Committee of Visitors Report

Advanced Scientific Computing Research FY13-FY15

Draft Date December 8, 2017

Programs Being Reviewed:	Applied Mathematics
	Computer Science
	Computational Partnerships called SciDAC
	ASCR basic research (open solicitation)
Fiscal Years being Reviewed:	2013 through 2015
Date of COV:	October 31st to November 1st, 2017
COV Chair:	Susan K. Gregurick, Ph.D.

Date Approved by Advisory Committee:

Committee of Visitors

Amber Boehnlein, Thomas Jefferson National Accelerator Facility

John Burns, Virginia Tech

Candace Culhane, Los Alamos National Laboratory

John Dolbow, Duke University

Susan Gregurick (Chair), National Institutes of Health

Leland Jameson, National Science Foundation

Hans G. Kaper, Georgetown University

Warren Pickett, University of California-Davis

Rajiv Ramnath, National Science Foundation

Amitabh Varshney, University of Maryland

Executive Summary

The DOE Advanced Scientific Computing Research Advisory Committee was charged by Dr. Murray to assemble a Committee of Visitors (COV) to review the management processes for the research programs in Applied Mathematics, Computer Science, and Computational Partnerships called Scientific Discovery through Advanced Computing (SciDAC) within in the Advanced Scientific Computing Research program during the fiscal years 2013-2015. In the charge, the COV was asked to consider and provide an evaluation for the following two major program elements:

1. **For both the DOE laboratory projects and the university projects, assess the efficacy and quality of the processes used to:**
    a. **solicit, review, recommend and document actions and**
    b. **monitor active projects and programs**
2. **Within the boundaries defined by DOE missions and available funding, comment on how the award process has affected:**
    a. **the breadth and depth of portfolio elements,**
    b. **the degree to which the program is anticipating and addressing emerging challenges from high performance computing and DOE missions, and**
    c. **the national and international standing of the program with regard to other computational science programs that are also focused on harnessing high performance scientific computing and utilizing massive datasets to advance science.**

In response to this charge, a ten-member COV was assembled including representatives from academic, non-profit, national laboratories and the federal government. From this group Drs. Burns, Dolbow, and Kaper focused on the Applied Mathematics program, Drs. Culhane, Ramnath, and Varshney focused on the Computer Science program, and Drs. Boehnlein, Jameson, Pickett, and Ramnath focused on the Computational Partnerships (SciDAC) program. The open research program was covered by Drs. Pickett, Boehnlein, Ramnath and Gregurick. A full list the COV membership is provided and a copy of the letter charging this committee from the chair of ASCAC are provided in Attachment 1.

Prior to the meeting, the COV was provided with a link to the DOE PAMS website. This website provided COV members with important information including background and program specific presentations, a number of prior ASCAC COV reports and DOE responses, information on the

ASCR merit review process and quick reference spreadsheets for the selected proposals available to the COV. The PAMS website also provided COV members access to a large number of selected proposals submitted to the programs under this review. The proposal documents included submitted proposals, peer reviews, and program staff recommendations.

The COV met on October 31 and November 1st, 2017 at the Rockville Hilton in Rockville MD. The meeting opened with a welcome on behalf of ASCR from Barbara Helland, Associate Director of the ASCR office. Following the welcome and introductions, Dr. Steven Lee presented an overview of the Applied Mathematics program, Dr. Lucy Nowell presented an overview of the Computer Science program and Dr. Randall Laviolette presented an overview of the Computational Partnerships program. These program specific presentations spanned a day and a half and provided information about the ASCR scientific research programs, including the organizational structure and staffing profiles, program specific funding announcements, application statistics and review details, active awards, and scientific highlights from the funded investigators.

The presentation sessions were interactive, with significant opportunity for questions from the COV members, answers from the program officers, and discussion between the COV members on process, content, and strategic vision for the programs under review. The COV met in executive session on the first day and requested follow-up information on questions related to process, program outcomes and measures, mission areas and forward planning. DOE program director Dr. Steven Lee followed up on the morning of the second day with the additional information that was requested by the COV members. The COV would like to acknowledge the ASCR staff both for their outstanding support and their willingness to provide additional materials for this COV to accomplish its task.

Following the presentations, the COV continued to meet within program specific subgroups to review the materials and draft early findings and recommendations. The COV met in executive session at the end of the second day to develop an outline for the findings and recommendations. Further communications between COV members include a googledocs website and email conversations. The PAMS website was open for COV members to continue to review materials and to assist with deliberations until November 9th. A teleconference phone call was held to finalize the report.

Because of the breath of the charge and the complexity of the programs under review, the report is structured as follows: The report first outlines summary findings and recommendations which cross multiple ASCR research programs and are sufficiently significant to raise to a higher level. Following the summary findings, each scientific program is separately discussed with program specific findings and recommendations which should be addressed by DOE ASCR staff within that program.

Advanced Scientific Computing Research Program

The COV endorses the mission statement of the ASCR Research Office: *The ASCR Research Division underpins DOE's world leadership in scientific computation by supporting research in applied mathematics, computer science, high-performance networks, computational partnerships (SciDAC) and educational programs.*

The overarching recommendation of the COV is for the creation of a strategic plan for the ASCR Research Division and its programs. Within each research area, ASCR has dedicated program staff, resources, materials, and information to support robust research in the critical areas of applied mathematics, computer science and computational partnerships. While each scientific domain has made great strides in the years covered by the COV review, and even beyond those years, the ASCR Research Office would benefit from an overall office wide strategic plan. A strategic plan will allow the ASCR office to not only plan for new scientific opportunities but also to address critical needs in management, review, and operational goals. The COV recognizes that strategic plans need to be living documents, with a finite lifetime, and that address agency priorities in a timely manner.

The COV observed that during the period covered by the review some program areas saw a higher than expected turnover of program managers. This raises a concern about long term programmatic stability and scientific direction. Fields of science benefit from DOE institutional memory as well as continuity of processes. A strategic plan could address how the ASCR office will handle recruitment, hiring, and retention of scientific staff as well as propose unifying principles or goals for processes and management.

While the ASCR office has specific programmatic research in applied mathematics, computer science and SciDAC computational partnerships, there are a few research programs, including the EXPRESS and Early Career Research Program (ECRP), that span the all of the missions areas of the ASCR office. Therefore, the COV would like to highlight these cross-cutting programs.

In 2014, the ASCR Office established a new program on *Exploratory Research for Enabling Extreme-Scale Science (EXPRESS)*. This program was intended to provide seed levels of funding for high-risk projects of two years in length. Solicitations were sent out in FY 14 and continued in FY 15. The first year of this solicitation resulted in 11 proposals being submitted to the ASCR Office, resulting in two awards, and 7 proposals were submitted in FY 15, resulting in one award. All of these proposals were evaluated using mail reviews. The committee is supportive of the EXPRESS program and would like to see this program expanded. In addition, efforts should be made to advertise EXPRESS more broadly within the community to garner the strongest possible submissions and to enhance the variety of topics.

The Early Career Research Program is critical for the development of the next generation of leaders in the fields of Applied Mathematics and Computer Science. The ASCR Office received 54 proposals in FY12, 64 proposals in FY14, and 40 proposals in FY15. The proposals were evaluated using panels, which were convened in Washington, DC. ASCR Program Managers recommended three awards in FY 12, three awards in FY14, and two awards in FY15.

In addition to these specific programs, the DOE ASCR office accepts unsolicited proposals. These proposals are initiated by "white papers" submitted by interested researchers. However, funding of unsolicited proposals is contingent on available funds within a given fiscal year. The risk to this *ad hoc* process is that exploratory research is coupled with yearly discretionary budgets and **not coupled** to strategic scientific directions of importance to the ASCR office. In recent years, the result of flat or decreasing budgets has resulted in a noticeable decrease in funding of unsolicited proposals from one year to the next.

Finally, the Committee felt that PAMS is an important new tool for helping program managers evaluate and process annual progress reports. It would be helpful for the DOE national laboratories to use PAMS to document ongoing activities and outcomes from their funded research projects.

Summary Findings and Recommendations:

- The COV recommends that new programs like EXPRESS, which emphasize short term, high-risk projects, be advertised more broadly and with increased available funds, to initiate new fields of ASCR-related science.
- The COV recommends that DOE track Early Career Research awardees as a means to build ASCR research communities.
- The COV recommends that the DOE national laboratories be encouraged to use PAMS to document ongoing activities and outcomes from funded projects.
- The COV recommends that implementation of pre-proposals into PAMS be considered as a means to document and archive information about the decision-making process.
- The COV recommends that ASCR develop a consistent process for evaluating and tracking pre-proposals.
- Given the breadth of the research portfolio, the COV recommends that program managers be given the opportunity to attend a wider range of professional meetings.

Attendance at meetings is an effective and efficient way to monitor ongoing projects and assess how sponsored work is being perceived by the research community.

- The COV encourages ASCR to develop strategies for the hiring and retention of experienced scientific program managers.
- Given the rapid and changing landscape of scientific areas that ASCR covers, the COV encourages ASCR to develop a five-year strategic plan. With a strategic plan, the advanced scientific computing research program is in a better position to maintain its vitality; without it, ASCR runs the risk of becoming reactive.

Applied Mathematics

The Applied Mathematics program in ASCR has a long history and focuses on mathematical research and software that impact the future of high-performance computing. The program supports research on vital areas important to creating and improving algorithms including numerical methods for solving ordinary and partial differential equations, computational meshing, numerical methods for solving linear and nonlinear equations, optimization, multiscale computing, multiphysics computations and mathematical software and libraries.

Efficacy and Quality of the Program's Processes

1a) Processes to solicit, review, recommend, and document proposal actions

Targeted Solicitations (UQ, Data Science, EXPRESS)

In 2013, the ASCR Office issued two targeted solicitations for basic mathematical, statistical, and computational research on "*Uncertainty Quantification Methodologies for Enabling Extreme-Scale Science*" and "*Mathematical & Statistical Methodologies for DOE Data-Centric Science at Scale.*"

Over 60 pre-applications were received in response to the former solicitation, "*Uncertainty Quantification Methodologies for Enabling Extreme-Scale Science,*" of which 22 were judged to match the objectives of the solicitation. The program managers encouraged the PIs of these 22 pre-applications to submit full proposals. In response, the DOE/ASCR Office received 22 full proposals. The program mangers assembled a panel of 18 reviewers to evaluate these proposals on the basis of the criteria described in the solicitation. Supplemental advice was solicited through one mail review. Based on the evaluations by the reviewers, the program managers recommended awards for six projects at $5M per year for three years each. The awarded projects involved teams of researchers from DOE labs and universities. The review process was fully documented, and the complete documentation entered into PAMS.

Over 100 pre-applications were received in response to the latter solicitation, "*Mathematical & Statistical Methodologies for DOE Data-Centric Science at Scale,*" of which 34 were judged to match the objectives of the solicitation. The program managers encouraged the PIs of these 34 pre-applications to submit full proposals. In response, the DOE/ASCR Office received 29 full proposals. The program managers assembled a panel of 18 reviewers to evaluate these proposals on the basis of the criteria described in the solicitation. Based on the evaluations of the reviewers, the program managers recommended awards for six projects at $3M per year for three years each. The awarded projects involved teams of researchers from DOE laboratories and universities. The review process was fully documented, and the complete documentation entered into PAMS.

The committee notes that with both the Uncertainty Quantification and the Data-Centric solicitations, a large fraction of pre-applications were judged not to address the intent of the solicitation. This may indicate that additional effort needs to be made to make the solicitations clear, particularly for investigators at universities, as to what type of research projects will be supported.

Overall, the committee was impressed with the quality and depth of the reviews received and the decisions to recommend awards on the basis of the reviews and mission needs.

Base Math Renewal Projects

The base math funding in the applied mathematics program supports core mathematics activities at the DOE national laboratories. These activities are critical for maintaining the base knowledge necessary for further innovations to meet future computational needs of the DOE Science programs. In FY14, the ASCR Office received proposals for 12 projects, which were evaluated by a panel of 11 reviewers in Washington, DC.

Open Solicitation and DOE Lab-Invited Proposal

Unsolicited proposals from investigators at universities and industry are submitted in response to the ASCR Annual Notice, which is posted on the Office of Science website. In addition, the program managers in the ASCR Office may invite proposals from investigators at DOE labs to fill out the research portfolio and address mission-critical research needs. The ASCR Office received 34 proposals in FY 13, 26 proposals in FY 14, and 25 proposals in FY15. Each of the proposals was evaluated by mail review. This process resulted in 6 awards in FY 13, 13 awards in FY 14, and 8 awards in FY 15.

The committee is supportive of this program, but would like to see additional effort placed on advertising the program within the applied math community. To facilitate the review process, additional clarity should be provided concerning the type of research the program will support.

Recommendations

The Committee finds that the applied mathematics program has continued to implement a robust set of procedures for soliciting, reviewing, and recommending proposals for funding. New solicitations have been widely disseminated to the community, and all proposals are reviewed in a timely manner by well-qualified reviewers. Reviews have been well documented within PAMS. The Committee found its initial experience with PAMS to be positive.

- The COV recommends that targeted solicitations in the applied mathematics program be advertised more broadly to the community. The solicitations should also make it clear what type of work the program will support.

1b) Processes to monitor active awards, projects, and programs

The applied mathematics program managers use a range of activities to monitor active projects and programs. These include review of annual progress reports and site visits. The applied mathematics program also has an annual PI meeting that is attended by over 200 researchers, post-doctoral fellows, and graduate students.

Program managers attend one conference each year in the field of applied mathematics. The committee felt that, given the current breadth of the program, it would be very helpful for program managers to attend a much wider range of scientific conferences. This would help to monitor ongoing projects and give program managers a better sense of how DOE sponsored research is received by the community.

Overall, the Committee felt that program managers, given their limited resources, do an excellent job monitoring active projects and programs. One suggestion that was raised during discussions of the overall program was that early career awardees should be tracked to see whether or not they are subsequently responding to new solicitations. These awardees represent new investigators in the program, and their ability to acquire follow-up support would indicate a healthy rejuvenation and a potential increase of the breadth of the program.

Recommendations

- The COV recommends that the applied mathematics program encourage early career awardees to respond to subsequent solicitations.

Effect of the Award Process on Portfolios

2a) The breadth and depth of portfolio elements

Given the boundaries defined by the DOE mission, the overall breadth and depth of the Applied Mathematics portfolio is superb. The programs on "*Uncertainty Quantification Methodologies for Enabling Extreme-Scale Science*" and "*Mathematical & Statistical Methodologies for DOE Data-Centric Science at Scale*" initiated in 2013 are designed to enable extreme scale computing mission. These programs are excellent.

The quality of the six funded projects on "*Uncertainty Quantification Methodologies for Enabling Extreme-Scale Science*" is uniformly strong and involves research ranging from the mathematical foundations to specific computational algorithms. The quality of the six funded projects on "*Mathematical & Statistical Methodologies for DOE Data-Centric Science at Scale*" are highly interdisciplinary with teams that include world class researchers.

Recommendations

- Although the depth of the applied mathematics research program is excellent and the investigators on funded projects are first-class mathematicians, the COV recommends that the program seriously consider extending the breadth of its programs by seeking to cover a broader spectrum of topics and supporting a corresponding increase of scientific and technical expertise.
- The COV recommends that the program accommodate new and emerging areas of research in applied mathematics not specifically tied to extreme-scale computing.

2b) Anticipating and addressing emerging challenges

The DOE "*Early Career Research Program*" (ECRP) supports the development of outstanding scientists early in their careers and stimulates research careers in the disciplines supported by the DOE Office of Science. This program is essential to ensure success of future ASCR research programs.

The "*Exploratory Research for Extreme-Scale Science (EXPRESS*)" program is an important mechanism to identify new areas and to broaden the existing applied mathematics portfolio. Overall, the breadth is good, given the limited budget and the focus on extreme-scale computing. The three projects funded by this program are already having an impact on the program and have resulted in new areas of research that are relevant to the DOE mission.

Recommendation

- The EXPRESS program brought new ideas and people into the DOE Applied Mathematics program. This program should be continued and expanded if possible.
- The DOE should provide sufficient travel support for the program managers to attend multiple professional conferences. This allows the program managers to remain at the leading edge of the technical developments and provides an economic mechanism for the DOE staff to talk with several existing investigators during one trip.
- Workshops and panels are very effective mechanisms to identify challenges and the emerging areas required to develop programmatic responses to support future DOE missions. These workshops should be continued and expanded.
- Applied mathematics program managers should make regular visits to facilities to identify needs of the scientific community and to anticipate future opportunities in mathematics.
- The Applied Mathematics program managers should continue and expand efforts to interact with the application offices. The work with EDER produced several successful activities and there are opportunities for additional interactions with EERE.

2c) The national and international standing of the portfolio elements

The applied mathematics program enjoys an excellent standing in the scientific research community. The ASCR office has a long and enviable record of supporting strong research programs in applied and computational mathematics, going back more than 50 years (if we include its predecessors in DOE). The COV commends the office and its program managers for this commitment to long-term basic research, especially in times of significant pressures on the budget. The applied mathematics program is recognized worldwide for its contributions to the advancement of science and engineering. It stands out as a successful model for transferring fundamental concepts and results from mathematics into algorithms and software that have been incorporated in the infrastructure for high-performance computing.

Several examples of initiatives taken by the DOE/ASCR applied mathematics program (or its predecessors) that have contributed to its reputation of excellence are worth mentioning:

- Early research in numerical linear algebra and optimization has been incorporated into commercial software products like Matlab, which are the workhorses of numerical simulations in science and engineering.
- More recently, the applied mathematics program has taken the lead in promoting research in two important research areas, namely uncertainty quantification and

algorithm development for multiscale problems. The results of these efforts are being integrated in application disciplines in the physical, life, and social sciences.

- In the timeframe under review, the applied mathematics program, anticipating the needs of the DOE users community, has undertaken significant initiatives to rethink the design of algorithms to take advantage of the exascale computing architectures of the future.
- New challenges will arise for the research program as increasing volumes of data are being generated and machine learning is taking its place as an instrument of exploration in science. The program is in the process of developing strategies to meet these new challenges.

The COV envisions an exciting future for the applied mathematics program and is confident that the program will continue to advance the state of the art of scientific computing through innovative ideas and techniques. The reputation of the ASCR applied mathematics program enhances its ability to connect with application scientists, both nationally and internationally, to engage in large research projects addressing issues of global importance. A good example is the CAMERA project (Center for Advanced Mathematics for Energy Research Applications, PI Sethian), which was set up during the reporting period to analyze challenges emerging at national and international light source facilities. This project counts about 30 facilities in the U.S. and abroad among its customers.

The national and international standing of the applied mathematics program is confirmed by several metrics, including the number of publications in peer-reviewed professional journals and conference proceedings, especially those in highly ranked international journals and proceedings of international conferences; invited presentations by investigators at meetings of professional societies; prizes awarded for scientific accomplishments; and membership of learned societies. Information provided by the ASCR office indicate that the program supports 10 National Academy of Sciences members, 44 SIAM Fellows, 2 AAAS Fellows, and several winners of prestigious prizes such as the ICIAM Lagrange Prize, the ICIAM Pioneer Prize, and the George B. Dantzig Prize.

Maintaining international leadership in any area of science requires both a long-term commitment and a sustained level of effort. The applied mathematics program office has done a commendable job in the past of maintaining a good balance of a long-term vision for mathematics research and a focus on the DOE mission. Assuming that the office will keep a full complement of program managers, the COV is confident that the program will maintain its leadership position in the applied mathematics community.

Recommendations

No specific recommendations.

Computer Science

The computer science (CS) program in ASCR supports research in utilizing computing at extreme scales and to understand extreme-scale data from both simulations and experiments. The emphasis changes from year to year, as reflected in the language of the Office of Science budget request. During FY 2013-2015, the primary focus of the CS program was on the requirements for the Exascale Computing Initiative, with a growing emphasis on data intensive science.

This program supported research across areas including:

- Application Foundations (specifically, co-design of applications, workflow tools, applications that incorporated novel applied mathematical techniques, and data analytics and visualization tools),
- Application Experiences (an area that includes resiliency, developer productivity and system performance, cyber-security and application integrity), and
- Computing Systems and Systems Engineering

Solicitations included XSTACK, Operating System and Runtime, Data Management, and Visualization, XSTACK renewals, and Resilience. The program also funded early stage research through the EXPRESS solicitation, and early-stage researchers through the Early Career Research Program. A small number of unsolicited proposals were also funded.

The SCR program in general has significant stature within the community, and this was implicitly recognized in DOE being named as one of three (DOE, Department of Defense (DOD) and the National Science Foundation (NSF)) lead agencies for the National Strategic Computing Initiative (NSCI). The program also boasts several highly awarded PIs with stature in the scientific community.

Efficacy and Quality of the Program's Processes

1a) Processes to solicit, review, recommend, and document proposal actions

Based on the presentations and interviews with program officers and management, and on examination of project folders in the Computer Science (CS) program in PAMS, the COV considered the CS program to be generally effective and well managed. For the programs that had records the solicitation and review processes appeared to be effective and well administered.

The programs appeared to be structured to attract a wide range of proposals. Thus, awards spanned a wide range of topics and levels. For example, in the XSTACK program, awards ranged from domain specific languages, new programming paradigms that used global address spaces, to performance auto-tuning. Similarly, projects in the Scientific Data Management, Analysis and Visualization at Extreme Scale program ranged from an ethnographic study to understand the abstractions scientists use to design their computational experiments (PI Deb Agarwal) to in-situ processing of adaptive mesh refinement (AMR) data (PI Maria Sematova).

Solicitations were crafted based on information gathered regarding the direction of the relevant fields gained by avenues such as workshops, program officer attendance at conferences, laboratory only summit workshops, broad DOE-wide initiatives (such as Exascale), and DOE-wide strategic plans. However, solicitations appeared to have no specific review criteria with respect to expected software quality, and software engineering processes (across the SDLC, from requirements analysis to testing).

The reviews were intended to be advisory only. Program managers had considerable (and appropriate) discretion with respect to awards made. For example, at least within the XSTACK program, program officers could choose specific aspects of proposals to be funded.

Lack of representation of under-represented demographics and early-stage investigators within review teams (panels as well as *ad hoc* reviews) appeared to be an ongoing issue. Program officers aimed to have representation of these demographics in review panels equivalent to the representation in the overall researcher population.

The COV notes that one of the awards (randomly sampled from the list of awards provided) was incomplete (as captured in the “Selection Statement”). The COV believes however that as the PAMS system matured in the subsequent years to this review, it is likely correctly capturing data and information. The COV believes that was an early phase of PAMS, which may be resolved its current implementation.

1b) Processes to monitor active awards, projects, and programs

The computer science research program managers use generally effective mechanisms, including site visits, meetings, and progress reports, to monitor ongoing awarded projects. Program managers had considerable (and appropriate) discretion with respect to how they managed awards. Thus, larger awards received more oversight than smaller awards. DOE program staff have laid out a clear set of management expectations for PIs to follow. In addition, PIs (likely new PIs) need to be educated regarding these expectations.

The ASCR program managers clearly put an impressive and significant effort into maintaining effective oversight of the current awards. The lack of availability of travel funds appears to impose constraints on program manager to attend conferences or to make programmatic site visits. In fact, the COV had concerns about the many limitations on travel for the ASCR program managers. The ability of these scientific managers to travel to conferences and to research sites hampers the DOE staff's ability to see facilities first hand and to interact with peer researchers and colleagues at valuable conferences.

Impact of the ASCR funded awards in computer science was measured by two broad metrics. Incorporation of the research into production systems was one long-term metric of impact, while shorter term metrics included publications in high-quality venues emerging from the projects. However, the artifacts of projects have been hard to capture. Attempts have been made to create a website with links to these artifacts. The program manager noted that software was typically not preserved by DOE for durations of greater than 5 years, so archival and long-term access to software not was considered an issue. The assumption of not needing software artifacts after 5 years may be something the ASCR program needs to revisit.

The COV also discussed how to take strategic advantage of the EXPRESS program, such as by using these awards to explore new scientific areas, and identify future areas of focus.

Turnover amongst program managers has been significant and is a cause for concern. The COV considered whether there are issues in ASCR that encourages mobility of program managers relative to staying with programs through 2-3 funding cycles and to enhance corporate memory. However, the COV did not have the necessary information to make specific recommendations.

Findings:

- Efficiency of the oversight process: The 2012 COV report wrote: "*The effectiveness of the program managers could be enhanced by considering mechanisms that do not rely on such frequent face to face meetings.*" The COV believes that this comment continues to hold true, in that more effective and efficient oversight mechanisms may be possible.
- Completeness of award documentation: Previous reports have noted that award documentation has become significantly more complete. Program officers are encouraged to continue this improvement.

Recommendations:

- The EXPRESS program was valuable, giving DOE program staff the ability to do what might be described as seedlings – high risk, short duration investigations to see if a topic is ripe for expansion into a full-blown program. Thus, the EXPRESS program and open calls for unsolicited proposals should be explored as strategic tools in the ASCR program toolbox.
- ASCR programs should examine a means of increasing the participation of under-represented demographics and early-stage PIs in review panels, and potentially in proposals and awards, especially within lead PIs.
- Mitigation strategies for personnel turnover should be explored. The COV encourages ASCR to develop strategies for the recruitment, hiring and retention of trained program managers.

Effect of the Award Process on Portfolios

2a) The breadth and depth of portfolio elements

The COV finds that this is an outstanding program of research, which is defining the landscape of truly big data management, visualization, resilience, and programming, related to emerging exascale architectures.

The CS portfolio has been shaped by taking into account productivity from the point of view of application developers and productivity of the system; how much science it can yield. Success is measured by whether the software is put into use and whether it is having an impact. For example, success can be defined as whether or not the developed software is run at ASCR computing facilities, the user base reached, and demonstrable measurable impact on applications that use them.

Some metrics for scalability on visualization software include statements from application scientists on the value of software. Progress is measured towards environments on which applications have to run (say compression needs to be 1000x, so how much towards that goal). With the above in mind, there is already a number of historical success stories (MPI, Adios, VTK/VisIT/ParaView) and many more being incubated in the portfolio.

X-Stack Research Portfolio spanning scalability, energy efficiency, programmability, and portability across future machines. Facilities did not require x-stack programming to be part of the platform acquisition RFP process. Even so, vendors picked up jewels of this portfolio and built it in their software. The research portfolio is targeting closer collaboration with software vendors to have sustained impact.

There were nearly a dozen PI meetings over the three years and different projects to build cross linkages; meetings were fairly frequently (6 months to a year). In general, the COV finds the PI meetings to be useful both for the community of ASCR CS researchers and for the ASCR program staff and allow both to keep abreast of fast moving research directions.

Findings:

- Excellent range of projects addressed - from HCI to systems. With the exception of the EPCR program, it appears that young PIs are under-represented in proposals

Recommendations:

- The COV recommends that ASCR continues to highly value and prioritize basic computer science research to build a foundation for the groundbreaking activities that will be require in the future.
- The ASCR program managers continue diversity consideration (such as geographic diversity, ethnic diversity, gender) to ensure balance for meetings, reviews, and funding decisions for the portfolio.

2b) Anticipating and addressing emerging challenges

In order for ASCR to keep pace with emerging challenges, a number of working sessions were conducted where program managers asked about gap analysis of the researchers' portfolio. The ASCR office also held workshops to bring scientists in a field together to discuss new directions. As an example, briefings were held about the exascale architectures on the machines that were going to be built and then the researchers were asked on how their work needed to change based on those briefings.

The emerging challenge of exascale class systems were considered by the program, including, for example, resilience.

Also admirable was the idea of using EXPRESS to solicit new and significant research directions over the horizon. Quantum information sciences is one such area that was identified in this period.

Findings:

- The CS program managers have done a commendable job in the past. There is a need to ensure strong and sustained funding to provide support for basic research.

Recommendations:

- The COV recommends the consideration of using EXPRESS systematically (if possible) to identify new areas of research
- The COV recommends a study of methods and ways to made DOE assets, in particular, software, findable, available and accessible. The DOE CODE repository (https://www.osti.gov/doecode/) is a good start.

2c) The national and international standing of the portfolio elements

The review of the FY13-FY15 CS program addressed funded programs in the areas of SSIO, X-Stack, Resilience and Scientific Data Management. The COV was told that much of this work was in preparation for getting ready for Exascale. The previous COV report on the CS program included the following recommendations:

> *"ASCR should do all that it can to ensure that it receives sufficient funding for the Exascale Initiative for the US to remain internationally competitive. The program should maintain its leadership role in high end computing by continuing to engage with the international community."*

Findings:

Since the last COV report, funding has been identified for the Exascale Computing Project and plans are moving forward with platform procurements to deliver an Exascale capability by 2021.

The COV has observed that national peers for ASCR include the ASC (Advanced Simulation and Computing) program and platform procurements, as well as NSF and their Blue Waters platform. International peers include Europe, Japan and China. In June 2013, China held the #1 position on the Top 500 list and it held that position over the rest of the time period reviewed here. Another measure of leadership is the Gordon Bell Prize (for scientific computing). This prize was awarded to US teams (running on ASC, NSF and private US industry platforms) in 2013, 2014, and 2015. In 2016 and 2017, the prize went to a Chinese team. While it remains important to monitor this prize, the future of supercomputing will surely emphasize many aspects beyond sustained peak FLOPS.

Recommendations:

- ASCR should do all that it can to ensure that it continues to invest in high quality, enabling computer science research to enable U.S. efforts to maintain leadership both nationally and internationally in areas whose impact continues to increase and support DOE missions.

Computational Partnerships

Computational partnerships bring together experts from the computer science and applied math programs with domain scientists in order to advance computationally based scientific frontiers. This partnership program (SciDAC), begun in 2001, has become a premier program for the ASCR office. Computational partnerships in ASCR encompasses SciDAC-3 Institutes competed in 2010, SciDAC-3 Partnerships competed in 2012, and Co-Design Centers competed in 2010. The SciDAC program is widely known within the computational sciences communities. The COV finds that SciDAC is highly respected nationally and is highly admired in Europe and in Asia.

Efficacy and Quality of the Program's Processes

1a) Processes to solicit, review, recommend, and document proposal actions

In its 14 years of existence at the end of this review period, SciDAC had assembled a robust and effective procedure to solicit, review, and recommend proposals. The persistence of the SciDAC program has contributed to the stability and continued optimization of the ASCR proposal process. A survey of several proposals and their reviews has supported the committee's evaluation of the proposal handling process as thorough and consistent. The practice of using a pre-proposal or Letter of Intent serves to focus and to limit the number of full proposals that need to be processed, reviewed, and decided upon. This is an important early step in the process.

The long-term service of the program managers (Randall Laviolette and Ceren Susut during this review period) has contributed to the impressive success of SciDAC. This sustained leadership is a welcome attribute of SciDAC that makes it different from other ASCR research programs, where program manager turnover has challenged the stability of planning, execution, and institutional memory. The COV advises that the extended service of program staff will be important for ASCR to successfully carry out its mission.

In order to prepare for the Exascale Computing project, a one-term partnership activity called Co-Design was awarded in 2012. Co-design refers to a computer system design process where scientific problem requirements influence architecture design and technology and constraints inform formulation and design of algorithms and software. In practice, large variations in computer architecture design in order to maximize the effectiveness of a computational platform for a given scientific task might not be an acceptable option for hardware vendors from a business point of view. Thus, there are usually tight limits that a hardware vendor can find acceptable such as adding more memory per node etc. Whatever is designed for the DOE must also be marketable to a fraction of the business community. With respect to software, in order for legacy codes, which are very important in many parts of scientific computation for the government, to be adapted to a new architecture in a timely fashion it is advantageous for the scientific programmer if each generation of hardware has some common features with previous generations. If the jump from a given generation of hardware to a new generation is too radical then the time required to adapt or rewrite legacy codes to the point where they are robust and reliable can become a significant fraction of the hardware's lifetime, i.e., if a platform has a 3 year lifetime and it takes 1.5 years for a legacy code to become robust then there is very little net benefit for the computational science. To ensure that future architectures are well-suited for DOE target applications and that major DOE scientific problems can take advantage of the emerging computer architectures, major ongoing research and development centers of computational science need to be formally engaged in the hardware, software, numerical methods, algorithms, and applications co-design process. Co-design methodology requires the combined expertise of vendors, hardware architects, system software developers, domain scientists, computer scientists, and applied mathematicians working together to make informed decisions about features and tradeoffs in the design of the hardware, software, and underlying algorithms.

Recommendations

No specific recommendations.

1b) Processes to monitor active awards, projects, and programs

Monitoring practices across the ASCR programs seem to vary noticeably. On one hand, flexibility in monitoring conforms with best practices if consistency is considered as well. Flexibility in ASCR is appropriate, considering that the computer science and applied math programs are focused in discipline, whereas computational partnerships is intentional and profoundly interdisciplinary.

Project monitoring across ASCR and computational partnerships consist of several requirements and activities. Most prominent are the annual and final reports which are collected in PAMS. These progress reports include important measures of success such as publications, software releases, patents, etc., and are listed with appropriate citations; descriptions of progress, impact, changes in directions, and immediate future plans. PAMS is viewed positively by the community and by this COV. In some cases, mid-term reviews were also conducted.

The COV notes that program managers are proactive in soliciting feedback from the investigators, a prominent example being the highlight slides that provides research advances and their impact for a science-literate but non-expert viewer. Quadcharts of the project are expected and should be updated annually by the PIs or program staff. For SciDAC partnerships, some program managers conduct substantive mid-year reviews. However, this is not uniformly the case in all partnership programs. In addition, websites are required although flexibility in presentation of project activities and progress is allowed. These activities are all representative of best practices in monitoring government sponsored research.

The COV believes that program managers should meet with investigators on a regular basis as a crucial aspect of project monitoring that will also contribute to a closer awareness of new directions as well as planting seeds for future programs. The severe restrictions on travel of program managers to conferences and for program specific site visits greatly hinder crucial avenues of scientific monitoring of awardees research, scientific advances, and emerging areas. Attendance at a conference results in (i) learning first-hand of the latest developments in the field from presentations and posters, and the conversation that can follow, and (ii) one-on-one meetings with PIs for real conversations that provide the flavor and intensity of a project that complements the more formal monitoring. Personal visits to the PIs' institutions, and meetings with group members and visits to group facilities, provides feedback and input that cannot be obtained via any other route. ASCR and DOE should continue to make this case whenever and wherever it may be heard and finally heeded.

In terms of more conventional modes of monitoring, ASCR program managers should collect quantitative data on programs. Such data could include numbers of scientific publications, and in what sort of journals (high visibility, high impact, archival) have resulted from the funded work. Similarly, information could be collected for invited talks at major conferences. The COV would like ASCR to quantify how the office evaluates success within its programs. Are highlight slides provided by PIs enough and can their impact be quantified or at least described? These sorts of self-reflective measures of success will enable the DOE ASCR office to not only be reactive to request for success metrics but to also become more proactive in monitoring its own programs and projects. This is particularly acute in a partnership program such as SciDAC where stakeholders are not only in the ASCR mission areas but also in the partner offices as well.

Lastly, considering the mission of ASCR, viz. forefront applications of HPC hardware and software, active use of digital object identifiers (DOIs) for software and data should be implemented wherever appropriate. Obtaining DOIs is relatively straightforward and available through the OSTI Data ID Service (https://www.osti.gov/home/doe-data-id-service), although not all of the necessary technologies are in place for using these effectively. Using DOIs in software and data will assist in gauging program impact, similar to how paper citation counts are currently used.

Recommendations

- Due to the importance of the SciDAC partnerships program the COV recommends that DOE determine and accumulate measures of success within this program.
- The COV recommends that DOE consider using identifiers, such as DOIs, as one method to gauge use and reuse of DOE software, data and other research products.

Effect of the Award Process on Portfolios

2a) The breadth and depth of portfolio elements

Computational science is a critical element in the portfolio for the Office of Science. Thus, by definition, the SciDAC partnerships is a broad and deep program, given the different goals that each program office brings to the partnerships. Perhaps more to the point, the partnerships have fostered the community and expertise that are developing applications as part of the Exascale Computing Project (ECP), which will give the US a tremendous advantage in scientific productivity in the 2020s.

With the SciDAC partnerships and Institutes program, as to be expected for a long-standing program, there has been a gradual evolution in the award process, incorporating experience, and changes in scientific focal areas and in the computational landscape. Over the course of the program the partnerships have changed in nature to be more collaborative (rather than transactional) with respect to the interactions between the domain experts and ASCR researchers. This is entirely appropriate.

For the computational partnerships programs, the award process and solicitation criteria was determined in collaboration with the partnering DOE offices. Continuity has been critical in forming the relationships and the understanding of the priorities and constraints from the partnering offices. The resulting flexibility was used to good effect, for example, allowing the

DOE Office of High Energy Physics (HEP) to synchronize the SciDAC term of performance with the HEP Particle Physics Project Prioritization Panel (P5) report. While the actual awards were not within the scope of this review, it was noted that prior experience had been used to guide the formation and number of the second period of SciDAC Institutes. A positive model has developed with the Institutes, involving targeted projects, funded through unsolicited proposals, that demonstrate that the Institutes are filling an important niche for the SciDAC program. One lesson from the SciDAC program is that the relationships between the researchers in the field and within DOE are a valuable resource worth fostering. However, developing the relationships has taken time and commitment.

The ASCR Co-Design awards were motivated by a need in the portfolio to look at areas where tight coupling between hardware and software design could have an influence on emerging architectures. Addressing this need with a formal solicitation is a commendable example of proactively identifying and addressing a gap in the portfolio.

One area of concern involves the seeming lack of proposals submitted by new principal investigators. In many cases, the SciDAC scientific partnerships facilitate long term programs of work undertaken by world leading teams on legacy community codes with broad applicability. While this is an essential goal of the program, long term programs led by experienced and established PIs well versed in writing proposals combined with funding constraints can create circumstances under which it is difficult for younger PIs to obtain awards and gain leadership. Similarly, the national laboratories have well-established roles in the SciDAC Institutes. This factor has resulted in surprisingly few proposals submitted to SciDAC FOAs, relative to other programs such as applied mathematics.

As has been noted in past COV reports, computer time at the Leadership Class Facilities (LCFs) and National Energy Research Scientific Computing Center (NERSC) is handled through a different process than the SciDAC awards, leading to concerns of the PIs about having adequate allocations on these infrastructures. This situation is the same for all scientists whose research is funded separately from allocation time granted on a facility instrument (such as a light source). Despite the constant nature of this concern, we are in an era of increasing demand on the computational hardware resources due to the success of the partnerships and the needs for experimental and observational data programs. The COV believes that this point remains worth considering as a factor in the overall success of the PIs.

Recommendations

- The COV recommends that SciDAC develop opportunities for new and younger PIs to participate in its programs.

2b) Anticipating and addressing emerging challenges

The COV generally finds that for the SciDAC partnerships, much of the leadership in anticipating and addressing emerging challenges is driven by the field and by other DOE offices, often through requirements gathering reviews or priority setting mechanisms such as workshops or meetings. Insights from these activities have led to some important directions and interactions, particularly in the area of observational and experimental data, and should be considered a success of the partnerships. However, it is sometimes the case that the emerging scientific needs are out of phase with SciDAC solicitations and therefore may not be competitive for a SciDAC partnership. The COV recommends that DOE consider awarding limited scope pilot projects and a process by which new initiatives can be fostered and funded. Finally, as noted above, attendance at major international conferences is an essential element in proactively managing a portfolio, especially to identify new techniques and emerging talent.

While the SciDAC partnerships are a well-established and well-regarded program, the Early Career Research Program is also fostering development in DOE relevant application areas. The COV believes that the ECRP should create an opportunity for the next evolution of SciDAC. The COV acknowledges that there is tension between fully exploiting the Exascale era with established computational community codes, in fostering computational methods for experimental and observational data, and for laying the groundwork for scientific computation post-exascale. These tensions are genuine and worthy of strategic discussion within the ASCR office and especially with the partnering program offices. An essential next step is for ASCR to clearly articulate their strategic goals for these partnerships areas in the next round of solicitations.

Recommendation

- The COV recommends that ASCR consider leveraging a program, such as EXPRESS, to fund limited scope projects to develop new research areas within the SciDAC partnerships program.
- The COV recommends that ASCR clearly articulate a strategic goal for the SciDAC partnerships

2c) The national and international standing of the portfolio elements

In the committee's experience, SciDAC is globally recognized as the first organized effort of its type among scientists who depend at least in part on computational science. The strength in bringing together multiple disciplines to work together can lead to scientific results that are far more likely to be correct than if the disciplines worked independently.

SciDAC projects usually have a domain scientist, an applied mathematician, and a computer scientist that jointly conduct research that has a computational component. As a counterpoint example in Japan, the field of applied mathematics does not exist as it does here. Certainly, there are researchers who know a great deal of applied mathematics and computational science but the numbers in Japan are much smaller. One tends to find domain scientists who have picked up computational science on the side. Hence, there tends to be much less rigor in the development of computational methods to the point that long-time simulations can have critical errors at the boundaries, discontinuities, etc. SciDAC and all programs that SciDAC evolved into have addressed this weakness by bringing together groups of researchers who could address the various issues at hand. Many Japanese scientists lamented the fact that such a program did not exist in their country. Similarly, the European Commission has been fostering cross-disciplinary programs in computing and computational areas, recognizing the power of the SciDAC model. As a way to increase the standing of this program, the COV encourages DOE to include more international reviewers for mail-in reviews.

Recommendations

- ASCR should consider including international reviewers for its programs.