
Los Alamos Computer Science Institute Review

Welcome, Review, and Charge

Ken Kennedy
LACSI Co-Director

http://lacsirice.edu/review/slides_2006/



LACSI Review Web Site

<http://lacsirice.edu/review>

Requires user and
password:

- User name: lacsirice
- Password: review



Los Alamos Computer Science Institute Review 2004

November 15-16, 2004

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



LACSI Review Charge I

- Is the research relevant to the short-term and foreseeable long-term needs of the LANL weapons program and the Laboratory as a whole?
 - balance between long-term, high-leverage research and short-term product development for the weapons program and LANL?
- Responses (quotes from the report):
 - 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
 - LACSI is actively conducting high-quality research in these areas:
 - 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
 - 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



LACSI Review Charge Ia

- **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?**
- **Response:**
 - 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.
 - Longer term research should continue to be emphasized within LACSI
 - 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
 - **Example: HPCToolkit**



LACSI Review Charge II

- **Is the research funded by LACSI of the highest quality?**
 - Is the project engaging the best minds in the nation on problems of relevance to LANL's overall goals in computer and computational science?
- **Response:**
 - 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
 - Examples:
 - **Telescoping languages**
 - **Fault tolerance**
 - **Computational science**
 - **Number of best papers**



LACSI Review Charge III

- **Is LACSI meeting its original goals as laid out in the original statement of work?**
 - Do these goals remain appropriate metrics of success for LACSI? 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.**
- **Response**
 - Yes to the last two goals
 - **Still insufficient interest nationally, but LACSI is holding on to some of the best talent**
 - The first two may have been overly ambitious (de-emphasize?)



LACSI Review Charge IV

- **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?**
- **Response:**
 - 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
 - **Annual planning meeting**
 - **Workshops**
 - **LACSI Symposium**
 - Planning is excellent and well connected to ASC and LANL needs and there is good flexibility built into the planning.
 - **Need to continue to set aside money for new starts**



Relationship to WSR

- **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 proposal be submitted to the WSR process as one entity.
- **Actions**
 - This review is the first under the new system
 - The proposal was developed and submitted under the new WSR system
- **Mixed Results**
 - Proposal was successful
 - Budget was severely constrained (to \$2,060,000)
 - All of the LANL participants dropped out



LACSI 2006 Review Charge

- **Is the research relevant to the short-term and foreseeable long-term needs of the LANL weapons program and the Laboratory as a whole?**
 - balance between long-term, high-leverage research and short-term product development for the weapons program and LANL?
- **Is the research funded by LACSI of the highest quality?**
 - Is the project engaging the best minds in the nation on problems of relevance to LANL's overall goals in computer and computational science?
- **Is LACSI meeting its goals?**
 - Do these goals remain appropriate metrics of success for LACSI?
 - 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.
- **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?**



Los Alamos Computer Science Institute

Overview

Ken Kennedy
Rice University
LACSI Co-Director

http://lacs.rice.edu/review/slides_2006/



Prologue

- **1988: NSF Science and Technology Center for Research on Parallel Computation**
 - Charter members: Rice, LANL, CalTech, ANL
 - Later added Syracuse, Tennessee, Texas
 - Planning by Executive Committee (NAPA approved)
 - Eleven (11) year sunset
- **1999: Los Alamos Computer Science Institute**
 - Weigand vision
 - Enhance CS research at LANL
 - Keep parallel computing expertise focused on ASC problems
 - Preliminary planning (Reynders, Kennedy, et al)
 - Los Alamos, Rice, Illinois, Houston, New Mexico, Tennessee



LACSI 2006 Objectives

- To pursue computer science research that is focused upon the **long-term goals of the HPC programs** at Los Alamos
- To ensure that there remains a **strong focus on high-performance computing** in the academic computer science community.

Observation: LACSI Academic Partners complement the strengths available in CCS and CCN

Examples: performance analysis tools, source-to-source transformation systems

New Focus: Direct collaboration with weapons program



LACSI Partners

- **Los Alamos National Laboratory**
 - Andy White, Bill Feiereisen, Adolfo Hoisie, Rod Oldehoeft, John Thorp
 - Jim Ahrens, Hank Alme, Jeff Brown, Rich Graham, Dave Higdon, Chip Kent, Ken Koch, Ron Minnich, Dave Montoya, Craig Rasmussen
- **Rice University**
 - Ken Kennedy, Keith Cooper, Rob Fowler, John Mellor-Crummey, Linda Torczon
- **University of North Carolina**
 - Dan Reed
- **University of Tennessee**
 - Jack Dongarra
- **University of Houston**
 - Yuri Kuznetsov, Lennart Johnsson
- **University of New Mexico**
 - Deepak Kapur, Patrick Bridges, Barney Maccabe




Partnership Model

- **Partnership of six major institutions**
 - Partnership means multi-year commitment to the institution
 - Partner responsibility:
 - focus research based on shared vision and plans
 - Ensures attention of important senior researchers on problems of importance to ASC and LANL
- **Vision and planning by the Executive Committee**
 - Annual review of priorities
 - Plans adjusted based on need
 - Budget adjusted based on need and responsiveness
- **Research leaders expected to redirect their research program**
 - Multi-year funding critical to focusing effort
 - Example: component integration



LACSI Structure and Organization



New Annual Planning Cycle




LACSI

Research Strategy

- **Focus on Long-Term Investigations**
 - Address difficult, long-term problems
 - Select targets based on what will be important in five years
 - Plan and adjust judgments each year
- **Drive Research by Problems of Importance to ASC**
 - Identify CS challenges from ASC problems
 - CS research should not be simply support for applications
 - Study of application challenges should define the research agenda
 - Integrate ideas from multiple disciplines
 - Compilers, systems, computational math and applications
- **Long-Term Progress for Short-Term Results**
 - Harvest results ripe for strategic application areas

LACSI

Research Strategy



LACSI

Sources of Hard Problems

- **Collaborations with Application Developers**
 - Performance Workshops
 - Sequence of meetings between tool developers and project teams
 - Tools applied to real code
 - Components Workshops
 - Topic: component-based approaches to code development
 - Individual collaboration visits
 - Enables achievement of specific goals
- **Vision of Senior Researchers**
 - Informed by application collaboration and interpretation of ASC priorities
- **Strategic Planning**
 - Priorities and Strategies Review
 - Plan for all of LACSI (LANL and academic components)
 - Developed at an annual workshop

LACSI

Example: LACSI HPCToolkit

- **Goal: Effective Tools for Performance Analysis**
 - Intuitive, top-down user interface
 - Provide information crucial for analysis and tuning.
- **Platform and language (compiler) independence**
 - Emphasis on LANL ASC Platforms
 - Multiple data sources ↔ Cross Platform Comparisons
 - Extract hierarchical program structure from binaries.
 - Handle multi-module, multi-language (F77, F9x, C, C++, ...)
 - No requirement for recompilation
- **Eliminate manual labor from the analyze–tune–run cycle!**



Impact on LANL Project Teams

- **HPCToolkit Deployed on Origin**
 - SAGE improvement by 2x on one example
- **Performance Workshops**
 - Feedback: Needed on Q and other secure machines, smaller DB on large codes, binary analysis too slow
 - Improvements: Sophisticated support for Alpha/Tru64 platform, new Java browser using compact database, speedup performance analysis tool on large codes improved by a factor of 30
- **General Deployment**
 - Working with CCN-8
 - Feedback: not on Clustermatic
 - Response: obtained Clustermatic system at Rice; implementation in progress
- **Kennedy visit in December**
 - Worked with Hank Alme and Mike McKay (with help from Chip Kent) to process two major X-Division codes
 - Discovered installed binaries were not up to date (now fixed)
 - Identified needed additional functionality
 - Cumulative call graph profiling (nearly ready)



Rapid Response to Need

- **DRC (Kennedy) learned that some codes exhibit memory problems**
 - Causes of this are difficult to understand
- **Mellor-Crummey designed an enhancement to HPC Toolkit to track all memory events in a run**
 - Ideas presented during October–November 2005 visit
- **Tool now nearly ready for trial**
 - Screen shot in Mellor-Crummey talk

Moral: Flexible research infrastructure permits rapid prototyping of new tools in response to need

Caution: Not enough support for making prototypes into robust tools



Component Integration and Optimization

- **Advanced Component Integration Systems**
 - Driven by need for modular approach to software development
 - Important consideration: high performance
 - High overheads for crossing component boundaries
- **Technologies Developed with LACSI Support**
 - Telescoping Languages
 - A strategy for precompiling component collections
 - In-advance optimization to expected run-time contexts
 - New strategies for optimization of object-oriented languages
- **Changes of Direction**
 - In response to discussion at first P&S meeting
 - Focus on component integration systems (later Marmot)
 - Additional focus on Ajax programming system after December visit by Kennedy



Community Interaction

- **LACSI Symposium**
 - Annual event in Santa Fe (mid-October)
 - Venue for workshops and panels on topics of interest to the ASC research and development community
- **Focused Workshops**
 - Research Planning Meetings
 - Performance Workshops
 - Components Workshop
 - Autotuning Workshop
 - Annual LinuxBIOS Workshop
- **Fellowships**
 - LACSI Fellowships: Now DOE HP Computer Science Fellowships
- **Web Site**
 - <http://lacsi.rice.edu/> — 259 publications listed (71 since last year)



LACSI Symposium

- **6th Symposium: LACSI 2005, October 11-13, 2005**
- **Approximately 240 registered in each of the last two years (record attendance)**
- **Workshops and tutorials**
 - High Availability & Performance Computing
 - Advanced Numerical Methods for PDEs
 - Performance and Productivity of Extreme-Scale Parallel Systems
 - Models & Simulations for Large-Scale Socio-Technical Systems
 - High Performance Computing in Beam Physics & Astrophysics
 - Automatic Tuning of Whole Applications
 - Algorithm Acceleration with Reconfigurable Hardware
 - Parallel Programming with Charm++ and AMPI
 - LinuxBIOS Summit
 - Application Development Using Eclipse & the Parallel Tools Platform



What Have We Done Lately?

- **Moved into WSR**
 - Positives and negatives
- **New Focus on Direct Collaboration with Weapons Program**
 - Recognition that LACSI academic researchers often complement the capabilities of CCS and CCN
 - Example: Compiler and performance tool technology
 - Different from performance prediction
 - Many visits this year the fall
 - New features for HPC Toolkit
 - New performance challenges
 - Move to get export-restricted codes for closer study
 - RAGE, PARTISN, FLAG
 - Redirection of effort
 - Focus on Ajax programming system
 - Builds on our strengths in telescoping languages and object-oriented programming support

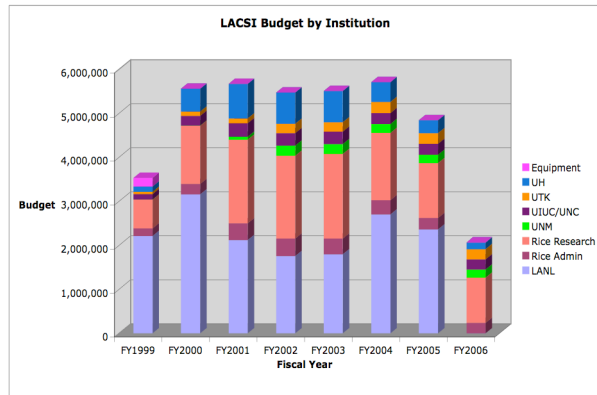


Concerns

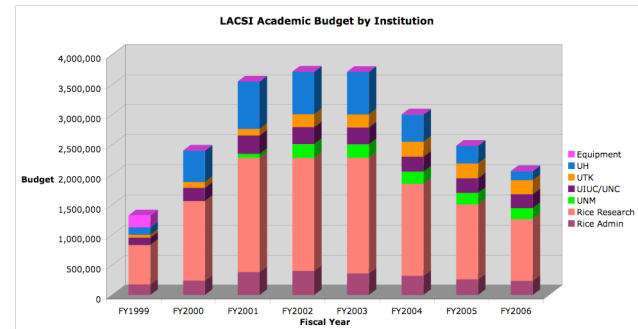
- **Budget Issues**
 - Continuing trend of budget reductions
 - No funded LANL participants
 - Unclear how WSR will treat the program this year
 - More, smaller proposals
 - Will we lose integrated planning advantages?
- **Collaboration Issues**
 - Research should be driven by real LANL applications
 - Easier when we have more Q-cleared researchers
 - We also need to streamline methods for getting access to export-restricted codes
 - Getting attention of application developers is sometimes tricky
- **Contracting difficulties**
 - Working on a no-cost extension from FY05 (soon to be funded?)



LACSI Budget History by Institution



LACSI Academic Budget History



Opportunities

- **LANL High-End Computer Science Institute**
 - Collaboration with A&M Institute
 - Large-Scale Coupled-Physics Simulations
 - Focused on graduate education
 - Could provide a home for LACSI-originated WSR projects
 - Budget up in air after LANL rebid resolved
 - Could be done more cheaply than SFITL
- **SFITL**
 - Still a great vision

Summary

- **Focus on long-term problem-solving**
 - Identify and address difficult, long-term problems
 - Evolve short-term deliverables from long-term investigations
 - LACSI has provided critical proof-of-principle funding
- **Research driven by importance to ASC and high performance computing**
 - Identify CS challenges from ASC problems
 - CS research should not be simply support for applications
 - Integrate ideas from multiple disciplines
 - Compilers, systems, computational math and applications
- **Revised planning process for WSR**
 - Formally reexamines progress and priorities every year
 - New proposal submission based on feedback and new planning
 - However, need to restore LANL participation
- **Better, closer relationships between academics and LANL staff and users**
 - Workshops, symposia, on-site visits