

Parallel Computer Needs at Dartmouth College

Parallel Needs (PaN) committee: *

David Kotz (chair)

Fillia Makedon

Matt Bishop

Scot Drysdale

Don Johnson

Takis Metaxas

Dartmouth Computer Science Technical Report PCS-TR92-176

June 29, 1994

Abstract

To determine the need for a parallel computer on campus, a committee of the Graduate Program in Computer Science surveyed selected Dartmouth College faculty and students in December, 1991, and January, 1992. We hope that the information in this report can be used by many groups on campus, including the Computer Science graduate program and DAGS summer institute, Kiewit's NH Supercomputer Initiative, and by numerous researchers hoping to collaborate with people in other disciplines.

We found significant interest in parallel supercomputing on campus. An on-campus parallel supercomputing facility would not only support numerous courses and research projects, but would provide a locus for intellectual activity in parallel computing, encouraging interdisciplinary collaboration. We believe that this report is a first step in that direction.

*All of the committee members can be reached by email at *fullname@Dartmouth.edu* where dots replace spaces in the name, e.g., *David.Kotz@Dartmouth.edu*.

1 Introduction

As part of its expanding research and educational activities in parallel computing, the Graduate Program in Computer Science formed a committee (called PaN, for *Parallel Needs*) in October, 1991, to identify the research and educational needs for parallel computing at Dartmouth and in New Hampshire. This began with an assessment of our own program's needs. It also coincides with the establishment of a new summer institute called DAGS (Dartmouth Institute for Advanced Graduate Studies), which in 1992 will focus on Parallel Computation.

The survey was conducted between December 10, 1991, and January 17, 1992. We used two questionnaires. The first was short and was intended to get a rough idea of who was interested. The second was sent to respondents of the first questionnaire, to get some details. The text of the questionnaires, and the list of people surveyed, are included in the appendices.

We hope that the collected information is useful to all groups on campus. In addition to supporting our program's parallel computing initiatives and the DAGS parallel computing institute, it provides valuable data for Kiewit's efforts regarding the NH Supercomputer Center. It is also a convenient campus reference for people interested in collaborating on parallel computing research and applications, and for industries looking for researchers involved in parallel computing.

1.1 Updates

This report will be updated periodically. Thus, please feel free to contact David Kotz with updated information at any time (particularly if your interests are not represented here). The new information will be included in future versions of the report. In the future, we hope to include other colleges in northern New England (New Hampshire, Vermont, and Maine), so please pass this invitation on to your colleagues there.

2 Responses

There were four classes of respondents:

1. Those who need high performance, but are not interested in a parallel computer *per se* (typically engineers and scientists).
2. Those interested in parallel computing in its own right (typically computer scientists).
3. Those hoping to use parallel computers for education, either at the undergraduate or graduate level.
4. Those who had no interest in parallel computers.

Some people fall in several categories. Each category is discussed below. There is no particular order within each category.

2.1 Performance

Jim Waugh (Engineering): “The arctic ice ocean dynamics lab at Thayer School¹ is involved in numerically modeling the formation and movement of ice throughout the arctic ocean region. The team members require computers with fast floating point performance, as well as 100s of Mbytes

¹Engineering.

of memory (physical and virtual). Parallelism isn't necessarily required, except in that it quickens the throughput of programs."

This project, funded by NSF and ONR, is headed by **Bill Hibler**. Others involved include **Greg Flato** (postdoc); **Jinlun Zhang**, **Cathy Geiger**, **Billy Ip**, **Jay Perry** (graduate students); and **Jim Waugh** (Computer administration and user services).

"The group currently uses an Alliant FX/80, which was selected because optimizations for concurrency and vectorization resulted in a fast computer for its time. I suggest an Alliant FX/2800 or Campus/800. Results are frequently presented using the NCAR² science support graphics software."

Jim Waugh (Engineering): "The numerical methods group at Thayer School is involved in numerically intensive computational research utilizing the Finite Element Method. The team members require computers with fast floating point performance, as well as 100s of Mbytes of memory (physical and virtual). Parallelism isn't necessarily required, except in that it quickens the throughput of programs."

"The group currently utilizes two Silicon Graphics workstations, among others.³ Silicon Graphics offers hosts configured for throughput enhanced via parallelism. Results are frequently presented using Silicon Graphics GL graphics applications programs. The group heavily utilizes the Dartmouth College software package known as *Northware* for visualization, plotting, editing, *etc.*"

This project is funded by NSF, NIH, ONR, Seagrant, and the Whitaker Foundation. The people involved include **Dan Lynch** and **Keith Paulsen** (co-PIs); **Matthias Johnsen**, **Xilin Jia**, **Curt Thalken**, **Chris Naimie**, **Mike Moskowitz**, **Paul Meaney** (graduate students); **Robert Lussky** (undergraduate student); and **Jim Waugh** (Computer administration and user services). (All are at Engineering).

Pete Sandon (Computer Science): Peter is interested in computer vision, image processing and neural networks, and would like fine-grained parallelism. He currently uses a special image-processing system called *Datcube* at the University of Rochester to solve simple vision-based robotics tasks in approximately real time. Although he may someday be interested in architectural issues, currently his "interest is the practical one of getting my algorithms to run fast..."

He would be storing large medical images, requiring about 1/2GB of disk space. He would like to be able to attach frame grabbers, frame buffers, and cameras. C and Unix are preferable.

Barry Fagin (Engineering and Computer Science): Barry is interested in gene sequence analysis, large integer arithmetic, and the investigation of architectures for computationally intensive engineering problems. He collaborates with **Bob Gross** (Biology). He would use a parallel computer for these scientific computing applications, performance analysis, and in computer architecture research.

He is funded by NSF, NIH, Whitaker Foundation, Actel, Xilinx, Altera, and Viewlogic.

He has some old Connection Machine code, but otherwise has no specific preferences. Currently he uses Connection Machines and Crays remotely. C and Unix are preferable. Floating point performance is important.

²National Center for Atmospheric Research in Boulder, CO.

³Three IBM RS6000 model 32H workstations, Silicon Graphics 4D/80GT and 4D/210VGX workstations, and two microVAX workstations running VMS.

Bob Gross (Biology): Bob collaborates with **Barry Fagin** (Engineering) on biological database searching, and enhanced algorithms for structure and pattern prediction. He is interested in collaborating with people on parallel solutions to these problems.

Steve Pieper (Engineering): Steve is interested in interactive computer graphics and real-time updates of complex numerical models that use finite element analysis. These are aspects of surgical simulators for the VA Hospital. The VA funds this work, which is joint with **Joe Rosen** (Medical School). Currently they use workstations.

“I think the MasPar is a good choice for my problems. I’d like to [have support for] parallel finite element analysis and image generation. But, I would want to make sure that each node could have several Kbytes of local memory and floating point support. Also, [I need] high bandwidth communication to external computers.” He also needs a sparse matrix solver, and a general set of numerical support routines. He would like to have support for graphics, and a good programming environment.

Thomas Roos (Biology): Computationally speaking, Tom basically does image processing. He collaborates with CSP Inc. (Billerica), the University of Kentucky, and Bristol-Myers (Buffalo). “I have developed an analytical system that uses [a local special-purpose] array processor as an integral part of its hardware. Data scanned from a electrophoretic gels is entered into the array processor memory and variously massaged to compute and compare image features. Not an awful lot of truly parallel processing goes on, but enough to handle the 1.4 Mbyte image files. I am using the scanner to examine protein and nucleic acid gels for studies in physiology (biochemistry) and evolution.”

John Winn (Chemistry): John does many-body molecular mechanics calculations, which essentially integrate the equations of motion for many atoms interacting by known pair-wise force fields. He plans NSF funding for this work. Currently he uses a Mac Iici. He may be using the Convex and/or Pittsburg Cray sometime next year.

FORTTRAN, graphics output (especially animation), and floating point performance are important. Speed is the most important, with memory and disk space secondary.

Ralph Lewis (Physics): Ralph is interested in plasma simulation — specifically, a variational approach to derivation of time-advance algorithms. He uses computers at NERSC (National Energy Research Supercomputer Center).

Keith Paulsen (Engineering): Keith does large-scale computation in biomedical applications, typically involving numerical solution to PDEs. This is funded by NIH. “I would be a user who might devote some effort to writing parallel code, but would prefer this to be handled by the [compiler and operating system]. We do develop (write) most of our software.” He is clearly interested in performance, as long as it is relatively easy to obtain.

He has huge memory requirements, which is limited in the machines he currently uses (IBM RISC 6000 workstations and Silicon graphics IRIS). Thus high performance, high capacity machines are important. They also need lots of disk space for 20⁺MB output files. They would also need graphics displays. They would want FORTTRAN and C, an operating system that stays out of the way, and graphics and matrix libraries.

Jane Lipson (Chemistry): Jane is interested in simulation studies on polymers. Recently she has been studying structure-property relationships in linear and branched polymers. The level of sophistication of these studies would (could) increase significantly given the kind of facility being proposed here.

She currently uses a DEC 5000/200. She has not used parallel computers because it did not seem like it was worth the effort given the high demand for time on those machines.

She needs an excellent random number generator, and various IMSL subroutines.

Mark Franklin (Engineering, NORTHSTAR): Mark is interested from an administrative standpoint. He supports faculty research at Thayer, NORTHSTAR scientific computing, high-end numerical computing for graduate students, and mechanical CAD for students. Currently they use NORTHSTAR workstations, faculty workstations, Convex, Alliant, VAXen, and remote supercomputers. Faculty research funding sources include IBM, Sun, HP (NORTHSTAR), NSF, DOE, and DOD.

Tom Bubolz (Med School): Tom is not sure how they could use actual hardware, but would be interested in collaborating with other researchers on their database issues. They are interested in converting sequential tasks to parallel tasks in order to speed search and retrieval of medical database information. This would involve the redesign of a database system to parallelize retrieval operations.

“We have an operational database with a data volume of approximately 100GB. Depending on the investigation, we may retrieve information about a person or an entire population of persons. A characteristic of our research is that once a person has been identified for inclusion in a study, we want to retrieve all of their relevant historical information. Currently, this information is neither keyed (as with an inverted list for example) or contiguous. We are looking for database designs and retrieval algorithms to make our processes more efficient. For example, some credit card validation databases are partitioned so that simultaneous parallel searches may be employed.”

They are currently funded by Health Care Financing Administration and the Agency for Health Care Policy and Research. The project involves faculty and staff in the Dept. of Community and Family Medicine. They currently use a network of H-P UNIX workstations.

Benoit Cushman-Roisin (Engineering): Benoit is interested in environmental fluid turbulence: mostly water (oceans and lakes), but also air (planes and jets). Benoit's applications require high-performance computing. He thinks that massively parallel computing will provide him opportunities to perform finer-resolution runs. He collaborates with **David Lynch** (Engineering).

He needs vectorization for most of his codes. His models need little input, and use a workstation for postprocessing the output data. He would need to use NCAR graphics routines, developed on a Cray Y-MP, and routines for matrix multiplication, inverse, root finding, *etc.* He would primarily be writing new programs to solve his particular partial differential equations.

He is currently funded by NSF, Sea-college grant program, ONR. In the future, NOAA, DOE, or EPA are possibilities. He uses an IBM RS/6000 workstation, and a Cray Y-MP at San Diego supercomputer center.

Mary Hudson (Physics): Mary currently uses remote Cray2, Cray XMP, and Convex computers in simulating space plasma and solar-magnetosphere interactions. She works with post-docs **Richard Denton** and **Anthony Chan**, and graduate students **Xinlin Li** and **Victor**

Marchenko (all in Physics). Other relevant people include **James LaBelle** (Physics), **Bill Lotko** (Engineering), **Bengt Sonnerup** (Engineering), and **Charlie Hitchcock** (Engineering).

For a parallel computer, she suggests that a Connection Machine CM-2 with FORTRAN would be useful.

She submitted an NASA proposal last year with Charlie Hitchcock and Bill Lotko to build a special-purpose computer to do high-speed space plasma simulations. She says that apparently “NSF has too much computer time to give away to fund special-purpose computer development.... We are nonetheless interested in parallel processing long term....”

Harold Frost (Engineering): Harold is interested as he may be doing some simulations that require parallel computing.

Horst Richter (Engineering): Horst is interested if it can give him higher performance. “In research and teaching I use mainly large computer codes such as **fluent**, a computer program which allows the study of flows in complicated channels, mixing of streams, *etc.*; or **CEC**, a chemical equilibrium solver for complicated chemical reactions.”

He is currently funded by DOE, EPRI, and the Army. He collaborates with **Paul Queneau** (Engineering) and **Colbeck**. They currently use SUN, IBM-RS6000 and IBM RT workstation. In the past they have found the VAX and CONVEX too slow for the effort.

Eric Hansen (Engineering): Eric is interested in image processing and large (2000x2000) matrix computations, with applications in medical and microscope image processing. Performance (floating-point) is his primary interest in parallel supercomputers (he would like an order of magnitude improvement over his RS/6000), although 3-D graphics support would be nice. He is currently funded by NIH, using on-campus RS/6000 and Convex computers. Software to support FFT, eigenvalue/eigenvector solutions, matrix decomposition, and matrix inversion would be helpful to him.

2.2 Parallel computing

Ken Bogart (Math): Ken is interested in the development of parallel algorithms for combinatorial problems and in combinatorial problems that arise in the design of parallel systems. He is especially interested in algorithmic questions in ordered sets and in developing parallel algorithms for solving them. He is currently supported by ONR, and collaborates with **Ann Trenk**, **Garth Isaak**, and **Larry Langley** (all in Math). No specific needs.

David Kotz (Computer Science): David is interested in computer architecture, operating systems, file systems, I/O, performance evaluation, and programming environments. Most of his previous work has been in file systems for parallel computers. He would prefer a MIMD multiprocessor with a large number of disk drives, to support his research in multiprocessor file systems. He feels that the best examples of this kind of architecture are currently offered by Intel and nCUBE. “Mostly, it is important to me to find a manufacturer that is willing to work with us, give us operating system source code, and so forth.” He expects that his research would be synergistic with applications research.

He has used a BBN GP1000 for prior research, and has remote access on Encore Multimax, BBN TC2000, BBN GP1000, CM2, MasPar, Sequent Symmetry, and Convex computers. He prefers C/Unix environments.

Scot Drysdale (Computer Science):

- Parallel versions of Numerically Controlled (NC) Machining. Peter Su (grad student) has implemented a version of this on the Connection Machine.
- Practical parallel implementations of geometric algorithms (Sweep tree, Voronoi diagram).

Scot is funded by NSF and Ford Motor Company. He would prefer a SIMD massively parallel computer, with C in some version. He has access to a Connection Machine at Thinking Machines Corporation, and to some machines at University of North Carolina.

Fillia Makedon (Computer Science):

- Offline packet-routing algorithms.
- Visualization of parallel algorithms.
- Experiment with parallel VLSI testing and simulation.

She is currently funded by NSF, SRC (Semiconductor Research Corporation), and the Design Automation SIGDA group. She collaborates with **Takis Metaxas** and **Don Johnson** (both in Computer Science).

She thinks that a MasPar would be good for visualization. A Cray or other vector machine would be useful for experiments on VLSI simulation and testing. She would like software for visualization on any parallel machine. She currently uses a Cray and a Maspar remotely.

Matt Bishop (Computer Science): Matt is interested in multilevel security, multilevel integrity, and cryptography. He is interested in possible parallel implementations of multilevel and commercial security and integrity models. He is also interested in the potential for these machines to cracking ciphers and passwords; this would also involve some performance analysis. The metrics would be based on the speed of bit operations rather than on floating point performance, which contrasts with the needs of many other researchers in this report (although many architectures may be good at both).

For the cryptography, either a SIMD or a MIMD machine would work. For the other parts of his research, a MIMD machine is by far more useful. He has no specific vendors in mind.

Currently he uses remote Cray YMP and Cray 2 supercomputers for the cryptography, but has so far not been able to pursue the other research because no suitable machine is available.

He is currently funded by NASA Ames Research Center. A future possibility is Trusted Information Systems.

Dennis Healy (Computer Science): Dennis is interested in parallelism in numerical analysis, scientific computation and visualization, and signal and image processing. He is currently part of the wavelet project, which “involves many subprojects which (I believe) would benefit from certain types of parallel machines”:

- compression/filtering work based on wavelet expansions and statistical decision theory
- simulations of nonlinear [ODE's] (Bloch equations) in MR
- simulations of Doppler ultrasound signal processing
- simulations of limited angle tomography algorithms

“We need something that is on the useful end of the spectrum. I can certainly understand the desire to get something fine grained and trendy for experimentation, but I don’t think our group would benefit much from that.”

They would need Mathematica, Matlab, IMSL scientific libraries, and their own Signal Processing software. They currently use five Decstation 5000s, a Sun3, and a MacIIci (all workstations). They are currently funded by DARPA and AFOSR. Also possibly NIH (pending). Other people involved are

- Math&CS faculty: **Reese Prosser, Jim Driscoll, Dan Rockmore, Tim Olson, and Bill Heller.**
- Radiology (DHMC) faculty: **John Weaver and Yansun Xu.**
- CS graduate students: Joe DeStefano and Sean Moore.
- Math grad students: Joe Gregorio and Harry Chen.
- Engineering grad students: Jian Lu.

Takis Metaxas (Computer Science): Takis is interested in the implementation and testing of parallel algorithms, especially graph algorithms, and the visualization of algorithms. He works with **Don Johnson** and **Fillia Makedon** (both in Computer Science).

He has no specific needs, but knows there is some relevant software already for the MasPar.

Jim Matthews (Computer Science student, and Kiewit employee): “I’m thinking about a system that involves text relevance analysis and information-browsing user interfaces. A parallel machine that allowed more brute-force techniques (such as used on the Connection Machine by Dow Jones) would be interesting to try and compare with less performance-intensive approaches. So this project would want a highly parallel machine, and would not care much about vector or floating point performance.” Presumably I/O would also be important, both to the disk system and to the network. TCP/IP would be important. Unix/C is the preference.

Cliff Walinsky (Computer Science): Cliff and his students are currently researching programming languages for massively-parallel scientific computers. They hope to solicit advice from faculty and students in other scientific disciplines within the college, to broaden their narrow base of programming examples. In return, they expect that their new programming languages will simplify application programming. They currently use a remote Connection Machine CM-2.

2.3 Education

The following faculty would take advantage of a local parallel computer in their classroom teaching:

Fillia Makedon is teaching a special-topics class (CS 187) on CAD (computer-aided design) in Spring 1992 that will consider parallelism. She is also teaching a special-topics visualization course CS 188 in Winter 1992.

David Kotz would be able to use a parallel computer in teaching undergraduate and graduate courses in Operating Systems, Parallel Computing, and Architecture. A MIMD machine would be most useful for these classes.

Matt Bishop would use a parallel machine to demonstrate concurrency to his undergraduate and graduate operating systems classes (CS 58 and 108), and to have the students do some concurrent programming to become familiar with the principles. The use of the parallel computer would be more involved in the graduate course.

Cliff Walinsky has taught a seminar on programming parallel computers. He expects that his new parallel programming language may be useful in similar courses in the future.

Bill Hibler is teaching the course *Applied Mechanics: Dynamics* at Thayer in Winter 1992.

Dan Lynch and **Keith Paulsen** are teaching the courses *Water Resource Planning and Management* and *Computational Methods for Partial Differential Equations*, respectively, at Thayer School in Winter 1992.

Barry Fagin could use a parallel computer for teaching ENGG138/COSC107 (computer architecture).

Thomas Roos plans to teach a course in Spring 1993 that examines programming for his special-purpose array processor and, if available, other parallel processors, in image-processing applications for biochemistry.

Ralph Lewis is teaching a special topics course in Computational Physics in Spring 1992.

Horst Richter uses large computer codes such as **fluent**, a computer program which allows the study of flows in complicated channels, in his teaching.

2.4 No use for parallel computers

We only surveyed people on campus that we heard might be interested. Of course, not everyone was. Steven Venti (Economics), Dick Birnie (Geology), and Robert Cantor (Chemistry) were not interested.

Some surveyed people did not respond (directly or indirectly): William Culp, Robert Ditchfield, Dan A. Longnecker, Joseph Henderson, David Blanchflower, Vijay Gupta, Ulf Osterberg, John Collier, Marcelo Gleiser, Janet Metcalfe, Sam Bent, Don Kreider, and Burt Rosenberg. In many cases, we believe that their lack of response probably has more to do with busy schedules than to lack of interest.

Finally, of course, we have certainly missed some people. If you know someone who is interested in these issues, but not included, please have them contact David Kotz.

3 Summary

There are 25 researchers listing projects above, involving 61 named researchers⁴ (and many unnamed), which would benefit from a parallel computer on campus. Two computer administrators from the Thayer school and Project NORTHSTAR pointed out a general need in their areas. In Math and Computer Science, nine projects are interested specifically in a parallelism, and two are interested in the performance possibilities of a parallel computer. At least eleven courses (some special topics) would be able to use a parallel computer and its applications directly in the course material, not to mention the support of both undergraduate and graduate student research projects,

⁴Students and faculty.

and institutes such as DAGS. Due to the voluntary nature of this survey, there are certainly more research projects and courses that could take advantage of such an on-campus resource.

There is clearly significant interest in parallel supercomputing on campus. An on-campus parallel supercomputing facility would not only support numerous courses and research projects, but would provide a locus for intellectual activity in parallel computing, encouraging interdisciplinary collaboration. We believe that this report is a first step in that direction.

3.1 Current on-campus resources used:

Researchers currently use the NORTHSTAR equipment and Alliant FX/80, IBM RS/6000, DEC 5000, Silicon Graphics, Macintosh II, microVAX, VAX, IBM, Sun Workstations, IBM RT, and Convex computers. Tom Roos uses a special on-campus image processor.

3.2 Current off-campus resources used:

Five researchers use remote Connection Machine CM-2 multiprocessors. Four researchers use remote Cray supercomputers, and one plans future use. Three use remote MasPar computers. One uses a special University of Rochester *Datacube* Image processing system. One uses remote Encore Multimax, BBN TC2000, BBN GP1000, Sequent Symmetry, and Convex computers. Several others use unspecified computers elsewhere.

3.3 Funding sources available:

Current projects are funded by NSF, ONR, NIH, NASA, DOE, DOD, DARPA, AFOSR, ARO, Seagrant, Whitaker Foundation, Actel, Xilinx, Altera, Viewlogic, Veterans Administration, IBM, Sun, Hewlett-Packard, EPRI, Health Care Financing Admin., Agency for Health Care Policy and Research, Sea-college grant program, Ford Motor Company, Semiconductor Research Corporation, and the Design Automation SIGDA Group. Future funding possibilities include NOAA, EPA, and Trusted Information Systems.

3.4 Hardware needs:

There are several applications that need hundreds of megabytes of physical memory, and gigabytes of virtual memory. Many applications need large amounts of disk space (large files of 20MB or more). Combined, the researchers would need disk space of at least tens of gigabytes, and possibly as high as 1 terabyte. In addition, for one researcher, the ideal is a large number of independent disk drives that together provide this large capacity and performance.

To attain the kind of performance these people need, and to compete with their existing on-campus resources, the computer should be able to sustain several hundred MFLOPS; in fact, most of the performance-minded people cited floating-point performance as their primary concern. (Second was the programming simplicity, and third was memory size, disk space, or networking performance).

Jim Waugh suggested an Alliant FX/2800 or CAMPUS/800, as an upgrade to Thayer's Alliant FX/80. Several people have used remote Connection Machine CM-2 SIMD computers, and would thus suggest one of those. Our research has found that the MasPar MP-1 would be a cheaper and easier-to-use alternative; it was suggested by several people. Thus, some people preferred the fine-grain parallelism of SIMD computers. This would also best fit the algorithm and language researchers in the Computer Science program. On the other hand, several researchers preferred MIMD computers (for example, David Kotz, Matt Bishop, and the Alliant proponents). We suspect

that the performance users are more dependent on the quality of the programming environment (to help them attain the rated performance of the machine for their applications) than on architectural details.

High-quality, fast graphics output would be essential for the visualization needs of both the scientific and computer science researchers. High-speed networking to connect the supercomputer to key areas on campus (for example, Thayer, Bradley, Kiewit, and the Medical School) and off-campus (to the rest of New England, at least) is important, as well as powerful workstations for development, pre- and post-processing, and graphics display.

Desired special-purpose equipment includes frame grabbers, frame buffers, disk arrays, and cameras.

3.5 Software needs:

The Unix operating system and the C and FORTRAN languages were the overwhelming preferences. Most manufacturers support these options today. Several researchers required support for matrix operations, IMSL routines, random-number generators, and graphics support such as the NCAR package. PDE and ODE packages would also help. One group needs Mathematica, Matlab, IMSL scientific libraries, and their own Signal Processing software. For all users, a quality programming environment, with good debuggers and performance analysis tools, would be useful. Convenient networking support for graphics, shared file systems, and remote login would be critical. Access to system software source code would boost research in operating systems, and provide benefits back to the manufacturer.

4 Appendix: First Questionnaire

- Your name:
- What is your interest in parallel computing? (Please give just a few key words).
- Briefly, what are your relevant current or planned research/teaching projects?
- What funding agency, or affiliated industry, is involved, if any?
- What Dartmouth people are involved?
- Do you have specific architectural needs, or suggested vendors?
- Do you have specific software needs?
- What computers do you currently use to fill these needs? Do you have access to parallel computers elsewhere?
- Would you like a copy of the report?

5 Appendix: Second questionnaire

- Name:
- Phone:
- Fax:
- Department:
- Email:

5.1 General

- We see two (overlapping) communities: those interested in parallel computers solely for their supercomputer performance, and those interested in parallel computing in its own right. Which are you?
- If we had a new supercomputer (parallel or vector) on campus, what would convince you to use it? Higher performance? Interest in a new architecture? Better programming environment? Special features?
- Do you have any comments on the possibility of a NH Supercomputer Center, based on what you've heard?

5.2 Research

- Short description of project(s) that could use a parallel machine:
- Funding agencies or industry that support this project (might they help us fund the super-computer center?):
- Names of other investigators or institutions in the project:

5.3 Education

- Short description of uses for parallel machine in your courses or other non-research educational activities:

5.4 Hardware

(Answer anything that is relevant; this is here for people who know about these details. Don't worry if you don't know or don't understand.)

- Do you have any particular vendor or architectural preferences?
- Do you have specific performance needs (eg, "1 GFLOPS or more")?
- Roughly how much memory will you need?
- Roughly how much disk space?
- Roughly how much computer time will you use, eg, per month?
- Do you have special hardware needs (disk, frame buffers, data-collection equipment, graphics output, networking, *etc.*)?
- Have you had experience with MasPar, Intel, nCUBE, or VAX9000 computers?
- Is floating-point performance a key concern?

5.5 Software

- What computer languages do you prefer?

- What operating systems do you prefer?

- If we had a new supercomputer (parallel or vector) on campus, would you be writing new programs or porting existing programs? If the latter, what language/system are they coming from?

- Do you need special software support, e.g., libraries for matrix computations?

- Have you written any supercomputer programs that others might find useful? We might ask you more about this.

6 Appendix: List of People Surveyed

Sam Bent
 Richard Birnie
 Matt Bishop
 David Blanchflower
 Kenneth Bogart
 Malcolm Brown
 Tom Bubolz
 Robert Cantor
 John Collier
 William Culp
 Benoit Cushman-Roisin
 Robert Ditchfield
 Jim Driscoll
 Scot Drysdale
 Barry Fagin
 Mark Franklin
 Harold Frost
 Marcelo Gleiser
 Robert Gross
 Vijay Gupta
 Eric Hansen
 Dennis Healy, Jr.
 Joseph Henderson
 Bill Hibler
 Charles Hitchcock
 Nancy Hossfeld
 Mary Hudson
 Charles Hutchinson
 Donald Johnson
 David Kotz
 Don Kreider
 James LaBelle
 David Lagomarsino
 Larry Levine
 Ralph Lewis
 Jane Lipson
 Dan A. Longnecker
 William Lotko
 Daniel Lynch
 Fillia Makedon
 Takis Metaxas
 Janet Metcalfe
 Ulf Osterberg
 Keith Paulsen
 Steve Pieper

Bruce Pipes
 Horst Richter
 Tom Roos
 Joseph Rosen
 Burton Rosenberg
 Peter Sandon
 Punch Taylor
 Steve Venti
 Cliff Walinsky
 James Waugh
 Karen Wetterhahn
 John Winn

And, all graduate students in Computer Science:

Chris Armen
 Deb Banerjee
 Pichet Chintrakulchai
 Preston Crow
 Joe DeStefano
 Perry Fizzano
 Dimitri Kagaris
 Clifton Kussmaul
 Jim Matthews
 Scott McElfresh
 Mel Menezes
 Sean Moore
 Jerry Quinn
 Carrie Quinton
 Larry Raab
 Pushpa K. Rao
 Mark Valence
 John Vanmeter
 Leonard Wisniewski
 Berrin Yanikoglu

All of the above can be reached by email at
fullname@Dartmouth.edu

where dots replace spaces in the name, e.g.,
 David.Kotz@Dartmouth.edu.