models to multiscale/multiphysics models. The algorithmic issues include the need for multi-resolution methods, particularly for multiphysics coupling (e.g., radiation hydrodynamics, general relativistic fluid dynamics), and the need for implicit methods for stiff time scales. How to make sense of the big amounts of data, from both observations and simulation is also an issue.

In atmospheric sciences, there is a near-term need for a 10× increase in capability, with a concomitant effort in scaling codes to 10^4 processors (for numerical weather prediction). The next major increment in model fidelity for climate could be obtained by a transition from statistical models of clouds and precipitation (valid down to 25 km) to cloud-system resolving models requiring 1-km resolution in the tropics. This transition leads to a major reconsideration of many components of the model, algorithm, and software space. There are new opportunities and requirements in data assimilation due to the vast increase in sources of data associated with numerical weather prediction or the application of data assimilation to climate modeling.

In evolutionary biology, the current use of computing is mainly discrete mathematics and statistics for data mining. These areas are transitioning from workstation-based activities to high-end computing. Modeling is in its infancy. There is an enormous range of spatial scales (from the cellular level to populations) and temporal scales (from seconds to geological time scales). There is no first-principle model in sight; everything has to be bootstrapped.

In chemical separations, the current state of the art in computational chemistry is adequate for investigation of the qualitative behavior of separations processes, but not for end-to-end quantitative prediction and optimization. There are severe tradeoffs between fidelity and computability, and there is the combination of model and algorithmic issues. The focus on process optimization requires new mathematical and algorithmic tools. This is a large, industrial process, making workforce issues a concern.

Looking at the crosscutting issues, in simulation there is further movement away from first-principles models to ones that are bootstrapped from experimental/observational data and more detailed auxiliary computer simulations. Moving in this direction produces more of an emphasis on hybrid stochastic–deterministic models. In algorithms, overlapping requirements include multi-resolution methods, methods for stiff systems, and high-performance particle methods. These algorithms must have strong interaction with new models. And in software, Moore’s law is dead. Complexity of multiscale, multiphysics simulation codes are prominent concerns here, with major impacts on software productivity. The scientists in all four fields see the handwriting on the wall. To various extents, they are willing to adopt the use of shared software infrastructure. There are, however, tradeoffs between domain specificity and shared software that are different for the different fields. The question of where the software will come from has not been resolved.

The crosscutting issue of data is important because astrophysics, atmospheric sciences, and evolutionary biology are data-intensive. They all have high-throughput experiments or observations and simulation data and are racing each other to use up all the available resources. The challenges are knowledge discovery from data, sharing of data, analysis and mining of data, and high-performance input/output. There is strong variation among these communities in the interaction of data with simulation.

In the crosscutting issues of education and training, all four fields have, to varying degrees, been successful in integrating computational science into their disciplines. They use and value computing. There is a need for better training in the fundamental mathematics and computer science that underpins the use of high-end computing. There should be a greater emphasis on software development in education of computational scientists and engineers. Most PhD dissertation projects are "proof-of-principle" implementations with no persistent software artifacts. There are several institutional problems. One is that there are two approaches to the integration of computer science and engineering (CSE) into scientific disciplines: disciplines can take ownership, or a new CSE disciplinary area can be formed. There is also the problem of a career track for application-software developers.

The specific conclusions that the study arrived at were that (1) the emergence of new hardware architectures precludes the option of just waiting for faster machines and then porting existing codes to them. The algorithms and software in those codes must be reworked. (2) All four fields will need new, well-posed mathematical models to enable high-end computing approaches to their major challenges. It is not just about hardware or algorithms; it is about models. (3) To capitalize on high-end computing's promise for overcoming the major challenges in many fields, there is a need for students in those fields, graduate and undergraduate, who can contribute to high-end-computing-enabled research and for more researchers with strong skills in high-end computing.

He closed with some final, personal comments. The operational realities of selling science programs to policymakers requires an airtight detailed science case placed in a larger context and dealing with timing issues ["Why this? Why now?"]. Cheerleading by computational scientists is not credible; outside validation is needed. Models (to some extent), algorithms, and facilities all have well-established funding mechanisms, and big data is getting a strong push. One might ask where the funding mechanism for software is. A CSE ecosystem is needed that includes training, career tracks, and persistent institutional commitments that conserve the skill base in these areas.

Manteuffel asked how Colella would restructure a university to get CSE training and research. Colella responded that, for training, that process has to be a collaboration between applied mathematicians and computer scientists. The research problem is harder. There are not a lot of choices for a stable career.

Stevens asked what other area would have been included in the study. Colella answered, structural biology.

Petzold said that, to every group that has studied this problem, universities and industries have said that the bottleneck is people, worldwide. The bottleneck is not just in PhDs but in incoming graduate students. Engineering schools have taken computing out of the undergraduate curriculum. Something has to be done to reintroduce computing to the engineering curriculum. Historically black college and university students are not getting computing experience. Colella responded that this is a jurisdictional problem. Ed Seidel is the person to whom to address that question. How to grow the teams to address the big problems mentioned by Koonin needs to be figured out. Research teams in computational science at the national laboratories are capable of rapidly training practitioners through their postdoc programs and could double the number of computational scientists in those institutions within 3 years to staff up such large new projects.

Hey said that it is difficult in universities to introduce computing. Universities are disciplined based. Arizona State University has abolished engineering disciplines and organized teams around themes and problems. Colella said that the California Institute of Technology is organized that way, also.

Chen asked what the rate-limiting factors are with data. Colella responded that it varies from community to community. With some communities, the problem is data sharing.

Berger asked about the long-term outlook for software. Machines come and go faster than quality software can be written. Colella answered that NSF has centers, the Accelerated Strategic Computing Initiative (ASCI) has centers, and SC has centers. Whenever such software centers go away without any mechanism for providing continuing support for the software, scientists that have depended on the software are left stranded. They will not trust these agencies in the future, and will refrain from using professionally developed software.

A break was declared at 2:36 p.m. The meeting was reconvened at 2:50 p.m., and **Roch Archambault** was asked to discuss the IBM Power Systems' compiler strategy.

The IBM Compilation Technology Group has more than 300 development, test and service engineers who deal with C, C++, and Fortran compilers. They try to standardize as much code as possible to minimize the need for testing. They are responsible for Java just-in-time compilers and have developed compilers for commercial markets, C/C++, COBOL, and PLX for zSeries during the past 25 years.

The Toronto compiler group was involved in the first system to achieve 1 TF sustained in the late 1990s (ASCI Blue) and supported the ASCI collaboration on the White, Purple, Blue Gene, and Roadrunner systems. It also participated in the High-Productivity Computing Systems (HPCS) collaboration with delivery of UPC [Unified Parallel C] and programmer-productivity improvements and the Blue Water and Sequoia collaborations. It ships its compilers every 2 to 2.5 years, driven by new hardware.

Common Fortran, C, and C++ features include 64-bit Linux, debug support, full support for debugging of OpenMP programs, etc.

In addition to the language front ends, the IBM XL compiler architecture includes the Toronto Portable Optimizer, partitions, TOBEY (which has a million lines of code, all in C++), interprocedural analysis (IPA) objects, libraries, profile directive feedback information, optimized objects, other objects, and a system linker. The compilers currently have manual intervention. They are moving toward semi-automatic parallel languages. The dream is automatic optimization. Future compiler releases will provide analysis of MPI applications because there is a lot of message-passing-interface (MPI) code out there.

Some of the new stuff in Fortran includes object-oriented extensions (inherited type extensions and type-bound procedures). In OpenMP 3.0, a major addition was task support.

The XL UPC compiler is built on top of the C++ compiler. In a hybrid execution environment, one can have a partitioned global address space. These language extensions make UPC more effective at optimizing a wider range of scientific applications.

For a Coarray Fortran Compiler (CAF), a user community has to be built up. The Coarray syntax uses codimensions, producing a ray on all the images so one can address all of the images. In Coarray Fortran, image numbers and threads are much simpler than they are in UPC.

The Toronto Group has been working on an OpenCL compiler for CELL and POWER processors. It seems to have some advantages. OpenCL C has less noise and is easier to read.

Delinquent load identification deals with bandwidth and latency issues. A very small number of *delinquent* loads are responsible for the vast majority of cache misses. There are three ways to identify delinquent loads: user annotation, static analysis, and dynamic profiling. In dynamic profiling for delinquent load identification, one first wants to do coarse-grain cache-miss profiling and then fine-grain cache-miss profiling. One instruments the code and then re-instruments the code, allowing one to focus in on the problem. But one needs to know where to put the prefetch and when to run it. There are a number of processes for delinquent-load-driven optimization.

Moving on to assist threads, there are chip multiprocessor threads and multicore threads. The level of cache shared is important. But not all the threads can be used or are available for prefetching. Therefore, one does delinquent load identification, loop selection, region closing, and back slicing. Assist threads code identifies available threads and dispatches them. Some provisioning results have been seen. Metrics used for different cache miss rate for delinquent loads only are: high when the miss rate is about 90%, medium when the miss rate is about 40%, and low when the miss rate is about 20%. With Assist Threads, one does not get a lot of speedup in the benchmarks, but there is some there. User annotated and automatic assist thread generation will be implemented under the –qprefetch=AssistThread compiler option.

Compiler transformation report infrastructure (generating XML from the compiler) can be fed back to the user or to a tool to do something about it. This information goes into the database at linking and at runtime. It then produces an XML schema or a style sheet that can be read on the browser.

The Group has had a lot of experience with automatic parallelization; lately, in Japan, enablement and cost analysis have been very important.

Hey asked how one knew that the scientist's codes and results had not been changed by automatic optimization. Archambault answered that the user is allowed to control the level of optimization. Also, the places in the code where results may have been changed are fed back to the user.

Stevens asked Archambault if he had thought how to push the prefetch-assist threads off the main central processing unit. Archambault said that he was all for it and has been asking for that. However, he had been told that it is too expensive to do that. Stevens asked if there is an external specification for the prefetch assist and whether there was any traction there. Archambault replied that he was only aware of what was done at IBM. There had been some experiments, and there was architecture but no metal. The pin people could do it if they had a specification, but the cost would be high.

Sarkar asked what the major game plans for compilers were as expressed by users. Archambault said that the question was how one is going to deal with all those processors and other elements.

Tang asked if there were any documentation on how effective this optimization has been. Archambault replied that some improvements have gone into specifications. There are benchmarks at each layer of improvement. Tang asked what was coming out of work in Japan. Archambault answered that array privatization would go a long way to automate parallelization. That is good for some users; others would like to parallelize the code themselves.

Vincent Dattoria was asked to describe some of the uses of ARRA funds.

All of this information has been vetted through all channels up to OMB. Under the ARRA, ASCR is pursuing Magellan, a leadership computing facility upgrade, advanced computer architectures, SciDAC-e, and an advanced networking initiative.

The request for Magellan is for $33 million. Magellan is a research and development effort to establish a worldwide, scientific, midrange, distributed computing and data-analysis testbed to see if the costs of midrange computing can be lowered by eliminating duplication from facility to facility. It will have two sites (NERSC and the Argonne LCF) with multiple tens of teraflops and multiple petabytes of storage, as well as appropriate cloud software tuned for moderate concurrency. It is hoped to get some statistics from this effort and an idea of where such midrange computing can be used.

The request for the Oak Ridge Leadership Computing Facility upgrade is for $19.9 million. It will replace 37,376 quad-core Barcelona processors with 6-core Istanbul processors, producing a 70% increase in node peak performance and a 21 to 27% increase in node memory bandwidth. This equipment should be accepted in early winter and then opened to users.

The request for advanced architectures is for $5.2 million. It is a new effort to provide early access to DOE researchers to use technologies emerging from the IBM Productive, Easy-To-Use, Reliable Computing Systems (PERCS) effort and an enhancement of the University of California at Berkeley's Research Accelerator for Multiple Processors (RAMP) effort to provide focused research on flexible simulations of performance of scientific applications on next-generation microprocessors.

The request for SciDAC-e is for $30 million; it has three elements at $10 million apiece. The three components are wholly within the ASCR mission. They are (1) applied mathematics research to enable bigger, better, and smarter electrical grids; (2) supplemental awards to SciDAC centers and institutes to support collaborative research with BES-supported EFRCs to develop a high-performance computing capability relevant to the goals of the EFRCs; and (3) enhanced user support (postdocs) at NERSC, the Argonne LCF, and the Oak Ridge LCF for SciDAC-e and energy users awarded allocations through the ASCR Leadership Computing Challenge and a SciDAC-e summer school for training and experience with leadership computing resources. The expected outcome includes (1) algorithms to simulate the performance of electrical grids over a full range of operating conditions and (2) the provision of a software environment and intellectual resources to EFRCs to meet computational goals.

The request for the advanced networking initiative is for $66.8 million. It is an ambitious effort in research and development to establish a nationwide demonstration prototype with a 100-gbps throughput capability. It will connect four geographically dispersed sites, including NERSC, the Argonne LCF, and the Oak Ridge LCF, and a peering point in New York City. It will constitute a distributed testbed and associated research that will support advanced network research topics. All proposals were in hand and peer reviewed. This initiative will explore the challenges and solutions for installing and operating a 100-gbps optical backbone, will demonstrate the capabilities of 100 Gbps, and will map a strategy for an ESnet implementation.

Chen asked if Magellan would support MPI applications. Dattoria said that the architecture specifications cannot be discussed now. There is a desire to support MPI applications.

Giles asked what the timescale was. Dattoria replied that the money has to be allocated by the end of next year and then has to be costed in the following year. It should be possible to complete the upgrade at Oak Ridge in less than a year.

Giles asked whether the scientific community will be invited to use Magellan during the ARRA phase. Dattoria replied affirmatively. The leads are being scheduled to seek out potential users.

Tang asked if they would be linking with the Office of Nuclear Energy (NE) or FES in SciDAC-e. Dattoria said that there are no plans to expand the currently specified ARRA-funded activities.

Stevens asked what the impacts of the reporting requirements were. Dattoria said that there are unprecedented requirements placed on ARRA participants, including weekly reports and summary reviews of plans. This reporting function has taken a lot of manpower. These projects get a lot of visibility from this reporting. Zacharia added that there is a lot of reporting required, and it requires a lot of transparency and resources. He asked if these requirements would change the Department's culture in a positive way. Dattoria stated that these projects were required to be self-contained. Also, funds were not comingled with core-research funds. These requirements are changing the culture at DOE. Whether that change extends beyond the ARRA is unknown.

White asked what thinking went on behind the Magellan project decision. Dattoria said that it is an R&D project. It will have a definite end. Midrange computing is not part of ASCR's mission. It is not intended to roll it into the ASCR program. It is intended to give some information to people at the facilities. Strayer added that a workshop had been run on midrange-computing requirements. That workshop identified midrange cloud computing as requiring a bigger backbone than that supported by commercial cloud computing. Thus, this R&D project was formulated.

Giles asked if there were a role for ASCAC to reflect on this decision. Strayer responded that the requirements placed on ARRA funding precluded ASCAC review. As these projects progress, their goals, success, etc. could be the topic of ASCAC assessment.

Stevens said that this is a good list, given the limitations placed on the Office.

Giles opened the floor to public comment. There being none, the meeting was adjourned for the day at 4:19 p.m.

Wednesday, August 12, 2009
Morning Session

The meeting was reconvened at 8:32 a.m. Possible dates for future meetings were announced by Giles. **Dona Crawford** was asked to present the report of the Committee of Visitors (COV) to the Computer Sciences Division, which had been asked (1) to assess the efficacy and quality of the processes used for soliciting, reviewing, recommending, and documenting application and proposal actions and for monitoring awards and (2) to comment on how the award process has affected the breadth and depth of the portfolio and the national and international standing of the portfolio components.

The COV found that the solicitation and review processes appear to be effective and fairly administered. The program managers are dedicated and effective. However, the documentation of these processes and associated summary statistics are not very readily available and this impedes effective presentation of the competitive nature of the Computer Science research portfolio. There are one general and several topical solicitations each year. There were no observed inconsistencies in the application of procedures for reviewing the responses to these solicitations. The COV recommended that the program should automate the archiving of material related to each of its solicitations, including the call, letters of intent, full proposals, reviewer comments, and selection/rejection statements, in a single easily-accessible repository. This could be done ASCR-wide. The program should also collect and maintain statistics related to each of its calls in a consistent format that would facilitate analysis of the number of responses, reviewed proposals, and funded proposals. These statistics should be presented at an ASCAC meeting each year.

The COV found that the Computer Science program managers use generally effective mechanisms, including site visits, meetings and progress reports, to monitor ongoing awarded projects. Control is provided by annual financial decision points. Documentation of these visits and meetings does not exist. The COV recommended that the program should exploit ASCR's team approach with its planned increased staff to improve the frequency and depth of monitoring efforts.

The Division's portfolio has two main components: (1) making today's and tomorrow's leading-edge computers tools for science and (2) extracting scientific information from petascale experimental and computational data. The COV found that the Computer Science programs have broadly engaged the high-end-computing community. They have provided the depth necessary to facilitate research into effective use of leadership-class capability computing. There has not been a comparable engagement with the data-intensive aspects of the CS mission. However, initial efforts to recruit talent in data management are commendable. The COV recommended that the Division should launch strategic initiatives in all mission-relevant aspects of data-intensive computing, data management, and analysis. In addition, the entire Division should further engage the broader computer-science community in its reviews and workshops to define future research activities.

In terms of national and international standing, the COV found that the Division contributes to DOE's leadership role in capability computing and is internationally recognized for the impact of its research results. Its software libraries and tools are used worldwide. The Division's support for the International Exascale Software Project (IESP) demonstrates global leadership. The COV recommended that the Division should continue its leadership in high-end computing and expand its work on the data side and its collaborations with the international community.

In general, the COV found the number of permanent staff currently allocated to the Computer Science program to be insufficient to the task at hand. In recognition of this, SC has approved three additional staff. The COV recommended that ASCR fill the approved Computer Science vacancies as quickly as possible by working with Human Resources to streamline the hiring process. The Computer Science program is having a large impact on the DOE mission and on the broader scientific community.

Berger asked what the size of the program was. Strayer replied that, in FY09, it was $33.6 million. For FY10, $46.8 million has been requested. Berger asked how many projects it funded. Crawford replied that they could not tell precisely because of a lack of documentation. Johnson stated that there were about 100 awards in the CS program, but the number of projects is smaller because of partnering in the research community.

Voigt asked about the balance between high-risk research and tool development. Crawford responded that the COV could not tell because of a lack of single-source documentation. It had been told that the ratio was 80:20 in some calls and 50:50 in the overall portfolio, but that was a judgment call. Johnson added that one of the of success of R&D was the need for long-term maintenance, which puts pressure on the program staff.

Hey noted that that is one of the strengths of the DOE program, but it skews the balance between R&D and support.

Simon asked if there were any indication of the percentage of women principal investigators (PIs). Crawford replied that that was hard to assess; there were no pertinent statistics but cited a few key female PIs. There are fewer women going into mathematics and computer science. Giles pointed out that there may be data-collection constraints. Strayer affirmed that DOE is not allowed to keep such statistics, but there are a large number of women who have been outstanding in the field.

Sarkar asked if the COV were recommending something like NSF's system. Crawford replied that the COV is saying that a variety of information should be kept, not specifying how. It believes that letters of intent and other documents should be retained and that statistics on solicitations should be summarized.

Giles said that the 2004 COV pointed to a DocuShare system and asked if that system existed. Strayer replied, yes, it did. The Office has implemented all this in DocuShare and is now in the process of migrating to a different record-keeping system imposed from above.

Tang asked if there were any discussion of lessons learned from current activities in international collaboration and data analysis. Crawford replied, no. Tang said that this is tied to the networking area, which deals with highly integrated activities.

Giles called for a motion to accept the COV report. White moved to accept, and Berger seconded the motion. The question was called. The motion passed unanimously.

Edward Seidel was asked to speak on the transformation of science through cyberinfrastructure.

Two trends should be used as a guide: data-driven science is growing tremendously, and the size of the collaborations required to do this science is also growing. (The culture changed 20 or 30 years ago.) During that time, the amount of data has grown by 12 orders of magnitude. Mechanisms for sharing and using these massive amounts of data need to be developed. These large communities are producing demands on the infrastructure of science.

The technology crisis demands basic research in programming models, fault tolerance, etc. The data challenge is that more data are now developed in one year than in all of prior history. Something needs to be done in both hardware and software to respond to this crisis. Software outlasts hardware, and people outlast software.

The NSF published *Cyberinfrastructure for 21st Century Discovery*. The agency is trying to develop a national-level, integrated system of hardware, software, data resources, and services to enable new paradigms of science.

The TeraGrid has 11 resource providers. Within this set of services, NSF has been deploying Track 1 and Track 2 facilities. Four of these awards have been made or are in progress. The first of these was at the University of Texas, about 500 teraflops. The second was the Oak Ridge National Laboratory (ORNL)–University of Tennessee data initiative, which is a central model for such awards and efforts. The Blue Waters Project at the University of Illinois is based on IBM PERCS and is designed to achieve 1-petaflop/s sustained performance on real applications..

The TeraGrid, Phase III, is now up for proposals. Awards up to $150 million over 5 years will be made. That grid will connect the computing environment being set up by NSF. The vision is to have open data available for sharing, but the details raise problems, which will be resolved by a peer-review process. The users need to be supported, and the methodology needs to be built up.

Some questions that have to be dealt with are

- Where are my simulation data?
- How can I compare those data with experimental results?
- How do I visualize my results or the results of others?
- How do I see the simulation code that produced the data?

NSF is dealing with such questions with DataNet, which is being funded at $100 million over 5 years. The goals are (1) to catalyze development of multidisciplinary science and engineering data collections that are open, extensible and evolvable, and sustainable for more than 50 years and (2) to support development of a new generation of tools and services that facilitate data acquisition, data mining, integration, analysis, and visualization. Two awards have just been made to the University of New Mexico and Johns Hopkins University.

Another area of NSF investment is petascale software development (peta-apps). It started with a small amount of money and has grown into a $60 million program.

NSF's Office of Cyberinfrastructure (OCI) and Directorate for Mathematical and Physical Sciences (MPS) both support the Open Science Grid. The cyberinfrastructure blueprint integrates campuses and connects them to the national network.

Campus networks are typically bad. NSF is trying to catalyze the improvement of campus network infrastructures. DataNet is not the TeraGrid, but the two will be integrated in the future. Remote data producers [e.g., the Laser Interferometer Gravitational Wave Observatory (LIGO)] could be integrated and should work with the national programs to standardize software so that university-based researchers can make use of these remote facilities. Virtual organizations for solving very large problems are emerging and need to be served. These issues, the need to develop a new architecture, and the need to educate the community have been discussed. Software is the new language of science. We must be much more aggressive about training people to use this cyberinfrastructure environment, to deal with this data, and to work in this collaborative environment.

Computational science needs to find a home. Many components are being integrated into computer science. The Office of Cyberinfrastructure can take ownership and run programs for the entirety of NSF. Through grants, it can catalyze the changes that are needed in academia. Six task forces have been formed on campus bridging, software, educational workshops, grand-challenge virtual organizations, high-performance computing (including clouds and grids), and data and visualization. These task forces are working for 12 to 18 months and are (1) led by an NSF advisory committee on cyberinfrastructure, (2) holding workshops, and (3) developing programs. DOE participation would be appreciated in these workshops and in other task-force activities.

Other DOE–NSF cyberinfrastructure opportunities include software (already partnering on IESP, the exascale software activities), computing (the NSF/Oak Ridge partnership is an interesting experiment and the exascale is in the future), and applications (through peta-applications, future activities, and SciDAC).

In education and workforce development, NSF is doing a lot of things that it is not doing very well and could benefit from copying or partnering with DOE, other agencies, and international partners (in Europe, Asia, etc.) in cyberlearning, data, and networks.

Hey stated that this is exactly what is needed.

Manteuffel asked what could be done to encourage universities to establish departments in computational science. Seidel replied that NSF has a bully pulpit and, working with EDUCAUSE, it can influence universities. A small program at NSF on curriculum development is influencing undergraduate education. A computer sciences program might be made a prerequisite of grant awards.

Sarkar noted that basic research is needed to make data persistent. Seidel agreed that a whole system is needed to store, archive, check, migrate, curate, etc. data.

White stated that the success of leveraging small OCI investments across NSF is interesting and asked why or how has that investment worked across a broad agency. That is to say, how do you get program offices to pony up to support the initiative? Seidel replied that there is a growing awareness that the different directorates are spending too much money on software-development tasks and infrastructure that could be shared with others and that, if they did share, they could do more science. Also, in NSF, cyberinfrastructure reuse is set up to allow people to share software with other programs and projects. They cannot reinvent their own wheels all the time.

Chen asked where NSF analyzed and visualized its petascale data. Seidel replied that it used clusters of dedicated leadership-class machines. Some Executive Director awards are being made in visualization. Chen asked who was responsible for data preservation. Seidel responded that there is a whole spectrum of responsible parties, including publications. Agencies seem to want their own archives. The long-term financial implications have to be considered. Libraries are taking on some of these responsibilities (e.g., John's Hopkins is the curator for the Sloan Digital-Sky Project data).

Stevens asked what NSF's thoughts were on investing in hardware research. Seidel answered that there is a program on "Beyond Moore's Law." It is investing a small amount in that program.

Voigt asked what the right model was for software maintenance and for rewarding people who undertake this task. Seidel replied that NSF has a policy of sharing data, but enforcement is a problem. Sharing software is hard to do. Journals can have software environments incorporated into them. NSF has a build-and-test facility for some software environments. It is hoped that such a facility is implemented for all NSF grants.

Tang asked what the best path forward would be for partnership with DOE in graduate education. Seidel suggested a boot camp for supported graduate students at

DOE facilities. Such programs could be run at NSF's TeraGrid sites. Career awards might be another mechanism for supporting computer science.

Berger pointed out that there is also a need for software funding and asked how that could be addressed. Seidel suggested incentivizing universities to promote young faculty who do computer science. Centers might be another approach to provide the critical mass to develop and support software. Berger asked how one got that benefit back into the disciplines. Seidel replied, by pointing out to the directorates how cyberinfrastructure can revolutionize their programs.

Giles asked what mechanisms for collaboration could evolve and what role advisory committees could play. Seidel suggested having joint meetings of advisory committees. Michael Strayer and he meet periodically. There are also discussions at the program-manager level with monthly meetings planned.

The floor was opened to public comment. There being none, a break was declared at 10:13 a.m. The meeting was called back into session at 10:30 a.m. **Warren Washington** was asked to speak about the challenges in climate-change science and the role of computing at the extreme scale.

The challenge is to model the complexity that makes up the climate system, including sea ice, glaciers, ocean circulation, solar activity, and many other variables. To that end, a workshop was held to review and identify the critical scientific challenges, to prioritize the challenges in terms of decadal or annual timelines, to identify the challenges where computing at the extreme scale is critical for climate-change science success within the next two decades, to engage international scientific leaders in discussing opportunities to shape the nature of extreme-scale scientific computing., to provide the high-performance-computing community with an opportunity to understand the potential future needs of the climate-change-research community, and to look for breakthroughs in climate-change research.

There were 95 workshop attendees. A white paper was prepared in advance of the workshop to focus the discussion. Breakout panels were offered on

- Model development and integrated assessment of CO_2 emissions and energy usage;
- Algorithms and the computational environment;
- Data, visualization, and productivity; and
- Decadal predictability and prediction of climate.

The community is working on global models to capture small-scale features, such as hurricanes, and simulate how they interact with oceans

There were two previous reports

- *Identifying Outstanding Grand Challenges in Climate Change Research: Guiding DOE's Strategic Planning*, published by BER
- *Report on Computational and Information Technology Rate Limiters to the Advancement of Climate Change Science*, published by ASCR.

Priority research directions (PRDs) were established by each of the breakout sessions. For the model development and integrated assessment breakout session, the PRDs were: How do the carbon, methane, and nitrogen cycles interact with climate change? How will local and regional water, ice, and clouds change with global warming? How will the distribution of weather events, particularly extreme events, change with global warming? What are the future sea-level and ocean-circulation changes? In responding to the first question, the participants pointed out that a number of issues need to be resolved (e.g., cloud-resolving models will need to be developed and applied to improve traditional climate models used for climate projection. in the next 2 to 5 years). One thing recommended was the description of the important processes governing ice-sheet melt. This topic is now being addressed by a research team. Other recommendations were that important vertical mixing in the ocean needs to be more accurately represented and how mixing eddies and surface forcing combine to affect the stability and variability of the meridional overturning circulation needs to be determined. In regard to sea-level and ocean-circulation changes, this breakout session pointed out that sea-level rise will lead to mass migrations, social unrest, and national-security concerns.

In algorithms and the computational environment, SciDAC has supported performance improvement for climate-change models. Numerical algorithms still need to be developed to efficiently use upcoming petascale and exascale architectures; an international consortium should be formed to deal with the topics of parallel input/output, metadata, analysis, and modeling tools for regional and decadal multimodel ensembles; multicore and deep-memory languages should be developed to support parallel software infrastructure; and scientists should be trained in the use of high-performance computers.

PRDs for decadal predictability and prediction include identifying sources and mechanisms for potential decadal predictability and developing strategies for tapping into this predictability, ultimately realizing predictions that have societal benefit. The oceans have to be initialized, land vegetation and soil moisture have to be taken into consideration, etc. Substantial computing resources are required for decadal climate prediction because of resolution, data, complexity, and duration and/or ensemble size.

The recommendations for data visualization and computing productivity include developing new, robust techniques. The Earth-System Grid is very helpful in preparing the IPCC assessments. Each simulation for those assessments uses a petabyte of data, which is used by thousands of researchers around the world. Help is needed for those researchers to understand how to find and use those data. This application will grow in the future.

Crosscutting issues raised at the workshop include the need to educate the next generation of climate scientists in how to use extreme computing and to train current scientists in the use of high-performance computers. Computer architectures have become increasingly complex, so it is important to have machines that are easier to use. The ability to predict changes in land cover, vegetation types, oceanic biology, and atmospheric and oceanic chemistry needs to be improved, which will require dealing with people with whom climate scientists have not interacted in the past. Scalable algorithms that can use upcoming petascale and exascale architectures efficiently need to be developed. New, robust techniques must be developed to enhance the input/output, storage, processing, visualization, and wide-area transport demands of exascale data sets.

Since the mid-1960s, climate models have developed from (1) atmosphere/land-surface models and ocean models developed and run by individuals or small teams to (2) ensembles that also include sea ice, coupled climate, sulfate aerosols, the carbon cycle, dust/sea-spray/carbon aerosols, interactive vegetation, biogeochemical cycles, and ice sheets developed and run by large teams (up to several hundred people).

New developments that are anticipated include coupling to glacier models (the newest estimates of sea-level rise are 39 cm by year 2100); it would be desirable to couple the models to extremely high-resolution hydrological models to obtain future water management information; and full coupling to integrative assessment models thus include economics, carbon emissions, population, and climate-system change.

Hey asked if the changes Washington was talking about are going to change current climate-change predictions dramatically. Washington replied that the desire is to reduce uncertainty with this new generation of models. One variable (e.g., vegetative cover) at a time will be changed to see the feedbacks more clearly. Most of the biases in the models have been eliminated, but some issues (e.g., sea-level change) have not yet been sorted out.

Manteuffel pointed out that there have been significant advances in the past few years and asked how much of that advance has been due to mathematics, machines, etc. Washington answered that there are better models and more satellite data. There are problems with modeling convection, but progress has been made. The Atmospheric Radiation Measurement (ARM) Program's data allow one to validate model results. Better algorithms and much better gridding are being obtained from SciDAC for the models. NSF and DOE want to support scientists with new ideas. The availability of computer time is crucial to run models iteratively, changing one variable at a time. When one adds something to a model, the model looks worse for a while until one figures out how the new capability should fit to the data.

Tang asked what the apportioning was between near-term, mission-oriented activities and seed-plant activities. Washington responded that DOE and NSF have done this well. It is desirable to bring in a new generation of researchers with new ideas. At the same time, the simulations for the next IPCC assessment need to be done. It takes about a year of production mode to do those simulations.

Douglas Kothe was asked to describe some petascale early science at the ORNL LCF.

The goals of the early science period are to

- Deliver early breakthrough and impactful science,
- Harden the system for production, and
- Embrace a key-science user community.

The plan is to document lessons learned and science delivered, producing many benefits for the community, such as application readiness and hardening, OLCF system infrastructure testing and maturation, educating current computer and computational scientists prior to INCITE general availability, and educating OLCF staff prior to INCITE general availability. The plan is to use this time as a "science-acceptance" period.

This effort started in May 2008 when a call was issued. Plan-of-record system specifications were provided, and applications were sought that needed more than 25% of the resource for 1 to 4 weeks. The initial targeted call was e-mailed directly to about 125 scientists involved in DOE Programs [the Innovative and Novel Computational Impact on Theory and Experiment (INCITE), SciDAC, etc.]. The initial deadline of July 11, 2008, was extended to September of 2008. A broader open call was issued from October through December of 2008. The application form was available online at the National Center for Computational Sciences (NCCS) website. The Jaguar/XT5 acceptance occurred on Dec. 29, 2008; the Jaguar early science period started on January 12, 2009; and it ended on July 14, 2009. The requested proposals were very similar to INCITE applications. Thirty-three proposals were received. They were reviewed by subject-matter experts prior to submitting recommendations to ASCR (similar to the INCITE process, except streamlined). The philosophy was to bring as many projects on as possible, to fully utilize the system, and to open the door to as many interested users as possible.

Users asked for one billion hours, 650 million hours were available, and 500 to 540 hours were set aside as deliverable. Projects were brought in in four phases. At the time of the early science transition to operations, 27 projects were executing early access science at scale. The usage so far has been 334 million hours, and some “background” usage is still occurring. Scalability was pushed, and some applications used the full capability. There were a lot of fundamental-science and energy-science projects; materials science dominated.

About 25 codes got on the system in this early science period. Projects exercised the system to the fullest extent. There were three Gordon Bell finalists among the users. The Jaguar’s peak performance was 1.645 petaflops with a system memory of 362 TB, a disk space of 10.7 TB, and a disk bandwidth of 240 GB/sec.

One crucial aspect that was seen was the centralized, high-bandwidth (240 GB/sec) filesystem; the bandwidth was accepted in July. As more codes were brought online, there was a drop in peak usage, but the utilization was satisfactory: 66% of the jobs used more than 30,000 cores, and 25% used more than 90,000 cores.

GEOS5 weather forecasting resolution has gone from 28 to 3.5 km, which has turned out to be a quite interesting game changer, especially for hurricane prediction.

CM 2.5 simulations allow projections for decades rather than centuries by driving the resolution down by a factor of 4 and by showing agreement of winds with observations.

PFLOTRAN, which showed leakage of uranium into the Columbia River, got much more predictive.

The simulation of combustion is the first 3-D simulation to fully resolve flame and ignition features including chemical composition, temperature profile, and turbulence flow characteristics.

The XGC1 code became the world’s only whole-volume turbulence simulation in a realistic Tokamak geometry; it is now operating in a hybrid mode.

GTC, a fusion-energy code, saw its performance increase linearly with scale-up.

Simulations with the Gyrokinetic Tokamak Simulation (GTS) code confirmed the presence of electron-scale fluctuations associated with electron-temperature-gradient turbulence but suggest their contribution to observed anomalous electron transport is insignificant, which is consistent with preliminary experimental evidence.

In nuclear energy, the UNIC code has produced an explicit-geometry representation of the Zero Power Reactor (ZPR) core with 120 billion degrees of freedom.

LS3DF became the first ab initio calculation for a II-VI rod with realistic surface structure.

The OMEN code gave coherent transport simulations in band-to-band tunneling devices with simulation times of less than an hour, allowing one to rapidly explore the design space.

CP2K performed the first quantum-based simulation [using density functional theory (DFT)], probing both the thermodynamics and chemistry of ions at the aqueous air–water interface.

Direct numerical simulation (DNS) produced the world's largest and highest Reynolds number simulation of turbulent dispersion with accurate tracking of trajectories.

The Joule codes were put on the machine during the early-science period. Four FY09 Joule codes were exercised extensively. Raptor had an MPI halo communication redesign/refactor, VisIt had an MPI all-to-all communication redesign/refactor, CAM moved to a hybrid MPI/OpenMP node model, and XGC1 moved to a hybrid MPI/OpenMP node model. A Joule report is forthcoming.

Breakthrough science is being delivered on the Jaguar/XT5 leadership system at the OLCF. A special science-highlights document will be issued this fall that will provide ample evidence. There have been many publications and invited talks, and more are forthcoming. Yet, science demands and goals outstrip the current resources. Jaguar/XT5 today is stable, available, usable, and in high demand. The hardware and software environment is proven and hardened for production. Science applications are scaling well across the full system and exhibiting excellent performance. Leadership usage is already occurring. The most aggressive Joule exercise ever is successfully proving its worth. Early indications are that the petascale science era will more than deliver on expectations of the community and stakeholders.

Giles asked how the support teams were doing. Kothe responded that it depends on the application team and the code. The people at ANL are really participating in the team. Hundreds of clients cannot be supported. This model of support will always be needed, even as the applications and architectures change.

The floor was opened to public comment. Archambault asked how many applications were hybrid ones. Kothe replied three or four could handle open MPI. Archambault asked whether any applications go to full scale with MPI. Kothe replied, yes. Archambault asked if any applications use coarray Fortran. Kothe was not aware of any. Archambault asked whether Fortran was used. Kothe responded, yes, in more than half of the applications.

There being no further public comment, the meeting was adjourned at 12:04 p.m.

Respectfully submitted,
Frederick M. O'Hara, Jr.
Recording Secretary
September 8, 2009