# White Paper

# Investigating the Utility of High-Performance Computing Capabilities At Six Small Liberal Arts Colleges

February 9, 2010

A project supported by the
Teagle Foundation

Submitted by
Valerie Barr
Union College

## Introduction

### Overview

Liberal arts colleges want to provide the greatest academic opportunities to our students, and provide faculty with the tools necessary to present these opportunities in their undergraduate courses. In this context, liberal arts colleges must continually investigate and assess the importance of technology to the educational and research goals of our institutions. Recent and ongoing technological developments pose challenges and questions that we must address. What facilities do we need to make available to faculty for use in their research? What technology use should be encouraged by our students? How are new digital tools changing interdisciplinarity and the ways in which faculty collaborate? What impact do these changes have on our institutions?

A number of recent publications have addressed aspects of these technologically motivated shifts in the academy. These include discussion of institutional change and participatory learning [1], the role of digital scholarship [2], and the specific impact of new technologies on the humanities and social sciences [3]. However, one area not specifically addressed is the potential role of high performance computing at liberal arts colleges.

Supported by a grant from the Teagle Foundation, a consortium of six liberal arts colleges in central New York State (Appendix A) undertook an investigation of whether students and faculty at small liberal arts colleges would benefit from having available powerful, high-performance computing (HPC) technology. We also endeavored to determine what would be the best way to provide this capability, taking into account all costs, including those associated with infrastructure and support. The participating liberal arts colleges, close in size, form a natural cohort connected by similar academic curricula in the sciences, humanities, and social sciences, a strong commitment to undergraduate research, and geographic location (central and eastern New York State).

HPC uses powerful computers to solve problems that require significant computational resources. Generally these machines are referred to as supercomputers or cluster computers, though the distinction is constantly shifting as hardware develops and prices drop. Some supercomputers use thousands of central processing units (CPUs) with very high-speed interconnections and massive storage. Others, more often referred to simply as cluster computers, use tens, hundreds, or thousands of CPUs with less specialized interconnections and less massive storage. In both cases work is distributed among the processors, allowing faster computation than if a single computer were used. The significant costs associated with a supercomputer (physical space, air conditioning, campus bandwidth, maintenance, plan for dealing with obsolescence) usually make this an infeasible option for a small college. Smaller cluster computers tend to be a more reasonable way to provide local HPC capability, in part because an institution can start small and increase capacity later as demand grows. Increased capacity can be achieved by either adding on to an existing cluster or installing a new one. It is now also possible to achieve the computing capabilities of a small cluster by using a multicore desktop machine. As of this writing (early 2010) one can buy an 8-core (8 processor unit) Mac Pro for about $3300.

Different machine configurations provide the ability to address different kinds of problems. On the one hand, consider the large scale simulations used to address problems from the sciences, social sciences, and engineering. These problems, found in climate prediction, energy research, biological modeling, etc., are best handled using supercomputer systems or large clusters with many processors and high speed inter-processor communication. By contrast, problems that arise in the humanities are likely to involve processing the very large amounts of data involved in extensive image and text analysis. These problems are often best tackled with smaller cluster computers or multicore desktop systems with high-capacity storage.

With the increasing capability and declining costs of hardware, faculty across an ever wider range of disciplines can begin to explore applications of computing, including HPC, to their work. Doing so will enhance and strengthen the academic offerings of liberal arts colleges and our students' learning and research experiences. Each college within our group strives to provide students with the most opportunities and experiences possible to give them knowledge and technical expertise they will take with them into the workforce after graduation. For example, Hamilton College successfully developed and implemented a facility (now two cluster computers and two shared-memory systems, totaling 220 CPUs) to support the **Molecular Education and Research Consortium in Undergraduate computational Chemistry** (MERCURY), a resource for chemistry students and researchers at eight liberal arts colleges.

The goal of our Teagle Foundation funded planning grant was to investigate the utility of HPC by exploring answers to questions such as:

- How is HPC being utilized in various academic disciplines?
- What degree of HPC can be effectively utilized in the liberal arts setting? How will it enhance student learning at these institutions?
- What is the best way to provide HPC access to faculty and students?

We chose to address these questions by holding a series of three workshops, followed by a final summary meeting of key consortium participants. We encouraged broad participation at the workshops from faculty, administrators, and IT staff. We invited representation from faculty in disciplinary and interdisciplinary programs, information technology experts, and academic administrators, recognizing that each would bring a different perspective on the benefits, issues, and concerns regarding HPC capabilities at small liberal arts colleges. They would also make it possible for us to consider the academic programs and faculty at each college, including their existing applications of HPC. The grant web site includes speaker presentations and workshop summaries.

**Summary of the workshops**

In formulating our workshop agendas, we considered changes in computing technology and in various disciplines over the past twenty years. In the past, fields such as the sciences and social sciences had smaller data sets than those available today, or had data that was largely qualitative. At that time the majority of computer analyses were done at the graduate or post-doctoral level. Today much more data is readily available and data sets have grown in size. With wider availability of sophisticated computing power and software packages, analysis of large data sets can be carried out at the undergraduate level. This includes work in numerous fields such as: genomics and protein folding in biology, analysis of terabytes of physics data, macroeconomic analysis and prediction using current data (such as the effects of the economic stimulus package), economic and statistical analysis using massive simulation, and image analysis. The development of digital humanities represents a significant paradigm shift as scholars begin to apply computational methods to the analysis and presentation of material. In many instances this does not necessitate HPC capabilities, but there are applications for which the computing power of HPC is appropriate, such as image analysis, video, and graphics rendering. The National Endowment for the Humanities is promoting digital humanities through the Office of Digital Humanities. The Digging Into Data Challenge provides examples of specific projects that require use of HPC within a humanities context. Other organizations that provide resources for those interested in digital humanities and social sciences are the Institute for Computing in Humanities, Arts, and Social Science (I-CHASS), the Humanities, Arts, Science, and Technology Advanced Collaboratory (HASTAC), the Maryland Institute for Technology in the Humanities (MITH), and MATRIX: The Center for Humane, Arts Letters and Social Sciences Online.

The goal of our first workshop was to gain a common understanding of the spectrum of HPC technologies. The second and third workshops focused on applications of HPC across the disciplines. Workshop two explored humanities and social science applications, while workshop three looked at various science and engineering applications. At our fourth and final meeting we discussed the possible ways in which we could, as individual institutions and as a consortium, avail ourselves of HPC. This included discussion of infrastructure and maintenance issues as well as how to adequately support faculty who want to use HPC resources. Appendix C includes the speaker information from each workshop. All summary reports and many of the speaker presentations are available at http://cs.union.edu/~barrv/Grants/Teagle/teagle-overview.html

**Review of existing usage of HPC on each consortium campus:**

*Bard College:* At present Bard College has no dedicated cluster computation. Ad hoc clusters have been created for teaching some computer science courses, primarily to study networks, not HPC. There is no current use of off-campus HPC, although there is an expectation that there will be a need for HPC access within the next five years. There is currently work with spiking neural simulations and work in climate modeling that may eventually lead to use of cluster computing.

*Colgate College:* Colgate has two cluster computers which are managed by individual faculty members. One is used by a faculty member in Chemistry for research, though he would like access to a more powerful HPC resource that could be used for classroom assignments as well. The other cluster is used in the Computer Science department for teaching. There is currently no use of external HPC resources.

*Hamilton College:* Hamilton's experience scaling up HPC use over a 7 year period provides a sense of the experience that other schools may have as they start out with HPC. Hamilton has seen a gradual growth in the use of HPC during this time. Initially one scientist began using HPC, joined subsequently by a few more scientist and then by some interested social scientists. Some faculty now have time on Teragrid, a resource funded by the National Science Foundation. Hamilton's experience indicates that it can be important to initially have a local resource, particularly for students, so that they are exposed to HPC before they start using an external resource. It has also been very important to the faculty at Hamilton that they have a dedicated systems administrator. In addition, they would also like to have someone who can help faculty move their research onto HPC platforms.

A group of four Hamilton faculty, from chemistry, physics, biology, and anthropology, have recently been awarded a grant of $177,950 through the National Science Foundation's (NSF) Major Research Instrumentation (MRI) program to fund a shared-use state of the art computing cluster. The project is entitled "MRI-R2: Acquisition of a High Performance Computing cluster with a fast interconnect to enable shared-use, college-wide computational investigations at Hamilton College". This cluster will primarily support research programs in enzymatic catalysis in RNase A and drug-target interactions of influenza neuraminidase; the evolutionary biology of eukaryote supergroups; the nature of Ultra-Luminous Infra-Red Galaxies; and examination of social networks through an analysis of projectile points and gravestone raw materials. This expansion of the computing resources will also allow greater access to computational tools for curricular projects. In addition, the cluster will be available to other Hamilton community users.

This brings the total at Hamilton to eight cluster computers. The college is currently in the process of developing a governance structure which will provide central support as well as guidelines for access.

*Skidmore College:* Skidmore College has a 20-CPU Xserve cluster in the Mathematics and Computer Science department. This cluster has been used for various CS initiatives over the past five years. The MCS Linux lab has also been used on occasion for HPC clustering, but student use of the lab limits the degree to which it can be dedicated to clustering.

*Union College:* Union has two cluster computers, an 8-CPU machine belonging to Prof. Mike Vineyard in Physics and a 50-CPU machine belonging to the Computer Science department. The latter machine is occasionally used by Prof. Janet Anderson (Chemistry) for computational chemistry work, though she more often uses equipment at Wadsworth Laboratory (University of Albany). Prof. Steven Romero (Psychology and Neuroscience) investigated the possibility of using the supercomputer at RPI, but has not yet done so. However, he has recently acquired an 8-core Mac for his research. Prof. John Rieffel (Computer Science) is using a GPGPU box (general-purpose computing on graphics processing units) for simulation of soft robotics. His machine consists of three high power graphics cards which account for 720 parallel cores.

*Vassar College:* The Vassar CS department currently shares with biology two dedicated clusters with 16 cores each along with one dedicated 4-core machine. They also have software installed to use their Computer Science lab machines as a cluster which they have done for course work involving parallel processing and non-course projects such as a grid program. They expect to soon be purchasing, under an HHMI grant, a minimum 32-core cluster with approximately 20TB of storage. This system will also supply the "head node" for the CS lab machine cloud that should begin to see regular use. All of this equipment is maintained by a system manager who is dedicated 9 months to Computer Science and 3 months to the Vassar Center for Collaborative Approaches to Science (CCAS), funded by their HHMI grant.

**Key issues to take into consideration**

A number of key issues arose during the workshops, as well as at our final summary meeting. We present these here, both because they inform our recommendations and because they are issues that should likely be taken into consideration by any institution considering a move to HPC.

*The Evolution of HPC in the Liberal Arts Environment:* The evolution of HPC usage on a campus will be driven by research and curricular needs, which will vary considerably by discipline and by faculty member. There are generally two scenarios: (1) a faculty member already has significant experience from graduate school and needs HPC facilities in order to continue his or her research; (2) faculty members will want to experiment with HPC, perhaps as part of a project involving students, perhaps in the context of new research projects. Supporting faculty members who are already well-versed in the uses of HPC for their research will require a greater investment than is necessary for the case where faculty and students are experimenting with HPC. Certainly the evolution of HPC usage on a campus will also require faculty development, and should be focused initially around a few key faculty members who already have interests or experience in this area. Investment in resources must initially be based on the existing faculty interests and strategically focused on areas that may promote additional faculty interest in HPC.

*Consider existing computing resources:* All of our schools have a significant number of "wasted computing cycles", the result of computers that are running but are not actually in use (for example, lab computers when students are not present, and faculty computers when they are teaching class or are not on campus). Rather than waste this computing power, it can be harnessed across the campus network for large computing projects. For example, there are very large scale projects such as SETI@home and Folding@home which capture idle cycles from registered users. There are systems administration issues associated with capturing computing cycles from existing machines over a network, but it is an approach worth considering, and good software is available that will harness unused cycles for general use. It is also important to be aware of resources that are controlled by departments and individual faculty, and determine whether any of those resources could be used more widely by other faculty on campus.

*Needs vary within an institution:* Teaching needs are different than research needs. Generally teaching **about** HPC and teaching **with** HPC do not require as much computing power as research applications do. Teaching **about** HPC in a CS department, while not requiring high absolute performance, does generally require a system that can be "broken" by students on a regular basis without disrupting work being carried out by others. That is to say, students have to be able to make changes to the operating software, adjust the network that is internal to a cluster, and otherwise make adjustments that immediately affect the realtime performance and operation of the system. On the other hand, teaching with HPC in a non-CS department will require a system with high reliability and may require high performance.

Furthermore, the type of HPC facility required will depend on the type of problems being addressed. Some problems require extensive data storage capacity, while others require significant processing power and high rates of interprocessor communication.

*Needs vary across a consortium:* Among any group of institutions there will be differences in the extent to which they already utilize HPC. It is important to be aware of the resources each institution is already using, who maintains those resources, whether they are storage intensive or computing intensive, and what sort of work is being carried out on them. Schools that are already using HPC will be interested in larger, faster systems, primarily in the form of external resources. Schools that have no or few clusters now and are doing little with HPC may be interested in purchasing small local clusters as a way to get started.

*Needs vary across disciplines:* The use of computing in general, and HPC in particular, varies across disciplines. Many faculty in science, engineering, and some social science disciplines have already been using computing as a component of their research and teaching. For these faculty members, the step to using HPC is not necessarily very large if proper support is present. For example, they may need help adapting an existing application to a multicore system or learning how to use the operating system on a cluster computer. However, the use of computation in research and teaching is much newer to many faculty in the humanities and other social science fields. They have not necessarily been inclined in the past to collaborate with research institutions, to work with multimedia data, and to engage as members of consortia. Therefore, the support needs may be greater for faculty in these areas.

*New young faculty will increasingly expect HPC resources:* Increasingly graduate students across a wide range of disciplines are going to have exposure to HPC resources. If liberal arts colleges want to compete with research universities for these faculty members, we must offer access to these resources.

*Support is critical:* We will see an increase in the numbers of faculty using HPC, or an increase beyond existing usage, if we make it easy for faculty to take the next step. Faculty will generally need help identifying projects that are appropriate for HPC, getting their projects running on an HPC platform, and moving projects to the next level of HPC resource. This requires HPC support similar to that many of us currently receive from instructional technology staff. This can help faculty identify the appropriate computing resource to use and show them how to use it and how to configure their projects to run on that resource.

*Infrastructure and support are expensive, but sharing costs and personnel can help:* The costs of support personnel and hardware purchase, maintenance, and upgrading can be considerable. For example, a 36-node 288-core system (as of early 2010) will cost approximately $177,000 (see Appendix B for sample configurations). In addition, a suitable HPC facilitator who can work with faculty from a range of disciplines and be familiar with many HPC options and systems could easily cost $75,000 - $100,000 per year. It can be very difficult for a single institution to provide both a high end computing resource and the support personnel who will facilitate its use by faculty. However, sharing these costs across a group of schools can help make them feasible. Sharing support personnel makes sense financially, and can also help build connections across campuses between faculty working in related areas.

*Role of computer science department:* Partnerships between computer science faculty and those in other disciplines work well, and can help move a problem onto an HPC resource, in cases where there is an interesting CS component to the project. But support must come from elsewhere in other cases. Institutions cannot expect that computer science departments will take on the job of bringing HPC to the campus or will take on system maintenance.

*Faculty incentives may be necessary:* Some form of incentive may be necessary to encourage faculty to learn to use HPC. These incentives could be in the form of:

- money or release time for course development work that incorporates HPC
- the designation of faculty fellows who get release time to advance their HPC related work and advocate for HPC on their campuses, and would also receive support to travel to the other campuses and work with faculty there
- summer research money
- coverage of expenses incurred to attend relevant conferences (digital humanities, digital social sciences, other computational science related conferences)

**Hardware options**

There are a number of hardware options for small schools wishing to access HPC. One option, largely independent of all the others, is to develop individual school facilities, such as local cluster computers. We note that in this case costs would likely become part of the IT services budget line, whereas the costs for other options would likely be explicit line items. With adequate faculty interest and IT support, it is also possible to build small clusters out of castoff computers on campus.

1. *External Resources*: There are a number of external HPC options. In many cases there are statewide resources. For example, in New York these include the supercomputers at the University at Buffalo Center for Computational Research, the New York Center for Computational Science at Stony Brook University, the College of Nanoscale Science and Engineering at the University at Albany, and the Computational Center for Nanotechnology Innovations at Rensselaer Polytechnic Institute. There are also supercomputing facilities nationwide that award CPU time on a competitive basis, as well as government granting agencies that award computing time. Cloud computing resources are available on the commercial market from companies such as Amazon and Google. Finally, the TeraGrid offers startup or educational allocations of up to 200,000 Service Units and up to 5 TB disk storage (up to 25 TB of tape storage) as well as larger research allocations.
2. *Common resource at each school:* One option is to install the same cluster hardware at each institution within a consortium. This option would allow the IT support staff to share knowledge across the consortium, including faculty and student training and computer support. However, this is not likely to be a viable approach across a consortium. If the schools differ in the extent to which their faculty already use HPC, some schools may already have clusters and will be ready to move up to larger ones. For schools that are just starting out with HPC, a small cluster is a low cost entry approach and a way to demonstrate on campus the value and utility of HPC, but they may have to take this step independent of the other members of the consortium.
3. *Single large computing resource, owned by the consortium:* An HPC resource could

be housed at one of the member institutions. However, this puts large space, infrastructure, and maintenance burdens on one institution and requires financial arrangements between institutions within the consortium. Where possible, the HPC resource should be located at a central facility (such as NYSERNet) rather than on any single campus. Then no single campus has to deal with issues of space, A/C, power, etc., and the consortium will be better able to divide the hardware costs as well as the costs for space and system administration.

4. *High Throughput Computing (HTC):* Some problems do not necessarily require high computational speed, but do require large amounts of processing capacity (computing throughput) over long time periods, such as months or years. These needs can be met by using existing available computing resources, such as on an entire campus, as discussed above. There are tools, such as *Condor,* which handle task allocation and workload management across all available computers (see discussion above of using existing resources).

**A proposal for incorporating HPC at liberal arts institutions:**

Our workshop series made it clear that there are opportunities to use HPC across all disciplines. Greater incorporation of HPC in teaching and research at our institutions will expose our students to the most current material possible and give them hands on experience with the most modern computational tools and approaches. However, institutions must provide support in this area for faculty. In addition, we believe that the infrastructure and resource costs are much more manageable if shared among a group of institutions. While hardware costs alone may be reasonable, it may be very difficult for a single small institution to absorb the costs of installation, maintenance and ongoing upgrades, and additional IT and instructional technology support. We present here a proposal for how HPC could be incorporated among a group of liberal arts institutions.

1. Create a formal consortium of schools.
2. Inventory the existing HPC facilities at each school and identify current research and teaching projects that utilize HPC. It is also important to develop a timescale that matches projected need to available and projected resources. This will help the consortium develop a timetable for purchasing, hiring, and exploration of external resources.
3. Consortium schools should consider investing in a shared machine. It is possible that this machine could initially be housed at one of the consortium schools if one has adequate facilities and technical support. However, it is likely that eventually the shared resource would be housed at a central facility. This allows systems administration responsibilities to be handled by a consultant, paid for by the consortium, and carried out at the central facility as needed. Institutions are, of course, also free to install their own individual cluster computers, but will have to support and maintain those themselves.
4. All schools in the consortium should eventually be on Internet2 and have adequate videoconferencing facilities.
5. The consortium should hire a facilitator who will work with faculty at each institution to identify projects suitable for HPC, and help move projects onto HPC (or move projects that are already using HPC to the next level of HPC resource, where appropriate).

This person must be knowledgeable about external resources such as state-wide resources, national resources such as Teragrid, etc. as well as work with the instructional technology specialists and the faculty on each campus (particularly visiting HPC professor and campus HPC fellows, see below). He or she will be appointed as a Campus Champion to the Teragrid, on behalf of the consortium. The facilitator should have a budget that supports HPC seminars held for consortium faculty and travel to schools within the consortium.

6. Each year the consortium should invite a visiting professor with HPC expertise. This person will be engaged in research which uses HPC, and will work with existing faculty to identify projects that are suitable for HPC.
7. Each consortium school should each year identify an HPC fellow from among the faculty. The group of fellows will work with the facilitator and the visiting professor to get projects running on HPC resources, they will hold seminars on each campus where they will discuss HPC with other faculty, and they will screen proposals and make decisions about allocating institutional summer research support money. The fellows will receive course release.
8. The consortium should participate in organizations such as iCHASS, MATRIX, HASTAC and make a concerted effort to send key faculty (the fellows, or others as appropriate) to relevant conferences and workshops. The goal is to create an ever larger presence on each campus of faculty who are committed to using HPC and are developing experience using HPC in their research and teaching. This will also increase the presence of liberal arts faculty in larger discussions about uses of HPC across disciplines.

**References:**

[1] The Future of Learning Institutions in a Digital Age. Cathy N. Davidson and David Theo Goldberg. MIT Press. 2010.

[2] Current Models of Digital Scholarly Communication. Results of an Investigation Conducted by Ithaka for the Association of Research Libraries, November 2008. Nancy L. Maron and K. Kirby Smith

[3] Our Cultural Commonwealth: The report of the American Council of Learned Societies Commission on Cyberinfrastructure for the Humanities and Social Sciences

## Appendix A

The consortium, led by Union College, includes:

- Bard College, Annandale-on-Hudson, NY.
- Colgate University, Hamilton, NY
- Hamilton College, Clinton, NY
- Skidmore College, Saratoga Springs, NY
- Union College, Schenectady, NY
- Vassar College, Poughkeepsie, NY

## Appendix B

Sample HPC configurations:

| | Number of Cores | Memory (GB) | Disk | Cost |
|---|---|---|---|---|
| **High** | 288 | 864 | 2.8TB | $177,000.00 |
| **Medium** | 156 | 80 | 5.4TB | $95,000.00 |
| **Low** | 80 | 160 | 800GB | $36,000.00 |

## Appendix C

**Workshop 1:**

- Geoffrey Fox, Indiana University, discussed the general landscape of HPC.
- Scott Kaplan, Amherst College, discussed the experience of setting up a large cluster and making it available to faculty across the campus.
- Peter Pacheco, University of San Francisco, discussed an undergraduate course in parallel computing.
- Brian Macherone, University at Albany and NYS Grid, discussed the NYS Grid and opportunities available for involvement by liberal arts colleges in the state.
- Michael Ridley, NYSTAR, discussed the resources available state-wide within New York and ways in which liberal arts colleges can tap into them.

**Workshop 2:**

- Orville Vernon Burton, Burroughs Distinguished Professor of Southern History and Culture at Coastal Carolina University; Founding Director and currently Chair of Board of Advisors, Institute for Computing in Humanities, Arts, and Social Sciences (I-CHASS); Associate Director for Humanities and Social Sciences at the National Center for Supercomputing Applications at the University of Illinois. "A framework for cyberinfrastructure for the humanities, arts, and social science"

- Dean Rehberger, Director, MATRIX: Center for Humane Arts, Letters, & Social Sciences Online. "The New Curiosity Box: High-Performance Computing and the Humanities"

- Geraldine Heng, Perceval Fellow and Associate Professor English, Middle Eastern Studies, and Women's Studies, The University of Texas at Austin; Founder and Co-director of the Global Middle Ages Project (G-MAP), the Mappamundi digital initiatives, and the Scholarly Community for the Globalization of the Middle Ages (SCGMA). "Futures of the Past: the Global Middle Ages, and the 'Clash' of Civilizations"

- Doug Reside, Assistant Director, Maryland Institute for Technology in the Humanities (MITH). "Ways in which Digital Humanities centers make use of the power of high performance computing."

**Workshop 3:**

- Fran Berman, VP for Research, Rensselaer Polytechnic Institute, former director of San Diego Supercomputer Center. "Driving 21st Century Research and Education with 21st Century Tools: Cyberinfrastructure and Today's Universities"

- Carol Parish, Professor of Chemistry, University of Richmond. "Thoughts from a computational chemist: getting 'high performance' out of high performance computing in a liberal arts setting".

- Janine Shertzer, Professor of Physics, College of the Holy Cross. "Big computing at a small school: how to 'beg, steal or borrow' CPU cycles"

- Ray LeBeau, Assistant Professor, Physics and Astronomy, University of Kentucky. "Experiences with Computational Fluid Dynamics on Cluster Computers"

- Bob Panoff, Executive Director, Shodor. "Computational Thinking at the Speed of Right"

*We gratefully acknowledge the generous financial support of The Teagle Foundation for the work of our consortial group.*