Minutes for the Advanced Scientific Computing Advisory Committee Meeting, April 5-6, 2004, Hilton Washington Embassy Row, Washington, D.C.

ASCAC members present:

- John W. D. Connolly, Vice Chair
- Jill P. Dahlburg (Monday morning and Tuesday only)
- David Galas
- Roscoe C. Giles
- James J. Hack
- Helene E. Kulsrud
- William A. Lester, Jr.
- Thomas A. Manteuffel
- Gregory J. McRae
- Karen R. Sollins
- Ellen B. Stechel
- Stephen Wolff
- Margaret H. Wright, Chair

ASCAC members absent:

- None

Recusals

- Stephen Wolff

Also participating:

- Steve Ashby, Director, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory
- Melea Baker, Office of Advanced Scientific Computing Research, USDOE
- James Decker, Principal Assistant Director, Office of Science, USDOE (Monday morning only)
- Daniel Hitchcock, Senior Technical Advisor, Office of Advanced Scientific Computing Research, USDOE
- Fred Johnson, Program Manager, ASCR, USDOE
- Gary Johnson; Program Manager; Advanced Computing Research Testbed; Mathematical, Information, and Computational Sciences Division; ASCR, USDOE
- Norman Kreisman, Office of Planning and Analysis, USDOE; ASCAC Designated Federal Officer
- Buff Miner, Program Manager, ASCR, USDOE
- Thomas Ndousse, Program Manager, ASCR, USDOE
- Frederick O'Hara, ASCAC Recording Secretary
- C. Edward Oliver, Associate Director, Office of Advanced Scientific Computing Research, USDOE
- Jeannie Robinson, Oak Ridge Institute for Science and Education
- Adam Rosenberg, AAAS Fellow, Senate Energy and Natural Resources Committee
- Rachel Samuel, Office of Management, Budget, and Evaluation, USDOE
- Rick Stevens, Director, Mathematics and Computer Science Division, Argonne National Laboratory
- Linda Twenty, Program Analyst, Office of Advanced Scientific Computing Research, USDOE
- John van Rosendale, Program Manager, ASCR, USDOE
- Thomas Zacharia, Director, Center for Computational Sciences, Oak Ridge National Laboratory

About 45 others were in attendance in the course of the two-day meeting.

Monday, April 5, 2004

Chairwoman **Margaret Wright** called the meeting to order at 8:30 a.m. She thanked Gloria Sulton of the office of Advisory Committee Management, Office of General Counsel, Department of Energy, for clarifying the rules under which the Advanced Scientific Computing Advisory Committee (ASCAC) operates. Sulton had stated that the "representational" advisory bodies bring together people who may have competing views and interests that may inure to the benefit of their employers or others to which they may be affiliated. Because the work of these committees is advisory, the Department wants to hear their opinions on the direction of programs, new initiatives, changes, or other recommendations to improve the work of the Department.

To allow the Department the benefit of their diverse views and expertise, one cannot put too tight a muzzle on the ability of committee members to be forthcoming on issues that may remotely impact their employer or other financial interests. At the same time, the Department wishes to protect the integrity of the Committee's deliberations and recommendations from undue criticism arising from information that could affect how other members may process/view or weigh the comments of a colleague. Sulton believes that a committee's discussions and deliberations can be open and robust if all its members take a moment to review their personal stake in the outcomes of the matters under review.

Sulton suggested that the following steps be taken at the beginning of each meeting:

1. Ask if there are any additions, corrections, or changes to the agenda as circulated.
2. Inquire if any member needs to recuse himself or herself on any item on the agenda.
3. Inquire if any member wishes to disclose a potential conflict of interest, as it may relate to an item on the agenda, based on a financial interest; affiliation as owner, partner, stockholder, board member, or officer in a particular entity; or such affiliation as a spouse or dependent child. If participation in a discussion and/or vote on the particular matter will not have a direct and predictable effect on that interest, then the member need not disqualify him/herself from speaking or voting.
4. Sometimes, it may be difficult to determine initially if a discussion will present a conflict (e.g., the discussion of a subcommittee report that has not been previously circulated or of an agenda topic that is vague). Let the members know that, at any time, a disclosure will be entertained if it becomes evident to the member that such a disclosure would be appropriate, indeed prudent, to dissipate after-the-fact critiques.

Wright asked if there were any changes to the agenda. There being none, she asked if any committee members would recuse themselves from any part of the meeting. Stephen Wolff responded affirmatively. She asked if there were any other conflicts of interest represented; there were no responses.

She reviewed the role of the Committee as stated on the Web page of the Committee:

The Advanced Scientific Computing Advisory Committee (ASCAC), established on August 12, 1999, provides valuable, independent advice to the Department of Energy on a variety of complex scientific and technical issues related to its Advanced Scientific Computing Research (ASCR) program.

ASCAC's recommendations include advice on long-range plans, priorities, and strategies to address more effectively the scientific aspects of advanced scientific computing, including the relationship of advanced scientific computing to other scientific disciplines, and maintaining appropriate balance among elements of the program.

Wright quoted the FY05 President's Congressional Budget:

"The [PART] assessment found that ASCR has developed a limited number of adequate performance measures. [The Program Assessment Rating Tool (PART) is an analytic tool used to identify performance goals and link them with the budget process.] However, the Office of Management and Budget (OMB) noted concerns regarding the collection and reporting of performance data. To address these concerns, ASCR will work with its Advisory Committee to develop research milestones for the long-term performance goals. ... OMB also found that the ASCR Advisory Committee is underutilized. ASCR will meaningfully engage the Advisory Committee in thorough assessments of research performance and in regularly revisiting the strategic priorities for the program."

She asked each member of the Committee to introduce himself or herself. Afterward, she introduced **James Decker** to give an overview of the activities of the Office of Science (SC).

Decker welcomed the new members of the Committee (Galas, Hack, and Manteuffel) and noted that they were joining the Committee at an exciting time in the Office. It recently produced *Facilities for the Future of Science: A Twenty-Year Outlook*, a facilities plan for SC for the next 20 years. He reviewed the process that Raymond Orbach, Director of SC, used in preparing that facilities plan. The Office has also produced the *Office of Science Strategic Plan*.

Prospective funding for the Department is down about $68 million from the previous year's appropriation. This is not bad because there were many one-year Congressional directives in the previous year. Overall, the Office's budget has a 2.3% increase. ASCR has about a $2 million increase or 1%. The SC budget request is up by $30 million over FY04. Among the increases in the FY05 budget request (some of which are offset by decreases elsewhere in the budget) are

- International Thermonuclear Experimental Reactor (ITER) negotiations and supporting R&D ($38 million, an increase of $30 million)
- Next-generation computational architecture and continued development of leadership-class computation ($38 million, sustaining a FY04 Congressional increase of $30 million)
- Nanoscale science, engineering, and technology ($211 million, an increase of $8 million)
- Hydrogen production, storage, and use ($49 million, an increase of $21 million)
- Genomics: Genomes to Life, including project engineering and design for the Protein Production and Tags Facility ($80 million, an increase of $9 million)
- Climate change science program ($134 million, an increase of $1 million)
- Scientific Discovery through Advanced Computing (SciDAC) ($64 million, an increase of $2 million)
- Workforce: laboratory science teachers professional development ($1.5 million, an increase of $0.5 million) and minority-serving-institution faculty sabbatical program ($0.5 million)
- R&D for new facilities: Rare Isotope Accelerator, BTeV (Fermilab), 12-GeV upgrade (Thomas Jefferson) to explore the fundamental nature of energy and matter ($15 million, an increase of $5 million)
- Linac Coherent Light Source R&D, project engineering and design and long-lead procurements ($54 million, an increase of $45 million)

Facility operations served about 19,000 users this past year and will be up to 95% of optimum in FY05, increasing from 92% during FY04.

On April 1, Orbach announced a restructuring of SC. The Office has huge responsibilities: 10

laboratories and contracts for the operating workforce, field offices, and 1000 employees (260 in Washington, D.C.). The structure of the field offices was designed during World War II with modifications by the National Nuclear Security Administration (NNSA) and others. To clarify roles and responsibilities, reorganization was undertaken. Operational Offices between field offices and HQ were eliminated, removing a thick layer of management. Field offices now report to a Chief Officer located in Chicago, assisted by Oak Ridge. At DOE Headquarters, a layer of management was removed in High Energy Physics and Nuclear Physics. The planning office and budget office were combined. A department was created for programs. Decker will be Principal Deputy and Deputy for Programs.

Kulsrud noted that the Committee had not seen the facilities report. Decker said copies for the Committee members would be arriving later in the day.

Dahlburg asked how the Office viewed the leadership-class machine. Decker replied that Orbach had decided to move forward aggressively. The Department is retaining an add-on in the FY05 request and believes that the United States will be able to field a leadership-class machine with the requested funds.

Giles asked whether, in moving the scientific agenda forward, the Office wanted the advisory committee members to be more aggressive. Decker replied that all of SC's advisory committees have to deal with COIs to one degree or another. If one has a conflict, one should not be part of the discussion. The Office does have good PART scores in general, and they are better than other agencies' in part because of how SC uses its advisory committees.

Wright commented that, a few weeks ago, the advisory committee chairs visited Congress, where they found great support for SC's programs and encouragement to inform Congress. But it is not clear how the advisory committees should support the budgetary process. Decker suggested that they work with their related professional societies, among other efforts. Manteuffel noted that Society for Industrial and Applied Mathematics (SIAM) is a strong advocate and a strong supporter of SC. Decker added that the American Chemical Society, American Institute of Physics (AIP), and others have become quite active and that an Energy Science Coalition has been formed.

Dahlburg agreed that the AIP encourages members to contact their members directly, but the reports from the advisory committees are very well read on the Hill. She asked if there would be any future charges from the Director for studies and reports. Oliver responded that he was sure that the Committee will be getting some new charges.

McRae asked what has to be done to get the strategic importance of computing recognized. Decker answered that Orbach is a great supporter and has increased support for the past two years. The more the advances possible from advanced computing can be demonstrated, the better the case that can be made for increased funding. The results from SciDAC and ITER computing will bring rewards. McRae followed up by asking how the Committee can help the Office move that case forward. Decker responded that a lot of effort has been put forward by the community in the past few years. He was not sure what else can be done.

Sollins asked how one trades off the long-term vs. short-term efforts. If one does not do the long-term, one does not have the pipeline for new ideas. Decker said that the Office tries to protect the long-term research, and it has been successful in supporting such long-term funding. Sollins responded that she did not see a lot of long-term computer science research being done. Decker pointed out that DOE's emphasis is on high-end computing. That is its niche. It will not compete with the National Science Foundation (NSF).

Connolly asked about making the Los Alamos Q machine available for large projects. Decker said that that strategy is being adopted for all the high-end resources of SC. Connolly noted that there are a lot of big machines that are not available to the research community. Decker agreed that the weapons-program computers are not available to the general researchers. Connolly said that he believed that the Q machine is being made available, and he would try to get more information on it.

Wright asked **Edward Oliver** to speak about the ASCR program.

A $30 million increase was made by Congress to the FY04 budget request, leading to a $25 million leadership-class computing solicitation and to $5 million for the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL). The leadership-class solicitation responses are being evaluated this month. The President's request was up $2 million over the Department's request, which included a continuation of the $30 million annual funding for leadership-class computing. Some 28 people are working on the ScaLeS (Science Case for Large-scale Simulation) report to make the science case for high-end computing. Volume 1 came out in September 2003, and Volume 2 will come out soon. The federal government has put together a multiagency task force on high-end computing [the High-End Computing Revitalization Task Force (HECRTF)] under the auspices of the Office of Science and Technology Policy (OSTP). A draft of their report is under review now. In FY04, a $7 million addition was made over the FY03 amount for next-generation architecture (including the Cray X-1 evaluation and other activities).

The networking solicitation is on the street and will be described later in this meeting. This current fiscal year is going to be busy. The FY05 request for Mathematical, Information, and Computational Sciences (MICS) Division includes funding of $8.5 million for multiscale mathematics. It also includes $1.5 to 2 million for partnerships with the fusion program. In March, a Committee of Visitors met with ASCR staff; Wolff's report on that activity, which was presented later in the meeting, is described below. This is the third year of the young investigator program in mathematics, computer science, and networking.

The Division's current mission statement is to deliver forefront computational and networking capabilities to scientists nationwide that will enable them to extend the frontiers of science, answering critical questions that range from the function of living cells to the power of fusion energy. With an annual budget of $200 million, ASCR parallels the other SC offices: Basic Energy Sciences ($1 billion), Biological and Environmental Research ($0.5 billion), High-Energy Physics ($0.75 billion), Nuclear Physics ($350 million), and Fusion Energy Sciences ($250 million).

Oliver reviewed the staffing situation of ASCR. The position of MICS director has been vacant for 7 years. Chuck Romine is on detail to OSTP as of March 15 with Gary Johnson filling in as manager of applied mathematics. Several other vacancies need to be filled.

Oliver graphically analyzed the ASCR program, enumerating the activities that contributed to the enablement of science through computing, the next-generation architecture initiative, a high-performance network environment for science, and core capabilities [high-performance production computing and the Energy Sciences Network (ESnet)].

Scientific Discovery Through Advanced Computation (SciDAC), an SC-wide program that is described at length in a presentation later in the meeting, was initiated in FY01 under Thom Dunning. It is an ambitious program designed to bring the power of terascale computing to science. It started with $60 million for research. The next year's ASCR budget was cut, but SciDAC was kept at $60 million.

The Next Generation Computer Architecture (NGA) initiative is funded at $37.8 million for

FY04. It will lead to the development of potential leadership-class machines. In FY05, the NGA will

- Complete the evaluation of the Cray X1 at Oak Ridge National Laboratory (ORNL);
- Continue work with the scientific application community on performance evaluation in partnership with the Defense Advanced Research Projects Agency (DARPA), NNSA, and the National Security Agency (NSA);
- Continue the academic, laboratory, and vendor research program on scalable technologies for future generations of operating systems and runtime environments ($4 million);
- Continue the strong partnership with the DARPA High-Productivity Computing Systems program; and
- Initiate development of a leadership-class computer through peer reviewed competition.

Multiscale mathematics is described at length later in this meeting. It is important because current models assume separation of length/time scales, a property that fails to hold in many critical application areas.

The ASCR/MICS budget has shown a modest increase from FY04 to the FY05 budget request, including the addition of the Multiscale Math program. Wolff commented that something had to be reduced in that budget and asked what it was. Oliver said that a lot of people have a stake in this budget and that he would like to hear suggestions from this committee of where to cut. Sollins said that the alternative is to ask for more money. Oliver said that they were always asking for more money. One can get increases in the presidential budget or one can get them from Congress during budget deliberations. The Office goes with the President. Wright noted that the opportunity exists to go to Congress to increase the budget. Oliver noted that one can lose money in the Congressional deliberations, also.

The new SC program, Innovative and Novel Computational Impact on Theory and Experiment (INCITE), makes 10% of NERSC's computing capability available to the scientific community for high-impact science. This year, 53 proposals were received, asking for 130,508,660 CPU hours when only 4,500,000 CPU hours were available. There were three winners (one was NSF-funded) in astrophysics, turbulence, and photosynthesis.

Software is an important component of the timeline for compute facilities and testbeds. Important issues include complex and multiscale systems, petascale systems, and transforming petabytes of data into knowledge. Oliver's analysis of the research included the term "long-term support," and Wolff questioned what that was. Oliver said it is long-term support for software produced by SciDAC and other programs. People at the laboratories could probably provide that support. Wright asked how this works at facilities. Oliver responded that that is part of their budgets. Wright suggested that there could then be an operating budget for software, also. Fred Johnson said that this is a well-recognized problem and that it will be addressed in the HECRTF report.

Oliver listed a few strategic issues for ASCR:

- Providing high-performance computing and network facilities (cycles and good networking);
- Effective partnerships with applications scientists in all of the Offices in SC;
- Accelerating transition from research to applications; and
- Long-term support of software.

ASCR is well coordinated with other federal information-technology (IT) research, including direct DOE involvement with the coordinating groups of OSTP's interagency working group, specifically the groups on High-End Computing and Communication; Large-Scale Networking; Human Computer Interface and Information Management; and Social, Economic, and Workforce

Implications of IT and IT Workforce Development. ASCR staff cochair the interagency High End Computing Revitalization Task Force (HECRTF), the interagency Large Scale Networking Coordination Group, and all three teams: Joint Engineering (JET), Network Research (NRT), and Middleware and Grid Infrastructure Coordination (MAGIC). They also participate in a lot of activities with NNSA, DoD-DDR&E (Department of Defense Directorate of Defense Research and Engineering), NSA, and DARPA. They have participated in a number of workshops and reports, including the workshop on the Science Case for Large Scale Simulation, which produced the ScaLeS report (Volume 1), and the Workshop on the Road Map for the Revitalization of High-End Computing. A number of reports came out at the right time for use in the budget process this year.

Connolly asked what was new with the memorandum of understanding (MOU) with DARPA. Daniel Hitchcock replied that it puts in place a framework for working with other agencies on high-end computing, especially performance-evaluation benchmarking. A plan is produced each August on high-end-computing development.

Manteuffel asked what would be received in return for the $25 million dedicated to leadership-class computing. Kreisman replied that there were other SC offices that could use lots of computational resources. High-performance computing is extremely important, but ASCR's whole budget is just half the cost of the machine in Japan. Gary Johnson noted that the solicitation was issued to the ten national laboratories; the solicitation closed the previous Friday. Four proposals were received. It is expected that the $35-million funding would continue for 5 years. Kulsrud asked when the money would be available. Gary Johnson said that it would be FY04 money. Giles asked where this solicitation came from and what role this Committee played in its crafting. Gary Johnson responded that the work of the Committee and several studies were used in crafting the solicitation and are cited in it.

A break was declared at 10:36 a.m.

Wright called the meeting back into session at 11:08 a.m., and Adam Rosenberg of the Senate Energy and Natural Resources Committee mentioned a bill (S 2176, sponsored by Bingaman, Democrat of New Mexico, and Alexander, Republican of Tennessee) that authorizes $150 million for high-end computing in SC: $100 million per year for facilities and a high-end-computing software-development center and extra funds for ASCR.

Wright returned to the Committee's question period. Sollins asked what was going on about networking. Thomas Ndousse answered that the Office was holding a series of workshops as well as seeking comments from the computing community. Part of the response called for an effort addressing networking issues. The staff is identifying specific requirements: computational biology, climate, and moving data from Europe to the United States. To optimize resources, the community is putting together a nationwide networking testbed for new transport technologies and network capability on demand. When the testbed is in place, a solicitation will be put out for projects to test it, and the technologies will be rolled into ESnet.

Sollins asked how one decides the balance of where to put the dollars (e.g., testbed development vs. off-the-shelf purchases). Oliver replied that one has facilities and R&D. Usually, researchers want more cycles today. It is a zero sum game. Perhaps the Office should get out of networking research and put the money into ESnet. Orbach ultimately makes those decisions.

Wright asked how the Office could regularly revisit the Strategic Plan with the Advisory Committee, as called for in the budget request. Oliver said that he puts it on Orbach's radar once a week, but Orbach is a busy person.

Kulsrud noted that there are no reports from the national-laboratory directors and asked if the Committee could hear from them about their needs and what has not worked that needs to be fixed. Oliver replied that NERSC and ESnet are being run at Berkeley, and they want to keep doing it. The national laboratories all have institutional plans, but what they do in applied mathematics, computer science, and networking tracks the ASCR budget. Those institutional plans influence ASCR's plans for the out-years. Kulsrud asked if the national-laboratories' computer centers are oversubscribed and what the status of their machines was. Oliver responded that NERSC hopes to get money for a cluster to offload smaller users to make room for large users. The national laboratories execute the program that we write the checks for.

Sollins asked if, in evaluating proposals, the Office considers facilities and management at a location. Oliver answered, yes; the possibility of executing the work proposed is very important. Once a year, the Office evaluates each entire national laboratory. The resulting rating of a laboratory translates into an operating fee for the laboratory's contractor.

Wright asked **Gregory McRae** to present the report of the High-End Computing Subcommittee of ASCAC.

ASCAC has consistently been focused on high-end computing issues within ASCR, but many issues apply across SC. The High-End Computing Subcommittee has now completed a report the purpose of which is to take a "big-picture" view of computing across SC; the draft report has been distributed to the Committee. The key message in the report is that high-performance computing and computational science are, and should be viewed as, key enabling technologies in meeting the missions of the Department. This is true across all offices of SC as well as across the federal government. While Acting Director of SC, James Decker issued a charge to ASCAC calling for

- A composite panel to assess the quality of high-performance computational needs and capabilities throughout SC;
- Benchmarking the quality of these capabilities by comparison with similar facilities, both domestic and international;
- An investigation of (1) the impact and effectiveness of interactions and resource-sharing among SC high-performance computational and network facilities, (2) the level and adequacy of funding provided by all SC programs for high-performance computing and networking facilities compared with other needs, and (3) the effectiveness of the current distribution of high-performance computational and networking resources across SC;
- A projection of the roles of these facilities and/or their distribution during the next 3 to 5 years so that SC programs can meet their high-performance computational needs and maintain their national and international scientific leadership; and
- Useful metrics to measure progress and guide investment decisions in the area of computing and networking.

Previous ASCAC subcommittees have looked at facilities, the role of large facilities, nanoscience, fusion simulation, biological research and high-end computing, the science case for large-scale simulation (ScaLeS), and strategic planning for ASCR. Two important reports that have come out of these efforts are

- Science Case for Large Scale Simulation, June 2003 (www.pnl.gov/scales/) and
- ASCR Strategic Planning Workshop (www.fp-mcs.anl.gov/ascr-july03spw).

The ScaLeS report brought together the recommendations of 300 scientists and engineers and focused on the scientific and engineering accomplishments achievable through modeling, simulation, and high-end computing..

The ASCAC High-End Computing Subcommittee came to the following conclusions:

1. The current SC high-end computing (HEC) capabilities are among the best worldwide but are eclipsed by the Earth Simulator in Japan.
2. Current and expected near-term SC resources are far from adequate to meet current and near-term needs.
3. The SciDAC program is world-class, but it is underfunded.
4. Opportunities abound in every mission area of SC for mathematical modeling and computational simulation.

The High-End Computing Subcommittee recommends that SC should

1. Invest serious new money to return the United States to leadership in high-end computing: $300 million in new funding for leadership-class computing and $200 million in new funding for capacity computing.
2. Coordinate SC with other offices in science, high-end computing, and applied mathematics.
3. Manage its advanced computing resources as a single, coordinated facility.
4. Guide investment decisions by the importance of the associated science along with the opportunities for cross-fertilization and common usage among scientific disciplines, mathematics, and computer science.

McRae pointed out that the Lax report (issued 22 years ago) had many of the same recommendations. The President’s Information Technology Advisory Committee (PITAC) report shows how important science problems have been solved by investments in HEC. The payback of these advances has been extraordinary.

Wright added that there was nothing startling about the messages put forward by the report. In fact, several of its conclusions and recommendations are exactly those of the ASCAC Large Facilities Subcommittee, chaired by Helene Kulsrud last year. Wright asked for discussion of the report.

Connolly commended the Subcommittee for its work and pointed out that the response to the Lax report came from Congress.

Stechel commented that the report points out that the conclusions and recommendations of this study are the same as those of prior reports and it is time to act on those recommendations. Also, multiscale mathematics underlies most of the treatment of complex problems today. Wright added that it is important to recognize that a big piece of hardware is useless without the mathematics, programming, and science to use that hardware. This report says that $300 million is needed and $25 million is being requested.

Kreisman commented that one pays for a new machine even if one does not buy it. What needs to be asked is, “What is the cost if we do not follow those recommendations?”

Kulsrud said that she had tried to construct a cost envelope and found that $150 million would not provide adequate resources. She used numbers that she could not talk about to cost systems that had been successful. In this way, she arrived at the $300 million figure as one that would likely be effective.

Wright asked for a vote to accept this draft report. Galas said that a substantial effort on computational biology was being made by the National Research Council (NRC). An appreciation is building that HEC is important to the scientific enterprise. Wright agreed and said that the report to needs to mention the other efforts and reports that are being produced. Stechel said that the main point to make is that all these reports are coming to the same conclusion.

Lester moved to accept the report, and Sollins seconded. The motion passed unanimously. A break for lunch was declared at 12:10 p.m.

The meeting was called back into session at 1:45 p.m. Karen Sollins asked **Stephen Wolff** to present the report of the Committee of Visitors (COV). The charge to the COV, the first-ever COV held for ASCR, was to evaluate how effectively the program adheres to the accepted policies, procedures, and management of major program elements and to provide an assessment of the processes used to solicit, review, recommend, and document proposal actions and monitor active projects and programs. Particular attention was paid to commenting on future directions proposed by ASCR management and on opportunities that might not have been considered. The key questions to be addressed were

- Are the best people in the associated areas, and the best proposals, being funded by ASCR?
- Have any proposals been funded that did not receive good reviews? If so, why?
- Are the most knowledgeable and informed reviewers being chosen?
- Have competition and peer review been used appropriately to guide major research expenditures?

The COV was not able to complete its task in full detail in the 1.75 days available, but it reached several important overall conclusions. The programs reviewed were research in applied mathematics, research in computer science, and research in collaboratories. The major findings were

- The programs are generally effective and reasonably well managed, given their complexity, diversity, and scope.
- The program officers are dedicated, competent, and know their portfolios and communities.
- Each program has achieved significant successes.

The COV found no decisions that were obviously incorrect or unjustified, and no instances where a program was skewed or where there were marked deviations from the normal standards of peer review.

In terms of the review process, the reviewers were clearly knowledgeable, although the reviewer pool seemed small. The COV was sometimes uncertain of the representativeness of the reviewer pool, and no reviewer pool data or statistics were kept. Occasionally, panels seemed too small to ensure thorough representation of all subdisciplines, or the panel's size was incommensurate with the requested funding. Program managers exercise considerable appropriate discretion and methods of proposal evaluation. Significant differences occur in the handling of and review techniques employed among programs and between national-laboratory and university proposals.

In terms of folders, there was considerable unevenness in detail, format, content, and organization of the documentation. Large-team awards were documented in individual folders, one per co-investigator institution; as a result, it was difficult to obtain an overview of the team effort. Significant differences existed between awarded and declined proposals and between the handling of national-laboratory and university proposals. These process differences make comparisons difficult. In addition, the geographic and demographic distributions for awards and declinations are very difficult to obtain because of the way the folders were organized. In various cases, the material in the folders was insufficient to evaluate the complete train of events leading to the final decision. The information was often available but was located elsewhere; the folders system does not support comprehensive and unified program management. A significant amount of data requested by the COV had to be drawn from program officers' private file space. It was the sense of the COV that the folders were designed more for fiscal management than for program management.

ASCR is developing guidelines for program officers, and the COV commends this effort.

The COV notes that laboratory and university processes, including decision processes, need not be identical. However, both folder types should contain enough information to detail the initial funding decision and to track progress for renewals. It recommends that ASCR develop a more comprehensive and consistent approach to program documentation, one oriented toward program management.

The COV was split into three teams to look at the program areas of applied mathematics, computer science, and collaboratories.

Research in applied mathematics was found to have excellent management, good documentation, and clear decision memos based on analysis of reviewer comments. It employed distinguished reviewers, and the size and representativeness of the reviewer panels were adequate. The intrinsic quality of funded proposals was high. Some projects with strong initial reviews were renewed with less enthusiastic reviews. At least one national-laboratory renewal was denied. The program officer is to be commended for the new initiative in multiscale mathematics. Although much funding came from SciDAC, it was unclear how that funding was acquired. The impacts of this program were evident in the computational modeling of combustion, computational fluid dynamics, laser-plasma simulation, shock-wave theory, and prizes awarded to the principal investigators.

Research in computer science was found to have folder information that was insufficient to assess the solicitation, review, and award. The COV spent significant time questioning the program officer, whose patience and openness was very much appreciated by the COV. Additional data had to be obtained from the program officer's personal file space. Research in computer science was found to have folders that did not allow the COV to assess how each proposal was ranked relative to the others evaluated. However, there were no clearly incorrect decisions. This is a newer discipline, so the process is somewhat different. The principal investigators (PIs), co-PIs, and reviewers seemed to be drawn from a very small pool, and the efforts to expand that pool have been disappointing. Mail reviews are asynchronous; as a result, the folder process was broken, and the program was being managed by another method. The following recommendations were made to improve the processes used in computer science:

- Consider synchronous mail reviews to allow comparative evaluations, or maybe the mail reviews should be collected and evaluated concurrently.
- Persist in attempts to widen the reviewer pool.
- Implement and formalize a consistent and uniform documentation process.
- ASCR should establish guidelines on the number of reviewers in relation to the magnitude of opportunity.

The impacts of this program were difficult to evaluate because of its limited mission and the small community served. The impacts of this program are not comparable to those of the NSF program; it supports specific types of research. Some impacts include the development of the Message-Passing Interface message-passing model and reference implementation, toolkits (e.g., SCAR for managing Linux clusters), and the Fernbach award to Jack Dongarra.

Research in collaboratories was found to be well-managed and innovative, with significant impact inside and outside DOE. The review processes and documentation were very good. The COV was impressed by the quantity of papers and reports from workshops that had resulted. It is not clear that the program is open to researchers without existing ties to national laboratories. There is also a concern about the stability of SciDAC funding, which is a significant fraction of the program's budget. The impacts of this program reflected accomplishments that are more than reasonable, given the age of the program: the Access Grid (>150 worldwide), the GLOBUS middleware toolkit, and the Electronic Notebook.

The general programmatic findings and recommendations of the COV include a concern over the talent pool and “openness.” ASCR is to be commended for its Early Career PI Program, which has produced 20 new PIs in 2 years. Ways to increase the pool of investigators include ASCR’s working with the Computing Research Association (CRA) to publicize its activities via *Computing Research News* and the Forsythe list, expanding workshops at open conferences, and establishing a “Research Corner” at the Supercomputing conferences. There is room for greater interaction between national-laboratory and university researchers. DOE should use collaboration technology to minimize travel for its review panels. And ASCR should assist in strengthening the relations between the national laboratories and academe.

In terms of strategic planning, SciDAC is a valuable basic function that should be sustained. However, a strategic review is needed. There is a concern over openness and follow-through (software maintenance). The COV commends ASCR for its advocacy for leadership-class computing systems. However, ASCR should pay more attention to professional development and growing the community. A long-term focus is needed. Multiscale mathematics is essential to support real applications. Planning needs to consider all possible high-end computing architectures as well as grids. ASCR should facilitate multiagency approaches for efficiencies and market impact. High-end-computing communities are small; to gain the attention of the commercial market, they should work together.

In terms of the COV process itself, regular visits will benefit DOE. The COV is grateful for the accessibility of key personnel during the meeting. The dinner meeting was very productive. Additional material should be distributed well before the meeting, and staff should invest more time in preparation (e.g., preparing demographics by state, type of institution, diversity, and success rate). Additional material is needed at the meeting. Logistical improvements that could be made in include

- Standardizing presentations
- Addressing COV charge issues in the talks
- Expediting physical entry to the site
- Providing an on-site dirty net with printers
- Providing enough paper handouts

Sollins said that the NSF deliberated whether documents should be paper or electronic. Their visitations are now done electronically, including access to all requests for proposals (RFPs), proposals, evaluations, and award information and data. Oliver said that ASCR agrees 100%. It has an electronic portfolio-management environment (EPME). However, ASCR’s budget was not sufficient to implement EPME; all of DOE is moving in that direction. Sollins cautioned that one has to be careful to make sure the system does what the program manager needs and what the bean counters need.

Sollins asked for a motion to accept the report; Kulsrud asked to postpone the vote until the next day to give Committee members a chance to read it.

Giles asked if there was any priority among the recommendations and whether there will be another COV in the future. If so, that COV will want to see if the recommendations were addressed. Oliver responded, yes; there will be a COV each year. Each will cover half the programs, and the topics will be revisited every other year.

Wolff suggested that network research’s being reviewed along with the facilities program does a disservice to both. They are too large to do together.

Wright said that the COV was an extraordinarily distinguished group, and she thanked the members of the COV and the staff that assisted them. Oliver noted that ASCR's relationships with other offices of SC are reviewed by COVs.

Wright introduced **Michael Strayer** to report on the SciDAC PI Meeting and SciDAC plans. SciDAC is DOE's pilot program in computational science. It is strongly interdisciplinary, involves all of the offices of SC, and reports to Orbach. Its purpose is to create a scientific culture to use high-end computers effectively. The goal of advancing scientific discovery is central to the SC research missions. SciDAC is in its third year of five. The third annual PI meeting was held March 22-24, 2004, in Charleston, S.C., and was very well attended. There were 23 invited talks, 2 poster sessions, and 1 panel discussion. The two plenary talks were exceptional: System and Application Performance at Extreme-Scale by Adolfy Hoisie (LANL) and The Grid: Essential Infrastructure for DOE Science by Ian Foster of Argonne National Laboratory (ANL).

The talk on Lattice Quantum Chromodynamics (QCD) from the Nuclear Physics Perspective by John Negele (MIT) excited a lot of discussion across the program. All the other talks were great, as well, and reflected outstanding science.

One of the highlights of the SciDAC program is the outstanding performance of the three Applied Mathematics integrated software infrastructure centers (ISICs):

- Applied Partial Differential Equation Center (APDEC) at Lawrence Berkeley National Laboratory (LBNL) and Lawrence Livermore National Laboratory (LLNL)
- Terascale Simulation Tools and Technologies (TSTT) at Brookhaven National Laboratory (BNL), ANL, LLNL, ORNL, Pacific Northwest National Laboratory (PNNL), and Sandia National Laboratories (SNL)
- Terascale Optimal Partial-Differential-Equation (PDE) Solvers (TOPS) at Columbia University, ANL, LBNL, and LLNL

APDEC has developed and deployed an adaptive-mesh analysis technique that has been applied to hydrocarbon flames, centimeter-scale turbulence in a runaway thermonuclear reaction in a supernova, and cell modeling. An example from the fusion community on the application of this adaptive mesh is embodied in detailed 3-D adaptive-mesh-refinement (AMR) simulations of pellet injection into a tokamak using the magnetohydrodynamic equations. In this treatment, the pellet ablates with an analytic model of a high or low field.

The TSTT has developed an unstructured mesh that allows the analysis of long-term beam heating of the PEP-II (Positron Electron Project) beam. Restructuring of the mesh and the heating leads to an understanding of the behavior of microbe flocs in oxygen-rich environments, a complex process, but modeled to great detail. Optimal mesh generation and adaptive methods have reduced error in climate applications.

The TOPS center is a leading example of the success of the centers. Its software has been employed in two Bell Prizes, in 1999 and 2003. Many researchers now use this TOPS software, and a new software initiative is planned.

The SciDAC PI meeting was very successful. The program is funded through FY05 and will be recompeted in FY06. Discussions are ongoing about the future of SciDAC with Orbach and Oliver and with the offices of Biological and Environmental Research (BER), Basic Energy Sciences (BES), Fusion Energy Sciences (FES), High-Energy Physics (HEP), and Nuclear Physics (NP). At the SciDAC PI meeting, a call was made for community input on how SciDAC should go forward. The present progress needs to be preserved, and new initiatives need to be planned, even during difficult budget years.

Stechel asked him to comment on the panel at the meeting that considered if SciDAC was changing how science was done. Strayer replied that that was an engaging panel of Bader, Mezzacappa, and Colella. They had three distinct points of view: Tony Mezzacappa said that SciDAC is like World Cup soccer, like a religion. Dave Bader cited pragmatics: the program needs investments and broad funding support. Phil Colella asked how one manages SciDAC. The discussion did not decide anything but it laid issues on the table.

Connolly asked if there was any discussion of Grid computing. Strayer said that there was in Ian Foster's talk. The NSF is leading Grid computing. Everyone would like to do computing on the Grid. Virginia Tech has put together the world's number-three supercomputer out of Macintoshes. Connolly noted that they still need a closely coupled common system; the SciDAC program can contribute middleware. One should not rely on just one project.

Kulsrud asked if the PIs have enough computing time available. Strayer replied, no; the amount available is probably off by more than a factor of 2. It is deploying science application codes. SciDAC alone will soon outstrip SC's computing resources. Kulsrud asked if they had any other problems. Strayer said, no; this was a very positive meeting.

Giles asked if the need for software distribution and support had arisen. Strayer replied that no one raised that as an issue. The PIs have already undertaken an extensive outreach program. Buff Miner commented that the program gets daily requests from PIs asking for computational techniques to put in their codes. That will be a new paradigm and a need for support.

McRae asked how much of this is getting out into industry. Strayer replied that he did not know.

Lester asked if anyone else was doing work like this. Strayer said that the team led by Robert Harrison, working on the NWChem computational chemistry package, uses radical and innovative approaches similar to those of SciDAC.

Kulsrud asked when more SciDAC will be put into fusion research. Oliver responded, in 2005; it would be a partnership with FES, depending on the availability of funding from FES. It may or may not be called SciDAC.

Wright asked how to train the next generation of researchers to work in teams. Human resources is not part of the SciDAC charter, but SciDAC seems to be a great way to support that objective. Strayer said that the Computational Science Fellowship Program has had great success in doing that. Oliver said that the Department used to support education from kindergarten to postdoctoral programs. It was great. Then the Department was instructed to get out of education and to limit activities to postdocs. If funding for that fellowship program were to be increased, the money would have to come from somewhere.

Wright brought the discussion back to talking about SciDAC and its involvement of graduate students and postdoctoral fellows. Oliver suggested putting in words requiring the participation and support of graduate students and postdoctoral fellows. Fred Johnson noted that, certainly within the ISIC Program, there have been outreach components (e.g., programming camps for researchers) for training and transfer.

McRae stated that the SciDAC Program and Computational Science Graduate Fellowship Program are true jewels, transforming how computational science is done. He expressed an interest in seeing these programs preserved and grown.

Wright asked why SciDAC is not getting more funding if it is so good. Strayer replied that everything Oliver said about ASCR applies to SciDAC, too. He described a cooperation and collaboration with the Executive Office of the President to put forward a budget that is then reviewed by Congressional committees, shaping what next-generation science should be. That is not a perfect process. Sometimes one wins, and sometimes one loses. All of the programs in

ASCR are wonderful, but the budget is flat. One cannot go out and print new money. One has to persevere over the long haul. Oliver commented that Orbach wants Strayer to have an equal say in SC's FY06 budget with all the assistant directors. Orbach will make the final decision on funding within SC.

Strayer noted that the original SciDAC had hardware in it. SciDAC will need leadership-class computing, which is lacking now.

Connolly said that sometimes good science is not good enough. He asked if any industrial spinoffs and contributions to the economy could be identified. Strayer replied that Oliver has always emphasized the benefits of advanced computing. Orbach has gone out to industry. This is a major thrust in going forward with HEC.

Stechel asked if SciDAC was changing the culture of scientific simulation. Strayer answered that Mezzacappa's researchers did not think of an integrated computer-science team until they joined SciDAC. His new simulation moves terabytes of data around, uses a visualization wall, and has discovered a new type of shockwave. This all came out of the team research. That statement and paradigm resonated with the PIs at the meeting. Stechel noted that the PIs are already converts and asked if the message was being taken outside the SciDAC community and if there are barriers in SC to the self-assembly of these multidisciplinary teams without SciDAC funding. Oliver responded that these teams can certainly self-assemble, but it is very expensive. Stechel said that it sounded as though there were some very high barriers.

Wright noted that it is important that all SC offices do something about this. She declared a break at 3:31 p.m.

The meeting was called back into session at 4:01 p.m. with the introduction of **Daniel Hitchcock** to talk about the measurement of software "effectiveness."

OMB wanted an efficiency measure for SciDAC's effect on computing software. In the budget request for SciDAC, it says that SciDAC will improve computational science capabilities by achieving an average annual percentage increase in the computational effectiveness (either by simulating the same problem in less time or simulating a larger problem in the same time) of a subset of application codes within SciDAC. The target for FY05 is to increase efficiency more than 50%, as determined by running a selected group of benchmark codes.

OMB wanted at least one year of baseline data, so the following four codes were selected:

- AGILE-BOLTZTRAN: time to solution for a baseline run of size 8 x 12 x 100 (8 angles, 12 energy groups, and a spatial resolution of 100)
- BEAMBEAM3D: total computational effort achievable in the simulation, represented by the product of the number of particle-turns, the number of collisions per turn, and the square of the number of longitudinal slices
- AMR Combustion: the product of the number of chemical reactions simulated times the number of effective zones, a measure of the effective resolution achieved throughout the computational domain
- NIMROD: the time required for simulation of the nonlinear tearing evolution over the transport timescale

These applications cover a broad range of SC offices. These codes were run on the same machine (no Moore's Law effects); all changes were produced purely by changes in algorithmic software. The results were surprisingly large improvements:

- AGILE-BOLTZTRAN: Simulation time in 2002: 4 weeks. Simulation time in 2003: 4 days. Improvement in effectiveness: 700%.
- BEAMBEAM3D: Maximum effective problem size in 2002: 10^{11}. Maximum effective

problem size in 2003: 10^{12}. Improvement in effectiveness: 1000%.

- AMR Combustion: Reactions x zones in 2002: 3.5×10^8. Reactions x zones in 2003: 3.7×10^{10}. Improvement in effectiveness: 10,571%.
- NIMROD (a fairly mature fusion code): Simulation time in 2002: 5.5 days. Simulation time in 2003: 2 days. Improvement in effectiveness: 275%.

OMB's staff was initially skeptical about these results, but the changes were eventually attributed to an anomaly for the first year.

The codes proposed for use in FY04:

- Global Climate CCSM
- Shell Model Monte Carlo (SMMC)
- Virginia Hydrodynamics 1 (VH-1)
- The MILC Code (QCD)
- Parallel R-Matrix with Pseudostates

The communities that are running these codes have agreed to do comparisons between the beginning and end of the year.

CCSM is an obvious choice for a reasonable code to benchmark. VH-1 is a 3-D, piecewise parabolic method (PPM) hydrodynamics code capable of simulating three-dimensional turbulent stellar flows with high accuracy and little dissipation (in other words, it is able to track features very well over long simulations). The MILC code is publicly available, 15 years old, and fairly portable. Parallel R-Matrix supports plasma science, such as fusion energy, atmospheric research, and technology processing of semiconductors.

Giles cautioned that one wants to be careful that mature, highly optimized codes do not give up any more efficiencies. Hitchcock responded that the program staff thought that 50% was an aggressive but realistic goal. If those expectations are not met, that situation will have to be dealt with. This is not a perfect measure, but it is better than everything else out there.

Sollins pointed out that being able to do something that was not able to be done before is not amenable to such measurements. Hitchcock replied that such advances have to be measured by technology transfers. The hope is that multiscale mathematics will open up things that could not be done before. Other metrics proposed included tracking unexpected discoveries; these metrics were obviously bad ideas. Whether efficiency increases by 50% or not tells us something about the flexibility of algorithm space.

Kulsrud noted that this method would use the same machines, but going to new architectures would be a significant advance. Hitchcock responded that one has to hold something constant. Here, architecture is held constant, and only algorithms are varied.

Connolly asked if some of the efficiency was gained by going to a higher level of parallelism. Hitchcock responded that some was and that some was gained by going to sparse algorithms. Connolly asked whether ISIC or MILC people do the work. Hitchcock replied that the Office is trying to figure out what is the best way to go. The 50% value is the average, so some could do worse than others.

McRae asked if they would track the same codes next year. Hitchcock replied, maybe and maybe not. Some codes will no longer be of interest, and their use will be supplanted. Other codes will take their places in the comparisons. McRae asked if they had thought about turning the experiment around. Hitchcock responded that this turns out to be very expensive for the scientists. McRae pointed out that there are very few data and suggested that the staff might think of investing in these metrics and measurement results.

Wright asked if they were profiling these codes. Hitchcock answered that they had done crude measurements. *Why* the code sped up was not looked at. Wright commented that so much more could and should be done. Hitchcock noted that such an investigation is a much more complex and expensive activity. Wright pointed out that one would like to see *why* things work. Is it message passing, linear algebra, matrix computation, next generation, or what? Hitchcock agreed that that would be wonderful to do; however, scope creep can be overwhelming. Strayer pointed out that there is a SciDAC PERC (Performance Evaluation Research Center) that has developed a number of evaluation tools. Many application codes have these tools built in and tell the user exactly what was being asked for here. The NNSA has used these types of tools to predict what will happen with different improvements on specific platforms. Manteuffel stated that, once these tools are put in place, they will make the case for advanced scientific computing. John van Rosendale said that, for many years, it has been claimed that algorithmic advances have been outstripping Moore's Law and that he had been skeptical about that statement. He pointed out that this approach offered an opportunity to answer this question and urged Hitchcock to pursue it.

Wright asked if one new code was going to be used each year. Hitchcock replied that the staff would like to hear from the Committee about what would be a good way to do it. If the Committee wanted a second year of data on these same codes, it would be prepared. Wright asked if anyone would like to formulate a recommendation. Manteuffel said that it needs some thought.

McRae said that these were great data, that they would have a tremendous impact, and that the staff should think more broadly about how these data could be used to support the program. .

Giles pointed out that whatever codes were used as benchmarks should be ones that are in use by the research community.

Wright asked what the changes were that produced the improvements. Hitchcock replied that that data had not been collected but that they could be; the data are available only at the mercy of the researchers. A statement is needed about the proposed codes stating that they are fair; any secondary recommendations that the Committee would care to make would also be appreciated.

Lester asked if these target codes would change each year. Wright responded that they may. It depends partly on what advice the Committee gives. Kulsrud noted that, over 10 years, one would have to change hardware. Hitchcock pointed out that one could change machines and just show the change from year beginning to year end. Connolly noted that such a method would benefit from the poor tuning of new systems.

Hack stated that it would be wonderful if climate codes could be speeded up by 50% and asked how these codes were picked. Hitchcock said that a rather random sample of codes being used in SciDAC were selected. Strayer noted that climate is a high priority in SciDAC, so the climate scientists picked a new code coming out. ITER is very important and has a new, immature, interdisciplinary code that came out of England. The QCD code is more mature; they have a collaboration that could investigate how to improve the mathematics of that code. VH-1 has been around a while and has been tuned for a particular machine. It has not been vetted by the SciDAC process, but it will be seen what they can do to improve it. Hack commented that it sounded as though these codes were being selecting for their scientific importance.

Wright opened the floor to general discussion by the Committee. There being none, she opened the floor to public comment.

Rick Stevens asked McRae about the strategies for acquiring large-scale facilities by using procurements vs. construction projects. The question is partly determined by the projects' being unique. McRae replied that the High-End Computing Subcommittee did not consider that aspect.

Kulsrud replied that the Large Facilities Subcommittee considered whether a building was being built or not. Giles added that the Large Facilities Subcommittee did not take into consideration federal regulations relating to construction. Stevens stated that the choice determines how one pays for the project (over longer periods) and how many management and approval processes have to be imposed. Procurements are much simpler than construction projects. But procurement of items worth tens of millions of dollars may not be the way to go. Kulsrud stated that the Large Facilities Subcommittee's consideration was necessarily superficial, given the time pressure under which it worked. Wright pointed out that the charge to the Large Facilities Subcommittee simply referred to readiness, not actual construction. All of these facilities were to be projects costing more than $50 million. Stevens commented that they would, de facto, have to be construction projects. McRae said that the High-End Computing Subcommittee did not consider construction, but the report from the Large Facilities Subcommittee did discuss this.

Roy Whitney pointed out that the facilities report for the Office of Science has one computing project and 19 other projects. Perhaps it should be asked what potential impacts computational science might have on the other facilities. Many, like ITER, depend on simulations. Computing will probably be 10 to 15% of its operations. It is 1 to 2% of construction projects and probably should be pushed up to 5 to 10%. Wright commented that that was a very interesting idea.

Steve Ashby noted that all present were members of professional societies. The message of computing science can be carried to other society members, and the societies can be pushed to make the case for more computing to Congress.

There being no further comment, Wright adjourned the meeting for the day at 5:04 p.m.

Tuesday, April 6, 2004

Wright called the meeting to order at 8:28 a.m. She recognized the staff that helped organize and conduct the meeting and introduced **Gary Johnson** to speak about testing the Cray X1.

A review of the Cray X1 activities was recently conducted at the Center for Computational Sciences (CCS) at ORNL. The charge was to review and analyze what is currently known about the performance of the Cray X1 on capability-limited science applications, drawing on all available information from all Cray X1 installations. Based on this analysis, ASCR was to be informed about the relative suitability of the Cray X1 to the science problem set of SC: For which applications is it well-suited? For which applications is it ill-suited? ASCR was also to be informed about the expected performance of a larger X1 installation and the potential impact on larger installations on capability-limited science applications. The evaluation process currently being used by the CCS was also to be reviewed, and improvements were to be recommended, if necessary.

He listed the review committee members who participated in the 1.5-day review, which ended with a writing session. The review included overviews of the Cray X1 system and of X1 software; comments on users, both internal and external; and a question-and-answer session.

The major finding was that the Cray X1 should be regarded as a one of the most powerful high-performance computer systems for enabling capability-limited science. In general, the machine does very well on codes that have been written as vector codes. Most codes that solve PDEs can be cast in this way, including important codes from climate modeling, fluid flow, plasma simulation, and molecular dynamics. A key question is, how much effort would it take to make the Cray X1 work on things it currently is not good at, such as computational chemistry?

The performance problems with those codes can be attributed in large part to the designs of many of these codes, which do not take advantage of the vector nature of the simulations they perform. Expanding this facility to more fully explore the capabilities of this architecture is regarded to be of strategic national importance. The committee was impressed by the openness of the evaluation process and its progress to date. The staff members work closely with the vendor and application community. The speed of the progress is remarkable.

Manteuffel asked what would be entailed in scaling up the machine. Gary Johnson replied, adding more processors and memory. From the science perspective, it makes sense to make the machine bigger. One must also consider the problem from the architecture perspective.

Giles noted that “capability-limited” might mean bad applications. Gary Johnson replied that the review panel was trying to denote capacity-limited.

Dahlburg asked how other owners’ experiences compared with ORNL’s. Gary Johnson said that it was interesting to watch the dynamics. ORNL is in a leadership position. A lot of cooperation is expected to occur between the other centers and ORNL.

Stechel asked how these applications were capability-limited. Zacharia answered that, from the point of view of the climate community, this machine is the most powerful (256 processors) available for their calculations, getting 14 times the previous performance (in years of simulation per real-time day). Some fusion problems can be solved *only* on the Cray X1. The materials scientists have been able to discern between two competing theories in high-temperature superconductivity and expect to prove one or the other with a larger machine.

Kulsrud cautioned that there are problems that one should never put on the X1 (e.g., highly scalar branching codes). Also, one can add more memory, but the speed will slow down. Users at Princeton have done an evaluation that can be shared.

Lester asked what computational-chemistry codes were used. Zacharia replied that Robert Harrison looked at NW Chem. There is a fine machine at PNNL that is highly tuned to that code. The code would not run well on the Cray X1. Codes that would not run well on this machine were not sought out and did not have a lot of activity expended on them.

Dahlburg asked Gary Johnson what he would want to see in the perfect computer center. He replied that it would depend on the purpose of the center. One should use commodity processors as far as they can go. One might want a mix of machines that would host clusters of appropriate applications.

Wright asked whether, with the leadership-class competition coming up, there are other machines that should be evaluated in a similar manner. Gary Johnson noted that ASCR has been in the business of evaluating early architecture for a long while. This is the current activity being carried out. The leadership-class machine is to be a new resource for capability-limited science applications. This evaluation and the leadership-class machine are two distinct activities.

Galas asked what the software and hardware costs were for the machine. Gary Johnson said that it depends on what “fully configured” means. DOE has been spending $20 to 25 million at the CCS, which has a couple of IBM systems and an SGI machine as well as the X1. Another way is to calculate the cost per cabinet of the X1, which is a sensitive number. Fred Johnson replied that both DARPA and the National Academy of Sciences (NAS) have cited the need for architectural diversity. These programs create architectures; they do not sell applications. That is why early evaluations are conducted. SC plays the role of helping these architectures mature.

McRae asked if the X1 is a leadership-class machine or whether we are talking about the X1 *and* a leadership-class machine. Gary Johnson said that there is an open procurement at this point. Stechel asked what other architectures are out there that have not been evaluated, noting that the IBM Power IV has been evaluated. Oliver replied that, as these machines come out, they are evaluated and put in an appropriate place. Stechel went on to ask if the diversity of available architectures had been evaluated. Gary Johnson replied that he believed so. Zacharia added that the Power 3 machines have been evaluated and that a lot of experience had been gained with the Power 4. The SGI systems are being evaluated now at NASA Ames and LANL. As a community, a lot of experience has been accumulated with clusters. DOE needs to work with vendors to ensure that architectures that are useful to users are developed. The centers try to gain as much experience as possible on all architectures.

Wright suggested that it would be good to explore what capacity-limited means. The term sounds as if there is a flaw in the software. If one does not have examples of how science will be significantly advanced, there is no impetus for larger machines. Gary Johnson said that there are many reports that have covered this topic (e.g., the ScaLeS report).

Sollins stated that some of the message from OMB was that the messages are often buried under pages of computer science. The information that OMB needs in order to act needs to be stated up front. Gary Johnson replied that what are needed are better writers and editors.

Giles said that the report should assess the relative costs and impacts of various alternatives. In running a center, one wants to make the best use of the hardware available. There are multiple points of view that can, and perhaps must, be taken. Gary Johnson answered that the purpose of this review was to respond to the charge, which did not call for the tradeoff analysis referred to. Productivity metrics is another topic that could be addressed. The right way to approach these problems at is at an interagency level. Evaluation is an ongoing activity. These evaluations tell us what applications work well on what architectures but very little about what information is handed down to future practitioners. We are working with all offices in SC to develop metrics that are useful to users of a broader range of applications. We could talk about this project in a future meeting.

Hack stated that the climate problem is moderately tractable on commodity systems now. It is 3 to 4 times faster on the Cray X1. Other researchers made compelling arguments for doing science that cannot be done on current systems if they had larger capabilities. The sense was that this is a promising direction to explore. Gary Johnson agreed and said that the most promising problem is fusion. Full-size simulations of fusion devices are needed, but they cannot be done on current-generation machines.

Wright said that it would be helpful to go through a variety of applications and list the limitations for each application. Stechel suggested that it be called “Science beyond the high-end frontier.”

Miner called attention to the fact that other agencies have developed benchmarks that would be useful to ASCR and have evaluated a number of other machines that ASCR has not looked into.

Wright thanked Gary Johnson for the presentation on the X-1 evaluation and asked him to speak about the Multiscale Mathematics Initiative.

The President’s FY05 Budget Request includes $8.5 million for the new “Atomic to Macroscopic Mathematics” (AMM; also known as the Multiscale Mathematics Initiative) effort to provide the research support in applied mathematics needed to break through the current barriers in our understanding of complex physical processes that occur on a wide range of interacting length- and time-scales. It is believed that progress in AMM will best be achieved through a combination of investments, including (1) funds for innovative approaches to multiscale mathematics at universities throughout the country, (2) investments in partnerships between university researchers and investigators at the national laboratories, and (3) additional investments in multidisciplinary teams at the national laboratories. ASCR is currently awaiting approval of the FY05 budget so it can investigate this topic. In the meantime, it is holding a DOE Multiscale Mathematics Workshop at the Sheraton Pentagon South Hotel in Alexandria, Va., on May 3-5, 2004. It will be a three-day meeting with plenary talks, breakout sessions, panel presentations, and writing sessions. About mid-May, a preliminary report from this workshop should be available to guide the conduct on a multiscale-mathematics initiative (if approved). Both workshops would probably be conducted to engage the mathematics community in planning and executing this program. This $8.5 million increase is the largest ever in the DOE mathematics program.

Giles asked if there is fertile ground for doing multiscale mathematics across many disciplines. Gary Johnson said that this program will resonate with the mathematics community. Manteuffel added, yes; one can take the mathematics and use it in many types of applications and in many disciplines.

Stechel asked if people from industry had been included in the workshop. Gary Johnson responded that he believed they had. The organizers cast a broad net for participants.

Connolly noted that turbulence is one of the toughest questions to address. Gary Johnson replied that it is addressed under combustion and also would be addressed in future workshops.

Kulsrud commented that she did not see any mathematicians among the workshop leaders. Gary Johnson replied that a lot of the mathematicians work in application areas. If people are doing mathematics, they are welcome. Manteuffel added that the organizing committee has tried to get a mathematician as a co-leader for each breakout session.

Dahlburg noted that almost every computational scientist has been working on this problem for decades and asked how this history was going to be addressed. Gary Johnson agreed that this is a long-standing problem, but noted that one does not see an organized body of literature on it. Dahlburg suggested that perhaps this is a good time to define this area (i.e., simulation) as an integrated discipline.

Wright asked what ASCAC could do. Gary Johnson replied that the committee members could attend the workshop.

Stechel noted that, while applications people have made a lot of progress in this area, a lot of application communities have wrestled with the same problems independently. A lot of advances could be made if these efforts were pulled together and made available in a generalized manner. Gary Johnson agreed and noted that collaboration seems to work well in the laboratory environment; it works less well in the university setting, which has a different reward system. There are forces and activities in the world that cannot be controlled.

Giles asked about what would happen after the workshop. Gary Johnson said that one or more solicitations in multiscale-mathematics activities would be issued, their exact shape and form to be determined.

Stechel asked if collaboration would be encouraged or would the awards go to individuals. Gary Johnson responded that SciDAC was actually mentioned in the language, and he believed that multidisciplinary collaboration will be highly encouraged.

Wright declared a break at 9:54 a.m. She called the meeting back into session at 10:30 a.m. and introduced **William Johnston**, manager of ESnet to speak about the future of that network. Steven Wolff recused himself from the ensuing two discussions and left the room.

ESnet is a communication network designed to move huge amounts of data among a small number of sites. It provides

- High-bandwidth peering for access to the U.S., European, Japanese, and other research and education networks;
- Access to the global Internet (managing 150,000 routes at 10 commercial peering points, which accounts for about 20% of the engineering effort involved in this activity); and
- Comprehensive user support with 24/7 coverage, including "owning" all trouble tickets involving ESnet users until they are resolved.

The science services it provides include Grid and collaboration services, which involve trust, persistence, and science-oriented policy to make the network work in the science environment. It is primarily SC focused, but supports NNSA/Defense Programs, including SecureNet as an overlay network.

It is truly a community endeavor. It receives strategic guidance from the SC programs, and network operation is a shared activity with the community, which ensures the right operational "sociology" for success. Both the network engineering and the network management are complex and specialized to provide network services to the national laboratories in an integrated-support environment. The network is extremely reliable in several dimensions. Taken together, these points make ESnet a unique facility supporting DOE science that is quite different from a commercial Internet service provider (ISP) or university network.

ESnet connects DOE facilities and collaborators with a high-speed ring around the country. That ring has six hubs; all national laboratories are connected; peering points are scattered around the network, with international connections at four places. There has been a 100% increase in traffic each year since 1993 when ESnet was established.

Monitoring tools have been put in place to determine who generates the traffic and where it goes to. Between January 2003 and February 2004, traffic from DOE sites onto the net decreased from 72% to 68% of the traffic; at the same time, traffic coming off the net and going to DOE sites decreased from 53 to 49% of the traffic. Interchanges between ESnet and the commercial sector and research and engineering sites (mostly universities) showed similar declines. The increase occurred in international traffic; traffic between ESnet and international peering points increased from 9 to 26% of the traffic and traffic from international peering points increased from 4 to 6% during the study period. All of this change results from BaBar data going from SLAC to Germany.

ESnet is a visible and critical piece of DOE science infrastructure. It requires high reliability and high operational security in both the network and the ESnet infrastructure support. A 24/7 help desk (joint with NERSC) and 24/7 on-call network engineers support operations.

Most of its physical facilities are at LBNL; a persistent-state replication is being set up at Ames and BNL.

The first lines of defense against attack are the laboratories' firewalls, but these are backed up by defense mechanisms in the core. All hubs have core routers, peering routers, and associated equipment (secure servers etc.).

The public-key infrastructure (PKI) service provided is trusted in the science community; it has negotiated compatible policies for transoceanic data exchange. Providing this PKI support for grids internally produces a huge cost savings to DOE. About 31 employees run this network.

The ESnet manager oversees the network engineering group, the infrastructure services group, and the science services group. LBNL business services provides resource management, accounting and bookkeeping, and contracting. The overall FY04 project is $23.3 million, of which MICS contributes $19.6 million.

ESnet is on a cusp of change. It needs to evolve fairly rapidly from this point forward because the rapidly changing requirements for supporting science require increased capabilities. (Bigger science requires larger networks.) These requirements will increase as the Spallation Neutron Source (SNS) and Large Hadron Collider (LHC) come online. The current availability of fiber provides an opportunity for a new ESnet architecture that addresses several major requirements that were identified in the roadmap workshop. However, the opportunity will decrease as the fiber market hardens after several soft years. ESnet's current budget is flat thru FY06.

A workshop held in August 2002 asked scientists about their data-communication requirements. A range of answers was received in network and middleware capabilities. The participants were also asked what would be needed in 5 to 10 years from now for end-to-end throughput. The information gained from this workshop was used to write the *DOE Science Networking Challenge: Roadmap to 2008*. The NSF has not done anything like this. This report is a good analysis of the requirements.

ESnet was ranked number 7 out of 20 in DOE's *Facilities for the Future of Science: A Twenty-Year Outlook*. Based on the requirements of the SC High Impact Science Workshop and the Network 2008 Roadmap, ESnet must maintain capable, scalable, and reliable production Internet protocol (IP) networking and provide network support of high-impact science with high, dedicated bandwidth. Evolution to optical switched networks will almost certainly come to pass. Still in question is what types of science services are needed to support Grids, collaboratories, etc.

ESnet needs connectivity between any DOE laboratory and any major university that is as good as the ESnet connectivity among DOE laboratories and the Abilene connectivity among universities. It needs to establish high-speed connections at several points in the networks [e.g., a 10-Gb/s ring in New York City to MANLAN (for 10-Gb/s ESnet to support an Abilene cross-connect and for international links) and a 10-Gb/s ring to StarLight (for a CERN link)]. Both of these connections are in progress. Major experiments will not be connected to the network as long as there are single-circuit failures. Because of the ring structure, the ring can be broken in one place, and data will still flow (in the opposite direction). No single-point failure can disrupt this configuration. What is needed is a redundant backbone so the removal of a hub connection would not bring down the network.

DOE needs to leverage and amplify non-ESnet network connectivity to the national laboratories. When ESnet has not been able to afford to increase the site bandwidth, the national laboratories have sometimes gotten their own high-speed connections. When possible, the existing non-ESnet connections must be incorporated into the new ESnet architecture to provide a better and more capable service than the national laboratories can provide on their own. The ORNL-contributed circuit plus the existing ESnet circuit effectively incorporate ORNL into a secondary ESnet core ring. The same thing is to be done at PNNL.

Johnston noted that Harvey Newman had asked, "What about increasing the bandwidth in the core?" The answer is technology progress; by 2008 (when the next-generation ESnet core will be in place) dense wavelength division multiplexing (DWDM) technology will provide 40 Gb/s per lambda, and the core will be multiple lambdas. The network will have multiple backbones and higher speeds.

The primary problem raised by high-impact science is the need to dynamically provision private "circuits" in the MAN to provide high-impact-science connections with quality of service guarantees while supporting the massive movement of data for extended periods. A proposal has been submitted to the MICS Network R&D Program for the initial development of a basic-circuit-provisioning infrastructure in the ESnet core network (site to site). Such a strategy would bypass the site firewall to obviate firewall-bandwidth problems. This bypass can be done because of the existence of a common security policy.

Johnston outlined how an optically circuit-switched network would be integrated with the traditional packet-switched networks. To clarify how to accomplish this integration, a workshop would be held at the Jefferson Laboratory the following week. In addition, ESnet co-organized a federal networking workshop on the future issues for interoperability of optically switched networks, and ESnet is participating in the Internet2 HOPI design team (in which the UltraScience Net also participates).

In summary, the roadmap workshop identified 12 high-priority middleware services, and several of these fit the criteria for ESnet support that would be required to meet anticipated needs:

- Long-term PKI key and proxy-credential management (e.g., an adaptation of the NSF's MyProxy service);
- Directory services that virtual organizations can use to manage organization membership, member attributes, and privileges;
- End-to-end monitoring for Grid/distributed-application debugging and tuning;
- Perhaps some form of authorization service; and
- Knowledge-management services that have the characteristics of an ESnet service, likely to be important in the future.

In conclusion, ESnet is an infrastructure that is critical to DOE's science mission. It is focused on the SC national laboratories but serves many other parts of DOE. ESnet is working hard to meet the current and future networking needs of DOE mission science in several ways: by evolving a new, high-speed, high-reliability, leveraged architecture and by championing several new initiatives that will keep ESnet's contributions relevant to the needs of our community.

William Wing was introduced to describe the Ultra Science Net, an ultra-high-speed research network for large-scale science, which is developing the technologies that the ESnet will be growing into in the coming years.

DOE's funding produced the early adoption of advanced network technologies [e.g., ESnet, the first wide-area network (WAN) deployment of ATM (asynchronous transfer mode); Van Jacobson's TCP (Transmission Control Protocol) slow-start and congestion control; and Mbone tools)]. It also produced the Next-Generation Internet (NGI) Quality of Service (QoS) Testbed, the first network-wide testbed for QoS technology used in the Qbone; the Clipper Project, which performed high-speed data-transfer testing of dedicated OC-12 for high-energy applications; and MORPHnet, a formal framework for conducting experimental networking activities (but which was never implemented).

The need to develop cost-effective technology with high reliability and the ability to observe nonperformance led to a series of workshops on science network requirements. Those workshops identified several classes of requirements. Diverse-domain networking requirements include

- Guaranteed QoS, best-effort, real-time capabilities, batch services, email, local/national/intercontinental large-file transfers
- On-demand bandwidth, dedicated bandwidth, shared network capabilities, SANs-LANs (storage area networks–local area networks) challenges
- Transparent cybersecurity, distributed trust, laboratory-based cybersecurity architecture, a DOE cybersecurity policy, etc.

The time is coming when distributed terascale computing facilities and petabyte data archives will need seamless access to resources located across the country and to users located at universities, national laboratories, and industry.

Critical and timely networking requirements for large-impact science that currently face the networks include advanced and deployable networking capabilities for LHC experiments, the SNS, climate modeling, astrophysics, and computational biology. What is needed is to leverage current opportunities in the telecommunication industry and mature optical network technologies to build advanced networking infrastructures for science.

The Office of Science Networking Roadmap developed by the 2003 workshop noted that capabilities and technology migration have occurred as we moved from an advanced research network, where we developed tools and techniques needed for a production network, to a high-impact science network with high reliability traded for high speed, to today's production networks like ESnet.

It is not known whether there will be switching at the lambda level, the sonic level, or at the multiprotocol label switching (MPLS) level. Today, ESnet carries the bandwidth from end to end.

The current areas of network research are

- Ultra-high-speed data-transfer protocols, TCP enhancements, and alternatives for ultra-high-speed networks;
- Dynamic provisioning of DWDM optical networks, providing on-demand bandwidth and dedicated channels;
- Network measurement and analysis, providing end-to-end performance monitoring, prediction, and fault diagnosis;
- Advanced research networks, including experimental network prototyping, testing, and deployment (building the coast-to-coast plumbing); and
- Scalable cybersecurity systems (i.e., firewalls and intrusion-detection systems).

Aggregation of small flows does not do DOE any good; it needs to transport large amounts of data in a single package. That is what the UltraScience Net is intended to do. The initial implementation of it is a multiple-lambda network. When the Virginia network gets built, the Jefferson Laboratory will be able to participate. BNL is not served yet. There is a lot of excitement. Laboratories are bringing up their own systems to get connected to this network.

As the UltraScience Net moves from prototyping and testing to production deployment, some of the technology can migrate transparently, some is partly disruptive, and some is very disruptive.

He quickly compared the features of the UltraScience Net with those of ESnet and outlined the UltraScience Net operations and management.

In summary, the UltraScience Net is an advanced experimental network testbed that is demanded and supported by large-scale science applications of SC. Large-scale science applications require extreme networking, which, in turn, requires in-house efforts because it is not cost-effective for industry. The UltraScience Net is an integrated infrastructure that is developing, testing, and deploying advanced network technologies for next-generation science applications.

Sollins noted that the Net Vision 2080 report from DARPA will be out soon; it presents another vision of where networking should be going. She asked what considerations are being given to multiscale behavior on the network. Johnston said that data will be collected and published and an archive will be established for those data; use policies will cover these activities; there will be an end-to-end monitoring capability; ways are being looked at to feed back packet-flow status to the originator. Wing said that several proposals have come in for new types of monitoring for testing an ultrascale net.

Sollins said that she hoped that SC recognized the opportunity to do multiscale mathematics as well as networks. Nagi Rau said that large-scale simulations are being run. But, when one tries to steer a visualization, device, or robot over the network, one sees the jitter in the network and loses control. This problem needs to be solved. Sollins stated that not everyone is going to get five-nines (99.999%) reliability because of physical limitations. Including these effects in network architecture is important. Johnston said that almost all applications want to see caching and storing in the network.

Stechel noted that this network is similar to the power grid in reliability and stability. Johnston replied that that is not the case. The physics of the two networks are quite different. Wing added that the economics are very different, too.

Wright initiated Committee discussion. Steven Wolff reentered the room and rejoined the discussions. Sollins asked for questions or comments on the COV report. Dahlburg asked to whom the report went. Sollins replied, Orbach. Giles said that a cover letter and a summary of the report should be written. Sollins agreed but said that she did not believe that anything else is expected to be added to the report. Dahlburg asserted that ASCAC should not accept a summary that it has not seen. Sollins pointed out that some of the appointments of members of ASCAC end on April 10. Wright said that this *is* the COV report; ASCAC could accept the substance. The hard work of the COV should be recognized.

Giles said that ASCAC should vote to accept the report with the understanding that a cover letter and summary will be added. He moved the acceptance of the COV report as is with a recommendation to add a summary and cover letter, to be approved separately, and to thank the COV members for their work. The motion was seconded by Stechel and approved unanimously.

A motion to commend the Subcommittee and the Office of Science for their fine work was made by Manteuffel, seconded by Dahlburg, and passed unanimously.

Wright asked Hitchcock to list the codes to be used in benchmarking efficiency. They are global climate change, shell model Monte Carlo, Virginia hydrodynamics, MILC, and parallel R matrix. Stechel moved that, inasmuch as ASCAC believes that the criteria for selection were sound and that the set of codes fulfills the criteria, it approves the set of codes that is to be used for the 2004 performance metric. The motion was seconded by Giles. Wright noted that ASCAC will be producing a follow-up protocol for the selection of codes and that, for now, the ASCAC is considering only the codes for 2004. The motion passed unanimously.

Wright opened the floor to Committee discussion. Manteuffel said that DOE should come up with a consistent way to carry out these performance evaluations.

Kulsrud stated that ASCAC should routinely have the national laboratories report about what their plans are and what they are doing.

Sollins asked what happens next with ASCAC because the terms of many of the members end in a few days. Oliver said that DOE is trying to get the membership up to 15 members with members having three-year terms and five members coming on the Committee and five members going off each year. Orbach, the Secretary, and the General Counsel will decide who will stay, who will go, and who will be added. That decision is yet to be made. Sollins said that such a decision should not have been delayed until the last days of the current terms of so many of the members. Oliver said that the Office has been beating the drum as best it could. Wright asked who would be the chair next week because duties, responsibilities, and requests fall upon the chair all the time. Oliver said that perhaps Orbach can designate a chair. Kreisman noted that, without a chair, one does not have a functional committee.

Wright opened the floor to public comment.

Hitchcock said that, after the multiscale workshop, a minisymposium will be presented at the SIAM meeting in July.

Rosenberg noted in response to an earlier question, that the Senate bill to provide funding for high-end computing is S 2176.

There being no further public comment, Wright adjourned the meeting at 12:03 p.m.

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