LOUISIANA STATE UNIVERSITY
Center for Computation & Technology

CCT Technical Report Series
CCT-TR-2010-1

Computational Science, Infrastructure and Interdisciplinary Research On University Campuses: Experiences and Lessons from the Center for Computation & Technology

Gabrielle Allen, Department of Computer Science and Center for Computation & Technology, Louisiana State University, Baton Rouge, LA USA

Daniel S. Katz, Computation Institute, University of Chicago

Posted January 2010.

cct.lsu.edu/CCT-TR/CCT-TR-2010-1

COMPUTATIONAL SCIENCE, INFRASTRUCTURE AND INTERDISCIPLINARY RESEARCH ON UNIVERSITY CAMPUSES: EXPERIENCES AND LESSONS FROM THE CENTER FOR COMPUTATION & TECHNOLOGY

GABRIELLE ALLEN[1,2] AND DANIEL S. KATZ[3,1,4]

Abstract. Computational science, interdisciplinary research and the use of information technology are transforming education, research, and institutional culture in academic institutions. In the last decade, Louisiana has targeted significant resources to these areas, one outcome of which has been the establishment of the Center for Computation & Technology (CCT) at Louisiana State University (LSU). This white paper, a contribution to the NSF Workshop on HPC Center Sustainability, aims to catalyze a broader investigation into the motivation for interdisciplinary, computational science research centers, and best practices for establishing, running, and evaluating such centers in academic environments. This paper is the work of Gabrielle Allen (CCT Assistant Director, 2003 to 2008) and Daniel S. Katz (CCT Director of Cyberinfrastructure Development, 2006 to 2009), and does not reflect the views or opinions of the CCT or LSU.

1. Introduction

In recent years, numerous distinguished national panels (e.g. [9]) have critically examined modern developments in research and education and reached a similar conclusion: computational science, as the third pillar of research, standing equally alongside theory and experiment, will radically transform all areas of education, scholarly inquiry, industrial practice, as well as local and world economies. The panels also similarly concluded that to facilitate this transformation, profound changes must be made throughout government, academia, and industry. The remarks made in the 2005 Presidential Information Technology Advisory Committee (PITAC) report are still relevant [13]: "*Universities...have not effectively recognized the strategic significance of computational science in either their organizational structures or their research and educational planning.*" Initiatives at universities to address these issues and provide the coordinated infrastructure and academic culture to promote computational research are expensive and require sustained effort over years.

Computational initiatives associated with universities have taken various forms: supercomputing centers that provide national, statewide, or local computing facilities and encourage research involving computation; faculty hiring initiatives focused on initiating research programs to change the university's expertise and culture; establishment of academic research centers on campuses that include formal involvement of faculty, for example, through joint positions with departments; and multi-university or other partnerships where the university is represented by a single entity.

We believe that any academic institution wishing to advance computational science needs to first examine its status in three areas: cyberinfrastructure facilities, support for interdisciplinary research, and computational culture and expertise (Figure 1). Cyberinfrastructure facilities refers to the compute, storage, network, and visualization resources (both local and national) to which researchers have access; to the technical and professional support for these services; and to the connection of these services to desktop machines or experimental instruments in an end-to-end manner.

[1] Center for Computation & Technology, Louisiana State University, Baton Rouge, LA 70803, USA.
[2] Department of Computer Science, Louisiana State University, Baton Rouge, LA 70803, USA.
[3] Computation Institute, University of Chicago & Argonne National Laboratory, Chicago, IL 60637, USA.
[4] Department of Electrical and Computer Engineering, Louisiana State University, Baton Rouge, LA 70803, USA.

Support for interdisciplinary research refers to the university's policies on joint appointments between units and associated promotion and tenure, policies and practices for university-wide curricula, and the academic appreciation of computational science that could rate, for example, software or data development in a similar manner to publications and citations. Finally, computational culture and expertise relates to the existence and prominence of faculty across a campus who develop or use computation as part of their research, and the provision of undergraduate and graduate courses that will train and educate students to work on research projects in the computational sciences.

Once the status of these areas has been reviewed, there are additional questions in designing a computational initiative. Should the cyberinfrastructure resources be state-of-the-art to enable leading edge research in computational science? Should faculty expertise in computational science be pervasive across all departments or resources concentrated in a few departments? Will the university administration continue to back a long term agenda in computational science and have the sustained desire to implement policies for changing culture? What is the reasonable timescale for change?

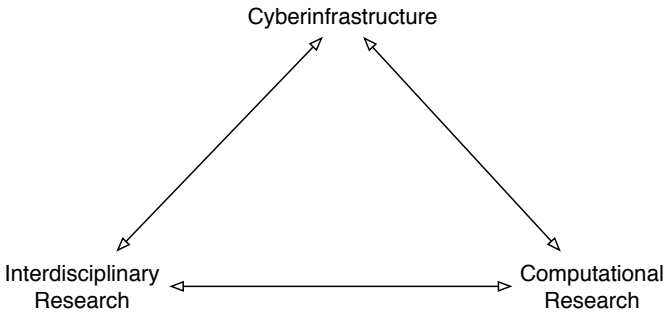

FIGURE 1. Advancing a comprehensive computational science program requires coordinated initiatives in developing and supporting interdisciplinary research, enabling cyberinfrastructure, and underlying research and culture in computation.

The required change in culture and practices for interdisciplinary research should not be taken lightly and needs the full support of university administration. Although there is some literature on issues relating to general interdisciplinary research (e.g., a National Academy review [7]), there is little written on the underlying visions, strategies, issues, practical implementations and best practices for computational initiatives. Further, what exists was usually written for a specific purpose, such as justifying an initiative for a state legislature, funding agency, or campus administration.

This position paper (which is a shortened version of a longer, more detailed report) summarizes the history and strategies of the Center for Computation & Technology (CCT) [6], established in 2003 at Louisiana State University (LSU). The CCT was created with a broad mission, to further research in the computational sciences, to enhance the education of undergraduate and graduate students in relevant fields, and to contribute to economic development in Louisiana. At LSU, all three of the areas identified in Figure 1 needed improvement, and the CCT focused on all three. We believe that our experiences at CCT over the last seven years are very relevant to the question of addressing sustainability of HPC centers, and further, that the CCT's emphasis on computational science introduces important issues that need to be addressed, including what the motivations and aims of campus HPC centers actually are and how they should be measured and assessed.

2. Establishment of CCT

In April 2001, Louisiana Governor Foster asked the state Legislature to fund an Information Technology Initiative as a commitment to the 20-year Vision 2020 plan adopted in 2000 to grow and diversify the state's economy. The legislature authorized a permanent $25 million a year commitment, divided among the state's five designated research institutions. LSU created the Center for Applied Information Technology and Learning (LSU CAPITAL), targeting funds in education, research, and economic development, with the intent that this investment would result in the creation of new businesses, increased graduates in IT areas, and increased patents and licenses.

After soliciting input from LSU faculty, an external advisory panel, the Governor's staff, the commissioner of higher education, and the president of the LSU System, the campus planned for an interdisciplinary search committee to carry out a multi-year process to hire a cluster of faculty into a small number of defined focus areas (core information technology, business, biocomputing, and materials science). Also, a search was conducted for a permanent director for the initiative. In January 2002, LSU began the process of recruiting Edward Seidel from the Max Planck Institute for Gravitational Physics (AEI) in Germany. Seidel worked with LSU to formulate a vision and detailed plan [12] to structure LSU CAPITAL into a research center related to computation and informational technology, with a physical presence on the campus and a broad mission for interdisciplinary research at LSU and across the state. In August 2003, Seidel took up the position of director of LSU CAPITAL with an associated endowed chair in the Departments of Physics and Computer Science.

The new plan detailed the build-out of focus areas in core IT, business, biology, and materials as already identified, along with funds allocated to the Laboratory for Creative Arts and Technology (LCAT). The plan envisioned LSU CAPITAL funds being used by focus areas as leverage towards large federal grants, and for postdoc, programmer, graduate student and visitor support. The plan included faculty positions split 50/50 with departments to ensure strong interdisciplinary collaborations, with reduced teaching loads for those faculty leading new programs or community initiatives. In October 2003, the center was officially renamed the LSU Center for Computation & Technology, or CCT.

3. Vision and Strategy

LSU was lacking in all the three areas identified in Figure 1: cyberinfrastructure; support for interdisciplinary research and education; and computational research, which necessitated a three-pronged approach for the center's strategy. Strategic planning has been an important activity at CCT, through yearly retreats and other mechanisms. The results of this planning are summarized in the 2006—2010 CCT Strategic Plan [4] and the CCT Faculty Plan [3].

To address LSU's cyberinfrastructure needs, CCT worked to develop campus and regional networks, connect to the national high-speed backbone, and build sustainable computational resources on the campus (Section 4.2). A major strategy towards funding these activities and promoting Louisiana as a hotbed of computation was to compete for a national center in Louisiana. A negative side effect of including a focus on the provision of cyberinfrastructure resources is that some people tend to label the center as (just) an HPC resource provider, rather than a research center, and this has proved to be an issue with how the center is represented and seen by the LSU administration.

To address the computational research expertise on the campus, CCT's core strategy has been to recruit faculty to LSU who have the skills to lead and take part in world-class interdisciplinary research groups related to computation. The fundamental group (in the CCT Core Computing Sciences Focus Area), centered around the Departments of Computer Science, Electrical and Computer Engineering, and Mathematics, was to have the necessary skills needed to build and sustain any program in computational science, including computational mathematics, scientific visualization, software toolkits, etc. Application groups were to be built following the recommendations of the external advisory board, and to leverage strength on campus, hiring possibilities, and new opportunities.

The final component of the triangle, interdisciplinary research support, was to be addressed more slowly, requiring engagement of senior administrators and a body of CCT faculty able to drive curricula change on the campus.

4. Components

4.1. **Organization.** The organization of CCT has changed over time to adapt to the needs of the growing center and it's evolving strategic plans. The CCT director reports directly to the LSU vice chancellor of research and economic development. The involvement of the vice chancellor of

research and economic development and the provost in the strategic planning and operation of the center has changed as the holders of these positions have changed.

CCT research activities are organized into focus areas, which span broad areas of research activities to facilitate cross cutting interactions and target larger problems that could be addressed by individual faculty members. Each focus area has a faculty or senior research staff lead, responsible for working with departments to coordinate hiring of faculty, building interdisciplinary research programs, administration, and hiring and coordination of staff within that group.

The broad mission of the center and large numbers of activities and programs require a number of administrative staff positions that report to the director of operations, including an events office, visitor program, public relations, technical support, business office, human resources, project management and administrative support. The center organizes various programs for supporting research at CCT, originally organized by the Assistant Director for Computing Applications and now by the Associate Director for Academic Affairs. These include graduate, undergraduate, and postdoc programs, the CCT colloquia series, and the CCT Technical Report series [5].

4.2. **Computational Infrastructure.** One of the first issues facing the new CCT in Fall 2003 was improving the network and compute resources in Louisiana to better support computational science at a nationally competitive level. A late-2003 white paper that outlined the need for, and a potential plan for, a state-wide high-speed network that would connect state research institutions with multiple 10 Gbps optical lambdas [1] led to a meeting of campus administrators in December, then to a presentation to the Board of Regents in early 2004, where the LSU Chancellor helped make the case for the sites to work together for a state-wide initiative in this area. Within a few days, Governor Blanco mentioned LONI as a priority in her State of the State address. At this time, National LambdaRail (NLR) was emerging as a high-speed optical national backbone without a plan to connect to Louisiana.

At a September 2004 workshop at the Board of Regents, national leaders, industrial representatives, federal program officers, and state researchers discussed how such a network could advance research, education, and industry in the state. Governor Blanco unexpectedly attended the meeting and committed $40 million over a 10-year period to fund the Louisiana Optical Network Initiative: lighting dark fiber already in place across the state, purchasing and deploying initial compute resources at five sites, and supporting technicians and staff. The state also funded a membership in NLR to connect the state to computational power available throughout the nation and world.

When the CCT was formed, LSU had recently deployed what were then significant compute resources: 128-node and 512-node dual-processor clusters. These machines were managed by staff from the physics department, not by the university Information Technology Services (ITS) unit [10], which provided its own limited HPC resource: a 46-node IBM Power2/Power3 machine. LSU's staff did not have experience in using or running advanced scientific computing environments. These were the only major compute resources in Louisiana, and there were also very few computationally oriented faculty in the state, which hindered research in computational science, state collaborations, and LSU's involvement in national or international projects involving computation.

CCT initially envisioned building a campus and national center for advancing computational sciences across all disciplines, with these groups' research activities integrated as closely as possible with the research computing environment. In this way, the services provided by the computing environment to the campus and nation would be the best possible, and the research output of the faculty, students, and staff would be advanced. CCT faculty would be able to lead nationally visible research activities, being able to carry out research program that would not be otherwise possible, providing exemplars to the campus, catalyzing activity in computational science approaches to basic sciences, engineering, humanities, business, etc. This was a key component of the CCT vision, one that has been successful at other centers (e.g. NCSA, SDSC, AEI) around the world.

During the creation of CCT, it was immediately apparent that existing staff had neither the experience nor the breadth to run an advanced scientific computing environment, and intense recruitment was undertaken. Experienced HPC, networking, visualization, or data archiving staff

are in high demand worldwide; bringing them to Louisiana proved very difficult. As recruitments proceeded, CCT began to work more closely with the campus computing environment, hiring its Director of HPC from within LSU ITS in the spring of 2004.

Brian Voss was recruited from Indiana University as LSU CIO in 2005, and CCT worked with the new LSU ITS to create the HPC@LSU group [8], funded 50-50 by CCT and ITS, and organized as part of CCT. This led to an immediate tension, as CCT's goals for computing systems was for them to be both development environments for new software, tools and applications that push computational science forward, as well as production systems for existing computational science applications, but the overall university user community and ITS were supportive primarily of just the latter goal. To try to better satisfy the user community, the HPC@LSU partnership was moved into the ITS organization in September 2005, still jointly funded by CCT and ITS.

Around this time, IBM Power5 systems were being introduced as LONI computing systems, and the Louisiana Board of Regents, recognizing that HPC@LSU was best able to administer these resources, contracted with LSU to run these systems around the state. This led to growth in the HPC@LSU group, adding new responsibilities such as user training, outreach, and the need to allocate a large set of systems across both state and local users. (Allocation processes are now in place for both the LSU resources and the LONI resources, with peer reviewing of proposals, to both transparently assign resources and to provide a path for users to move up to the competitive procedures for national systems.) HPC@LSU is also funded by NSF (through LONI) for several staff members to run Queen Bee (the largest LONI system) as part of the TeraGrid. HPC@LSU comprises three groups, HPC Operations, HPC Systems and User Services.

4.3. **Cyberinfrastructure.** In 2007, Cyberinfrastructure Development (CyD) was introduced as a CCT division, to better integrate research and HPC activities at CCT with the campus and national initiatives. Daniel S. Katz was hired to lead this effort. CyD's mission is to design, develop, and prototype cyberinfrastructure systems and software for current and future users of LSU's supercomputing systems, partnering where possible with the research groups at CCT to help professionalize prototype systems and support and expand their user base. A core goal of CyD has been to reach out to the scientific communities at LSU and across Louisiana, expanding existing partnerships and creating new ones in an effort to develop new computational tools to support research statewide.

CyD grew to six computational scientists (five with PhDs) and one unfilled position; a training/education group responsible for coordinating training in CCT products and HPC systems across the campus and the state with one staff member; and a visualization group, primarily providing outreach on campus with one PhD staff member and a few students. The CyD computational scientists are expected to cover 30-50% of their time through proposals led by scientists elsewhere at LSU or LONI, and spend the rest of their time on computational science activities that lead to new funding or projects and internal support of HPC and LONI activities. In 2009, Katz and two CyD computational scientists left LSU. Overall, the future of CyD within CCT is unclear, due to different visions for what the center should be doing, how much work should be faculty led from within the center vs. outside the center, and what the optimal organization for the center is.

4.4. **Contribution to National Cyberinfrastructure.** One of the CCT strategies for having LSU make a national impact in computational science was to catalyze, encourage and support the development of enabling software: to foster the use of locally developed cyberinfrastructure (e.g. the Cactus Framework, PetaShare data management, SAGA grid application development toolkit) to enable domain specific research on the campus and in the state, so that the software would be developed to satisfy local needs and ensuing federal grants would have a broad local impact. Additionally, CCT has contributed to the regional and national cyberinfrastructure through four projects: TeraGrid [15], Open Science Grid [11], Blue Waters [2], and SURAgrid [14]. CCT played a role in the management of TeraGrid and SURAgrid, with Katz as the TeraGrid GIG Director of Science and as a member of the SURAgrid Executive Committee until he left LSU.

4.5. **Research and Education.**
Faculty: CCT generally hires faculty through joint 50-50 positions with departments, with tenure residing in the departments. This model has been discussed at length and has continuously been seen as the best model for strengthening departments in computational science, and encouraging real buy-in the overall initiative from the departments. The faculty holding positions with 50 percent or more funding from CCT form the voting members of the CCT faculty. CCT also implements other strategies for associating faculty with the center, both for encouraging and supporting the participation of faculty already on the campus to take an active role in the center's programs and research, and for helping to attract and recruit faculty whose research interests overlap with CCT. For example, associate faculty positions at CCT are short term (initially three years and then renewable), where some portion of salary or research support is contributed by CCT. Joint faculty positions, whether through the 50-50 model or as associate faculty, bring up multiple problematic issues related to recruitment, promotion and tenure, division of overheads, credit for funding, etc.

Focus Areas: The research activities of the CCT are organized into five broad focus areas: Core Computing Sciences, Coast to Cosmos, Material World, Cultural Computing, and the recently created System Science and Engineering. Each focus area has a faculty lead responsible for building cross-cutting interdisciplinary research programs, administration, coordinating the hiring of new faculty and staff, and organizing their unit. The different focus areas were coordinated by the Assistant Director for Computing Applications between 2003 and 2008, and since then, Focus Area leads report to the Director. The Focus Areas have small discretionary budgets for travel, visitors, equipment and seed money for new projects. As the center's budget has become more constrained these discretionary funds have been reduced. These funds were envisioned to be used strategically by the focus area leads and were not targeted at replacing grant funds.

The interdisciplinary research of the center is driven by activities in strategically motivated, large scale projects in the focus areas, faculty research groups and the Cyberinfrastructure Development division. Several focus areas are associated with large, often statewide, funded projects: Core Computing Sciences, CyberTools and the Computational Science MHI; Coast to Cosmos, the CFD IGERT; Material World, the LONI Institute and Louisiana Biomedical Research Network; and Cultural Computing, the AVATAR MHI. These projects provide support (students, postdocs, direction) to the Focus Areas as well as broad outreach for education and training across the state.

Research staff are an essential component of the center, making it possible to quickly bring in expertise in a particular computational area as a catalyst and tool for faculty recruitment, to form a bridge from center activities to the LSU campus, to provide consistent support to strategically important areas, and to facilitate production level software development. Originally hired with state funds, these researchers were envisaged to be later funded at least 50% from grants.

Education: CCT's education goal has been to cultivate the next generation of leaders in Louisiana's knowledge-based economy, creating a highly skilled, diverse workforce. To reach this goal, objectives were set to assist in developing curricula and educational opportunities related to computation, to help hire faculty who would support an integrated effort to incorporate computation into the curricula, to offer programs that support activity in scientific computing, to attract and retain competitive students, and to advance opportunities for women and minorities in the STEM disciplines.

4.6. **Economic Development.** Particularly from the state legislature's point of view, CCT should catalyze and support new economic development in the state. In fact, the initial metrics for success provided for LSU CAPITAL included the number of resulting new businesses and patents. Economic development needs to be carefully planned and is a long term initiative, where success can be hard to measure, particularly in the short term. At CCT, one of the first hires was for a business manager to instigate inroads in this area, though CCT hasn't had good success yet in this area, other than LSU's relationship with Electronic Arts Inc. (EA), one of the world's leading independent video game developers and publishers. In September 2008, EA announced that they would place their North American quality assurance and testing center in LSU's South Campus complex. This center planned to create 20 full-time jobs and 600 half-time jobs, most of which would be occupied by

LSU students, with an annual payroll of $5.7 million throughout the next two years. EA noted that strong education and research efforts at LSU, including the AVATAR initiative and Baton Rouge Area Digital Industries Consortium along with other CCT research areas, were a strong factor in the company's decision to locate this center in Louisiana.

5. Recent Developments and Concluding Thoughts

A major change occurred at CCT in 2008 when Seidel was recruited as the Director of the Office of Cyberinfrastructure at the National Science Foundation, a role which he assumed in September 2008. In response to this, LSU appointed an interim director (Stephen Beck) and co-director (Jorge Pullin) and began a search (currently on-going) for a new permanent director. Since this time, five senior personnel (including Seidel) have left CCT. Starting in 2009, LSU has faced several significant and currently ongoing budget cuts that are currently impacting the CCT.

The issues faced at LSU are similar to those at other institutions responding to the nation's call for an advancement of computation, computational science and interdisciplinary research. We believe it is important to carefully analyze the experiences of centers such as at LSU, as we have attempted to begin in this paper, in order to establish best practices for new initiatives or to lead to more fundamental change. From our experiences at CCT, we can highlight four key points that we feel are crucial for the success and sustainability of computational research centers such at CCT:

(1) The three facets of computational science shown in Figure 1 have be taken seriously on the campus at the highest levels and seen as an important component of academic research.
(2) HPC facilities on campuses need to be integrated with national resources and provide a pathway for campus research to easily connect to national and international activities.
(3) Education and training of students and faculty is crucial; vast improvements are needed over the small numbers currently reached through HPC center tutorials, computation and computational thinking need to be part of new curricula across all disciplines.
(4) Emphasis should be made by the funding agencies on broadening participation in computation, not just focusing on high end systems where decreasing numbers of researchers can join in, but making tools much more easily usable and intuitive and freeing all researchers from the limitations of their personal workstations, and providing access to simple tools for large scale parameter studies, data archiving, visualization and collaboration.

References

[1] Gabrielle Allen, Charles McMahon, Edward Seidel, and Tom Tierney, *The 2003 Louisiana Optical Network Initiative (LONI) Concept Paper*, CCT Technical Report Series (2009), URL http://www.cct.lsu.edu/cct-trs/download.php?file=CCT-TR-2009-3/CCT-TR-2009-3.pdf.
[2] Blue Waters project, URL http://www.ncsa.illinois.edu/BlueWaters/.
[3] CCT Faculty Plan, 2006, URL http://www.cct.lsu.edu/~gallen/Reports/FacultyPlan_2006.pdf.
[4] CCT Strategic Plan, 2006 — 2010, URL http://www.cct.lsu.edu/uploads/CCTStrategicPlan20062010.pdf.
[5] CCT Technical Report Series, URL http://www.cct.lsu.edu/CCT-TRS.
[6] Center for Computation & Technology (CCT) at Louisiana State University, URL http://www.cct.lsu.edu/.
[7] Committee on Facilitating Interdisciplinary Research, National Academy of Sciences, National Academy of Engineering, Institute of Medicine, *Facilitating Interdisciplinary Research*, 2004, URL http://www.nap.edu/catalog/11153.html.
[8] HPC@LSU website, URL http://www.hpc.lsu.edu/.
[9] 2009, International Assessment of Research and Development in Simulation-based Engineering and Science, URL http://www.wtec.org/sbes/.
[10] LSU Information Technology Services unit, URL http://www.lsu.edu/its/.
[11] Open Science Grid, URL http://www.opensciencegrid.org/.
[12] Original plan for the LSU CAPITAL center from 2003, URL http://www.cct.lsu.edu/~gallen/Reports/LSUCAPITAL_OriginalCenterPlan.pdf.
[13] President's Information Technology Advisory Committee, *Report to the president of the us, june 2005. computational science: Ensuring america's competitiveness*, 2001, URL http://www.nitrd.gov/pitac/reports/20050609_computational/computational.pdf.
[14] SURAgrid, URL http://www.sura.org/programs/sura_grid.html.
[15] TeraGrid, URL http://www.teragrid.org/.