Minutes
Advanced Scientific Computing Advisory Committee Meeting
Aug. 14-15, 2007, American Geophysical Union, Washington, D.C.

ASCAC members present:

F. Ronald Bailey
Jill P. Dahlburg, Chair
David J. Galas
Roscoe C. Giles
James J. Hack
Thomas A. Manteuffel
Horst D. Simon
Ellen B. Stechel
Rick L. Stevens
Robert G. Voigt, Co-Chair (Tuesday only)
Thomas Zacharia

ASCAC members absent

C. Gordon Bell
Virginia Torczon

Also participating:

Melea F. Baker, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Christine A. Chalk, ASCAC Designated Federal Officer
Jerry W. Elwood, Office of Biological and Environmental Research, Office of Science, USDOE
Barbara J. Helland, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Daniel A. Hitchcock, Senior Technical Advisor, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Frederick C. Johnson, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Gary M. Johnson, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Frederick M. O'Hara, Jr., ASCAC Recording Secretary
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
William J. Valdez, Office of Workforce Development for Teachers and Scientists, Office of Science, USDOE
Jeffrey S. Vetter, Computer Science and Mathematics Division, Oak Ridge National Laboratory

About 35 others were in attendance.

Tuesday, August 14, 2007

Preliminaries

Chairperson Jill Dahlburg called the meeting to order at 9:00 a.m.

Michael Strayer: Advanced Scientific Computing Research (ASCR) Overview

Although ASCR's FY07 budget had been marked up at $318 million, the Office got only $283 million under the continuing resolution. As a result, the IBM Power 5 at NERSC [National Energy Research Scientific Computing Center] was not upgraded, SciDAC [Scientific Discovery Through Advanced Computing] activities were delayed, university partners were put on no-cost

extensions, and leases for the leadership computing facilities (LCFs) had to be stretched to avoid changing their scope. The FY08 request of $340 million has been approved by the House and reduced to $334 million by the Senate. The budget now awaits a joint-committee conference. The request would increase funding for applied mathematics, computer science, and research and evaluation prototypes. The House appropriations report was very complimentary of ASCR. The Senate report specifies $7.7 million for the Oak Ridge leadership computing facility, shifts $13 million from the Office of Science (SC) budget to the National Nuclear Security Administration (NNSA), and cancels a transfer of $19 million from DOE to the Department of Defense.

The President signed the Protect America's Competitive Edge (PACE) Act, establishing a DOE Director of Science, Engineering, and Mathematics Education (reporting to the Under Secretary of Energy for Science) and Advanced Research Projects Agency–Energy (ARPA-E).

Strayer distributed the Facilities Division's High-Performance Computing (HPC) Strategic Plan, which calls for support of a broad range of scientific disciplines; the tight coupling of pioneering scientific applications and of computer science and applied mathematics with future HPC resources; the conquering of barriers to success; and a balance between high performance production and leadership-class facilities.

Town-hall meetings were held on modeling and simulation at the exascale for energy, ecological sustainability, and global security. Other workshops were held on software for petascale computing, cybersecurity research, SciDAC, fast operating systems, visualization and analytics, applied mathematics, and petascale-system integration. Most of these meetings were reported on in detail later in the meeting.

At NERSC, the 100+ teraflop Cray XT4 (Franklin) is undergoing acceptance testing, the scalability program has been a great success, a 4-petabyte mass-storage system is in use, and the request for proposals (RFP) for the NERSC-6 upgrade is expected in FY08. Argonne is acquiring three IBM Blue Gene/P systems. The LCF at Oak Ridge National Laboratory (ORNL) is on budget and ahead of schedule with Jaguar upgraded to 119 teraflops and was recognized as the fastest open-science computer in the world as a result of a # 2 listing on the Top 500 list. An upgrade of Jaguar to 250 teraflops is scheduled for December 2007. In the Energy Sciences Network, ESnet4 is now operating from coast to coast; it is below cost and ahead of schedule.

More than 80 proposals were received for Innovative and Novel Computational Impact on Theory and Experiment (INCITE) allocations.

Christine Chalk: The Federal Advisory Committee Act (FACA)

The basic purposes of FACA are to enhance public accountability of advisory committees and to reduce wasteful expenditures in their conduct. There are more than 900 such committees with more than 45,000 members holding nearly 7,000 meetings per year at an annual cost of $400 million. Inter alia, FACA requires agencies to file a charter, maintain a balanced membership, hold open meetings, document all meetings and make those documents publicly available, scrupulously follow ethics constraints, and appoint a designated federal officer (DFO). DOE advisory committees are appointed by the Secretary of Energy. Subcommittees are, generally, not subject to FACA.

Discussion: A website for comments and feedback was seen as a good idea. It will be up and running as soon as the firewall issues can be resolved. Each advisory committee is rechartered every two years, and each one's charter will be changed to have the committees report to the Under Secretary of Energy for Science.

Action: A break was declared at 9:51 a.m., and the meeting was called back into session at 10:18 a.m.

Daniel Hitchcock: Network Update

The ESnet management structure has been simplified. Several program management tools have been put in place: workshops with the offices of Basic Energy Sciences (BES) and Biological and Environmental Research (BER) on program-specific network requirements, site visits, a consolidated network-requirement plan, an ESnet service plan, and a roadmap.

The Interagency Task Force on Advanced Networking had its preliminary plan out for review. That plan calls for providing network services anywhere, anytime; making secure, global, federated networks possible; managing network complexity and heterogeneity; and fostering innovation through development of advanced network systems and technologies. The comment period closes Sept. 30, 2007, and the final report is due to the Office of Science and Technology Policy (OSTP) in April 2008.

Discussion: It is hoped that the issue of the R&D roadmap will be addressed in the FY09 budget. The draft roadmap is out for review. It is good for the next 3 to 4 years. This is a roadmap for ESnet; however, the long-term problem is stitching together large networks around the world. Work is still waiting to be done on the overall strategy among agencies in regard to roadmap development. The agencies addressed the definition of the problem rather than dividing up the topics. Actually, only a few agencies support network R&D. Others are interested in specific applications. Federation issues are of interest mainly to the National Science Foundation (NSF) and DOE. The roles of agencies will be defined by their interests and missions.

Walter Polansky: Scientific Discovery Through Advanced Computing, An Update

A group of program managers from ASCR, BER, BES, Fusion Energy Sciences (FES), High Energy Physics (HEP), Nuclear Physics (NP), and NNSA coordinates the SciDAC activities. The National Science Foundation (NSF) also participates in SciDAC, but NSF program managers do not participate in the Coordinating Committee.

In FY 2007, new awards were made in accelerator modeling and simulation and in computational biology, broadening the scientific coverage of SciDAC. SciDAC funds science-application partnerships (SAPs) in climate change, petascale communications, materials modeling, energy density functions, accelerator-design modeling, fusion simulations, subsurface flows and biogeochemistry, astrophysics, turbulence and shock waves, and computational biology.

SciDAC has established nine Centers for Enabling Technology with 30 university partners and four Institutes with 12 university partners. In addition, a pilot-program outreach center has been established as a central resource for inquiries and technical information about SciDAC. The start of some awards within the Centers and Institutes were delayed because of the Continuing Resolution at the beginning of FY 2007.

A very successful SciDAC conference was held in Boston with more than 300 attendees, 36 plenary talks, 76 posters, and 2 panels. In addition, a tutorial in using SciDAC software was held at Massachusetts Institute of Technology.

Discussion: The funding for SciDAC is flat in the FY 2008 Budget Request. Delaying funding for certain universities was seen as a dangerous precedent, and it was recommended that it should not be repeated. In another budget cycle, it would involve a funding cut rather than a ramp-up delay. Another delay would also affect those who would use the centers' and institutes' output.

Action: A break for lunch was declared at 11:25 a.m. The meeting was reconvened at 1 p.m.

Ellen Stechel: Role and Efficiency of Networking and Networking Research

The Subcommittee on the Role and Efficiency of Networking and Networking Research has held biweekly teleconferences, held panel discussions and heard testimony, attended or reviewed pertinent workshops, and referred to informative websites. The major findings are (1) ESnet has done an excellent job of deploying an optimized and cost-effective network and of staying ahead of the bandwidth curve; (2) a strategic vision for the network and networking research is needed with full stakeholder buy-in; (3) the network and networking research are important to the enabling of DOE research; (4) petascale science is producing unmet needs in archiving, remote control of experiments, remote data access, re-usability of data, collaboration, and data sharing; and (5) the gaps between research and test-bed deployment and production need to be bridged. The Subcommittee also found that the current networking R&D program in ASCR is seen to be neither well defined nor aligned with strategic needs. A major question is, how should the research get matured through deployment to production, through testbeds or other mechanisms? The network is one element of a very complex whole, and managerial responsibility often stops short of moving a technology to the next step of maturity. The Subcommittee recommends a system-of-systems approach that concentrates on the interfaces and boundaries of components and on their connections and arrangement to make the elements work together. The Subcommittee is now preparing its preliminary report.

Discussion: There will be a preliminary report in November to be considered for acceptance at the November meeting. Scoping out a testbed is beyond the Subcommittee's capability. The Subcommittee will not prescribe a model.

Rick Stevens: Subcommittee on Computational Biology and GTL

The Subcommittee on Computational Biology and GTL, a joint effort with the Biological and Environmental Research Advisory Committee (BERAC), is charged with examining the issue of computational models for GTL. The Subcommittee is gearing up for input meetings and teleconferences this fall with a report to be written and submitted to ASCAC and BERAC before their November meetings. BER and ASCR need a common set of science goals to drive advances in bioinformatics, computational and theoretical biology, and systems biology. The existing Program Assessment Rating Tool (PART) goal is vague, and progress toward it is not easily measured. The goal is to describe the many levels of biological organization from genome to whole organism. Current techniques include microbial-cell modeling, functional behavioral assessment (FBA) modeling, flux modeling, and regulatory modeling. Another approach is to build a computational model of a cell and use it as a testbed organism. The modeling of simple microbial communities (e.g., the acid-mine-drainage-sediment community) is producing lots of data, and reconstruction of these communities has begun. Such reconstruction can shed light on the biochemistry of an environment. Extending the models to include the environment is the key to progress. The Subcommittee is focusing on a short list of exemplar goals to show what progress might be possible: microbial communities associated with carbon sequestration, microbial communities associated with bioremediation, communities associated with cellulose degradation, and a synthetic model (perhaps a minimum structure) for analyzing measurement needs for systems identification. The goal should probably be to develop integrated predictive models relating cell processes, phenotypes, and response to the environment.

Discussion: The report will be ready for consideration at the November meetings of ASCAC and BERAC. There is a feedback mechanism for changing the goal, but this is a PART goal, so it would have to be changed by the Office of Management and Budget (OMB). This recommendation from ASCAC and BERAC would start that process.

Jerry Elwood: Role of High-Performance Computing in BER

BER uses HPC for computational biology, contaminant transport and fate in subsurface environments (from the cellular scale to the field scale), and climate and Earth-system modeling. Computational biology uses will ramp up with the opening of the new Bioenergy Research Centers. Climate and biology are already important areas in the INCITE allocations at ORNL and LBNL/NERSC. Computational biology includes investigating enzyme actions, assigning gene functions, and simulating complex biological systems to develop enzymes for cellulose hydrolysis, fermenting organisms, and engineering feedstock crops. The biology roadmap leads from comparative genomics to dynamic simulations of complex molecular machines. The climate roadmap calls for the development of dynamic vegetation, biogeochemistry, tropospheric chemistry, and cloud-resolving models. In its various programs, BER needs HPC to quantify model uncertainty, develop test and apply high-resolution models, conduct modeling experiments to develop a predictive understanding of the functioning of complex biological and environmental systems over a wide range of scales, and develop multimedia predictive models.

Discussion: Asked if the allocation model is serving the needs of BER, Elwood replied that some users are not getting the needed allocations. The principal investigators of SciDAC proposals are being put at double jeopardy in the review process - resulting in their proposals being subjected to two separate, independent reviews, one for merit review of the science and a second review that evaluates the request for HPC cycles needed to conduct the proposed research. The Earth Systems Grid is important because it is used to send large amounts of data between laboratories. It is essential to the BER program. If the model data are not analyzed, the models cannot be improved.

Michael Strayer: New Charges to ASCAC

Two new charges to the Committee were received from Under Secretary Raymond Orbach. The first is to convene a joint subcommittee with BERAC to investigate bottlenecks and barriers to joint ASCR-BER research, specifically in climate modeling. A report of the Subcommittee will be due at the November ASCAC meeting. James Hack has agreed to be the lead for ASCAC, and Eugene Bierly the lead for BERAC. The second is to assess the strategic priorities of ASCR, focusing on the balance between core research and HPC facilities and the balance between near-term research needs and long-term investments. A subcommittee report is due in February to inform the FY10 budget process that begins in that month. The Subcommittee will be co-chaired by Jill Dahlburg and Robert Voigt, will be formed by September, and will have a week-long meeting in late December or early January.

Action: A break was declared at 2:37 p.m. The meeting was called back to order at 2:53 p.m.

Robert Voigt: Committee of Visitors (COV) Preliminary Report on SciDAC-2

The SciDAC-2 COV was held July 17-18, 2007, in Germantown with presentations by ASCR, BER, HEP, NP, and FES. The COV received excellent support from all program offices involved, but particularly from ASCR. The SciDAC-2 RFP was very complex, had only general review criteria, and led to some confusion on the part of both proposers and reviewers. The 240 proposals were down-selected by a series of panels, and ultimately remarkable balance in the portfolio of awards was achieved via significant effort on the part of program directors. Documentation was spread out among many DOE offices, and the jackets exhibited significant inconsistency. Analysis of decisions was very spotty, and communication with PIs was inconsistent. The program requires close scrutiny and external review. Recommended are: a more clear and complete RFP; separation of the selection of science applications (SAs) from Centers, and Institutes; a more detailed letter of intent; more-focused review panels; inclusion of

data on prior SciDAC performance in the review process; more-detailed analysis in the documentation; submission of reviews to PIs; electronic archiving of jackets; and annual peer review of funded projects. Overall, the COV found that SciDAC is an excellent program, and SciDAC-2 resulted in an impressive portfolio of activities despite intense time pressures on the selection process.

Discussion: The broad range of inconsistency in the way proposals responded to the RFP was striking. More programmatic guidance should be given in the RFP. It is puzzling that only one proposal was offered in some technical areas of interest; however, one cannot get a response if there is not a community willing to support that area. The COV did not consider whether fewer awards should be made more frequently for science applications and associated partnerships (SAPs), only for Centers and Institutes. Having different start times would be advantageous. The concern was raised that an annual review of the program would be a tremendous burden on the participants, and the possibility was offered that a midterm review might be adequate. The COV had talked about this issue at length and was unanimous about the benefits of an annual review. The question was raised whether the same objective could be met by annually reviewing only a portion of the awarded projects. The COV considered that to be possible but not as effective. Technical reviews get all the parties together and force them to work as a group. Also, they provide an opportunity to make changes in the project and establish a metronome for the project that sets deadlines for papers, runs, etc. Because 35 to 40 projects would require review annually, bringing all the team members together to be reviewed by an outside panel would be very complicated. Compromises might be necessary, but annual face-to-face meetings would be the ideal situation. NSF has done virtual site visits, which limited costs, travel, and time needed to conduct the review.

Action: Dahlburg polled the committee on accepting the COV report as represented by the substantive presentation by Voigt. The Committee strongly supported the report with four members expressing reservations about the potential effects on the staff of annual reviews. The Committee also expressed support for staggering funding. The consensus was to adopt the report. Voigt pointed out that these are general recommendations. The Office can support and adopt these recommendations or not.

Frederick Johnson: DARPA, FASTOS, and Petascale Tools Workshop

The Defense Advanced Research Projects Agency (DARPA) High-Productivity Computing Systems (HPCS) Program, which is to provide a new generation of economically viable HPCSs for the national-security and industrial-user community, has selected two vendors [Cray and IBM] for Phase III of the program. The funding of productivity work in Phase III is the responsibility of the vendors. ASCR has been involved in this program since day one. There is a difference between the House and Senate versions of the FY08 Appropriation with regard to DOE funding of this activity.

The program to develop fast operating systems (FASTOS) is being recompeted, and final proposals were due June 11. Thirty proposals were received and reviewed, covering operating-system framework/kernel/virtualization, input/output/storage, fault tolerance, and runtime among other topics. Awards will be made in FY08. Strong proposals exceed available funding.

A workshop on tools for petascale computing was held in early August, focusing particularly on tools for measuring and managing performance of petascale computers. The point of the workshop was to stimulate a community dialogue regarding the challenges of petascale tools. A workshop report is due by the end of August. Petascale computing will require performance tools with increased automation, abstraction support, runtime adaptation, and heterogeneity in

programming models and hardware. These tools must handle 100,000-core architectures, 2-GB executables, a large number of libraries, and multiple architectures in a single node.

Discussion: The applications will be a major challenge to scaling up to the petascale. Quite a bit of discussion at the workshop had been devoted to taking research results and getting it into production. That issue will be considered further. Another set of SciDAC workshops in July also looked at tools and runtimes. That discussion might help get these ideas into production.

Barbara Helland: Innovative and Novel Computational Impact on Theory and Experiment, An Update

The Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program had 88 proposals asking for 184 million processor hours on the ORNL and Argonne National Laboratory (ANL) leadership computers, NERSC and high performance computing resources at Pacific Northwest National Laborator; 95 million hours were allocated to 45 projects.. A quarter-billion hours will be available for allocation under this program in 2008; proposals are due Sept. 5, 2007; awards will be announced in December; allocations start in January. Major accomplishments of INCITE include the first-ever control runs of the community climate-system model CCSM 3.5 at groundbreaking speed, insight into the molecular basis of Parkinson's disease, and the first 3-D simulated detonation of a white dwarf. The success of INCITE is based on the hard work of the Office of Science high performance computing centers.

Discussion: A second round of INCITE awards in midyear had not been considered but could be done. Committee members were struck by Orbach's involvement in allocations, but he has always had the final say. Turned-down startups can get other allocations; proposal reviewers point such applicants to other resources at the centers, and the Under Secretary also has some discretionary allocations.

Action: Dahlburg asked for public comment. There was none, and the meeting was adjourned for the day at 4:39 p.m.

Wednesday, August 15, 2007

Preliminaries

The meeting was called to order at 9:05 a.m. Chalk announced that the next meetings of ASCAC are tentatively scheduled for November 6–7, 2007, and February 26–27, 2008.

Gary Johnson: European and Asian Supercomputing

In the top 12 computers in the world, there are currently no Japanese machines. The United States is back on top. The only European machine is in 12th position; it has a French manufacturer - Bull.

However, the HPC hardware race is on. The Japanese Ministry of Education, Culture, Sports, Science, and Technology (MEXT) has selected supercomputing technology as one of the nation's key technologies needed, and the country has three machines running at more than 40 teraflops. MEXT is calling for continuous development of supercomputers with no more one-shot systems (like the Earth Simulator). A next-generation supercomputer R&D center has been established at RIKEN (Rikagaku Kenkyusho), and it is partnering with the Japan Agency for Marine-Earth Science and Technology, the owner of the Earth Simulator. They have announced 21 target applications in pharmaceuticals, semiconductor materials, structural analyses of cars, seismic analysis, and storm prediction. These appear to be mainly for applications benchmarking purposes. The shorter list of priority applications appears to be biology and

nanoscience. The applications focus of the new Japanese effort is very significant. The Next-Generation Supercomputer Project is, in large measure, about the use of supercomputers – not just their design and construction.

The European Union (EU) has set up the Partnership for Advanced Computing in Europe (PACE). Its roadmap is similar to that of the United States, and its strategy is similar to that of the Japanese. Europe plans to move into leadership-class computing. The EU wants to see a unified proposal from the PACE partners to develop Tier 0 capability, a sustainable infrastructure, and complementarity with the national HPC centers and services. Education in computational science is an important part of their plan – as is partnering with industry.

Around the world, applications are currently run predominantly on U.S. hardware. Both Europe and Japan are building out their HPC infrastructures. There is a strong feeling in Europe that the EU needs to develop an independent capability to produce HPC hardware. Both Europe and Japan are focusing on an international HPC competition for the biggest/fastest. But, most importantly, both Europe and Japan recognize the importance of ensuring that HPC hardware is effectively used in science and industry.

Discussion: China could ramp up quickly. The European Union and Japan do not seem to be linking computer science to theory as SciDAC does, although in Germany, they have Technomathematik (the sense that mathematics is an equal partner with technology in the conduct of R&D). G. Johnson did not see anything like SciDAC in Japan right now. The French manufacturers currently do not use European components; they use Intel processors. However, they will not always rely on American parts. In Japan, the split between domestic uses of applications as opposed to international uses is still evolving. The focus of the Earth Simulator was on domestic usage. RIKEN leadership has a different approach; it is pursuing collaborations and is more open to a broader involvement. It was pointed out that a program in Japan is trying to launch SciDAC-like institutes for international teams; ten institutes are being funded at $5 million per year apiece and are bringing in foreign nationals for lengthy tenures. INCITE is available internationally, also; and the Europeans and Japanese recognize the benefits to competitiveness of openness.

Horst Simon and Rick Stevens: Modeling and Simulation at the Exascale for Energy and the Environment (E3)

The three ASCR scientific computing facilities organized three town-hall meetings on opportunities introduced by exascale computing and how it will revolutionize the approaches to global energy, environmental-sustainability, and security challenges, establishing computation as a critical tool along with theory and experiment. The town hall meetings were held in Oak Ridge, Berkeley, and Argonne this past spring to see what would be scientifically possible beyond the next 3 or 5 years in modeling and simulation. Breakout sessions explored potential advances in climate and biogeochemical modeling, energy, microbial life, socioeconomic modeling, astrophysics, mathematics and algorithms, software, cyberinfrastructure, and advanced computer architecture.

In climate/biogeochemistry modeling, improvements in representing biogeochemical cycles, spatial resolution, and quantification of uncertainty were expected. In energy, computational contributions to nanoscience and material science for renewable energy, fission systems, and fusion energy were expected. Microbial life is at the base of all ecosystem functions on Earth, and computational techniques will contribute to reconstructing genomes, modeling metabolisms, and modeling multispecies communities. Global socioeconomic modeling offers huge opportunities for exascale computing in detailed treatments, coupled models, and statistical analysis. Astrophysicists want to model and analyze large-scale structure formation, galaxy

formation, stellar evolution, supernovae, and compact objects. In mathematics and algorithms, the big themes were complex systems and high-dimensional spaces. The major software challenges will be computer architectures, a replacement of message passing, fault tolerance, and analysis of the tsunami of data produced. Cyberinfrastructure needs will include workflow management, collaboration frameworks and techniques, data management, authorization and authentication, and performance analysis.

By 2017, systems with 10 to 100 million processors, 1000 cores per socket, 3-D packaging, large-scale optical interconnects, up to 100 petabytes of memory, and ubiquitous fault tolerance are likely. Even at the stretch goal of 100 gigaflops per watt, an exascale system would still consume 10 MW.

A summary report has been produced.

Discussion: The path forward on hardware is clear, but software is driven by applications. One needs to get buy-in from those doing those applications. These town meetings were a first step to get application developers thinking. It is hoped that further interactions will occur this fall. In regard to the range of predictability of uncertainty in these models, an aphorism was quoted: "The current state of the art is to build a model and to try an experiment. From these one might learn that whatever policy is proposed, the predicted situation is going to be worse, making one re-evaluate the policy." Role-playing is used as a validation step for the models, but the computer models are a better predictor than the people because people forget their roles and let externalities govern their behavior.

It is easy to map out the hardware, and some candidate applications can be identified. The question is, how does one come by the mathematics (algorithms) for breakthroughs that may happen? The opinion was offered that the algorithmic knowledge will be sufficient; the problems encountered will be in fault tolerance etc. The mathematics section of this report is well done *and* calls for exploring algorithm futures and making sure that the mathematics are adequate through speculative mathematical inquiry.

The domain of hybrid computing is mostly in communications, like in cell phones. Some experiments are being done by researchers with access to hybrid systems, but progress in exporting those experiences has been slow. One should be cautious in setting hybrid computing as a target.

In the United States, DOE is in charge of driving HPC, which is recognized as an important part of future competitiveness. Biomedical science is also very important to this country, a $30 billion government-funded R&D effort. But DOE leadership emphasizes energy and the environment and does not consider biomedicine its business, even though the impact of HPC on biomedicine is going to be a very important part of the future. That is a strategic issue that needs to be thought about. The aggressive separation of biology into biomedical and nonbiomedical components needs to be moderated. The leaders of the town-hall meetings tried to stay in the DOE mission space, but the two trains of biology do need to be reconciled. Many biomedical applications are not ready for HPC, but they are bearing down on the field. When the synchrotrons were started up, biologists made very little use of them. Now, about 80% of synchrotron usage is devoted to protein structure.

Action: A break was declared at 11:35 a.m. the meeting was called back into session at 10:48 a.m.

Jeffrey Vetter: The Future of Performance Engineering in HPC

Performance engineering is being developed and used to improve the performance of high-performance computers so that actual performance more closely matches predicted (theoretical) capabilities as these machines become more architecturally complex, scale up, and run more

sophisticated applications. The useful lifetime of a machine has decreased from 17 years to 5, and new machines push the envelope on power consumption, heat production, parallelism, latency, and bandwidth. Performance engineers use measurement, prediction, and optimization to increase the productivity of computing systems by factors of more than 700. The Performance Engineering Research Institute (PERI) of SciDAC at ORNL accomplishes success through frequent interaction with application teams, the provision of tools, the development of automatic optimization, and the presentation of feedback to architects and system-software designers.

Discussion: Asked about best practices for working with an application team, Vetter replied that tutorials are provided to make performance personnel accessible to the teams. Liaisons are also sent out. The performance engineers themselves exchange information in all-hands meetings and in the literature. In regard to the relationship between performance engineers and algorithm engineers, it was noted that a piece of code can be optimized without changing the runtime; hardware use needs to be optimized *as well as* code. It is a symbiotic relationship. PERI has 13 participating organizations, including laboratories, universities, and companies.

William Valdez: Accelerating Innovation in DOE Through R&D Integration

DOE is attempting to accelerate innovation through R&D integration (i.e., through having division directors talk to each other, share program plans, identify resources, and coordinate activities). Integrated planning, investment, and assessment would reduce risk, accelerate achievement, capitalize on R&D investments, and create economic advantage. The creation of the Under Secretary for Science position by the Energy Policy Act of 2005 (EPAct) now provides a senior leader for R&D integration. A Science and Technology Council has been formed to identify areas ripe for integrated investment. Already, a Portfolio Working Group has reviewed opportunities and resources that could contribute to such integration in a number of DOE program areas: advanced mathematics, energy storage, carbon dioxide sequestration, characterization of radioactive waste, predicting high-level waste-system performance, and high-energy-density laboratory plasmas. That working group has also made a series of recommendations, including joint roadmaps, coordinating councils, coordinated budgeting, and independent review of integration efforts. Recent ASCR workshops have highlighted some of these opportunities for integration, and ASCR is being looked to for leadership and help to the other programs of SC. It needs to discover its appropriate role in knowledge transfer, improve the state of computation in DOE, explain its role to OMB and Congress, and realize the 10-year vision for exascale computation.

Discussion: Committee members observed that ARPA-E could address the problem of integration by shortening the time between basic and applied research. Valdez, in response to a question, said that discussions about the integration of technology programs focused on climate change have become a priority within DOE.

Concluding Actions

Dahlburg called for public comment. There was none. She thanked the support staff and adjourned the meeting at 11:48 a.m.

Respectfully submitted,
F. M. O'Hara, Jr.
Recording Secretary
Aug. 24, 2007