MINUTES
Advanced Scientific Computing Advisory Committee Meeting,
March 15-16, 2006, American Geophysical Union, Washington, D.C.

ASCAC members present:
- Jill P. Dahlburg , Chair
- John W. D. Connolly, Vice Chair
- F. Ronald Bailey
- Gordon Bell
- David Galas
- Roscoe C. Giles
- James J. Hack
- Ellen B. Stechel
- Virginia Torczon
- Stephen Wolff

ASCAC members absent:
- Thomas A. Manteuffel

ASCAC Ex Officio Members Present:
- George Cotter, National Security Agency
- Dimitri Kusnezov, National Nuclear Security Administration
- Tsengdar Lee, National Aeronautics and Space Administration
- Jose Munoz, National Science Foundation
- Doug Post, Department of Defense

ASCAC Ex Officio Members Absent:
- Daniel Harrad, Defense Advanced Research Projects Agency
- David Keyes, Society for Industrial and Applied Mathematics
- Peter Freeman, National Science Foundation

Also participating:
- Melea Baker, Office of Advanced Scientific Computing Research, Office of Science, USDOE
- Richard Colton, Director, Institute for Nanoscience, Naval Research Laboratory
- Donna Crawford, Director of Computing, Lawrence Livermore National Laboratory
- Anil Deane, Mathematical, Information and Computational Sciences Division, Office of Advanced Scientific Computing Research, Office of Science, USDOE
- David Goodwin, Mathematical, Information, and Computational Sciences Division, Office of Advanced Scientific Computing Research, Office of Science, USDOE
- Barbara Helland, Office of Advanced Scientific Computing Research, Office of Science, USDOE
- Daniel Hitchcock, Senior Technical Advisor, Office of Advanced Scientific Computing Research, Office of Science, USDOE
- Fred Johnson, Program Manager, Mathematical, Information and Computational Sciences Division, Office of Advanced Scientific Computing Research, Office of Science, USDOE
- Thomas Ndousse, Program Manager, Mathematical, Information and Computational Sciences Division, Office of Advanced Scientific Computing Research, Office of Science, USDOE
- Frederick O’Hara, ASCAC Recording Secretary
- Raymond L. Orbach, Director, Office of Science, USDOE

Joel Parriott, Program Examiner, Office of Management and Budget
Walter Polansky, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Jeannie Robinson, Oak Ridge Institute for Science and Education
Mary Anne Scott, Mathematical, Information and Computational Sciences Division, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Yukiko Sekine, Mathematical, Information and Computational Sciences Division, Office of Advanced Scientific Computing Research, Office of Science, USDOE
George Seweryniak, Mathematical, Information and Computational Sciences Division, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Rick Stevens, Director, Mathematics and Computer Science Division, Argonne National Laboratory
Michael R. Strayer, Associate Director, Office of Advanced Scientific Computing Research, Office of Science, USDOE
William Tang, Princeton Plasmas Physics Laboratory
Suzy Tichenor, Council on Competitiveness
Chris Yetter, ASCAC Designated Federal Officer
Thomas Zacharia, Director, Center for Computational Sciences, Oak Ridge National Laboratory

About 60 others were in attendance in the course of the two-day meeting.

Wednesday, March 15, 2006

Preliminaries

Before the meeting began, each of the attending ASCAC members was sworn in as a special government employee by a staff member from Human Resources, Office of Science (SC), U.S. Department of Energy (DOE). The Committee members were given an orientation session describing their responsibilities.

Chairwoman **Jill Dahlburg** called the meeting to order at 10:30 a.m. **Michael Strayer**, Associate Director of the Office of Advanced Scientific Computing (ASCR), introduced himself, welcomed the members, and thanked them for their efforts. The Designated Federal Officer, members of the Committee, and ex officio members of the Committee introduced themselves.

Department of Energy Perspective

Raymond Orbach reviewed the current activities and status of DOE’s Office of Science (SC). He termed “historic” the state of the union address in which the President announced the American Competitiveness Initiative (ACI) and a doubling of the federal commitment to the most critical basic research programs in the physical sciences over the next 10 years. Orbach called this commitment a testimony to ASCAC’s leadership. The ASCR program received the largest percentage increase in funding in SC. The SC budget went up 14% and has the potential to double in the next 10 years along with the budget of the National Science Foundation (NSF) and the core research program of the National Institute of Standards and Technology (NIST). SC is the steward of national science facilities that maintain U.S. world leadership in the physical sciences. It conducts forefront research and accelerates innovation through virtual prototypes. High-end computing (HEC) will ensure for the United States an order-of-magnitude dominance in key scientific fields that will transform the 21st century global economy.

Some of the highlights of the proposed FY07 budget are

- The International Thermonuclear Experimental Reactor (ITER) is fully funded. This is a self-standing international agreement for a major scientific facility. It will be the model for all future large-scale collaborations.
- High-end computation will provide more than 250 teraflops in Oak Ridge in 2008, and 100 teraflops on the Blue Gene/P at Argonne; the capacity of the National Energy Research Scientific Computing Center (NERSC) will be increased to 100–150 teraflops.
- Linac Coherent Light Source (LCLS) construction continues; it will be the world's first X-ray free-electron laser and will reveal the electron cloud during a chemical reaction.
- The Spallation Neutron Source (SNS) is being completed on-time and on-budget. It will be the world's leading neutron source for at least a decade.
- Four of five DOE nanocenters will begin operations in 2008.
- The International Linear Collider (ILC) would have up to 30 times the luminosity of Fermilab, and the United States has expressed interest in hosting the ILC at Fermilab.
- The Continuous Electron Beam Accelerator Facility (CEBAF) and Relativistic Heavy Ion Collider (RHIC) are operating at optimum running times.
- The National Synchrotron Light Source-II (NSLS-II) will be the first fourth-generation X-ray imaging machine with a 1-nm spot size. It opens a lot of opportunities and challenges for HEC, without which one cannot do the scaling.

ASCR is slated to get a 35.7% increase in budget from FY06 to FY07 as part of the 14% SC increase. About half of the SC budget will be for research, and half for facilities.

The President has committed to doubling funding for the physical sciences over 10 years. However, every reduction from the FY07 SC budget request would be compounded over 10 years. A 1% cut could mean as much as a $550 million reduction over the period. This funding increase is taking place in a period of budgetary stringency. These funds must be used wisely and well. That responsibility is in the hands of ASCAC and on the broader scientific community.

Discussion: Asked about the effects of congressionally directed funding, Orbach replied that, if those funds are redirected, SC's FY08 budget will be reduced; it is hoped that Congress will support the President's budget.

About the computational resources needed by the ILC, Orbach noted that we will need to figure out how to handle very large databases. Data mining is a possibility, but that is the wrong direction. New mathematics will be needed to understand the data. New computer facilities will be needed to deal with the data from the ILC, and the network will need to be upgraded.

Asked about the need for a community effort to analyze data from a large-scale program, Orbach said that DOE plans to do just that. Scientific Discovery Through Advanced Computing (SciDAC) does it already and will be used as the model for future efforts. Shared infrastructure is the future, and two solicitations are out now for such intermediate facilities and software centers.

In response to a comment that the scientific community is two decades behind in how it deals with large databases, Orbach said that we now know why neutron stars rotate and collapse because we can analyze them in three dimensions. That knowledge was not available when we were working only in two dimensions. A time of great discovery is ahead of us.

Asked if the Department was planning to procure one of the petaflop machines being assembled in partnership with the Defense Advanced Research Projects Agency (DARPA), Orbach replied, yes, otherwise DOE would not be investing in their development.

Office of Advanced Scientific Computing Perspective

Michael Strayer reviewed the status of ASCR, the vision of which is to produce best-in-class computational science for advancing science and technology through modeling and simulation by reinvesting in existing facilities while investing in new facilities, enabling technologies, and computational partnerships focused on petascale computers. He emphasized the uniqueness of SciDAC, which, together with ASCR facilities, could provide powerful resources and the nexus of a new global village for computing that could take computational science and scientific discovery to wholly new levels.

The SC 20-year plan lists ultrascale scientific computing capability as SC's second priority and upgrades for the energy sciences network (ESnet) and NERSC as SC's seventh priority. In accord with this plan, and given the language in the President's state of the union address, ASCR intends to

- Upgrade the Oak Ridge Leadership Computing Facility (LCF) to provide more than 250 teraflops peak capability by the end of FY07 and a petaflop by the end of FY08
- Acquire a 100-teraflop IBM Blue Gene/P high-performance computer system at the Argonne National Laboratory (ANL) in FY07, creating the Argonne LCF and increasing capability to 250 to 500 teraflops by the end of FY08
- Upgrade NERSC to a peak capacity in the range of 150 teraflops by the end of FY07 and to 500 teraflops peak capacity by the end of the decade
- Evolve ESnet over a 5-year period to dual backbone rings at 40 Gb/sec with fault-tolerant 10 Gb/sec connections to most major SC laboratories and higher bandwidth connections to NERSC, the LCFs, and other sites with exceptional data requirements

In addition, ASCR will

- Conduct research in applied mathematics and computer science to deliver the operating systems, programming models, software tools, and mathematical algorithms and libraries needed by scientists to make effective use of petascale computing
- Strengthen activities at the software centers initiated in FY06 for petascale computing and initiate research investments in applied mathematics and computer science to accelerate efforts in modeling and simulation at the petascale
- Coordinate the Research and Evaluation Prototype effort with the National Nuclear Security Administration (NNSA) and focus on the DARPA High Productivity Computing Systems (HPCS) program R&D partnership

The road to petascale computing will be played out in fusion, combustion, and astrophysics. Various laboratories are developing algorithms, and ASCR will develop the middleware and collaboration for a distributed network environment that will require advances in distributed data management, security, end-to-end performance, high-performance middleware, and partnerships.

The ASCR budget shows a sudden increase from FY06 to FY07 and steady increases each year until FY11. These increases will be shared by both facilities and research. In the FY07 ASCR budget request, funding for the LCFs is increased aggressively; NERSC receives a substantial increase for the buildout of leadership-class computing and an increased number of users; research and evaluation prototypes will receive $13 million to support the R&D partnership with the NNSA; and applied mathematics, computer science, and SciDAC will receive increases. ASCR currently has partnerships with seven other DOE offices, NNSA, and industry. Expansion of these partnerships is being actively pursued.

He pointed out that the full-committee meetings of ASCAC are public; that the Committee provides advice to ASCR by means of written reports in response to charge letters; the reports are written by subcommittees of ASCAC; reports are discussed at full committee meetings and voted on for acceptance; upon acceptance those reports are made public; and ASCR supports the activities of the Committee. Two charges are currently before the Committee. The first deals with performance measurement. All investments must be reputable and productive. In this effort, ASCR partners with the Office of Management and Budget (OMB) and the Office of Science and Technology Policy (OSTP). An interim report on performance measurement is due July 30, 2006, and a draft report is due October 30, 2006. The second charge deals with ESnet. In SC, $8 billion of facilities are producing data that has to be moved by ESnet. The Subcommittee's study should consider the proposed evolution of ESnet, its appropriateness and comprehensiveness in addressing the data-communication needs of SC that will enable scientists nationwide to extend the international frontiers of science. Furthermore, the Subcommittee needs to make suggestions and recommendations on the appropriateness and comprehensiveness of the networking research programs within ASCR with a view to its meeting the long-term networking needs of SC. An interim report is due October 30, 2006, and a draft report is due October 30, 2007.

ASCR is undergoing large changes in staffing. The budget has doubled while the staff has decreased and programs have become more complex. This problem needs to be addressed.

Discussion: Asked about the involvement with other offices of SC, Strayer replied that SciDAC alone reflected $70 million in commitments and contributions from the ASCR partners will increase from $38 million to $50 million between 2006 and 2007. In addition, there are investments from other agencies (e.g., NSF). In regard to new models for the integration of computation into the sciences, he replied that SciDAC will be the model for implementing computing science and will involve teams of computer scientists, scientists, and technicians. NASA did such collaborations in the 1980s with people having multiple intellectual capabilities pulling these collaborations together. This team approach is vital, and upcoming solicitations will call for unifying institutes. Ties to applications will be broadened. Computational biology is being pursued on several fronts (e.g., through institutes, pilot programs, and the new-math-tools initiative). Application development and buildout are tightly integrated, and applications will be running on the new facilities a short time after those facilities are completed.

SciDAC Status

Walter Polansky noted that planning for SciDAC started 5 years ago with an initial focus on software and the subsequent emergence of new scientific application areas and advanced computer architectures and systems software like modeling and simulation. Emphasis on university-based institutes and experimental science were added. The next stage, SciDAC-2, will create comprehensive, scientific computing software; infrastructure-integrating applied mathematics; computer science; and computational science for scientific discovery at the petascale and a new component of data management and knowledge discovery for large scientific data sets. Partners in SciDAC-2 include SC, NNSA, and NSF. The first successes in petascale computing will likely come from the SciDAC program. A strong coupling is expected between the applications and infrastructure. Significant results are expected across areas that reflect the research interests of SC.

Discussion: Asked how much of the ESnet and computing-facility resources SciDAC used, Polansky replied, about 50% at Oak Ridge National Laboratory (ORNL) and about 25% at NERSC. At times, SciDAC projects overstrained computer facilities and networks. In response to questions about meshing SciDAC with the programmatic needs of SC, Polansky replied that several mechanisms were used, primarily at the level of office director and associate director.

The peer-review process judges the intrinsic scientific merit and use of petascale systems. In the selection process, there are haves and have-nots; the haves recognize the advances that could be made within such a program. However, new work should not be excluded but enabled.

High-Performance Facilities and Testbeds

Barbara Helland presented the ASCAC Facilities Subcommittee report, which concluded that (1) current SC computer and networking facilities are among the best in the world and (2) current and expected near-term SC high-end computing resources are far from adequate to meet the anticipated needs of its science and engineering missions and associated scientific communities. The report recommended that SC

- Guide its investment decisions in high-end computing by the science drivers,
- Over a 3-year time frame, make substantial investments to return the United States to leadership in high-end computing, and
- Manage its advanced computing resources as a single, coordinated facility.

ASCR has two types of facilities in its portfolio: high-performance production computing and leadership computing facilities (LCFs). NERSC has a standard upgrade path with a 3-year cycle. NERSC-5 is being acquired to alleviate the current backlog of meritorious requests for resources; delivery is scheduled for FY07. The LCF at ORNL currently has an 18.5 teraflop Cray X1E and a 25-teraflop Cray XT3. The philosophy is to get applications online early. Staffers are assigned to help uses get the best use of the facility. The LCF at ORNL also has an upgrade path, and architectural diversity will be provided with the establishment of an LCF at ANL in FY07. Research and evaluation of prototypes will be initiated in the next year and will be focused on the DARPA High-Productivity Computing Systems (HPCS) program partnership.

NERSC and the LCFs are subject to regular external reviews. All review committees stressed the need for performance metrics. Each of the facilities is slated for a significant increase in funding in FY07. 85% of the resources goes to SC-mission-related work; 10% goes to the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program; and 5% is reserved for special projects (e.g., the analysis of hurricane surges along the Gulf Coast).

Performance metrics are currently focused on NERSC and call for 40% of the computing time to require at least one-eighth of the total resource (768 processors). The major challenges are to deliver petascale computing hardware by the end of FY08 for scientific discovery; to revise performance metrics to reflect the differences between capability and capacity computing; and to define a coherent allocation strategy, especially as capabilities are scaled up.

Discussion: Asked about allocation strategies, Helland replied that the program offices are used rather than a mechanism like the National Resource Allocation Committee. An effort is made to tie the needs for time and labor to the needs for high-performance computing.

High-Performance Networks and Associated Research

Daniel Hitchcock noted a significant increase in the budget for ESnet to accommodate the output from the Large Hadron Collider (LHC). A number of workshops are being held to see what the needs of the user community will be by the end of the decade. The current architecture of ESnet connects SC assets to scientists worldwide with a total downtime of <1% due largely to the replication that is engineered in. A redundant European connection is being installed in Washington, D.C., and engineering is being carefully coordinated with Abilene.

The Ultra Science Net testbed is being built largely on the Lambda Rail to explore means of distributing data at the petabyte scale. It is not used for production.

On ESnet, the top 100 flows account for 30% of the traffic, with big flows going to a small number of places. This large consumption of bandwidth must be gotten off the Internet because it looks like a denial-of-service attack. There will have to be a greater integration between the DOE networks and those that service universities.

An external review of ESnet in February endorsed the ESnet approach and called attention to the increased complexity and need for outreach to partners. In the future, sites will need hybrid networks, network peering, cybersecurity defenses, and next-generation DMZs (demilitarized zones). Needs include interdomain interfaces, end-to-end performance tools, trust and authentication tools, high-performance middleware, and integrated testbeds and networks.

Discussion: Questions centered on other recommendations from the review panel and who was working on these questions. Hitchcock replied that a strategic vision was never stated and should be made explicit. The ESnet management structure has become less connected to the clientele as the program has grown. A circuit-based network should be put in place. Research has been conducted in collaboration with Oak Ridge, Internet 2, Shiboleth, and others, and ASCR is partnering with the National Lambda Rail and Abilene.

DOE's Distributed Network

Mary Anne Scott noted that, to handle the very large experimental and simulation data sets it is producing, DOE needs to develop a scalable, secure, integrated, distributed infrastructure. To get there, it needs to research, develop, test, and deploy advanced network technologies, scientific collaboration tools, and frameworks that enable scientists to use unique and expensive enterprise research facilities. In the short term, it needs to do integration, prototyping, testing, and accelerated deployment of advanced computing, communications, and middleware technologies. In the long term, it needs to conduct fundamental research for advanced collaborative and network capabilities.

During the past 10 years, DOE has learned how to fully integrate real partnerships, develop leadership in distributed computing technology, and develop an effective pipeline from research through development. Difficulties encountered include cultural inertia, technology barriers, and organizational barriers. The opportunities and goals sought include cost-effective, agile network infrastructures to support high-impact science applications; gigabit networks and services for interconnecting data analysis and management centers; partnerships between scientists and network/middleware researchers; research network testbeds; and experimental network testbeds.

Funding decreased after 2003 and was flat from FY06 to FY07 at $13 million. Research investments are needed in middleware, network research, pilot programs, and research testbeds.

From this program has come the Earth System Grid (ESG), which has two portals serving climate data; non-onerous cybersecurity services; distributed data sharing; large-scale data processing and analysis; increased collaboration among scientists; and middleware that has enabled a sustained 10-day flow of 500 terabytes from CERN to its associated sites.

Discussion: Asked what budget would be needed in the other years, Scott replied $25 million. In the near term, that would probably be used on high-performance middleware, testbeds, and data-management issues. Middleware is the most difficult thing to do, but remote operations are impossible without it. However, cybersecurity and data management are the most important issues.

Processes Involved in Funding Science and Technology

Joel Parriott of the Office of Management and Budget explained that OMB has a hand in the development and resolution of all budget, policy, legislative, regulatory, procurement, e-government, and management issues. Its four program associate directors (PADs) oversee programs in natural resources, human resources, general government, and national security, respectively. In natural-resource programs, the competition is among DOE, NSF, National Aeronautics and Space Administration (NASA), U.S. Department of Agriculture (USDA), U.S. Geological Survey (USGS), Environmental Protection Agency (EPA), and the Smithsonian Institution. NNSA is in National Security.

In the budget process, agencies assess and prioritize their needs, they prepare a final budget request and submit it to OMB, OMB considers it and responds, appeals are settled by OMB and the President, a final budget request is formulated, and that President's budget request is submitted to Congress. In the course of this process, each PAD is expected to take care of the President's priorities, take care of other Administration priorities, be cognizant of Congressional priorities, fix other miscellaneous problems, and present a recommended program that clearly identifies where problems remain.

In addition, the OMB administers the President's Management Agenda with quarterly scorecard ratings. The Program Assessment Rating Tool (PART) has formalized questions on which recommendations can be based. The R&D Investment Criteria call for science programs to be evaluated on the basis of quality, relevance, and performance metrics. These evaluative tools look at how awards are distributed; how funds are distributed among national laboratories, universities, and other performers; whether there is a strategic plan in alignment with the mission and national needs; and milestones.

The OMB PART review found ASCR to be moderately effective and ASCR's external advisory committee to be underutilized. It recommended that ASCR engage its advisory panel and respond to past expert reviews.

In the President's FY07 budget, high-end computing is a priority. However, ASCR needs to formulate a set of clear, challenging, outcome-oriented performance metrics for its machines. "Getting up to speed" does not cut it; there should be codes that are ready to run when the switch is flipped on. The contractor should be held accountable for scientific outcomes. These machines also need to be open to the very best proposals from the broad non-DOE community with time allocated on the basis of a competitive, merit-based process 100% of the time.

ASCR should ask itself how to realistically implement its own recommendations within a fixed budget envelope and use the R&D Investment Criteria to drive arguments. ASCR does not yet appear to have a culture of engaging interested stakeholders in order to receive balanced, transparent, public advice. They should improve their consensus reports: Put together a compelling narrative and do a better job on executive summaries and navigational elements. Reveal the assumptions and context. Admit limitations. Workforce arguments are typically weak ones; let the science drive the case. Well-grounded constructive criticism adds to one's credibility. Strong participation by outsiders adds to a report's credibility.

Discussion: It was noted that NIST does not fall under the same PAD as DOE and NSF do, but the director of OMB can move money where he wants; in such a case, the two PADs have to get together to talk about it. Economic development is the driving force behind this budget. That produces a problem for curiosity-driven research. The OMB tries to make sure discovery science is supported as a stewardship issue. These computers should be treated like the light sources, where everyone can compete as long as everything is going to be openly published. However, national laboratories cannot compete to use NSF facilities, and that is not fair.

It was asked how SciDAC interacts with the program offices. Strayer said that DOE does not have a mechanism across all the program offices; implementation is done within five program offices, and each is different. The SciDAC experience has resulted in a number of high-level computing capabilities. In the future, other programs will be involved, as will other offices and agencies. This is a great opportunity to build a collaboratorium tool, which would be a great advance. Donna Crawford noted that NNSA has a wealth of experience in high-performance computing and should work more closely with ASCR.

The meeting was adjourned for the day at 5:20 p.m.

Thursday, March 16, 2006

Dahlburg called the meeting to order at 8:33 a.m.

ASCR's Computer Science Research

Fred Johnson said that the program's strengths lie in the co-array FORTRAN on the Crays, which uses memory-based virtual message passing; coordinated research activities; and ongoing partnerships with DARPA, Department of Defense (DOD), NNSA, National Security Agency (NSA), and High-End Computing Research Task Force (HECRTF). The challenges lie in data representation, exploration, and understanding for terabyte to petabyte datasets and in the development of petascale systems by 2008. At the petascale, there are too many nodes and processes to use brute-force approaches to debugging and performance analysis. Multiple system failures will be normal, and scalable software test beds are essential. The time from idea to robust, reliable software realization is at least 10 years. Operating systems were identified as a major topic to be addressed, and a number of workshops were held.

Three systems are being brought up simultaneously: NERSC-5, the LCF at ORNL, and the LCF at ANL. These systems need to be ready for full-scale applications and acceptance and to run fully capable, scalable, and reliable vendor and DOE software. This will be done by having applications drive the efforts; focusing on application-specific critical-path software; developing partnerships with vendors, national laboratories, and academia; leveraging NNSA experience and expertise; understanding scalability through testing; conducting comprehensive progress reviews; and adapting the computer science research program to support this effort. The ultimate goals are to produce facilities with 3.2 petabyte per second bisection bandwidth and 64,000 giga updates per second (GUPS), which is 10 to 10,000 times business as usual. The system must be able to handle one trillion files in a single file system

ASCR and HPCS have been involved in this process from day one. For Phase III, we are making a four-year $13 million per year budget commitment that will focus on the system software ecosystem and petascale application development. No commitment has been made to purchase any of these machines; competitive procurements will be based on cost/performance ratios.

Discussion: Algorithmic development will be included in software development and largely funded out of the Mathematics Department. Nothing on databases is currently included; that area could bear a lot more exploration. Also, computer science research needs to be extended to the entire user base. That is done through workshops and other mechanisms. Architectural changes are anticipated, but this will be a multimode program for a long time.

Blue Gene/P and Q

Rick Stevens said that the Blue Gene approach focuses on low-power, communications networks, and simplicity. A Blue Gene recently exceeded 150 teraflops sustained, Blue Gene swept all four HPC-challenge class-1 benchmarks, and large-scale applications are running on Blue Gene/L. Fourteen Blue Genes have been installed plus four at IBM. A consortium has been formed to rally the community. For Blue Gene, IBM selected several applications, looked at the computational needs, and mapped those needs onto the machine's architecture. This machine does communication very well, and it exhibits high power efficiency. But scaling it out beyond 1 petaflop produces power-consumption problems. What is needed is a new technology that depowers the circuits when they are not being used (active power management).

IBM plans to develop the Blue Gene/P, which is scalable to 1 petaflop, by 2008, and the designers are looking at Blue Gene/Q with power architecture and targeting 10 petaflops. In Blue Gene/P, the network topology is the same as in L. Performance will be doubled by doubling the number of processors per node. The functionality of the compute node will be enhanced. The Blue Gene/L packaging will be retained as much as possible. Software will be able to run both MPI [message-passing interface] only and MPI plus open MP, and the compute node kernels will add support for threads and limited dynamic linking. It will use the complete HPC software stack from the beginning. It will be built the same way Blue Gene/L was built, with an R&D phase, a build phase, and a review process. The development program addresses both P and Q development, with overlapping milestones set for both hardware and software.

SC and NNSA have agreed to support the continued development of Blue Gene/P and Q. Lawrence Livermore National Laboratory (LLNL) and ANL have formed a partnership to co-manage the development contract with IBM. The program will provide for technology demonstrations on Blue Gene/P in FY07 and Blue Gene/Q in FY10. Risk is managed through incremental milestones, reviews, progress payments, and multiple go/no-go decision points. Deployment of systems will be via separate build contracts focused on specific systems.

Discussion: Discussion brought out that the machine did much better than expected. The community rallied and put a lot of applications on the machine, and the machine is more reliable than expected. It was built on a shoestring budget. A number of strategies will be used to limit power usage. "Scheduled discoveries" are being postulated to maintain progress, which is risky and may lead to unintended discoveries as well as planned discoveries. There is cross-fertilization with the Cray platform; in the build contracts, the vendor is allowed to keep innovating up to the end. Life-science problems were postulated as being major users of previous high-performance computers, but these postulates turned out to be false. Developers turned their attention to large genomic databases because the machine has great storage capacity. Blue Gene will be useful to biologists because it is like a large Linux cluster.

Performance Measures

Thomas Ndousse said that the Government Performance and Results Act and the OMB Criteria for Assessing R&D Investment drove the adoption of performance measures. For PART, ASCR solicits advice from advisory committees, uses committees of visitors, conducts peer reviews, conducts Lehman construction reviews, develops strategic plans, conducts periodic external reviews, requires annual progress reports, reviews staff performance, and holds open workshops and conferences. PART performance measures are applicable to capability computing and to computational science capabilities.

For capability computing at NERSC, the number of processors varies from year to year. A target of 1/8 of the processors is set, and a target percentage of usage is set. Then the actual percentage usage by applications is calculated and compared to the target percentage. In the past 3 years, the actual usage has exceeded the target percentage.

For computational science capabilities, a target set of code and a target system on which to run the code were identified. The performance of the target code at the beginning of the fiscal year was known. The code was then tuned or new mathematical algorithms were developed during the year. At the end of the year, the code was executed on a target system with the same configuration. To be deemed successful, the annual improvement in the code, when measured in time to solution, must be 50%.

Performance measures encourage dialogue with OMB, force the evaluation of program progress and effectiveness, enable programs to set higher performance goals, and improve program management and strategic planning. The challenges of performance measures are that (1) quantifiable and sensible performance measures of R&D activities are difficult to define and (2) scientific discoveries are not predictable.

Discussion: The “magic number” (1/8 usage) that must be achieved is reviewed with OMB each year. One cannot see what each user is doing as a function of time; one has to do statistics and deal with ensembles because some projects are going to fail.

SciDAC Conference

William Tang, 2006 SciDAC Conference Chair, noted that the first three SciDAC conferences were structured as annual principal-investigator (PI) meetings. The SciDAC 2005 conference was run as a computational science conference with broader goals and more diversity in attendance. This conference helped publicize what can be done in this new collaborative mode. All presentations were invited by the organizing committee, and they reflected a broad computational science community. There was a large vendor participation and press presence. Proceedings were published and made available online, and the meeting was netcast. The new format also stimulated growth in the scientific material produced and preserved.

The SciDAC 2006 Conference, which will be held in Denver on June 25-30, will continue this successful focus with a launching of the SciDAC-2 program. It will include 39 invited talks, a vendor panel covering architectures, more than 60 invited posters, and cross-participation by the DOD High Performance Computing Modernization Program (HPCMP) users meeting being held in Denver at the same time.

For the SciDAC 2007 conference, a conference chair has been selected, an infrastructure team is in place, and the conference hotel is being selected.

ASCR’s Applied Mathematics Research

Anil Deane stated that the Applied Mathematics Research (AMR) Program contributes to the overall ASCR strategic goals by extending the frontiers of science through mathematical programming. Investment areas are numerous and involve algorithms development in each of the areas. These refined algorithms should be used and not replicated. The program started more than 50 years ago and is carried out in national laboratories and at universities. It has two components: the base and SciDAC. Historically, its activities have been divided into (1) linear algebra, optimization, and predictability; (2) partial differential equations and high-performance computing; (3) computational fluid dynamics, advanced meshing, and others; (4) multiscale mathematics; and (5) SciDAC activities.

The program's budget has been flat at $30 million per year, split about evenly between laboratories and universities. Funding for SciDAC in FY05 was $9,921,000. It is believed that new initiatives can continue to be supported while broadening SciDAC to include institutes.

Multiscale mathematics overcame the separation of scales common in large-scale modeling, as in climate, biology, high-energy physics, combustion, and fusion. Current models that assume the separation of length and time scales fail in practice. Therefore, in 2005, multiscale mass awards were made to 15 collaborative projects. In the future, the program has to maintain a balance between short-term and long-term horizons. It must also be in alignment with the ASCR priorities and look at the barriers to scientific progress.

Discussion: Asked about AMR's involvement with other programs, Deane said that the Global Nuclear Energy Partnership (GNEP) is like another application, but a big one; its key problems include the fuel cycle, reactor design, and waste disposition. The program will address large-database analysis, but it is not known what shape it will take. In nanoscience, the scientific problems are electronic structure and electrostatic dynamics, which also contribute to biology. AMR is developing a suite of tools that are applicable to solving density functional theory. The program has $4 million per year in nanoscience partnerships and SciDAC. A lot of the program's activities are done in close cooperation with the Computer-Science Program.

SciDAC Impact on Computing Facilities

Walter Polansky noted that about 40% of the cycles in FY06 at the Oak Ridge Center for Computational Sciences and about 25% of the usage at NERSC is attributable to SciDAC. The FY06 SciDAC proposals were reviewed for scientific merit and for computational needs, as they have been since FY04. Globally, SciDAC is having a negligible impact on ESnet because of the high usage by high-energy physics.

ASCR Partnerships with Other Offices

David Goodwin stated that, of the 100 million CPU hours at ORNL and NERSC in FY05, Basic Energy Sciences used 28%, Fusion, 25%, Biological and Environmental Research 16%, Nuclear Physics 16%, High-Energy Physics 11%, and ASCR 4%. NERSC underwent an interagency review in May 2005, the overall findings of which were favorable.

SciDAC received more than 350 letters of intent, of which 270 were encouraged to submit proposals. Of those, 230 proposals were submitted in 15 science areas.

ASCR expects to participate in GNEP by providing 100-teraflop computing resources for reactor design and for simulations for reprocessing, fuel fabrication, and waste disposal. The purpose of GNEP is to provide energy security, mitigate climate change, enable a hydrogen economy, increase the storage capacity of Yucca Mountain, and provide energy for nation building. Discussions indicate that 17 or more reactors will be built at 11 U.S. sites by 2015 to produce electricity at $35/MWh with a 93% online reliability and a 60-year service life.

Education Programs

George Seweryniak noted that education comprises about 4% of ASCR's budget. At the undergraduate level, the Historically Black Colleges and Universities (HBCU) Program seeks to develop and expand research and education relationships by directly funding HBCUs and operating a research alliance at ORNL. The program has doubled in the past four years. At the graduate level, the Computational Science Graduate Fellowship Program is designed to ensure an adequate supply of appropriately trained scientists and engineers to carry out DOE's mission in computational sciences. Costs are split between ASCR and NNSA; applications for the 15 fellowships increased from 337 last year to 410 this year. The program currently supports 64 students at 30 universities in 22 states. A fellowship conference will be held in June 2006; an external program review is scheduled in June 2006; a fellow will be sent to Lindau, Germany, for the Nobel winners lectures; and additional funding has been requested for FY07 to increase stipends and students. At the doctoral level, the Early Career Principal Investigator Program seeks to support research in applied mathematics, computer science, and high-performance networks performed by exceptionally talented scientists and engineers early in their careers. The average award is $100,000, and total funding is $2.6 million in FY06; 66 applications have been received for 15 awards. An external program review is scheduled for August 2006.

Discussion: Asked about the level of funding, Seweryniak replied that it could be doubled; there are qualified students out there. In actuality, the number of applications goes down each year as the call varies and becomes more focused. A letter of intent is now used to cut down applications from those who are not qualified.

COVs, Congressional Actions, and ACI

Yukiko Sekine reported that the first committee of visitors (COV) had evaluated the ASCR programs in Applied Mathematics, Computer Science, and Collaboratories. It found that the programs were effective and reasonably well-managed, the national-laboratory proposal process needed standardization and a more formal review process, more documentation was desired, and document management needs improvement. In response, ASCR has implemented a more formal laboratory proposal review process, implemented a new laboratory process, updated the Selection Statement Form, started using the SC Selection Statement Form for both grants and laboratory projects, and started using DocuShare to store and communicate forms. The second COV evaluated the facilities and network research program. It found the program highly effective and well-managed, especially the services provided to users. It stated that ASCR needs capacity and capability planning consistent with the size of its budget and programs.

The DOE High-End Computing Revitalization Act of 2004 requires the Secretary to carry out a program of research and development to develop high-end computing using leadership systems and to establish at least one software development center. The LCFs have been established at ORNL and ANL, and the center will be implemented as a SciDAC institute later this year. Another bill calls for the establishment of the Advanced Research Projects Agency–Energy (ARPA–E) to take on high-risk, high-payoff research to move advanced energy technologies into the marketplace faster. Funding of $3 billion is proposed for FY07. ASCR could be a major resource to ARPA–E for computing, simulation and modeling, and data management and analysis. And the ACI is designed to encourage American innovation and strengthen America's ability to compete in the global economy. It would double the federal commitment to the most critical basic research programs in the physical sciences and engineering over the next 10 years. Again, ASCR could be a major resource.

Closing Activities

Dahlburg announced tentative dates for upcoming meetings (Aug. 8-9 and November 1-2). In August, the Committee will hear an interim report from the facilities COV; and in November, it will hear interim and final reports from the ESnet Subcommittee and a final report from the Performance Measures Subcommittee. Committee members said that they would like to hear reports from the field like that from Rick Stevens, input on how the users are being served by the

ASCR user facilities, information on how ASCR has had an impact on nanotechnology, an explanation of ASCR's real challenges, and presentations summarizing the SciDAC projects.

Dahlburg thanked to support staff for a wonderful meeting and opened the floor to public comment. There being none, she adjourned the meeting at 3:00 p.m.

Action Items

Bell and Hack will head a subcommittee to address the charge on science-based performance metrics. An interim report is due July 30, 2006; a draft report is due October 30, 2006.

Wolff and Stechel will head a subcommittee to address the charge to examine the role and efficacy of networking and networking research within SC. An interim report is due October 30, 2006, and a draft report is due October 30, 2007.

Respectfully submitted,
Frederick M. O'Hara, Jr.
Recording Secretary
April 25, 2006