

Los Alamos Computer Science Institute (LACSI) Review Panel
Final Report
January 4, 2002

Purpose and Scope of Review

This review examined the long-term research and development activities within the Los Alamos Computer Science Institute (LACSI) as they pertained to the contract with Rice University. Within that context, the review panel considered the quality of the partnership in the context of the best research nationwide and the relevance of the work within the context of the Los Alamos National Laboratory ASCI program. The Review Panel received a full day of presentations on the LACSI academic program on October 15, 2001 and attended most of the LACSI Symposium that followed. Additional written materials were provided to the panel and reviewed in detail, such as the FY00 and FY01 Project Report for Academic Participants.

Integral to the review process was consideration of the two criteria for the review provided in the LACSI SOW as well as the stated goals of LACSI. The criteria are:

- Quality of the research
- Relevance to Los Alamos application domains

The goals are:

- 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;

This written report begins with an executive summary, followed by specific findings and recommendations for each of the two criteria and goals listed above.

Executive Summary

The review panel unanimously agreed that the quality of the LACSI research is excellent. There is a good mix of short- and long-term focus in the research topics discussed, as well as proper alignment with HPC goals and requirements. The annual LACSI Symposia have been of high quality, with excellent internal (Los Alamos) and external speakers, papers, and attendance.

To benefit even more from the research activities, additional efforts should be made in the area of technology transfer. While the quality of the individual research and researchers is outstanding, further opportunities for communication with one another as well as with requirements within Los Alamos should be actively sought. Los Alamos has

DRAFT

not yet devised effective ways to take advantage of the LACSI academic work, and LACSI has not yet put in place processes to strengthen ties to the application requirements within the program in general. There are many contributing factors to this issue, including the “security barrier”, the lack of a coherent code development vision and strategy within Los Alamos, and the lack of a suite of representative unclassified codes. The latter was cited numerous times during the review, and we encourage Los Alamos to resume the compact applications effort (or similar effort) within the ASCI program to create a suite of such unclassified codes representative of the problem domains of interest.

To address this problem in general, LACSI and Los Alamos need to create a joint plan for making these direct connections to Los Alamos research and for transferring technology to Los Alamos projects. The ASCI Technology Prospectus (ATP) document¹ should be addressed directly by the research portfolio of LACSI and by this integrated plan.

Finally, the committee recognizes that the technology transfer issue is a difficult problem in general, representing a significant cultural change for Los Alamos. Nevertheless, we encourage Los Alamos management to stay the course and surmount the obstacles so that the benefits of this very strong and viable research program can be reaped.

Criterion 1: Quality of Research

Findings

There is no question that the quality of both the individuals involved in the research and the research itself is high. LACSI academic partner research has led to the publication of 85 papers and over 30 new algorithms and software tools. Most of the investigators serve on important external committees and are recognized leaders in the HPC community. LACSI themes include software for future platforms, as opposed to platforms of today; re-targetable applications, making it easier to migrate from computer to computer; simple, high-level programming systems -- bringing back the goal of having scientists develop programs, not professional programmers. These themes are appropriate for the current set of computational and computer science issues facing the HPC community.

The work on Compilation, Systems, and Performance Evaluation of Large Scale Parallel Machines is clearly needed, relevant, and well connected to the application requirements. HPCview has been deployed with and for a number of Los Alamos applications, and other applications (Telluride, etc.) are viable partners. The Chaitin-Briggs register allocator appears promising and should also be deployed and tested in some of the key LANL applications. The prismatic time skewing could greatly help key implicit Los Alamos applications that contain iterative algorithms with large problem size (e.g., Chad,

¹ This document, DOE/DP/ASC-ATP-001, July 2001, is a joint NNSA, Los Alamos National Laboratory, Sandia National Laboratory, and Lawrence Livermore National Laboratory publication which is periodically updated by subject matter experts at all four institutions.

DRAFT

Telluride, etc.). This is another opportunity for collaboration and exportation of technology. We note and support the fact that some of this work has been released through KAI (Kuck and Associates Inc., now owned by Intel, are one of the dominant commercial compiler companies) and is available to the community at large.

The Component Architectures work, while more speculative, may lead to changes in code development strategies within ASCI and improvements in the portability and reliability of critical software components within the HPC community. It is important, however, that that this work be conducted in close collaboration with the Common Components Architecture (CCA) Forum.

The focus of Software Tools for Parallel Computing has been to develop tools to analyze the performance of large software systems from ASCI, NSF, and SciDAC and to ensure the tools are available for future architectures. This work is particularly important as a focus on advanced architectures and future platforms at the petaOPS scale and beyond.

In the area of Enhancing Performance for Kernels on Super-Scalar RISC Processors, automated methods are being created for determining the memory architecture of a particular platform and then applying appropriate strategies to algorithms to take into account the architecture to enhance performance. The analysis of the behavior of collective operations on a particular cluster and the subsequent implementation of the most efficient collectives on that cluster is of critical importance to all HPC cluster applications, including those relevant to the ASCI program.

The Adaptive High Performance Numerical Components (e.g., UHFFT) are performing well on a number of different architectures (e.g., IA-64), hence they appear ready to deploy at Los Alamos. A clear strategy for carrying out this deployment needs to be developed. Extension of the components to support wavelets and convolutions would be in direct support of current ASCI activities (e.g., analyzing large datasets) at Los Alamos.

The work in Computational mathematics is being carried out by well-known and established researchers in this field. The researchers have begun collaborations with staff members in the Mathematical Modeling group in T Division on a number of interesting problems. However, not enough detail was given in the Parallel Numerical Methods for the Diffusion and Maxwell Equations project to discern whether this work is value-added relative to current Los Alamos technologies, and in the ASCI program in particular. If it is, many Los Alamos ASCI applications (Telluride, Antero, radiation transport) could be customers of this technology if the appropriate connections and collaborations are made.

The Eigenvalue Methods and Software work in the Computational Mathematics project is world-class. Non-ASCI connections have been made with Los Alamos (e.g., climate), but ASCI connections are still pending. Given that knowledge of eigenspectra for linear systems can greatly accelerate iterative solvers (through more intelligent preconditioners), this work needs to tie into the current Los Alamos ASCI Solvers Project (focused on linear and nonlinear solvers). Opportunities for connecting this work with the ASCI Alliances centers should also be explored

DRAFT

The ADIFOR-based project is useful, practical, and timely. However, given that most Los Alamos ASCI projects are Fortran 90 based, it is unclear how the ADIFOR work will tie in or if it can evolve to be applicable to F90. Additionally, the ADIFOR tool might be useful in approximating a Jacobian for nonlinear (e.g., Newton-Krylov) solvers; this should be pursued in concert with the LANL ASCI Solver project.

Recommendations

While the current LACSI themes are appropriate, the time may be right to re-examine the research portfolio and consider replacing some areas of existing research with activities in other areas of emerging mutual interest. Such an examination should include a careful consideration of the recently published ATP document. The questions that need to be addressed are:

- How is the mix of research topics chosen and prioritized?
- What are the unifying concepts?
- How does one research activity interact with another?
- How do the research activities interact with Los Alamos?
- What requirements are placed on the research activities from other LACSI projects or other projects at Los Alamos?

Emerging interest areas may include (but are certainly not limited to) data mining, grid computing, and so on.

Criterion 2: Relevance to Los Alamos Application Domains

Findings

The technical issues that are of importance to Los Alamos applications include large-scale parallelism, performance evaluation and enhancement, complex data structures and their mapping to massive parallel systems, and so on. The current LACSI activities in application domains, algorithms, compiler technology, and programming environments are therefore clearly relevant. From a high-level perspective, the general direction and relevance of activities is appropriate. Additionally, we note that the computer science research and development activities presented by the LACSI researchers are relevant not only to ASCI but to projects throughout Los Alamos. (The discussion of Criterion 1 included a number of specific instances of LACSI research relevance to Los Alamos projects.) Furthermore, contributing to achieving this relevance are the numerous visits among participants, over forty – six of which were for an extended period. Also, the LACSI Fellowship Program that supports computer science graduate students and provides opportunity for those students to spend time at Los Alamos has been successfully launched. A LACSI program for undergraduates at Los Alamos has not yet been initiated.

On the other hand, the relevance of some of the LACSI activities is less clear. Some of the research is more similar to what one would find in an NSF funded activity, as opposed to a program focused on application needs at Los Alamos. Although LACSI research should not necessarily be tied closely to already recognized Los Alamos needs, it would be useful to identify at least the long-term impact of the research. In some cases we believe that the relevance exists but the researchers have not yet taken the steps that would demonstrate it. For example, the adaptive FFT software might be useful for analyzing turbulent flows and some of the SANS experiments could have been run on Blue Mountain or the early Q Machine. Also, more interactions with industry, such as the technology transfer to KAI, would help.

It should be recognized that security requirements have without question hindered interactions with Los Alamos technical staff, the use of Los Alamos machines, and technology transfer to Los Alamos application domains. This issue was discussed many times throughout the day of the review. A representative example of such a discussion is worth noting. The current work in Software Tools for Parallel Grid Computing is both relevant and well connected to application requirements, and the committee applauds the direct work on the SWEEP, SAGE, RAGE, CHAD, and PARTISN applications. However, a theme that was strongly apparent here, and was recurrent throughout the day, was the severe limitation placed on progress in direct connections to applications as all such work (in this specific area) had to go through two researchers at Rice. Furthermore, the Rice research team improved the performance of SWEEP by a factor of two, yet there is no indication that Los Alamos has picked up these improvements and integrated them into relevant codes inside the laboratory. This issue was often cited as a frustration on the part of the researchers, in that research output (products, information, etc) passed back to the laboratory disappears “behind the fence” and researchers never know what, if anything, was learned from their work or what could be improved. There is a strong desire on the part of the researchers involved in this review to redress this situation. An example of an opportunity we wish could be more easily taken advantage of would be to deploy and test SvPablo/PAPI on a key Los Alamos ASCI application, whereby a determination could be made as to its efficacy for such applications.

Recommendations

Since we believe that much of the LACSI research is indeed yielding relevant technologies and/or software, we recommend that more effort be devoted to communication between the LACSI researchers and potential user groups at Los Alamos as well as with commercial vendors who are working with Los Alamos. For the projects with shorter-term objectives, it may be appropriate to institute regular briefings at Los Alamos on the progress of the research and discussions aimed at identifying its relevance to particular application domains. For these briefings to achieve their intended purpose, Los Alamos staff will have to ensure that appropriate Lab personnel participate in these discussions. This suggested process is not appropriate for all research areas, as some will look into the future and address areas of interest for application domains in general. While a formal split of projects into “direct application domains” and more “general research” is not necessarily required, each project should certainly be aware of its

DRAFT

applicability on these differing timescales and model its communications and deliverables to Los Alamos appropriately.

A number of our recommendations require actions within Los Alamos as well as by the LACSI academic partners.

LACSI should create a formal plan for outreach to the rest of the world in general and Los Alamos in particular. This plan should address the barriers to technology transfer into Los Alamos and articulate a strategy for eliminating or at least reducing these barriers. Formulation of this plan should be a joint LACSI and Los Alamos effort and should identify supportive actions that Los Alamos management could undertake. In addition, Los Alamos staff should create a set of relevant, unclassified, software kernels that are readily accessible to LACSI researchers. This software would then be used by LACSI staff to advance the transition of research technology into relevant application domains at Los Alamos. It is also worth examining the SciDAC application codes that are available outside of Los Alamos to help formulate this suite. Where appropriate, Los Alamos should work to support analysis and research on Blue Mountain and Q machines as these are the largest relevant machines available for application of the research and developments. LACSI should be proactive in helping to strengthen this tie and Los Alamos needs to build mechanisms for identifying and seizing opportunities for collaborative efforts with LACSI. This will require Los Alamos management endorsement of the use of staff time for such activities as well as support and attention to alleviate the “security barrier” problem.

This plan should also leverage the current “LACSI Los Alamos Fellows” (which are the current technical staff receiving LACSI funding and engaged in research activities) in the very important role of identifying connections between LACSI research and Los Alamos projects that could benefit from them. To enhance these interactions, Los Alamos staff involved in critical Los Alamos projects should be both invited and encouraged by their management at Los Alamos to attend and provide overview talks at the LACSI annual symposia highlighting areas where issues exist and where Los Alamos needs some technical help. This will help set research directions and goals for LACSI.

In addition, on a yearly basis, LACSI staff should meet with the NNSA and tri-laboratory technical staff responsible for the ATP and ensure that their research projects are aligned with the requirements defined in the ATP, where appropriate.

While it is admittedly difficult to define metrics to quantify the relevance of LACSI-sponsored computer and computational science research to the needs and deliverables in key Los Alamos programs (e.g. ASCI), it is nevertheless useful and necessary to do so as long as a principal goal of LACSI remains conducting research “relevant to the goals of Los Alamos”. Quantifiable “Los Alamos relevance metrics” should therefore be defined and articulated to the LACSI-sponsored investigators. These metrics should help to guide decisions about current and future project funding as well as strategic directions, program planning, etc. The metrics will not be static, but evolve according to needs and wishes as communicated by Los Alamos staff and management.

DRAFT

Suggestions for Los Alamos relevance metrics (for a particular project) range from the obvious to the not-so-obvious:

- documented Los Alamos use of a particular LACSI-developed software tool or library (e.g., who, how many, for what);
- the number of joint (Los Alamos/Academic Fellows) publications/presentations/visits (these are already well documented in the LACSI report);
- results of regularly scheduled solicitations of user opinions on the relevance of existing and proposed LACSI projects (much like the UCDRD Program); and
- size (staffing, funding) and impact of the Los Alamos program affected.

Appropriate metrics should be defined with the help of Los Alamos representatives (Los Alamos LACSI Fellows, etc.) and documented as an important factor for decisions made on all current and future LACSI projects.

A merging of the “internal” and “external” LACSI components into a single LACSI is required. For example, many of the LACSI projects could be populated from both Los Alamos and the university partners. The executive committee should have purview over all of LACSI and should address this issue. This may warrant a review of the membership of the executive committee and possible expansion to include additional Los Alamos members.

Findings and Recommendations on the Four LACSI Goals

Goal 1: To build a presence in computer science research at Los Alamos that is commensurate with the strength of the physics community at Los Alamos

Findings

It is clear that the current PI's are well qualified and connected to the computer science community and are therefore well suited to build such a presence. However, findings in this area were hindered by the fact that the scope of the review was the LACSI contract, not the entire LACSI program. Therefore, it was difficult to judge progress toward this goal. It should be noted, however, that goals like this take several years to achieve, and insufficient time has passed from the onset of this program to complete this goal .

Recommendations

DRAFT

One of the factors that would contribute to achieving this goal is assistance in recruitment. Three senior researchers, in particular Lennart Johnsson, Ken Kennedy, and Dan Reed, have served on DOE advisory and review committees. Direct involvement of these individuals and others from other academic partners in recruiting computer scientists to Los Alamos should be strengthened.

Some of the excellent ideas discussed in the LACSI SOW should be pursued more aggressively. Ones that caught our attention include:

- LACSI positions might be made more attractive by including an appointment in either the Rice Computer Science Department or the Computational and Applied Mathematics Department;
- Academic faculty at Rice can serve on search committees for LACSI, providing an academic perspective and knowledge of the computer science community drawn from years of faculty recruiting.
- Both Rice and Los Alamos will provide space for long-term visitors from other LACSI sites. In addition, LACSI will establish programs for visiting faculty (including sabbaticals), graduate students, and undergraduate students at Los Alamos.

A particularly attractive action that would contribute to building a very strong computer science presence at Los Alamos is the creation of the Santa Fe Information Technology Laboratory (SFITL). The SFITL would be created jointly with Rice University and would conduct fundamental, open, unclassified research on information technology related to high-performance computing and networking, topics of long-term interest to the Department of Energy. It would provide a unique, world-class intellectual environment focused on computer science and information technology issues critical to the future of high performance computing. One of the goals for the SFITL would be to provide an exciting environment for DOE scientists to collaborate with academic and industrial counterparts, providing additional incentive for top young computer scientists to work for the DOE labs, and for existing scientists to remain at the labs. We strongly recommend that the SFITL should be supported, as it will contribute significantly to the achievement of the goal of greatly strengthening the computer science research presence at Los Alamos. If the SFITL is created (as we hope it is), we suggest that its goals and those of the LACSI academic partners be reviewed by a group such as ours with a charge of ascertaining that the two activities are complementary and beneficial.

Finally, senior management at Los Alamos needs to engage in an active manner to help LACSI achieve this goal. Senior management should take the time to get briefings on LACSI; perhaps 30 minutes each quarter with senior management including the Laboratory Director, the Deputy for Science, the Associate Director for Weapons Physics, and the Deputy Associate Director for ASCI.

Goal 2: 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

Findings

As is the case for the first goal, it is clear that the current set of PI's already have strong reputations as individuals. Therefore, they are well suited to help LACSI gain a high level of prestige in the computer science community. As for the first goal, it will take several years to achieve this goal; gaining prestige takes time.

Recommendations

It was clear from the review and subsequent discussion that the research has generated numerous papers and presentations in the computer science community. To enhance the visibility of LACSI, it is important to ensure that all LACSI-sponsored presentations and papers advertise LACSI in some appropriate manner (such as acknowledging funding sources in papers and presentations). As has already been noted, the symposium is a successful endeavor. It would gain even more visibility by publishing the symposium proceedings in an appropriate journal or book series.

Metrics are available for ranking and evaluating various academic departments in the United States. Some of these metrics could prove valuable for the evaluation of the LACSI program – for example, number of type of publications, quality of journals in which they are published, reputation of individuals, and so on. Los Alamos and LACSI should also carefully track the number of individuals that come to Los Alamos as students or as permanent staff through LACSI, as is done with the ASCI Academic Alliances program.

Finally, since it takes a number of years to gain prestige and recognition in scientific efforts, we recommend that the contract for the academic side of LACSI be renewed for at least three years, and preferably for five years.

Goal 3: To pursue CS research that is relevant to the goals of HPC programs at Los Alamos

Findings

As previously noted, the alignment of the research portfolio with Los Alamos HPC programs is good at the present time. As we noted in the discussion of Criterion 1, while the current LACSI themes are appropriate, the time may be right to re-examine the research portfolio and consider replacing some areas of existing research with activities in other areas of emerging mutual interest. There are a number other research areas that are candidates for partnering with Los Alamos HPC activities and with other agencies. Examples from the Office of Science alone include:

- Terascale Optimal PDE Simulations
- Terascale Simulation Tools and Technology
- Algorithmic and Software Framework for Applied Partial Differential Equations

DRAFT

- Scalable Systems Software
- Performance Evaluation Research Center
- Center for Component Technology for Terascale Simulation Software
- Scientific Data Management

We recognize that only a few research areas at a time can be tackled adequately and that LACSI staff cannot be expected to have expertise in all areas of interest. We include the list above as an indication of how many opportunities there are for contributing to Los Alamos projects.

Recommendations

Both Los Alamos and LACSI need to address the issue of relevance regularly. Los Alamos needs to articulate its HPC vision and requirements more clearly and engage in discussions with LACSI participants aimed at conveying clearly the intent of the requirements. The publicly available RFP documents for the procurement of ASCI-class computational platforms are a good place to obtain some of the technical requirements. Creating a sanitized, unclassified code suite is one action that would help alleviate this problem. A one day workshop could also be held at Rice at which a subset of LACSI fellows at Los Alamos would come and give talks on the current set of problems to stimulate research interest in specific, focused, areas. This too would help alleviate this problem. LACSI leadership should review its research portfolio and consider phasing out some activities in order to address some new areas.

Finally, increased interactions with the ASCI Alliances and DOE SciDAC programs are encouraged because they present opportunities for leveraging scarce resources to address research topics of interest to all parties.

Goal 4: To ensure that there remains a strong focus on HPC in the academic CS community

Findings

Although limited access to Los Alamos supercomputers has been a hindrance, the achievement of this goal is aided by the fact that many of the members of the LACSI academic partnership are well-known national and international leaders in high performance computing. They already play leadership roles by serving on numerous advisory and program review committees; those roles give them opportunities not only to help make the case for the importance of HPC research but also to help define the national research agenda so that it is responsive to HPC needs.. As LACSI participants they are also in a position to communicate to the broader CS community both short- and long-term research needs of the ASCI program. One important communication step that LACSI has taken is holding the annual LACSI Symposium followed by the LACSI Workshop.

Recommendations

LACSI Los Alamos and Academic Fellows (i.e., the technical staff at Los Alamos and Rice or other participating institutions engaged in LACSI-sponsored research) should assume the responsibility of working with members of the ASCI tri-laboratory team that develops and updates the ASCI Advanced Technology Prospectus to prepare a presentation on future HPC technology needs of the ASCI program. They should then identify appropriate meetings, conferences, and publications for presentation of these needs to the broader computer and computational science community. Such an effort not only allows LACSI to consider these long term needs in setting their own agenda but also, as academic leaders, to communicate these needs to the broader academic community.

Los Alamos should complete the Open Collaboration Network (OCN) to make available access to Los Alamos supercomputers thereby allowing LACSI to carry out its research agenda. This would allow academic and industrial partners to experience first hand the strengths and weaknesses of current HPC technology and to consider these strengths and weaknesses in executing their own research. These partners would gain first hand experience that they could use to inform other academics.

LACSI leadership should work with other government-sponsored programs such as the NSF PACI and the DOE SciDAC to define and advertise future HPCC research needs.

The Department of Energy, the ASCI program, and the three laboratories (Los Alamos, Livermore, and Sandia) should better inform the academic community about HPC, and in particular, do a better job of socializing the ATP in a variety of venues.

Concluding Remarks and General Recommendations

We recommend that LACSI move ahead with the formation of a small external advisory committee, the Computer and Computational Science Advisory Committee, called for in its Statement of Work. If the committee were chosen wisely, it would provide additional visibility to LACSI and enhance the ability to make connections to other programs and projects. For example, individuals that chair external review committees for ASCI, in particular the PSE and Alliances program elements would be good candidates for this advisory committee as might Coordinating PI's for some of the DOE SciDAC awards. Although the review scope was the contractual portion of LACSI, it remains clear that there are still "two" LACSI's – the Los Alamos internal one and the external, Rice, one. We recommend that a yearly external review of LACSI as a whole be scheduled, similar to that which was done here for contractual purposes. We further suggest that this review immediately follow the annual symposium itself, and that the review panel be encouraged to attend the symposium. This will pave the way for a more unified approach, and should mitigate the issue of an "internal" and "external" LACSI.

DRAFT

Review Committee Members

Paul Messina, Chair

Fred Johnson

Doug Kothe

Stephen Lee

Madhav Marathe

Merrell Patrick

DRAFT