

Homework 6 – Proposals

Due: April 22, 2008 / 3 Ordibehesht, 1387

ECE 667 – Technical Communications in English, Spring 2008

Dr. Soltanianzadeh, Dr. Rahmani

Please review Lecture 13 “Proposals 1” as well as your notes from the two lectures on proposals. Below you will find an actual proposal submitted to the National Science Foundation. Please read it, and answer the following questions. You do not have to read the proposal in depth. Simply reading over its main points should enable you to answer the questions.

1. Please comment on how effectively the summary was written. Make sure you answer the following
 - a. What is the problem statement?
 - b. Was the problem statement described well? Why?
 - c. What is the proposed solution?
 - d. Is the proposed solution described well? Why?

2. In addition to the Summary, this proposal includes 7 sections (A-G). List the 7 sections. For each section, explain
 - a. What is the purpose of this section? What is the main message the author of this proposal is trying to send in this section to the reviewers?
 - b. Why is this section important to the proposal?
 - c. In your opinion, was the section effective? Meaning, what the author able to clearly and convincingly communicate the material necessary for the section? Why?

NATIONAL SCIENCE FOUNDATION GRANT

Integrating Parallel Processing as a Tool Throughout the Undergraduate Computer Science Curriculum

DUE-9552245

William E. Toll, Principal Investigator
Timothy C. Diller, Co-Investigator
Henry D. Voss, Co-Investigator

Most of the text of the grant proposal is included here.

Summary

The laboratory provided by this grant will allow the introduction of parallel processing into additional courses in a Computer Science curriculum where it is already established. Parallel processing will be used as a tool as well as an object of study. In addition, it will be used to support term projects and independent student research. Seven additional courses will include parallel processing bringing the total to 10 (48%) of the department's 21 major courses that include parallel processing. The department has a strong Artificial Intelligence track within the major and has recently begun a Graphics track. The lab provided by this grant will include a SIMD parallel processor which is ideally suited for many of the assignments and projects in areas such as image processing, speech and language processing and neural networks. In addition, a MIMD processor will be acquired which will be used in some of the AI courses, such as Natural Language Processing and Machine Learning, and in Graphics courses. The Computing and System Sciences department views term-long student projects and independent research as an integral part of the curriculum. The parallel processing lab will be used for these projects. Students often undertake research projects as part of an independent study or as part of a summer research program. This program provides funds for student wages while they conduct research projects under the direction of a faculty member. Some of the on-going projects are at a stage where the additional processing power provided by this lab is necessary. The Science Division has a Research Professor who has several projects involving the processing of images from satellite data. Students from other disciplines, such as physics and mathematics, will have access to the parallel processing lab for such research. This project will demonstrate the feasibility of using parallel processing on a large scale in the undergraduate computer science curriculum.

Narrative

- A. CURRENT SITUATION
- B. DEVELOPMENT PLAN
- C. EQUIPMENT
- D. DEPARTMENTAL RESOURCES
- E. FACULTY EXPERTISE
- F. CURRICULAR NEEDS
- G. DISSEMINATION AND EVALUATION

A. CURRENT SITUATION

Institutional and Departmental Context

Taylor University is a private liberal arts college located in Upland, Indiana. The 1994 fall enrollment is 1825. The school has a stable enrollment pattern with the number of applicants more than twice the number that can be accepted.

The Computing and System Sciences department offers both majors and minors in Computer Science as well as a Systems program. The department enrolls about 100 CS majors and approximately 250 students who are receiving a BS in Systems with a variety of majors, including Business and Computer Science. The Computer Science major requires 64 hours and consists of a core of courses together with the choice of one of five tracks: Artificial Intelligence, Business Information Systems, Graphics, Scientific Computing or Integrated. The largest enrollments are in the AI and Business tracks although the relatively new Graphics track (1993) is attracting an increasing number of students. The Artificial Intelligence track in the major has been in existence since 1982 and has been well received by students and employers. The specialized courses are taught by knowledgeable faculty as listed below in the section titled Faculty Expertise.

The tracks primarily affected by this proposal are the AI and Graphics tracks. Assignments in AI courses do not require the use of parallel processing. Both tracks have courses requiring projects which include student research and/or application to solve new, practical problems. Many of these projects are developed in cooperation with industry. Courses that include term-long multi-student projects are Introduction to Artificial Intelligence, Computer Vision, Natural Language Processing, Knowledge Based Systems, Machine Learning, Advanced Computer Graphics and Directed Research. In addition, students often continue to work on projects or start new projects during independent studies. Taylor has also been able to fund a small number of students for summer research projects with the Summer Research Training Program. Many of the research projects involve image processing and/or neural networks.

Appendix 1 contains catalog descriptions and recent enrollments for the courses most affected by this project. Both the courses where parallel processing is currently taught and those where it will be added are listed. Appendix C contains current research projects that would be enhanced by this project. A listing of other recent projects is also included.

The Division of Natural Sciences hired a Research Professor in the fall of 1994. This individual has extensive experience with large grants and contracts from his previous employment with Lockheed and working relationships with such schools as Stanford University and the University of Chicago. He will be funded by continued and new contracts and will be able to provide research projects for several Taylor students in such majors as physics and mathematics as well as computer science. His current research emphasis is evaluation of satellite data that requires image processing.

The Computing and System Sciences Department has 7 full-time faculty plus a technical support person. One of the faculty is on leave each year in industry gaining current experience. The research professor

mentioned above is also included in this proposal although not part of the Computing and System Sciences Department. His primary duties are conducting research, managing research projects and supervising student projects. Three of the CSS faculty have extensive experience managing student research projects. Several have experience managing large projects in industry.

B. DEVELOPMENT PLAN

While many parallel processing curricula emphasize more theoretical issues, the Taylor program also emphasizes practical applications. With the introduction of this project, parallel processing will be used as a tool throughout the curriculum besides being studied as a topic. Because of the present program and experience in AI, graphics and parallel processing, Taylor is in a unique position to demonstrate that a small college can afford and make good use of practical parallel processing facilities. As shown by the list of courses that follows, about 48% of the COS courses will use parallel processing either as a topic of study or as a tool as a result of the funding of this project.

The goals of the proposed parallel processing facility are:

1. to provide adequate parallel hardware for the courses where parallel processing is currently taught - most of these courses teach parallel processing as a topic rather than using it as a tool
2. to allow several additional courses to make use of parallel processing as a tool
3. to enhance the already strong AI program by providing parallel processing tools in the areas of image processing, speech and text processing and neural network use
4. to enable Taylor University CS graduates to have a large amount of actual, practical parallel processing experience
5. to provide adequate equipment to support the strong emphasis on undergraduate research already in place - many of the projects are in need of additional processing power and lend themselves well to a parallel approach

Appendix 2 lists details of the courses where parallel programming projects are currently assigned: Algorithm Design, Computer Organization and Language Structures. Each of these courses will be enhanced by the facilities provided by this project.

A separate section of Appendix 2 lists present course descriptions where parallel assignments will be given based on the availability of this lab. A short summary of these courses and the proposed modification is:

COS 250 Data Structures - parallel sort used with complexity discussion

COS 350 Computer Graphics - modify assignments, such as producing a ray traced image, to use a parallel ray tracer generating more realistic images

COS 351 Computer Vision - parallel image processing on SIMD and option of doing the course project on parallel hardware.

COS 380 Natural Language Processing - SIMD processing for recognition and MIMD processing to have independent agents doing phonological, syntactic and semantic analyses in parallel and communicating the results with each other via a blackboard

COS 423 Advanced Computer Graphics - required independent project on parallel hardware

SYS 411 Machine Learning - parallel processing for neural networks, genetic algorithms, case based learning and decision trees for rule induction - will use both SIMD and MIMD

All of the above courses are taught by Professor Toll or Dr. Diller. Professor Toll has experience in Graphics and in parallel graphics code while Dr. Diller has experience in AI.

Appendix 4 lists current research projects. Some of these projects will be greatly enhanced by the availability of parallel hardware. Most of the research projects in the department are directed by Prof. Toll, Dr. Diller or Dr. White. While Dr. White has no parallel processing experience, he does have experience in image processing research and in managing large corporate research projects.

In addition, Dr. Hank Voss, research professor in the Division of Natural Sciences, has much experience in directing research projects and in image processing and data visualization of satellite data. He is very interested in applying parallel processing techniques to the student projects he directs.

The modifications to the above courses and research projects cannot be done without the funding requested on this proposal. The present parallel environment is not adequate to support any expansion of parallel processing. Taylor has a unique and well-recognized AI program and a growing Graphics program. Taylor has a good history of undergraduate student research. Taylor has introduced parallel processing into the present curriculum. The availability of this lab will allow Taylor to be a real leader not only in the individual areas but in their integration and will enable Taylor to serve as a model of what a small liberal arts school can do with AI, Graphics and Parallel Processing.

The project calls for obtaining two pieces of parallel hardware along with necessary support hardware and software. Two units are requested because some of the proposed projects and assignments require the use of SIMD architecture while others require a MIMD architecture.

The timetable for the project includes obtaining the SIMD machine during the first year of the grant and the MIMD machine during the second year. This plan will require learning one new platform each year rather than two new platforms at the same time. A Sparc host will be acquired during the first year.

C. EQUIPMENT

The goals of the proposed equipment acquisition are stated earlier. To meet these goals, two parallel architectures are needed. Several systems, all in the \$50,000 to \$65,000 range, were investigated. The only SIMD system seriously considered was the CNAPS system from Adaptive Solutions. This system is designed for exactly the types of AI applications needed at Taylor. No other SIMD systems were discovered during the search process that seemed appropriate enough to investigate carefully. Both shared memory and distributed memory MIMD systems were investigated.

The MIMD systems researched were:

1. PowerXplorer distributed memory system from Parsytec. This system runs the Parix operating system and is based on 8 PowerPC chips with transputers for communication links. The system is quite flexible in configuration.
2. Gigacube shared memory system from Microway. This system is based on the i860 and consists of 20 nodes
3. A multi-transputer system from Computer System Architects. This system has a great deal of flexibility in that parts of the network of transputers can be used by different users and the network topology can be configured dynamically rather than requiring re-wiring.

The PowerXplorer has the advantage that the individual processing nodes are very fast and the communication links are user programmable which could lend itself to some data communications projects. The Gigacube has the advantages that it contains more nodes and is a shared memory system. The Gigacube is based on the i860 that is a good graphics processor. The CSA transputer system has the advantages that more processors are available, the processors can be shared and reconfigured easily and that the students and faculty already have transputer experience.

After analysis of the above advantages it was decided the best choice for a MIMD system is the PowerXplorer because of its raw processing power, its well developed and stable PARIX operating system, and the ability for different users to share the processing power during code testing.

Therefore the requested equipment is: (Details and price quotations for the selected systems are found in the budget section.)

1. CNAPS Server II with 512 DSP style SIMD processors. There is 16 Mb of data storage memory. A 68030 processor with 4 Mb of memory controls the interface between a Sparc host and the Server II. Software provided includes a C compiler, a set of libraries, the CodeNet assembly language development system, the BuildNet neural network development system, and a third-party image processing library LNK_ImageLib. The system connects to a controlling host through an ethernet connection. The total price of this system is \$56,300.
2. Parsytec PowerXplorer distributed memory MIMD system. Each of the 8 processor nodes consists of an 80 MHz MPC601 and a 30 MHz T805 Transputer for communication. A C compiler is included. The system runs the PARIX operating system that allows for several systems to be used such as PVM, P4

and Linda. In addition, the system supports multiuser partitioning allowing multiple students to do development work simultaneously. The system is controlled by a Sparc host connected through an S bus card and costs \$58,718.

3. Tatung model 2050B Sparc 20 system with 64 Mb of memory and 1 Gb of disk. This system will be used as the system controller via ethernet for the CNAPS system and the controlling system via S bus interface for the Parsytec system. This system was chosen based on the attractive price relative to an equivalent Sun system and the positive experience the department has had with current Tatung systems. The model chosen was selected because it has sufficient processing power and memory to support large host data sets and multiple simultaneous compiles. The cost of this system is \$10,160.

D. DEPARTMENTAL RESOURCES

The CSS Department has good computing facilities to support its program. The hardware and software available to CSS students is listed in Appendix 1. The equipment which currently supports parallel processing and which will continue to be used includes:

- 4 Sparc workstations
- 8 Sun workstations
- 6 Transputers that can be connected to a Sparc or PC host
- Various parallel simulators and network environments including P4, Linda, PVM, hypercube simulator and parallaxis

Funding for these facilities has been provided by a combination of University special allocations and generous alumni support. The Sun workstations were funded by an NSF grant in 198X that provided facilities to begin the AI program. This previous grant was very important to the success of the AI program.

Research activities in the Science Division as a whole and the Computing and System Sciences Department have been greatly stimulated in the past 5 years by a joint effort between Eli Lilly, Inc. and Taylor University. Additional funds from other sources and a small endowment have allowed the Research Training Program to fund several student projects. Including funds secured by the Research Professor during the current school year, the RTP has had available about \$435,000 which has been used to fund 78 projects involving 90 students and 12 faculty. Forty-two papers have been presented as a result of the work. The CSS Department has been active in the RTP since its inception and has had 3 faculty direct 7 projects involving several students. In the future significant funding will come from the grants and projects directed by the research professor as well as RTP. The RTP funds have been restricted to pay students, small faculty stipends and supplies.

E. FACULTY EXPERTISE

Curriculum vitae of the principal investigators are attached. A short summary of experience relevant to this project is included here:

Dr. Timothy Diller, Professor of Computing and System Sciences and Director of the AI Program: MA and PhD in Linguistics. He has industrial experience in speech and natural language processing, expert system design and project management.

Professor Bill Toll, Associate Professor of Computing and System Sciences: MS in Computer Science, nearing completion of PhD in Computer Science with research in Graphics. His research includes implementation of a parallel algorithm on transputer and Sequent shared memory parallel processors. He has attended an NSF workshop on parallel processing at Illinois State University and has returned to this workshop as an invited speaker. He will present a paper at the 1995 SIGCSE conference on decisions in the introduction of parallel processing into the CS curriculum. He is responsible for implementation of parallel processing in the CS curriculum.

Dr. Hank Voss, Research Professor: Ph. D. in Electrical Engineering. He has considerable experience in space physics and image processing projects while working at Lockheed's Palo Alto Research Laboratory. He has been principal or co-investigator on numerous space projects.

Dr. Arthur White, Associate Professor of Computing and System Sciences: MA in Computer Science and EdD in Biology. He has experience in biologically-oriented image processing research projects and experience in managing and technical development of a large expert system for diagnosis of faults in medical analyzers. Dr. White will be involved with the student research projects.

Other faculty members have industrial or research experience in systems analysis and design, project management, development of large expert system, floating point numbers and compilers, satellite data analysis and geographic information systems.

Based on the above qualifications, the faculty of the Computing and System Sciences Department have the necessary background, experience and interest to effectively use parallel processing hardware in the AI and Graphics curriculum and to direct student research projects. The experience of Dr. Voss will also be valuable in the types of projects envisioned for these facilities. Professor Toll will have the responsibility, along with students experienced in parallel processing, of helping the other faculty learn to use the parallel processing capability provided by this project.

F. CURRICULAR NEEDS

The introduction of parallel processing into the CS curriculum has been quite successful. Parallel processing is currently taught in 3 courses and has been used as a tool in a few research projects. The one exception has been the AI curriculum. Although AI has many natural applications of parallel processing, the courses neither teach applications of parallel techniques nor require the use of parallel processing in assignments, because of the lack of practical hardware.

The overall success of the parallel emphasis can perhaps best be measured by the fact that students are choosing to use parallel processing in course projects and independent research projects. The major current limitation to the use of parallel processing is the lack of true parallel hardware and software. The only parallel hardware available is the 6 transputer network and P4 running on the workstations. The transputers are accessible to only one user at a time and the workstations cannot be made exclusively available to network parallel processing. Real speedup cannot be measured in class assignments because of the use of simulators and the small number of transputers available. Even with these limitations, students have used the transputers for research projects.

The most pressing need in the CS curriculum is for practical parallel hardware and software that will not only function well in courses where parallel processing is currently taught but will allow for the introduction of parallel processing into additional courses and will be adequate and appropriate for the research part of Taylor's curriculum.

G. DISSEMINATION AND EVALUATION

It is very important that results of this grant be disseminated as widely as possible and in as many ways as possible. The bibliography lists many references where parallel processing is demonstrated to be of great value in AI and graphics. In addition many references report on the use of parallel processing from a theoretical perspective. Others report on the integration of parallel processing into many areas of the curriculum. This project will demonstrate that parallel processing can be integrated throughout an undergraduate curriculum in such a way that it is studied both as a topic in itself and as an important tool to solve problems. Furthermore, that knowledge can be applied by the students through the use of parallel processing in various research projects. Such results have not been seen in the literature.

Since equipment such as that requested in this proposal is beginning to be available at a reasonable cost, this project will demonstrate the feasibility of adopting a comprehensive integration of parallel processing into the undergraduate CS curriculum. Further it will show that such integration can be accomplished at a small liberal arts college. Therefore, the target audience for dissemination of the results of this project will be undergraduate CS educators and research project leaders. The following avenues of dissemination will be appropriate.

- Papers will be prepared for submission to undergraduate research conferences based on the use of parallel processing as a tool.

- Papers will be prepared for submission to undergraduate research conferences based on research into the potential use of parallel processing in new areas. An example would be the paper currently being prepared on Parallel Feature-Based Image Metamorphosis.
- Student poster presentations may be prepared for ACM based on the success of student projects. Similar presentations will be submitted to the annual Argonne Research Symposium.
- At least one paper will be prepared and submitted to the annual SIGCSE conference based on the integration of parallel processing into the undergraduate curriculum, the integration of parallel processing into the AI curriculum or the widespread use of parallel processing as a tool.
- Cooperative research projects with industry will be continued. Since the hardware requested is fairly inexpensive, such projects could be practical for adoption in industry.
- Presentations at conferences such as that sponsored by NSF at Illinois State University will continue.
- Participation in an informal consortium of schools interested in undergraduate CS curriculum development and student research has begun. This group of schools is investigating ways to obtain an anonymous FTP file server. Such a resource would be used to share assignments and projects as well as software written using the hardware and software obtained from this grant. It is anticipated that there will be interest in software developed and results obtained from the use of both of the parallel processors requested.
- Consideration will be given to the preparation of a future UFE grant proposal to fund a workshop on using parallel processing as a tool in the undergraduate curriculum. Also considered would be a proposal to fund an undergraduate summer research program to include students from many universities in the research projects conducted using the parallel facilities. As the equipment will arrive during the beginning of the 1995-96 and 1996-97 school years, a target date for funding workshops and external student research projects would be the summer of 1997.

In order to measure the success of a project such as this one and to create credibility for some of the presentations and papers discussed above, an evaluation of the changes caused by this project must be done. One obvious measure of the success of the project is the success of the attempts to disseminate results as listed above. The number of submitted papers and proposed presentations based on research projects that are accepted will be one measure of the success of the use of parallel processing in such areas. The number of accepted papers and presentations based on the curricular issues involved will measure the interest by other universities and colleges in the project and its results. The funding and interest in the possible workshop for undergraduate faculty will also measure this level of interest.

The above evaluation criteria are of an external nature. To prepare some of the above material, internal evaluation must also be conducted. Such evaluation should be done on its own merit. Examples of such evaluations will include:

- the number of students involved with parallel processing in new ways
- the number of students using parallel processing as a tool in course or research projects
- the number of students using parallel processing for class assignments
- the number of students presenting papers, posters, etc. which make use of parallel processing results
- the number of students choosing to use parallel processing when it is not required - this will help to measure student attitude
- the number of faculty and students still using parallel processing 4 years after the beginning of the project
- the effect of the parallel processing emphasis on student recruitment and attitudes of incoming freshmen