Minutes
Advanced Scientific Computing Advisory Committee Meeting
Aug. 5-6, 2008, DoubleTree Hotel and Conference Center, Berkeley, Calif.

ASCAC members present:

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

ASCAC members absent:

Jill P. Dahlburg, Chair	Thomas Zacharia

Also participating:

Melea F. Baker, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Christine A. Chalk, ASCAC Designated Federal Officer
Phillip Colella, Computational Research Division, Lawrence Berkeley National Laboratory
Emlyn Hughes, Physics Department, Columbia University
William Johnston, Computational Research Division, Lawrence Berkeley National Laboratory
Philip Kegelmeyer, Informatics and Decision Sciences Department, Sandia National Laboratories
Arnold Kritz, Chair, Physics Department, Lehigh University
Ron Minnich, Sandia National Laboratories
Frederick M. O'Hara, Jr., ASCAC Recording Secretary
Lenny Oliker, Computational Research Division, Lawrence Berkeley National Laboratory
Walter M. Polansky, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Michael R. Strayer, Associate Director, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Michael Wehner, Computational Research Division, Lawrence Berkeley National Laboratory

About 45 others were in attendance.

Tuesday, August 5, 2008

Preliminaries: Cochair Robert Voigt called the meeting to order at 9:03 a.m. He will be vacating the position of deputy chair at the end of this meeting. A new chair will be named before the next meeting. He announced the death of Edward Oliver, and a moment of silence was observed. Christine Chalk announced changes in the agenda. Public comment would be allowed immediately after each presentation.

Michael Strayer: Status of the Office of Advanced Scientific Computing Research (ASCR)

SciDAC 1 saw upgrades to ASCR's scientific computing infrastructure, established seven applied-mathematics Integrated Software Infrastructure Centers, and funded four very successful collaboratories, improving the effectiveness of scientific application codes by up to 10,000%.

The program was reorganized in FY06 to create a comprehensive, scientific computing software infrastructure, and more than 230 proposals were received, requesting approximately $1 billion. The landscape for facilities was changed to follow a path to the petascale. Applied math has been a major objective of SciDAC 2, new centers for technology have been established, and some existing centers have been realigned. SciDAC 2 also funds data-management and visualization centers and institutes. The SciDAC 2 collaboratories were stalled by budget challenges, have re-energized, and are being supported by the Earth Systems Grid, the Center for Enabling Distributed Petascale Science, and the Open Science Grid. SciDAC funding has increased from $57.3 million in FY01 to $79.9 million in FY08. The SciDAC 2 open-science annual conferences have grown in participation. A new publication, *SciDAC Review* has been established to highlight scientific discovery and the role of high-end computing. The Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program provides the Office of Science (SC) resources to a small number of projects that are computationally intensive, large in scale, and high in impact. The program is open to national and international researchers, including industry. There is no requirement of DOE funding, and funded projects are chosen by peer and computational reviews. In 2008, more than 265 billion hours were awarded to 55 new and renewal projects. The call for INCITE 2009 has gone out. At Oak Ridge National Laboratory (ORNL), 17 petaflop cabinets are in place ahead of schedule, on scope, and under budget. Argonne National Laboratory (ANL) has the world's fastest open science computer, Blue Gene/P (BG/P), ten months ahead of schedule. Of the top-ten computational-science accomplishments this past year, six were funded by SciDAC. ASCR's response to the E3 initiative is Transforming Energy, the Environment, and Science Through Simulations at the eXtreme Scale. However, it is not just extreme scale but also extreme complexity that requires figuring out how to deal with the unpredictable evolution of hardware, multi-level and heterogeneous parallelism, scaled-up programming models, data management, communications, and reliability. SciDAC X will address these challenges by wrapping up several problems, socializing the investigation, carrying that through to social impact that compels funding, and providing a focus.

Discussion: The rollup of SciDAC accomplishments will list the number of publications, citations, and awards, all appearing in a community report in a presentation to this Committee. Computing and simulation will play an important role in setting research agendas and formulating research questions in climate science, chemistry, radioactive waste, nuclear energy, bioenergy, and anthropology. These challenges will not be prioritized; the science should be socialized, collaborations should be developed, and a larger portfolio and resource base should be built for computing. This Committee allows a broad range of communities to interact with ASCR.

Administration: A break was declared at 10:19 a.m., and the meeting was called back into session at 10:36 a.m. Voigt expressed the Committee's appreciation of Jill Dahlburg's services to the Committee. The statement was followed by a round of applause from the members. Giles cited Voigt's service, and he received a round of applause, also.

Philip Kegelmeyer: Mathematics for Analysis of Petascale Data Workshop

The Workshop was held in Rockville, Maryland, June 3–5, 2008. Without capable analysis tools, the flood of data produced by science is locked away and unusable. The goals of the workshop were to understand the needs and to identify gaps that need to be filled to produce good analyses. Of the 66 attendees, half were from the application world, and half were from mathematics. The workshop found: (1) Algorithms must be re-engineered to scale with the size of the data. (2) More and more data are distributed, often non-uniformly, requiring algorithms that can come to the data and can adapt to the lack of a global perspective on the data. Validation and metrics are as important as the algorithms themselves. (3) New algorithms should make effective use of new computer architectures. Those new architectures differ widely, and their memory behaviors are different. (4) One way to attack scalability is data and dimension reduction. Metrics for estimating the accuracy and uncertainty induced are important here. Standard metrics do not work. (5) Models requiring data analysis are pervasive. Models should be driven by real-world data collection. Evaluating models is important: noise and skew are problems. (6) Visualizing uncertainty is difficult because of its statistical nature and human factors. (7) Detecting outliers is made difficult by the fact that what one is often looking for is new and rare. An analysis tool often can look at a data point only once. Furthermore, what is novel today may be common tomorrow. (8) An effort should be made to make cross-disciplinary discussion more prevalent and to encourage persistent storage and development of sample data sets. (9) DOE has made good investments in architecture, simulation, etc. Similar investments need to be made in mathematics.

Discussion: Social networks on the Web were considered, and people from Google and other such companies were represented at the workshop. Follow-up workshops would be interesting but are not in the purview of the current workshop leadership. Data curation kept coming up, but the issue was not sufficiently resolved; it is discussed in the published workshop summary. Visual analytics is an important tool in understanding data, and DOE recognizes that; some computational issues may not scale to the petascale, and those issues were addressed at the workshop. The implementation of methods for shuffling and re-analyzing data were considered at the workshop.

Marsha Berger: INCITE Committee of Visitors (COV) on Innovation and Novel Computational Impact on Theory and Experiment

The INCITE Program was conceived specifically to seek out computationally intensive, large-scale research projects with the potential to significantly advance key areas in science and engineering. It was originally based at the National Energy Research Scientific Computing Center (NERSC), where three projects used 10% of the cycles. In 2004, the Leadership Computing Facilities (LCFs) at Oak Ridge and Argonne were created, where 80% of the cycles were made available to INCITE. 10% of the cycles are allocated by the Assistant Secretary for Science. There has been a tremendous growth in INCITE submissions and allocations. INCITE provides the largest national leadership-class capability to the national science community and support for the DOE community. In 2008, 16 of 20 proposals received allocations. Program managers prioritize INCITE proposals, and recommendations are made to the Director of SC, who makes the final decision. Proposals are reviewed for scientific merit and for computational readiness. The COV found: (1) There is a tension between the DOE and non-DOE user communities. (2) The review process is nonuniform and nontransparent. The readiness reviews are perfunctory and do not necessarily identify projects ready to go on day one. (3) The INCITE portfolio has evolved so that there are a large number of projects taking up a small fraction of the facilities, stretching the capability of the support staff and devoting allocations to projects that are outside the leadership class. (4) The program aptly balances computing and the support services needed to effectively use complex computing environments. Both the users and the facilities believe that the 1-year proposal cycle is too long, but increasing the frequency would increase the workload for the program office. The review process almost always resulted in a continuation of project allocations despite minimal previous-year accomplishments. (5) The INCITE selection process attempts to balance conflicting requirements and lacks a transparent set of selection criteria. The scientific reviews are accomplished differently for different disciplines, making the comparison of scores problematic. However, no evidence was found that the outcome was unfair or biased. The COV recommended that: (1) The selection processes for leadership-class and high-end DOE capability class computing should be separated. A significant portion, but less than half, of INCITE competition resources should be allocated to high-end DOE capability-class computing. (2) INCITE awards should be fewer in number and larger in size, with the expectation of demonstrated concurrency across a very large number of cores. For projects deemed important but not ready, some resources could be reserved for development. Renewals should meet an achievement threshold below which projects are rejected. (3) INCITE should continue to provide robust expert assistance to the science teams performing leadership-class computing. (4) The selection process should be made as transparent and uniform across disciplines as practical. Consideration should be given to increasing the frequency of INCITE calls for proposals. The computational-readiness review process should adopt a more descriptive outcome, and an appeals process for allocation decisions should be implemented. (5) In approximately 5 years, a formal review panel should be convened to assess the impact of the INCITE program.

Discussion: The COV did not look at how light sources allocate access but did talk about different strategies of allocation. Many of these projects may not seem innovative, but that is because a lot of work must already have been done when the machines come online. The allocation of 80% of the LCFs to INCITE seems to have been overkill; one could use the SciDAC award process to select projects. Perhaps 60/40 would be a better ratio. Item 1 could change what LCFs do; they should address projects that need large-scale computing. The staggering of calls for proposals is a good idea. The DOE program managers need to be strict in their prioritizations. The original vision was for about 10 projects; now there are 55. The 80% allocation figure came from a request by the Office of Science and Technology Policy (OSTP), Office of Management and Budget (OMB), and SC. OSTP wanted 100% of the centers to be open to the general public, and OMB insisted that the open portion not go below 80%. OMB stated that DOE computing could go elsewhere. Not all program managers feel that this should be so; rather, DOE computers should be used for DOE science. SciDAC evaluations do not have any way to judge computational readiness, and some SciDAC projects do not adequately exploit the capabilities of LCFs. The question of endstations was not addressed. Other DOE offices allocate resources purely by scientific merit, but ASCR seems to make use of facilities on a programmatic basis. The COV looked at the allocation process and concluded that programmatic questions might be more appropriate for other subcommittees. However, the COV *has recommended* that improvements be made in the area of allocations. There is a growing need for larger allocations; at the same time, the number of users and proposals is growing rapidly. Industry is getting ready to use this capability.

Administration: Voigt called for an acceptance of the report subject to revisions made to address the issues raised at this meeting. A vote to accept the report was unanimously favorable. A break for lunch was declared at 12:20 p.m. The meeting was called back into session at 1:37 p.m.

William Johnston: Planning for, Status of, and Future Issues Related to ESnet

ESnet's primary mission is to enable the large-scale science of SC. The requirements of SC are determined by exploring the plans and processes of the major stakeholders and observing current and historical network traffic patterns. The major facilities require adequate bandwidth, connectivity, network services, and an architecture that can provide those services. Large-scale science uses systems that are data intensive and high-performance, have a high duty cycle, are widely distributed, and depend on network performance and availability. More than 85% of ESnet's traffic goes to and comes from outside of ESnet. ESnet4 is being built to address specific program requirements with more complex and much higher capacity than previous versions. The core network will have many rings and will be connected to metro-area-network rings to which data are shifted, where possible. One consequence of the ESnet's new architecture is that site availability is increasing. The current strategy was based primarily on 2003 information. With 1442 Gb/s budgeted for 2011, demand seems to be covered in the aggregate. However, surprises happen. Several Tier-2 centers are capable of 10 Gb/s now. Many Tier-3 sites are also building 10-Gb/s-capable analysis infrastructures. It is possible that the U.S.-installed base of Large Hadron Collider (LHC) analysis hardware will consume significantly more network bandwidth than was originally estimated. Technology advances will provide 100-Gb/s optical waves with comparable network interfaces, but there are some reservations about the affordability of those interfaces, and the planned 1 to 2 terabits of aggregate capacity may not be adequate to meet future needs. The planned ESnet capacity is growing more slowly than are the High Energy Physics experimental data and ESnet traffic. The current estimates from the LHC have the ESnet 2011 wave configuration operating at capacity, and other major sources will be generating significant data in that same time frame. Traffic will eventually overwhelm the capacity. With completely flexible traffic management extending down to the optical transport level, it should be possible to extend the life of the current infrastructure by 5 to 7 years by moving significant parts of the capacity to the specific routes where it is needed. To do that, the optical transport must be integrated with the network and provide for dynamism and loop flexibility at the optical level. The tools are there to make this process dynamic, but R&D is needed. Partnership R&D is a successful approach that is used with the OSCARS virtual circuit system that provides bandwidth reservations and integrated Layer 2/3 network management. DOE has recently informed ESnet that funding for OSCARS R&D will end this year. To best utilize the total available capacity, the optical (Layer 1) transport must be integrated with the network (Layer 2 and Layer 3) and provide for dynamism and route flexibility at all layers and to transparently direct routed Internet protocol traffic onto the Science Data Network (SDN). Meaningful information on the state of end-to-end paths is needed.

Discussion: In end-to-end service, one has to understand what services are wanted and when, which takes a lot of work. Why not wait until the instrument comes on line? Layers 2 and 3 will take care of this. Management of these flows will require R&D. Some telecoms are interested in such integration, and Nortel is one of DOE's partners because they are beginning to see some of the flows that ESnet has experience with.

Administration: A break was declared at 2:36 p.m. The meeting was called back into session at 2:49 p.m.

Michael Wehner and Lenny Oliker: Ultra-Efficient Exascale Scientific Computing

Exascale computing will revolutionize approaches to global challenges in energy, environmental sustainability, and security. Disciplines and problems that need ultra-efficient computing should be targeted for special machines that would be relevant to specific questions.

A feasibility design study is being prepared to influence the high-performance computing industry by evaluating a prototype design. A deficiency in climate science is resolving the global cloud system, but such simulation (at 1 km) will require higher resolutions than those now in use (200 km). Such a high-resolution simulation would require 10 petaflops sustained per simulation and 10 to 100 simulation runs. Such models would require advanced dynamics algorithms, about 20 billion cells, 100 TB of memory, and about 20 million total subdomains. A cooperative science-driven system architecture approach that would be applicable to numerous scientific areas is proposed. The primary design constraint is power, with estimates for exascale systems ranging up to 179 MW. A fundamentally new approach to computing design is needed, one that would identify high-impact exascale scientific applications; tailor systems to the requirements of that problem; and tightly coupled hardware, software, and science development. Wasted power from extra transistors, wasted computation, extra bandwidth, and serial designs would need to be eliminated. Design for low power means more concurrency. The Research Accelerator for Multi-Processors (RAMP) will be used to develop a highly programmable core architecture with a 10,000% power efficiency improvement that could be used to run a 1-km-scale model. A Tensilica XTensa system design is projected to produce 2.7 gigaflops at 3 MW for $75 million. This approach would produce multiple architectures, each different in the details, employing auto tuning, which automates the search across a complex optimization space to achieve performance far beyond current compilers. Traditional new architectures have a cycle time of 4 to 6 years, but the proposed new architecture could be done with a cycle time of 1 to 2 days. This approach is applicable to a broad range of exascale-class DOE applications.

Discussion: The dynamics are going to dominate from the computational view. The Navier-Stokes equation is the solution of that part of the model. One climate model is different from another, but all share a lot of similarities. One needs to map the algorithm and architecture together, but that does not mean that one will get an efficient algorithm. The best way to find out is to do it. One can leverage power efficiency by looking at the nearest neighbor. But that is not an efficient way to do it. The code being used is one developed by David Randall, which is trying to become the standard. The climate-change community wants to optimize for that model. However, there are other factors that will need to be dealt with at the 1-km resolution. One can look at the best available production codes, look at their parameters, and leverage their similarities. RAMP will let one test one horizontal subdomain with hundreds and thousands of processors. But with all the interactive models, one will never be able to find a fastest model. That is what is troublesome about specialized machines. One needs to look at whether the methodology is appropriate. Researchers prefer machines that look like what they are already using. Porting the application to the target machine has to go on in concert with the architecture development; a lot of programming avenues could be used. Finding the right program for exascale computing is a difficult task. A lot of memory bandwidth is needed. The important point is that the application has to have a lot of parallelism if 25 million subdomains are used. The size of the subdomains is the same as what is used today; so the memory bandwidth is tractable.

Administration: Public comment was asked for. There being none, the meeting was adjourned for the day at 3:33 p.m. so the participants could tour the Molecular Foundry and the Advanced Light Source at Lawrence Berkeley National Laboratory.

Wednesday, August 7, 2008

Preliminaries: The meeting was called to order by Cochair Voigt at 8:38 a.m.

Ronald Bailey: ASCR's Participation in the Fusion Simulation Project

The Fusion Simulation Project is a 15-year effort at approximately $25 million per year to get an integrated, whole-device framework for the simulation of tokamak plasmas to (1) support the International Thermonuclear Experimental Reactor (ITER) diagnostics and design after the first 5 years, (2) produce a simulation facility for meeting scientific and engineering objectives of ITER after the first 10 years, and (3) attain a validated simulation facility for a demonstration reactor world-class simulation after 15 years. The panel strongly recommends that ASCR seize the opportunity to participate in this project with the Office of Fusion Energy Sciences (FES). The project needs simulation, petascale computing, and a knowledge and software base. Its scope is not yet fully defined, nor are the stakeholders, customers, and users. A requirements specification needs to be documented by ASCR and FES. The simulation of ITER requires a scaleup from current devices by a factor of 25 in the magnetic field strength and by a factor of 10^{12} in space-time points from current codes. A lot of work would be needed to reach the ITER goals for calculational speed, and ASCR should focus its research efforts on applied mathematics, computer science, and computational science for the project. In applied mathematics, critical technologies include multi-physics, multiscale; data–model fusion; uncertainty quantification; realistic models; and data visualization and analysis. This project is rich with opportunities for developing new mathematical models, analysis techniques, and algorithms. In computer science, critical challenges include sustained performance, data management, visualization, and collaboration. FES and ASCR should (1) develop technology that ensures that performance meets expectations as simulation applications move among diverse petascale to exascale platforms; (2) develop and employ distributed data-management technologies, visualization technologies, and collaboration technologies to make this project's implementation more productive. In computational science, the critical challenges are frameworks, workflow, and verification and validation. ASCR has considerable experience and expertise in each of these areas and should take the lead in recommending appropriate solutions. ASCR should be a collaborator and a provider of basic research, technology, and infrastructure for this project. Its scientists should participate in integrated task teams while project management remains on the physics side. The SciDAC model may be employed in developing partnerships at the base research program level. Frequent interaction should be encouraged among the researchers involved. ASCR should ensure that applicable results from its other applied research are made available and, if necessary, tailored to the project. The facilities provided should include hardware infrastructure, software infrastructure, user services, and network infrastructure. A partnership between ASCR and FES should determine an appropriate set of hardware requirements for the project. The project will offer some special challenges: project pressures, lack of publications, technology insertion, software engineering, and the maintenance of productivity, many of which are outside the research culture of ASCR. ASCR should develop and implement methods to attract, motivate, and reward researchers to participate in the project. Early on, it should address the challenge of fostering a culture in which research creativity exists alongside project engineering. It should establish and enforce good software-engineering standards. It should establish a software-quality activity within the project. And it should leverage existing software-development productivity tools and further efforts to improve the software-development environment.

Discussion: This is a big engineering project and may not be able to be done without being laboratory directed. Not all the needed technology exists; new knowledge is required. A discipline-oriented project office needs to be set up, and new technology needs to be injected into the process. The exascale workshops led to the conclusion that programming languages will not be evolutionary. Algorithmic development will be necessary in addition to exascale computers' being developed. The scale change may transcend what ASCR can do, producing a management problem. However there is considerable experience in SC in dealing with sociological challenges. The computing and fusion communities must work together. The expertise is distributed, but there must be a lead institution. The Community Climate Model uses core groups clustered in particular institutions; that might be a good model to follow. Because there would be other users of the research, the risks for the fusion community are greater than those for ASCR. $25 million per year does not look small, but the fusion community has experience to leverage and may need co-investment.

Administration: Voigt called for a vote on accepting the report. The vote was unanimously in favor of acceptance.

Rick Stevens: GTL Subcommittee's Activities

The charge was to address the issue of computational models for GTL and how progress could be accelerated through targeted investments in applied mathematics, computer science, and computational biology. This revision of the Subcommittee's report recommends that: (1) The 10-year OMB Program Assessment Rating Tool (PART) goals should be accompanied by a specific set of metrics of progress. Such metrics could include (for a given organism) the number of correct metabolic phenotype measurements predicted, the fraction of an organism's genes and gene products included in a model, the number of transcription regulatory elements in a model, the number of correct gene-expression experiments predicted, the fraction of correct predictions of essential genes, and the number of organisms for which predictive models can be generated. (2) DOE should develop an explicit research program aimed at advancing the state of the art of cell modeling directly. That program should include equal participation from biologists and mathematicians, computer scientists, and engineers and be directly coupled to the more applied goals of bioenergy, carbon cycle research, or bioremediation. (3) DOE should establish an annual conference on the progress in predictive modeling in biological systems. (4) The GTL modeling and simulation research program should be supported by an explicit series of investments in modeling technology, databases, algorithms, and software infrastructure needed to address the computational challenges. The lower eukaryotes and plants should be included as targets in longer-term modeling and simulation goals. Such inclusions will advance the goals of Genomics:GTL by strengthening efforts to integrate the modeling and by advancing systems-level and synthetic knowledge for microbes and plants. (5) DOE should establish a mechanism to support the long-term curation and integration of genomics and related data sets to support biological research in general and specifically the needs of modeling and simulation in the areas of energy and the environment that are not well supported by the National Science Foundation and the National Institutes of Health. (6) DOE should work with the community to identify novel scientific opportunities for connecting modeling and simulation at the organism level to modeling and simulation at other space and temporal scales. A table of funding levels and the complementary potential program scopes was presented. At the higher end, $30 million per year would provide a significant implementation of the goals outlined in this report and enable partnerships with other agencies to broaden support for software for systems and synthetic biology. $50 million per year would allow full delivery for DOE of the opportunities for GTL from modeling, including intensive collaborative modeling and simulation efforts; this would create a major opportunity for interactions to accelerate discovery and minimize false or inferior directions for any project within GTL. At this level, SC could play a leadership role and significantly catalyze the growth of GTL and its impact.

Discussion: The questions raised seem to have been fully addressed. The learning factor is unknown, raising an interesting issue. Including eukaryocytes is a huge challenge. There is a groundswell of modeling efforts that use small-scale-computer clusters in the biological community. This practice raises a cultural issue. What is needed is an engagement between biological and computational scientists who can show the benefits of the methods of petascale computing. The communities are at different levels of development in using large-scale machines. Building tools would have a huge benefit. The authority of the Subcommittee to redefine the PART measures was questioned. The recommendation can be considered by DOE and used in their negotiations with the OMB in setting PART roles. The Biological and Environmental Research Advisory Committee (BERAC) has accepted the suggestion, and the program office is considering it. The data for the funding-level table came from discussions during the workshop. Of all the science fields, biology is the most uncertain of its models. There is a variation in maturity. In trying to do some practical task, the current options are limited and random in approach. Predictive skills *are* getting better, but they are narrowly defined. Global prediction is very ambitious, but engineering an organism with a minor change is possible and important. This is a hard problem, and now is the time to address it. A concerted effort by high-end computing and by biology could make real breakthroughs. The problem with the discussion is one of tone, not substance. A given model can only capture a small number of variables. The fact that development will be starting small and then dealing with massively complex systems will take great effort and time.

Administration: Voigt called for a vote on accepting the report subject to any small changes that BERAC might make. The vote was unanimously in favor of accepting the report. A break was declared at 10:36 a.m. The meeting was called back into session at 10:54 a.m.

Ron Minnich: Basic Input/Output System (BIOS) Issues

The BIOS is a small software system that resides in a flash memory chip on the PC motherboard and boots up the computer upon startup. Today, most computers have a BIOS with hundreds of system calls. This memory chip is large, and its unused portion can be used to hide an exploit. As the BIOS self-protection technologies improve, bad software is easier to hide, and measures taken to improve BIOS security actually degrade system security. Compromises of the BIOS cannot be remedied by reloading the operating system, removing the file, removing a malicious user, auditing the disc, or putting in a new disc. Operating systems now incorporate an advanced control and power interface, which is like a virtual machine in a virtual machine. It has the ability to modify system memory and thereby to allow an attacker to deploy a rootkit. Even the extensible firmware interface, which was a replacement for BIOS, allows an attacker to modify the bootloader directly, modify bootloader variables, modify and reflash firmware, or exploit an implementation flaw in the driver. To fix the BIOS security problem, an open-source BIOS could be used in which the customer supplies the BIOS image to a vendor and, upon receipt, verifies the BIOS image. INCITE supports the Plan 9 project on BG/P at Argonne. One application has been ported to Plan 9 with a small MPI [Message-Passing Interface] Lite library. A CNK personality mode was written for Plan 9 that can start a CNK [compute node kernel] binary and run it. A new network model is being developed for BG-style machines. It has very low latency but still runs through the kernel. It is not an OS-bypass interface but will have lower through-kernel latency than operating-system-bypass MPI.

Discussion: With a million nodes, cooperative checkpointing takes up a lot of time. Caching can improve performance by at least a factor of 4. Plan 9 allows one to talk to all resources in the same way, decreasing bandwidth and latency. Aggregates can be put together that allow easier checkpointing. The question was raised whether Hurd has anything to offer. One does not solve anything by shimming in Linux. Plan 9 puts only necessary things in the kernel. Linux has 300 system calls, and Plan 9 has 30.

Administration: Gordon Bell, Virginia Torczon, and Ellen Stechel are leaving the Committee, and they were thanked for their service and received a round of applause. A call for public comment received no response, and the meeting was adjourned at 11:41 a.m.

Respectfully submitted,
Frederick M. O’Hara, Jr., Recording Secretary
September 3, 2008