

Minnesota Supercomputing Institute (MSI)
External Review
April 2007

Jeremy Kepner (MIT Lincoln Laboratory)
John Quackenbush (Harvard / Dana-Farber Cancer Institute)
Dan Reed (Carolina/Renaissance Computing Institute)

Executive Summary

In early April, 2007, we conducted a two day review of MSI. During this time, we met with MSI staff, the institute leadership, advisory committee members, faculty and selected university leaders. The committee appreciates the hospitality and the insightful comments from all who participated.

It is now widely recognized that computing is now the third pillar of scientific discovery, as a complement to theory and experiment. Given the centrality of computing and MSI's history, the MSI is the only university-wide organization that can provide substantial benefit to every major department and research center. However, just as computing technologies have evolved, so have university needs for computing infrastructure and support. Hence, we believe the structure and priorities of MSI must evolve accordingly.

First, the MSI should align its mission clearly with university research priorities and should report directly to the Vice-President for Research. Second, the university should conduct a national search for a director. Third, the MSI committee and oversight structure must be simplified. Finally, the MSI budget is inadequate for projected university needs – additional funding is needed, notably from the Academic Health Center (AHC) if the MSI is to support biomedical research adequately.

1. National Profile and Context

The MSI has a strong and distinguished history in supercomputing. The external research revenue that can be directly attributed to MSI involvement is impressive. The MSI also has a strong base of computing capacity, the staff members are highly professional and their contributions to research support are widely recognized throughout the campus. In addition, the MSI has taken the initiative to reach out to the high growth biomedical research area and plays a key role in enabling collaborations across the University and with outside organizations.

The MSI and the University of Minnesota are uniquely positioned with respect to local industry. Minnesota has the strongest supercomputing industrial base in the world. Exploiting these relationships has been a critical part of MSI's historical success. As recently as ten years ago, the MSI could have made the claim of being one of the Top 5 academic supercomputing centers in the nation. However, the increasing importance of computation to the research enterprise has led many institutions to invest significantly in this field. The result is that although MSI has consistently made progress, the rest of the field has been more aggressive, and the MSI is probably no more than a top 25 center.

The MSI should foster strong relationships with local industry focused on research on core technology bottlenecks, new applications (e.g. medical imaging) for current technologies, testing of new products at large scale, and creating an academic community around these industries to support their needs. Such relationships can create external funding opportunities by helping define the national research agenda and demonstrating strong research transition partnerships,

which are important to many government funding organizations. By exploiting the strong local industrial base, it is a realistic goal for MSI to regain its former leadership position.

Finally, given the importance of the national security mission to the supercomputing industry and to several government funding agencies, the MSI should determine what role it can play to support the research needs of these organizations.

2. User Services and Structure

The MSI operation's can be made significantly more efficient by adopting current best practices. Most importantly, the MSI must align services, use cases and policies to best meet user's needs. This alignment will significantly decrease friction in the system.

The MSI team leads should visit the leading supercomputing centers to adopt current best practices (e.g. integrated incident tracking systems). Negative user experiences should be analyzed to evolve practices and policies to minimize or eliminate these experiences. We believe many can be eliminated when precise user needs are understood in detail by the entire team. The MSI also needs a clear service model that includes concise descriptions of its services for different user classes and a clear set of use cases for satisfying different user needs.

The current 16 committees should be replaced with a single advisory board. A user's primary contact with the MSI should be via well defined interactions (not committees or out-of-band channels). Moreover, MSI policies are a tool by which the director improves user service. Hence, policy and exception handling must be controlled by the director.

Finally, the MSI should develop a "Return on Investment" model to optimize dollar allocations to the most effective activities.

3. Biomedical Engagement

Although high performance computing has traditionally played a large role in engineering and the physical sciences, the completion of the human genome and the technologies it has spawned, together with advances in imaging techniques, are rapidly turning the biomedical sciences into information sciences. There is a clear need for sophisticated computational approaches to data analysis, for high-performance data storage, and for the development of novel data management strategies.

Given both the physical and intellectual resources available in the MSI, supporting biomedical research represents a great opportunity both for advancing the field and for securing non-traditional sources of funding. The National Institutes of Health, the National Science Foundation, the Department of Energy, and the Department of Homeland Security all have initiatives that support computational biology approaches or require them as essential elements of large research projects.

However, meeting the needs of the biomedical research community presents challenges that are distinct from other communities. In particular, the majority of faculty members in biomedical departments lack the level of quantitative and computational training common in MSI's more traditional partners. Consequently, support for research projects in the biomedical sciences requires greater direct involvement in developing the basic computational methods and instantiating them in code from the start, rather than simply porting and optimizing existing code for high-performance computing applications. This requires highly trained individuals who are computationally skilled yet conversant in the biomedical sciences and aware of existing methods and resources in the field.

To its credit, the MSI support staff has recognized the opportunities available in the biomedical sciences and has begun to expand into this area, providing support for a wide range of projects. Most of this has been through purchase of site licenses for commercial software packages and through serving as consultants in data analysis. There are, however, opportunities for more direct involvement in the research enterprise that would include development of new methods for data analysis and modeling, database research, database development, encryption and data security, and problems associated with data integration. The recent biofuels initiative grant application, in which MSI played an integral part, is one such example.

Although there is a growing partnership in biomedical applications with the Mayo Clinic, the MSI should look closer to home. There are significant opportunities open through collaboration with the Academic Health Center (AHC). The AHC's recently rejected Center grant application pointed out their weakness in bioinformatics and computational biology and this is an area where the MSI could provide meaningful support. The MSI administration should look for opportunities to build a strong relationship with the AHC and to partner with them in meeting the challenges they are facing.

The MSI should also look for other opportunities that address its strengths and those of the local high-performance computing and storage industry, including applications in biodefense, biofuels, and in medical diagnostic image processing.

4. Summary

The MSI is critical Minnesota resource. To exploit this resource effectively and fully, the university must restructure and refocus the MSI. This will require additional funding to support burgeoning biomedical research opportunities and coupling with the university's strategic research agenda. Concurrently, the MSI must sharpen its policies and processes to bring them in line with current best practices.

Finally, the committee expresses its appreciation to the MSI and the university for their hospitality during our visit.