Minutes

Advanced Scientific Computing Advisory Committee Meeting

November 6-7, 2007, Hilton Washington D.C. North, Gaithersburg, Maryland

ASCAC members present:
Robert G. Voigt, Co-Chair
F. Ronald Bailey
Gordon Bell
Marsha Berger
David J. Galas
Roscoe C. Giles
James J. Hack
Thomas A. Manteuffel
Horst D. Simon
Ellen B. Stechel
Rick L. Stevens
Virginia Torczon
Thomas Zacharia

ASCAC member absent:
Jill P. Dahlburg, Chair

Also participating:
Melea Baker, Office of Advanced Scientific Computing Research (ASCR), Office of Science (SC), U.S. Department of Energy (DOE)
David Brown, Lawrence Livermore National Laboratory
Lisa Carroll, ASCAC Recording Secretary
Christine Chalk, Designated Federal Official for ASCAC
Phil Colella, Lawrence Berkeley National Laboratory
Giulia Galli, University of California, Davis
Fred Johnson, ASCR, SC, DOE
David Keyes, Columbia University
Douglas Kothe, Oak Ridge National Laboratory
Dimitri Kusnezov, National Nuclear Security Administration, USDOE
Rachel Smith, Oak Ridge Institute for Science and Education
Walter M. Polansky, ASCR, SC, DOE
Michael R. Strayer, Director, ASCR, SC, DOE
Homer Walker, ASCR, SC, DOE

About 25 others attended all or part of the meeting.

Tuesday, November 6, 2007

Preliminaries

Robert Voigt, ASCAC co-chair, called the meeting to order at 9:00 a.m. He welcomed new member Marsha Berger to ASCAC.

Michael Strayer: Advanced Scientific Computing Research (ASCR) Overview

Recent Office of Science (SC) appointments include George Malosh as Deputy Director for Field Operations, Patricia Dehmer as Deputy Director for Science Programs, and John Alleva as acting Deputy Director for Resource Management. These three will assure continuity of operations when the administration changes in 2009. Within ASCR Fred Johnson is acting Director of Computational Science Research and Partnerships.

SC's partnership with the National Nuclear Security Administration (NNSA) resulted in acquisition of the Cray XT3 computer by Oak Ridge National Laboratory (ORNL), development of the next generation of IBM's Blue Gene supercomputer at Lawrence Livermore National Laboratory (LLNL) and Argonne National Laboratory (ANL), and support of research in turbulence, materials, astrophysics, nuclear structure, and petascale computing.

The fiscal year 2007 (FY07) appropriation is $283.4 million compared to the President's request of $318.7 million. The President's request for ASCR for Fiscal Year 2008 is $340.2 million, which is also the level in the House Bill but the Senate Committee on Appropriations recommended $334.9 million (the net of an increase of $7.7 million for LCF and reallocation of $13 million from SC to NNSA's Advanced Computing and Simulation Program). Additionally, the 2008 Defense Authorization Bill requires an independent evaluation of the NNSA plan for advanced computing, including assessment of the NNSA support of high performance computing (HPC) research and development (R&D).

In FY07 ASCR held several workshops with partners in computational science. The 156 registrants at the September workshop on computation research needs in alternative and renewable energy found a need for better methodologies and computational tools in the priority research directions of hydrogen, biomass conversion, solar energy conversion, wind energy, and grid futures and reliability. For example, advances in cellulose conversion require simulation codes to describe plant cell wall architecture, and wind energy research needs accurate models of wind turbine aerodynamics.

Carbon sequestration and storage were addressed in two workshops, held in collaboration with the Office of Fossil Energy (FE), and two more are planned for FY08. The high cost of prototypes makes models and simulations important in sequestration research.

A recent DOE Senior Leadership Summit on Cyber Security launched a comprehensive planning effort within the Department to define the threat, the values to protect, the protective systems, and ways in which to strengthen DOE's defenses. SC was charged with the preparing a cyber security R&D plan, which must cover both classified and unclassified systems across the DOE community and build on similar work by other agencies. The size, scope, and range of its research portfolio make DOE the logical leader of this effort. Moreover, SC manages advanced mathematical and computational research programs that can generate transformational improvements to cyber security.

In 2007, ASCR celebrated the successes of it's researchers such as David Keyes, who received the Sidney Fernbach Award. Computers at ORNL and Lawrence Berkeley National Laboratory (LBNL) provided more than half of the simulation data for the DOE/National Science Foundation contribution to the Intergovernmental Panel on Climate Change (IPCC) Fourth Assessment Report for which IPCC shared the 2007 Nobel Peace Prize with Al Gore.

Discussion: In response to a question about the cyber security plan, Strayer said that all analyses and plans will roll up to the Secretary's office. Asked which independent body would conduct the evaluation of NNSA's computing activities, he replied that SC had offered ASCAC but a DOE advisory committee may not be seen as independent enough for the work.

After a break beginning at 9:38 a.m., the meeting resumed at 9:55 a.m.

Walter Polansky: Scientific Discovery through Advanced Computing (SciDAC) Update

The ASCAC Committee of Visitors (COV) was charged to review the SciDAC efforts within the Office of Advanced Scientific Computing Research. The review focused on the process that led to the recent SciDAC awards. The COV conducted their review July 17-18, 2007, at the DOE facility in Germantown, MD. Overall, the COV found the resultant SciDAC portfolio to be strong and well positioned to address program goals. The proposal selection process was successful despite severe time constraints and the program's scope and complexity. Finally, the COV observed that coordination among programs offices in SC, NNSA, and the National Science Foundation was remarkable throughout the SciDAC recompetition. The COV made 19 recommendations in six general areas: call for proposals, review of proposals, cross-cutting review, selection, documentation of awards and declinations, and management of awards.

ASCR responded to all recommendations. ASCR understood the COV's rationale for recommending an annual review of all SciDAC projects. Although an annual review is not feasible, annual reports are required from all projects and reviewed by program managers who assess performance and progress toward goals. Any issues are then discussed in the coordinating group. For 2008, ASCR may review the investment in the Centers for Enabling Technologies (CETs), institutes, and Science Application Partnerships—collectively about 75% of SciDAC funding, to assure that this key component of the overall portfolio is positioned to meet SciDAC goals.

The FY08 SciDAC request is $50 million, but spending under the current CR is at $42 million. If the FY08 funding bill is not signed in mid-November, ASCR may need to repeat no-fund extensions of many SciDAC awards, perhaps extending 21 university awards for up to 90 days.

The pilot SciDAC Outreach Center provides a central resource for inquiries and technical information about the program. The Center has had several initiatives underway: dialog with Computational Sciences Graduate Fellows on their computing needs and resources, organizing a meeting of SciDAC PIs from the institutes and CETs, and expanding SciDAC's central web hosting and collaborative development services. The Center is also developing industry collaborations and working with the Council on Competitiveness on to enhance its industrial outreach.

SciDAC CETs and institutes are exchanging information and beginning to coordinate efforts. At the SciDAC PI meeting in 2007 meeting, the CETs and Institutes reported on efforts to upgrade their websites. The CETs are addressing the need for responsible software engineering practices in HPC and for maintaining the balance between support of current software and new development. Institute websites are being redesigned to reach broader audiences and bring new participants to SciDAC. The curved mesh correction tool developed by the Interoperable Technologies for Advanced Petascale Simulations (ITAPS) Center is a good example of the contribution the centers are making to research, as well as advancing the SciDAC mission.

Discussion: As a member of the SciDAC COV, ASCAC member Robert Voigt was pleased that ASCR is contemplating a cross-cutting review of the CETs and Institutes in FY08.

Giulia Galli: Quantum Simulations of Materials and Nanostructures (Q-SIMAN)

This research attempts to predict and design molecular and materials properties using the fundamental laws of quantum mechanics, with no input from experimental data, and to do so with controllable accuracy. Q-SIMAN uses high-performance computers and codes such as

QBox and QMCPack to describe electron-electron and electron-ion interactions in molecules, condensed systems, and nanostructures. Recent results obtained using quantum simulations, such as those for heat transfer in nanostructures and optical properties of hydrogenated diamond clusters, directly compare to experimental data, demonstrating the value of simulations to materials science.

Other Q-SIMAN simulations model the properties of water in different environments, including water confined in carbon nanotubes (CNT). Model predictions, such as the decrease in dipole moment of water molecules at the CNT surface, which leads to enhanced lateral diffusion and faster reorientation dynamics, are consistent with experimental observations of the rapid flow of water in CNTs. Understanding the properties of liquid water in confined media is useful to fields ranging from biochemistry to corrosion inhibition.

Q-SIMAN models may also help to answer outstanding questions about phase transition in elements and “simple” compounds. A long-term goal is “quantum engineering” in which complex simulations of integrated materials with specific functions become part of an integrated materials science. Current issues for Q-SIMAN include validation and verification (V&V), to address issues of simulation replication and analysis.

Discussion: The effort to quantify uncertainty in predictions remains an important part of the V&V work in Q-SIMAN. Other challenges include such problems as the orthagonalization of wave functions, and the choice of techniques for scaling up.

The meeting adjourned for lunch at 11:50 a.m. and resumed at 1:07 p.m.

David Keyes: Fusion Simulation Project (FSP) Workshop Report

Keyes provided an overview of the Fusion Simulation Project and the report of the May 2007 FSP workshop. FSP will develop predictive capability for integrated, realistic modeling of current and planned toroidal fusion devices and create high-performance software for these simulations. It combines resources of Office of Fusion Energy (FES) programs with ASCR’s resources in applied math and computer science, software development, and HPC. FSP is poised to use near-term petascale computing capabilities to support burning plasma simulation capability at ITER (c. 2012) and to prepare for ITER’s experiments (c. 2017). FSP’s success will promote U.S. leadership in ITER and confer a competitive advantage in fusion energy R&D. The ultimate goal is a commercial energy source that is essentially inexhaustible, carbon-neutral, and free of proliferation risks.

The report identifies scientific and technical challenges that can only be addressed by integrated modeling and the research areas where more investment may yield the tools to achieve these models. It also lays out a plan for improving the fidelity of the models’ physics components. Integration of many codes and models from various plasma science disciplines will require petascale computing.

Analyzing current gaps, the report defines critical scientific and technical issues for which computation is required; the state of the art for each issue and the current deficit in capability; and new capabilities needed to produce simulations to address the critical issues. ASCR’s contribution of tools and expertise are essential to closing these gaps. The FSP applied math and computer science agenda leads ASCR’s 10-year vision statement, and the list of ASCR research topics valuable to FSP is extensive in both fields.

ASCR expertise also supports validation and verification. If simulations are used in planning ITER experiments, there must be confidence in their accuracy, and greater rigor and formality in V&V are among FSP's 5-year deliverables. The FSP report addresses other technical deliverables for 5, 10, and 15 years.

In June, the Fusion Energy Sciences Advisory Committee (FESAC) reviewed the report. , FESAC gave conditional "yeses" to the five charge questions, noting the need for more specificity, but commended the collaboration by the FES and ASCR communities and asserted that the report made a case for pursuing simulation in conjunction with theory and experiment to benefit plasma science. FESAC recommended that FSP move to the project definition phase.

Concluding with a personal view, Keyes reckoned that scaling fusion simulations up to levels needed for ITER, performed naively with uniformly refined, non-fitted meshes and explicit algorithms, would require a capability of 10**24 floating point operations, 12 orders of magnitude over what is required to simulate today's tokamaks, such as CDX-U. In order to achieve this within a decade, algorithmic improvements, including geometric and solution-driven mesh adaptivity and implicit solvers, must be employed to bring the level down to about 10**15. Three areas with potential to close the gap are in the mainstream of SciDAC research.

Discussion: The ambitious goals of the FSP report sparked concerns about the feasibility of closing the gap between needed and current capabilities in the time available. The race to do so may also exacerbate tension between international collaboration and competition. Comments on the relative importance of qualitative and quantitative results and experimental validation highlighted the need for better communication between the simulation and experimental design communities. Working in isolation could lead to "doing things right," but not "doing the right thing." Time pressure and the extreme scaling issues also raise the risk that FSP could proceed with inadequate codes.

Questions about budget constraints and the level of national commitment to fusion energy research led Michael Strayer to observe that ITER is one of the administration's highest science priorities. In DOE's 20-year facility plan, ITER is the top priority.

Michael Strayer: New Charge to ASCAC

Pursuant to the FSP Workshop Report and the FESAC review, DOE's Under Secretary for Science, Raymond Orbach, charged ASCAC to consider what is being proposed in the report, focusing on the most critical challenges in applied mathematics, computer science, and computational science. ASCAC should then recommend an appropriate, mutually beneficial role for ASCR in FSP. How can ASCR best work to make FSP successful? The preliminary report is due at the February 2008 ASCAC meeting, and the final report at the August 2008 ASCAC meeting. Robert Voigt said that Ronald Bailey has agreed to chair the effort.

Homer Walker: Applied Math Program Update

ASCR supports research in such areas as partial and ordinary differential equations (PDEs and ODEs), dynamical systems, linear and nonlinear solvers, optimization, meshing, computational fluid dynamics, and multiscale mathematics (MSM). A highlighted research achievement is the work of Paul Fischer of ANL, who used Nek5000, a spectral-element multigrid code with a parallel coarse-grid solver, to simulate magnetorotational instabilities, thermal hydraulics in reactors, and turbulent flow in blood vessels.

The current FY08 plan is based on $22.4 million, which does not include $5 million for the Computational Science Graduate Fellowships or awards delayed by the continuing resolution. It allocates 81% of the funds to the national laboratories; however, this is likely to be closer to the final FY07 fraction of 70% when the budget is settled and all awards are made.

About 25% of the funding is currently allocated to MSM research and education. Currently, 13 projects at 28 institutions have MSM funds; lmost of all of these projects end in August 2008. Notable MSM accomplishments include the work of Russel Caflisch and Richard Wang of the University of California at Los Angeles (UCLA) with Andris Dimits and Bruce Cohen of LLNL on accelerated hybrid methods that effectively address multiple time and space scales in simulating "near-continuum" plasmas, in which some plasma subspecies can be treated as fluids while others must be treated as aggregations of particles.

The program plans to announce new awards in MSM and in a new area, optimization of complex systems, such as power grids and the nuclear fuel lifecycle, with multiple components, complex connectivity, and hard-to-predict behavior. Participants in a December 2006 workshop identified research likely to generate new insights needed to optimize these systems.

Discussion: Walker assured a questioner that the FY08 budget request was adequate both to continue multiscale math and to fund complex systems optimization research.

The committee took a break at 2:37 p.m. and resumed at 3:05 p.m.

Ellen Stechel: Role and Efficiency of Networking and Networking Research, Preliminary Report of the Subcommittee

The subcommittee is charged with examining networks and has focused on the implication of network technologies and advanced networking research strategies for the DOE science community. The group concurred with recent reviews of ESnet and chose to define the basis for innovation rather than prescribing specific changes in networking.

The subcommittee sees modeling, simulation, and analysis as major drivers for advanced cyber resources. Increasing amounts of data collected and archived, the diversity and variety of data, and the use of electronic data in new ways of doing science mark a radical change in data use and distribution. Cyber systems transcend traditional boundaries (time, location, discipline) and demand a "system of systems" approach to infrastructure. New network technologies can and should transform the way science is done.

The subcommittee had 7 findings: (1) DOE must continue supporting ESnet and (2) cannot depend on industry for its network technologies, which should grow out of DOE's leadership in developing petascale networking. (3) Systems-level science will require new network architectures and service models. (4) "Cyber environments" must be managed so that end-users can readily refine them to meet their needs. (5) Data-intensive systems may require a petabit per second application in 10 years. (6) New network research needs help in continuing through development to prototypes and production. (7) Domain-specific advances in networking must be leveraged so that they can serve other fields.

The subcommittee made 6 recommendations: (1) SC should apply the approach used for ultrascale computing to advanced networking. (2) ASCR should pursue the implications of viewing DOE science, facilities, and programs as a "system of systems." (3) ASCR should convene an external review committee for its networking research. (4) ASCR needs to form, fund, and implement a plan to keep new concepts moving. (5) In its network research program, ASCR should explore automated, dynamic, intelligent agents to take advantage of network architecture. (6) The network research program should also include an integrated data management component.

Discussion: Responding to a question about the cost of greater bandwidth, Stechel acknowledged that the subcommittee was not assessing economic feasibility. Questioned about the state of networking research, Stechel and subcommittee member Ronald Bailey said they had addressed networking services, rather than research per se, and that improving the quality of service applications is an active area. Thomas Zacharia noted the effects of increasing cyber security requirements on network applications. The subcommittee is identifying research for improving service quality and volume, but did not re-examine topics already covered in other, recent reviews. The committee agreed that there was no need to duplicate other work, but that it was important to define the scope of the report carefully.

Rick Stevens: Update on Joint ASCAC-BERAC Panel on Modeling for GTL

Stevens and John Wooley of the Biological and Environmental Research Advisory Committee (BERAC) co-chaired a panel charged with examining the status and feasibility of ASCR's 2015 goal for computational modeling in genomes to life (GTL) research. The panel is revising its findings and recommendations, taking into account two recent National Research Council (NRC) reports that describe how modeling and simulation are advancing biological research at levels from molecules to populations and ecosystems.

Models and simulation have become closely coupled with experimental biological research, offering new insights and course corrections. They play a crucial role in integrating understanding of biological mechanisms in areas important to DOE's mission in energy and the environment. Big advances in model scope and application are imminent. However, no long-term program of appropriate scale currently offers a single source of support for developing these capabilities. ASCR does not currently support applied math or computer science projects that primarily target developing integrated modeling and simulation capabilities for microbes or plants. More support is needed to curate data for organisms important to DOE's mission since modeling capabilities also depend on genomic and other databases. Also, for these organisms there are few integrated models covering multiple components. Obstacles to a predictive model for GTL include the lack of (1) integrated genomics databases, (2) robust mathematical frameworks and software to integrate metabolic and gene regulation models, (3) MSM and related codes and tools for integrating processes at disparate scales, and (4) a computational and analytical theory incorporating evolution in the interpretation of analytical results.

In light of these findings, the panel recommends modifying the 10-year goal of ASCR/BER joint modeling and simulation activity to demonstrate capability in predicting phenotype from genome "through advances in genome sequence annotation, whole genome scale modeling, and simulation- and integrated model-driven experimentation." DOE should also establish an explicit program for research in predictive biological modeling and hold an annual conference tracking progress. ASCR and BER should jointly support curation and integration of genomics and related data, especially for organisms of interest to energy and environmental research. Finally, DOE should work with the larger scientific community on new approaches to multiscale coupling of models and simulations (e.g., microbial systems to ocean ecology to global climate).

Discussion: Questions about metrics and funding led to a discussion of barriers to implementing the panel's recommendations such as the large investment needed, format differences in the databases and datasets of interest, the diversity of the biological research community, and the absence of current incentives for sharing data.

Dimitri Kusnezov: The New Frontier of Computational Science: Applications to High Leverage Decisions

Rapid advances in computational science raise expectations for answers to big scientific and policy questions, but expectations may be far ahead of current capabilities. Even so, computational science, supercomputing, and simulation are being applied to important decisions where little may be known about how well the predictions can be trusted. However, the structures for this work often do not encourage valuable cross-disciplinary collaboration.

Because testing is no longer done, managing the nuclear stockpile to keep it safe, secure, and reliable is data poor. As decades-old weapons are handled, questions arise about their condition, ranging from as-built problems to aging effects to needed upgrades. With no current experimental data and a dwindling number of scientists experienced with past testing, can computational science be pushed to give definitive, timely answers to these questions?

NNSA's Advanced Simulation and Computing Program (ASC) offers an environment where computational science can address problems with complex processes and dynamics. Like ASCR, ASC is looking toward the new exaflop frontier and has devised a strategy for HPC investment and allocation with advanced systems used for advanced science, capability systems to run long-term large integrated weapons simulations, and capacity systems for less demanding, routine simulations. The strategy includes bulk buying of capacity. ASC saved $26 million this year in its contract to provide 3 labs with 8 Linux clusters.

ASC also recognizes that, to do leading-edge simulations and meet broader scientific needs, it must work with other DOE offices, other agencies, other nations, industries, and universities. For example, the recent build-out of Blue Gene left Blue Gene L unclassified for university users. NNSA's collaboration with SC is strong and includes co-funding work on Blue Gene P and Q as well as SciDAC-2.

Collaboration is also needed in quantifying uncertainty, a major barrier to trustworthiness of predictive simulations. Related efforts are also needed in verification and validation (V&V). To stimulate research in these areas, the new Predictive Science Academic Alliance Program (PSAAP) focuses on multi-scale, multi-disciplinary unclassified applications. ASC expects that its leadership in computational science will pay off across research fields and in industry because the skills and knowledge readily transfer to research in other large systems.

Discussion: Answering a question about the shape of possible V&V collaborative efforts, Kusnezov said that a plan is in development but acknowledged the difficulties of joining data across multiple scales and the current lack of clarity as to where research is most needed.

Public Comment

Phil Colella, an FSP committee member, contrasted FSP workshop report's emphasis on integrating existing codes with Keyes' view of the large gap between current and needed capability. How, he asked, can program planning and project management assure these advances are available when they are needed? High-risk investment and speculative research are contrary to the way things are organized. Rick Stevens said that the FESAC panel agreed that marshalling resources and disciplines would require creative program management.

The meeting adjourned for the day at 5:30 pm.

Wednesday, November 7, 2007

Preliminaries

Robert Voigt called the meeting to order at 9:08 a.m.

James Hack: Computational and Informational Technology Rate Limiters to the Advancement of Climate Change Science, Draft Report on the Charge to the Joint ASCAC-BERAC Subcommittee on Climate

James Hack and Eugene Bierly of BERAC co-chaired this subcommittee, which was charged with investigating barriers to successful outcomes of the ASCR and BER complementary investments in climate change research.

The panel has found research opportunities in such areas as predictive regional modeling, modeling climate extremes and abrupt climate change, sea level rise (e.g., incorporating land ice melt), understanding human-induced impacts on the carbon cycle, and improving models by reducing systematic modeling errors and by upscaling to resolve resolution disparities. Rate limiters to these opportunities range across every aspect of climate theory, observation, and computation. For example, assimilating the ocean into models will require the HPC to explore design space at resolutions 100 km or smaller. Multiscale mathematics is needed to model climate variability where interactions range widely over time and space scales. Other needs include model evaluation tools and greater investment in algorithms, production-quality software, and computational capacity and capability.

Preliminary recommendations to reduce rate limiters include DOE investment in (1) collaborations to create algorithms and scalable software for the field and (2) high-end HPC computational facilities, data storage facilities, and supporting environments, tools, and technologies. DOE should also develop (1) greater predictive capability of higher-risk, lower-probability events like abrupt climate change, (2) the foundations for new modes of climate simulation, and (3) stronger understanding of the carbon cycle.

Discussion: When asked about specific short term changes, Hack said that there was no single obstacle to progress, but an overall need for more resources in many different areas, in particular for more computational resources. Discussion followed on the gap between current and needed computational capacity and capability and the extent to which these constraints limit advances. ASCR brings unique facilities to use as well as needed expertise. A tightly coupled ASCR/BER partnership could fill some of the empty niches. Smarter use of networking may be another area of promise. Climate change science also needs human capital. Even if there were no funding limits, there are not enough scientists with the right training for the field, Hack also fears that the rewards structures of academic science may discourage modeling research. Recruiting and training new scientists to pursue the wealth of climate change research opportunities will take innovative approaches and sustained effort.

David Brown: ASCR Applied Mathematics Research (AMR) Program: Strategic Planning Activities, Preview of the Applied Math Panel Report

ASCR’s half-century of support to world-class mathematics investigators has delivered an impressive record of mathematical models, analysis, and algorithms useful to many scientific fields. What should the next 50 years look like? A panel of distinguished leaders in math, statistics, and computer science met in August 2007 to give input to AMR’s strategic planning. From a review of DOE planning documents, they extracted the mathematical challenges DOE faces in the future. They found that advances in DOE mission areas—energy, environment, and national security—will require understanding complex systems and their behaviors over wide ranges of parameter space. In addition to developing better models and analysis of complex systems, the mathematics for integrating experimental data with simulations (“data-model fusion”) will become increasingly important. Interestingly, while the impact of future developments in the mathematics of complex systems will have broad impact on high-end computing, many of the mathematical challenges that must be addressed will not directly involve the use of high-end computing. Mathematicians will be the most successful at contributing to DOE mission challenges in the future if they begin collaborating with applications scientists early in the research and learn to ask the right questions.

The panel drafted goals for future research in the AMR program in three thrust areas: (1) predictive modeling and simulation of complex systems, (2) mathematical analysis of the behavior of complex systems, and (3) use of complex systems models to support decision-making. The mathematical research infrastructure developed by AMR over the past 50 years forms a strong foundation for the research that AMR must invest in for the future.

Discussion: Gordon Bell questioned the absence of goals for developing parallelized algorithms. Both Brown and Thomas Manteuffel asserted that mathematicians still have work to do even with everything parallelized and running at maximum clock speed. Brown added that many of the detailed strategies presented in the final report will include developing scalable algorithms.

Several committee members questioned the scope of the goals, noting the omission of areas not related to complex systems and of tools, like data analysis, not yet in the base program. Brown responded that mathematics for data analysis will be an important component of the final recommendations of the panel.

Zacharia questioned the omission of science from the list of DOE missions. Brown observed that the panel chose not to do a simple list of every possible topic and that some strategies he did not cover are under discussion. Science was viewed as an enabler of the other three missions, and SC programs had the broadest coverage by the panel.

There was interest in seeing proposals on what more ASCR can do to jump start collaboration between mathematicians and scientists.

Stechel asked if all the systems identified were all complex, rather than merely complicated since the former require different approaches to statistical predictability and uncertainty. Brown said the panel had discussed probability and uncertainty at length, but questioned whether complex and complicated systems were easily distinguished.

After a break beginning at 10:50 a.m., the meeting resumed at 11:05 a.m.

Douglas Kothe: Early Access Pioneering Applications for the 250 TF Leadership System at the ORNL Leadership Computing Facility (LCF)

Following the ORNL roadmap for HPC, in December 2007 the LCF will upgrade its Jaguar Cray XT4, already the world's fastest open computer, taking it from 119 to 250 teraflops. After acceptance tests are completed, documented, and reviewed, the new system will enter the transition to operations (T2O) period during which it will run six applications, which have been chosen for exclusive early access. During the six weeks of T2O, real, at-scale science can be done with these applications, subjecting the system to an actual production workload, and addressing any outstanding problems found during acceptance. The six applications were chosen a year beforehand from 20 that applied for their potential to challenge the system. An LCF liaison worked closely with each code team to prepare them for their run. A similar process will be followed when the Cray Baker petaflop machine comes on line in early FY09.

The applications are (1) ORNL's CHIMERA a multi-dimensional radiation hydrodynamics code designed to study core-collapse supernovae; (2) Princeton Plasma Physics

Laboratory (PPPL) GTC, which simulates turbulence in a large tokamak; (3) Sandia's S3D, which simulates turbulence in diesel and gas flames; (4) LANL's POP, an ocean simulation model that includes ocean biogeochemistry; (5) ORNL's DCA++, used to study high-temperature superconductivity; and (6) ORNL's MADNESS, with two applications for large-scale, all-electron, density functional simulations, and one for dynamics of few-electron systems.

Each applications has different science goals and impacts, physical models, algorithms, readiness activities, and scalability and performance issues, as well as varying requirements and simulation specifications. The LCF has an aggressive plan for getting all six applications through, with some portion, if not all, of their run during the six weeks of T2O.

Discussion: In response to questions about ongoing performance engineering of these codes, Strayer said that 3 had met the criteria to be Joule metric codes for SC in FY07. Kothe noted that benefits are already apparent from the cross-talk among these 6 teams, which are in disciplines that do not usually talk with one another. Although several codes are expected to do well in the 2008 runs, improving utilization performance of codes can take many years.

Michael Strayer: New Charge to ASCAC for an INCITE Committee of Visitors

Dr. Orbach charged ASCAC with forming a Committee of Visitors to review the process used to manage the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. INCITE was launched in 2003 to help research communities tap into the capabilities of current and future Office of Science supercomputers. The report is due at the Fall 2008 ASCAC meeting. Robert Voigt indicated that Marsha Berger has agreed to chair the COV.

There was no public comment. With thanks to the committee and to the presenters, Robert Voigt adjourned the meeting at noon.

Respectfully submitted,
Lisa Carroll
Recording Secretary
January 14, 2008