

COMMITTEE OF VISITORS

REPORT

Advanced Scientific Computing Research

April 2004

Date of COV: March 9-10, 2004
Programs: Applied Mathematics, Computer Research and Collaboratories
Research
Office: Advanced Scientific Computing Research (ASCR)
Agency: United States Department of Energy

Committee Membership:

Randolph E. Bank, University of California, San Diego
Marjory Blumenthal, Georgetown University
Howard Elman, University of Maryland
Sally Floyd, ICSI Center for Internet Research
Sharon Glotzer, University of Michigan
Robert Grossman, University of Illinois at Chicago
Cray Henry, HPCMO
Timothy Killeen, Chair, National Center for Atmospheric Research
Joseph McCloskey, NSA, Columbia, Maryland
José Muñoz, National Science Foundation
Linda Petzold, University of California, Santa Barbara
J. Mark Pullen, George Mason University
Steven Richardson, Howard University
Valerie Taylor, Texas A&M University
Linda Torczon, Rice University
Stephen Wolff, Cisco Systems

0.0 Executive Summary

The Committee of Visitors (COV) for the Office of Advanced Scientific Computing Research (ASCR) programs in Applied Mathematics, Computer Science, and Collaboratories met March 9-10 2004 at the DoE facility in Germantown, MD.

The COV is extremely grateful to the program officers and other ASCR staff who gave unstintingly of their time and knowledge to help the committee in its deliberations.

Findings:

Based on presentations by, and interviews with program officers and management, and on examination of project folders of the three research programs, the **COV considers the ACSR programs under review to be generally effective and reasonably well managed**, with allowance given for the complexity of the program, the differential character of the various constituencies involved, and the differing but related missions of the individual programs, office, and agency.

The committee found no cases in which the ultimate decision of the program manager was obviously incorrect or unjustified, nor any instances where the program was skewed in a questionable direction or where the process deviated markedly from the normal standards of peer review.

There was some concern expressed by members of the COV regarding the representativeness of the reviewer pool, owing to the repetition of names of various well-known scientists in the proposal folders and the relatively small number of panel members used for some quite large solicitations.

The COV folder review revealed significant differences in handling and review techniques from program to program and between sub-categories (e.g., between declined and awarded proposals, between initial proposals and renewals, and between national lab and university proposals), which made intercomparisons and “roll-up” statistics difficult to obtain. It was not possible, for example, to readily gain a summary of geographic or demographic distribution for awards or for declinations.

The material in the folders was, in various cases, insufficient to track the events leading to the final action; folder information in these cases had to be supplemented by data drawn from program officers’ personal filesystems in order to evaluate the complete train of events. In general, it was the view of the COV that **the folders appeared to be designed more for “fiscal management” than for “program management,”** and that **this constituted a limitation of the current process** and led program officers to employ ad hoc and idiosyncratic procedures to manage their programs effectively.

Recommendations:

The COV applauds the development of program management guidelines for ASCR (presented on the first day of the review) and believes that such guidelines should be used in the future to make the review processes more comprehensive, uniform, consistent, and easy to evaluate.

The COV was particularly pleased to learn of the ASCR commitment to professional development of new scientists and applauds the Early Career PI Program, which has led to funding of more than 20 new investigators over the past two years. The COV believes that the ASCR would benefit from a formal longitudinal evaluation of the efficacy of this important and apparently very successful new program, which was started in 2002.

The COV noted that the cross-cutting Scientific Discovery through Advanced Computing (SciDAC) program has had a major influence on all three programs under review. There was concern within the COV over SciDAC's future, and **the COV believes that ASCR should prepare a strategic plan that covers the future of SciDAC.**

At the start of the meeting, the COV was informed that newly developed guidelines would be promulgated to ensure, amongst other things, that “university and lab proposals will be handled equally”. **The COV believes that ASCR need not mandate complete equality between the lab and university processes**, since ASCR clearly has a major and unique mission to support cutting-edge lab science. However, the COV feels that **both sets of the folders should contain sufficient information to detail the decision-making process and track continued progress for renewed awards, and that similar (or identical) robust standards for tracking decision processes and documentation should apply to both types of proposals.**

A major recommendation of the COV is that ASCR develop a more comprehensive and consistent approach to program review documentation. Formal records should include the key information that document the process and justify the resulting decision in context.

The COV made a number of specific recommendations for the conduct of future COV reviews.

TABLE OF CONTENTS

	<i>page</i>
0.0 Executive Summary	2
1.0 Introduction	5
2.0 Overview of Findings	6
3.0 Integrity and Efficacy of Individual Program's Processes & Management	9
3.1 Applied Mathematics Research	9
3.2 Computer Science Research	12
3.3 Collaboratories Research	13
4.0 General Programmatic Findings and Recommendations	13
4.1 Balance between National Laboratory and University Participation	13
4.2 Evaluation of Program Impacts	14
4.2.1 Applied Mathematics Research Highlights	15
4.2.2 Computer Science Research Highlights	15
4.2.3 Collaboratories Research Highlights	15
4.3 Strategic Planning	16
4.4 Committee of Visitor Process	17
A1 Full List of COV Members	19
A2 Letter Charge to Committee from ASCAC Chair	20

1.0 Introduction

The Advanced Scientific Computing Advisory Committee (ASCAC) for the Office of Science, United States Department of Energy (DOE), was charged by Dr. Raymond Orbach, Director of the Office of Science, with assembling a Committee of Visitors (COV) to review the research programs of the Office of Advanced Scientific Computing Research (ASCR) in Applied Mathematics, Computer Science Research, and Collaboratories Research. The list of participants in the COV is provided in Attachment 1 and the letter charge to the committee from the chair of the ASCAC is provided in Attachment 2.

This report presents the findings and recommendations of the COV. The committee members recognize that this is the first time that the respective DOE programs have been evaluated in this manner and, as a consequence, paid special attention to recommendations for future improvements in the review process.

The COV met at the DOE Germantown location on Tuesday, March 9 and Wednesday, March 10, 2004. The meeting began with a series of presentations by the DOE program officers. Dr. Gary Johnson provided a summary of the charge to the committee, discussed the DOE conflict of interest policy, and presented helpful summaries of the ASCR programs under review. Each program manager then made presentations describing their specific program's goals, accomplishments, and the approach used to evaluate proposals. These program managers were Drs. Fred Johnson and John Van Rosendale (Computer Science Research), Dr. Charles Romine (Applied Mathematics Research) and Dr. Mary Anne Scott (Collaboratories Research). The COV also received a presentation from Dr. Martin Rubenstein of the DOE Division of Grants and Administration describing how the ASCR programs comply with federal policies and the relevant OMB circulars.

Following the program manager summaries, which included information on program growth, proposal success rates, PI demographics and other topics, the COV split up into three subgroups to review proposal folders for each program. These folders included documentation for both DOE national laboratory and university-led proposals. Following the initial review of the program folders, the COV met with the program managers to discuss preliminary findings and to request further information to aid in the development of this report.

On the second day of the review, the requested additional information was presented and discussed by the program managers. The additional material included an overview of budget data, a spread sheet of all proposals in the awarded and declined categories, lists and descriptions of program highlights, and additional statistical information on proposal pressure, success rates, and review panel membership. The COV continued the analysis of the documented review processes and expanded its work to include an evaluation of program outcomes and plans. A synopsis of the draft report was discussed with the program managers before the end of the meeting and the final report itself was prepared using subsequent e-mail exchanges with COV members. The COV members were

grateful for the active and helpful engagement of the ACSR program managers throughout the review process.

The specific charge to the COV included the following four elements:

1. An assessment of the efficacy and quality of the processes used to solicit, review, recommend, and document proposal actions and monitor active projects and programs.
2. An assessment of how the award process has affected the depth and breadth of portfolio elements and the national and international standing of the portfolio elements.
3. Commentary on future directions proposed by ASCR management and on opportunities that might not have been presented
4. Commentary on the proposal review process, and on how the COV evaluation process might be improved in the future.

The charge also requested that the COV address such questions as:

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

2.0 Overview of Findings

2.1 Efficacy and Integrity of Processes

The COV considers the ACSR programs under review to be generally effective and reasonably well managed, with allowance given for the complexity of the program, the differential character of the various constituencies involved, and the differing but related missions of the individual programs, office, and agency. The program officers are clearly dedicated and competent public servants who have considerable knowledge of their respective portfolios and communities of practice. Each program has achieved significant successes and some of the programmatic highlights of the respective programs are mentioned in Section 4.2.

The format and time allocated for the COV review process allowed for a substantive, but not complete, evaluation of the ASCR proposal review processes and enabled the committee to address many, but not all, of the elements of the charge summarized in Section 1.0.

The COV was presented with a subset of the proposal folders in each of the three subject programs. The COV also requested, and was provided with, additional folders during the meeting and was given access to a complete list of Principal Investigators (PIs) for both awarded and declined proposals. Some of the program managers listed the names of panel reviewers for specific solicitations, but a complete list of panel and mail-in

reviewers was not available across all programs. All requested documentation present in the central file system was quickly provided to the committee, though some relevant data (such as comprehensive reviewer lists) were apparently stored on personal filing systems and were not reviewed for that reason. The initial set of folders did not include declinations, although a sampling of such folders was later made available on request (though not for all three programs).

The committee found no cases where the ultimate decision of the program manager was obviously incorrect or unjustified. Similarly, the COV found no instances where the program was skewed in a questionable direction or where the process deviated markedly from the normal standards of peer review (although COV concerns with regard to the peer review process are mentioned below). Proposal reviewers were clearly well informed and knowledgeable of their respective areas. There was some concern expressed by members of the COV regarding the representativeness of the reviewer pool, owing to the repetition of names of various well-known scientists in the proposal folders and the relatively small number of panel members used for some large solicitations. In the case of mail-in and panel reviews, a minimum of three reviewers was the standard across all three programs. A full evaluation of the adequacy of the reviewer pool was not possible given the lack of readily available and comprehensive reviewer pool statistics and associated data.

The COV folder review revealed significant differences in handling and review techniques between programs and between sub-categories (e.g., between declined and awarded proposals and between national lab and university proposals). There was considerable unevenness in the level of detail, format, content, and organization of the documentation, which made it difficult for the COV to gain a thorough picture of processes. Each folder contained the subject proposal, but, in some instances, only the university-led proposal folders had detailed decision memos written by the program managers. A variety of techniques were used for proposal evaluation, and the program managers clearly exercised significant discretion in deciding which technique was appropriate for specific solicitations. The COV felt that this level of discretion was appropriate. These methods included “pre-proposal” evaluation by program managers and subsequent encouragement/discouragement of a full proposal, large and small panels for targeted solicitations, and mail review. The COV did not see or review any documentation relevant to the pre-proposal screening processes and no complete lists of pre-proposals were in the folders, making it difficult to comment on the quality of the procedures used to screen out proposals at the pre-proposal stage.

The COV discerned several significant differences between how proposals from the national labs and universities were handled, including the style and format of the respective proposal folders. There were also significant differences between the processes used in each of the individual programs, which made comparisons and “roll-up” statistics difficult to obtain. It was not possible, for example, to readily gain a summary of geographic or demographic distribution for proposer awards or for declinations. Renewal proposal folders (particularly for lab proposals) generally contained less information concerning the decision processes, relative to first-time or “new” proposals. A substantial

amount of time at the meeting was consumed by supplementary information requests and associated dialog with the program managers. A recommendation for the future will speak to the need for improved database systems to enable such queries to be more readily addressed ahead of, as well as during, the COV meetings. One complicating factor was that, in some instances, large team proposal activities were documented in different (and separately filed) folders, apparently one per Co-PI institution. This multiplicity of folders for the same collaborative activity made it difficult to gain a complete overview of the team effort. Specific issues with regard to the folder review are described in Section 3.0 under the respective programs.

The COV applauds the development of program management guidelines for ASCR (presented on the first day of the review) and believes that such guidelines should be used in the future to make the review processes more comprehensive, uniform, consistent, and easy to evaluate.

The COV was particularly pleased to learn of the ASCR commitment to professional development of new scientists and applauds the Early Career PI Program, which has led to funding of more than 20 new investigators over the past two years. The COV believes that the ASCR would benefit from a formal longitudinal evaluation of the efficacy of this important and apparently very successful new program, which was started in 2002.

The ASCR program uses the panel review mechanism extensively for the larger program elements. There was a concern expressed by members of the COV that, in certain cases in some programs, the panels were not sufficiently large to ensure thorough representation of the various required areas of expertise; the COV suggests that specific guidelines be developed for such panels. In one instance a (roughly) \$20M program was reviewed by a panel comprising only 4 expert referees. The COV believes that, in general, the number of expert reviewers comprising a panel should be (in some way) commensurate with the magnitude of the funding level for the specific research opportunity.

The COV discussed the DOE's Scientific Discovery through Advanced Computing (SciDAC) program in considerable detail throughout the visit. It is clear that this cross-cutting DOE activity has had a major influence on ASCR plans within all three programs under review. There was concern (discussed in greater detail in section 4.3, below), however, that the SciDAC program's future was in some jeopardy, and that strategic planning within ASCR related to the future of SciDAC was urgently needed.

Despite some of the procedural issues discussed above (which varied considerably from program to program), the COV believes that each ASCR program manager demonstrably and regularly used inputs from the peer review process (sometimes a panel and sometimes mail-in reviews) and exercised proper judgment in arriving at the final funding/no funding decision. The level of detail in the folders documenting these decisions differed from program to program and specific issues are identified below in the appropriate sections. The material in the folders was, in various cases, insufficient to evaluate the complete train of events leading to an action. In general, it was the view of the COV that the folders appeared to be designed more for fiscal than for programmatic

management And, furthermore, that this constituted a limitation of the current process that should be addressed

At the start of the meeting, the COV was informed that newly developed guidelines would be promulgated to ensure, amongst other things, that “university and lab proposals will be handled equally”. The COV noted that the documentation in the lab folders was less complete than for the university folders. The COV believes that the ASCR program need not mandate complete equality between the lab and university processes, since the ASCR program clearly has a major and unique mission to support lab science. However, the COV feels that both sets of the folders should contain sufficient information to detail the decision-making process and track continued progress for renewed awards. The COV feels that the proposal decision processes for lab and university proposals can differ, if warranted by the nature of the particular solicitation. However, it is the strongly held view of the committee that similar (or identical) robust standards for tracking decision processes and documentation should apply to both types of proposals.

A major recommendation of the COV is to suggest that ASCR develop a more comprehensive and consistent approach to program review documentation. Formal records should include the key information that document the process and justify the resulting decision in context. The COV is very optimistic that such improvements in process will quickly lead to measurable gains in ASCR program efficiency.

3.0 Integrity and Efficacy of Individual Program’s Processes and Management

3.1 Applied Mathematics Research

In some ways, the Applied Mathematics Research program was the easiest one of the three for the COV to assess, in part because its objectives appeared to be relatively straightforward and readily understood. The COV believes that the program manager of the Applied Mathematics program, Dr. Charles Romine, has done an excellent job in managing this program, and has demonstrated significant initiative and leadership for the program and its community, notably in his development of the exciting new program in multi-scale modeling.

The long-lived Applied Mathematics Research program has achieved significant impact across a broad spectrum. Program highlights include progress in developing mathematics for simulation of combustion physics, computational fluid dynamics and laser-plasma interactions. In general, the documentation provided by the program manager was in very good shape, with clear decision memos in the folders, based on an analysis of the referee and/or panelist comments. The program manager provided the COV with lists of distinguished panelists used in the review process and the COV felt that the size and representativeness of the panels were both adequate for the specific opportunities evaluated. The COV understands that Dr. Romine left ASCR shortly after its visit, and it urges ASCR to find a similarly effective replacement.

A set of approximately twenty proposals funded by the Applied Mathematics Program was reviewed by a subcommittee of the COV. The proposal evaluation procedure was generally well documented, with at least three reviews (collected by email) used for each proposal and rationales for funding spelled out by the program officer. In most cases the timing of communication to PIs was reasonable (with decisions within one year), although there were examples of reviews that took considerably longer; the reasons for delays were not transparent.

It was felt that the intrinsic quality of the funded proposals was high, and that this is consistent with a long tradition of quality within the Applied Mathematics Program. However, there were projects whose initial reviews were very strong but for which the evaluations of later (funded) renewal requests were less enthusiastic, leading to a COV concern about sustained standards of excellence for long-running efforts. On the other hand, COV found at least one case where a national lab renewal proposal was declined and the members felt that the program manager was exercising an appropriate level of quality control and initiative in reaching this (probably difficult) decision.

Some projects were funded under the auspices of the SciDAC program. It was not clear how the percentages of funding coming from SciDAC were determined. More importantly, it was felt that the SciDAC program itself was not as well publicized or as open to the general community as it could be. There was also an interest on the part of COV members in reviewing more general program information, including percentages of proposals that are funded, demographics of funded proposals, for example, by state, diversity, and types of institutions. In the future, the COV recommends that such information be made available prior to the review.

3.2 Computer Science Research

The ASCR Computer Science Research program is supporting leading work in various areas of computer science and has also achieved significant impact. A stated purpose of the program is to conduct computer science research with a view to support application needs. Key accomplishments include the development of the MPI message passing model, and various compilers and toolkits, including the OSCAR toolkit for the management of Linux clusters. The program managers are knowledgeable and experienced.

Dr. Fred Johnson provided a summary of the processes used. There are two types of solicitations in the Computer Science Research program: Open and Specialized. Both solicitations are announced on the web. Typically, panels are convened for specialized announcements (with at least 3 panelists). Open announcements are typically reviewed by mail-in referees (at least three per proposal). The individual reviews are augmented by a summary decision memo written by the program manager and this memo is included in the university folders (but not the lab folders). The PI is required to write an annual report, which is also included in the documentation. PI's for declined proposals are given access to edited versions of the reviewer comments.

The information provided for COV review was not sufficient to allow the COV to reach a complete assessment of the solicitation, review and selection process in this program. Dr. Johnson, however, was very helpful in addressing resulting queries and a significant amount of time was spent in oral question and answer sessions to fill in gaps in the COV's understanding and appreciation of the process. In this regard, the COV subcommittee was greatly appreciative of Dr. Johnson's patience and willingness to address any and all questions. The primary documents provided to the COV included a selection of funding/contract management folders. The initial set were all for successful proposals, though declination folders were later provided and reviewed. Further discussion with the program manager clarified that there were additional ad hoc documents that were significant in the conduct of the process, including reviewer and panelist lists and summary reviewer comments (for lab proposals).

It proved difficult for the COV subcommittee to navigate through the proposal folders for various reasons involving non-uniform formats and the lack of navigation aids (tabs, table of contents, etc.). The university folders contained summary decision memos written by the program manager, but no such memos were routinely found in the lab folders. Lab renewal folders contained surprisingly little evaluative information.

In some cases, a lot of supplementary verbal information was needed to determine what the COV was reviewing. For example, the program management decision to support the "Center for Programming Models for Scalable Parallel Computing" resulted in 10 independent folders, one for each participating university and laboratory. These folders were filed alphabetically according to the Co-PI institution and, for that reason, proved to be difficult to assemble as a group. The individual folders captured the review of a university's or lab's portion of the proposal, the associated review comments, and the financial actions taken to fund the work. Each folder appeared to document an independent decision as opposed to documenting a partial implementation of a larger integrated team effort. From further discussion, it was clear that the complete proposal was reviewed and evaluated as a unit and that other (ad hoc) records existed to document the composite review and decision process. However, these other records were not available for review.

The information included in a folder typically did not permit the COV members to determine how the funded proposal was ranked relative to other evaluated proposals, so it was hard to determine if the "best" research was funded (as per the COV charge). This was also true for declinations. However, each decision was obviously contingent on the program manager's analysis of the referee comments in the file and, for the sample of folders examined, the COV found no examples of clearly incorrect or questionable decision making.

In their review of the Computer Science Research program, the COV members tended to see the same names time and again, as Co-PI's and/or PI's and as reviewers. While there were not sufficient data to fully evaluate this issue (e.g., pre-proposals were not listed or the screening process available for detailed review), it was the sense of the COV that the ASCR program would benefit from methods (e.g., newsgroups, open workshops, etc.) to

increase the size and diversity of the proposer and reviewer pools (this last comment applies to all three programs reviewed)

The COV found that the mail review for the Computer Science program occurs asynchronously. Some members felt that a plan could be developed for each solicitation that would enable proposals to be collected and reviewed more regularly and synchronously as a set. This might make more effective use of reviewers and result in more focused proposals for such opportunities. In addition, the COV found that the number of panelists was not always matched to the size of the opportunity and recommends that ASCR provide guidelines for panel size, commensurate with the magnitude of the specific opportunity.

The Computer Science Research program should implement a uniform process that includes consistent record keeping. Formal records should include the key information that document the process and justify the resulting decision in context. There should be sufficient information retained so that the outcome of an individual proposal can be compared with others in the same competition (including pre-proposals). Information as to how a particular proposal was ranked in comparison to all the other proposals should be retained. The program manager's summary of individual reviews should be retained as part of the proposal folder in all cases. In the case of renewal proposals, the annual report from the PI appears to be the principal formal mechanism for determining if funding should be continued and for monitoring progress. The COV recommends that ASCR review this component of the process to determine if additional means could be used to monitor progress and determine continuation funding. Finally, the COV recommends that ASCR investigate whether review templates should be required in all cases.

3.3 Collaboratories Research

The COV found that the Collaboratories Research program is well-managed, innovative, and has had significant impact both within and outside of the DOE. Significant accomplishments given for the Collaboratories Research are mentioned in section 4.2.3, below. The program manager, Dr. Mary Anne Scott, is knowledgeable and closely tuned to community developments and needs. In general, the review processes and quality of the documentation are very good. Dr. Scott does a fine job managing the program and interacting with the broad external research community. The COV was impressed with the number of written documents and polished published reports from workshops emanating from this program.

The COV subcommittee was also impressed by the evident high national and international standing of the Collaboratories Research portfolio, particularly given the relative newness of the program. The program manager suggested that a key metric for success is the degree to which the research products are used by external communities. By this measure, the program has had considerable success, since many of the innovative developments have already seen broad community use. Nevertheless, COV members

were concerned that, with diminishing funding, and absent some intervention or aggressive planning process, program sunset or transition looms ahead.

In reviewing the proposal folders (no declination folders were provided or reviewed), the COV developed a particular concern about whether the program is fully open to university researchers that do not have existing collaborations with DOE laboratories. Broadening the proposer and reviewer pools might extend the reach of this already influential program. The COV was appreciative of the inter-agency interactions and of the support for standards development.

A significant portion of the Collaboratories Research program is derived from SciDAC funding. The COV expressed some concern about the continuation and stability of these funds.

4.0 General Programmatic Findings and Recommendations

4.1 Balance between National Laboratory and University Participation

The DOE program directors for Applied Mathematics, Computer Science, and Collaboratories provided significant information to the COV about the proposal solicitation and review process, as detailed in Section 3.0. A particular concern of the committee centered on how lab and university proposals were treated and whether or not differential considerations applied. In general, it became clear to the COV that all proposals are, in fact, reviewed for intellectual merit in much the same way, regardless of whether they originate from a lab or a university. Final programmatic decisions, however, clearly take into account the degree of focus on the DOE mission. In this regard, proposals from national labs often seem to have a competitive advantage and this was felt by the COV to be reasonable, given the key mission of the programs to support innovative research at the national labs. However, in general, the documentation for lab proposals was less rigorous than for university proposals and made a full evaluation of the quality assurance for lab proposals difficult to perform.

Some members of the COV also raised concerns about the degree of openness of the ASCR programs and the committee recommends that the ASCR leadership consider approaches to increase community awareness of the research opportunities and, as a consequence, augment the size of the proposer and reviewer pools. With respect to awareness of DOE Programs, information is publicly available on the DOE website at <http://www.er.doe.gov/grants/>. The COV noted, however, that, while this site is very-well known by the small community that has DOE funding, the program managers should work to implement ways to make this information known to a broader community. Some additional approaches worthy of consideration might include:

- Disseminating information about the initial posting of programs by leveraging the relationship with CRA to publish an article in the CRN and sending email to departments heads via the Forsythe list.
- Continue giving workshop updates at conferences such as SC, SIAM, ACM SIGs, and other conferences as appropriate for the DOE programs. It is recommended

that this be done on an annual basis so that the communities expect the updates. Further, having the updates at well-attended conferences allows for dissemination of information to new communities.

- Institute (and staff) a “research corner” at DOE SC booths.

With respect to increasing the involvement of different or a wider range of researchers in the review process, it is important to address the travel limitation issue as well as the development of new communities. It was noted that ASCR program managers have reached out to new communities with the review process for the Computer Science area, only to find the reviews to be disappointing in terms of depth and breadth of knowledge about a particular proposal topic. While the initial results may be discouraging, a persistent policy of this kind will eventually bear fruit. The following are some suggestions:

- Utilize collaborative technologies to eliminate the travel days for panelists to participate in the review process. It is noted that such technologies are not currently ready for normal use for the review panels; effort and resources are needed to make this technology stable for a panel meeting involving four to six panelists. It is important that such technologies be made a part of a critical path, to insure advancement in the technologies.
- Increase the involvement of different researchers in the workshops, which provide input on the development of programs. While it is important to continue inviting the top researchers in the field, it is important that new researchers with excellent ideas also be identified. One suggestion is to continue with what was done with the HECRTF, that is, to have an open call for white papers related to the issues to be addressed.

The COV felt that the interactions between university and lab scientists could be strengthened. Such interactions are fruitful when there is obvious benefit for all parties involved. Not only is it important for universities to become more aware of ASCR research programs, it is vital that ASCR assist the national labs to expand their circles of university interactions – especially among younger, less well-established researchers. In addition to collaborating on research, the DOE labs provide leading-edge computing, networking, and collaborative resources that are critical for experimental work in high performance computing.

4.2 Evaluation of Program Impacts

The COV spent some time reviewing the major accomplishments of each program mentioned by the program managers and elected to highlight several important outcomes of each program, illustrating the overall impact of the research. It notes, however, that prior knowledge from the fields of expertise of the COV members represented the biggest source of insight. The time constraints of the visit and the emphasis given to the review of proposal handling processes limited the ability of the COV to conduct a comprehensive evaluation of ASCR program impact.

4.2.1 Applied Mathematics Research Highlights

The Applied Mathematics program has had significant impact in numerous areas, including:

- Computational modeling of combustion
- Computational fluid dynamics
- Laser-plasma simulation
- Shock wave theory
- Statistical mechanics and combustion

The impact of the program in these areas has been enabled through its support of specific targeted projects, as well as in fundamental algorithmic methodologies such as optimization and algebraic multigrid for discrete partial differential equations.

An indication of the international recognition of its programs is the set of awards given to its PIs, including a National Medal of Science (J. Glimm, 2003), Norbert Wiener Prize in Applied Mathematics (J. Sethian, 2004, A. Chorin, 2000), a SIAM/ACM Prize in Computational Science and Engineering (J. Bell and P. Colella), and election to the National Academy of Engineering (L. Petzold).

4.2.2 Computer Science Research Highlights

It was more difficult to evaluate the impact of the Computer Science research program, because the program's goal is to support the needs of DOE scientists using HPC facilities, and this is rather small community in the computer science world. However, it is clear that the Computer Science research program is having significant impact in this community. Important noteworthy work includes the development of the MPI message passing model and the MPICH reference implementation, compilers, systems software, and toolkits such as the OSCAR toolkit for management of Linux clusters.

The COV also noted that the 2003 Sidney Fernbach award was given to Dr. Jack Dongarra, for work conducted under the auspices of the Computer Science Research program.

4.2.3 Collaboratories Research Highlights

Major accomplishments of the Collaboratories Research program include the development of the Access Grid, a group-to-group collaboration tool used by over 150 institutions worldwide; , a world-wide de facto standard for base software for grids; Grid2003, a recent production grid used by multiple science communities; and the Electronic Notebook, a record of ideas, data, and events for joint experiments and research programs. The committee felt that the record of accomplishment for Collaboratories Research program was more than reasonable, given its stage of maturity,

4.3 Strategic Planning

The COV believes that the ASCR programs should carry forward a strategic planning process, designed to identify and address several important emerging opportunities for the DOE. The following paragraphs mention some of the key considerations that should be discussed in a planning context.

The SciDAC Program represents an important investment on the part of the DOE in advanced software and hardware developments to address important scientific challenges for the ASCR program. Clearly, the SciDAC activity has had a significant impact by improving the effectiveness and efficacy of key DOE mission applications. The COV discussed the SciDAC program and its role within the ASCR programs extensively and believes that its basic function is extremely valuable and should be sustained in the future. A strategic review is needed to consider adjustments to the activity and plan for further developments after the initial 3-year period of performance. There is concern that the SciDAC effort needs to be more open and better integrated to support both the core research program and applications. For example, SciDAC's own success creates the need to maintain the application codes that have been developed and adapted through the program.

The development of "leadership class" computing systems, presently under discussion within the DOE, represents an important opportunity to support new, cutting edge scientific research that can only be enabled by formidable terascale computing capabilities. As documented in numerous studies, this class of system is needed to support national science-based modeling and terascale simulation needs. The COV commended ASCR for its leadership in the definition and advocacy for improved national capabilities in high-end computing infrastructure.

Efforts to develop new expertise in computational and computer science are important to assure continued US scientific and technical leadership into the future. Professional development activities, such as those taking place within the ASCR programs, can have significant economic benefit on the economy over time and can help create jobs. The existing effort to cultivate US citizen interest and capability in high performance computing and computational science needs to be expanded and creative methods need to be found to encourage further growth in the fields. The COV commended the ASCR program officers for their evident interest and commitment to this important role of the DOE.

A long-term research focus within the ASCR program is important and should be fostered. Clearly there is a need to support important ongoing projects, as well as to enable new ideas to surface and be explored. This dual need creates a tension between support of on-going long-term work that leads to incremental improvements and new research that might lead to revolutionary change. The COV believes that careful planning to balance the research portfolio will pay dividends in the future.

The COV believes that multiscale algorithm research, such as proposed within the Applied Mathematics Research program, is increasingly essential to support the

advancement of multidisciplinary applications targeted to solve real problems. For example, climate science, combustion science, radiative transfer, and laser simulation science are all examples of important DOE areas that will benefit from the multiscale approach. ASCR could exhibit significant national leadership in this arena.

The strategic planning process needs to consider all possible high end computing architectures and solutions, as well as grid computing approaches.

Finally, the COV encourages the ASCR program managers to work to facilitate appropriate multi-agency efforts across the federal government. For example, the commercial supercomputing community is small enough that greater cooperation across the federal agencies could create greater market impact. Open source “Infrastructure” HPC software (e.g. OS, compilers, performance tools, libraries, middleware) could be shared and jointly supported across several HPC lead agencies (e.g. DOE, NSF, DOD, etc.). These types of software tools will not be developed for or by the commercial market; they will only exist tailored for the HPC market if the Federal Government supports their development and on-going maintenance. This burden should be shared across several agencies. (This is more properly viewed as a mix of infrastructure support and advance development.)

4.4 Committee of Visitor Process

The COV spent time towards the end of the meeting discussing the nature of the review process, with a view to making recommendations for possible modifications to the process in the future. The COV certainly found the review process to be both interesting and informative and believes that a regular COV review of the ASCR programs will be of significant benefit to the DOE. The committee members were very pleased with the accessibility of key program managers. The mix of plenary and breakout sessions was efficient and productive. Accessibility to multiple rooms for the breakout sessions was helpful. The dinner discussions between ASCR and COV members on the evening of the first day were particularly useful and enabled the COV to gain a deeper appreciation of the program challenges. Members felt that the composition of the COV was balanced and well suited to the task.

For future COVs, we would like to suggest modifications to facilitate the review process. In particular, for charge to be more completely addressed, future COVs should require the following information and data, some of which should be made available prior to the next visit, as indicated.

- ❖ Materials and information to be made available prior to the review (over password-protected Web site):
 - Strategic plan(s) and related documents that can help to explain the ASCR mission.
 - Short overview of the ASCR program and its sub-elements including facilities.
 - Summary information about proposals received and complete statistics on funding, including:

- Breakdown of awarded funding for all of ASCR and by each of the three research programs (SciDAC vs. base, lab vs. university)
 - Complete list of proposals and PIs (both funded and declined)
 - Success rates by categories (SciDAC vs. base, lab vs. university)
 - Lists of pre-proposals received, their status, and the approach(es) used to screen them
 - Demographics of funded proposals
 - Number and description of “new starts”, including ECPIs
 - Detailed mapping of funded proposals onto strategic thrust areas of ASCR.
- ❖ Materials and information to be made available at the review:
- Copies of all slides available at the start of the meeting
 - *Representative* sample of proposals to review including:
 - Funded and unfunded proposals
 - Strong, weak and borderline proposals
 - Lab and university proposals
 - Individual investigator proposals and multi-institutional proposals
 - For multi-institutional proposals, a single jacket that gives a comprehensive view of the project, including a comprehensive annual report, results of project-level reviews, etc.
 - Complete, detailed description of the research funding and monitoring process
 - Solicitation
 - Preproposal review
 - Proposal review
 - Selection
 - Reporting
 - Monitoring
 - Comprehensive list of all reviewers used, categorized as panel or mail-in reviewers.

To best utilize the COV’s time during the review the COV suggests the following logistical improvements:

- ❖ Improve overall standardization of information and presentations.
- ❖ Tailor talks to more specifically address the issues in the COV charge.
- ❖ Prepare and circulate a fixed schedule early in the process.
- ❖ Expedite the process of getting reviewers to and into DOE.
 - Provide centralized transportation to/from hotel and DOE.
- ❖ Provide onsite access to local networks and printers, including access to ASCR proposal database.
- ❖ Ensure that a sufficient number of each handout is available.

Attachment 1. Full List of COV Members

<p>Professor Randolph E. Bank Department of Mathematics University of California, San Diego La Jolla, CA 92093-0112</p>	<p>Dr. Joseph McCloskey 5385 Woodnote Lane Columbia, MD 21044</p>
<p>Dr. Marjorie Blumenthal, Associate Provost Office of the Provost Georgetown University 37th and O Street, NW Washington, DC 20057</p>	<p>Dr. Jose L. Munoz National Science Foundation CISE/SCI, Suite 1122 4201 Wilson Boulevard Arlington, VA 22230</p>
<p>Professor Howard Elman Department of Computer Science University of Maryland College Park, MD 20742</p>	<p>Professor Linda R. Petzold Department of Computer Science University of California, Santa Barbara 124 Engineering III Building Santa Barbara, CA 93106-5070</p>
<p>Dr. Sally Floyd ICSI 1947 Center Street, Suite 600 Berkeley, CA 94704-1198</p>	<p>Professor J. Mark Pullen Computer Science Department George Mason University, MSN 4A5 4400 University Drive Fairfax, VA 22030</p>
<p>Professor Sharon Glotzer Material Engineering & Chemical Engineering University of Michigan 2300 Hayward Street Ann Arbor, MI 48109</p>	<p>Professor Steven L. Richardson Department of Electrical and Computer Engineering Howard University 2300 6th Street, NW Washington, DC 20059</p>
<p>Professor Robert Grossman Director, National Center for Data Mining (M/C 249) University of Illinois at Chicago 851 South Morgan Street Chicago, IL 60607</p>	<p>Professor Valarie Taylor Department of Computer Science H. R. Bright Building, Room 305B Texas A&M University College Station, TX 77843</p>
<p>Dr. Cray Henry HPCMO Suite 510 1010 Glebe Road Arlington, VA 22201</p>	<p>Dr. Linda Torzcon Department of Computer Science Rice University 6100 Main Street, MS 132 Houston, TX 77005-1892</p>
<p>Dr. Timothy L. Killeen (COV Chair) National Center for Atmospheric Research 1850 Table Mesa Drive Boulder, CO 80305</p>	<p>Dr. Stephen Wolff (ASCAC Representative) Cisco Systems 5116 Rockwood Parkway, NW Washington, DC 20016-1916</p>

Attachment 2. Letter Charge to Committee from ASCAC Chair

December 19, 2003

Dear [COV Committee Member],

The Advanced Scientific Computing Advisory Committee (ASCAC) for the Office of Science, United States Department of Energy, has been charged by Dr. Raymond Orbach, Director of the Office of Science, with assembling a Committee of Visitors to review the research programs of the Office of Advanced Scientific Computing Research (ASCR) in applied mathematics, computer science, and networking. On behalf of ASCAC, I thank you very much for agreeing to act as [Chair / a member] of this Committee of Visitors – the first ever for ASCR.

Dr. Orbach's formal charge includes the following four elements:

1. For both DOE laboratory and university projects, the COV should assess the efficacy and quality of the processes used to
 - a. solicit, review, recommend, and document proposal actions, and
 - b. monitor active projects and programs.
2. Within the boundaries defined by DOE missions and available funding, the COV should comment on how the award process has affected:
 - a. the breadth and depth of portfolio elements, and
 - b. the national and international standing of the portfolio elements.
3. The COV should comment on future directions proposed by ASCR management and on opportunities that might not have been presented.
4. The COV should comment on how the proposal review process, and the COV process, might be improved.

The COV is expected to consider, as well as process, broader issues of program quality, relevance, and performance. Thus the COV should answer questions like

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

COV members will be provided with background material about the ASCR research program before the review. As further background, two Committees of Visitors have been organized by the Basic Energy Sciences Advisory Committee (BESAC), and their reports are available online. See, for example, <http://www.sc.doe.gov/bes/besac/COV%20Report.pdf>.

The ASCR COV is scheduled to take place on Tuesday, March 9 and Wednesday, March 10, 2004, ending at around 2pm on March 10, at the DOE Germantown location in Germantown Maryland. The complete list of those who have agreed to serve on the 2004 ASCR COV is provided in a separate attachment.

The COV report will be presented and discussed at the spring 2004 ASCAC meeting, April 5-6, 2004. Following acceptance of the report by ASCAC, the COV report will be presented to the Director of the Office of Science.

Travel and accommodation will be handled by the Oak Ridge Institute for Science and Education (ORISE). Details concerning logistics and expense reimbursement will be sent to COV members by ORISE well in advance of the meeting.

Linda Twenty will be the point person in the ASCR office for questions concerning legal or logistical issues. Her phone number is 301-903-4087, and her e-mail address is Linda.Twenty@science.doe.gov. I am also happy to answer general questions about the COV.

Finally, thank you very much once again for helping the DOE with this important exercise.

Sincerely, Margaret H. Wright; Chair, Advanced Scientific Computing Advisory Committee