Minutes for the
Advanced Scientific Computing Advisory Committee Meeting
October 31 to November 1, 2000,
Key Bridge Marriott Hotel, Arlington, Virginia

ASCAC members present:

- John W. D. Connolly, Vice Chair
- Jill P. Dahlburg (Tuesday and Wednesday afternoon only)
- Roscoe C. Giles
- Helene E. Kulsrud
- William A. Lester, Jr.
- Gregory J. McRae (Tuesday and Wednesday morning only)
- Juan C. Meza
- Karen R. Sollins (Tuesday only)
- Ellen B. Stechel
- Stephen Wolff
- Margaret H. Wright, Chair

ASCAC members absent:

- Warren Washington

Also participating:

- David Bader, Office of Biological and Environmental Research, USDOE
- Gregg Burgess, Acting Deputy Assistant General Counsel for Standards of Conduct, USDOE
- James Decker, Assistant Director, Office of Science, USDOE
- Steven Eckstrand, Office of Fusion Energy Sciences, USDOE
- Jennifer Gregory, Oak Ridge Institute of Science and Education
- Rich Hirsh, National Science Foundation
- Daniel Hitchcock, Office of Advanced Scientific Computing Research, USDOE
- Dale Koelling, Office of Basic Energy Sciences, USDOE
- Arnold Kritz, Department of Physics, Lehigh University
- Frederick O'Hara, ASCAC Recording Secretary
- Edward Oliver, Acting Director, Office of Advanced Scientific Computing Research, USDOE
- Rachel Samuel, Federal Advisory Committee Management Office, USDOE
- Vicky White, Oak Ridge National Laboratory
- Andreene Witt, Oak Ridge Institute of Science and Education
- Melea Baker, Office of Advanced Scientific Computing Research, USDOE

About 30 others also were in attendance.

Tuesday, Oct. 31, 2000

The meeting was called to order by Chair **Margaret Wright** at 8:30 a.m. She welcomed everyone to the inaugural meeting of the Advanced Scientific Computing Advisory Committee. She asked those at the table to introduce themselves. The purpose of this meeting was to get the membership up to speed on the role and function of the Committee. She introduced **Rachel**

Samuel to review the mandating legislation and guidelines covering the conduct of the committee.

Samuel noted that the committee is convened under the Federal Advisory Committee Act of 1972 (as amended in May of 1999), which was enacted to avoid government managers' getting self-serving advice from biased sources, and is governed by Title 41 of the *Code of Federal Regulations*, Part 101-6, and by DOE regulation M 510.1-1 (Advisory Committee Management Program). All 21 DOE advisory committees are overseen by James Solit, and ASCAC is guided by the Designated Federal Officer, Edward Oliver, whose responsibilities include setting agendas, attending meetings, and moderating discussions. The committee's role is strictly advisory (lobbying is prohibited); it must conduct its business openly; and it should provide advice on the development, implementation, and evaluation of policies and programs in a defined DOE subject area. The members were selected for their expertises and to reflect a balanced point of view; their membership reflects a geographic, ethnic, institutional, and public diversity. They must be sensitive to conflicts of interest and were provided with a conflict-of-interest charge letter. Requirements of members include commitment (prepare for meetings and ask questions), frankness (make candid and objective observations and recommendations), and the avoidance of even the perception of conflict of interest.

Sollins asked if DOE principal investigators (PIs) or contractors could serve on an advisory committee, and Samuels said they could. Sollins asked if Committee members can talk to Congress, and Samuel replied they could as long as they speak as individuals and not represent themselves as speaking for the Committee. Speaking as Committee members should be done as a group through DOE. Wright noted that other advisory committees had "educated" Congress without going through the Department. Samuel responded that Gregg Burgess would deal with the fine details of Congressional relations.

Wright introduced **Edward Oliver** to speak about the role of the Committee. He noted that the research portfolio of the Office of Advanced Scientific Computing and the balance of that portfolio should be an item of interest to the Committee. Other items of interest include evaluating the quality of the research and facilities of the Office, reviewing future plans and funding, and giving advice and recommendations and initiative ideas.

A number of other questions of interest include the role of individual investigators in teams of research. Is the National Science Foundation (NSF) doing enough, or should DOE support these research teams more? Congressional staffers are interested in such subjects. Is there a good balance between the national laboratories and universities? Generally, the laboratories do not exceed 20% in "work for others" (work performed under funding from agencies other than DOE), an unwritten rule.

Some of the objectives of the Office are to develop software tools for the DOE research community to make those researchers more competitive. A perennial question is whether DOE should pursue high-risk goals. Also, the infrastructure vs research issue must be dealt with and reviewed periodically. Another item that the Office is thinking about is how to work better with other agencies, such as the NSF.

Wright noted that advisory committees normally have an agenda of questions such as these to deal with, but because of the initial nature of this meeting, it is geared toward background information. Meza commented that DOE's research is often referred to as mission oriented, and NSF's work is often considered more likely to be basic research and he asked if there was such a split. Oliver replied that it is a matter of perspective; some would say that *all* research was basic research; however, all of DOE's work has to be relevant to its mission..

Dahlburg asked about the relationship between the Office and the Accelerated Strategic Computing Initiative (ASCI) run by DOE's Office of Defense Programs (DP). Oliver answered that a lot of interaction occurs between the two entities, pointing out that Office staff and ASCI participants often attend the same meetings and that some DOE contractors and even ASCAC members are on ASCI.

Wright introduced **Gregg Burgess**, Acting Deputy Assistant General Counsel for Standards of Conduct, to speak about conflict-of-interest issues. He noted that advisory-committee members who are not federal employees have certain policy restrictions placed on them. No activities shall be tainted by personal interests (including those of a spouse, children, employer, or other boards on which the member might serve). If a question of impartiality might come up, it should be brought to the attention of the Designated Federal Officer (DFO; in the case of ASCAC, Ed Oliver). In cases of potential conflict of interest, an advisory committee member might be asked to recuse himself or herself from those particular discussions. Also, advisory committee members should refrain from using the position for personal gain because the government wants to be sure that the Committee's advice and recommendations are credible. He went on to note that federal employees have additional responsibilities (covered by criminal statute) concerning conflict of interest and that all of these basic restrictions on federal employees also apply to their activities as members of an advisory committee.

Wright noted that it might be helpful if the committee members had copies of the policies and guidelines, and Burgess said that he would see that they did. He went on to note that sometimes an individual had to be removed from an advisory committee and the committee had to redo its entire work without that member present to ensure that the considerations and recommendations were impartial.

Sollins stated that she was an IPA (an Intergovernmental Personnel Act detailee) and was going back to academia in December. She expected that the research she was involved in will eventually be the topic of discussion of this advisory committee. She asked if her experience and expertise will disqualify her from participating in those discussions. Burgess said, probably not. That experience and expertise is exactly why she was wanted on the committee. The question would be whether her participation might have a favorable influence on her future financial status.

Wright asked Burgess if he could say more about advisory committee members not lobbying Congress. He began by noting that his dealings with lobbying interests are limited but that a basic distinction must be made between going to Congress as a representative of the Advisory Committee and going representing yourself, your employer, or your professional society. Under the Constitution, individuals and corporations have the right to address their concerns to Congress; what one has to be careful about is addressing a personal concern to Congress as though that concern was advice from the Advisory Committee. Decker pointed out the cases where one might be requested by Congress to testify on behalf of an advisory committee or where one might be accompanying DOE staff from the Hill to provide advice to Congressional committees.

Burgess closed by stating that DOE cannot abridge the First Amendment rights of anyone; what the policies seek to do is to make clear what activities can be carried out as a member of an advisory committee. He promised to get more information to the Committee about this subject.

McRae returned to the subject of Oliver's presentation and asked how this program related to the NSF program and others. Oliver responded that there was a lot of coordination and copublished documents among DOE's DP, the NSF, and the Office of Advanced Scientific

Computing Research (OASCR or ASC). Decker suggested that a briefing be given at a future meeting of this committee on the linkages between this office and other agencies.

Wright stated that several of the other 21 advisory committees may touch on topics common with the interests of this Committee and asked if ASCAC could get some idea of what these other advisory committees are doing. Decker responded that a lot of synergism occurs among the different DOE program offices and that such interaction among the advisory committees may be helpful, also. Oliver said that some overlap occurs in the advisory committees' memberships, noting that Warren Washington was not only on ASCAC but also on BERAC (the Biological and Environmental Research Advisory Committee).

Wright asked what a Committee member should do if he or she wished to know what ASC is doing on a specific topic. Decker said that the members could ask DOE and that the Office will try to put them in touch with the appropriate person in DOE or elsewhere to provide anything the Committee (or its members) wishes to learn about. Sollins noted that many advanced research topics are of interest to all advisory committees and asked if there should be a joint advisory committee meeting for such cross-cutting issues. Decker said that DOE has had joint advisory committee meetings, although not a meeting of all the committees.

Wright noted that another aspect to be considered is the size of the efforts conducted by the different offices of DOE. For example, the high-energy-physics projects are very large and require teams of workers and much coordination of those teams, To see the difference between DOE offices, one has only to ask the question whether software can be written by a 50-member team.

A break was declared at 9:42 a.m. The meeting was called back into session at 10:08 a.m. with the introduction of **James Decker** to provide an overview of the operations of the Office of Science (SC). He started by delineating the DOE mission areas: energy resources (to foster a secure and reliable national energy supply), national security (to maintain the safety and reliability of the nuclear stockpile), environmental quality (primarily to repair the environmental consequences of the cold war), and science. Most of the science is performed in and funded by SC. He reviewed the DOE organization chart to show the relationship of SC to the rest of the Department. He noted that, with the recent establishment of the National Nuclear Security Administration (NNSA), a long-term issue currently facing the Department is whether security activities at the weapons laboratories [Sandia National Laboratories (SNL), Los Alamos National Laboratory (LANL), and Lawrence Livermore National Laboratory (LLNL)] will be a barrier to good research. He said that for FY 2001, the budget of the Department was split among its four mission areas as follows:

	Overall (billion $)	R&D (billion $)
Environmental quality	6.8	0.29
National security	6.6	3.20
Science	3.2	3.02
Energy resources	2.2	1.15

This devotion of funding to R&D makes the Department a major player in the conduct of research in the United States. It is in the top five federal agencies in funding basic and applied research, underwriting academic research, and operating R&D facilities. In addition, DOE is in the top five federal agencies in funding the physical sciences (for which it is the largest supporter), environmental sciences, mathematics and computing, engineering, and life sciences.

SC has five program offices:

- OASCR provides large-scale computing capabilities.
- Basic Energy Sciences (BES) funds research on materials, chemistry, physics, biosciences, engineering, and geosciences and it constructs and operates some of the large user facilities.
- Biological and Environmental Research (OBER) investigates global climate change, the human genome, carbon sequestration, and environmental cleanup.
- Fusion Energy Sciences (FES) researches magnetic and inertial confinement and plasma science.
- High Energy and Nuclear Physics (HENP) builds and operates the major reactors and accelerators in the country.

He showed a map that displayed the number and locations of SC's user facilities and the universities that use these facilities. He noted that 50% of the facilities' 17,000 users each year are university researchers; the other 50% come from national laboratories, other federal agencies, and industry. He also pointed out that some of the computing centers that support these user facilities are the only ones like them in the country.

He reviewed the SC organization chart to show the relationship of the advisory committees to the Office and the relationship of OASCR to the other program offices and the field operations.

He then reviewed the history of the FY-2001 SC budget, which increased from $2788 million to $3151 million, including an additional $161 million for the Spallation Neutron Source (SNS), $70 million for high-performance computing, $6 million for user-facility upgrades, $36 million for nanoscience engineering and technology, and $25 million for the life sciences. This budget received broad-based support from the scientific, university, industrial, and legislative communities; however, the overall budget is not as good as it sounds because most of the large increase was for the construction of the SNS. Nonetheless, high-performance computing did receive dramatic increases because of computing's contributions to such programs as climate modeling, the human genome, and protein dynamics. After displaying an extensive list of major research areas pursued by SC, he showed a graph prepared by the American Association for the Advancement of Science (AAAS) that showed that funding for the physical sciences has remained constant from 1970 to 2000 while funding for the life sciences has increased significantly (in constant dollars). The Director of SC, Mildred Dresselhaus, is concerned about large increase in life sciences not being matched in the physical sciences. This imbalance of the nation's research portfolio is troublesome because a large part of the U.S. economy depends on advances in the physical sciences. Because of inflation and the flatness of funding, the buying power in SC's budget has decreased with time. Some large facilities (e.g., the Advanced Photon Source) have been built, and the operating costs of those facilities have increased, taking away from the funding for research.

He listed SC's advisory committees and noted that they review large portions of the program, assess the balance among programs, conduct program-balance reviews, and develop long-range plans. The Basic Energy Sciences Advisory Committee (BESAC) was given as an example. It reports to the Director of the Office of Science, who provides guidance about needed input through charge letters to the Committee. It reviews elements of the program. It provides advice on long-range plans, priorities, and strategies to address more effectively the scientific aspects of energy-related research. It gives advice on the appropriate levels of funding to develop these plans, priorities, and strategies and to help maintain the appropriate balance among competing elements of the program. Much of its work occurs between Committee meetings and is done by subcommittees. Those subcommittees can pull in expertise from outside the Committee, and their meetings are not necessarily open to the public.

Lester asked Decker what concerns for the future he would recommend that the Committee look at. He responded that the Committee might look at computing-related issues in the programs of other divisions of the Office of Science. Also, it might look at the future of computing and how DOE does it. DOE managers have to prepare their budgets about two years before those budgets are executed, requiring those managers to make decisions with very long lead times. Those managers need advice on where the Department should be going.

Dahlburg asked if he could comment on the mission-oriented nature of DOE. Decker noted that many of the programs of SC are cross-cutting and have a component of advanced computing.

McRae asked how big SC's R&D budget is in comparison with those in other agencies. Decker said that he did not have the budget data for any other agencies and therefore could not answer the question. McRae asked if technology transfer to industry was part of the SC portfolio. Decker said that it was but noted that that budget, which used to be $45 million per year, had taken a beating at the hands of Congress. Wolff asked what vehicles were used to carry out technology transfer. Decker said that the national laboratories are the main players in this activity and that they use cooperative research and development agreements (CRADAs), technical assistance procedures, technical staff exchange, work for others, and other mechanisms.

Meza asked if the bump-up in the budget for advanced computing was a one-year phenomenon. Decker said that it certainly was hoped not. SC is hoping for a 15% per year increase in budget, but a new administration will have a lot to say about it. Wright commented that, as more of science relies on computing, the country should fund more computing research. The approach in the past has been to hire more physicists, under the impression that physicists can do anything. After noting that he, himself, was a physicist, Decker agreed that this has been a problem and that this Committee could make a contribution by talking with the High Energy Physics Advisory Committee about how they do research and how computer science can contribute to that effort. He noted that the Defense Advanced Research Projects Agency (DARPA) has traditionally put a lot of money into computer science and asked if DOE's focus on computational activities should be continued. DOE certainly recognizes the need for greater participation by computer scientists. McRae noted that most congressmen do not understand the difference between computer science and computational science and that the Committee should look at how the budget is split between these two areas. He suspected that computation at SC has been going down in relation to computer science. Decker said that they had spent a good deal of time on the Hill trying to explain this difference and the need for computation research and development.

Meza pointed out that recruiting and retention in information technology (IT) is difficult and asked to what level this issue had risen at SC and whether this Committee should look at that, also. Decker said that it is a concern at SC and even at the Secretarial level. The Department has heard a lot from the national laboratories about the difficulties in this area. If this Committee has any suggestions, DOE would be glad to hear them.

Wright then introduced the Committee's DFO, **Edward Oliver**, to give an overview of the Office of Advanced Scientific Computing. He pointed out that much of the advanced research conducted by DOE requires the very highest in computing: big machines that are not easy to use, so enabling tools have to be built. Referring to the map of DOE laboratories, facilities, and user groups, he noted that SC is the landlord for ten of these national laboratories, that DP is the landlord for three of them, and that SC is involved in many of the smaller laboratories, also. He displayed an organization chart for SC that was the same as that shown by Decker, except that this one showed an Assistant Director for Scientific Simulation (Thom Dunning) reporting directly to the director of SC. An organization chart for SC's OASCR showed the Mathematical, Information, and Computational Sciences (MICS) Division, the Technology Research Division, and the Office of Scientific and Technical Information (OSTI) reporting to OASCR's director. The two divisions are coupled budgetarily, while OSTI is a separate line item whose funding has declined from $20 million per year to $8.5 million this year.

A graph of MICS's budget history from FY 1991 to FY 2001 showed slow but steady growth from FY 1991 [when the High-Performance Computing Centers (HPCC) Program started] to FY 1999, a sharp decrease in FY 2000, and a large increase in FY 2001. A final organization chart for MICS showed eight teams reporting to the director. Note was taken that many of the personnel (including the director) of MICS were "acting" leaders. All of those people have other, full-time positions. These positions have been offered to top applicants, but they have all gone elsewhere. As a result, the division does not have the personnel to start new activities and needs advice on how to cope with this situation.

Connally suggested using IPAs to fill slots. Oliver pointed out that he was an IPA and that a number of the other staff positions were already filled with IPAs.

Wright asked how decisions had been made in the absence of an advisory committee. Oliver said that the problem had been easy because funding had been flat and no decisions about new programs have had to be made. Within existing programs, the staff considers how to balance the level of available support on a regular basis. Wright asked how flexible or inflexible the Office's budget was to respond to new, hot topics. Oliver said that it is flexible enough to respond to such opportunities. McRae asked how the Graduate Fellowship Program in Computational Science was funded. Oliver responded that DP and other divisions contribute to the funding for that program. Connolly asked Oliver what he had meant when he said that the budget numbers presented were being adjusted. Oliver responded that the numbers are still being tweaked in the second decimal place as fine adjustments are made among programs' funding.

Wolff asked if there was an education program. Oliver said that the only education program was the Graduate Fellowship Program, but it was not filling the demand for trained personnel. The educational program for students from kindergarten through high school had been eliminated throughout the Department several years ago. Dahlburg asked why, and Decker answered that Congress had not viewed the program to be within the DOE mission but that the NSF is now getting an additional $100 million per year from fees collected from telephone usage. It is hoped that some of those funds will be used in DOE for promoting science education. Giles said that the NSF would be a natural target to try to get funding for education. Sollins noted that bringing students into a research lab is a great investment. Wolff stated that the best thing about such a program is that it is a program director initiative. It could be done at the graduate and faculty level, too. Connally said that, unfortunately, it is a minuscule effort at this point. Wolff said that such a program could make teaching a lot easier; endowed chairs are a waste of money, and five or six researchers or teachers could be supported for the same amount of money put into an endowed chair. Lester asked what the level of support was for these programs, and Wright said that certain laboratories have limited numbers of slots for which students compete. However, the number of slots is small, and DOE could do a lot more.

The meeting was adjourned for lunch at 11:44 p.m. An executive meeting was held during the lunch break. After the executive committee meeting, the full committee was called back into session at 1:16 p.m. with the introduction of **Vicky White** to speak about the computing needs in high-energy and nuclear physics. That community's experiments involve a large number of collaborators producing huge amounts of data. In addition, it uses computers for accelerator design, for theory development, and for simulation.

Accelerators are a fundamental tool for high-energy physicists. Large-scale simulation for accelerators is needed

- to study existing accelerators to expand the operational envelope, to optimize performance, and to increase reliability;
- to design next-generation accelerators, requiring design modeling and reducing cost and risk; and
- to develop new accelerator technologies.

To meet these needs, a terascale accelerator-simulation environment is being developed. This effort grew out of the Accelerator Grand Challenge Project sponsored by OASCR/MICS and supported by HENP; it now involves personnel from six laboratories, two universities, and two supercomputing centers. The group is now modeling the 2-GeV H^- linac of a conceptual design for a neutrino factory. The effort involves simulation of the beam system and the electromagnetic system and the integration of the two. Simulations require 2.5 to 75 teraflop days to run, and modest problems involving grids would require 4 TB of memory per array. Right now, 100-million-particle runs require 5 to 10 hours of machine time. General-purpose codes have been developed for accelerator simulation. Two examples are IMPACT, a three-dimensional parallel particle-in-cell code, and OMEGA3P, an electromagnetic code.

The goals in accelerator simulation are to

- develop a comprehensive simulation infrastructure and set of long-lived and portable codes,
- build an interdisciplinary team,
- improve the precision of simulations,
- optimize performance and employ new techniques,
- benchmark codes against experiments,
- leverage IT developments elsewhere in the industry, and
- advance visualization and animation.

Theoretical physicists are major users of computational and memory resources for analysis of lattice-quantum-chromodynamics (QCD) studies of weak interactions of strongly interacting particles, the hadron mass spectrum, and the internal structure of hadrons. In response to this great need for computational and memory resources, theorists are prepared to use and even build special-purpose computers in addition to using the National Energy Research Scientific Computing Center (NERSC) and other supercomputing centers. In 1998, Columbia completed the QCDSP (Quantum Chromodynamics on Digital Signal Processors) machine, which is capable of sustained operation at about 300 Gflops. The Japanese also have a 300-Gflops machine, and the Italian theorists have built special-purpose 4 × 125-Gflop machines. In comparison, the largest allocation to a single PI at NERSC or an NSF center is less that 10 Gflops. U.S. researchers are now looking at the possibility of designing and building clusters and/or QCDOC (QCD-on-a-chip) machines.

Another area that requires additional computer resources is nuclear astrophysics simulation. For example, ascertaining the explosion mechanism for neutrino-driven supernovae and the accompanying process of nucleosynthesis is a terascale application that requires (1) the development of a scalable large-sparse-system solution for radiation-transport and nuclear-structure computation; (2) scalable, multidimensional, multifrequency radiation hydrodynamics; and (3) collaborative visualization. The programming community for this project has as its goals (1) to move from MPI (the Message-Passing Interface) to MPI + OpenMPI hybrid parallel-programming models; (2) derive memory-placement strategies and cache-aware codes; (3) produce parallel tools for modeling, predicting, and visualizing; (4) develop customized instrumentation of terascale codes for performance analysis; and (5) write enhanced parallel debuggers.

Experimental physics presents challenges of its own: petabyte data sets, collaborations among a thousand individuals distributed worldwide, a model for distributed hierarchical computing for data analysis, very high network bandwidth and performance, the use of long-lived and evolving codes built by huge teams, and complex pattern-recognition problems and algorithms. Such experiments typically start with some basic physics process that produces fragmentation and decay, which then interacts with a detector material, and the detector imperfectly responds to give raw data (currently 250 kB per experiment; soon to be 1 PB per experiment). The physicists then have to work backwards through this to simulate what they thought they saw in the experiment. A large amount of computing has to be done at each of the steps, and the data have to be stored and transported to the collaborators. The compact muon solenoid project includes 1800 physicists at 150 institutions in 32 countries and uses a five-tiered worldwide computing plan: The data are collected by an online system, transferred to an offline farm, distributed to five national centers, redistributed to regional centers, passed on to institutional centers, and finally downloaded to people's desktops.

A number of studies have been performed on this massive amount of data transfer. These studies found that more bandwidth was not all that was needed, but projections of HENP bandwidth requirements for 2005 indicated a range of needed bandwidth from 10 Mbps for individuals to 10,000 Mbps for central laboratories housing one or more major experiments. The Grid concept of distributed resources offers some hope of assistance in meeting these requirements. MICS has funded a workshop on the establishment of a particle physics data grid, a proposal to NSF has been funded at $11.9 million over five years to establish the GriPhyN network, and the European Union has funded the EU DataGrid coordinated by CERN [Conseil Européen pour la Recherche Nucléaire; now the European Organization for Nuclear Research (Organisation Européenne pour la Recherche Nucléaire)] for three years at a total of 30 million euros.

In summary, HENP has enormous needs for computing. The scientific simulation programs need cycles, algorithm development, frameworks, applied math, visualization, collaborative and problem-solving environments, and data management. The theory program has huge demands for raw computing and increasing demands for flexible environments and portable codes. HENP already is a heavy user of NERSC but needs help in processing and making petabyte data sets available to thousands of users. The Grid concept that is being adopted assumes a highly performant, adaptive, and controllable network infrastructure. Huge collaborations and the resultant networking will produce voracious demands. As a result, HENP will be looking to OASCR for help in addressing these needs.

Connolly suggested that a lot of this computing might be shifted over to Linux clusters. Giles asked if the reuse of software technology was occurring across the discipline or within research groups. White answered that it is happening mostly within a collaboration. Meza asked what kind of performance was being obtained from the visualization efforts in accelerator research. White responded that new algorithms and statistical techniques are being used. Meza noted that the ASCI also has a goal of 1 million messages, and White responded that one way of achieving progress is through collaboration. Rich Hirsh of the NSF asked if the codes are used only within a specific research group or if they are generalizable. White responded that they are generalizable and that people are moving toward community codes.

Wright noted that physicists seem to use “favorite” codes and algorithms, and she asked if there is a push to open this group to interdisciplinary teams and the use of best-available codes even though those codes might come from outside the physics community. White said that she thought that the situation was improving.

Connolly asked if a mechanism existed to get people in DOE together , and White responded that she hoped that this new initiative would take care of that concern.

Wright introduced **David Bader** to speak on the need for computational and simulation support by OBER. He noted that OBER was dealing with

- complex environmental systems that cannot be studied in the laboratory (e.g., climate change),
- massive volumes of biological data that need to be searched and analyzed (at the level of the genome, protein, molecule, etc.), and
- molecular and biomolecular simulations for environmental and biological applications.

Each of these tasks includes huge computational problems. OBER’s approach to these computational problems is collaboration with other agencies and other DOE divisions, such as the NERSC, Oak Ridge National Laboratory (ORNL) and LANL (where OBER is a part owner of the machines there), and the Environmental Molecular Sciences Laboratory (EMSL) Computing Facility.

Biology happens through instructions for making proteins (from a genome); the activation of on/off switches for genes; the coordination of gene expression; and the functioning, behavior, and interaction of proteins. Modeling and analyzing these processes will overwhelm the current computer capabilities in the next few years.

Knowledge of the genome will help predict health risk, clean up the environment, remove excess CO_2, and produce clean energy. High-performance computers will be needed for each component of cell-level simulations in solving the assembly and analysis of genome sequences, prediction of protein structure, determination of classical molecular dynamics, derivation of molecular dynamics from first principles, and simulations of biological networks.

The EMSL was established to conduct fundamental research on the physical, chemical, and biological processes that underlie contaminant fate and transport in the subsurface, processing and disposal of stored wastes, the cellular response to environmental contaminants, and atmospheric chemistry. The Molecular Science Computing Facility provides EMSL users and staff with production computing, advanced molecular modeling software, and production facilities for visualization and analyses of complex data sets. Its software suite includes the Northwest Computational Chemistry Suite (NWChem), Extensible Computational Chemistry Environment, and Parallel Software Development Tools. EMSL performs a broad range of computational-science research, supporting major facilities in high-field nuclear magnetic resonance (NMR) and mass spectroscopy (MS) as well as in molecular sciences. This computing capability is essential for all the programs at EMSL, and a large number of users are offsite. Dahlburg asked how many people worked there, and Bader said there were 300 people on staff and 1200 other users. EMSL is hoping to develop collaboratories to provide remote access to EMSL's NMRs.

Climate change has gone from a scientific question to being a source of modeled predictive information, with those models requiring thousands of high-resolution, time-dependent runs. Such climate prediction requires accurate and verifiable projections of climate change at regional resolution, statistically meaningful measures of natural variability, multiple scenarios to evaluate emission-reduction strategies, and completion of the Intergovernmental Panel on Climate Change (IPCC) third assessment by 2005. A main output is Warren Washington's parallel climate model, a coupled general-circulation model (GCM) that can be brought up on any system in a couple of weeks. That model produces predictions of global surface temperature anomaly that are in very good agreement with observed data. OBER also supports a coupled ocean GCM that produces projections of ocean-surface temperatures and currents that have a resolution that is an order of magnitude greater than was possible a decade ago. These GCMs can take as little as one hour of computer time for every five years simulated, with the great majority of the computational load being in the atmospheric calculations. Graphs of performance (simulated years per wall-clock day) showed a great difference from machine to machine (because different machines are optimized for different capabilities), by number of processors, and type of simulation.

However, a parallel climate model is not a meaningful benchmark because *resolution* is the problem here. The performance gap between the United States and its foreign competitors is real not because they have better software (the U.S. software is just as good) but because they have better machines to run the software on. In response to this gap, the Accelerated Climate Prediction Initiative (ACPI) will be the United States' entrée into high-performance computing. The goal is to predict weather on the ocean at a resolution of 10 km, in the atmosphere with a resolution of 30 km, and at the land surface with a resolution of 1 km with 10 scenarios and 10 realizations per scenario. These predictions would be based on calculations of cloud-radiation interaction, biogeochemistry and hydrology, aerosols and non-CO2 greenhouse gases, and surface-atmosphere exchange. ACPI activities during FY 2000 included an end-to-end pilot project and a computational collaboration on the National Center for Atmospheric Research's (NCAR's) Community Climate System Model (CCSM). The plan is to build model consortiums, an effort that would (1) require true collaboration among participants (not dominated by any one discipline or institution), (2) have few precedents in climate and weather (although it has worked in other disciplines, such as computational aerodynamics), and (3) work toward identified milestones.

Dahlburg asked if it was policy to buy only U.S. computers. Oliver noted that any computer can be bought, but with a NEC computer, a 400% tariff is paid. Meza called attention to the time it took for calculations to be made and suggested that, perhaps, the wrong algorithm was being used, Bader said that the code being used was the most accurate, and the quality of the result is very important, especially because the calculations were for annual simulations that were being iterated for tens of sequential years and any error is cumulated. Meza asked how the calculations are verified. Bader said that diagnostic runs were performed for past climate scenarios, and the results (or "predictions") were compared against historical data. However, even with all the improvements of algorithms, the improved performance of the machines cannot be matched.

Wright asked Bader what his message was. He answered that high-end computing in this country is behind that in other parts of the world. High-end computing is in a decline in this country, and it needs to be reinvigorated to advance physics, atmospheric sciences, etc.

McRae asked if the hardware was only a tiny part of the problem, and Bader agreed that it was but that it is a critical part of the problem.

A break was declared at 2:50 p.m. The meeting was called back into session at 3:07 p.m. to hear **Dale Koelling** speak about the computing done by BES. He used an organization chart to show that BES is an office of SC, as is OASCR, and that each has its own advisory committee. He highlighted the range of research conducted by BES by reviewing its organization chart and a list of major research areas participated in by the different offices of SC. He noted three broad areas of overlap between BES and OASCR: computing, networking, and applied mathematics.

In computing, BES uses just under a quarter of the resources provided by NERSC. That usage is heavily CPU intensive, modest in memory demands, and small in archival storage. BES researchers have participated significantly in the OASCR computing research centers at Argonne National Laboratory (ANL), LANL, and ORNL, as well as at others at Pacific Northwest National Laboratory (PNNL) and SNL. Much work is done on workstations and Beowulf clusters. Database sharing and archiving, however, has been provincially organized and may face some difficulties.

In networking, BES researchers use all the traditional network services: remote login, file transfer, e-mail, and the World Wide Web. The national facilities that BES operates have special communication challenges: data storage and transfer are getting larger, offline data analysis needs to be faster and more efficient, collaborative efforts are expanding, and remote operation is becoming more popular. This research is being conducted at many sites across the country, and the new generation of light sources are the most intensive and interesting sources of data.

In applied mathematics, BES researchers understand and use libraries. Joint projects occur at all levels, from individual investigators to grand challenges; in these projects, the biggest roadblock is often agreeing on the problem. However, opportunities do go knocking as the problems get larger in scale and complexity. If the hardware or software gets better, the BES researchers will be right there to use it.

BES work at the DOE laboratories is now dominated by world-class scientific facilities serving the nation, by collaborative research centers, by research associated with the themes of these facilities and centers, and by other research uniquely suited to the laboratories. This trend is supported by numerous blue-ribbon advisory panels. Work at universities is a critical component of our portfolio; it has remained a constant fraction of the research portfolio for more than a decade, and will continue so. Laboratory activities are increasingly linked to activities at other institutions. The constant level of funding that the physical sciences have received during the past several decades is actually decreased funding because of the effects of inflation.

These trends imply more networking, more collaboration, and more cooperative efforts among SC and DOE offices. One response to the increased complexity of the problems is the code sharing of quantum chemists. However, what one can do with a canned code is more limited than what one can do with a custom code. Other responses are the Computational Materials Sciences Network and cooperative research teams that are assembled to work on large-scale projects. But these large, dispersed research teams bring new requirements and challenges, such as the remote operation of instruments at the High-Flux Isotope Reactor (HFIR) and the SNS in Oak Ridge, Tennessee, by users at industrial and university laboratories across the country. One effect of these trends is that the initiatives are multiplying as the core research program diminishes.

Another way to perceive the variety of the research conducted by BES is to consider the vast range of length scales dealt with. These lengths go from the scale of electronic shells and orbitals, through the nanoscale where material properties are determined, to the continuum scale at which manufacturing occurs. These scales cannot be bridged by increasing the computer power used to study them. Computing will, however, broaden the steps from one scale to another so the importance of the transition can be studied.

Looking to the future, advanced computing will have to be done both on high-end machines and on workstations. In addition, a third category that will become more and more important is the dedicated engine to provide on-the-spot processing of the vast amounts of data that instruments produce too rapidly to be piped away. In addition, the networking resources will need to be enhanced. One such enhancement must be the interfacing of ESnet with VBSnet to reach the universities. And in applied mathematics, (1) libraries must be advertised and documented as well as created and (2) basic research (which encompasses ideas, discoveries, codes, and databases) must benefit from funding that will bring that research to the point where it can be developed into usable and applied products. This last point is applicable to database management, neural nets, validation, management, and interfaces and must be a highly leveraged endeavor.

Wolff noted that ESnet already had an interface with VBSnet. Koelling said that the network system needs continual care to exist along the way. Wolff said that Koelling had referred to some new light sources. Koelling pointed out that all of them are already in place, except for the SNS. Wolff asked what special challenge they presented for computing. Koelling said that each of them is a virtual fire hydrant for data. The data have to be partly digested onsite and partly have to go to the offsite researcher for analysis. These facilities bring in data so fast that one problem is getting the detectors to transfer the data. Wolff asked if the data-production rate was larger than 259 MB/s, and Koelling said that he did not know the exact rate. Wolff asked if the scientific community had a common understanding on how to deal with these large data sources. Koelling said that the community acknowledges the problem with the networks. If a telepresence is to be created for a researcher, that produces a huge bandwidth demand. Wolff said that bandwidth is not a network problem. Koelling said that it is a component of using the network. He said that he was not a network expert, but he knew that the data must be transferred from the research site to the user.

Sollins asked Koelling what he would like to see for a relationship between BES and OASCR. Koelling said that cooperation between BES and OASCR was essential; BES has to look upon OASCR as an opportunity for the solution of BES's problems.

Wright introduced **Stephen Eckstrand** to tell the Committee about the computing needs of FES. The mission of FES is to advance the knowledge base needed for an economically and environmentally attractive fusion energy source by understanding the physics of plasma, identifying and exploring innovative approaches to fusion power, and exploring the science and technology of energy-producing plasma as a partner in international collaboration. In technical terms, the goal is to produce a plasma at 100 million degrees at a density between 10^{20} and 10^{30} particles per cubic meter. Currently, 180 institutions (67 of which are in the United States) in 33 countries are working on fusion energy. The simplest problem is the deuterium-tritium reaction. The three traditional ways to keep the reacting particles together are gravitational confinement (which works well for stars), inertial confinement (which heats, compresses, and ignites the plasma before the constituents fly apart, a method that works well at the megaton level), and magnetic confinement (which uses the unique properties of ionized particles in a magnetic field). In magnetic confinement, major losses are incurred if the helical motion of the ions follows a linear path. These losses are avoided by bending the path to form a torus.

The challenges that fusion research faces are to

- develop physical models for plasma stability and transport, including
    - the motion of particles and fluids in three-dimensional (3-D) magnetic-field structures,
    - a wide range of spatial and temporal scales, and
    - free energy that drives strong turbulence;
- design large experiments;
- develop complex plasma-diagnostic instruments;
- develop plasma-heating and fueling methods; and
- acquire, analyze, display, and interpret large quantities of data.

All of these tasks are computationally and communicationally intensive because they are done in a collaborative environment. Despite the challenges, progress has been significant. During the past 35 years, fusion yield has increased by about 10^7, leaving about a factor of about 10^2 to go to attain practical fusion power.

However, before any large machines are built, they need to be modeled to predict how they will work. That task is made difficult because it entails a many-body problem with essential nonlinearities, turbulence, self organization, mesoscale phenomena, and a huge range of spatial and temporal scales. It also includes complex 2-D and 3-D geometries. Two ways to solve the basic dynamics are particle-in-cell simulations and turbulent-transport studies. These techniques each involve seven dimensional equations of motion. Realistic problems require many approximations to solve because they track, say, 400 million particles in a large radial structure with flow, which leads to large transport of particles and energy. If the flow is rapid enough, it breaks up, eddies form, and transport is very much reduced. Magnetohydrodynamic (MHD) codes can model large, complex fields that ultimately disrupt the plasma. We can now model 100 million to 400 million ions. (He showed the results of a 3-D simulation of 100 million ions with five nonlinear equations that were solved 1.2 trillion times.) What is needed is to simulate larger-sized plasma cross-sections and to add the electron dynamics to evolve the magnetic as well as the electric fields in the plasma. Although some problems will be solved onsite in real time on UNIX clusters, a 10-teraflop computer is needed in the future, and there is interest in high-end, massively parallel machines with specialized architectures.

In summary, the collaborative computing practiced by the fusion-energy research community means that enormous data sets must be shared, code development is increasingly the work of geographically dispersed teams, visualization and other distributed-computing tasks are being explored, and computing equipment needs to be coupled with the operations of large machines.

Arnold Kritz then described the National Transport Code Collaboration (NTCC), which was begun by Fusion Energy Sciences about two years ago. Cutting across laboratories, universities, and industries, it was designed to change the way fusion modeling codes are constructed and used. The central idea is to use flexible codes based on modern software engineering. A major objective is to develop new, Web-invocable codes that can be run from a browser and integrated to conduct modeling, to test models against experimental data, and to predict confinement in new experiments. Another objective is to design transport-modeling codes for use by experimentalists, theoreticians, and modelers.

The components of the NTCC are a module library, a client-server framework with a Java client, a data accessor to access the stored experimental data, a physics server that uses Python and C++, and an education aspect. The NTCC code can use three computers simultaneously.

(The user runs the client on a local computer, the physics and data server may run on a remote computer, and the data can be stored on and accessed from a third computer.) The client provides a graphical user interface that controls the physics and/or the data server and can produce plots of input and results that are written in Java script for Web access.

The integrated modeling code consists of a central framework that handles modules, such as those for core transport, plasma turbulence, neutral beams, rf (radio frequency) heating, neutral gas, radiative transport, nuclear reactions, external circuits, large-scale instabilities, equilibrium shape, plasma-wall interactions, and edge transport. Currently, 29 modules are available at w3.pppl.gov/NTCC. To solve the physics problems, the NTCC framework uses modern object-oriented computer techniques and is designed to generate transport codes that are easy to maintain, customizable, Web-invocable, and user friendly. This framework allows the rapid development of customized applications. All the software resides on and runs from the NTCC site; no software needs to be downloaded, even for a Macintosh. It also allows one to pull up and plot any of the data that have been stored on the site.

In summary, the NTCC uses five computer languages (Java, CORBA, Python, C++, and FORTRAN) in a client-server framework to provide integrated modeling code that has been validated and used to carry out new physics research. In the future, the NTCC will move to parallel processing for greater speed and will develop more physics for self-consistent integrated modeling.

McRae noted that all the major computer-intensive scientific questions have the same problem of scaling and asked why a focused effort was not made to solve this problem across the disciplines of those questions. Eckstrand said it is hoped that this is exactly what the Enabling Technology Centers will do and pointed out that many of the codes that have come from MICS are now using libraries that allow carrying solutions from problem to problem. These problems, although comparable, are not the same, however. The common problem is dealing with different rate systems simultaneously. Kritz noted that, clearly, there is an attempt being made to leverage all the efforts and use common technology and pointed out that one of NTCC people already spends a good deal of time on the ASCI project. Eckstrand acknowledged that DOE should look into the problem of scaling and hoped that, once the initiative is off the ground, some workshops will be held on the topic.

Dahlburg said that the two speakers seemed to be asking for more computer time. Eckstrand answered that that was part of the perceived solution, but tool development, the selection of the proper tools, and cross-disciplinary workshops and application teams are also needed. Dahlburg asked how much of a tokamak could be simulated now and how much will be able to be simulated in 5 years. Kritz said that the current code compares and tests transport models that are based on turbulence calculations. Under construction are an equilibrium module and an rf module for heating. What will be possible 5 years from now depends on the resources that are brought to bear on the problem. The idea is to go from wall to core, getting all the edge coupling and calculating all the turbulence. Given the requisite support, that is all possible in 5 years. Meza asked what type of processor was used for the simulation computation involving five nonlinear equations solved 1.2 trillion times, and Eckstrand replied, all 512 processors of a Cray T3E supercomputer. Meza asked what the ultimate goal was, and Eckstrand responded that a 10 teraflop computer running for a few days would be needed to simulate the electrons. Kritz noted that modeling and simulation are now separate from theory and experiment; what is important is that computerized treatments are recognized in their own right.

Giles asked them how they planned for the future. Eckstrand replied that it would be a long time before the resources are in hand to reach the goals. In the meantime, local computing and the available resources will be used to make, albeit slow, progress toward the goals.

McRae asked if the current codes can do plasma etching etc. Eckstrand replied that they could and that the ability to simulate these processes has been a significant breakthrough; similar codes are used for modeling plasma processing. Dahlburg commented that the radiation transport and equations of state are still needed but the modelers now have their hooks into the problem.

Wright called for public comment. There being none, she adjourned the meeting for the day at 4:34 p.m.

Wednesday, Nov. 1, 2001

Wright called the meeting to order at 8:30 a.m. and introduced **Daniel Hitchcock** of the MICS Division of OASCR to present a brief history of mathematics, computing, and information-sciences research at DOE. He started with the observation that computing is an essential tool in the conduct of research today. A map showed the hundreds of locations in the United States where DOE sponsors research, highlighting the essential nature of networking. Major milestones in the DOE computing activities include

- Mathematics Program (1954)
- National Magnetic Fusion Energy Computer Center (NMFECC) (1974)
- NMFECC became NERSC (1986)
- MFENET became ESnet and converted to TCP/IP (1986)
- High-Performance Computing Centers (1991)
- Next-Generation Internet (1997)

The computing initiatives that DOE has undertaken include

- High-Performance Computing Research Centers
- Advanced Computing Research Centers
- Grand Challenges
- DOE 2000
- Electricity Supply and Demand Management (which provided real-time electricity pricing over the Internet but was not continued after the second round of grants because it was seen as being outside DOE’s mission)
- DOE Next-Generation Internet Funding
- SciDAC

In FY 2000, 51% of the MICS budget was devoted to facilities; applied mathematics and computer-science research each received 17%; and the rest was distributed among computational-science education, scientific application partnerships, network research, collaboratory tools, and collaboratory pilots. The challenges the Office faces include terascale computers; terascale to petascale data handling; applied mathematics; and high-quality, reliable software. Part of the challenge of terascale computing is that so many processors are involved that code efficiency is not a problem but many CPUs will be sitting there turning electrons into heat. However, data transport is a major issue. At the same time, the users want to be working on and thinking about the science involved rather than learning to operate the computer system and its operating software. The data challenge consists of software and network problems that arise from the large number of CPUs and the high speed of input and output (I/O). The applied-math challenges come from the need for the algorithms to scale efficiently to the large number of processors. The software challenges are that computer-science researchers need to innovate and scientists need reliable software; scientific modeling and simulation codes must endure much longer than the life of a generation of computers (about 3 years; moving an application to a new generation of computers is time consuming and expensive); high-quality, supported software is expensive to develop and maintain; frameworks for software and compartmentalization are key but must maintain very high efficiency because of the large data size; and scientific research, as well as education, persistent testbeds, and software support, are required to advance the state of the art. Standardization from system to system would be desirable to allow researchers to move from one machine to another and still be able to use their codes. Stability of operating systems and libraries over time is also desirable to researchers.

MICS is an integrated program, and its budget is partitioned among applied mathematics, computer science, advanced-computing software tools, scientific-application pilots, and applications (in the material sciences, chemical sciences, combustion modeling, accelerator science, high-energy physics, nuclear physics, fusion energy, global climate, and other uses). A good deal of this work is done by NERSC and carried out over ESnet. MICS is also an integral part of the federal IT research enterprise.

In FY 2000, MICS devoted $19 million to applied mathematics by underwriting applied-mathematics research (on linear algebra, fluid dynamics, partial differential equations, optimization, and grid generation) and coupling applied mathematics to applications (e.g., in materials sciences, protein folding, combustion, high-energy physics, and string theory). In the future, it will develop additional math libraries, study the estimation of errors in simulations, and develop ways to describe features in data. Examples of its accomplishments are (1) the partitioning of meshes and assigning them to processors and (2) developing front-tracking methods for solving problems in Richtmeyer-Meshkov instabilities.

In FY 2000, MICS devoted $19 million to computer science research and tools, supporting work in systems software environment, data management, I/O, visualization, and advanced-computing tools. This work is closely tied to what happens in ASCI. Ongoing projects include work on indexing and query estimation, distributed data mining, advanced database concepts, rendering very large data sets in visualizations, high-performance storage systems, and network-attached disk caches. Notable accomplishments include (1) PNNL’s Aggregate Remote Memory Copy Interface, which optimizes the number of calls necessary to get noncontiguous data to the processor and (2) an access-coordination system that uses a caching-policy module that allows codes to run faster by determining when processes start and finish, what data they call, and whether that data are already in the cache and do not need to be read from tapes again. Future plans include work on scalable-systems management, performance measurement, benchmarks, modeling and prediction, microkernels for terascale systems, high-performance messaging, and remote memory access, all with an open-source emphasis.

In FY 2000, MICS devoted $8 million to collaboratory tools and pilots to accelerate the ability of DOE to accomplish its mission through advanced computing and collaboration technologies. Some ongoing projects that are supported are the Materials Microcharacterization Collaboratory (that allows the remote operation of electron microscopes) and the Diesel Combustion Collaboratory. These efforts require security that is managed by humans (which scales poorly) and a high quality of service and maintenance.

DOE has supported network research for a long time, devoting $2 million to it in FY 2000. It currently holds the record for data transmission, which occurs on a data link between LLNL and the Stanford Linear Accelerator Center (SLAC).

MICS operates three types of facilities: NERSC ESnet, and the Advanced Computing Research Facilities (ACRFs). With an annual budget of about $26 million, NERSC provides capability resources and professional, user-friendly services to computational scientists working on DOE projects (about 60% of them from DOE laboratories). It operates Cray, SV1, and IBM computers, and upgrades the network about every three years. ESnet provides a highly capable and reliable communications infrastructure and leading-edge network service that support DOE's missions. The ACRFs provide pioneer-capability computing for scientific applications relevant to the SC mission. They provide limited support, do not operate 7/24, and operate a program on testbeds. Their ongoing projects include a 512-CPU Linux cluster at ANL, a Compaq AlphaServer at ORNL, a 2045-processor SCI at LANL, a "Probe" high-performance storage system (HPSS) testbed at Lawrence Berkeley National Laboratory (LBNL) and ORNL, a TERA evaluation at the University of California at Santa Barbara in collaboration with the NSF and the National Security Agency (NSA), and a prototype topical center at ORNL that is connected to NERSC.

In addition, MICS devotes $2 million a year to its Computational Science Graduate Fellowship Program to support 50 doctoral students each year. It is hoped that this dollar amount will be doubled next year.

Stechel noted that the partnerships had been part of the grand challenges, which are now defunct, and asked how such partnerships will be handled now. Hitchcock said that the partnerships will be recompeted, with a special call for small-scale partners. This activity will be viewed as a new program. McRae asked how much interaction existed between DOE and financial industries and others with an interest in data storage. Hitchcock responded that commercial organizations have concerns that DOE does not have; they get small chunks of data, and DOE gets large. Therefore there is not much interaction. With the seismic (i.e., the petroleum companies), high-energy-physics, and biology communities, there has been a fair amount of interaction. McRae asked if any use was made of commercial database management systems. Hitchcock responded positively; BaBar is one such usage; however, size and structure of DOE's data tend to break or cannot be handled by such SQLs (structured query languages) as Oracle. Oliver noted that commercial relational databases are sometimes used to hold metadata.

McRae asked what happens to the graduate research fellows. Hitchcock said that a list of them appears in the backup materials distributed to the Committee. DOE has hired a number of them, and they have subsequently moved on to industry. DOE has lots of data on the students.

Kulsrud asked what DOE's long-term commitment was to HPSS and if HPSS itself is required for large-scale data problems. Hitchcock said that the installed base of HPSSs was significant and that it is designed to move large files around fast. What is needed now is the ability also to move around large numbers of small files. Kulsrud asked if any statistical studies were being conducted to see how the system is being used. Hitchcock relied that there were but that there was a tension between maintaining the current capabilities and opening new inquiries.

Meza asked if the Enabling Technology Centers were limited to SC research projects. Hitchcock replied that they would be funded by SC; if others wanted to participate, SC would be happy to talk with them.

Wright asked what the Office was planning to do and what the Committee could offer advice on. Hitchcock responded that his presentation included only what was in the President's budget, but the funds appropriated did not include everything that was asked for. Where to redistribute the available funds would be one area in which help could be provided. Connolly commented that the Office must have proposed budgets for FY 2002 and FY 2003 in the works. Oliver responded that they did not, that this was a special year, it being an election year.

A break was declared at 9:52 a.m. The meeting was called back to order at 10:14 a.m., and **Thom Dunning** was introduced. He spoke on the topic of scientific discovery through advanced computing, an activity started by the Undersecretary of Energy several years ago. Dunning began by noting that (1) microprocessor speeds double every 18 to 24 months and (2) more and more processors are being used on a single problem. These trends will continue, and new, innovative designs will be seen [e.g., processors in memory (PIM) and hybrid technology–multithread technology (HTMT)]. That notwithstanding, each of the other offices within SC has identified computing challenges that hinder scientific advance. A crisis in high-performance computing has occurred because industry was building computers that did not meet the needs (in data analysis and visualization, mathematical libraries, programming environments, scientific-data management, and problem-solving environments) of the scientific-research community. The Department's strategy to deal with this problem is to

- create a scientific-computing software infrastructure that takes full advantage of terascale computing capabilities for scientific research,
- establish a scientific-computing hardware infrastructure that supports scientific research in the most efficient and effective manner possible, and
- enhance collaboration and access to facilities and data through advances in networking technologies and the development of electronic collaboratories.

In the nineties, peak performance increased by a factor of 100; in the next five years, it will increase by a factor of 1000. But peak performance does not mean anything; efficiency has declined from 50% on the vector supercomputers of the nineties to as little as 5% on the parallel supercomputers of today. The real research challenge is software. A new generation of scientific codes is needed to model and simulate physical processes and systems; new computing and mathematics software is needed to enable the use of advanced computers for scientific applications; and this will be a continuing challenge as computer architectures undergo fundamental changes. Some of the needs in scientific computing are for

- high-fidelity mathematical models;
- better-designed computational modeling and simulation codes;
- increased functionality in vendor operating systems;
- computing-system software that accelerates the development and use of terascale scientific codes, facilitates porting of software codes among high-performance computers, and manages and analyzes massive data sets, both locally and remotely; and
- algorithms that scale to thousands or even millions of processors.

Very few people understand the process of developing software. Teams that address theory, applied mathematics, computational science, computer science, and science and engineering simulation codes are needed to tackle these problems. That such teams work is attested to by NWChem, a software package that models molecular electronic structure and molecular dynamics. It currently consists of 750,000 lines of code and is still growing. It runs on the Cray T3D/E, IBM SP2, SGI, SMP, NOWs, Sun and other workstations, and X86 PCs (Linux) and scales to more than 2000 processors. It was developed by a core group of 15 people and 20 worldwide collaborators. More than 100 person years at PNNL alone have been devoted to its development. What has this meant to the chemistry that can be done? In 1992, one could model molecules with about eight atoms; today one can model molecules with about 18 times that number of atoms.

One way to promote this teamwork would be to establish Enabling Technology Centers, where such teams would create mathematical and computing systems software to enable scientific simulation codes to take full advantage of the extraordinary capabilities of terascale computers. They will work closely with scientific-simulation teams to ensure that the most critical computer science and applied mathematics issues are addressed in a timely fashion, and they will support the full software life cycle. Out of this effort will grow dispersed collaborations that will require collaboratories and networks. Four to six teams will be competitively selected to begin development of advanced computational modeling and simulation codes. In addition, another four to six teams will be competitively selected to begin development of mathematical and computing-systems software. Three to four teams will also be competitively selected to continue development of collaboratory software. It is hoped that new funds will allow these software-development efforts to be strengthened and broadened in FY 2002 and beyond.

In the hardware infrastructure, a need was felt for

- a flagship computing facility to provide robust, high-end computing resources for all SC research programs;
- topical computing facilities with machines tailored to a particular problem to provide the most effective and efficient computing resources for a set of scientific applications and to serve as a focal point for a scientific research community as it adapts to new computing technologies; and
- experimental computing facilities to assess new computing technologies for scientific applications and to work with industry to ensure that the features needed are delivered.

Topical facilities are seen as necessary because of the wide variation in requirements for time, memory, storage, and node I/O that different scientific applications have. In such situations, the configuration of the machine must be optimized for the particular application. Such facilities would also provide the framework needed for multidisciplinary activities, allowing long-term collaborations among disciplinary computational scientists, computer scientists, and applied mathematicians.

Experimental facilities are seen as necessary because of the need for an organized approach for the evaluation of new computing technologies (processors, switches, etc.). Although computer technology is currently on a plateau vis-à-vis computer architectures, this will not last through the decade. Examples of new approaches include PIM and HTMT. An organized approach for interacting with computer designers as early as possible is also needed. Computer designers have many variables to consider; some are beneficial for scientific computing, and some are not. The earlier the scientific community can provide input, the more likely its advice will be heeded.

In FY 2001, OASCR plans to upgrade the existing flagship facility (NERSC) to 5 teraflops and to upgrade ESnet and establish a network testbed and to upgrade the ACRFs. In FY 2002 and beyond, it plans to competitively establish a number of topical computing facilities and establish up to two experimental computing facilities by recompeting the ACRF program.

Giles asked about the $10 million reduction in OASCR. Oliver responded that the President's request was not met, so where the money goes has to be reconciled. Giles commented that a lot of what was spoken about in this presentation has applications to other agencies and to industry and asked if the value of this research was obvious. Dunning replied that it was and that DOE was talking with the NSF and others to leverage this effort and to get additional funds from these other agencies in the future. He noted that the Department is open to all possible partnerships. Meza asked if there was money for the vendors. Dunning replied that one can often invest minor amounts of money and get major paybacks. This is what is being done in ASCI. Wright asked how he expected this money to be spent. Dunning said that it will go out through upgrades and through solicitations to the community for research proposals, almost all of which will be for software development and will reflect people costs. Responses would be expected from universities as well as from national laboratories. The problem areas will be identified in the solicitation, and the topics will be kept as general as possible.

McRae noted that the vision that was presented here cuts across SC but the Committee is focused on just a small portion of this vision. He asked if the Committee's charter should be changed to cover the broader subject of simulation. Oliver said that it makes sense to expand the purview of this committee or to set up links with other advisory committees.

Lester noted that, in considering peak performance, scaling is very different in, say, Monte Carlo calculations and asked if any thought had been given to alternative calculational approaches and how they might fit into the framework. Dunning said that there are two ways that one could become part of a team focused on the computational chemistry software that includes multiple approaches to the problem. One could promote the alternative approaches; if there are fundamental advances to be made, one could come in as a PI. We do not want to ignore these other approaches just because they do not fit the general model we are using. Lester noted that market forces do operate and asked if vendors are really eager to build machines for specific applications. Dunning replied that they are not interested in a lot of special-purpose machines, but they are interested in having a basic design that they can then tweak in various ways that will make it more appropriate (or less appropriate) for a particular application.

Stechel commented that it was not obvious how Dunning's office interfaced with other offices within SC. Dunning said that, when BES, OBER, and others write up their budget reports that call for computing sciences, the money will come out of the MICS budget. McRae asked what it will take to implement the vision put forward in the presentation. Dunning said that his office tries to make its request as robust as possible to the possibility of no future funding; that said, the answer to the question is a couple of hundred million dollars. The potential payback is enormous, as in the design of accelerators, to say nothing about increases in scientific knowledge.

Wright then introduced a series of three research talks, the first by **Ian Foster**, who spoke about research on Grid computing. The Grid, he explained, is the Web on steroids. The Web has made possible the uniform access to HTML (Hypertext Markup Language) files. The Grid is a set of protocols that allow high-performance access to *all* significant resources (e.g., software catalogs, computers, colleagues, data archives, and sensor nets). Brought together, these resources produce an on-demand, powerful, virtual computing system.

Grid computing will enable communities to share geographically distributed resources as they pursue common goals in the absence of central control, omniscience, and trust relationships. Grid computing seeks to find out what new applications become possible when resources can be shared in a coordinated way. Such sharing will be facilitated by protocols, algorithms, and persistent infrastructure.

DOE runs unique and expensive facilities, such as accelerators, microscopes, and supercomputers, that cannot be replicated. Its research is large-scale, multidisciplinary science on climate, materials, high-energy physics, and other computationally intensive applications that are rarely geographically collocated. The question is not *whether* to Grid-enable DOE science but *how*. Data grids for high-energy physics are central to their ability to analyze data from such places as the Advanced Photon Source (APS).

Grid R&D started at ANL in 1995 with the I-WAY experiment, which created a national-scale grid infrastructure. This project led to formulating the Globus R&D project, much of it conjointly with the Information Sciences Institute of the University of Southern California (USC/ISI). These efforts have developed

- innovative security, resource-management, data-access, information, communication, fault-detection, etc. technologies;
- a large user base among tool developers;
- widespread adoption in "production" grids (e.g., the National Aeronautics and Space Administration's IPG, NSF's NTG, and DOE's DISCOM); and
- exciting application demonstrations.

Another set of relevant achievements are the access Grid collaboration technologies that incorporate designed spaces for group interactions, hands-free audio, multiple video and audio streams, and a wide field of view.

Some cross-cutting technical issues that need to be addressed include the development of

- Grid protocols and services, such as authorization-protocol-mediated access to remote sources, resource brokering, and speak Intergrid protocols and
- Grid application-program interfaces and software development kits (APIs and SDKs) to facilitate application development by supplying higher-level abstractions.

These protocols are mostly extensions to existing protocols. The model here is the Internet. But the Grid is not a distributed operating system.

A layered Grid architecture has been developed and is being used. It starts with controlling things locally (the fabric that allows access to and control of resources), protocols for talking to things (the communication and security capabilities), the sharing of resources (including negotiating access and controlling use), ubiquitous infrastructure services (that manage multiple resources), and user- or application-specific specialized services (which are distributed).

In this hierarchy, the connectivity-layer protocols and services take care of not only communication but also security, providing

- uniform authentication, and authorization mechanisms in a multi-institutional setting and
- single-sign-on, delegation, and identity mapping mechanisms.

The resource-layer protocols and services include

- Grid resource allocation management that provides for remote allocation, reservation, monitoring, and control of computer resources;
- ongoing work on a Grid file-transfer protocol (Grid FTP) to produce high-performance data access and transport;
- a Grid resource information service that provides access to structure and state information; and
- network reservation, monitoring, and control.

All of these resource-layer protocols and services are integrated with the Grid security infrastructure. The ubiquitous infrastructure services include

- index servers (i.e., metadirectory services) that provide custom views on dynamic resource collections assembled by a community,
- resource brokers,

- replica catalogs, and
- coreservation and coallocation services.

Data-grid R&D is designed to enable a geographically distributed community to pool its resources to perform sophisticated, computationally intensive analyses on petabytes of data. The framework of the network quality-of-service (QoS) research is designed to produce secure, policy-driven bandwidth allocation for high-end applications, to provide immediate and advance reservations, and to distinguish between coreservation and coallocation of multiple resources for end-to-end flows. All of these QoS activities are supported by the Globus Toolkit, a modular architecture widely used for Grid applications.

The team conducting this QoS research has experience in differentiated services, sees future opportunities in all-optical networks, and is engaged in a collaborative effort with LBNL. QoS controlled experiments are being funded by Cisco Systems, DOE, and the Next Generation Internet Initiative and are looking at bulk transfer on local-area networks (LANs) and wide-area networks (WAN, ESnet). The results of this testing indicate that the development team can deliver the product that was promised to the sponsoring programs. Some new capabilities are needed, but a good deal of existing technology can be reused. An application that is being worked on is making climate-model data sets available to a wide user group.

Grid computing represents a significant success story for DOE research and an opportunity for accelerating DOE science. It is now timely to take this project to the next level by

- focusing on making key DOE applications (climate, physics, combustion, etc.) Grid enabled;
- looking at new efforts focused on security, next-generation optical technologies, Grid tools, etc.;
- developing new infrastructure, such as faster networks, protocols, and security certificates; and
- envisioning the evolution of ESnet to ESgrid.

Wolff noted that GARnet does not have a connection to QBone. Foster said that was correct, but that the proposal was to use Globus to link physics researchers and that using some QoS protocols is being discussed. As a result, a fair amount of work must be done before they can be worked with. When the Internet2 protocols are set, that will force some of these issues. Wolff asked how the ESnet would be converted to Gridnet. Foster replied that a certificate authority would be set up, and then QoS would be transitioned into ESnet. Connolly noted that a lot of these recommendations are being adopted by PACI (NSF's Partnership for Advanced Computing Infrastructure) and that should make it easy for ESnet to do that. He asked if any plans were being made, and Foster answered that they were. McRae asked how Foster saw commercial software fitting in. Foster responded that one way was to focus on this as a protocol problem and then to take these improvements to industry standards groups like the IETF (Internet Engineering Task Force). McRae asked if XML was made use of, and Foster said that it was and that it is one way to deliver such resources.

Wright asked Foster if he was happy with the OASCR plans for this work. He said that there is great promise in funding the computer-science, computational-science, and application researchers in parallel. The question is whether it will be done with the proper balance. Certainly, more can be done, but we are off to a good start.

Wright declared a break for lunch at 11:51 a.m. The meeting was called back to order at 1:20 p.m. with the introduction of **Marsha Berger** to speak on high-resolution adaptive methods for complex flows, the goal of which is to automate the computation of high-resolution simulations in realistic engineering applications. Enhancing and speeding up these simulations can be accomplished by geometric specification, mesh generation, algorithm development, numerical discretizations, adaptive techniques, robust software, and adaption to high-performance (parallel) computing.

One technique that is used is adaptive mesh refinement, in which recursively nested, locally refined, block-structured grids are used to raise the level of accuracy. In this technique, one uses the same integrator to advance the solution on all grids (both coarse and fine); one employs smaller time steps; and one has to develop stable, accurate, and conservative interface connections across the grid. One may also employ an automatic error estimator, an automatic grid generator, and/or a simple data structure. This technique can save orders of magnitude in the size of the task and the time to complete it. A lot of work is being performed on other partial differentials related to incompressible flows, reactive flows, etc., and that work could be applied to error estimation, implicit schemes, directional refinement, and software and parallelization issues.

A major question that arises is how to extend such work to complicated geometries. Berger and her colleagues have looked at developing automatic methods for rapid-turnaround flow computations in complex geometries. In such computations, one needs a closed, watertight surface description, which involves the use of

- computer-aided-design (CAD) definition and surface-definition software, which together take months of work;
- volume-mesh generation, which takes a day to a week;
- flow computation, which takes less than a day; and
- postprocessing.

The main bottleneck here is geometry acquisition.

A number of alternative approaches are available. Structured meshes are the most accurate and most efficient means of calculating flow over a body but are the most difficult to generate because they must be mapped to Cartesian coordinates. Unstructured meshes (which are body fitted to the geometry) handle complex domains and are easy to program but incur high overheads in memory and CPU usage. Cartesian non-body-fitted grids (which use regular Cartesian grids with solid objects cut out of the underlying grid) are more accurate, efficient, and easier to develop than the other types of meshes and they transform the problem into a simple geometric calculation. Cartesian meshes simplify the description of the surface, and because irregularities are confined to lower-dimensional space, the user does not pay an efficiency penalty over the entire domain. In addition, grid generation is easier because the surface grid is not the computational grid; the surface description resolves the geometry, and the Cartesian mesh describes the flow. The drawbacks are that

- Cartesian grids lack the resolution of body-fitted or unstructured grids, requiring the use of adaptive mesh refinement;
- irregular cells produce a loss of accuracy at the boundary; and
- small-cell instabilities are produced that call for new numerical discretizations.

What is needed are stable, accurate schemes with a CFL (Courant-Friedrich-Levy) number based on regular cells. However, one rarely gets closed, water-tight surfaces in multiple-component geometries (e.g., the wing, pylon, and engine of an airplane). The question is how to put the pieces together. It would be desirable to allow separately defined, water-tight component triangulations as input to mesh generation. The problem then becomes how to compute topologically consistent surfaces in the presence of floating-point roundoff error. Degeneracies make this task very complicated. In essence, one has two (irregular) surfaces and must come up with an intersection algorithm. One strategy is to

- find the intersect-component triangulations,
- retriangulate the intersected triangles,
- remove the internal geometry, and
- resolve degeneracies with adaptive precision-determinant computations and virtual-perturbation tests.

With such a strategy, a triangle intersection boils down to multiple computations of a four-by-four determinant of the signed volume of a tetrahedron. That determinant can be computed with floating-point arithmetic with a floating-point filter (if the result is less than the error bound, recompute with adaptive-precision exact arithmetic; if exact arithmetic gives a value of zero for the determinant, resolve the degeneracy with a tie-breaking algorithm), which is a virtual-perturbation approach.

Space-filling curves can be used to linearly order a multidimensional mesh by Peano-Hilbert or Morton ordering, and domain partitioning can be carried out by using work estimates based on cell types (cut or full-body) to partition a given mesh into load-balanced domains. This last technique can be done on the fly and can produce any number of partitions.

At this point, modelers are in a position to achieve large payoffs in scientific and engineering applications. Some of the remaining needs are more algorithm development, dealing with some modeling issues, and anisotropic refinement.

Giles asked if she saw the possibility of doing mesh generation dynamically as an extension of CAD. Berger replied that the CAD is just not good enough at this time, but work is being done on achieving that goal; some intermediate technique may need to be developed.

Wright asked Berger how she was connected to DOE and SC. Berger responded that she had a long-term collaboration with Phil Colella and John Bell at LBNL. A close connection exists in applying this technique to DOE problems. Additional work is being done in the Overture group at LLNL that uses several of these techniques and tools. Results coming out of all these efforts are being shared.

Meza stated that he thought that unstructured meshes were the method of choice for very complex geometries and asked if the technique Berger described would replace unstructured meshes. Berger replied that it certainly can. Users of unstructured meshes now use hexes instead of tetrahedrons for many of the same reasons that her group prefers regular grids. But the two approaches have many of the same difficulties. In the Cartesian framework, anisotropic abundance is difficult and is still one of the open problems. Similarly, with the Cartesian framework, heavily anisotropic triangles or tetrahedrons are very hard to generate with the high aspect ratios of, say, aerodynamic boundary layers. Unstructured meshes cannot do that, either; they have many of the same floating-point-error issues. It is also very hard to make methods that have tests that do not enter diffusion. One gets a diagonal that is not quite aligned and tries to get rid of it by going to hexes or prisms. So, for anisotropic flows, Cartesian grids have the same difficulties; and for inviscid flows, they have had a lot of advantages. In the past 10 years, a lot of people have adopted the approach, so a body of experience is being built up that shows the combination of embedded boundaries and Cartesian grids to be a good alternative.

David Greenberg was introduced to speak about bridging gaps. He said that computer science should be bringing better architectures to scientists. The current emphasis on peak performance is great for certain applications but not for others. What is needed is more precision for larger problems; larger parameter space; adaptive, irregular meshes; applications that are resource aware; and parallel processing. For typical science codes, parallelism is easy to find but harder to express. The challenge is to balance the work and, at the same time, to not move the data. He did not think that the right expression had been found yet. The question remains, how does one choose the next task to do, decide what processor to use, encapsulate input and output, and monitor progress, all while keeping work near its data, performing I/O, and scalably dealing with fault tolerance.

The computing community has had some successes in developing the libraries PVM (Parallel Virtual Machine) and MPI, the distributed-memory model, independent processes, explicit data movement, and library interfaces. It has also had successes in standardization and codification, vendor acceptance, portability, bandwidth for large messages, latency [compared with TCP (the Transmission Control Protocol)], collective routines, and communicators.

Compilers should be looked at. A step forward has been the adoption of Unified Parallel C (UPC) by some vendors. This compiler is also available in open source. This tentative first step needs to mushroom. UPC adds to C a "shared" keyword that tells the compiler to distribute. It still has the problem that allocations and barriers are collective operations. A library, Global Arrays, is being established that is data parallel and fits into the SALC (shared-address, local-caching) memory model; it matches particular programmer needs and provides multiple ports through the use of a layered model.

To meet the computational needs of the scientific community, the design of a computer should start with those needs and build a machine from commodity parts, filling in with custom parts. Such a machine would grow like a plant with the needs of the computing community, would employ hierarchical storage space, and would always be at the leading edge. Like a power plant, it would fuel a revolution in simulation. In addition to the DOE laboratories, Cplant, Chiba City, HPTi, and TurboLinux are also working in this area. People have realized that the high-performance computing industry is precarious and that it needs to leverage commodity parts and build on the Beowulf trend. However, future machines must also have much more scalability and manageability. Scalability is key because the trend is toward lots of users, nodes, and types of programs. The response to this issue must be to go to inexpensive parts, hot-swappable parts, and interoperable parts, where everything plugs into the network. For service nodes, compute nodes, and I/O nodes alike, some of the system must be supplied by the vendor, and the performance modules must be custom made.

Ongoing efforts are being pursued in the Extreme Linux Forum, where some major endeavors have come up: surmounting BIOS/startup (the Basic Input/Output System) issues, scalable system management, and reasonable I/O. But Linux has issues, too, a lot of which have been solved before. Linux allows one to plug in project-specific capabilities, such as customizing partitions for compute vs I/O vs login and customizing for high-performance scientific computing through the use of big pages, page coloring, and Myrinet and Quadrics drivers to make network interfaces faster.

In summary, the architecture, programming model, applications, and libraries all have to be developed together.

Dahlburg asked if the hardware was developing faster than the theory. Greenberg responded that having a lot of possibilities going on is good. However, a lot of machines are out there along with a lot of ideas of how to do systems management. It is time for this situation to coalesce into a smaller number. Kulsrud asked if this Office should be supporting radical new languages and compilers. Greenberg said that someone should be funding that but, for the scientific programmer, someone should be adding to the available tools and setting guidelines for syntax so the desirable changes can be made. Some of the problems are language problems, and some problems are in the underlying infrastructure.

Meza asked what the top three barriers to scalability were. Greenberg responded:

1. 32-bit addressing,
2. keeping a table of things coming from different processors, and
3. managing I/O without a lot of buffering.

Wright declared a break at 2:55 p.m. The Committee was called back into session at 3:10 p.m., and the floor was opened to discussion by the Committee members of the presentations that had been made. Meza asked what a topical center would look like. Dunning said that DOE would look at whether an organization would benefit from customized architectures. In other instances, such centers might have customized configurations rather than customized machines. Others might be formed around a software base. A peer-reviewed competitive process would be used to determine who would get awards to establish or operate such centers. Meza asked if there would be a one-to-one mapping of disciplines and centers. Dunning replied that that might happen sometimes, but several disciplines might use the same machine sometimes, and one community might use several machines sometimes. DOE wants to put in place the scientific teams that can exploit the available resources.

Dahlburg noted that there was NERSC and there were special machines and that somewhere there had to be a cutoff where a center should not spend money and the funds should go into a pool. Dunning agreed that there will be a tension in the process of allocating funds. Dahlburg asked if there were any standards to govern that process. Dunning said that part of that standardized approach will be the review process itself. Part of the challenge will be making the applicants go through an analysis and defense of their proposals. Giles asked if others would be able to use these facilities. Dunning responded that DOE would probably require each of the centers to make a certain amount of time available to others. Making OASCR the center for funding would assure that these new facilities could not be provincially shielded from others' use. Also, with the grid, users will not have to be located at the site of the facility. The problem is, who should be included. Lester commented that the term "topical" may be restrictive and that there may be a better term to describe these centers. Meza said that having a center available will free up time to develop new and better scientific techniques but there is the danger that centers might limit the techniques used in the future to techniques that match the capabilities of the machine and software.

Giles asked if NWChem was developed under an open-source model. Dunning replied that it used a variation of open source in that it can be altered but then the altered code has to come back and be reviewed by the controlling community before it is released to the public. When it comes to the application domain, one has to worry about the quality. Thus, under this model, anybody can change the code, but they have to go through a validation and verification step.

Wright noted that open-source software was evolving, so one needs to be aware of the issues. She asked if there was any public comment. **Daniel Hitchcock** said that, in the plans for FY 2001, there is no place for funding of topical centers, so there will be plenty of time to get these software teams up and running. **Robert Marianelli** commented that interactions between this committee and other offices of SC was very important. He wanted to make sure that the Committee understood that the increment for advanced computing is part of the larger IT research funding with a well-defined interagency working group that includes NSF, NASA, DOD, etc.

There being no further public comment, Wright adjourned the meeting at 3:32 p.m.

Respectfully submitted,
Frederick M. O'Hara, Jr.
Recording Secretary
Revised by Margaret Wright, 1/19/2001