Final Report of the
Advanced Scientific Computing Advisory Committee
Office of Science, Department of Energy

High-Performance Computational Needs and Capabilities in the Office of Science

April 6, 2004

Summary

In August 2001, Dr. James Decker, then Acting Director of the Office of Science, asked the Advanced Scientific Computing Advisory Committee (ASCAC) to assess high-performance computational needs and capabilities throughout the Office of Science, commenting on quality, impact, adequacy, management issues, funding requirements, and measurements of progress. This report, ASCAC's response to that charge, is based on a variety of activities related to high-end computing, including meetings, workshops, and reports in which ASCAC has been involved.

Our four conclusions are:

- Although current Office of Science computing facilities are among the best worldwide, the Japanese Earth Simulator is likely to remain the world's most powerful and effective high-end computing facility for several years.
- Current and expected near-term high-end computing resources in the Office of Science are far from adequate to meet the anticipated needs of its science and engineering missions and the associated scientific communities.
- The Office of Science SciDAC (Scientific Discovery through Advanced Computing) program is a world-class exemplar of the success that can be achieved through creation of focused multidisciplinary teams of application scientists, applied mathematicians, and computer scientists.
- Opportunities abound in every mission area of the Office of Science for major scientific advances from effective application of mathematical modeling and computational simulation.

Our five recommendations are:

- Substantial investments are required to begin to return the United States to leadership in high-end computing. (Suggested funding levels are discussed in the body of the report.)
- ASCAC strongly supports coordinated Office of Science programs that tightly couple science, high-end computing, and applied mathematics.

- The Office of Science should manage its advanced computing resources as a single, coordinated facility.
- Investment decisions in high-end computing should be guided by the importance of the associated science drivers along with the opportunities for cross-fertilization and common usage among scientific disciplines, mathematics, and computer science.
- All programs in the Office of Science should explicitly consider the roles of high-end computing and mathematical modeling in current and emerging research challenges.

1. Charge, Conclusions, and Recommendations

This report of the Advanced Scientific Computing Advisory Committee (ASCAC), Office of Science, Department of Energy, is one of a sequence of ASCAC reports and statements concerning the needs and capabilities for high-end computing in the Office of Science.

1.1. The High-End Computing Charge

The specific charge addressed by this report arises from a letter dated August 26, 2001, to Dr. Margaret Wright, ASCAC chair, from Dr. James Decker, then Acting Director of the Office of Science. That letter requested

> that the Advanced Scientific Computing Advisory Committee (ASCAC) form a composite panel to assess the high-performance computational needs and capabilities throughout the Office of Science (SC). In addition to ASCAC, the panel should draw its membership from the other five Advisory Committees to the Office of Science, namely, the Basic Energy Sciences Advisory Committee, the Fusion Energy Sciences Advisory Committee, the Biological and Environmental Research Advisory Committee, the DOE/NSF High Energy Physics Advisory Panel, and the DOE/NSF Nuclear Science Advisory Committee. The panel will provide recommendations and advice to the Director, Office of Science, on high performance computing needs, management issues, and funding requirements.

Dr. Decker's letter made the following observations:

> High performance computing has become an increasingly important element of the Department of Energy's Office of Science, and its predecessor organizations, for the past 50 years. Today, high performance computing is widely recognized and regarded as an essential tool for enabling science and is an indispensable, integral part of virtually all SC research endeavors. . . . Your assessment will help ensure that SC's high performance computational needs and capabilities are fully integrated into our process of analyzing and planning programs and operations. The panel should carry out a broad assessment of the status and the prospects for high performance computing with respect to all five SC programs.

The charge letter asked specifically for a detailed discussion of the following topics:

Topic 1. The overall quality of high-performance computational and networking facilities throughout the Office of Science;

Topic 2. Benchmarking this quality by a comparison with similar facilities supported by organizations that support science in the context of a mission agency, both domestically and internationally;

Topic 3. The impact and effectiveness of interactions and resource sharing among Office of Science high performance computational and network facilities; the level and adequacy of funding provided by all Office of Science programs for high-performance computing and networking facilities compared with other needs; and the effectiveness of the current distribution of high performance computational and networking resources across the Office of Science complex;

Topic 4. The evolution of the roles of these facilities and/or their distribution over the next 3–5 years so that SC programs can meet their high performance computational needs and maintain their national and international scientific leadership;

Topic 5. Useful metrics to measure progress and guide investment decisions in the area of computing and networking.

Dr. Decker's letter requested a report of findings and recommendations by September 1, 2002. Dr. Gregory McRae, a member of ASCAC, agreed to chair the ASCAC subcommittee to fulfill this charge. (The subcommittee will hereafter be called the High-End Computing Subcommittee.) Dr. Decker's letter asked each of the chairs of the other five advisory committees for the Office of Science to designate a member of the subcommittee. These subcommittee-designated members were:

- Charles DeLisi, Biological and Environmental Sciences Advisory Committee (BERAC);
- Bill McCurdy, Basic Energy Sciences Advisory Committee (BESAC);
- Vincent Chan, Fusion Energy Sciences Advisory Committee (FESAC);
- Paul Avery, High Energy Physics Advisory Panel (HEPAP);
- Tony Mezzacappa, Nuclear Sciences Advisory Committee (NSAC).

Since the formation of the High-End Computing Subcommittee, its work has been affected and delayed by related and overlapping events in ASCR and in the Office of Science, some of the latter are described in Section 4. As a result, the other five programs in the Office of Science have been involved with the High-End Computing Subcommittee only implicitly, through the activities discussed in Sections 4.2, 4.3, 4.4, and 4.6, rather than explicitly as in Dr. Decker's original charge.

1.2. Conclusions

Conclusion 1. *The current Office of Science high-end computing facilities, including computers and networking, are among the best worldwide. However, the Earth Simulator in Japan is, and is likely to remain for several years, the most powerful and effective high-end computing facility in the world.*

Conclusion 1 addresses Topics 1 and 2 of Dr. Decker's charge; also see Section 2.

Conclusion 2. *Current and expected near-term high-end computing resources in the Office of Science are far from adequate to meet the anticipated needs of its science and engineering missions and the associated scientific communities.*

Conclusion 2 addresses Topic 3 of Dr. Decker's charge, and responds to his question about high-end computing needs and funding requirements; also see Section 3.

Conclusion 3. *The Office of Science SciDAC (Scientific Discovery through Advanced Computing) program is a world-class exemplar of the success that can be achieved through creation of focused multidisciplinary teams of application scientists, applied mathematicians, and computer scientists.*

Conclusion 3 responds to Dr. Decker's question about management of high-end computing in the Office of Science; also see Section 4.1.

Conclusion 4. *Opportunities abound in every mission area of the Office of Science for major scientific advances from effective application of mathematical modeling and computational simulation.*

Conclusion 4 responds to Dr. Decker's question about the needs for high-end computing; also see Sections 4.2, 4.3, 4.4, and 4.6.

1.3. Recommendations

Recommendation 1. *Substantial investments are required to begin to return the United States to leadership in high-end computing.*

Over a three-year timeframe,

- for leadership-class computing we recommend an investment of $300 million;
- for capacity computing (from a combination of leadership-class facilities and continuous upgrades of user support software), we recommend an investment of $200 million; and
- for essential experimental research in networking, computer hardware, and software for applications and system support, we recommend a separate conceptual "facility" funded at $100 million.

Recommendation 1 responds to Dr. Decker's question about funding requirements; also see Section 3.

Recommendation 2. *ASCAC has previously urged, and we once again strongly support, a coordinated Office of Science program that tightly couples science, high-end computing, and applied mathematics.*

Building on the traditional strengths of the Department of Energy, such an enterprise would be certain to produce important scientific advances. The SciDAC program, which already has a solid track record of success, would be an excellent foundation for such a program. (See Section 4.1.) New funding of at least $30 million per year is needed across the programs of the Office of Science to ensure that such a program can achieve the requisite levels of innovation and originality, which cannot be created by piecemeal funding in small increments.

Recommendation 2 responds to Dr. Decker's question about the needs for, management of, and funding requirements for high-end computing in the Office of Science.

Recommendation 3. *The Office of Science should manage its advanced computing resources as a single, coordinated facility.*

In particular, users should be able to treat high-end computing facilities as a seamless entity, and the Office of Science should develop an integrated allocation strategy for all of its high-end computational resources.

Recommendation 3 addresses Topic 3 of Dr. Decker's charge as well as his question about management; also see Sections 2 and 3.

Recommendation 4. *Investment decisions in high-end computing should be guided by the importance of the associated science drivers along with the opportunities for cross-fertilization and common usage among scientific disciplines, mathematics, and computer science.*

Before making a significant investment in special-purpose hardware tuned to today's algorithms or codes, we recommend a careful study of whether general yet application-specific techniques—for example, very recent research on core matrix calculations that exploit hierarchical memory structures[1]—can be applied.

Recommendation 4 addresses Topic 5 of Dr. Decker's charge; also see ASCAC's statements on a major computing initiative, made in May 2002, and on topical computing, made in January 2003.

Recommendation 5. *All programs in the Office of Science should explicitly consider the roles of high-end computing and mathematical modeling in current and emerging research challenges.*

This kind of examination has already been done, or is being done, in Basic Energy Sciences Research (Sections 4.2 and 6.2), Fusion Energy Sciences Research (Section 4.3), and Biological and Environmental Research (4.4). The SCaLeS report (Section 4.6) summarizes an array of applications across the Office of Science, along with the expected benefits from advanced simulation capabilities. Even so, a more detailed study would be useful for the topics that are most dependent on simulation, real-time data collection, data analysis, and visualization.

Recommendation 5 addresses Dr. Decker's question about the needs for high-end computing.

2. ASCAC Subcommittee on Facilities

In April 2002, a formal announcement was made of the Japanese "Earth Simulator" computer, then (and still) the world's fastest computer. The Earth Simulator's substantial superiority in sustained speed to even the fastest machine available in the United

[1]E. Elmroth, F. Gustavson, I. Jonsson, and B. Kågström, Recurive blocked algorithms and hybrid data structures for dense matrix library software, *SIAM Review* 46, 3–45 (2004)

States made headlines around the world, and raised many questions for the high-end computing community in the United States.

At the time of the announcement of the Earth Simulator, the ASCAC Facilities Subcommittee, chaired by Dr. Jill Dahlburg, had almost completed its comprehensive report on Office of Science computing facilities. Because of the obvious but still poorly understood impact of the Earth Simulator, ASCAC requested of the facilities subcommittee that its final report should include an assessment of the implications of the Earth Simulator for American leadership in computational science. To provide information for this assessment, an "Earth Simulator Rapid Response Meeting" was held at the request of Dr. Raymond Orbach, Director of the Office of Science, on May 15–16, 2002. Complete details of this workshop may be found at `www.ultrasim.info/esrr_meeting`.

The final report of the ASCAC Facilities Subcommittee, submitted to Dr. Orbach on May 31, 2002, contained five major conclusions concerning the Earth Simulator, which we summarize as:

ES-1. The Earth Simulator simulation capability is real, and provides significantly enhanced performance on real science applications;

ES-2. The Earth Simulator is not a surprise, but rather results from the decisiveness, commitment, and accountability of the Japanese government;

ES-3. The Office of Science is well positioned to lead the United States back to the front rank of computational science;

ES-4. The Office of Science research community, in partnership with domestic vendors, is ready to respond to the Earth Simulator challenge;

ES-5. The Office of Advanced Scientific Computing Research is significantly underfunded to address this goal adequately.

Topics 1–5 of the charge to the High-End Computing Subcommittee are primarily concerned with facilities and are answered in large part by the final report of the Facilities Subcommittee, especially by that subcommittee's careful attention to issues associated with the Earth Simulator. For details, see

`www.krellinst.org/esinfo/ASCAC_facilities_final.mhw.doc`.

3. ASCAC Subcommittee on the Role of Large Facilities

In a letter dated December 18, 2002, Dr. Orbach requested that Dr. Wright form an ASCAC subcommittee to consider several questions related to the construction of *large* facilities, especially new ones—specifically, which new or upgraded facilities would be necessary to position ASCR at the forefront of scientific discovery. (An analogous request was made of all Office of Science Advisory Committees.) Dr. Orbach asked that the subcommittee consider only facilities costing at least $50 million, and that it discuss the importance of the science supported by the facility. Dr. Orbach's charge asked the subcommittee to classify each facility as one of "absolutely central", "important", or "don't know enough yet"; he also asked about the readiness of the facility. For ASCR, the subcommittee was to consider discovery and potential accomplishments in applied mathematics and computer science as well as in connections with core DOE scientific missions.

Helene Kulsrud of ASCAC agreed to chair the ASCAC subcommittee, and convened a workshop on March 2–3, 2003, in Princeton. The preliminary report of that subcommittee, available at

`www.sc.doe.gov/ascr/ASCACRecommendations%20report.doc`,

contains the following conclusions:

LF-1. It is clear from current events that today's resources cannot support the Office of Science computing needs. For example, SciDAC researchers requested twice as many computer hours for their projects as they were allocated, and they expect to require more as the projects mature. Furthermore, new initiatives such as the Fusion Energy Simulation Initiative, the Nanoscience Master plan, and the Genomes to Life program envision the need for substantial new computational resources that far exceed any currently planned facilities.

LF-2. The Office of Science should invest in the creation and maintenance of at least one "Leadership Class" computing facility that will support high-capability computing. Leadership class systems, which will enable and sustain computation at a world-class level, must be reinvested and renewed with approximately a three-year lifecycle. This facility, which require advanced computation hardware and software support, is required to answer the difficult scientific questions posed by the old and new scientific initiatives in the Office of Science. We expect that acquisition, maintenance, and operation costs for such a facility would amount to about $300 million in new funding over a three-year timeframe, and that at least this same amount would be needed in the subsequent three-year periods. We consider this facility to be "absolutely central" and to depend on the availability of "leading edge hardware".

LF-3. The Office of Science needs to invest in capacity computing to satisfy the needs of many researchers who individually do not need the dedicated capabilities of the leadership-class systems but whose aggregate need is substantial. It is essential to provide such capacity in order to be able to sustain the effective use of the leadership-class systems on those problems for which they are uniquely suited (instead of timesharing them out into infinitesimal fragments to many people). This need can be addressed by an appropriate combination of continued use of leadership-class facilities and continuous upgrade of user support software. We estimate this facility at $200 million in new funding over a three-year timeframe. We consider this facility to be "absolutely central" and "ready to begin".

LF-4. It is essential for the Office of Science to support experimental research in networking, computer hardware, and software for both applications and system support. To support the needed petaflop environments, it is especially important that we provide experimental systems at a scale that enables researchers to assess how algorithms, protocols, applications, etc., scale. A separate facility for this research is advised. We estimate the cost of this activity at \$100 million over a three-year timeframe. We consider this facility as "important" and "ready to begin".

LF-5. The advanced computing facilities, though they may be physically distributed among the labs and centers, must be managed and coordinated as an Office of Science resource as are other major DOE facilities. In addition, it is quite possible that these facilities could be part of a larger national open computing infrastructure that could be used by researchers across the country. DOE's ability to provide such facilities is what makes it unique among science agencies. This is nowhere more important than in computing.

LF-6. From the viewpoint of users, high-end computing facilities should be viewed as a coordinated, seamless entity. A user or group of collaborators should be able to have an integrated view of computing resources, linking components at leadership class systems, development systems, and distributed systems. Also, all offices under the Office of Science should participate in allocation and planning of high-end computing, including any special-purpose computers.

These conclusions of the Subcommittee on Large Facilities address not only Topics 3 and 4 of the charge to the High-End Computing Subcommittee (see Section 1.1), but also the language in that charge requesting advice on "high performance computing needs, management issues, and funding requirements".

4. Computing and Science in the Office of Science

Dr. Decker's August 2001 letter requests significantly more than an assessment of SC facilities (the topics of Sections 2 and 3). He asks for "a broad assessment of the status and the prospects for high performance computing with respect to all five SC programs". Since August 2001 (and even before that date), several Office of Science workshops and reports have looked at the likely progress in science that could be achieved with a balanced, strong program in high-end computing. We next describe a selection of those activities, noting any involvement by ASCAC. (To keep this report reasonably short and focused, we have chosen not to discuss numerous recent studies by the National Research Council highlighting the importance of computing in an array of scientific and engineering fields; the associated reports can be found on the NRC Web site, `www.nationalacademies.org/publications/`.)

4.1. The SciDAC program

Starting in December 2000, the Office of Science announced a new and ambitious five-year program called "Scientific Discovery through Advanced Computing" (SciDAC); see `www.sc.doe.gov/production/grants/ilb01_06.html`.

The goal of SciDAC is, broadly speaking, to develop advanced mathematical modeling and computational simulation capabilities across the five programs of the Office of Science. Specific examples of scientific problems for which high-end computing is essential to progress include:

- combustion, environmental fate and transport, catalysis, chemical processing, and nanoscale phenomena (Basic Energy Sciences);
- prediction of climate from decades to centuries (Biological and Environmental Research);
- macroscopic stability and microscopic turbulence, basic plasma science, and inertial fusion energy applications (Fusion Energy Sciences);
- electromagnetic field and beam dynamics in particle accelerators, understanding the Standard Model of Particle Physics, and predicting the structure of nuclei in energetic events (High Energy and Nuclear Physics); and
- mathematical models and methods, simulation codes, data management and visualization, and scalable systems and software for high-end computing (Advanced Scientific Computing Research).

The SciDAC program has been, in the view of ASCAC, a phenomenal success. Its philosophy of creating deliberately multidisciplinary teams has been a visible proof of concept for a longstanding style of research in the Department of Energy—a style advocated in numerous high-level reports on science policy from the National Research Council.

The major features of SciDAC-supported research can be seen from the instructions for the most recent meeting (March 23–24, 2004) of SciDAC principal investigators. Some of the questions that speakers were asked to address were:

> How has the opportunity to be part of the SciDAC initiative enabled you to accomplish things that would otherwise not have been possible?
>
> What specific algorithmic or software tools or technology are you now able to provide to scientists to enable them to advance their research through the use of terascale computing?

See `www.csm.ornl.gov/workshops/DOE_SciDAC/` for further details.

ASCAC has been enthusiastic about SciDAC since its inception. A resolution passed unanimously at the ASCAC meeting of April 10, 2003, states:

ASCAC strongly commends the Office of Science, especially the Office of Advanced Scientific Computing Research, for conceiving, initiating, and implementing the SciDAC program, which serves as an exemplar of twenty-first century partnerships connecting theoretical, experimental, and computational science and engineering.

Bringing this resolution up to the present, we wish to state for the record that we believe that the success of SciDAC clearly and unambiguously deserves additional new funding—not funding at the expense of other programs. One of the recommendations of this report is for added funding for SciDAC or an immediate descendant of SciDAC; see Recommendation 2 in Section 1.3.

4.2. BESAC–ASCAC Workshop on Nanoscience

The National Nanotechnology Initiative (NNI) is a federal multiagency program designed to support research to exploit the full potential of nanotechnology. Within the Office of Science, progress in nanoscience has been seen as closely related to high-end computing and mathematical modeling. In particular, a workshop called "Theory and Modeling in Nanoscience" was held in May 2002, , cosponsored by ASCAC and BESAC, cochaired by Bill McCurdy (BESAC) and Ellen Stechel (ASCAC).

The executive summary of the workshop report states:

> During the past 15 years, the fundamental techniques of theory, modeling, and simulation have undergone a revolution that parallels the extraordinary experimental advances on which the new field of nanoscience is based. . . . A clear consensus emerged at the workshop that without new, robust tools and models for the quantitative description of structure and dynamics at the nanoscale, the research community would miss important scientific opportunities in nanoscience. The absence of such tools would also seriously inhibit widespread applications in fields of nanotechnology ranging from molecular electronics to biomolecular materials.

The report details numerous ways in which high-end computing is essential to progress in nanoscience. See `www.sc.doe.gov/bes/Theory_and_Modeling_in_Nanoscience.pdf`.

4.3. The Fusion Simulation Project

High-end computing has been explicitly recognized as having an important role in fusion research as well. The Fusion Energy Sciences Advisory Committee (FESAC) received a charge in February 2002 asking it to assist the Office of Fusion Energy Sciences Research in preparing a roadmap for a joint initiative with the Office of Advanced Scientific Computing Research. The stated purpose of the initiative was to develop an improved capacity for integrated simulation and optimization of fusion systems. The proposed 5–6-year program would "build on the improved computational models of fundamental processes in plasmas that are being developed in the theory program and in the SciDAC program".

The resulting ISOFS (Integrated Simulation and Optimization of Fusion Systems) subcommittee was chaired by Dr. Jill Dahlburg, a member of FESAC as well as ASCAC. Dr. Dahlburg presented the preliminary report of the ISOFS Subcommittee at the ASCAC meeting in May 2002.

The final report of the ISOFS Subcommittee, sent to Dr. Orbach in December 2002, recommends a major initiative to create (among other things) "the algorithmic and computational infrastructure that enables the [disparate physics] models to work together". The executive summary of the ISOFS report notes that

> The characteristics of fusion plasma make the goal extremely challenging. These characteristics include the presence of multiple time scales, ranging

> over fourteen orders of magnitude, and multiple spatial scales, ranging over eight orders of magnitude... The computational domains are geometrically complex, and the solutions severely anisotropic.

The cover letter from Dr. Dahlburg to Dr. Orbach accompanying the ISOFS report notes that FESAC considers "the advanced computation frontier to be one of the most exciting in its purview"; it also highlights the belief of FESAC that this initiative would bring "huge benefits to fusion research and to the fusion energy goal". Complete details may be found at `www.isofs.info`.

4.4. Biological Research and High-End Computing

The Office of Science "Genomes to Life" program, now called "GenomicsGTL", is the result of a strategic alliance between the Office of Biological and Environmental Research and the Office of Advanced Scientific Computing. The first awards under this program, made in the summer of 2002, were preceded by a series of workshops. The report from one of these, a May 2002 workshop on computer science for Genomes to Life, concludes that "high-performance computing has fundamentally changed the way that biologists do science". A detailed discussion of GenomicsGTL, including its history, is available at `doegenomestolife.org/index.html`.

As part of the Office of Science's examination of the role of computation in biology, a letter from Dr. Decker in August 2001 asked ASCAC to provide advice on "the computational side of biotechnology". The final report of ASCAC Biotechnology Subcommittee, chaired by Dr. Juan Meza, was sent to Dr. Orbach on December 16, 2002. That report, like the reports mentioned above on nanoscience (Section 4.2) and fusion (Section 4.3), stressed the importance of mathematical modeling, computational simulation, and computational tools in the still-emerging field of computational biology.

The recommendations of the Biotechnology Subcommittee report most relevant to the present report are the following:

- ASCR should expand upon the SciDAC program to bring biologists, mathematicians, and computer scientists together as a team.
- The ASCR program should continue to invest in biophysics and biomolecular simulations, which are already having an impact in the underlying science;
- Computational biology will drive new fields of mathematics and computer science. The ASCR program should address these new areas through investments in fundamental mathematical and computer science algorithms;
- The ASCR program should develop new database and scientific data management infrastructures that can be used for computational biology.

4.5. DOE Science Computing Conference 2003

On June 19–20, 2003, an informal workshop sponsored by the Office of Advanced Scientific Computing Research and the National Nuclear Security Administration, organized by Dr. Gary Johnson of ASCR, was held of "experienced researchers, working at the forefronts of science and its enabling technologies". The focus of the conference was on the next generation of computational resources to be provided by the Office of Science, examining the needs of representative science applications as well as computer and network resources. ASCAC was not involved in organizing this workshop, and no ASCAC members were able to attend.

The intended outcomes of the conference were:

- Definition of the role and structure of the next generation high-end computing and networking facilities for the Office of Science;
- A vision of how multidisciplinary geographically-dispersed scientific research will be enabled by these facilities;
- Defining the management challenges for the Office of Science that arise from the above.

Further details are at `www.doe-sci-comp.info/index.html`; the report from this workshop is not yet available.

4.6. The Science Case for Large-Scale Simulation (SCaLeS)

On April 2, 2003, an initiative was begun by Walt Polansky of the Office of Advanced Scientific Computing Research to encourage the scientific community to "identify rich and fruitful directions for the computational sciences from the perspective of scientific and engineering applications" and to "build a strong science case for an ultrascale computing for large-scale simulation. This activity, dubbed "SCaLeS", was chaired by Dr. David Keyes; the SCaLeS workshop was held on June 24–25, 2003. Approximately 325 people from eleven areas of science and engineering, eight areas of applied mathematics, and eight areas of computer science participated in the workshop. ASCAC was not formally involved in organizing the workshop, but three ASCAC members (Roscoe Giles, Gregory McRae, and Margaret Wright) attended.

The SCaLeS report presents a very large array of scientific applications central to the Office of Science, and discusses in detail the expected advances from increases in computational and algorithmic power. The first and second recommendations of the SCaLeS report are:

- Major new investments in computational sciences are needed in all of the mission areas of the DOE Office of Science, so that the United States is the first, or among the first, to capture the opportunities presented by the continuing advances in computing power;
- Multidisciplinary teams, with carefully selected leadership, should be assembled to provide the broad range of expertise needed to address the intellectual challenges associated with translating advances in science, mathematics, and computer science into simulations that can take full advantage of advanced computers.

The first volume of the SCaLeS report is at `www.pnl.gov/scales/index.stm`, with the second volume expected soon.

4.7. The INCITE Program

On July 31, 2003, a new competitive, peer-reviewed program designed to support "innovative, large-scale computational science projects"—*Innovative and Novel Computational Impact on Theory and Experiment* (INCITE)—was announced by Secretary of Energy Spencer Abraham, who noted:

> The capabilities of terascale computing are transforming the conduct of science, bringing scientific simulation through computational modeling to parity with theory and experiment as a scientific tool. . . The INCITE initiative will make Lawrence Berkeley's NERSC facility available to all qualified researchers for grand challenge calculations—and in the process bring us closer to achieving the full potential of scientific simulation to solve outstanding scientific and industrial problems of major significance.

To achieve these goals, the Office of Science decided that 10% of the computational capabilities of the NERSC (National Energy Resource Scientific Computing) Center should be made available for grand challenge calculations, with no requirement of any connection to the Department of Energy. INCITE proposals needed to demonstrate not only major scientific impact, but also evidence that the project team could effectively use a major fraction of the NERSC Center.

Fifty-two proposals were received, more than 60% of which were from academic users, requesting a total of more than 130 million hours of processing time. On December 23, 2003, announcing INCITE awards to three projects, Secretary Abraham commented:

> The number and quality of the proposals we received show that this promise [of scientific discovery through high-end computing] is shared by our colleagues in the scientific community world-wide. We are delighted by their enthusiasm, and only wish that we had more resources to provide.

The response to the INCITE program supports Conclusion 2 (Section 1.2) and Recommendation 1 (Section 1.3) in this report.

4.8. Strategic Planning for ASCR

On July 22–23, 2003, a strategic planning workshop for the Office of Advanced Scientific Computing Research was organized by Dr. Daniel Hitchcock of ASCR. ASCAC members were explicitly invited to participate; Jill Dahlburg, Roscoe Giles, Gregory McRae, and Margaret Wright attended the workshop. (See `www.fp-mcs.anl.gov/ascr-july03spw`.)

Several subsequent discussions, coordinated by Dr. Hitchcock, were held in August and September 2003 concerning the development of ASCR's strategic plan, in consultation with members of the ASCAC High-End Computing Subcommittee. The ASCR strategic plan has not yet been publically released, but the draft last seen by ASCAC members closely reflects the views expressed in this report.

5. The High-End Computing Revitalization Task Force

In this section we mention an important interagency activity that includes, but is broader than, the Office of Science. Because this activity is related to the need for high-end computing, it is relevant to ASCAC in general and the High-End Computing Subcommittee in particular.

In March 2003, responding to a statement in the *FY2004 Budget of the United States Government*, an interagency task force—the High-End Computing Revitalization Task Force, or HECRTF—was formed. Dr. Alan Laub, then the Director of SciDAC (see Section 4.1) was appointed as one of two cochairs of HECRTF.

The duties of HECRTF include planning activities to guide future investments in high-end computing—in particular, developing an interagency research and development roadmap for high-end computing core technologies, a federal high-end computing capacity and accessibility improvement plan, and a discussion of issues relating to federal procurement of high-end computing systems.

The final report of HECRTF, originally scheduled for publication in August 2003, has not yet been released; see `www.itrd.gov/hecrtf-outreach/`.

6. The Future

For completeness, we summarize several current and future activities related to high-end computing in the Office of Science.

6.1. Leadership-Class Computing

The Office of Advanced Scientific Computing Research recently announced its interest in applications for leadership-class computing capability in support of science. The deadline for applicants is April 2, 2004, and the decision is expected on April 15, 2004. The call for applications notes:

> DOE's Office of Science, in order to accomplish its mission, is faced with the need for computing capability that far exceeds what is currently available from commercial sources.... This solicitation is part of ASCR's response to the need for leadership-class computing for capability-limited science applications.

The focus of the proposed effort is intended to be leadership-class capability computing in support of high-end science, rather than enhanced capacity for general science users. See Section 3 for the recommendation on leadership-class computing by ASCAC's Subcommittee on Big Facilities.

6.2. BESAC Subcommittee on Theory and Computation

In January 2004, a Subcommittee on Theory and Computation was formed by the Basic Energy Sciences Advisory Committee (BESAC). The purpose of this subcommittee is

> to identify current and emerging challenges and opportunities for theoretical research within the scientific mission of Basic Energy Sciences, with particular attention paid to how computing will be employed to enable that research.

This subcommittee is chaired by Kate Kirby and Bill McCurdy; Margaret Wright of ASCAC is a member.

According to Patricia Dehmer, Associate Director of the Office of Basic Energy Sciences Research, the intent of this subcommittee is to "provide a philosophy for investment,... partnering with ASCR in computational sciences". The BESAC Theory Subcommittee has arranged for testimony to be given in April 2004 by leading scientists in areas related to Basic Energy Sciences. Dr. William A. Lester, Jr., a member of ASCAC, has been invited to testify on April 17, 2004.

A draft report from this subcommittee is expected by August 2004, with the final report in December 2004 or January 2005.