FINAL DRAFT: 7 FEB 2008

Report of the 2008 ASCAC Balance Panel

International Competitiveness: Facilities for Participation Research for Leadership

Submitted to the ASCAC on 7 February 2008 for consideration as a:

Report of the Advanced Scientific Computing Advisory Committee

March 2008

ASCAC Balance Panel Members

F. RONALD BAILEY, Computer Sciences Corporation
VINCENT CHAN, General Atomics
JILL DAHLBURG, Naval Research Laboratory (Panel Chair)
MICHAEL HEATH, University of Illinois/ Urbana
CHARLES McMILLAN, Los Alamos National Laboratory
ROBERT VOIGT, Science Applications International Corp. (Panel Co-Chair)

Executive Summary

The Department of Energy [DOE] Advanced Scientific Computing Advisory Committee [ASCAC] was charged with assessing the balance within the Office of Advanced Scientific Computing Research [ASCR] between investments in facilities (particularly Leadership Class Facilities [LCF]) and investments in research, and within the latter, the balance between research of near-term benefit to current scientific applications and research of a more fundamental nature with higher risk and longer-term payoffs. The overarching goal presumed as underpinning the Charge is the acceleration of scientific discovery on the most important scientific questions of our time, and particularly on maintaining and enhancing US leadership in scientific discovery.

To address the Balance Charge, the ASCAC convened a Balance Panel, which considered a substantial number of relevant reports and other documents in preparing this Panel Report, and also received input from a wide spectrum of interested parties including numerous researchers currently engaged in research sponsored by ASCR.

Perhaps the most fundamental finding of this ASCAC Balance Panel Report is that very high-end computing hardware is becoming increasingly common worldwide, and hence high-end facilities alone no longer ensure competitiveness in terms of scientific results. Moreover, as high-end architectures become increasingly complex, their effective utilization requires ever more sophisticated algorithms, tools, and software, which in turn depend on new research and new paradigms in applied mathematics and computer science. The ability to use LCF most effectively is what makes the crucial difference for continued leadership in scientific discovery. Thus, the most fundamental recommendation of the Panel is that concomitant investments in research must be made to ensure that investments in LCF will realize their intended goals. From the perspective of international competitiveness, **we must invest in facilities to stay in the game, but we must invest in research to win**.

More specifically, the Panel found that current DOE Office of Science [SC] investments in LCF have shifted the balance away from research, and the Panel recommends that the balance be restored to the approximately 50:50 relationship that had previously prevailed in recent years. The Panel also recommends that a similarly even balance be maintained between near-term research typified by the DOE SciDAC and longer-term, more fundamental research in applied mathematics, computer science, and networking, as both are necessary to meet the needs of current scientific applications while still enabling new paradigms for future scientific discovery. This report also contains a number of additional findings and recommendations concerning the most effective strategies for managing a balanced research portfolio that leverages the successful SC SciDAC model while still supporting more open-ended research with high potential payoff. Finally, the Panel also addressed the balance between LCF and more broadly applicable and widely available capacity computing facilities, and recommends here too a need for a closer balance and for careful planning of the transition of LCF into more general scientific use as they inevitably lose their leadership status over time but remain valuable assets.

Table of Contents

1. Introduction and Approach

1.1 Interpretation of the Charge

The CY2007 Charge letter from Dr. Raymond Orbach asked the Department of Energy [DOE] Advanced Scientific Computing Advisory Committee [ASCAC] to assess the strategic priorities of the DOE Office of Science [SC] Office of Advanced Scientific Computing Research [ASCR] program, with the focus on (1) the balance between high performance computing facilities and "core" research and (2) within research, the balance between the immediate needs of current scientific applications and longer-term investments. To address this Charge (reproduced in the Appendix), the ASCAC formed a Balance Panel with members: F. Ronald Bailey; Vincent Chan; Jill Dahlburg, Panel Chair; Michael Heath; Charles McMillan; and, Robert Voigt, Panel Co-Chair.

The Balance Panel views the first Charge focus as represented by the balance between DOE SC High Performance Computing [HPC] facilities investments, in Leadership Class Facilities [LCF] at Oak Ridge National Laboratory [ORNL] and Argonne National Laboratory [ANL] and the Lawrence Berkeley National Laboratory [LBNL] National Energy Research Scientific Computing Center [NERSC], and SC investments for research in applied mathematics, computer science, networking, and the DOE SC SciDAC program. The Panel views the second focus as concerning the balance between SciDAC (representing investment in immediate research needs of current scientific applications and program partners) and investments in longer-term research in applied mathematics, computer science, and networking. In addressing the charge, the Panel did not consider either Testbeds (ESNet) or Research and Evaluation Prototypes and Network Facilities, as these represent partnership activities with other agencies. Note: an ASCAC Networking Report, which addresses the Charge to "make suggestions and recommendations on the appropriateness and comprehensiveness of the networking research programs within ASCR with a view towards meeting long-term networking needs of SC" is expected from ASCAC in February 2008.

The Panel appreciates that achieving advances in scientific discovery requires a continuous spectrum of high performance computing resources, and wishes to commend ASCR for outstanding leadership in the LCF and also the SC INCITE allocation program. Through INCITE, the LCF currently provides extreme computing to a small number of projects selected from the general science community that have a reasonable probability of resulting in high-impact scientific discoveries. In addition, the NERSC ASCR facility provides a lower tier system serving a much larger user community. NERSC also contributes to high-impact scientific discovery and additionally provides for the more complete exploitation of previous scientific accomplishments. These two types of facilities are viewed as providing necessary and complementary resources. Therefore, the Panel decided to address an additional balance, namely, that between "Tier-1" or *capability* (currently centrally provided by LCF) and "Tier-2" or *capacity* computing (currently provided by NERSC). A pictorial representation of our assessment of the current balance among all elements of the ASCR program is shown in Fig. 1.1.

Figure 1.1. Current balance of programmatic elements in DOE SC ASCR.

1.2 Approach to address the charge

The ASCAC Balance Panel convened at the Naval Research Laboratory in Washington DC during the week 14-18 January 2008 to evaluate the more than 100 documents of input offered for Panel consideration by members of the US high performance computing community. The Panel also heard testimony from ASCR Associate Director Dr. Michael Strayer, and others from ASCR, on the ASCR budgets, plans, and activities, and from Prof. Jack Dongarra (University of Tennessee) on the planned developments of high-end computing systems around the world.

1.3 Shared Viewpoints of the Panel

As a backdrop to our findings and recommendations, the Panel arrived at some shared viewpoints regarding the advancement of high performance computing for scientific discovery and ASCR's role in accelerating that advancement. They have been shaped in part by the many helpful community inputs (*cf*. Section 1.2), which we here acknowledge and for which we offer sincerest thanks.

Leadership in advanced computing for open science is a competitive advantage the US currently enjoys, and it is one that must be sustained. Leadership is our only option, and aggressively advancing extreme computing is the correct strategy. The Panel takes the fact that LCF sustains a top tier position in the Top 500 as an important indication of ASCR's commitment to this strategy. Furthermore, since the basic hardware for high performance computers is readily available globally, it is our advanced application software for extreme computing that is by far the largest contributor to US leadership.

The Panel feels strongly that scientific importance is paramount. Thus, perceived system performance and achieved system efficiency are not sufficient measures of readiness to advance science. Furthermore, the choice of applications must be heavily weighted toward scientific importance. High-end computer hardware must be used effectively for high impact scientific applications, but the value of the scientific outcome ultimately must trump considerations of system efficiency.

Advances in software across a broad front are every bit as necessary for extreme computing as are advances in hardware. Computer science research is necessary to provide environments and tools that allow effective use of the hardware and underlying operating system. Applied mathematics research is necessary to provide problem solving schemes and algorithms to implement multi-scale, multi-discipline models on complex computer architectures.

In ASCR's approach to advancing extreme computing, LCF takes on two distinct pathfinder roles. The first is as a pathfinder in advancing modeling and simulation and thereby opening new scientific regimes and even new science for exploration. The second is as a pathfinder in the application of computing hardware and software at the extreme end of the performance range, thereby expanding both the number and size of scientific domains that can be addressed by modeling and simulation.

While the advancement of extreme computing and its application to scientific discovery is an essential goal, basic research in computer science and applied mathematics is still a fundamentally important mission. Such research provides the "seed corn" from which derives the tools and techniques to address future challenges as we enter into new regimes of scientific research using advanced computing technology made possible by our earlier accomplishments.

The Panel notes that two emerging modes of conducting science are accelerating discovery. The first is that DOE science is becoming increasingly collaborative with the formation of virtual science communities, often on a global scale, utilizing distributed cyber-environments. There is an increasing dependence on networks and related technologies to unite these communities, allow efficient exchange of data, and enable new modes of scientific discovery. The second is the emergence of data analytics as a mode of scientific inquiry. We are entering an age of massive, petascale data sets in which scientific discovery depends on new tools for data transformation, visualization, and mining. Thus, advanced networking, collaborative science, and data analytics are candidates for new or reinvigorated research programs.

Finally, the Panel feels that a leading and productive research organization must employ a strategic approach to developing a strong research portfolio that is both highly productive and responsive to its mission goals. Such an approach provides a basis from which to plan, evaluate, and adjust program balance. Suggestions on the formulation of such an approach are further discussed in the summary of this report.

2. Balance Between Facilities & Research

In the recent past the balance between facilities and research, as measured by resources expended, was approximately 50:50, as suggested by Fig. 2.1.

Figure 2.1. Recent past balance between DOE SC ASCR research and facilities.

For purposes of this discussion, and also as noted in Sec. 1, *facilities* includes the LCF and its associated support, the production facility, and the network; *research* includes the applied mathematics, computer science, the research component of networking, the research prototype systems, and ASCR's portion of SciDAC.

Current investments in LCF have shifted this balance away from research and towards facilities, as suggested by Fig. 2.2.

Figure 2.2. The current DOE SC ASCR balance between research and facilities.

Recommendation: ASCR should return to a balance between facilities and research of approximately 50:50.

This recommendation is motivated by two emerging phenomena: increasing international competitiveness, as measured by the raw computing power available, and increasingly radical architectural features of new high-end systems that require new efforts in software development, from operating systems through tools to algorithms, to insure effective utilization on challenging problems facing the science and engineering community.

Returning to a 50:50 balance between facilities and research may require some difficult decisions. In a growing budget, we suggest using that growth to achieve the desired balance; if the growth is slow or non-existent, however, then the challenge becomes much more difficult. In such an environment, we urge ASCR to achieve the desired balance by delaying hardware upgrades or the acquisition of new systems, if necessary. It does the US HPC community no good to have the fastest machines if they cannot be effectively utilized, and we believe that for systems on the horizon advancement of software will be equally as important as advancement of hardware. Delaying the availability of a useless system in order ultimately to provide useable cycles is preferable to the alternative.

Two years ago, all five of the top five systems as measured by the Top 500 ranking were in the US, as were eight of the top ten. In the latest release, only two of the top five systems are in the US, and seven of the top ten. The point is that the US no longer has a monopoly on high performance systems: they are becoming an international commodity, and competitiveness can no longer be characterized solely in terms of hardware.

The competitive advantage that we currently enjoy stems from our ability to utilize these systems to solve the most challenging problems of science and engineering effectively. We must maintain this edge as high performance hardware becomes even more pervasive internationally.

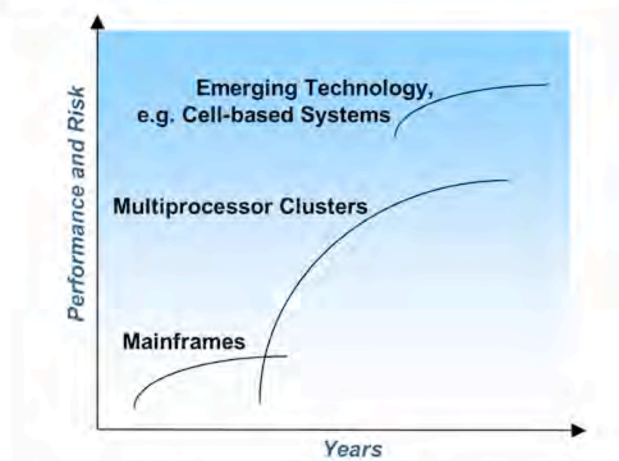

Figure 2.3. Over time, distinct architectures have provided significant advances for HPC functional capability.

Figure is based from communications with Dr. Michel McCoy (LLNL), per: Riding the Waves of Supercomputing Technology [https://www.llnl.gov/str/June03/McCoy.html]

Historically, as Fig. 2.3 suggests, there have been "knees" in architectural developments that required revolutionary thinking about the software required to make the systems effective. For example, going from scalar systems to vector systems required a large investment in research on system software and algorithms, and even sometimes on reformulating the problem. We are now at a new "knee" as we face the probability that upcoming systems will be heterogeneous, with nodes containing hundreds of cores (multiple processors on a single chip). Thus, if we are to continue our competitive edge, we must increase our investment in applied mathematics and computer science research in order to insure that the science and engineering community can utilize these systems effectively to solve the important problems facing the nation.

Any healthy research program must be cognizant of the potential for new paradigms for doing science and must make room for a fraction of the portfolio to be focused on such opportunities. Naturally, not all of these will be successful, but some will lead to revolutionary new ways of doing science. The following represent examples of areas where some research investment might pay large dividends in the future:

- Institutionalize multiscale-multiphysics in SC
- Develop programs in data analytics
- Support collaboratories to advance the pace of scientific discovery
- Support research in Verification & Validation and Uncertainty Quantification

Recommendation: Devote a fraction, for example 40%, of the research budget in applied mathematics and computer science to support high-risk activities that have the potential to make fundamental changes in how scientific discovery is conducted.

The DOE SC INCITE program has proven to be an important mechanism for the advancement of US science by making large allocations of computer time on very high performance systems available to a small number of applications that are judged to represent some of the most important and challenging problems in science and engineering. In order to demonstrate the efficiency of a new system, there may be pressure to favor those applications that can demonstrate high utilization of the system over those whose science is more important, but whose codes are less efficient.

Recommendation: The balance in the INCITE program between scientific importance and code efficiency should strongly favor the importance of the science.

3. Balance Between Leadership Class & Capacity Facilities

In 2001, the ASCAC was charged with evaluating how the roles of ASCR facilities might evolve to serve the missions of SC over the next three to five years. ASCAC provided a preliminary reply in March 2002 (see 11 March 2002 - ASCAC Facilities Subcommittee Report from Jill Dahlburg to Ed Oliver {[http://www.er.doe.gov/ASCR/ASCAC/Reports.html]}), which recommended that "ASCR should build on its present plan to develop a strategic plan for the next generation high-end ... multi-user mission-driven computing environment." To achieve this goal, the report noted that "a range of machines is required within the ASCR portfolio … from research and development of new architectures to the deployment of the multi-user computing facilities of the 21st Century." In briefing this response at the 2 May 2002 ASCAC meeting, the current and developing ASCR portfolio was schematically depicted as a function of priority and importance; see Fig. 3.1a.

The ASCAC Facilities Subcommittee's May 2002 briefing then concluded with the observation that between the time of the March report submission and the May ASCAC meeting, Japan announced installation of the Earth Simulator. In light of this development, the Subcommittee requested that the March 2002 report be considered interim, and that an extension be granted for the Subcommittee to re-address the future roles of ASCR facilities. The subsequent, revised Facilities Report (31 May 2002 - ASCAC Facilities Committee Report (including the impact of the Earth Simulator) from Margaret Wright to Raymond Orbach {website *ibid.*} recommended that "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 US mission-driven research of the DOE Office of Science" and that "planning should assume a funding increment on the order of $150 M/year for the near term." Through an Office vision of extreme computing, this recommendation is being acted upon outstandingly and has resulted in a fundamentally new ASCR facilities portfolio, as illustrated in Fig. 3.1b.

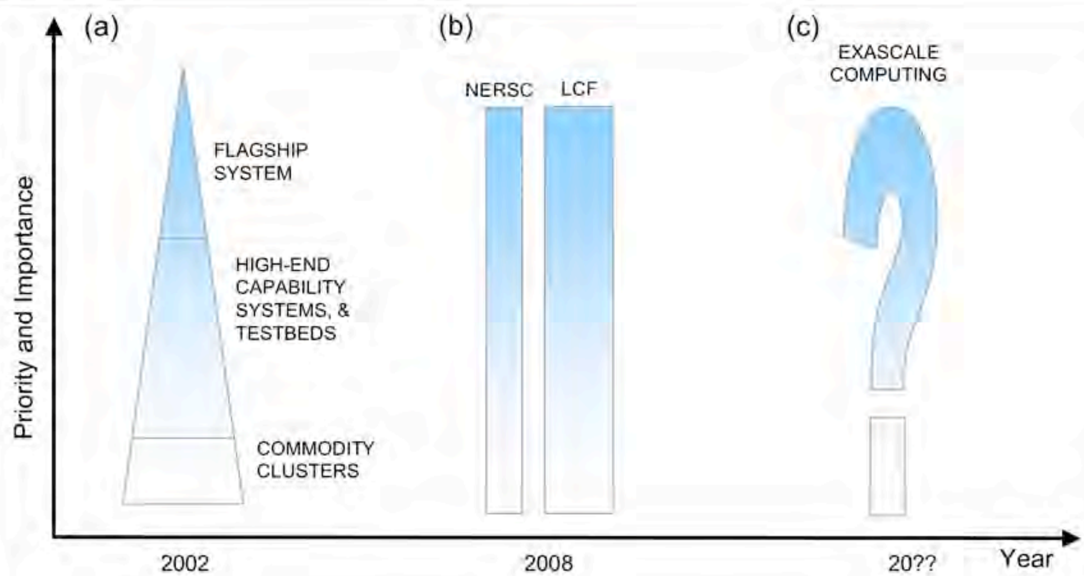

Figure 3.1. ASCR Facilities Portfolio in: (a) 2002; (b) 2008; and, (c) the future.

ASCR's overarching priority has continued throughout to be the use of high performance computing to enable science. The Office's vision of extreme computing, which was brought into focus by the arrival of the Earth Simulator, is enabling the advancement of science through advanced computing simulations both on the LCF at ORNL and ANL, and also on NERSC at LBNL. Beyond simulation and computing, the mission space of SC's intermediate capacity and leadership facility, NERSC, has been appropriately expanded to include the broad emerging HPC data-driven areas of informatics and visualization.

With LCF architectures changing dramatically on sometimes very short timescales, it would be optimal for ASCR to provide the scientific computing user community with access to three tiers of computing facilities:

Tier-1 (ASCR owned)	LCF - Gen 1
Tier-2 (ASCR owned)	Capacity- Gen 2 (unaltered LCF that is 3-5 yrs old)
Tier-3 (user owned)	Commodity- Gen 3

Such a distribution would reduce the risk that computer science and applied mathematics algorithmic advances and implementation know-how made within the architecture-focused LCF user community would be lost to the wider cadres of HPC researchers, and instead would stimulate efficient technology transfer of Leadership Class algorithms from the highest-end user pool to that of production users who tend to resist algorithmic change until solid benefits in both algorithms and infrastructure have been demonstrated. For lasting computational science advancement, this strategy would take best advantage of LCF, which remain truly cutting-edge for relatively short times and therefore must be exploited wisely.

While budget-driven infrastructure constraints may necessitate plans for upgrades rather than entirely new LCF every 3-5 years (thus limiting the plan described above that includes Tier-2 “graying” leadership class systems/highest-end capacity systems on that same time scale), the Panel believes that the ASCR facilities will continue to be in the forefront for the next 5-10 years, and will be exercised with appropriate deliberation. However, changes diagrammed between Figs 3.1a and b, coupled with consideration of Figure 3.1c above, leads to the following.

Recommendation: ASCR should plan for re-competing the LCF on a timeframe for exascale computing, in 7-10 years. Concomitantly, the Office should develop a plan to (a) migrate the then-LCF to Tier-2 facilities; and, (b) support the associated range of different types of core research required.

4. Balance Between SciDAC & Core Computer Science, Applied Mathematics and Networking Research

The invention of revolutionary technologies has historically led to breakthroughs in scientific discoveries: particle accelerators led to the discovery of subatomic particles, electron microscopes led to critical understanding of materials and biological properties, and the space telescope confirmed the evolution of the universe. As our quest for knowledge becomes deeper and more complex, it reaches a point where the most powerful experimental probes or our most sophisticated theories may no longer be adequate. LCF provide an unprecedented opportunity for advancement of science – one that will make it possible to use computation not only as a critical tool along with theory and experiment in understanding the behavior of the fundamental components of nature but also for fundamental discovery and exploration of the behavior of complex systems with extremely large components, including those involving humans. This fact is recognized worldwide as evidenced by the accelerating pace of development of

Leadership Class computing capabilities around the world. Just as advances in diagnostics are essential to extract scientific knowledge from experiments, advances in software are crucial for the productive use of LCF. Note: for purposes of this discussion, core networking research is considered as part of core computer.

PITAC 2005: "Today's computational science ecosystem is unbalanced, with a software base that is inadequate to keep pace with and support evolving hardware and applications needs."

For the US to stay competitive and to remain a leader in scientific discovery, SC must sustain its investment in building world-class simulation capability for large-scale science. ASCR can play an important role in sharing the responsibility for such development by providing near-term support in Computer Science and Applied Mathematics [CS/Math] to program partners in and outside SC, and by investing in long-term research in CS/Math with an eye toward disruptive technologies that will open up new frontiers for scientific discovery. In considering the investment in software research for maintaining our leadership, three evaluation criteria emerged: (i) Is it leading edge science? (ii) Is it large-scale science that takes advantage of HPC capabilities? (iii) Is it agency relevant? Preferably, a project worthy of funding should satisfy all three of these criteria.

Recommendation: ASCR should continue emphasis on support of core research in CS/Math/Networking that has the potential to advance scientific discovery.

In addressing the issue of balance between near-term and long-term research, one can classify research as incremental, disruptive, or fundamental. Incremental research addresses near-term needs, disruptive research can have near-term or long-term goals, and fundamental research is justified by long-term potential. Incremental research can be mission critical because it enhances the chances of success of an SC experimental facility. It has relatively low risk or its success will reduce the risk of a DOE mission critical project. The value of near-term research is illustrated by the success of SciDAC, which is an inter-office and inter-agency collaboration.

2007 COV Report on SciDAC-2: "The SciDAC-2 program is unique given the computational science goals of integrating science and simulation at the petascale level. But it is equally unique because of its broad intellectual scope and a broad administrative scope that cuts across multiple offices within the SC, and includes financial and intellectual participation by the National Nuclear Security Administration [NNSA] and the National Science Foundation [NSF] … the process was very successful despite time pressures and the ambitious and complex nature of the solicitation … there was a remarkable level of coordination amongst the various programs and offices within the Office of Science and partnering agencies … As a result, SC has in place a strong scientific portfolio that is well positioned to address the goals of the SciDAC-2 program."

Many SciDAC projects fall under the incremental category. Their methodologies are approaching maturation, and significant progress can be made with existing high performance computing facilities. As an example, in the past few years particle-in-cell simulation tools for plasma based accelerators have been verified against each other, against experiment and against theory. The availability of high performance computers will allow parametric studies in regimes that will not be accessible to experiments for years to come. However, SciDAC applications projects can also be disruptive research. They have higher risks and higher potential for revolutionizing simulation methodology or scientific discovery. They will likely require access to LCF. For example, there are no three-dimensional models of sufficient realism for core collapse of supernovae because of the computational limitations for simulating multi-scale physics. Only LCF will provide the capability to extend the models to three dimensions, including both macroscopic and microscopic models of stellar core phenomena. Based on the present INCITE allocation process, selected (relatively few) SciDAC projects that are deemed LCF ready receive computer time on LCF. We believe with some incentive and assistance from ASCR, some important applications that are on the borderline of LCF usage may be encouraged to cross the threshold to become LCF users, thereby broadening the base of LCF projects and opportunities for scientific discovery. In other cases, some important applications are ready to make effective use of LCF, but often application teams are not eager to modify their codes for fear that redirection of resources to achieve this would only lead to marginal return. Success with cutting edge, disruptive simulation and data analysis generally requires close interaction between applications expertise and tool developers to alleviate this fear.

Recommendation: Form "SciDAC-like" groups (pilot projects) to accelerate readiness, with attention to disruptive pathways (e.g., productive tools from one application brought to bear on another). For these projects, application developers should be tied directly with specific applied mathematics and computer science expertise in order to bring the application to the point where it is at least a candidate for INCITE allocation if not a SciDAC award.

Such efforts would have the additional advantage of building support within the other Offices for the ASCR program.

Based on information available from ASCR, the funding split between SciDAC related research and basic research is currently approximately 40:60 (*cf.* Fig. 1.1). We recommend a 50:50 balance to enhance the success of SciDAC to ASCR and partnering offices. The increase in SciDAC funding should go toward establishing closer collaboration between CS/Math and domain scientists. Existing and new SciDAC Centers for Enabling Technology [CETs] should be evaluated based on scientific productivity of the entire project and not separate CS/Math components.

Recommendation: Increase the base of support for SciDAC in other SC offices by establishing closer collaborations between CS/Math and domain scientists.

These collaborations need to be monitored closely for effectiveness.

Recommendation: Prune CET's or other collaborative activities that are not advancing applications.

As the SciDAC program is presently managed, winners of a SciDAC award must go through a second review process in order to obtain cycles on the LCF via the INCITE program. We believe this is an unnecessary burden on the HPC applications community. If the SciDAC review process is truly identifying the best science, then that science needs the tools required to meet their stated goals. These needs should be part of their submission and should be reviewed along with the science. If their science passes muster, then they should automatically receive the appropriate access; if their codes are not ready for the LCF, but the science requires that level of computing, then ASCR should establish partnerships with applied mathematicians and computer scientists to bring them quickly to a state of readiness. This interaction may require research, development, implementation, or a combination, and the partnership should be structured accordingly.

Recommendation: Streamline processes for SciDAC access to LCF.

CS/Math research can be incremental, disruptive, or fundamental. Incremental research provides a path for technology transfer to SciDAC and other applications. It requires close collaboration between CS/Math and SciDAC domain scientists. Disruptive research paves the way for the future, be it the next generation LCF architectures or scientific applications. It identifies the roadblocks, develops solutions, and transfers knowledge to future users. Because of its higher risk, possible long-term nature, and focus on critical SC missions, it is unlikely that other communities and other DOE offices will support the required disruptive research in CS/Math. This is the responsibility of ASCR. ASCR has a number of important long-term research and development opportunities in this area that are critical to its mission. Two of the most pressing examples are (a) programming and performance of heterogeneous many-core based high-end computing systems and (b) information science and technology. Heterogeneous many-core processors will likely be the future, with SIMD engines, graphics, network interfaces, encryption, and other specialized functionality integrated on-die. Message passing will continue for some time to provide system level parallelism, but a new programming paradigm must be developed to utilize processor technologies of the future effectively.

Information science and technology is evolving into a foundation for scientific advances in the 21st century. Explosive growth of new technologies will continue to generate increasing amounts of data in the experimental, observational, simulation, and analytical sciences. Analysis and integration of these data with multiple information sources have become critical components of the scientific process for national security, materials science, informatics, and the biological, earth, and social sciences. This shift is creating opportunities to accelerate discoveries in science and technology and to apply these discoveries to problems of national and global importance. Fundamental research is justified by its potential in leading to disruptive technologies. It has the highest uncertainty, and hence should be subjected to more frequent scrutiny. We do not have sufficient information to characterize the present balance between fundamental and disruptive research in basic CS/Math research. However, based on the higher uncertainty of fundamental research and critical importance of disruptive research, a balance of 1:3 seems reasonable, as illustrated in Table 4.1.

Table 4.1 DOE SC ASCR CS/ Applied Math Research Balance.

	Near-Term	Long-Term
Incremental	SciDAC, CS/ Math	
Disruptive	SciDAC, CS/ Math	CS/ Math
Fundamental		CS/ Math

Funding Balance — decreasing

Recommendation: Invest in new, important, long-term opportunities in disruptive R&D, which can lead to new CS/Math paradigms.

5. Summary: Toward a Closer Balance

In summary, the Panel is fully supportive of ASCR's strategy to be the leader in advancing open science through high performance computing. Our findings and recommendations suggest moving to a more even balance between investment in high performance computing facilities and in "core research" as a response to increased complexity in computer design and to provide the most effective use of the facilities. In addition, increased investments are needed to provide for pilot use of LCF, expand SciDAC partnerships, and develop a base for the exploitation of emerging new paradigms for advancing the pace of scientific discovery.

The Panel feels that the question of balance in implementing this vision is important and that ASCR should consider adopting a methodology to plan and evaluate its strategic research investment at regular intervals. The balance we envision is depicted in Fig. 5.1.

Figure 5.1. Recommended balance of programmatic elements in the DOE SC ASCR.

We suggest that to achieve this goal ASCR consider a strategic approach based on the formulation of a research portfolio aimed at achieving beneficial results for DOE and the nation (see, for example, Roussel, P. A., Saad, K. N. and Erickson, T. J.: *Third Generation R&D Managing the Link to Corporate Strategy*, Harvard Business School Press, Boston, MA, 1991). We note that in the FY2008 budget ASCR research areas are already tied to DOE Goals 3.1 and 3.2. However, we suggest that this be carried further, specifically to include and account for the basic research categories of: fundamental research, which creates new knowledge and can lead to new pathways for scientific discovery; disruptive (or revolutionary) research, which can lead to new paradigms or ways of conducting science; and, incremental research, which leads to improvements or extensions of existing tools or methods.

6. Acknowledgements

The Panel wishes to acknowledge with gratitude the many thoughtful comments provided for Panel consideration by members of the US high performance computing community. We offer deep thanks for insightful papers, to Dr. Thomas Zacharia (for the 8 Jan 2008 advice paper: Computing and Computational Sciences at ORNL, A Vision for Sustained Leadership and Scientific Impact) and to Dr Andrew White (for the 11 January 2008 advice paper: A Discussion of Balance in the ASCR Program).

7. List of Acronyms

ANL	Argonne National Laboratory
ASCAC	Advanced Scientific Computing Advisory Committee
ASCR	Office of Advanced Scientific Computing Research
CET	Center for Enabling Technology
COV	Committee of Visitors
DOE	Department of Energy
HPC	High Performance Computing
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
LBNL	Lawrence Berkeley National Laboratory
LCF	Leadership Class Facilities
NERSC	National Energy Research Scientific Computing Center
NNSA	National Nuclear Security Administration
NSF	National Science Foundation
ORNL	Oak Ridge National Laboratory
PITAC	President's Information Technology Advisory Committee
SC	Office of Science
SciDAC	Scientific Discovery through Advanced Computing
SIMD	Single Instruction, Multiple Data

Appendix: ASCAC Balance Charge Letter

June 26, 2007

Dr. Jill Dahlburg, Chair, ASCAC
Naval Research Laboratory, Code 1001
4555 Overlook Avenue SW
Washington, DC 20375

Dear Dr. Dahlburg:

I am charging the Advanced Scientific Computing Advisory Committee (ASCAC) to assess the strategic priorities of the Advanced Scientific Computing Research (ASCR) program, focusing on the balance between the "core" research efforts and high performance computing facilities and on the balance between more immediate research needs of current scientific applications and the long-term investments necessary to effectively utilize the high performance systems of the future.

The ASCR program is delivering leadership computing facilities and advanced networks critical to advancing scientific applications and Department of Energy missions. In that role, ASCR is an enabling partner to the other research programs of the Office of Science and, through INCITE, to American competitiveness. But ASCR also plays a critical role in advancing the underlying applied mathematics, computer science, and advanced network research necessary to effectively utilize the computing and network resources of the future. Because success is often built upon a decade or more of research effort, it is vital for ASCR to carefully balance investments in facilities and research and, within research, to balance the immediate needs of program partners with the long-term investments necessary for sustained progress.

In order to influence funding decisions in ASCR, I would like a full report on findings and recommendations of the ASCAC by the February 2008 meeting. I appreciate ASCAC's willingness to undertake this important activity.

Sincerely,
Raymond L. Orbach
Director
Office of Science