Minutes
Advanced Scientific Computing Advisory Committee Meeting
Nov. 3–4, 2009, Oak Ridge Associated Universities, Oak Ridge, Tenn.

ASCAC members present:

F. Ronald Bailey
Marsha J. Berger
Jacqueline H. Chen
Roscoe C. Giles, Chair
Thomas A. Manteuffel
Linda Petzold
Vivek Sarkar
Horst D. Simon
Larry L. Smarr
Rick L. Stevens
William Tang
Victoria A. White
Thomas Zacharia (Tuesday only)

ASCAC members absent:

James J. Hack
Anthony J. G. Hey
John W. Negele
Robert G. Voigt

Also participating:

Arthur Bland, Project Director, Leadership Computing Facility, Oak Ridge National Laboratory
Peter Cummings, Principal Scientist, Center for Nanophase Materials Sciences, Oak Ridge National Laboratory
Vincent Dattoria, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Katherine Evans, Computational Earth Sciences Group, Oak Ridge National Laboratory
Al Geist, Group Leader, Computer Science Research Group, Oak Ridge National Laboratory
Barbara Helland, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Alexander Larzelere, Deputy Assistant Secretary for Fuel Cycle Management, Office of Nuclear Energy, USDOE
Walter Polansky (via Internet), Office of Advanced Scientific Computing Research, Office of Science, USDOE
Frederick O'Hara, ASCAC Recording Secretary
Robert Rosner (via telephone), Director, Argonne National Laboratory
Jeremy Smith, Center for Molecular Biophysics, Oak Ridge National Laboratory
Rachel Smith, Oak Ridge Institute for Science and Education
Michael R. Strayer, Associate Director, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Jeffrey Vetter, Computer Science and Mathematics Division, Oak Ridge National Laboratory

About 30 others were in attendance in the course of the two-day meeting.

Tuesday, November 3, 2009
Morning Session

Michael Strayer welcomed members to the meeting at 9:05 a.m. He introduced **Barbra Helland** and noted that she was acting as the Designated Federal Official pro Tem. He also introduced **Marsha Berger**, who began the meeting on behalf of Roscoe Giles, who was late in arriving because of airline delays. **Rachel Smith** made safety and convenience announcements. The meeting was webcast.

Michael Strayer was introduced to present an update on the activities of the Office of Advanced Scientific Computing Research (ASCR). He noted that the afternoon portion of the meeting would be a tour of the major computing and scientific research facilities of Oak Ridge National Laboratory (ORNL).

Five positions in the Office were being filled with new personnel. Cybersecurity has been moved from Next-Generation Networking to Applied Mathematics.

The 7% overall increase in the budget of the Office is robust. Most of that increase has been devoted to research programs in the Office. A new fellowship program in applied math and high-performance-computer science has been proposed. In computer science, there is a new effort in advanced computer architecture design for science applications. In computational partnerships, the Office is going to support interdisciplinary teams focused on transforming critical DOE applications for extreme-scale computing. In terms of facilities, the increases support lease payments and site preparation at Argonne National Laboratory (ANL) for its proposed upgrade. The Jaguar is rumored to be the fastest open-science computer in the world. The budgetary increases also support Energy Sciences Network's (ESnet's) beginning to deliver 100 to 400 Gbps to Office of Science (SC) laboratories.

A new research initiative, Ice Sheet Initiative for CLimate ExtremeS (ISICLES), enables a breakthrough in climate modeling. The Mathematics for Complex, Distributed, Interconnected Systems call for national laboratories only looks to research into the behavior of large-scale, complex, distributed, interconnected systems not as decentralized component models but as integrated entities; 38 proposals have been received and reviewed; five to seven awards are anticipated at $3.5 million per year for 3 years. The call for Mathematics for the Analysis of Petascale Data Research sought proposals for the mathematics of extracting features from extremely large datasets and understanding these data; 81 proposals were received and reviewed, and 11 projects were awarded at $4 million per year for 3 years.

ISICLES seeks research on community modeling, adaptive algorithms, associated uncertainty quantification, and simulation of ice-sheet fracture and dynamics. Accurate near-term simulation of ice sheets at 1-km resolution with refined uncertainty-quantification techniques and solution frameworks for extreme-scale ice-sheet simulations are expected.

The Joint Mathematics/ Computer Science (Math/CS) Institutes are being set up to foster collaborative research in applied mathematics and computer science to bridge the gap between large, complex, scientific applications software and next-generation hardware. ASCR is interested in better partial-differential-equation solvers; multi-precision and architecture-aware algorithms and libraries; software stacks to support scalable, resilient performance on multicore architecture; and mathematical modeling and optimization of the use of input/output resources.

For petascale data, the mathematical challenges involve extracting insights from extremely large datasets and investigating fundamental issues in finding key features and understanding the relationships between those features. Research topics include data analysis with uncertainty, rare-event detection, machine learning for massive data sets, scalable graph decompositions, geometric analysis for data reduction, and scalable statistical analysis. It is important to look for anomalies in large data sets.

The Mathematics for Complex, Distributed Systems initiative will look at scalable methods for representing, characterizing, and generating large graphs; agent-based integrated model for complex networks; intrusion detection for high-performance computing; stochastic modeling of complex networks; and risk management and planning of complex networks.

In terms of facilities, at the National Energy Research Scientific Computing Center (NERSC), the quad-core upgrade to Franklin was accepted June 17, 2009, and the NERSC-6 contract was awarded to Cray for at least one petaflop Cray XT5. The Leadership Computing Facilities (LCFs) at ANL and ORNL won mission-need approval in January 2009 for upgrading to tens of petaflops by 2013. Follow-on Lehman reviews were held at each LCF in July 2009, and operational assessment reviews have been completed. The Oak Ridge proposal introduces a hybrid architecture with new mathematics and computer-science challenges. New prototype machines for these architectures are expected. The architecture is well understood; it presents a software challenge. The Blue Gene/Q will be stood up at ANL. Discussions with IBM are being held on the acquisition.

ORNL's Cray XT5 was upgraded from 2.3-GHz quad-core processors to 2.6-GHz six-core processors with ARRA funds. This upgrade increases system peak performance to 2.3 petaflops and allocatable hours to 1.5 billion hours. It will be the world's most powerful workhorse for scientific computing. About 12 areas of science are now operating on the machine at scale as part of the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) portfolio.

The new NERSC-6 system is named Hopper. The Cray was selected competitively with application benchmarks from climate, chemistry, fusion, accelerator, astrophysics, quantum chromodynamics (QCD), and materials to receive the best application performance per dollar and the best sustained application performance per megawatt. Phase 1 will be a Cray XT5 with 668 nodes, 5,344 cores (2.4-GHz AMD Opterons), a 2-PB disk with a 25-GB/s bandwidth, and air cooling. Phase 2 will be a Cray system (not an XT5) operating at more than 1 Pflop/s, peak; with about 150,000 cores, 12 per chip; a 2-PB disk with an 80-GB/s bandwidth; and liquid cooling. Phase 1 will be completed this year, and Phase 2 in the fourth quarter of FY10.

The upgrade at ANL, Mira, enables key science impacts, like predicting abrupt regional climate change; designing safer, cost-effective nuclear power reactors; enhancing the extraction of biofuels from biomass; and in silico design of nanostructured storage systems. The Blue Gene series is amazing. Lawrence Livermore National Laboratory [LLNL, which is a National Nuclear Security Administration (NNSA) facility] and ANL tailored an R&D contract with IBM and other contractors to integrate the application needs with the architecture to produce an incredibly effective system. The

Mira Blue Gene/Q system will operate at 10 Pflop/s, peak; have about 800,000 cores, 16 per chip; have a 70-PB disk with a bandwidth of about 470 GB/s; will be power efficient; and will be water cooled. It will be completed in FY12. In the future, power needs will be a limiting design feature. Mira builds on ASCR/NNSA investment and the LLNL Blue Gene/Q Sequoia competitive bid procurement.

ASCR asked the community to assess its decadal accomplishments in science, applied mathematics, computer science, and NERSC operations. The Committee conducting this assessment is looking at recently developed topics that contain new ideas as well as mature areas that have had substantial recent activity and achievement. It will report out in the coming year.

Another committee is looking at computer-science accomplishments. It is to identify the major ASCR-funded accomplishments in computer science in the past 8 years, considering the impact and difficulty. A preliminary list has been assembled. A report is expected later in the year.

Discussions with the Office of Management and Budget (OMB) must be based on science accomplishments. Breakthrough accomplishments in computational science are being assessed by another committee. A preliminary list has been put together. It includes the first self-consistent simulation of baryon acoustic oscillations in the intergalactic medium; large-eddy simulation of combustion instabilities in a gas-turbine combustion chamber; adaptive mesh refinement (AMR) simulations of pellet injection in tokamaks; hybrid quantum simulations of biomolecules; solving the Continuous Electron Beam Accelerator Facility (CEBAF) beam breakup with the shape-uncertainty-quantification method; simulation of lean, premixed, hydrogen flames; unique high-fidelity calculations of direct-injection processes for internal-combustion-engine applications; and reactor-core hydrodynamics.

Smarr commented that these machines are enabling a wonderful breadth of science. The National Institutes of Health (NIH) were urged to access networking and high-end computing. By and large, that has not happened. There is a huge opportunity here. He asked if any discussions about partnerships were being held by DOE and NIH. Strayer replied that the previous administration had held high-level talks with NIH, but the principals have changed, and the talks have not restarted but will come. The facilities will continue to be made available under the INCITE program.

Manteuffel stated that the jump in applied-mathematics and computer-science funding is exciting and asked what was seen for the future. Strayer answered that the cost of keeping a research program going is difficult to maintain. However, at the extreme scale, software will have to be more fault tolerant and be more reliable. The changes in software, networking, and computer science will be enormous. The hard thing will be to bring onboard a family of young investigators.

Sarkar said that programmability of extreme-scale computing will also affect smaller-scale machines and asked if there are any initiatives organized around education and lessons learned. Strayer responded that the first steps toward extreme-scale computing are just now being made. There will be many more workshops and discussions needed. A broad-based, integrated community will have to be developed.

Chen observed that moving to 20-petaflop machines will be a huge software challenge and asked if Scientific Discovery Through Advanced Computing (SciDAC) would play a large role in that progress. Strayer said, yes. SciDAC has been very successful and will be one instrument for pulling these communities together. Ten years from now, SciDAC teams will look very different from those seen today, involving very large (500- to 1000-person) collaborations.

Bailey asked if any changes were seen in the direction of the Office under the new administration. Strayer said that there are changes. There are no INCITE policy changes; that openness policy has very strong support and is remarkably successful. There is a strong emphasis on integrating science into applied technology. Advanced Research Projects Agency–Energy (ARPA-e) will be another avenue for integrating science into applied technology. The President has underscored the need to address the issues of energy and environment.

Manteuffel asked what a Lehman review was. Strayer explained that Lehman reviews are very rigorous processes to keep large projects on track in finances and schedule.

Manteuffel asked what non-aqueous liquid coolant was being used in the systems. Bland replied, R134 Freon coolant.

Berger asked how the Committee was to keep up with the awarding of the institutes and other initiatives. Strayer answered that workshops are held for community members. The administration actions are reported on the website and by other public announcements. They often take a little longer than is wished because of the need to notify the awardees and to complete contract negotiations.

A break was declared at 10:18 a.m.

The meeting was called back into session at 10:31 a.m., and **Jeremy Smith** was asked to discuss high-performance-computer simulation in the energy biosciences. The Jaguar can now screen 120 million ligands to see how they would dock to an organism.

The Center for Molecular Biophysics is designed to overcome the recalcitrance of cellulose to be broken down into sugars. Lignocellulosic biomass exhibits structural complexity, including hemicellulose, lignin, and cellulose. These molecules are being modeled to better understand the recalcitrance at the molecular-mechanical and quantum-mechanical levels. The simulation for exploring the energy landscape is Herculean. One has to look at atomic-level interactions, macromolecular levels, and the supermacromolecular levels (many millions of atoms).

Scaling used to reach only to less than 100 nodes but has now surpassed 7000 nodes.

Electrostatics were smoothed out with shift truncation (which is not accurate) and particle-mesh Ewald (which is slow). It is believed that reaction-field methods are needed; they would be fast, but their accuracy is unknown. Early applications indicate that reaction-field methods give the same results as the other methods do. This process has been scaled to 100 million atoms. If one projects to the exascale, this process would be able to simulate one-tenth of a living cell for 10 μs.

There is an interest in merging high-performance computing with the Spallation Neutron Source (SNS) to get structures of 500,000 to 3.3 million atoms at 10 to 1000 Å. What is desired is to pull cellulosic fibers apart to convert them to sugars. Simulation models are helping us to understand how to pull them apart. To conduct such simulations, one has to know what these molecules look like in three dimensions. These simulations are calibrated to the neutron-scattering results from the SNS.

It turns out that lignans reprecipitate onto the cellulose fibers, blocking the conversion of the cellulose to sugars. It will take awhile to understand the meaning of these physics results.

The cell is hydrolyzed by enzymes. Through simulation, how these enzymes operate and how to design more-catalytic enzymes are being investigated.

Cellulosome assembly is mediated by cohesin-dockerin interaction, and that process is being investigated at the atomic level. Those investigations show that different strains are mechanically constructed differently.

Atomic-level findings are being mapped onto mesoscale course graining.

Bacteria (e.g., MerR) convert methyl mercury into elemental mercury. Modeling is capturing the mechanics behind this biological machine and bond breaking during the process through protonalysis. It is hoped to be able to develop synthetic catalysts to do these jobs.

Manteuffel asked where the mercury goes. Smith replied that it goes through bonding to a series of enzymes and eventually leaves the cell as elemental mercury.

Tang asked what level of cross-verification was performed with other simulations. Smith answered that results are compared with density functional theory (DFT) and other atomistic approaches. They are also compared with experimental results. The results are also tested experimentally. Tang asked if there were any examples where feedback had been given to the experimental community. Smith replied, yes: joint simulation/experimental data papers are published.

Stevens asked if small molecular inhibitions had been investigated in the recalcitrance problem. Smith answered, no; however, that would be a good idea.

Larry Smarr was asked to discuss ESnet support for cosmology's supercomputing, super network, and super visualization needs.

Cosmology is a data-intensive science, but the researcher in the laboratory is trying to move 100 TB over a shared network on which it takes 10 days to download 100 TB. Project Stargate was established to explore the use of explore the use of OptIPortals as petascale supercomputer "scalable workstations"; exploit dynamic 10-Gb/s circuits on ESnet; connect hardware resources at ORNL, ANL, and the San Diego Supercomputing Center (SDSC); and show that data need not be trapped by the network event horizon. From the end-user point of view, the results are the intellectual property of the investigator, not the center where it was computed. Petascale high-performance computing (HPC) machines are not ideal for analysis/vizualization. Centers do not take advantage of local cyberinfrastructure resources on campuses (e.g., Triton) or at other national facilities (e.g., ANL Eureka).

A 10-Gbps data path needed to be opened up from ORNL's National Institute for Computational Sciences (NICS) to ANL to the SDSC. One challenge was that the memory at Kraken was not connected to ESnet.

The amount of data moved from ORNL to ANL data-transfer nodes was about 148 TB in 577 time steps at a peak bandwidth of about 2.5 Gb/s, sustained. In addition, data were stored in the ORNL High-Performance Storage System (HPSS) and had to be staged to a disk on the data-transfer nodes. Once data are moved to the HPSS partition, they cannot be moved back again. Because each time step was a file in the tar file format, the data had to be untarred after transfer. To move forward, one will need a direct high-bandwidth path from the Kraken to the Eureka.

There are a huge pile of NVIDIA cards that make this visualization possible. 65 nodes of Eureka are needed to run the hardware-accelerated volume-rendering library.

With a data size of 4096^3, it takes 129 processor/graphic-card loads in 5 minutes and 10 seconds to render the image in 4.51 seconds. The loading is the bottleneck.

The next experiment that the project is going to do is to stream a 4000 × 2000 movie from an ANL storage device to an OptIPortable on the SC2009 show floor. For supersonic magneto hydrodynamic (MHD) turbulence simulations for star formation, it will try to use the Jaguar as the data generator rather than the Kraken. The question is, can DOE make this new mode available to other users so large-scale calculations can be done? If the end user cannot access and use the data, the computational job is not done.

Giles assumed the chair at 11:05 a.m.

Chen asked if one could use other strategies, such as moving smaller amounts of data to the rendering engine and multi-resolution visualization methods. Smarr replied that the solution will be to link the supercomputer to the rendering engine. One needs to be able to shift in and out to be able to query and study the visualization. The key is to let the end user drive the software, not the other way around. Global context is wanted at the same time as fine detail. This is a new way to get science into the brain of the investigator. Collaborative efforts are also desired to simultaneously view results around the world.

Simon asked what the top technological and sociological challenges were. Smarr answered that the last mile from the supercomputer to the user's desktop is the top technological challenge. The community is stovepiped. End-users must recognize a right to this capability, and DOE needs to recognize the need to provide technical support.

Katherine Evans was asked to discuss the ISICLES initiative, which started in September 2009 and deals with extreme ice-sheet modeling issues.

There is an urgent need for advanced dynamical ice-sheet modeling. Models used to date do not include uncertainties in the full effects of changes in ice-sheet flow. Models need to identify the extent and role of sea-level rise from sea-ice melting.

Unlike sea ice, when land ice melts, sea level rises. There is a large population within the zone where rapid sea level rise would submerge land. If the aggregate ice amount over land were to melt, the calculated sea-level rise would be about 65 m. Current Intergovernmental Panel on Climate Change (IPCC) predictions of ice-sheet loss provide no upper bound. The parameters involved are ice thickness, bedrock characteristics, ocean circulation and temperature, ice-stream dynamics, grounding lines, and precipitation.

Observationally, measurements of the rate of change of surface elevation for Antarctica and Greenland are median filtered at 10 km, gridded to 3 km with a mean time of about 2 years for 2003 to 2007 (with only three data points).

Now there are several efforts to model ice sheets on a continental scale:

- DOE's Community Climate System Model (CCSM): Glimmer through the Community Land Model (CLM)
- Parallel Ice-Sheet Model (PISM)
- Community Ice Sheet Model (CISM)
- SImulation COde for POLythermal Ice Sheets (SICOPOLIS) (Greve, Germany)

There are three typical ways to represent ice sheets: The thermomechanical shallow-ice approximation (SIA) assumes that bedrock and ice-surface slopes are sufficiently small. A higher-order model gives a consistent approximation to the full Stokes equations that minimizes the stress-strain functional. In Non-Newtonian Stokes flow, viscous forces dominate and are not time-dependent except to readjust to boundary conditions. The Cadillac is non-Newtonian Stokes flow.

Bottlenecks to progress in climate modeling investments by ASCR and the Office of Biological and Environmental Research (BER) were studied. ASCR needs to provide the tools for ice-sheet modeling.

Ice-sheet modeling is in its infancy and can therefore adopt ASCR tools rapidly.

ISICLES has funded six projects.

The goal of the SEA-CISM [Scalable, Efficient, and Accurate Community Ice-Sheet Model] is to provide a parallel, scalable capability as soon as possible to (1) allow high-resolution simulations with code extensions with reasonable throughput and accuracy; (2) maintain consistency and interaction with the production-level CCSM; and (3) enable seamless inclusion of incremental developments, such as new parameterizations and higher-order flow equations and coupled simulations with other climate components.

Glimmer is connected to the CCSM through the coupler to the CLM serial, coarse-grid, SIA-based modular open source code. It computes the ice sheet surface mass balance on the coarse 100-km grid. Results are downscaled to the finer 10-km ice-sheet grid. Previous versions of CCSM have used a static representation of the Greenland and Antarctic Ice Sheets. Many extensions are planned to increase model realism and complexity in various stages of implementation. The climate community needs constant access to a basic CISM, with the ability to test and post model improvements. The goals are to develop parallel capability and a fully implicit solution method. Ice sheet modeling is going to undergo significant growth of complexity in the short term. Algorithm design must account for increased coupling and multiscale behavior.

Glimmer is built and is operating on the Jaguar; performance testing is ongoing. Currently, almost all work is in the solver. Two new tests using a new physics of ice sheets are now available. There is a Fortran interface to Trilinos hooks. It uses 1.5 million nodes, and each iteration takes 1 to 5 minutes, with an iteration count in the hundreds.

The Berkeley-ISICLES has adaptive gridding in regions with higher velocities that incorporate AMR because ice velocities toward the centers of ice sheets are much slower than near the edges. Much higher resolution (1 km vs 5 km) is required in regions of high velocity.

They will need to (1) improve the performance of high-resolution ice-sheet modeling because of the increase in problem size; (2) attain high performance via parallel computing, algorithmic improvements, and autotuning; (3) implement Glimmer-CISM in the Chombo framework with structured-grid AMRs; and (4) apply autotuning to improve the performance of computational kernels.

The deliverables are (1) completing basic algorithm and software design and implementing basic solver components for an ice-sheet model in the Chombo framework as independent software components, including testing and verification; (2) applying autotuning to key computational kernels in the existing Glimmer-CISM code; (3) investigating the impact of linear-equation solvers on the performance of Glimmer-CISM; (4) prototyping and validating AMR-based code; (5) investigating performance optimization of the AMR code with autotuning techniques; and (6) performing detailed algorithmic and software improvements.

The project for modeling the fracture of ice sheets on parallel computers seeks to examine, learn, and evaluate existing ice-sheet models as a base-code development platform and to seek new partners in the ice and climate communities. The problem is to define a geometry; set boundary conditions and loads; generate a mesh; model continuum damage mechanics (crack initiation and propagation); elasticity; produce an extended finite-element model (XFEM) for crack modeling; and develop specialized highly parallel multigrid solvers for the XFEM.

XFEM provides modeling flexibility and can be used to (1) predict fracture and collapse of ice sheets, (2) predict ice calving, and (3) explain accelerated ice-sheet flow. The drawback of XFEM is that it adds degrees of freedom, and accurate modeling will quickly result in billions of unknowns. An efficient parallel solver is needed.

The Lagrangian Model for Ice Sheet Dynamics has the objectives of (1) developing a 3-D Lagrangian particle model for ice-sheet dynamics and implementing it on leadership class computers; (2) developing highly scalable, meshless algorithms that are based on smoothed particle hydrodynamics; and (3) using the 3-D model to investigate assumptions in simplified but computationally more efficient ice sheet models for different types of ice sheets and glaciers.

Grid-based solutions of 3-D free-surface problems are very complex and are rarely sought in ice-sheet models. Most of ice-sheet models use (quasi) 2-D first-order SIAs and shallow-shelf approximations (SSAs). Under certain conditions, these approximations may lead to significant errors. Full 3-D solutions are also important for accurate simulations of tidewater glaciers, ice shelves, ice streams, surge dynamics, the influence of ice-shelf back-pressure on inland ice flow, the dynamics of flow across the grounding line, and the dynamics in the vicinity of ice-sheet divides. Lagrangian particle methods are very efficient for free-surface problems and for problems involving large material deformation and fracturing.

The advantages of smoothed particle hydrodynamics (SPH) are that

- Pure advection is treated exactly.
- Interface problems in free-surface flow simulations are trivial for SPH but difficult for grid-based schemes.
- Particle methods bridge the gap between the continuum and fragmentation in a natural way.
- There is a close similarity between SPH and molecular dynamics (MD), so complex physics can be included relatively easily.
- Highly scalable algorithms are employed.

This method has been applied to flow in flooding, lava flow, material fracturing, as well as sea ice.

The proposed research will improve the predictive ability of the climate. Because of the novelty, the model will create a new user base within the scientific community and become widely recognized as a unique and valuable capability. The research team has a vast expertise in the particle methods, and this makes the proposed research highly feasible.

The University of Texas project will address mathematical and computational challenges in large-scale ice-sheet-dynamics modeling. Some of the challenges are a wide range of spatial scales, from 10^6-m continental scale to 10^2-m scale of flow transitions; severe ill-conditioning of linearized systems because of about 5 orders of magnitude contrast in ice viscosity; severe nonlinearities caused by complex ice rheology; unknown ice parameters and basal boundary conditions require the solution of an ill-posed inverse problem; and the sparse and noisy data require a statistical approach to the inverse solution.

The project will build on a base of models, algorithms, and software for global mantle convection scaled to more than 60,000 cores. The research group has previously demonstrated strong scaling of an AMR library on problem sizes up to 2.24 billion elements on 62,000 cores with excellent scalability. Weak scaling of the AMR library with 131,000 elements/core indicates excellent parallel efficiency.

A Bayesian framework provides a solution of the inverse problem expressed as a probability-density function (PDF). The method of choice is to sample this PDF with Markov-chain Monte Carlo (MCMC) analysis. For inverse problems with expensive forward simulations, contemporary black-box MCMC methods become prohibitive. The goal is to develop methods that exploit the structure of the parameter-to-observable map, as has already been done successfully.

ANL's SISYPHUS proposes to develop techniques for solving the fully 3-D Stokes problem for continent-scale ice sheets integrated over hundreds or thousands of years on petascale computers. They are developing more accurate, high-performing ice-sheet modeling methods and a framework for constructing the models, connecting them to solvers, and coupling them to regional and global climate models. They are using mesh-motion and enthalpy-transport equations. They are using an unstructured, hexahedral grid, a discrete, mesh-based geometry for the bed, adaptive mesh near the coastline and bed, a method-of-lines modeling approach, a dual-order preconditioning scheme over space, and a solver approach that is also applicable to the adjoint from the start. They are going to use component-based solvers and tools to solve a challenging physics problem at scale with a higher-level interface to Portable, Extensible Toolkit for Scientific Computation (PETSc) for expressing physics and physics-based preconditioners in component form, which separates the overall solution strategy from the specifics of the physics models, allowing variations on either side of that interface. Their petascale integrated tools build on component services that are unified by common interfaces.

Some projects have an overlap of tasks. Projects range from short-term deliverables to longer-term impacts. Their mesh approaches vary. All are using iterative numerical methods, where appropriate. Interactions with climate scientists provide a link to the models and modelers.

The goals of the ISICLES projects are to address the importance and complexity of ice-sheet predictability, to leverage the computational science tools developed through related ASCR efforts, and to provide petascale-ready simulation capabilities for the ice-sheet modeling community.

Manteuffel asked whether solvers are the bottlenecks in several projects. Evans replied, yes.

Giles asked the Committee to study the two charge letters and additional information before the reading of those letters on the second day of the meeting

The floor was opened to public comment. There being none, the meeting was adjourned for lunch at 12:00 noon.

During the afternoon session, the Committee toured the Center for Nanophase Materials Science, Spallation Neutron Source, High-Flux Isotope Reactor, BioEnergy Science Center, and Leadership Computing Facility at ORNL.

Wednesday, November 4, 2009
Morning Session

The meeting was reconvened at 8:31 a.m. Giles welcomed **Peter Cummings** to speak on solving nanoscience problems with high-performance computing.

Theory is important because one cannot characterize everything experimentally. Nanoscience has come to the forefront because of the emergence of experimental tools. At short length and time scales, there are incoherent neutron scattering (INS), quasi-elastic incoherent neutron scattering (QINS), second-harmonic generation (SHG), sum frequency generation (SFG), and x-ray reflectivity, *inter alia*, to validate computational results. There are molecular-interaction techniques, coarse-grain techniques of particle-based methods, and *ab initio* molecular dynamics techniques to study reactions and interfaces. All of these tools are also used in computational nanoscience. An understanding of these processes has to be taken to the macro scale if one is going to control manufacturing processes and make anything. The ORNL Center for Nanophase Materials Sciences (CNMS) was colocated with the SNS, High-Flux Isotope Reactor (HFIR), and National Center for Computational Sciences (NCCS).

Theoretical and computational nanoscience have played, and continue to play, a central role in nanoscience. Theory, modeling, and simulation (TMS) play a greater role in nanoscience than in macroscopic materials and chemical sciences. Many experiments performed at the nanoscale can only be interpreted through theory. TMS can provide a convenient framework to isolate effects and phenomena in a way that may be difficult or impossible to achieve in an experiment. In a theoretical treatment, one can control the boundary and initial conditions, which may be impossible to achieve in an experiment. TMS is crucial in understanding emergent phenomena in nanoscale systems. TMS can be used to design new nanostructured materials, as well as systems based on nanoscale phenomena. By design and from the beginning, theoretical and computational nanoscience have played a central role at the CNMS. The Center supports a large number of user projects (about 25% of the total) and fosters a strong interaction with computational scientists.

There is a strong interconnection and collaboration between CNMS and the Computer Science and Mathematics Division (CSMD). There is a budget for CNMS to fund joint activities with CSMD and a staff for a computational endstation for nanoscience, affording unique capabilities (e.g., the Vienna *Ab-Initio* Simulation Package (VASP) code, which is efficient on petascale platforms]. The 2008 Gordon Bell Prize was won by the CNMS/CSMD project Multiteraflop Simulations of Disorder Effects on the Transition Temperature of the High-Tc Superconducting Cuprates. This collaboration has been a win-win for CNMS and CSMD.

Some of the science done at the center includes pulling monatomic gold wires with single molecules; determining the stability of rutile and anatase, two phases of the same compound (titanium dioxide); determining the properties (proton transfer) on water at titanium oxide surfaces; identifying the processes of libration, bending, hydroxyl stretch, and free hydroxyl formation; and investigating the electrical double layer, which is important in many phenomena. In the last experiment, the water was found to dissociate at the titanium dioxide surface. A multiplicity of techniques can be used to look at the adsorption of various species to the surface of the titanium dioxide. Proton relaxation times from quasielastic neutron spectrometry (QENS) and MD can show the amount of adsorbed water and the formation of three distinct layers by water adsorbed from the air.

The role of TMS in nanoscience includes the interpretation of complex phenomena in an experiment and the design of new nanostructured materials and systems based on emergent phenomena at the nanoscale. The theoretical limits to nanoscale manipulation are being pushed by recent theoretical developments, Evan's fluctuation theorem, and the quantification of the probability that a system will exhibit negative entropy production. This was experimentally verified by manipulation of a colloidal particle by optical tweezers and subjected to additional verification through the response of an electrical circuit to thermal noise. This study is an example of computational discovery; the control of boundary conditions enabled the isolation of phenomena.

Performing 540 simulations of pulling a gold wire to rupture allowed the theoretical description of force, temperature, etc. interdependencies.

In conclusion, TMS play vital roles in nanoscale science and engineering in the interpretation of experiments, design of experiments, characterization and design of nanostructured materials, and design and control of manufacture. TMS in nanoscale science and engineering typically require many different techniques. Future advances in the field will result from the development of additional methods, such as multiscale methods, electron transport dynamics, optical properties, self-validating forcefields, etc. Real advances in computational nanoscience need strong computational science/nanoscience integration. Experience has shown that the dividends paid far outweigh the investment.

Smarr asked about nanoengineering and whether nanoelectronics was something to consider. Cummings replied that the Center is thinking about that, particularly on the theoretical side rather than the engineering side. In replacing photolithography with self-assembly, one might consider how to deal with faults. Molecular electronics will also have to deal with noise.

Chen asked if there were any connectivity between *ab initio* and course-grained techniques. Cummings responded that one usually does upscaling as many-scaling, not multi-scaling. Hybrid schemes are also being looked at. It is planned to extend these hybrid schemes to additional scales.

Alexander Larzelere was asked to discuss the Nuclear Energy Advanced Modeling and Simulation (NEAMS) Program that is being conducted jointly by the Office of Nuclear Energy (NE) and SC.

Nuclear energy offers a lot of advantages. Nuclear-energy technology can decrease costs, improve performance, increase the pace of deployment, enhance innovation, responsively deal with nuclear waste, and promote nonproliferation. There are 104 reactors in the United States today. By doing power upgrades through modeling and simulation, 6 GW have been added to the fleet's capability. The field needs to go beyond the traditional test-based approach in understanding nuclear energy. When the current fleet was designed, the peak computing power was in the kiloflop per second range. Today, that peak computing power is in the petaflop per second range. When the current fleet was designed, scientific understanding was gained by fusing theory and experiment. Today, modeling and simulation adds something to the equation that was not there before.

In the understanding of complex physical processes, one might have well-understood initial conditions and well-characterized effects but not have any insight into the actual physical processes because the conditions are too small, too hazardous, too long, too far away, too expensive, too complex, lacking the requisite facilities, or not allowed by policy. Advanced modeling and simulation supplements theory and experiment to explain "how" things happen; it does not replace the need for theory or experiments.

A workshop was held on science-based nuclear-energy systems enabled by advanced modeling and simulation at the extreme scale. It was organized into four panels:

- integrated performance and safety codes (IPSC);
- advanced materials;
- verification, validation, and uncertainty quantification; and
- systems integration.

The IPSC panel was concerned with fluid/structure interactions; cross-section methods, covariance, and usage; 3-D fuel failure; evolution of pin and assembly failure; full-scale-plant radiation-field modeling; reactive flows; up-scaling microstructures to macro structures; and upscaling bench-scale technologies to the plant scale. Modeling and simulation can contribute in the areas of Monte Carlo simulation, 3-D fuel design, 3-D pin and assembly design, 3-D reactor design, and 3-D plant design.

Materials are important in this area, so the workshop considered material behavior in extreme environments, which must include verification, validation [together, V&V], and uncertainty quantification (UQ). First, one must focus on research into the critical issues and challenges and conduct a study using the V&V and UQ process to analyze a number of critical, integrated physics applications that would provide a problem focus and address the issues of coupled multiscale physics and UQ. The workshop recommended an open-source, flexible, and extensible energy-enterprise model. That is probably the way that the field will go forward.

There have been several calls by experts for a nuclear-energy advanced modeling and simulation effort: In 2003, a Massachusetts Institute of Technology (MIT) report called on DOE to establish a major project for the modeling, analysis, and simulation of commercial nuclear-energy systems. In a 2009 update to that report, the MIT panel said that the DOE program has moved in this direction but much remains to be done. In a 2008 *Bulletin of the Atomic Scientists* article, Robert Rosner said that high-fidelity (science-based) integrated simulations must form the core of design efforts, allowing for rapid prototyping that minimizes the need to experiment. A 2008 Nuclear Energy Advisory Committee (NEAC) report stated that an advanced modeling and simulation effort can lead to better understanding of nuclear energy systems and has the potential to resolve longstanding uncertainties associated with the deployment of nuclear systems.

The NEAMS initiative wants to rapidly create and deploy "science-based" verified and validated modeling and simulation capabilities essential for the design, implementation, and operation of future nuclear energy systems with the goal of improving U.S. energy security. Today, the research community is dealing with low dimensionality, test-based physical behaviors, low resolution, uncoupled systems, and workstation computing. In 10 years, NEAMS hopes to go to 3-D, 2-D, and 1-D models of science-based physical behaviors with high resolution for integrated systems with advanced computing. Modeling and simulation tools need to be built to do the design of the next generation of nuclear power reactors. These capabilities will have to be flexible so they can be applied to different types of nuclear-energy technologies.

The effort has to build on what others have done. Important lessons were learned from the Advanced Simulation and Computing Initiative (ASCI):

- Have a clear and compelling vision of the mission and develop a comprehensive program to create new capabilities.
- Headquarters need a "team of rivals" at the national laboratories for leadership of the program.
- Success requires the best from universities, industry, and national laboratories in the form of partnerships.
- Accomplishing the ambitious goals will take time and funding, but it must deliver increasing capabilities early and often.

NEAMS users are the research and development community, technology designers, regulators, and utilities and operators.

NEAMS is divided into two channels: (1) integrated performance and safety codes for nuclear fuels, reactors, safeguards and separations, and waste forms and (2) supporting program elements in the form of fundamental methods and models, V&V and UQ, capability transfer, and enabling computational technologies. For each of these elements, NEAMS needs central points of contact.

For both the current fleet of nuclear power plants and for a wide variety of potential reactor designs, one needs to predict the performance and safety of reactors over a 40- to 60-year lifetime.

One has to analyze fuels, which requires the development of a multiscale, multiphysics framework with appropriate scale-bridging techniques.

Safeguards and separations are important for fuel recycling. 95% of the energy content is still in the spent fuel. One can separate the good stuff from the bad stuff, process the bad materials so they are not so bad, dispose of the bad material, and reuse the good material. That path forward presupposes a waste repository, which would require one to predict the performance of waste forms under repository conditions for their expected lifetimes. In addition, one needs to develop cross-cutting models and methods for lower-length-scale materials performance. Methods need to be developed for V&V and UQ, and analyses need to be conducted of existing methods to select the best-in-class to serve as a basis. NEAMS will serve as the single point of contact for experimental data, so requirements for the experimental data that is needed to support those methods must be developed, and an understanding must be gained about the availability of existing data.

NEAMS has other components: Capability Transfer is the need to get new knowledge to the engineering level. Enabling Computational Technologies (ECT) includes software quality, application-development tools, problem setup tools (e.g., mesh generation), numerical libraries (e.g., solvers), and results analysis (e.g., visualization).

NEAMS has assembled the "A" Team of national laboratories, universities, and industry. Last year was the first year of significant funding, and that funding is growing.

Within DOE, NE is forming partnerships with the offices of SC, Environmental Management (EM), NNSA, and Fossil Energy (FE). It also has international interfaces with France for using the Salome code of Électricité de France (EDF) to create functional specifications and for collaborating on the open-source Saturn computational fluid dynamics (CFD) code, with Russia on sodium-fast-reactor design, and with Japan on seismic modeling.

NEAMS has a management structure, but it has not established an outside review panel or a national laboratory steering committee.

In FY09, NE received $20 million for advanced modeling and simulation as part of the Advanced Fuel Cycle Initiative and about $9 million as part of the Generation-IV (Gen-IV) Nuclear Energy Systems Initiative. In the FY10 budget appropriation, there is $19 million for modeling and simulation in the Fuel Cycle R&D budget and $5 million in the Gen-IV R&D.

DOE proposed to Congress that eight energy-innovation hubs be established. Only three were funded, including one on modeling and simulation for nuclear fuel cycles and systems. NE can benefit from modeling and simulation in conducting power uprates, achieving higher burnups, deploying Gen-III+ reactors, and dealing with long-term problems. Currently, in the science domain, there are a few users, few jobs, very big computers per job, and long runtimes. Engineering analysis is different; it uses short, high-pressure timelines; requires many short jobs; and has many users. What is needed is to adapt science-domain capabilities to address issues for the near, medium, and long term in the engineering domain. If one does that, one changes the game by accelerating the use of advanced modeling and simulation to address near-term nuclear-energy issues.

The hub should not operate in the abstract. It should be focused on a specific operating reactor (a virtual reactor). The hub's activities should be distributed. Ideally, it should have one physical location competitively awarded to a team of national laboratories, universities, and/or industries and have one colocated team of 75 to 100 multidisciplinary researchers. People will have access to (but not own) high-performance computing resources. The hub will deliver, in 5 years, an engineering environment that uses advanced modeling and simulation to address near-term nuclear-energy issues.

A competition will be conducted. Factors that will be considered include the proposed virtual reactor problem, the approach, and the team's qualifications. Site visits will be conducted, and the award will be made in June 2010.

The hub will be short-term focused. It will be funded for 5 years with an additional 5-year option. NEAMS will continue beyond the hub.

Berger asked what one has to do to upgrade an existing reactor. Larzelere replied that the analysis of 40 years of operations and the shrinking of the uncertainty allowed an increase in the flux of the reactor. Berger asked whether a success story like this helped. Larzelere answered, yes, it does, but it takes time. People are coming to the point that they appreciate the value of modeling and simulation.

Smarr noted that, at the Kyoto science and technology meeting, Japan, South Africa, India, and France expressed an interest in developing a new generation of reactors, mostly breeder reactors with plutonium recycling. He asked whether the United States has turned its back on that path. Larzelere responded that this administration does not see plutonium to be a problem that needs to be solved in the short run. It is a problem that has to be addressed in the long term. There is plenty of uranium in hand. That is not the case for France and Japan. China is standing up a sodium-cooled fast reactor and is building a mix of a large number of reactors.

Bailey said that, in 1915, the National Advisory Committee for Aeronautics (NACA) was formed to bring together aircraft designers, builders, and users. He asked why userengineers are not included in the NEAMS. Larzelere replied that they tend not to be interested in longer-term ventures. They will come around. DOE will not build a reactor. Industry needs to use the tools DOE develops. Talks are being held with GE–Hitachi, Westinghouse, Areva, etc.

Tang noted that, for validation, the United States is facilities poor and asked how it could bring foreigners to the table. Larzelere said, by modeling and simulation to analyze their data. It will take money and time to do it. Validating data will largely come from overseas, but there is a lot of historical U.S. data, also.

A break was declared at 10:16 a.m. The meeting was called back into session at 10:30 a.m. Rosner joined the meeting via telephone.

Michael Strayer was asked to present the new charges to the Committee from the Director of SC. There were two charges and a letter from Under Secretary Steven Koonin. The pertinent passages of the first charge letter are

"To help the research communities utilize the capabilities of current and future supercomputers, ASCR also supports a basic research program in Applied Mathematics. To ensure the integrity of this research program, I am asking the (ASCAC) to assemble a COV [Committee of Visitors] to review the management processes for the Applied Mathematics elements of the ASCR program. A report will be expected at the August 2010 ASCAC meeting.

"The COV should provide an assessment of the processes used to solicit, review, recommend, and document proposal actions and monitor active projects and programs. The Committee should assess the operations of the Applied Mathematics programs during the fiscal years 2007, 2008, and 2009. The panel may examine any files from this period for both DOE laboratory projects and university projects. The Committee will be provided with background material on the program prior to the meeting.

"I would like the Committee to consider and provide evaluation of the following two major program elements:

1. For both the DOE laboratory projects and the university projects, assess the efficacy and quality of the processes used to:
   (a) solicit, review, recommend, and document proposal actions, and
   (b) monitor active projects and programs.
2. Within the boundaries defined by DOE missions and available funding, comment on how the award process has affected:
   (a) the breadth and depth of portfolio elements, and
   (b) the national and international standing of the program with regard to other applied mathematics research programs that are also focused on the demands of high performance scientific computing and analysis of petascale datasets."

Strayer underscored that applied mathematics is at a point in its development that is critical to the Department. He read the second charge letter, the salient portions of which are

"Over the last few years, several workshops and subcommittee reports have identified and described the scientific opportunities for high performance computing. By this letter I am charging the Advanced Scientific Computing Advisory Committee (ASCAC) to assemble a subcommittee to look at the results of these activities and to analyze the opportunities and challenges for the Office of Advanced Scientific Computing Research

(ASCR) and the Office of Science associated with exascale computing. Specifically, I would like the subcommittee to deliver a report that:

- Assesses the opportunities and challenges of exascale computing for the advancement of science, technology, and Office of Science missions.
- Identifies strategies that ASCR can use to address the challenges and deliver on such opportunities.

"We would appreciate the committee's preliminary comments by March 30, 2010, and a final report by August 15, 2010. I appreciate ASCAC's willingness to undertake this important activity."

In addition to the charge letters, the Committee received a letter from Under Secretary for Science Koonin, to wit:

"As your subcommittee begins to work on Dr. Brinkman's new charge, I would like you to consider:

"Modeling and simulation at the extreme scale have the potential to span the entire Department and to forge lasting links between applied and basic science, technology and engineering. I would therefore ask the subcommittee to

- acknowledge and comment on the broader issues of opportunities and challenges for exascale computing to advance Department of Energy missions.
- include on the subcommittee members familiar with both NNSA and various applied programs that are amenable to extreme scale computing

"We live in times when the Nation is faced with important issues involving energy, environment, and national security, yet resources are constrained. It is therefore very important to have a clear justification for how new endeavors in extreme scale computing, ultimately reaching the exascale, will affect science, technology and society."

Giles announced that he had identified appropriate leadership for both of these subcommittees. Robert Rosner has agreed to chair the exascale subcommittee. Giles is in discussions with Linda Petzold to chair the COV. He initiated a discussion of the charges and what they cover.

Berger noted that the applied mathematics COV, Section 2b, covers only high-performance computing and petascale data sets. She asked why other aspects were not covered. Strayer responded that those should be considered as only examples. A clarification should be sent back to the Director's office. It should also cover multiscale and other topics.

Giles pointed out that an interim report should be presented at the next ASCAC meeting. The other charge is unusual in that it has been commented on by an under secretary. It was hoped to narrow the focus of the report and still to address the broad concerns of Under Secretary Koonin. The Subcommittee might develop some goalposts and milestones for the exascale initiative.

Manteuffel stated that this charge is very broad and would take years to complete. He asked how it could be narrowed down. Giles pointed out that a lot of work has been done on the exascale in workshops, by a cross-cutting panel, in Defense Advanced Research Projects Agency (DARPA) studies, and at town meetings.

Sarkar said that extreme scale refers to 1000+-core chips and asked if the focus here were all on the exascale or on better scalability. It does not mean that today's software could run on exascale machines. He asked if the Subcommittee were to look at the underlying implications for extreme-scale computing. Strayer replied that the charge letter asks about the challenges and opportunities of exascale computing.

Bailey raised another question on whether the application of exascale computing is to DOE mission needs.

White asked whether it would be unreasonable to expect that other parts of SC be made aware of ASCAC's charge and mission. Giles said that the Subcommittee will not interact with other SC offices without the Secretary's consent and instruction. That will need more discussion within DOE.

Stevens stated that the charge was broad but clear. How well the Department's plan for exascale computing maps onto the opportunities it offers is one question that the Subcommittee could ask. It could ask whether the strategy were reasonable. Giles pointed out that one could always ask the Designated Federal Officer (DFO) for clarification.

Manteuffel asked what Rosner thought about these comments. Rosner said that, with all the components of getting to the exascale [software systems; middleware; applied mathematics (algorithms); applications; and validation], it is not clear which is/are within the purview of the Subcommittee. Strayer offered that all of the above were included. Giles noted that the Subcommittee needs to work fast and to focus down to meet the interim-report deadline.

Strayer noted that the Office does not have an exascale initiative, yet. There have been a lot of community workshops. A strategy is needed on how to go forward. ASCAC has the opportunity to influence the makeup of that strategy.

Giles stated that this is an important responsibility. Committee members should make suggestions about elements to be discussed.

Tang asked what size of panel was desired and how this fitted into the international context. Giles replied that a size has not yet been set.

Smarr noted that talented people have worked on prior reports. In 2000, papers predicted what the petascale would look like; but in the end, commodity chipsets were used. The exascale will require breakthroughs. One group should look at the path forward for commercial technology, and another should look at revolutionary technologies.

Giles noted that one thrust is to look for the path to the exascale. Stevens agreed. Three petascale strategies were put forward, and all three were built. It was the software people who were wrong. It has to be an evolutionary development. He suggested that the Subcommittee include some people who had gone through this prior experience and had learned its lessons.

Simon said that the systems that achieved a petaflop were driven by targeted investments by DOE. That is a model for the exascale. The Subcommittee should be relatively small and made up of people with a lot of time to read all the material available.

Berger commented that many projects on Jaguar take a year to ramp up to using the machine well with a lot of software technical support.

Smarr stated that DOE has been instrumental in working with industry and supporting mathematics and computer science. The National Science Foundation (NSF) missed an opportunity to fund software development.

Rosner said that the effort required to make the March deadline will be large, and the Subcommittee will need to make sure it gets people who have the available time.

Vincent Dattoria was asked to present an update on ASCR's ARRA investments.

ASCR's ARRA projects total \$153.9 million and include the Advanced Networking Initiative (ANI; \$66.8 million), LCF upgrades (\$19.9 million), Advanced Computer Architectures (\$5.2 million), Magellan (cloud computing; \$32.8 million), and SciDAC-e (\$29.2 million). SciDAC-e supplements and leverages existing SciDAC investments to advance the high-performance computational capabilities of the BES Energy Frontier Research Centers (EFRCs); to provide extra user support for energy-related projects at the Leadership Computing and NERSC facilities; and to conduct applied-mathematics research in support of DOE's electricity-grid efforts.

The ASCR Magellan Project at NERSC and the Argonne LCF is about a 100-TF/s compute cloud testbed (across sites) and petabyte-scale storage cloud testbed. Funding was distributed to ANL and LBNL on the basis of a peer-reviewed proposal. ANL and LBNL have issued contracts to procure compute and the first stage of data hardware. Coordination with the ANI has begun. The first availability of cycles is expected in January 2010. A joint Magellan–ANI principal-investigator meeting is scheduled at Supercomputing-09. Cloud questions to explore on Magellan include: Can a cloud serve DOE's midrange computing needs? What part of the workload can be served on a cloud? What features (hardware and software) are needed of a "Science Cloud"? How does this differ, if at all, from commercial clouds?

The ANI addresses the DOE Science Network challenges: a 72% annual growth in traffic since 1990 (compared to 45% at AT&T), an accelerating rather than plateauing growth, and scientific collaborations' becoming more international. The ANI will evaluate transport technologies for optical fiber backbone, deploy a prototype 100-Gbps-capable network, and develop and support an experimental network environment allowing researchers to test new concepts in networking. The ANI will be going to the Bay Area, Denver, Chicago, the Nashville ring, and New York City.

Advanced Architectures is a new effort to provide early access to DOE researchers of technologies emerging from the IBM Productive, Easy-to-use, Reliable Computing System (PERCS) effort. Two proposals were received and peer reviewed. Funding is in place at ORNL for the IBM PERCS effort, and grant paperwork for the Berkeley Research Accelerator for Multiple Processors (RAMP) effort has been submitted to Chicago for processing.

SciDAC-e has three components:

1. Research grants and national laboratory projects to develop mathematical techniques and algorithms to enable a bigger, better, and smarter electric grid (\$8.3 million);
2. Approximately 30 new postdoctoral appointments at ASCR facilities to offer assistance to SciDAC-e projects and other energy users (\$10 million; see below); and
3. Supplemental awards to existing SciDAC efforts to support BES EFRCs to develop a high-performance computing capability relevant to the goals of the EFRCs (\$10.86 million).

Seven SciDAC-e Applied Mathematics projects support the development of a smart grid.

The SciDAC-e Argonne Computational Postdoctoral Fellows program is based on Argonne's existing Director's Postdoc Program. There is a 10-member search committee that requires a curriculum vitae, a 2-page research proposal, and a list of publications and significant presentations. $3.125 million became available in September 2009. Eighty applicants have submitted proposals, four have interviewed, and one has accepted an extended offer.

At NERSC, which received $3.125 million, most of the postdocs will be directly embedded in a specific application area while also spending approximately half time at the NERSC facility to directly interact with NERSC staff to ensure there is communication of new ideas and "what works." One postdoc has already been selected and is working.

The Oak Ridge Leadership Computing Facility (OLCF) received $3.75 million in stimulus money to support approximately ten SciDAC-e postdocs and to host a summer institute. No postdocs are on board, yet.

Smarr, addressing the last-mile problem, asked whether people specify their computer's bandwidth when they apply for INCITE. Dattoria replied, no; however their work is being done at the centers. Smarr said that he would like to see an end-to-end analysis of bandwidth. Strayer said that that is an excellent idea and should be a standard procedure. Dattoria said that the Joint Engineering Team (JET) will analyze end-to-end cross-domain bandwidth. Smarr replied that it has to extend into the laboratory where the user sits. Dattoria agreed. Smarr said that the 100 INCITE applications would provide a good topology of available networks. Companies should be made aware of the need to make investments.

Chen asked what the formal process was for SciDAC award. Dattoria said that there is not one right now; it is in process.

White pointed out that, to be a Tier-2 center for Compact Muon Spectrometer–Large Hadron Collider data, a university has to show how it is going to handle the data.

Vetter asked if all the ANL testbeds were occupied. Dattoria said that some have been awarded, and there will be follow-on awards. ANL and NERSC will also look at how outside users will use the cloud computing.

Al Geist was asked to comment on the Extreme-Scale Algorithms and Software Institute (EASI) Joint Math/CS Institute.

The petascale was achieved by going to multi-core chips and dual-socket machines. To provide the computational resources required to tackle critical national problems, it is expected that systems will be heterogeneous with nodes composed of many-core graphics processing units (GPUs) and central processing units (CPUs). Some challenges include

- Scalability (nodes/cores/threads)
- Resilience (if nothing is done, a mean time between interruptions of tens of minutes will occur)
- Power consumption (if nothing is done, power consumption will exceed 140 MW)
- Programming environment (data movement and heterogeneous architectures will drive new paradigms)

Two complementary ASCR Math/CS projects are paving the way to the exascale. They share a common goal: closing the application-architecture performance gap. The Institute for Advanced Architectures and Algorithms (IAA) Algorithms Project is focused on homogeneous multicore systems and extreme-scale system simulations. The EASI Joint Math/CS Institute is focused on heterogeneous systems with accelerators and application resilience.

Both projects share a common approach to success: (1) An integrated team of mathematicians, computer scientists, and application experts is working together to create new architecture-aware algorithms and associated runtime to enable many science applications to better exploit the architectural features of DOE's petascale systems. (2) Application team members immediately incorporate new algorithms providing a near-term high impact on science. And (3) numerical libraries are used to disseminate the new algorithms to the wider community, providing broader and longer-term impact.

The Institute for Advanced Architectures and Algorithms (IAA) was begun in FY09 as Joint effort between Sandia National Laboratories and ORNL. It was to foster the integrated co-design of architectures and algorithms to enable more efficient and timely solutions to mission-critical problems and to influence vendor roadmaps through partnership and joint research and development. The IAA Algorithms Project is funded through this Institute. IAA revolves around the science and includes algorithms, runtime, architecture, and simulation.

Multiprecision algorithms have advantages: they double the bandwidth to the socket, double the cache size, and double the peak flop rate.

The project team, a mix of mathematicians, computer scientists, and application experts, is serving the High-Order Method Modeling Environment (HOMME), Multiresolution Adaptive Numerical Environment for Scientific Simulation (MADNESS), and Charon applications. It also has cross-site teaming to address runtime and affinity.

The project has developed a new algorithm for climate application that provides a 20-times faster solution; has incorporated a parallel-in-time algorithm in several applications; and is overcoming key message-passing-interface (MPI) limitations on multicore processors.

The new Institute has an EASI project team on architecture-aware algorithms for scalable performance and resilience on heterogeneous architectures. It will deliver codes to the community through the Scalable Linear Algebra Package (ScaLAPACK), Trilinos, Open MPI, MPI Chameleon (MPICH2), MADNESS, and HOMME.

EASI has developed a heterogeneous programming application-program interface (API) for multicore CPUs and GPUs. It allows writing portable parallel linear algebra software that can use pthreads, OpenMP, CUDA, or Intel Threading Building Blocks (TBB) (even more than one of these within the same executable). The API is extensible to other programming models, as needed. The API has been used to demonstrate compiling and running the same software kernel with pthread, Intel TBB, and compute unified device architecture (CUDA). The Trilinos Tpetra and Kokkos packages will incorporate this API in Trilinos 10.0.

Stevens asked if the abstractions generated for Trilinos were usable in other applications. Geist said that an effort was being made to make it as transparent as possible, and it is hoped to make them able to be pulled out of Trilinos and used in other codes.

White asked if the Institute had communicated with members of the QCD community who have been doing similar things for years. Geist replied, no, but it is an interesting idea.

Petzold stated that it would make a huge impact on small systems. Geist said that that was what was hoped for.

Manteuffel asked whether it wasn't a big challenge to allow people to put these algorithms right into their applications. Geist responded that that has been done by getting the main people for each application involved, by putting them right on the team. They then put it right inside the source stream.

The floor was opened for public comment. There being none, the meeting was adjourned at 12:04 p.m.

Respectfully submitted,
Frederick M. O'Hara, Jr.
Recording Secretary
November 23, 2009