

High Performance Green Computing

A Proposal for New Investments in Faculty Hiring in the Departments of
Astronomy and Computer Science

Ronald Snell, Martin Weinberg, Neal Katz, Min Yun, Grant Wilson, Gopal Narayanan,
Houjun Mo, and Daniela Calzetti
Department of Astronomy

Eliot Moss, Prashant Shenoy, and Chip Weems
Department of Computer Science

Rationale

High performance computing (HPC) – using commodity compute clusters, special-purpose multiprocessors, massive data stores and high-bandwidth interconnection networks – has rapidly become an indispensable research tool in science and engineering. In application areas ranging from the design of drugs, to understanding and predicting nano-scale properties of photovoltaic and thermoelectric materials for energy conversion, to modeling and predicting climate changes, to modeling galaxy dynamics, computation is finding an increasingly important role beside theory and experimentation as a tool for inquiry and discovery. Beyond science and engineering, large compute clusters, massive storage, and high-bandwidth connectivity have also become mission-critical infrastructure in medical, industrial, and academic settings.

While the costs of computers, storage, and networking have fallen dramatically over time, the costs of the buildings, power, and cooling that support this equipment have risen dramatically. Indeed, the costs for power and infrastructure exceed the cost of the computing equipment they support. The environmental costs of information and communication technologies are also increasing; recent studies estimate that these computing and communication technology sectors are responsible for 2% of global carbon emissions and these emissions are increasing at 6% annually.

Green computing (*aka* sustainable computing) can be broadly defined as the problem of reducing the overall carbon footprint (emissions) of computing and communication infrastructure, such as data centers, by using energy-efficient design and operations. The area has garnered increasing attention in recent years, both from a technology and societal standpoint, due to the increased focus on environmental and climate change issues. Hence, there is a need to balance the dramatic growth of high-performance computing clusters and data centers in the computational sciences with green design and use so as to reduce the environmental impact. Technical issues in high-performance green computing span the spectrum from green *infrastructure* (energy-efficient buildings, intelligent cooling systems, green/renewable power sources) to green *hardware* (multi-core computing systems, energy-efficient server design, energy-efficient solid-state storage) to green *software* and *applications* (parallelizing computational science algorithms to run on modern energy-efficient multi-core clusters).

At the state level, Governor Patrick and the presidents of UMass, MIT and BU, together with Cisco and EMC, recently announced a Green High Performance Computing Center that will be built in Holyoke, MA at an estimated cost of \$100M. Holyoke was chosen due to the availability of inexpensive clean hydro-power as well as for the economic development benefit to this gateway city. The shared, green co-location facility for high performance computing will provide researchers with the infrastructure to undertake cutting-edge collaborative research in computational sciences and green computing. A key goal is to make the Holyoke facility a national showcase for designing green data centers for high-performance computing. The UMass Amherst campus is well-poised to benefit from and exploit this major state investment in a green high-performance computing facility. The campus has substantial strengths in computational sciences (an area that makes extensive use of high-performance computing) and computer sciences (which is leading the campus effort in green computing). This proposal argues that a modest additional investment of faculty lines to fill certain critical gaps can vault the campus into becoming a major national player in research and education in high-performance green computing.

This proposal couples the interests of the Department of Computer Science for research in the development of the hardware and software involved in green high performance computing with interests of the Department of Astronomy for the use of such resources for scientific research

involving immense data sets and theoretical models that require massive amounts of computation to evaluate. The two departments have a history of collaboration both in developing high performance computing capabilities and in developing statistical techniques such as the Bayesian Inference Engine (BIE, a parallel-computing-optimized research tool for statistical inference). Individually they have great depth that goes beyond the current relationships, and there is much potential for further collaboration, both between the departments as well as with other computational science disciplines on campus. However, those new relationships are inhibited by a gulf in understanding that separates their researchers. This proposal approaches the creation of these new ties by bringing in a specialist in each department whose complementary research areas develop intra-department computational science collaborations, and a third person whose research initially connects the two specialists, then their collaborators, and eventually expands to involve collaborations in additional departments and colleges.

Computational science is interdisciplinary by nature, lying at the intersection of science, engineering, computer science, mathematics and statistics. Within the College of Natural Science, many departments have faculty who are actively involved in computational science. In fact, there has been an effort for a number of years to establish a Computational Science Center on campus to bring together talented and productive faculty to focus on interdisciplinary computational issues with an emphasis on developing a coherent interdepartmental computational curriculum and providing a campus-wide forum and resources for computationally-based research. The hires proposed will form a nucleus of leadership for carrying this goal forward. They will build on the existing network of active researchers and enhance computational science on campus while exploring the novel paradigms of green computation. Funding for interdisciplinary computation and energy efficiency is increasing nationally and is currently a major theme at the NSF, NIH, DOE and NASA.

On the Computer Science side, there is a need for improved machine architecture, compilers, run-time support and parallel algorithms to fully realize the potential of green computing. Modern parallel processors can easily consume megawatts of power, occupy thousands of square feet of space, and need to have their hardware replaced on a regular schedule to maintain a meaningful level of performance. Thus, high performance computing centers can be environmentally damaging in multiple ways. On the Astronomy side, there is a strong need for increased computer performance, however, realizing the desired computational potential is no longer simply a matter of building bigger, faster, and more power-hungry computers every few years, green computing research endeavors to optimize its use to deliver the most *computation per watt* to applications. The success of this research depends on bringing together the right combination of researchers in computer science and computational science so that solutions can penetrate all the way from low-level hardware choices up to the algorithmic approaches taken in solving the scientific problems to which it is applied.

Our proposed 'cluster' will leverage the planned Holyoke facility, which will provide a home for the next generation of computers that will be used by researchers at the University, including the faculty in Astronomy and Computer Science, and will provide a challenge for faculty in Computer Science to develop the next generation of programming languages, compilers, run-time systems, and algorithms to optimize the performance and efficiency of the installed computer systems. This green facility and our proposed venture into green computing will help make the University and the Commonwealth leaders in next-generation computing.

In addition to the research activities of the hires envisioned in this plan, there is also an important education component. Training students at both the undergraduate and graduate level in both the architecture of many-processor computers and in computational techniques that use them effectively and efficiently is extremely important. Astronomy has developed computational techniques courses at the sophomore (P281 – Computational Physics, taught jointly with Physics) and graduate (A732 – Computational Methods in the Physical Sciences) levels. The proposed new hires will add to the expertise in computing on campus and will contribute to the development of an interdisciplinary computational curriculum on campus. Such a curriculum would benefit students in most degree areas, since this expertise is increasingly important in not only in science and engineering fields, but also in business and the humanities. An interdisciplinary curriculum will increase the student's value on the job market, as both academic and corporate surveys highlight information-technology trained personnel as a critical need.

Faculty Hiring Plan

We request faculty positions to enhance high performance / high efficiency computing, one in the Department of Astronomy, to build a computational astronomy group that is world-class and having critical mass, and two in the Department of Computer Science, one focused on the technology

of power-efficient computing through architecture, language, and compiler innovation, and one focused on issues of efficient computational science (parallel algorithms for cluster / grid / cloud computing, management of large dynamically changing collections of computing resources, etc.) to bridge to computational astronomy, and computational science more generally.

In astronomy, computational astrophysics is an essential component of research today and one of the major themes of the Astronomy Department. The existing computational theory group brings state-of-the-art numerical techniques to bear on otherwise intractable astronomical problems with complex geometries, multiple scales, and physical interactions. The work of this group has led to new paradigms in studying the large-scale structure of the universe, through comparison of cosmological simulations with observations, of which the Large Millimeter Telescope (LMT) will play a major role in the future, and to new insights into galactic interactions and mechanisms governing their underlying dynamics. On the observational side, the Astronomy Department will have a leading role in the science output of the LMT, a state-of-the-art facility unmatched in capability. Modern instruments on today's telescopes are capable of acquiring data at ever increasing rates, thanks to detector technology advances, requiring increasingly sophisticated computational solutions for data acquisition, processing, analysis, storage, and distribution to achieve their science goals. Terabyte-sized datasets at multiple wavelengths are now the norm and require advanced high performance computing for full analysis, not least appropriately to correlate, align, and weight observations acquired with different instruments at different wavelengths and having varied formats and properties. The observational group within the Astronomy Department is at the forefront of such multi-wavelength multiple dataset analysis, and brings unique expertise to the international astronomy stage. Both the theory and observational groups have been extremely prolific, with many citations to their work, and have been highly successful in securing grant funding.

Faculty in the Departments of Astronomy and Computer Science have jointly developed the Bayesian Inference Engine (BIE) system, which is publicly available and which combines state-of-the-art approaches in numerical statistics, data representation, and efficient computation to provide analysis tools for terabyte-class scientific datasets. This tool has been applied to theory-data statistical analyses using large observational surveys, including UMass's own 2MASS survey. The disciplines covered in the development effort for BIE specifically include astronomy and Earth-based geographical resource management such as forestry and ecology, but the approach applies to any inference problem with mapped data. The cluster hires would build on this existing collaboration to improve both department's capabilities and capacity in computational science generally, and efficient / green computing more specifically. This will lead to establishing strong collaboration with other departments in computational science as the new hires mature in their UMass connections.

In the Department of Astronomy there is a strong nucleus of faculty focused on computational studies and the statistical analyses of large datasets, and many more faculty that would greatly benefit from further computational expertise. We believe the addition of another computationally-oriented researcher in astronomy will provide the critical mass needed to put the UMass group in position to become a leading national and international program in computational astronomy. We plan to advertise broadly seeking a talented astronomer with a strong background in utilizing computational techniques in either theoretical studies or in the acquisition, processing and statistical analysis of large data sets. A theory hire would greatly complement the work of our present computational astrophysics group, and in addition work closely with observers to help interpret, synthesize, and model observational results. Likewise, an astronomer working on statistical analysis would benefit from the BIE development and position our Department at the forefront of the growing field of large data-base analysis. The increase of our faculty even by just one new hire, will have an enormous impact on the visibility of our computational program nationally and in success at grant funding with NSF and NASA.

Computer Science proposes to hire two faculty members to support this cluster. One will be dedicated to the technical aspects of green computing. The second will be the computational science "connector" who links the other two hires together, as well as taking a leadership role in developing broader collaborations among other departments and colleges around the common interests that they share in computational science and engineering.

Green computing is an emerging area of research that emphasizes reduction of the impact of computation on our environment. It requires a multi-pronged approach from manufacturing through procurement and operation to disposal of computing resources. Computer science focuses mainly on the procurement and operational phase of the lifecycle. In this phase, the choice of hardware and how the software makes use of it has significant impact on the consumption of material resources, space, and energy. The concerns naturally break down into two distinct, but overlapping, focus areas. One is

on the *technology* of green computing, which includes low-power / energy-efficient hardware, software techniques to reduce power consumption, and compilers and run-time systems that optimize software for efficient computation on a given platform. The other focus area is on the overall algorithmic approach of the computations (which involves an understanding of the domain scientists computational needs) and on the larger scale management of clusters, grids, and clouds of resources. Computer science lacks faculty with either of these foci. The new technologist would fit well with existing faculty that work on architecture, compilers, run-time systems, and operating systems, and we hope would also reach out to ECE concerning low-power design. Thus that new hire would not lack for close collaborators, but would build critical mass and bring unique emphasis on power-efficient computing. The computational science-oriented person would connect with a variety of Computer Science faculty with interest in management of larger scale systems and in data-intensive and large computations. However, they would bring a new and unique focus on computational science and of reaching out and bridging to the other sciences (and perhaps beyond), of which Astronomy would be the first. The existing collaboration of our two departments around the BIE would be a starting point from which this new person could expand.

The computational science person would also catalyze efforts to develop interdisciplinary educational programs in computational science, sorely needed across the scientific disciplines. The exact specialty of the person might range from parallel algorithms to cloud computing to visualization. What is more important is their commitment to bridging to the sciences on the one hand and to collaboration in green computing on the other.

In summary, the hiring goal of this cluster is to create the seed for a much larger and broader effort to establish the University of Massachusetts as a leader in developing and applying green computing techniques through a university-wide, on-campus center that has close ties to the facility in Holyoke. Furthermore, although we have avoided digressing to a discussion of five-college involvement, given the long history of collaborations that both departments have across the colleges, we also see this initial effort as quickly expanding to encompass researchers in the corresponding departments around the valley.

Investments and Funding Opportunities

Significant resources have been expended to support the activities related to these hires. In astronomy we have an older computer cluster that has been operated by the group for the past 9 years. This cluster was originally funded by an NSF MRI grant (jointly with the faculty in Computer Science and several other NSM departments), while the infrastructure costs were paid for by the College. In astronomy, this facility has supported the work of 18 federal research grants totaling over \$5 M. A new and much more powerful computer cluster has just now become operational and is located in space provided by OIT. The University has made a significant investment in renovating space in the low-rise of the LGRT for the co-location of HPC clusters. The facility supported by OIT provides power and thermal control for our current system much like the space being considered for Holyoke will provide for future systems. The astronomy computing cluster was largely funded by grants from NSF, but with contributions from the Office of the VCR. The new astronomy cluster has 1500 processing cores, 3000 Gigabytes of RAM, a 10 Gigabit per second interconnect rate, and 100 Terabytes of data storage. It is currently the most powerful computer cluster on campus.

Computer Science has also recently brought a new cluster on line, consisting of 480 cores, with 960 Gigabytes of RAM, 10 Gigabit per second network, 65 Terabytes of centralized disk, and 8.4 Terabytes of distributed storage. Funding for this cluster came from an NSF MRI award, with funds for renovation of the LGRC basement space to house it coming from the Vice-Chancellor for Research. It is used for a wide range of research studies, including data mining, information retrieval, networking, remote sensing, computer vision, and extreme-precision numerical algorithms. Much of this research is relevant to the massive datasets that characterize modern astronomy and most of computational science.

In the future, both departments will capitalize on the new Holyoke Green High Performance Computing Center for the next generations of high performance computers. The presence of this Center will facilitate funding efforts and provide an ideal environment for developing machines with even greater computing power and energy efficiency. Such machines will be needed to both provide the computing needed for future cosmological simulations and to analyze the terabytes of data that will be obtained with future telescopes, including the Large Millimeter Telescope (LMT).

In astronomy funding has traditionally come from either the NSF or NASA. The computational group in astronomy at UMass has a strong record of funding success. Over the past five years, grants

in astronomy that are computationally oriented total \$3.5 M. Funding has been steadily increasing, and in this past year alone four awards totaling \$1.2 M were obtained from NSF and NASA. In addition, significant funding has been provided specifically for instrumentation development (\$1.8 M over the past five years) that will benefit from this computational initiative. Since computational science is a major theme at these funding agencies, we believe that there will be many funding opportunities to support computational astronomical science. We believe that funding success by astronomers at UMass will be greatly enhanced by the addition of another computationally oriented faculty member. With the existing computational infrastructure available, the start-up costs for a new hire in astronomy will be quite modest, requiring a campus investment in the range of \$100 – 150 K.

Computer Science receives funding from a variety of agencies, including NSF, NASA, DOE, NIH, and DARPA. Of these, NSF is the most significant source. There are several funding opportunities in the area of green computing at the NSF and DOE. DOE has recently announced a new program in energy-efficient computing and communication (50M total funding). NSF is working on a new program on the “Science of Power Management” which directly targets the green computing area. A variety of other program funds various facets of green computing research:

- Data Intensive Computing, which is part of the Cross-Cutting Programs solicitation from the Computer and Information Science and Engineering Directorate. This program has \$10M allocated, to be divided among 10 - 20 \$500K small, \$1.2M medium, and \$3M large, awards. It specifically targets applications of computer science to large scientific datasets.
- High-End Computing University Research Activity (\$10M, divided into 10 - 20 awards of \$500K to \$1M) and the Information Integration and Informatics portion of the Information and Intelligent Systems: Core Programs solicitation (\$90M, 200 awards). These opportunities are particularly focused on computational science.
- The Software and Hardware foundations program and the Multi-core chip design and architecture program directly target various facets of computing systems research related to green computing; these program have funding levels of \$10M and \$6M, respectively.
- The Computing Systems Research and the Networking Technology and Systems programs also encourage applications-oriented research, in addition to more traditional CS areas such as operating systems and networking (both of which are key to green computing).
- Computer science is also well-placed for submitting to the Computing Expeditions program, which seeks to make awards of \$10M over 5 years for major new initiatives. This year, our Expeditions proposal in the area of high performance computing made it into the semi-final review stage, and we anticipate submitting again with a stronger proposal next year.

Hires in computer science have traditionally yielded a good return on investment, with nearly all receiving funding in the range of \$200K within the first two years, and approaching the departmental average of \$3M by the time of tenure. Startup costs for new Computer Science faculty already tend to be low in comparison with lab-science disciplines; typically on the order of \$250K, mostly for RAs and summer support for the first two years. The Computer Science Department is also prepared to cover 1/3 of the startup costs for each of its two hires in this cluster, thus requiring a campus investment of about \$160K, which we anticipate will be more than recovered from overhead return on grants. We expect that a cluster hire in this area may also involve purchase of some additional local cluster hardware to be shared among the new faculty without contention from other users. However, it is now reasonable to build a cluster suitable for developmental work for a fairly modest sum.

Summary and Assessment

This interdisciplinary faculty hiring proposal will create a new thrust in high performance green computing on campus by providing new faculty actively engaged in innovative research in computer science, computational science, and state-of-the-art computational astrophysics. These new hires will join and benefit from other faculty on campus and across the five colleges who are interested in high performance computing. They will amplify the value of prior investments by the University and various granting agencies and the future investment of the State in the Green High Performance Computing Center planned for Holyoke by forming the core of an on-campus computational science center. We also note that there is an important educational component to this plan, as we anticipate that these new hires will contribute to the development of an interdisciplinary green high-performance computing curriculum on campus. The effectiveness of this investment can be assessed by the increase in jointly funded proposals, joint papers and shared curriculum development, and in the longer-term, the expansion of computational science activities on campus.