Boston University College of Engineering
Department of Electrical and Computer Engineering

8 Saint Mary's Street
Boston, Massachusetts 02215
T 617-353-2811 F 617-353-7337

BOSTON UNIVERSITY

December 8, 2014

Dr. Patricia M. Dehmer
Acting Director, Office of Science
U.S. Department of Energy

Dear Dr. Dehmer,

On behalf of the Advanced Scientific Computing Advisory Committee, I am pleased to convey the report of our SciDAC Committee of Visitors.

At our November 15th ASCAC meeting in New Orleans, ASCAC unanimously accepted the committee's report, which is attached. ASCAC extends its thanks to the distinguished members of the subcommittee: Dr. Thomas Clune, Dr. Jeff Greeley, Dr. David Keyes, and Dr. Claudio Rebbi. Personally, was a delight for me to work with these individuals as chair of the SciDAC COV.

Sincerely,

Roscoe C. Giles

Dr. Roscoe C Giles
Chair, ASCAC
Boston University.

Scientific Discovery through Advanced Computing – 3 (SciDAC-3) Committee of Visitors Report

Program Being Reviewed: SciDAC-3

Fiscal Years Being Reviewed: 2011-2013

Dates of COV: ..... October 6-7, 2014

COV Chair: ..... Roscoe Giles, Boston University

Date of Approval by the Advisory Committee: November 25, 2014

Charge to the Committee of Visitors

I. Assess the efficacy and quality of the processes used during the past three years to:
   a. solicit, review, recommend, and document application and proposal actions, and
   b. monitor active awards, projects and programs.

II. 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 portfolio with regard to other computational science programs that are also focused on harnessing high performance scientific computing and utilizing massive datasets to advance science.

Committee of Visitors Membership:

Thomas Clune, NASA
Roscoe Giles, Boston University
Jeff Greeley, Purdue University
David Keyes, King Abdullah University of Science and Technology (KAUST)
Claudio Rebbi, Boston University

Executive Summary

The Committee of Visitors (COV) for the Department of Energy (DOE) Office of Science (SC) Office of Advanced Scientific Computing Research (ASCR) for the Scientific Discovery through Advanced Computing – 3 (SciDAC-3) program met at the DOE Germantown location on Monday, October 6 and Tuesday, October 7, 2014.

The COV is extremely grateful to the program officers and other ASCR staff who gave graciously of their time and knowledge to help the committee in its deliberations enabling the review process to proceed smoothly and effectively. We also want to extend our appreciation to the program managers from the other offices in SC - including BES, BER, FES, HEP, and NP - who are participants in SciDAC-3 and who provided extremely useful and candid briefings on their role in the process.

Highlights of Findings

- *SciDAC remains the gold standard nationally and internationally for fostering interaction between disciplinary scientists and HPC. The Program Managers are to be commended on continuing the excellence of the SciDAC "brand."*
- *The breadth and depth of the portfolio is appropriate.*
- *Critical to the success of SciDAC-3 is the active role played by PMs in selecting projects and evaluating proposals in light of the highly complex and multidisciplinary mission needs. We commend the program managers.*
- *The communication and interaction of program managers with the complex teams that are involved is essential. The level of interaction of the PMs with the teams is commendable.*
- *The ability of program managers to travel to project meetings and important conferences is important but is currently insufficient. Current travel support is inadequate.*
- *The fact that a PI cannot count on an INCITE allocation generates concern for SCIDAC PM's and PI's. This represents a risk if allocations do not in fact meet the need.*

Highlights of Recommendations

- *Coordination between science programs and ASCR priorities in timing decisions pertaining to future proposals should be maintained. Staggering the proposals between the Institutes and the Application Partnerships permits better informed and posed proposals for the second group, and therefore more efficient proposal evaluation.*
- *It is important that the program managers can impose the SciDAC priority filter over and above the peer reviewers, who (properly within their sphere) rank based on the traditional merits of quality and originality.*
- *ASCR should pursue synergisms between the SciDAC and Co-Design programs.*
- *In terms of demonstrating success for SciDAC collaborations, wide adoption in the field of code developed by the Institutes should be regarded as at least as meritorious as shared post-doctoral FTEs, in that it shows that the algorithmic and software technology has reached maturity.*
- *Be attentive that balance between ALCC and INCITE computing resources is tuned in light of SciDAC requirements.*

Introduction

The SciDAC-3 program continues the efforts of ASCR and the Office of Science to enable forefront mission science on the latest advanced computing systems at the Leadership Computing Facilities (LCF's) and at NERSC.

The goals, structure, participants and achievements of the SciDAC program, including SciDAC-3, are well represented on the SciDAC website and in reports – including the previous 2007 ASCAC Committee of Visitor's report. We will not attempt to reproduce this information here except insofar as it pertains to the findings and conclusions of this report. This report focuses on the execution and achievements of the SciDAC-3 program for the past 3 years: FY2011-FY2013.

The Committee of Visitors consisted of the following expert reviewers:

Thomas Clune, NASA
Roscoe Giles, Boston University (chair)
Jeff Greeley, Purdue University
David Keyes, King Abdullah University of Science and Technology (KAUST)
Claudio Rebbi, Boston University

The committee met with ASCR staff and with program officers from the BER, BES, FES, HEP, and NP divisions of the Office of Science at the DOE Germantown location on Monday, October 6 and Tuesday, October 7, 2014. In addition, the committee had full access to program related documentation including:

- Calls for participation to Labs and Universities
- Proposals – both accepted and declined
- Proposal reviews
- Summary Spreadsheets of proposal scores
- Correspondence with proposers
- Project Annual Reports and Presentations to the SciDAC PI's meetings

The COV is extremely grateful to the program officers and other ASCR staff who gave graciously of their time and knowledge to help the committee in its deliberations enabling the review process to proceed smoothly and effectively. We also want to extend our appreciation to the program managers from the other offices in SC - including BES, BER, FES, HEP, and NP - who are participants in SciDAC-3 and who provided extremely useful and candid briefings on their role in the process.

Program Elements

The SciDAC-3 program has two elements: SciDAC Institutes and Scientific Computation Application Partnerships (SAP's).

As described in the 2011 solicitation for the Institutes,

> "The mission of the SciDAC Institutes is to provide intellectual resources in applied mathematics and computer science, expertise in algorithms and methods, and scientific software tools to advance scientific discovery through modeling and simulation in areas of strategic importance to the Office of Science and the National Nuclear Security Administration (NNSA)."

The Science Application Partnerships program consists of 5 joint efforts by ASCR and an SC program office (BES, BER, FES, HEP, NP) each focused on DOE mission applications that benefit from forefront computation. In the language of the solicitations:

> "A successful Partnership will:
>
> 1. Exploit leadership class computing resources to advance scientific frontiers in an area of strategic importance to the Office of Science, and
> 2. Effectively link to the intellectual resources in applied mathematics and computer science, expertise in algorithms and methods, and scientific software tools at one, or more, SciDAC Institutes."

Our COV was charged to review all of these program elements -- involving 2 Institute solicitations and 5 Partnership solicitations. Because the Institutes and each of the Partnership Applications are tightly coupled, this report is organized as a series of responses to each part of the charge. For each part of the charge we present our comments, findings, and recommendations for the program as a whole followed by discussion of individual elements as needed.

Processes to solicit, review, recommend, and document application, proposal, and award actions

SciDAC Institutes

A challenge of the solicitation/review process is the interdependence of the institutes and the science applications. The computing and applied mathematics research/development at the Institutes needs to support the science applications (and so needs to be aware of them) while the science applications need to be aware of the new capabilities being supported by the institutes.

Experience shows that earlier incarnations of the SciDAC program also wrestled with this challenge.

SciDAC-3 planned a three-step approach that took nice advantage of close collaboration between SC application and ASCR program managers:

- ASCR formulates the Institutes solicitation informed by the application area needs discussed with area PM's.
- After the Institutes are awarded, the solicitations for each of the areas take place in the context of the known Institutes.
- Science Applications are reviewed and awarded.

The institute solicitation in Spring 2011 received 37 letters of intent and 27 proposals. The review process included panel reviews in in 3 parallel sessions – focused on the areas of Architecture-Aware Performance, Data Analytics and Visualization, and Scalable Solvers and Scalable UQ – and mail-in reviews for System Software. A peer review panel that included select reviewers from the separate panels followed. The peer review panel looked at management/organizational strengths and weaknesses of top rated proposals. Finally, input was sought by the ASCR PM's from SciDAC partners in the other offices. The final portfolio was built by the ASCR PM's in consideration of all the developed review materials.

The initial review process identified 3 fundable projects -- in the areas of scalable mathematics and solvers, uncertainty quantification, and architecture aware performance. It did NOT succeed in identifying a sufficiently highly reviewed project in the area of data/visualization. For each of the 3 fundable projects, the PM's had extensive interactions with the proposers about the scope, activities and funding levels and awards were made at the end of June, 2011.

In September 2011, ASCR issued another solicitation for a single data-focused Institute. This resulted in 7 letters of intent and, 8 submitted proposals. An award was made in December of 2011.

Thus the SciDAC-3 program created the following 4 institutes:

FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics.
Topic areas: Structured & unstructured mesh tools, linear & nonlinear solvers, eigensolvers, particle methods, time integration, and differential variational inequalities.

QUEST: Quantification of Uncertainty in Extreme Scale Computations
Topic areas: Forward uncertainty propagation, reduced stochastic representations, inverse problems, experimental design & model validation, fault tolerance.

SUPER: Institute for Sustained Performance, Energy and Resilience.

SDAV: Scalable data management, analysis & visualization.

Findings:

- *The timing of the calls for institute proposals and the interrelated partner proposals is a challenge. Asking the program managers in the science areas to define their areas of interest, followed by the institute competition with knowledge of those areas, followed by the actual science program completion, was a good process.*
- *The PM's are to be commended for having the courage to re-compete the data institute rather than accepting a suboptimal solution among the original proposals.*
- *Process Documentation has much improved since last review in 2007*

Science Application Partnerships

Each of the SAP solicitations focused on activities defined by the corresponding partner program. The partner office in close collaboration with ASCR led the review processes in each case.

There is no "one-size-fits-all" style of FOAs and review processes. Office-specific solicitations and flavors are natural. For example, in SciDAC-3, NP emphasized obtaining predictions of direct relevance to current experimental programs at the accelerators, by means of LQCD. In contrast, the BES Program Managers sought a broad portfolio of projects in the context of frontier electronic structure calculations that they had previously identified.

Communication within the SciDAC program seems excellent. Within Germantown, communication between the program managers of the science applications and the institutes is effective and demonstrates healthy interdependence for success. In particular, the SciDAC-3 applications award selection has been harmoniously coordinated.

The short turn around time between announcements and awards is remarkable.

Reviewers were very good despite numerous conflicts of interest that reduced the pool of available qualified people.

BER

The BER solicitation received pre-proposals from 11 teams of which only 6 were encouraged to submit full proposals based upon responsiveness to the SciDAC call. The three successful proposals all obtained average scores of 8 or 9 (out of 10) and were uniformly ranked higher than the declined proposals.

BES

The BES partnership solicitation received the most pre-proposals of any of the SC partners (89), and the large number necessitated a correspondingly significant down select before encouragement of full proposals. The down select, to 28 full proposals, was made largely on the basis of responsiveness to the SciDAC call, with effective collaboration with SciDAC Institutes being a primary consideration. The size and competency of the teams, as well as the consistency of the proposed scientific applications with the current BES science portfolio, were also evaluated. The 28 full proposals received were sent for external review (21 BES and 21 ASCR reviewers), and written comments were supplemented by a numerical score of 1-6 from each

reviewer. Successful proposals generally obtained average scores of five or greater, but specific comments of the reviewers, including, but not limited to, the reputation of the PI's, the direct experience of the PI's with mathematical and algorithmic analysis, and the overall balance of the portfolio, were also considered in the final decision-making process.

FES

The number of reviews per proposal in SciDAC-3 was very impressive in many cases, e.g., FES had 5-8 FES reviewers per Edge Physics proposal. This took major effort due to the smallness of the US community and required reviewers from abroad. However, this was important to the integrity of the process.

HEP

There were 7 proposals, reviewed by mail. Three were recommended for support. We noted that one proposal reviewed highly, but it appeared that the reviewers failed to see shortcomings in the proposal, giving it a rather good score; the PMs in HEP and ASCR exercised good judgment in noting the lack of relevance to HEP and ASCR goals and declined.

NP

There were 10 pre-proposals, 9 proposals reviewed by mail and 3 awards. The role of the program managers in properly identifying proposed activities that were both responsive to DOE mission needs and achievable was evident in our review.

Findings:

- *Critical to the success of SciDAC-3 is the active role played by PMs in selecting projects and evaluating proposals in light of the highly complex and multidisciplinary mission needs. We commend the program managers.*
- *The solicitation was complex and effective overall, with evident cooperation among various PMs. The communication about the nature of the data institute was not as effective – necessitating a second call.*
- *The requirements for pre-proposal and portfolio down selection vary substantially from one program to the next, but the PMs are managing those differences in a manner appropriate to their discipline.*

Recommendations

- *Preserve decision documents, even for declined proposals, and provide summary feedback in the declination letter.*
- *Coordination between science programs and ASCR priorities in timing decisions pertaining to future proposals should be maintained.*
- *It is important that the program managers can impose the SciDAC priority filter over and above the peer reviewers, who (properly within their sphere) rank based on the traditional merits quality and originality.*

Processes to monitor active awards, projects and programs

The active awards are monitored through weekly telecons, semi-annual reports, site visits, and annual SciDAC-3 PI meetings in the Washington DC area. The program managers take an active monitoring role.

For the institutes, reports focused on Institute-Awareness (interaction between institutes), Architecture-Awareness, and Application-Awareness. These interconnections are essential for success of the Institutes in supporting applications and each other.

For each of the application areas, the PM's in the area and the ASCR PM's worked cooperatively on project monitoring.

We noted that communication between the PMs and the PIs in the field is evidently much tighter in SciDAC-3 than in SciDAC-1 and SciDAC-2. Weekly or bi-weekly telecons with institute directors occur in some cases. Semi-annual reports are uniformly collected. Annual PI meetings in DC are also important and valued.

We note that partnership projects are generally on a successful path, though one (unspecified) has been relatively unresponsive. For example, one failed to produce a requested flyer for standard distribution.

Collaborations with SciDAC Institutes appear to be successful, especially with regard to the three institutes that were established prior to the partnership selections. Interactions with the data centered SDAV institute started developing later, and some needed connections have not been solidified.

Findings

- *Projects are well monitored by PM's through frequent telecons.*
- *PI meetings are an excellent mechanism for oversight.*
- *The program managers seem to be able to work together very effectively in supporting the projects. Negotiations among program managers were essential, and positive solutions were readily achieved*
- *The communication and interaction of program managers with the complex teams that are involved is essential. The level of interaction of the PMs with the teams is commendable.*
- *The ability of program managers to travel to project meetings and important conferences is important but is currently insufficient. Current travel support is inadequate.*
- *The program was adaptive to changing circumstances. For example, when one PI became ill, there was an intervention that resulted in a two-PI arrangement that worked very well.*

The breadth and depth of portfolio elements

We were impressed with the breadth and depth of the portfolio elements and find it excellent overall.

The science application partnerships seem well focused on partner mission and have achievable goals. The areas of emphasis for each of the partner programs are summarized below:

BER: earth systems modeling.

BES: chemistry & materials: first principles modeling of excited states; electron correlations in extended systems, "materials genome".

FES: edge physics, multiscale integrated modeling, device wall plasma-facing materials.

HEP: cosmic frontier simulations, lattice gauge theory research, accelerator science modeling.

NP: nuclear theory simulations relevant to experiments.

The applications projects have developed serious connections with the three institutes that were formed in the initial round of SciDAC-3. We look forward to an increasing impact of data analytics in the SciDAC program and growing interaction with the SDAV Institute.

SciDAC chooses projects with a high chance of success within five years. This is distinguished from the Co-design initiative, which is oriented towards influencing hardware for 2023. However, from the algorithms perspective (e.g., multiphysics formulation adaptivity), the two programs can have a similar appearance and even involve the same PIs. Therefore, though the risk tolerance and timescales are different, ASCR should pursue synergisms between SciDAC and Co-Design.

SciDAC has been highly successful in developing and deploying code, relevant to many applications areas, on leadership computing resources. This emphasis is wholly appropriate and is a distinguishing feature of the SciDAC program. While maintaining this focus, we also encourage SciDAC to develop and exploit science-based algorithms that increase the inherent efficiency of the physical analyses. The use of such algorithms on leadership facilities can only increase the impact of these facilities.

As a result of the success of SciDAC, there are now software library dependencies that many applications groups (expressed most strongly from the fusion community) regard as potential vulnerabilities. There is a concern that software developed under this program may not be able to be supported beyond the life of the program. This involves ensuring that software be of sufficient quality to be supported with limited or nonexistent overt funding.

Findings:

- *The Scientific Grand Challenge "town hall" reports undertaken recently by the Office of Science are being used to determine scientific priorities evidenced in the FOA calls in Fusion, Materials, and others areas.*
- *The aggressiveness of distributed memory scaling among the applications is much improved relative to previous SciDAC rounds. Now, software migration to many-core architectures is of strong interest to many of the applications groups, as well. This is accelerated by the adoption of a many-core architecture for NERSC-8 (Cori).*
- *There is a potential vulnerability in the dependence on software that may not be able to be maintained after the life of the program*

Recommendations:

- *Maintain or create an appropriately balanced emphasis on science-based algorithms and insights, mathematical/computational algorithms and HPC.*
- *ASCR should pursue synergisms between SciDAC and Co-Design*
- *In terms of demonstrating success for SciDAC collaborations, wide adoption in the field of code developed by the Institutes should be regarded as at least as meritorious as shared post-doctoral FTEs, in that it shows that the algorithmic and software technology has reached maturity.*

Anticipating and addressing emerging challenges from high performance computing and DOE missions

The SciDAC program has aimed to address leading challenges related to the DOE mission in a timely and effective manner. We have noted that this seems to be working well in the SciDAC-3 program.

Success or progress requires melding of HPC resources, mathematical algorithms and science-based algorithms. SciDAC explicitly has pulled in HPC facilities and mathematical/computing and relies on the programs for the science-based algorithm development.

We note that from the application perspective there is still a sense of "double jeopardy" by virtue of the need to apply separately for computational resources from INCITE. The INCITE program is regarded by some as a "crapshoot." The ALCC program has provided resources to help meet application needs, but is small by comparison to overall demand.

As a promising step towards ensuring the persistence of Institute codes by broad engagement with the scientific community, it is noted that at least two Institutes, QUEST and SDAV, have organized annual summer schools and/or extensive tutorial efforts.

Example Success Stories

- **BER**: There is considerable optimism among the PMs that the projects will lead to highly accurate resolution-independent techniques and significantly reduced reliance upon parameterizations.
- **BES**: The partnership projects are well on their way to producing impressive scientific and algorithmic advances, and high impact publications have already emerged. Notable examples include a recent *Science* article detailing a highly accurate electronic structure-based crystal structure for solid benzene, a new code to solve the Ginzburg-Landau equations on GPU's, a fully *ab initio* analysis of mean free paths of electrons in silicon, and a promising new, highly parallelizable strategy to solve the Schrödinger Equation with discontinuous functions. The portfolio thus appears to be well on its way to having a transformative scientific impact, and the success of the program will lay a firm foundation for a future SciDAC-4 effort.
- **FES**: Codes co-developed under SciDAC have successfully weak-scaled to the limits of several of the world's most capable supercomputers, yielding new insights that address the feasibility of ITER, particularly in the effect of turbulence in plasma containment. Fusion codes have been important "stressors" for the development of improved solvers and preconditioners.
- **HEP**: The "Dark Universe" SciDAC 3 project recently completed the world's largest high-resolution cosmological simulation on Titan at OLCF, covering more than 60 times the volume of previous runs. Galaxy distribution predictions from this run will enable dark energy investigations using current and future sky surveys. Lattice QCD calculations have achieved astounding levels of accuracy in first-principle predictions for several experimental observables, including some that may portend evidence for beyond-the-standard-model physics.
- **NP**: This project answers the question: What strongly - interacting mesons, (q-qbar, q-qbar-g, …) (q=u,d,s), including exotic mesons, does QCD predict that JLAB experiments will produce after the ~ $350M 12 GeV upgrade? A detailed spectrum of meson states has already been obtained by lattice techniques: this is just what experimentalists needs to know.

Findings:

- *The breadth and depth of the portfolio is appropriate.*

- *Summer schools and tutorial efforts by the Institutes to expand the science impact of their work are successful and important.*
- *The fact that a PI cannot count on an INCITE allocation generates concern for SCIDAC PM's and PI's. This represents a risk if allocations do not in fact meet the need.*

Recommendations:
- *Be attentive that balance between ALCC and INCITE computing resources is tuned in light of SciDAC requirements.*
- *The COV strongly encourages the Institutes to expand outreach efforts in the out years of SciDAC-3 to reach a larger scientific community.*

The national and international standing of the portfolio elements

Findings
- *SciDAC remains the gold standard for fostering interaction between disciplinary scientists and HPC. The PMs are to be commended on continuing the excellence of the SciDAC "brand."*
- *Informal conversations of the reviewers with overseas colleagues indicate that SciDAC is seen as a model program, which they wish could be replicated in their home countries.*