

31 May 2002

Dr. Raymond L. Orbach
Director, Office of Science
Department of Energy, SC-1
1000 Independence Avenue, SW
Washington, DC 20585

Dear Dr. Orbach,

In April, 2001, James Decker, the then Acting Director, United States Department of Energy (DOE) Office of Science, presented this charge to the Advanced Scientific Computing Advisory Committee (ASCAC), advisory panel to the Office of Advanced Scientific Computing Research (ASCR):

For ASCR facilities such as NERSC, ESnet, Chiba City at ANL, and the CCS at ORNL

- (a) What is the overall quality of these facilities relative to the best-in-class in the US and internationally?
- (b) How do these facilities relate and contribute to Departmental mission needs?
- (c) How might the roles of these facilities evolve to serve the missions of the Office of Science over the next three to five years?

To address this charge, the ASCAC formed a Subcommittee on Facilities whose members are: John Connolly (University of Kentucky); James Coronas (Krell Institute); Jill Dahlburg (General Atomics), subcommittee chair; Helene Kulsrud (IDA); Gregory McRae (MIT), Paul Messina (Caltech); Warren Washington (NCAR); and Stephen Wolff (Cisco).

A preliminary report of the subcommittee was presented by Jill Dahlburg at the most recent ASCAC meeting, held on May 2-3, 2002, in Washington DC. In light of the announcement in April 2002 of the Japanese Earth Simulator, the ASCAC requested of the subcommittee that the final report should include a suggested response to the implications of the Earth Simulator for American computational science leadership. Such a response was formulated following an "Earth Simulator Rapid Response Meeting", held at your request on May 15-16, 2002.

As chair of ASCAC, I am formally submitting the final report to you on behalf of the Subcommittee on Facilities.

The essential finding of the Subcommittee is that each of the four diverse and complementary ASCR facilities is among the best worldwide in its respective category. It is the opinion of the Subcommittee that these ASCR facilities and the related spin-off research efforts contribute outstandingly to the mission needs of the DOE, and profoundly and positively impact high performance computing activities worldwide. Further, the Subcommittee believes that the ASCR Office is uniquely positioned within the U.S. to spearhead a response to the Earth Simulator that well-addresses the computational science challenge posed by that facility.

Looking ahead, the Subcommittee offers five recommendations for the ASCR future:

1. ASCR should retain focused commitment to high end computing in the service of DOE Office of Science missions.
2. ASCR should build on its present plans, to formulate the response to the Earth Simulator:
 - The time is right for a major new initiative whose goal is to regain, and in some areas to retain, world leadership in scientific computing to advance the U.S. mission-driven research of the DOE Office of Science.
 - It must be noted that previous ASCR planning was carried out in a very constrained budget environment. For this reason, it would be detrimental to redirect existing, scarce program funds in order to meet this challenge. Planning should thus assume a funding increment on the order of \$150 M/year for the near term.
 - The planning program elements and activities ought to include: advanced architecture development; computational science and enabling technology research; and, focused technology deployment in support of the DOE mission applications.
3. ASCR should develop an integrated allocation strategy for its computational resources. This allocations process should seek to ensure that each machine is filled to the greatest extent practicable with high priority DOE Office of Science jobs that are not feasibly run on smaller machines.
4. DOE Office of Science researchers should be encouraged by ASCR to procure both mid-range and lower end machines/clusters with individual program funds.
5. ASCR should embrace a cohesive networking and resource allocation approach for computing infrastructure as a way to provide the most uniform interface to all the types of computing facilities that are encompassed by ASCR. To this end, ASCR should continue to incorporate advances in networking and resources integration as they develop, and should, further, encourage enhanced research efforts in new architectures and networking capabilities.

I. Background

During the period April through December 2001 the Subcommittee gathered information for the above-noted findings and recommendations, making primary use of two sources: presentations to the ASCAC Subcommittee on Facilities and associated institutional information, and query response. Facilities presentations to the ASCAC and/ or to the Subcommittee, included briefings: (a) on May 2-3, 2001, by: Walter Polansky, *DOE Office of Science* [overview of the MICS Facilities]; Horst Simon, *National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL)* [overview of

NERSC]; Richard Stevens, *Argonne National Laboratory (ANL)* [overview of Chiba City]; Thomas Zacharia, *Oak Ridge National Laboratory (ORNL)* [overview of ORNL CCS]; and, James Leighton, LBNL [overview of the ESnet]; (b) on August 16, 2001, by: David Schissel, *General Atomics (GA)* [overview of the National Fusion Collaboratory]; and, Paul Messina, *Caltech*, [overview of the Grid, internationally]; and, (c) on October 25-26, 2001, by: Richard Stevens, *ANL* [on high-performance computing (HPC) facilities: Grids, petaflops, and usage]; Robert Ryne, *NERSC* [on requirements of a HPC system user]; Dalton Schnack (talk given by Jill Dahlburg, *GA*), *Science Applications International Corporation (SAIC)*, [on a mission-driven perspective on the status and needs for DOE computing]; and, Stephen Wolff, *Cisco* [on networks, with particular emphasis on ESnet].

The Subcommittee, further, held an Earth Simulator Rapid Response Meeting on May 15-16, 2002. A summary of this meeting can be found in the Appendix to this report. The major conclusions from the meeting are:

- The Earth Simulator simulation capability is real; i.e., it provides significantly enhanced performance on real science applications.
- The Japanese machine is not a surprise, but is a result of Japanese decisiveness, commitment, and accountability. However, all, including the initial Japanese users, have evinced surprise at the sustained efficacy of the platform;
- DOE's focus on high-end science applications complements activities in other U.S. governmental agencies. The DOE Office of Science is well positioned to lead the Nation back to the front rank of computational science.
- The DOE Office of Science research community, in partnership with the domestic vendor community, is ready to respond to the Earth Simulator challenge.
- The DOE ASCR Office is significantly under-funded to adequately address this mandate.

With regards the first aspect of the Decker charge, facilities evaluation, the Subcommittee considered the facilities in isolation and comparatively. Findings from this study are summarized in Sec. II.

The second aspect of the charge, relation to and contribution of the facilities to DOE mission needs, is addressed in Sec. III. Several of the speakers to the Subcommittee described the fundamental role of advanced computing in solving DOE mission-relevant problems (*e.g.*, Schnack), the importance to their research of state-of-the-art ASCR HPC resources, and the excellent responsiveness of the ASCR facilities. The more subtle subtext of mission driven research as a 'requirements pull' for enabling fundamental advances was also indicated directly or indirectly by several of the speakers (*e.g.*, Stevens noting that much grid software has been developed for the solution of specific DOE problems). The essential observation is that mission-driven research, which inspires some of the best basic research extant, is exemplified superbly by the DOE ASCR.

Twenty years ago, high end scientific computing was performed on manufactured-on-request vector mainframes for which 30 Mflops sustained was considered to be good performance. Vector FORTRAN was the standard programming language for which few debugging and optimizing tools were available, and most jobs were submitted by remote batch processing using dumb terminals and 9600 baud telephone connections. In contrast, high performance computing of 2001 was typified by commodity massively parallel platforms on which 30 Gflops of sustained performance was easily possible using group-developed object orientated FORTRAN/C software that was coded with the assistance of automated parallel debugger and development tools on versatile desktop workstations. The higher speed connectivities such as ATM OC-12 enabled mainframe-driven visualization systems for tasks ranging from debugging to large database results processing. With this advent of the ability to routinely perform highly resolved, multi-dimensional, engineering-class computations, advanced scientific computing has become an enabling tool for first principles exploration. Computational science now is considered by most researchers within the DOE Office of Science, and in the physics and engineering communities at large, to be a third arm of research, ranking as equal with theory and experiment as a tool for discovery. Continuity of these advances will require a careful planning of resources: of facilities capabilities, of allocations, and of connectivity. Strategic directions for the next three to five years, part (c) of the charge, are topics of Sec. IV.

II. Facilities: *What is the overall quality of these facilities relative to the best-in-class in the US and internationally?*

The charge to this Subcommittee includes as first element the evaluation of four facilities funded by the OASCR. These are NERSC, ESnet and the Chiba City and CCS facilities at ANL and ORNL respectively. Each of these four facilities has a very different vision, and all serve complementary purposes within the DOE Office of Science.

NERSC is a computational production center providing state of the art access to computational resources to applications researchers. These researchers are interested in a stable productive environment where software can be optimized without the added complication of a constantly changing environment. The role of NERSC is to concentrate on the current generation of supercomputer, and provide service to a wide variety of computational scientists. Whether or not the facility should venture beyond the current generation, and overlap its role with the other facilities is a matter of ASCR policy. This Center, situated at LBNL, is an excellent, classic production center, with a long history of providing first class service to large scale applications users by means of state of the art commercial machines. With the shift to massively parallel platforms (MPP), NERSC consulting requirements became more involved and demanding. NERSC responded successfully, by using the personnel resources to provide in-depth collaboration, algorithm and visualization support. The current, well-used, high end machine at NERSC is ranked Number 3 in the world, in the November 2001 version of the Top 500 list, with the DOE classified machine ASCI White ranked top and the Pittsburgh Supercomputer Center Compaq machine ranked second. NERSC has made a proposal to

continue to move up the terascale ladder, which has drawn good reviews from the user base. It proposes to increase the facility's capacity by an order of magnitude over the next seven years. The NERSC allocation process gives priority to Office of Science "grand challenge" problems, but is also available to a large number of DOE grantees from all branches of the Office of Science. In FY2000 NERSC delivered 7,846,244 machine hours to production projects.

ESnet is a facility that provides networking and no cycles. It serves a similar purpose within the Office of Science as does NERSC. At first blush, it is a small, reliable ISP. It is governed by and is extraordinarily responsive to its users. It has relieved the user sites of the need to provide networking expertise and user services. In terms of performance, ESnet's connectivity is at present adequate for most current users. It is a lean, cost-effective operation, with good management tools and user services, but no central capability for networking research. In February 2001, ESnet carried a tremendous 45 terabytes of data, and continues to experience a 100% growth each year. The danger is that the net is required to meet the increasing demand even as the governance structure, which primarily supports current connectivity stability, is not well constituted to deal with the explosive growth. In order to meet the growing demand, ESnet needs to take on research tasks in strategic directions. New applications could add 1 Gbytes/sec or 300 Tbytes/month.

The ANL Chiba City Project, is a testbed for generating computer science tools. The goal of the project is to provide a series of parallel hardware and software testbeds for the computer science and applications community aimed at supporting research in software scalability. The work supported by the initial testbed has had significant effect on the high-performance computing community, and as the project evolves and expands, researchers at ANL believe that it will have increasingly broad impact by enabling more rapid progress in realizing the dream of scaleable systems. This impact will only be achieved if the community can exploit the capabilities of the testbed on a routine basis, an objective which is embraced by the project personnel.

The Oak Ridge Center for Computational Sciences (CCS) is a high-end capability computing center focusing on a few key timely topics in support of the Department's science mission. In support of this mission, the CCS is currently focused on providing specialized services to the biology, climate and materials sciences communities. In addition, the CCS is the principal resource for SciDAC projects. Another important role is to evaluate new architectures through specific applications benchmarks. This Oak Ridge system was set up on the model of a few groups of highly sophisticated users running very large and/or long-running jobs on large computer systems to "push the envelope" in computational science by performing tuned calculations that could not otherwise be easily carried out. This model has allowed the CCS to often take delivery of emerging, and unproven architectures, such as the Intel Paragon (or more recently the IBM Power4) to drive computational sciences at the leading edge. Today, CCS continues to cater primarily to a small number of applications groups that use the systems. These groups have both very sophisticated users and technician level users. The sophisticated users develop the codes to a state where they can be turned over to the technicians to

make many runs of the code to study a parameter space. The computers and the problems are so complex that it is nearly impossible for one person to understand both the science and the computers at a “world class” level; consequently, teams of users are required to effectively utilize the systems. It is anticipated that this trend will increase in the next five years. With movement to clusters of larger and larger SMP nodes, and, in particular to next generation "cellular" petascale machines, the applications will be forced to adopt more and more levels of parallelism to take advantage of the resources. The CCS is encouraging teams that, as they grow, include more specialists.

One way to characterize (and perhaps oversimplify) the three computational facilities is that NERSC is a computational production center providing access to multi-teraflop state-of-the-art computational resources to a broad set of applications researchers; the Center for Computational Sciences at ORNL is a special-purpose, high-end computational sciences facility providing multi-teraflop, focused resources on advanced architectures for a few key science topics; and Argonne is a teraflop-range computer sciences facility that is focused on enabling rapid progress in realizing the dream of scaleable systems. Each is a first class facility of its type and purpose when considering both US facilities such as those within the National Partnership for Advanced Computational Infrastructure and systems in the DOE Advanced Scientific Computing Initiative complex, and also international capabilities.

III. Mission Relevance: *How do these facilities relate and contribute to Departmental mission needs?*

In this section is examined the relationship between mission needs and facilities of the Department of Energy’s Office of Science. Large-scale computational modeling and simulation have become central to most scientific and engineering research and the missions of the DOE Office of Science are no exception. A number of mission statements to this effect have been issued by the Secretary, by the Office of Science, and Advanced Scientific Computing Research Program. A brief review of some of the most relevant mission statements is given below:

1) In the FY02 request to Congress the ASCR program articulated the following mission statement: “The research and facilities supported by ASCR are critical to the success of all the missions of the Office of Science because computational modeling and simulation have become an important contributor to progress in all SC scientific research programs. Modeling and simulation is particularly important for the solution of research problems that are insoluble by traditional theoretical and experimental approaches, hazardous to study in the laboratory, or time-consuming or expensive to solve by traditional means. All of the research programs in the U. S. Department of Energy’s Office of Science—in Basic Energy Sciences, Biological and Environmental Sciences, High Energy and Nuclear Physics—have identified major scientific challenges that can best be addressed through advances in scientific computing.”

2) At the joint DOE and National Science Foundation (NSF) National Workshop on Advanced Scientific Computation of July 30-31, 1998, there were a set of recommendations developed a long-term strategy for the two agencies. The report that contains the recommendations is often referred to as the Langer report, which is named after the chair, James S. Langer. “ The impact of Advanced Scientific Computing on industry, government, and national labs has been growing for several decades. However, in the future, we believe that will be a very rapid expansion of such techniques across a far broader segments as Advanced Scientific Computing will become an indispensable tool in understanding and managing our ever more complex and inter-related world. In industry, it will move beyond crash simulations, airplane design, and drug design to a whole new world of data intensive computing such as financial risk management, fraud detection, and supply chain optimization. In government, computing in the service of national defense will be extended to decision support for such societal issues as disaster planning and management, infrastructure investments in protection, and environmental and energy security. As one of the largest producers of data and reports, the digital age will employ scalable computers to help organize and deliver more cogent information products to our citizens. National labs will extend their missions from use of high-end computing in the service of national defense to national decision support for policy issues involving the environment on the energy economy.”

The above broad statements allude to the future for scientific computing with the Office of Science. Through the materials that have been presented to this Subcommittee and the testimony of many experts, we see a future of increasing reliance of scientific computation to accomplish the missions. The Associate Directors of each of the major components of the Office of Science strongly concurred with the need for a range of computing resources for their respective science missions.

As important as providing access to large-scale computing environments (powerful computers, large data archives, advanced visualization technologies) are efforts aimed at advancing the state of the art of computing technologies and making it easier to use them effectively and efficiently. The ASCR program and its predecessors in the Office of Science have a long and distinguished history of developing mathematical models, algorithms, software libraries, and software tools for high-end computing on advanced architectures.

The current manifestation of such efforts is found in the recently developed Scientific Discovery Through Advanced Computing (SciDAC) program, which has as its goal to produce the scientific computing, networking, and collaboration tools for DOE science. The program goals of SciDAC are aimed at addressing the computation needs more effectively across the major Office of Science programs and in collaboration with other government programs. DOE Office of Science has also contributed to the creation of new computer architectures that have led to major advances. These activities are also highly relevant to the Office of Science mission, because without them the available computer systems would be less capable and more difficult to use, and the Subcommittee urges that they be continued.

More recently, Secretary Abraham said in a speech on homeland defense, "Our world class scientific and engineering facilities and creative researchers helped make our nation more secure for over 50 years. These same resources have been trained on the threats posed by terrorism for some time and because of this foresight, technologies such as those are in deployment today." Also, the Secretary has stated that "program like the Human Genome Project, or the President's National Climate Change Technology Initiative support our mission." DOE senior management and OMB have made performance, planning, and accountability high priority. In addition, the House Committee on Science has held hearings December 5th on how the nation's research federal establishment can contribute to the war on terrorism. The Office of Science and the Office of Budget and Management is presently conducting an inventory assessment of federal research related to terrorism. The urgent current activities are likely to lead to more mission responsibilities for the Office of Science, some of which will involve the need for advanced computing capability. As examples of how advanced computing can help, there is a new need for more rapid DNA sequencing of microbial pathogens used in biothreat agents and for faster and more detailed holographic imaging devices. Over the next few months it is expected that this new mission responsibility will be better articulated.

Many of these mission statements and requirements will involve increased computing capability. The Subcommittee sees the role of advanced computing in the Office of Science's research program growing significantly in the future and becoming more integrated into the Office's program and projects. This position was reinforced by the recent deployment of the Japanese Earth Simulator.

IV. Strategic Directions: *How might the roles of these facilities evolve to serve the missions of the Office of Science over the next three to five years?*

In addressing possible future roles of ASCR facilities, there are two strategic points to consider. First, mission directed research, from basic to highly applied, is the orientation of ASCR. Second, high end computing is the unique charge of ASCR within the DOE Office of Science.

Testimony from researchers within the Office of Science computing constituency has established that DOE Office of Science missions require computing resources that range from a small number of processors (local clusters) to the highest end (petaflops and beyond). These resources are to a large part already available within the Office, and are successful at satisfying mission needs.

However, in the absence of appropriately evolving usage and procurement strategies for existing resources, HPC hardware and funding can easily be misdirected for computing which either neglects the ASCR mission of high end computing in the exclusive service of mission needs, or neglects the missions in the interests of a few very large computational science projects. The balance of using high end resources effectively both from the perspective of flop rate and also from the perspective of Office of Science

missions requires significant attention to allocation procedures and strategies. This balance, filling each ASCR machine to the greatest extent practicable with mission-relevant jobs that are not feasibly run on smaller machines, implies that (1) a range of machines are required within the ASCR portfolio, and (2) allocations will need to be considered globally across that portfolio. The evolution of the internet with grid technologies will further enable such a strategy, allowing ready integration of technologies from the highest end HPC resources to few-processor clusters across geographically diverse communities of interest.

In summary, the Subcommittee applauds ASCR for its current success, and urges the Office to mandate the future. With mission-directed research from the basic to the highly applied as the orientation and driving excitement of the research fostered by ASCR, and high end computing the unique charge of that Office within the DOE Office of Science, the Subcommittee believes that ASCR provides an ideal environment in which to harvest a future that spans the spectrum from research and development of new architectures to the deployment of the multi-user computing facilities of the 21st Century. The Subcommittee strongly encourages the Office to build on its existing plans, to develop a blueprint that will provide the framework for the synthesis of research from applications people, computer scientists, and computational scientists in a way which, when achieved, will produce the paradigm shift to the next generation of High Performance Computing. In light the Earth Simulator Challenge, such planning should be developed assuming significantly enhanced program resources.

Please do not hesitate to contact me for further information regarding this report.

Yours truly,

Margaret H. Wright
Computer Science Department
New York University
251 Mercer Street
New York, New York 10012
212-998-3056
mhw@cs.nyu.edu

Appendix: Report on the Earth Simulator Rapid Response Meeting, May 15-16, 2002

A.1 Preface

On May 2, 2002 the Facilities Subcommittee of the Advanced Scientific Computing Advisory Committee [ASCAC] reported to the full ASCAC regarding the draft subcommittee report. Due to the recent announcement of the successes of the Japanese Earth Simulator project, it was proposed that the presented subcommittee report be viewed as an interim report and that the subcommittee rapidly readdress one of its charges: “ How might the roles of ASCR facilities evolve to serve the missions of the Office of Science over the next 3-5 years?” in light of recent developments. Since at the same meeting Dr. Orbach asked for a quick response to the Earth Simulator issue, it was agreed that the Facilities Subcommittee would address Dr. Orbach’s request. To this end the ASCAC Facilities Subcommittee convened the Earth Simulator Rapid Response Meeting (ESRRM) on 15-16 May 2002 in Crystal City, VA. Attendees included representatives from the inter-agency high performance computing (HPC) community, SC program offices, DOE laboratories, and experts on the Earth Simulator and climate simulation applications. What follows is a discursive summary on the ESRRM which, it should be noted, included a great deal of discussion. This summary is framed in light of Dr. Orbach’s request. The subcommittee report will be available shortly and will be based on the ESRRM as well as previous work of the Subcommittee.

A.2 The Challenge

Dr. Orbach’s statements at the May 2, 2002 ASCAC meeting that “... there is a centrality of computation in everything we do,” and that “...large scale computation is the future of every program in the Office of Science,” perfectly frame the discussion that follows. The past decades have seen the rise of computational science as the third paradigm for scientific exploration, complementing the traditional approaches of theory and experiment. Surveys of computational challenges in many fields ranging from high energy physics, to fusion energy science to computational biology frequently show a range of spatial and temporal interest, each of which covers a dozen or more orders of magnitude. To meaningfully explore the astonishing breadth of these problems requires major commitments of focused scientific resources: interdisciplinary teams of scientists, computer scientists, and applied mathematicians enabled by state of the art computational resources. The breadth and complexity of these opportunities provides funding agencies with new challenges as well. The complexity and expense of forging these efforts requires a long-term strategic commitment to excellence in one or more of these areas. The Japanese response to these opportunities in one class of applications is the Earth Simulator initiative.

In summary, the Earth Simulator is a high-end general-purpose computer focused on a class of problems. It represents maximum capability computing applied to a targeted attempt at a scientific breakthrough. It is the result of a focused, long-term, top-down

Japanese design effort. The Earth Simulator (GS40) appears to be the first of a class of machines that will be manufactured in some quantity and deployed as a computing engine for other scientific applications shortly. On a series of real climate applications the Earth Simulator provides a 50-fold computational performance advantage. While results are still somewhat sketchy, similar performance is expected on other science applications. It should be noted that the plans of the ASCI program include hardware that can deliver *peak* performance equivalent to the Earth Simulator in 2004; however, projected *sustained* performance for actual simulations applications is not yet clear. U.S. access to the GS40 is only possible on site in Japan and through research collaborations with Japanese scientists. The Earth Simulator allows science that is currently beyond the reach of American scientists.

A strength of the Earth Simulator project is its focus on a particular important class of problems. The challenge that it poses goes well beyond climate and related science: the challenge is to American leadership in computational science. At this writing, climate research, fusion energy research, high energy physics and materials science are well positioned from purely scientific perspectives to take immediate advantage of a 50-fold increase in computing capability. Other science applications, such as nano-science and biology, are in the early stages of developing high performance computational tools. However their computational needs will be substantial and their longer-term requirements must be considered. In order to meet program intermediate-term objectives in biology, a computational capability an order of magnitude greater than the GS40 will likely be required. In addition, in the SciDAC initiative an additional number of applications are now ready for GS40-class computational science. In short, without a robust response to the Earth Simulator challenge, the United States is open to losing its leadership in any of these areas. Perhaps more fundamentally, we become open to losing our leadership in defining and advancing the frontiers of computational science as a new approach to science and so lose leadership in an arena increasing critical to both our national security and economic vitality. This loss has the potential for significant long-term detriment to American interests.

Of particular concern and comment at the meeting were workforce issues. Shorter term “brain drain” issues (including our inability to attract superior people to our shores) due to a lack of first class competitive facilities is a danger. However, more fundamentally American leadership in science is fueled over the long term by the bright young minds attracted to the enterprise. The pull of Wall Street in the 80’s and the “.coms” in the 90’s has drawn away many that in other circumstances would be at the early or mid stages of vigorous careers in science, particularly science strongly connected to computation. We are a competitive nation; our best and brightest will only be drawn to areas in which we excel, to areas in which we lead. They will not be drawn to a second class science enterprise.

A.3 DOE Plans and the Emphasis of Other Agencies

The DOE has a long-standing role in high performance scientific computing. This focus provides the Department with a unique blend of technical expertise and management perspective. In the past several years these assets were used by the ASCR Office as it undertook an extensive planning process which has led to the SciDAC initiative that was launched last year. From a planning perspective, the Earth Simulator challenge finds the ASCR Office with a well-developed and widely vetted set of planning guidelines for the enhancement of computation resources. To be clear, the SciDAC planning processes define both the scientific rationale and implementation guidelines for a robust response to the Earth Simulator challenge. However, this planning was carried out in a very constrained budget environment for the ASCR Office. The ASCR Office receives about 5% of the Office of Science budget. By comparison the CISE budget of NSF is about 10% of the total NSF budget. Like wise the ASCI budget is about 10% of the overall NNSA budget. The 1990s were a time of significant growth for computing at NSF and with the Defense Programs of DOE. Both of these entities utilized advances in computational science pioneered by DOE and the ASCR Office. During this time the ASCR Office budget has lagged behind in both absolute and relative terms. The Office has thus been put in the position of attempting to implement its plans to serve a need central to all aspects of the Office of Science missions with severely inadequate funds. With sufficient resources, the ASCR Office can provide a national computational science infrastructure – comprised of both enhanced *capacity computing* and a substantial component of maximum *capability computing*. This role fits well with the emphasis of other federal agencies.

Other agencies with an interest in high performance computing have of course also moved forward with planning. Congress has required the DoD to conduct a study of the computational needs and possible future computer architectures needed to support essential DoD applications. This National Security Agency-led activity is potentially an important element of a National strategy in this area and ASCR is participating in this activity. However, this activity is not focused primarily on the types of scientific questions that SC computing must address. Enhanced HPC productivity through novel hardware and software development is a recently initiated DoD activity, principally through DARPA. This activity has a goal of making systems available in 2008. NSF is in the final phase of an extensive planning process focused on cyber infrastructure, with emphasis on enhanced networking/ grid capabilities and data-intensive applications. It is quite clear that DOE's above-outlined focus on high-end science applications complements activities in other governmental departments. DOE is well positioned to lead the Nation back to the front rank of computational science.

A.4 Technology Readiness

While it is important not to view the Earth Simulator challenge as simply a matter of building a better and faster compute engine, the need to provide hardware that is GS40 capable and beyond is a critical issue. A crucial question then is the readiness of the domestic vendor community to meet this challenge. While this question requires

extremely careful examination before public money is spent with a particular vendor, it appears that there are no insuperable barriers to the domestic acquisition of appropriate hardware. It further appears that domestic vendors could provide systems that match or exceed the GS40. However, to do this, a strong partnership between the Government and the vendors will be needed. Lack of technology is not a barrier to success. In the short term, the Office of Science can reduce the science-gap created by the Earth Simulator by significantly enhancing its high end computing capability, incorporating the best available U.S. technology. In the mid- and longer-term, the Office of Science should embark on an aggressive, integrated program to regain and to sustain leadership in the areas of computational science important to the DOE mission. Program elements/ activities may include: advanced architecture development; computational science and enabling technology research; and, focused technology deployment in support of the DOE mission applications.

While no definitive discussion of the differences in the Japanese versus American approaches to developing high performance computers in the 90s emerged, the line between these approaches was sharply drawn. The Japanese approach has been to design the top end and have the advanced technology “trickle” down. The American approach has been to build up to the high end. In light of the need for a government/ vendor partnership the question of the technology approach for the next phase requires careful consideration before investments are made.

A.5 Conclusions

There is no doubt that DOE has the experience, the mission need and scientific expertise to lead a national effort in science-driven high-performance computing to regain national leadership in computational science. DOE Laboratories and domestic computer vendors have a long history of successful collaborations in high end computing. Both parties are ready and willing to respond to the Earth Simulator challenge. The mandate of ASCR is computational science in support of DOE missions. With adequate resources, DOE can compete in, win, and dominate this space.