High Performance Computing Systems and Scalable Networks for Information Technology

Joint White Paper from the

Department of Computer Science

and the

Department of Electrical and Computer Engineering

With the participation of the

Department of Industrial & Manufacturing Systems Engineering

Submitted to

Center for Applied Information Technology

September 25, 2001

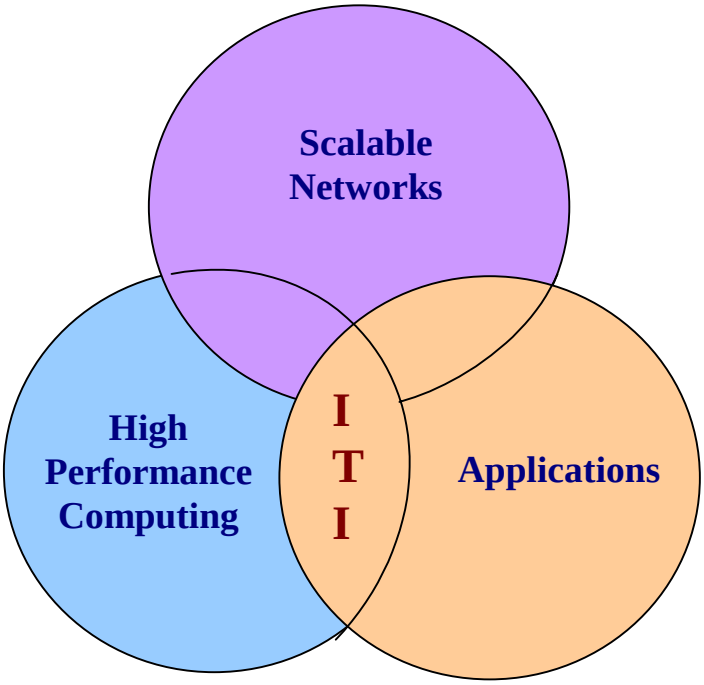

Table of Contents

1. Introduction

The disciplines of computer science and computer engineering form the foundation of information technology. One cannot have a strong program in information technology unless these disciplines are strong. This initiative seeks to strengthen these disciplines at Louisiana State University by integrating the efforts of the Departments of Computer Science and Electrical and Computer Engineering.

As applications become more sophisticated, there is an increased demand for high performance computing systems. As the performance of computing systems increases, the applications become more complex and require still more processing power. High performance creates its own demand. A scalable system provides for incremental increases in performance with incremental increases in computing and networking resources. Major advances will require networked computing systems designed to incorporate the latest research results in computing systems and scalable networks.

This initiative proposes to build upon existing strengths of faculty from the Department of Computer Science, Department of Electrical and Computer Engineering, and Department of Industrial and Manufacturing Systems Engineering to establish an Information Technology Computer Laboratory (ITCL). Academic and research programs in computing systems and information technology at Louisiana State University are the primary responsibility of the Departments of Computer Science and Electrical and Computer Engineering. The two departments currently have strong academic and research programs in high performance computer architecture, networking, parallel compilers, database systems, and intelligent systems. The Department of Industrial & Manufacturing Systems Engineering is strong in information technology applications and industrial outreach. The ITCL Laboratory will focus on scalable high performance computing and networked systems for information technology. This laboratory will form the core of a center of excellence to leverage the expertise of the faculty and enhance the national standing of Louisiana State University in information technology.

The Laboratory will enhance the university's ability to obtain increased federal funding. Laboratory personnel will collaborate with existing industry and assist in the incubation of new industry in information technology. The Laboratory and the associated departments will educate and train engineers and scientists in advanced information technology and provide an industrial workforce. The ITCL laboratory will facilitate cross department faculty collaboration and allow collaborative faculty expertise to utilize shared infrastructure resources to advance the academic and research agenda.

Objectives

The objectives of this initiative are to:

- Foster the development of new research collaborative cross-disciplinary research and academic programs in high performance computing systems and scalable networks. Provide cross-disciplinary infrastructure to promote research and academic programs in the Laboratory's focus areas.
- Develop a new collaborative unit of faculty to enhance academic and research programs in high performance computing systems and scalable networks.
- Provide Louisiana with the well-educated work force in Computer Science and Computer Engineering to meet the future challenges in information technology. Enhance undergraduate and graduate academic programs in the participating departments.
- Provide leadership and resources to move Louisiana State University to prominence and national competitiveness in information technology.

- Promote entrepreneurial enterprises and industrial collaboration in academic and research programs.

2. Scope and Applications

2.1 Processor Design for Scalable Networking

Major advances in processor design will be critical in realizing the sophisticated IT systems and applications of the future. Some advances relate directly to computation, such as processors for multimedia and specialized processors used, among other things, for network routing. A promising approach is to incorporate application-specific instruction sets in these processors. Aggressive compiler and runtime techniques targeting these specialized processors are essential for realizing their performance potential. Novel processor architectures will also be pursued. These include chip multiprocessors that take advantage of ultra-fast on-chip data transfer to tie together several processors on a single chip. Data-intensive IT applications stress the memory system. There are many memory research areas, all concerned with providing fast access to a vast amount of data using memory devices which are continually losing ground to quickly improving processor speeds. Extension of conventional techniques, such as caching and prefetching, and pursuit of novel techniques such as intelligent and software-controlled memory will be investigated to remove this IT bottleneck. Moving large amounts of data to and from disk storage has challenges of its own; architectures and system software for delivering high levels of performance will be investigated.

Power is clearly important for mobile IT devices; it is also important for server processors, as the problem of heat removal has become a design-limiting factor. Techniques such as low-energy computation and design of low-power processors that might include functional units that can be deactivated under program control will be researched. Not only must processors themselves be advanced, but so must the tools for designing them. The design of an individual component alone can be time consuming. IT processors can include several major components that are tightly choreographed; their design is limited by tools meant for simpler systems (tradeoffs not explored, for example). The development of electronic design automation tools meant for such systems-on-a-chip will be pursued.

2.2 Distributed and Networked Systems

The development of networking protocols and architectures is a necessary ingredient in the development of infrastructure supporting the seamless delivery of Information Technology content and related applications. At the most fundamental computing level, there must be substantial improvements in computing system architecture to support the intensive processing, communication, and system resource management demands of heterogeneous IT applications with differing quality of service requirements. Substantial improvements in IT application performance in areas such as distributed data mining can only be seen by the development of higher level networking protocols, real-time and distributed operating systems, resource management algorithms that are effectively co-designed to utilize the underlying architecture in the most efficient manner. Research on the design of scalable networking protocols and architectures is necessary to move from a local view of the computer to a broader understanding of how that computer system fits within a larger network optimized around different goals. Improvements to distributed operating systems are necessary to accommodate a wider array of file types, such as those that accommodate geospatial data, and enhancements to move such distributed systems to a more dynamically configurable setup. The design and evaluation of algorithms which take full advantage of advanced architecture and hardware features is an area that will continue to increase in importance as the disparity between processor speed and memory/network/disk speed grows. Another challenge is delivering IT content is a scalable and reliable system for media-on-demand (MoD) and live streaming (for the

uninterrupted delivery of high quality music and video). Shortcomings with acknowledgement protocols include excessive delay, inefficiency in one-to-many transmission, and massive reception of duplicate packets. Currently, 3GPP addresses mobile streaming standardization. We will extend these concepts and will also work on retrieval and indexing issues related to the streaming-media that has application to multimedia.

2.3 Applications

We envisage several important IT applications to be investigated in the proposed research program. Multimedia networking and its applications in distance learning will be researched. Distance learning requires the delivery of real time, uninterrupted audio/video data over the network and thus presents great challenges to networking. These will be addressed by developing new scalable networking protocols and architectures. Flexible course management software and intelligent tutoring tools will be developed that incorporate knowledge-based techniques for student modeling and automated learning support. Techniques for intelligent human-computer interaction will be pursued, including clustering-based methods for web user modeling, efficient algorithms for natural language/speech interface support in restricted domains, and smart web search tools based on soft computing approaches. Data mining aims at finding novel and useful patterns from large databases commonly used in many real-world applications in engineering, medicine, business and science (e.g., bioinformatics and geoinformatics). Due to its potential to make new discoveries and solve important problems in these domains, advances in data mining will impact many IT applications. Scalability is a critical issue here because of the complexity and huge size of databases. Novel data mining methods will be investigated that fully exploit the high-performance computing environment discussed in (2.1) and (2.2) of this white paper. These include parallel and distributed algorithms for data mining, new data structures for data mining in specialized areas (such as in bioinformatics), hybrid methods combining statistical and artificial intelligence techniques. Mining of the Web data and geospatial/coastal data will be explored. The related areas of database and information retrieval will also be studied. With more and more critical information transmitted over the networks, data security becomes an increasingly important area. We will develop new encryption protocols and supporting language and software environment for better data security. Many defense and civilian applications require the surveillance of a geographical region by a distributed sensor network (DSN) consisting of spatially distributed sensors each with limited sensing capability/coverage. Data fusion from all sensors in a DSN is necessary to provide reliable solutions. We will explore robust sensor field architectures and topologies for optimally deploying sensors such that maximum coverage of a surveillance region is obtained at minimum cost. Combinatorial formulation of this problem and its solution using a variety of optimization algorithms will be investigated.

3. Implementation Plan

The Laboratory will have an executive committee to facilitate and oversee the direction of the Laboratory. This committee will consist of the chairmen of the departments of Computer Science and Electrical and Computer Engineering, four faculty members from the departments, and representation from industry. The executive committee will report on the progress of the laboratory once a year to a university oversight team consisting of the Deans of the Colleges of Basic Science and Engineering, the Director of the Louisiana State University Center for Applied Information Technology, and industrial advisors. The executive committee will coordinate the direction of the academic and research programs of the Laboratory, will promote cross departmental research and academic programs, and monitor progress. The executive committee will conduct annual workshops with external speakers/consultants to promote and report on the progress of the Laboratory. The consultants will provide direction and evaluate the progress of the program.

In years one through four, we will hire four new faculty and eight new graduate students. In the fourth year we will hire three new faculty and six new graduate students. Six of the faculty will be in the processor design for scalable networks area. Three of these positions will be senior positions. Six of the faculty will be in the distributed and networked systems area. Three of these positions will be senior positions. Three of the positions will be in the applications area. One of these will be a senior position. The 30 graduate students will be equally distributed among the three focus areas.

4. Equipment

A total of $400,000 in equipment funds is requested, including $300,000 in computing equipment, $80,000 in software, and $20,000 in specialized hardware.

Equipment for use in all areas of this proposal includes a 12-way multiprocessor, twenty PCs, twenty workstations, and assorted software and network hardware. The 12-way multiprocessor, based on current prices, costs $200,000; the total for the other computing equipment is $100,000. The computing equipment will be used for running electronic design automation software, for running simulations and analyses, and for developing the various applications. Where possible, large simulations and analyses will be run on the proposed 1024-processor cluster. Complementing the cluster, the multiprocessor will be used for running "fine-grained" simulations and analyses for which the cluster is not effective.

Software costs are estimated to be $80,000. This includes support and license fees and purchase prices for design automation, simulation, and software development packages. Specialized equipment is estimated to cost $20,000. This includes power measurement instrumentation, test bed systems, and prototyping equipment and supplies.

5. Summary

Computer Science and Computer Engineering programs must be strengthened in order to advance information technology in Louisiana. Any state that has significantly advanced in technology has developed strong programs in computer science and engineering. Austin, Texas has developed a strong high technology industrial base over the last twenty years while the University of Texas at Austin computer science department has increased to 22 faculty in computer engineering and 37 faculty with 14 lecturers in Computer Science. Currently the computer engineering program at Louisiana State University has 8 faculty and the computer science department has 16 professors. The total number of faculty is 24 and the additional faculty would bring the number to 39. These requests are very reasonable if Louisiana State University desires to advance in information technology.

The Laboratory will enhance information technology. The ITCL Laboratory will have research and academic programs to increase the manpower base of engineers and scientists in these areas. There will be industrial outreach programs to industry to support and promote entrepreneurship and industrial development. The ITCL Laboratory will provide a framework that will assist in the development of information technology based industry in Louisiana.

The Departments of Computer Science and Electrical and Computer Engineering are the departments that perform the primary education of computer scientists and computer engineers in the state. The computer engineering program is the only accredited program in the state. This project will strengthen these programs by increasing the number of faculty and strengthening the infrastructure available to the faculty and students.

The educational outreach program will mirror the interdisciplinary approach of the ITCL Laboratory. The education thrust will be directed toward graduate, undergraduate, K-12, and the current workforce. The Laboratory will recruit and educate students and professionals that can apply computer science and engineering principles to the practice of information technology.

Graduates need to communicate and work effectively in teams and develop interdisciplinary approaches to information technology. The cross department educational effort will involve web based course offerings, graduate and undergraduate students in interdisciplinary research teams, and co-advising of students by faculty from different disciplines.

A consequence of this project would be to improve our competitive position for obtaining research support from Federal agencies as well as attracting the best students and faculty to the university. Our strategy is to utilize existing strengths in our departments to develop a strong research base in scalable high performance computing and networked systems for information technology that can compete for large NSF, DOD, and NIH research grants. Research competitiveness will be greatly increased with the addition of the new faculty and the cross department focus of the ITCL laboratory. This will be manifested in scholarly research and increased research funding from sources such as NSF, DoD, and NIH that all have strong research interests in information technology.

Many of the faculty have strong ties with industry. The new faculty will also be expected to have strong ties with industry. The faculty involvement together with a focus on industry involvement with the ITCL laboratory will increase the formation of information technology industries in Louisiana. The increased research funding and associated research advancements will be an attraction for industrial involvement and development. There will be industrial outreach programs to industry to support and promote entrepreneurship and industrial development. Furthermore, the ITCL Laboratory will provide a framework that will assist in the development of information technology based industry in Louisiana.

Appendix:

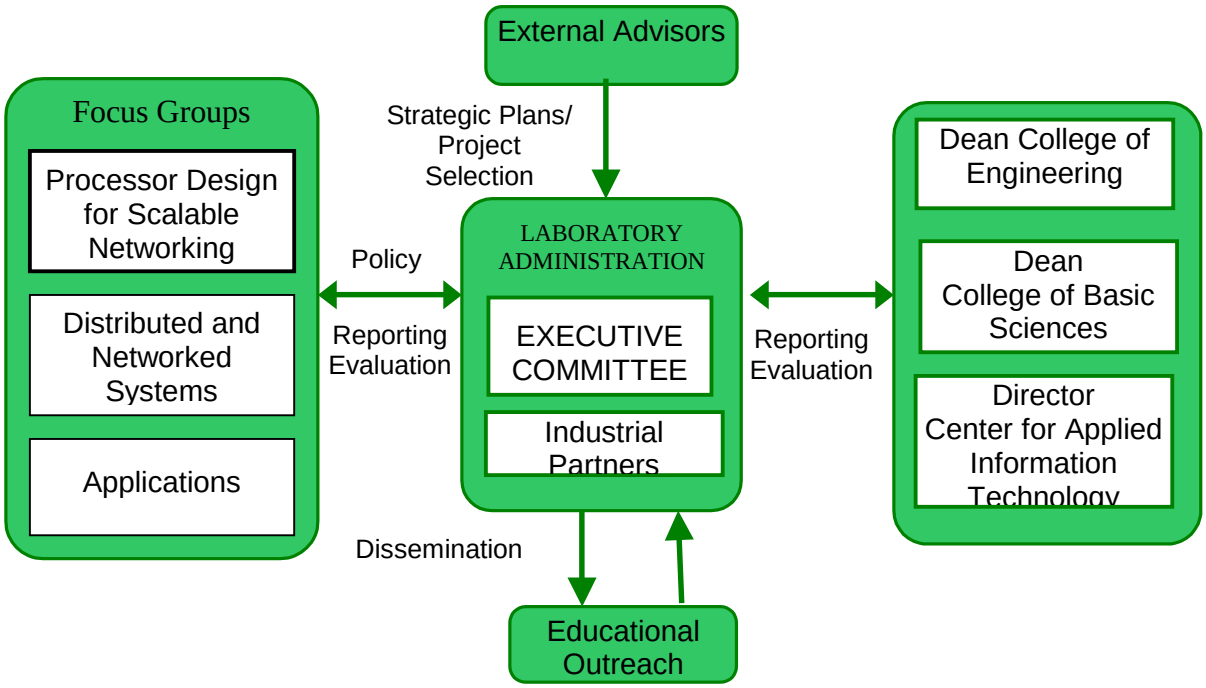

Management Plan