

**Los Alamos Computer Science Institute (LACSI)
Review Committee**

Final Report of Review conducted November 15 – 16, 2004

February 2005

Review Committee members

John Cerutti, LANL
Chris Johnson, University of Utah
Mike McCoy, LLNL
Dan Meiron, Caltech
Paul Messina, Chair
John Morrison, LANL
Burton Smith, Cray
Steve Wallach, Chiaro

Executive Summary

Our highest-level finding is that we consider LACSI to be conducting excellent high-risk, long-term CS research that is highly relevant to ASC and the Lab and that has already resulted in substantial benefits to ASC and the Lab.

Purpose and Scope of Review

This report conveys the findings and recommendations of the first external review of all the activities of the LACSI. In October 2001 a panel of experts reviewed the activities carried out under the LACSI academic contract but not the LACSI research carried out within Los Alamos National Laboratory by staff members. (The report of that review is in Appendix A.)

This Review Committee (RC) was asked to examine all of the LACSI activities and to assess the quality of the research and its relevance to the needs of the LANL weapons program, the Lab as a whole, and the NNSA Advanced Simulation and Computing (ASC) program. The Committee was also asked to comment on the balance between long-term, high-leverage research and short-term product development, on LACSI's progress towards achieving its original goals, and the effectiveness of LACSI management processes. (The charge to the committee is in Appendix B.)

The Committee met at Los Alamos on November 15 – 16 and received presentations on most of the activities of the Los Alamos Computer Science Institute. (The final agenda is in

Appendix C.) The RC was also given access to the LACSI web site and examined many LACSI documents, including the “*LACSI Impact on ASCI Projects at LANL*,” a list of publications, the yearly Statements of Work, and the Priorities and Strategies reports that are developed each year to guide the planning of LACSI activities.

Findings and recommendations

This report addresses all the questions in the charge to the Review Committee and includes comments and recommendations on a few additional topics that the RC believes are relevant.

The content of this report is organized around the questions that were posed in the charge to the Review Committee, followed by a summary.

1. Is the research supported by LACSI appropriately relevant to the short-term and foreseeable long-term needs of the LANL weapons program and the Laboratory as a whole?

All of the research that has been undertaken by LACSI is highly relevant to the Los Alamos weapons program and to the Laboratory as a whole. A representative list of needs that have been identified by LANL includes

- Fault tolerance and reliability
- Large scale parallelism
- Performance analysis leading to performance enhancement in high-leverage areas
- Tuning of applications
- Computational physics
- Component architecture research

LACSI is actively conducting high-quality research in all these areas. For example, in the critical area of component architectures, LACSI researchers are actively involved in the Marmot project.

The LACSI Directors regularly refocus the LACSI research portfolio, most recently in response to the January 2004 meeting. As is noted below, it is important to continue these nimble “course corrections” while staying the course on the long-term research directions that are LACSI primary mission.

Even the long-term research projects can and do yield short-term benefits to the ASC program and to the Laboratory, thanks to LACSI’s strategies (described below). Examples of major payoffs from long-term basic research efforts include

- Compiler-based tools technology research by Mellor-Crummey that resulted in a 4x increase in performance on SAGE, and an even bigger factor on an AMR module from BLANCA.
- Minnich’s Supermon package that dramatically reduces the effort to manage large clusters and is in use not only at LANL but also at Livermore and other outside sites, including some in intelligence agencies.

Is this program, within the entire portfolio of ASC activities, achieving the right balance between long-term, high-leverage research and short-term product development for the weapons program and LANL?

LACSI has done an excellent job of balancing long-term and short-term priorities within its research program. Given the largely short-term development agenda of most of the ASC Program, the emphasis of LACSI has been focused primarily and appropriately on longer-term research.

The ASC Program is strongly mission and deliverable driven, in all program areas: weapons applications and science, computer operations and computer science. Consequently, R&D, almost by default, is approached with a shorter term focus. For instance, Problem Solving Environment (PSE) and Data and Visualization Sciences (DVS) R&D has at most a 2-3 year window, usually much shorter. Funding levels for staff effort for each of these elements (PSE or DVS) are on the order of \$10M per Laboratory. While not all of this effort goes to R&D (some goes to services and operations), the order of magnitude is nonetheless large. This then leaves a question open, where does the longer term, riskier research get funded?

From this perspective, we suggest that longer term research should continue to be emphasized within LACSI. From the presentations, it is clear that this is and should remain the primary focus of LACSI. LACSI is not, however, insensitive to the value of providing some shorter-term results. The Institute addresses the need to balance long-term and short-term priorities within the program by establishing for each project a vision of a desirable long-term research outcome and then delivering intermediate results of that research to the laboratory. Although, as is appropriate, not every project yields near-term payoffs, this strategy has been largely successful, with examples in performance prediction, cluster management, and code optimization. Some of the clearer success stories, however, have yielded shorter term deliverables that have supported the program in timely ways. Production of near-term results has also acted to suggest that the overall project has ultimate value, removing some pressure on the researchers to show relevance. This too is valuable. In balance, given the largely short-term development agenda of the rest of ASC efforts, the emphasis of LACSI has been focused primarily and appropriately on longer-term research.

An example of the importance of long-term research is the work by Mellor-Crummey in compiler & tools technology. Interest in compiler technology for parallel programming, especially for scientific applications, has waned in the overall CS community and was never high in industry. Therefore, LACSI funding was instrumental for continuing that research, some of which started in CRPC over a decade ago. This long-term research is now paying off by providing very substantial performance increases in codes of great interest to ASC.

This LACSI research strategy relies for its success on multi-year project funding rather than annual competition. Termination of projects will doubtless occur, but

success in long-term research requires patience in the selection and review process.

2. Is the research funded by LACSI of the highest quality?

In Computer Science, it is very clear to the Review Committee that the research is of the highest quality. LACSI has chosen a set of topics (e.g. systems, fault tolerance, etc.) that are absolutely key to further progress in high performance computing and the group of faculty and Los Alamos researchers addressing these problems is top-notch.

The work on the use of telescoping languages of Ken Kennedy is particularly interesting and could go a long way towards improving the ability of applications programmers to construct rapidly high performance codes that will perform with good efficiency on massively parallel architectures, especially if integrated with other ongoing software architecture research. The Review Committee encourages LACSI to continue to engage the Common Component Architecture (CCA) researchers and both leverage off the CCA research and promote LACSI research in component architectures. The analysis that goes into this is very much associated with the strengths of Kennedy's group at Rice. Indeed, in Rice, LACSI has a valuable resource that can provide state-of-the-art direction in the use of compiler technologies for future ASC applications.

The work on fault tolerance is similarly of very high caliber. This is an absolutely central problem that must be solved if we are to advance to using tens of thousands of processors routinely. Particularly attractive is the mix of mid-term research contributions (e.g. fault tolerant MPI) coupled with long term analyses (work using SvPablo) to understand the fault tolerance problem more widely from an applications perspective.

In computational science, the work that was reviewed (development of new finite element methods for diffusion problems on polyhedral meshes) was first rate. It provides a very promising solution to a problem that has bedeviled researchers for quite some time. Professor Yuri Kuznetsov at the University of Houston has been working in conjunction with LANL technical staff in groups T-7 (Misha Shashkov), CCS-2 (Jim Morel), and X-3 (Scott Runnels) to devise, implement, and test a new set of high-fidelity discrete vector calculus operators for two- and three-dimensional polyhedral meshes. These operators must by design be "mimetic", or mimicking key properties of the actual operators. Kuznetsov has led the way to devise a new set of MFD (mimetic finite difference) and MFE (mixed finite element) operators that preserve accuracy, order of convergence (nominally 2nd order), vector identities, generality (e.g., for polyhedral meshes), without sacrificing efficiency (e.g., the resulting operator remains SPD). The research results in this area are of very high quality. X Division personnel are evaluating his approach in a key simulation tool in the weapons program.

An indication of its importance and relevance is that X division researchers are sufficiently impressed with this new approach that they are currently integrating it into their codes. While not all LACSI achievements need to have direct short term relevance to weapons program goals it is very encouraging to see this level of success. On the other hand, the research portfolio of the computational science program in LACSI is not as broad as that of the computer science component. This is by design and it is appropriate that LACSI concentrate on the core computer science issues. At the same time it is important to continue to support Kuznetsov's work if not through LACSI then through some other research opportunity.

The quality of LACSI's research is also indicated by the large number of Best Paper Awards received by the LACSI researchers at international conferences.

Is the project engaging the best minds in the nation on problems of relevance to LANL's overall goals in computer and computational science?

As indicated above, the answer here too is an emphatic "Yes." LACSI has engaged researchers that are preeminent in the field of high performance computing, especially with regard to HPC systems research. Regrettably there are not as many people going into high performance computation as a field in computer science as one might like and so a particular value of LACSI is that it can be a nexus for bringing together the nations best talent in this area. Should it prove possible to build the proposed Santa Fe Information Technology Laboratory (SFITL), then LACSI could emerge as a unique resource to the country in terms of its ability to bring together researchers in high performance computing.

We note again that there are fewer such contacts in computational science but this may be appropriate, since LACSI's budget is not very large and many of the best minds in computational science are already engaged through other programs. Again, the computational science research we reviewed was excellent.

Are the most significant high performance computing issues being addressed by LACSI?

There are many significant and challenging issues that relate to high performance computing, in particular in scaling to petaFLOP/s systems that will certainly be on the scene before the end of the decade. The recently released NRC (National Academy) study on "The Future of Supercomputing" provides a detailed analysis of these issues. We were pleased to note that LACSI is actively engaged in a number of these software problems. The 2004 LACSI *Priorities and Strategies* document lists Components, Systems, Computational Science, Application and System Performance as strategic thrust areas.

Among these efforts, areas of particular focus and high-quality work that we highlight are:

- Component architecture development to facilitate application development
- Efficiency and performance of the message passing library
- The scalability of the underlying OS (in particular) Linux Kernel

- Tuning and restructuring of important application kernels
- Providing a contemporary application and compiler environment. This environment emphasizes programmer productivity and efficiency of code generation.
- Characterization and modeling of the machine state model (i.e., the processor micro-architecture) as it relates to application performance. Predicting system performance
- Providing a Software Fault Tolerant environment for message passing and specific application libraries.
- Network performance and cluster interconnects
- Open Source software stacks for high performance clusters

All of these are highly pertinent HPC research areas. As was pointed out in the review, today the DOE/NNSA laboratories regularly deploy systems with thousands of processors. We consider these systems almost commonplace. This is due in part to the work sponsored by LACSI. As laboratories begin to deploy systems with tens of thousands of processors, the current LACSI efforts will be needed to allow scaling to these challenging environments.

3. Is LACSI meeting its original goals as laid out in the original statement of work?

LACSI has made excellent progress towards meeting the goals that were laid out in the original statement of work, especially good progress towards the third and fourth goals:

- To pursue computer science research that is relevant to the goals of High Performance Computing (HPC) programs at LANL.
- To ensure that there remains a strong focus on high-performance computing in the academic computer science community.

Regarding the third goal, as is described in more detail elsewhere in this report, all of the computer science research that has been undertaken is highly relevant to the goals of HPC programs at LANL and the entire tri-lab ASC program. LACSI projects have carried out impressive work in significant areas such as scalability of hardware and software, and in some cases the research results are already in use in production systems.

LACSI activities have contributed to motivating interest in research on high-performance computing in the academic computer science community, by holding a well-attended symposium every year, publishing widely, making presentations at major conferences, organizing special issues in journals, and publishing widely. Garnering best paper awards at conferences like Supercomputing XY also brings HPC research to the attention of computer scientists. The senior LACSI academic researchers are not only highly regarded by computer scientists worldwide, they serve on or lead key efforts and committees whose reports are influential and widely circulated. For example, Dan Reed chairs the HPC subcommittee of the PITAC and led the HECRTF effort, Ken Kennedy Chaired the PITAC for several years, and Jack Dongarra served on the National Academies panel on the Future of Supercomputing and many other HPC-related committees.

Unfortunately, as is documented in various reports [see for example “Getting Up To Speed: The Future of Supercomputing,” <http://www.sc.doe.gov/ascr/workshopreportspage.html>], there is still insufficient academic interest in high-performance computing research. However, LACSI-funded people, projects, and research results are keeping HPC research on the radar screen of computer scientists.

The goals of achieving the highest levels of prestige inside and outside the lab are very ambitious, long term, *enduring* goals, so while there has been good progress, there is still a long way to go. Those goals will take decades to reach or may not ever be fully reached and LACSI should not be faulted for not having fully met them.

Do these goals remain appropriate metrics of success for LACSI?

Appropriateness of LACSI goals

The original LACSI goals as stated in the charge to the committee are laudable, but some are more achievable than others. For example, the first goal suggests LACSI work to make the strength of computer science at Los Alamos comparable to that of physics. This will be very difficult to achieve generally, and even more so for LACSI. The second goal, to become as prestigious as the best U.S. computer science departments, is equally out of reach. We suggest these two goals be recast to seek more attainable objectives.

Although the computer science presence within the laboratory is not as strong as the LACSI goals would make it, Los Alamos is fortunate in that it has brought a very strong and prestigious group of external computer science researchers into the LACSI orbit to address problems the laboratory considers important. Other National Laboratories do not enjoy this level of interest from the computer science research community at large, and the relationship and level of interest that LACSI enjoys should be nurtured in years to come.

If and when the Santa Fe Information Technology Laboratory comes into being, the prospects for a stronger computer science presence will be greatly improved. Such an entity could become a world-class center for computer science research, with the possibility of an extremely strong visitor program very similar to the one ICASE enjoyed in computational fluid dynamics during the 1980's and 90's. There will be little chance of this kind of success without full disengagement from the rigors of NNSA security and DOE management, however.

Succinctly stated, those four original goals were:

- To build a presence in computer science research at LANL that is commensurate with the strength of the physics community at LANL.
- To achieve a level of prestige in the computer science community that is on a par with the best computer science departments in the nation.
- To pursue computer science research that is relevant to the goals of High Performance Computing (HPC) programs at LANL.
- To ensure that there remains a strong focus on high-performance computing in the academic computer science community.

4. Have the LACSI management structures and planning process been effective in ensuring the quality and relevance of LACSI activities and in supporting the original LACSI goals?

The LACSI co-Directors have put in place and used highly effective management structures and planning processes. After all, it is no accident that we have judged LACSI research to be of very high quality and relevance to its goals. Among the processes that deserve praise are

- the annual meetings that produce the Priorities and Strategy documents that guide the planning of projects and the content of the annual SOW;
- the workshops that bring together people from the weapons program and LACSI to determine the applicability of LACSI research products to Lab missions, transfer the technology as appropriate, and guide the future research efforts to ensure that the results will meet applications needs.
- The annual Symposium that brings in top HPC researchers, presents LACSI research results, and includes tutorials and workshops on HPC topics of interest to LANL and ASC HPC.

Planning is excellent and well connected to ASC and LANL needs and there is good flexibility built into the planning. Some projects have been terminated and new ones have been started, as is appropriate in a vigorous research institute. Given that LACSI is not expanding, indeed has recently shrunk, we recommend that LACSI leaders be vigilant to ensure that some of the funds are available for high-risk new starts every year.

We are also impressed by LACSI management's responsiveness to all but one of the recommendations of the October 2001 Review, including the recommendation that LACSI be managed and reviewed as one institute, instead of treating the academic and Lab portions separately. (The recommendation to conduct yearly external reviews has not yet been implemented.)

Since the October 2001 review, LACSI management has put in place appropriate vehicles for change. It has ensured that the quality of individuals and research continues to be excellent, the projects that have been selected are highly appropriate and relevant, research priorities are re-examined yearly in a process that involves many people inside and outside of LACSI, as noted previously workshops are held to promote more communication and technology transfer between LACSI researchers and LANL users, and there now is much better integration of internal and external LACSI activities

In short, LACSI enjoys good and effective management. Of course in the very difficult field of technology transfer, though some excellent progress has been made, there is always room for improvement and the LACSI co-Directors presented plans for improving the planning process and increasing the interactions between staff in the LANL weapons program and LACSI personnel.

Additional comments

Near the end of the review, we learned that starting in this fiscal year Los Alamos is phasing in a new “Weapons Science Research” mechanism for supporting long-term research for the nuclear weapons program. LACSI and its funding will be brought into the WSR process. Although the details of the implementation of the WSR have not yet been formulated, the RC has some comments regarding ways to nurture high-risk, long-term research.

LACSI is an Institute for high-risk, long-term research, not a pot of money for funding unsolicited proposals. As such, its research directions and activities must be guided in a coherent way by its top management, following a long-term vision; we urge LANL and NNSA to keep in mind. That is not to say that once a set of projects has been funded there are no new starts for many years; when compelling new ideas arise that promise to contribute to the research directions, Institute management should find ways to provide seed funding to explore them. However, that is quite different from opening the institute’s funds to competition based on proposals that are not aligned with the long-term research agenda. Piecemeal decisions about individual LACSI projects will destroy the coherence of the LACSI research agenda, especially if decisions are made by committees not familiar with the long-term vision. That this coherent, long-term approach pays off has been amply demonstrated by the benefits that LANL and ASC have already obtained from LACSI projects. When translated into dollars, just the gains in performance of the Q machine due to the LACSI systems software performance projects would fund the entire LACSI budget for a number of years.

Successful conduct of multidisciplinary, long-term, high-risk, high-leverage research requires careful management that promotes interactions while allowing investigator-initiated ideas to be pursued. Low turnover in funded personnel is appropriate because over time good researchers who work on different topics within the same institute stimulate each other and sometimes discover research projects on which they wish to collaborate, enhancing the results of both activities. This phenomenon occurs in many top research organizations and strengthens them. An example within LACSI is that Hoisie’s and Mellor-Crummey’s projects have had two workshops recently to explore merging their tools, given that they are complementary. While that merger may or may not take place, it is significant that LACSI provides a context in which such collaborations can be explored.

That is not to say that the LACSI participants should form a closed club, and indeed that has not been the case; several universities have been added to the LACSI academic contract, as well as new individuals from the participating institutions.

Specific recommendations on how to fold LACSI into the WSR process are that

- LACSI be reviewed on a yearly basis by a Review Committee;
- the LACSI Executive Committee, with guidance from the LACSI Oversight Board that is being formed, develop a proposal each year; and
- that the proposal be submitted to the WSR process as one entity.

Summary

The Review Committee wishes to thank the organizers and presenters for the excellent presentations and responsiveness to its questions.

Appendix A

Los Alamos Computer Science Institute (LACSI) Review Panel

Final Report

January 4, 2002

Appendix B

LACSI Review Charge, November 9, 2004

Appendix C

Agenda for November 15 – 16, 2004 Review of LACSI