


RL-TR-92-41
Final Technical Report
July 1992

AD-A259 627

NORTHEAST PARALLEL ARCHITECTURES CENTER (NPAC)

Syracuse University

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. 6263

DTIC
SELECTE
DEC 21 1992
S E D

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

92-31865



23607

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Rome Laboratory
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700

92 12 17 001

This report has been reviewed by the Rome Laboratory Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RL-TR-92-41 has been reviewed and is approved for publication.

APPROVED:



RAYMOND A. LIUZZI
Project Engineer

FOR THE COMMANDER:



JOHN A. GRANIERO
Chief Scientist for C3

If your address has changed or if you wish to be removed from the Rome Laboratory mailing list, or if the addressee is no longer employed by your organization, please notify RL(C3CA), Griffiss AFB NY 13441-5700. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

NORTHEAST PARALLEL ARCHITECTURES CENTER (NPAC)

Dr Geoffrey Fox

CONTRACTOR: Syracuse University
Contract Number: F30602-88-C-0031
Effective Date of Contract: 4 Jan 88
Contract Expiration Date: 30 Sep 91
Short Title of Work: Northeast Parallel Architectures Center
Period of Work Covered: Jan 89 - Sep 91

Principal Investigator: Dr Geoffrey Fox
Phone: (315) 443-4736

Project Engineer: Raymond A. Liuzzi
Phone: (315) 330-3577

Approved for public release; distribution unlimited.

This research was supported by the Defense Advanced Research Projects Agency of the Department of Defense and was monitored by Raymond A. Liuzzi, RL (C3CA) Griffiss AFB NY 13441-5700, under Contract F30602-88-C-0031.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE July 1992		3. REPORT TYPE AND DATES COVERED Final Jan 89 - Sep 91	
4. TITLE AND SUBTITLE NORTHEAST PARALLEL ARCHITECTURES CENTER (NPAC)				5. FUNDING NUMBERS C - F30602-88-C-0031 PE - 62301E PR - F263 TA - 00 WU - 01	
6. AUTHOR(S) Dr Geoffrey Fox					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Syracuse University Syracuse NY 13244-4100				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency 1400 Wilson Blvd Arlington VA 22209 Rome Laboratory (C3CA) Griffiss AFB NY 13441-5700				10. SPONSORING/MONITORING AGENCY REPORT NUMBER RL-TR-92-41	
11. SUPPLEMENTARY NOTES Rome Laboratory Project Engineer: Raymond A. Liuzzi/C3CA/(315) 330-3577					
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report summarizes the third and fourth year of operation of the Northeast Parallel Architecture Center (NPAC). NPAC promotes and explores advanced computing technology by providing parallel architectures and research support to university, corporate and government researchers nationwide.					
14. SUBJECT TERMS Computers, Parallel, Architecture				15. NUMBER OF PAGES 230	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL		

Contents

Report

Appendix A - User Research Profiles

Appendix B - Selected User Research Bibliography

Appendix C - *An Overview of NPAC*

Appendix D - *Parallel Computing News*

Appendix E - SCCS Bibliography and Abstracts

Appendix F - ACTION Brochure

Accession For	
NTIS CRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

DTIC QUALITY INSPECTED 8

**Rome Laboratory End of Contract F30602-88-C-0031 Report
January 89 - OCTOBER 91**

Executive Summary

This two-year report summarizes the research of the Northeast Parallel Architectures Center (NPAC), which was supported by the Defense Advanced Research Projects Agency of the Department of Defense. This research was monitored by Raymond A. Liuzzi (COES) Griffiss AFB, NY 13441-5700 under Contract F30602-88-C-0031. It encompasses NPAC's Scope of Work during its third and fourth years of full operation and covers other important developments at the center. An interim report was issued in December 1990, which covered NPAC's second full year of operation. This report is divided into the following sections:

Program Overview;
Contract/Scope of Work;
Program Development;
Implementation; and
The Future.

It represents an effort to measure NPAC 's effectiveness during its third and fourth years of operation in the fulfillment of the mission:

The Northeast Parallel Architectures Center promotes and explores advanced computing technology by providing parallel architectures and research support to university, corporate, and government researchers nationwide.

Program Overview

This was a period of rapid growth and development for NPAC. By the end of this contract the center had evolved as NPAC and the Syracuse Center for Computational Science (NPAC/SCCS) to reflect the addition of Geoffrey Fox , his staff of researchers, and the educational center that encompasses NPAC at Syracuse University. Although NPAC still relied on research faculty from other educational institutions, as well as the government; it also drew interest from corporations, such as Xerox, IBM, and many others. Adding and developing staff, software, resources, support and training programs, NPAC grew as a national parallel computing center. This is fully discussed in the Implementation section of this report.

To strengthen the research and development component of the center, NPAC recruited a renowned physicist and computer expert, Geoffrey Fox, from Caltech. He brought with him a profound expertise in computational science, the field of applying computers to scientific applications. Also known for his creation of the hypercube parallel architecture, he came with a proven track record in parallel computing research and development, an area in which NPAC needed greater expertise.

Contract/Scope of Work

NPAC provided a focal point for (a) testing existing parallel architectures, (b) developing new software tools, (c) providing the local and national research community with a rich environment to explore and exploit large-scale parallel computation, and (d) conducting applications research. For a complete listing of the kinds of ongoing research, refer to the User Research Projects section of this report.

The center's four parallel architectures were available to university, corporate, and government researchers around the world, giving priority to DARPA researchers for up to 60% of the computer time. NPAC's training facility offered up to twenty workstations connected via a file server to the parallel computers. Training workshops included those in all of the architectures and languages, as well as in parallel programming techniques. These workshops are detailed in the Implementation section of this report.

Some special software tools were developed to measure the performance of the parallel machines and to aid in software migration. Lending and reference libraries were established and included manuals from all of NPAC's computer vendors. The research consulting staff developed training materials to accompany workshops, and offered comprehensive telephone and email consulting in addition to creating an NPAC User Group.

During this period, under the direction of Geoffrey Fox, NPAC also created the ACTION technology transfer program. The ACTION program is detailed in the Implementation section of this report.

Program Development

In NPAC's third and fourth years, the center continued to build the infrastructure, staff and services planned during the previous two years of its existence. NPAC (1) acquired computer hardware and the related systems infrastructure, (2) recruited staff, (3) and developed an extensive user service group. A variety of hardware and software was also added to the existing parallel computers. NPAC's staff grew with the addition of consulting staff and management, support service personnel and a systems manager. NPAC also implemented a variety of consulting and training and development programs. These are outlined in the Implementation section of this report.

As a service organization, NPAC focused heavily on planning and implementing user services. These included collaborating with the Cornell National Supercomputing Facility on a scientific visualization project and implementing a new NPAC user information database, a necessary tool for both technical and administrative staff, and a new computer usage accounting system.

As part of its general study of parallel computing, NPAC examined IBM's RP3 architecture closely, and through IBM, was able to offer accounts on the RP3. Additionally, with Rome Lab, NPAC evaluated the distributed properties of the Mach operating system.

As the staff grew, NPAC created several mechanisms for in-house staff training, including workshops and informal discussions called staff forums where work in progress was presented to other personnel at the center.

Parallel Computing expert Geoffrey Fox took his position as Director of NPAC effective July 1990. With the assistance of NPAC staff, he designed ACTION, a technology transfer program, and is in the process of designing both a graduate and undergraduate degree program in computational science, the field of applying computers to scientific applications.

Implementation

Systems Work

The increase in the number of systems staff, and a Manager of Systems Engineering allowed the center to upgrade its parallel architectures and focus on beta and pre-beta testing of software from various vendors. In maintaining four different parallel computing architectures, the systems staff continually upgraded existing software and added to NPAC's existing software and hardware.

Networking

Additionally, the systems staff maintained NPAC's file servers and workstations, so that researchers and staff in general enjoyed local area network (LAN) support to link to networks on campus, and through NYSERNet, to the national network, the Internet. This was essential to the center in offering remote access to NPAC's parallel computers to researchers nationwide. NPAC was also able to provide research consulting services over the Internet, because of the excellence with which NPAC's networks were maintained by the systems staff.

The Connection Machine

Work with the Connection Machine included updating NPAC's existing DataVault data storage system and adding a second one. Installing a copy of the National CM Users' Library, cm-lights, NPAC made this software library available to all of its researchers, and in some cases, simplified access to the software through customized routines written by NPAC staff. Eventually, the Connection Machine system software was updated to version 6.0, with many new utilities that make it easier to use, including a new batch system and many enhanced operating system commands. Other software running on the Connection Machine includes the UNIX interface, Ultrix, and the Sun operating system; X Windows, a multiwindow display package; News, a UNIX news group package; and the GNU Emacs and vi editors. Programming languages include CM Fortran, *Lisp, C* and C/Paris.

Software tools developed by NPAC staff to make parallel computing on the CM easier, were installed (detailed in the Technology Transfer section of this report). A CM monitoring system and filter written by

NPAC staff allowed the systems personnel to better schedule CM sequencers and to monitor the system. Since the CM is the center's most popular architecture, time on this machine became quite competitive, and the filter permitted us to allow more equitable access.

The Alliant FX/80

There were many improvements to the Alliant FX/80 (an architecture that has since been discontinued) during this period, as well. The Alliant ran Concentrix, another version of the UNIX operating system, X Windows, News, and the Emacs and vi editors. Languages included versions of Ada, Fortran and C. The Alliant also featured a large software library, ALLUS, to which NPAC continued to add many numerical solver routines. The Alliant's eight Advanced Computational Elements were updated; the operating system was updated; and new versions of FX/Ada and FX/Fortran were installed. Additionally, FX/Spice, an analog circuit simulator, was installed on the Alliant and added to the ALLUS library.

The Encore Multimaxes

The systems staff also continued to improve the Multimax machines from Encore. The Encore 520 (AMAX) runs Parallel Fortran, Fortran 77, Parallel Ada, Franz Lisp, Multilisp, and C. Its operating system is UMAX, which offers a UNIX-like interface, and the UNIX editors GNU Emacs and vi. Additionally, users can access the Encore Parallel Threads Library. The Encore 320 (Mach 1) runs the Mach operating system, and has a UNIX-like interface. It offers Fortran 77, Parallel Ada, and C, runs the GNU Emacs and vi editors and offers a C Threads Library. In addition to adding the Parasight Debugger to the 320, the systems staff continued upgrading the Mach operating system, and staffed a national Mach hotline for new users. With the consulting staff they evaluated a test version of parallel Lisp called Top Level Lisp running on the 520. This version of Lisp was not found to be useful and was not continued.

General Systems Work

In addition to architecture-specific maintenance, installations and work, the NPAC systems staff developed a number of services and simplified administering these systems. These included developing a method to routinely back up the DataVaults; updating vendor

manuals; reporting on the reasons for all system downtime; and implementing an NPAC News Group, which serves as a bulletin board for all NPAC users to communicate about events and issues in parallel computing.

Operations

The operations staff became a vital part of NPAC, because keeping the computers up and running was crucial to the success of the center. During this period they established round-the-clock computer coverage, including weekends.

The operations group wrote an operations manual for each parallel architecture. They developed a system for backing up all computer systems on a regular basis and a library storage area for these tapes. Off-site tape archives were also established, so that in the event of a disaster, back-ups could be retrieved. One of their most important responsibilities was to prevent data loss due to power outages. The operations group was trained to prevent further damage by powering machines off so that a surge of power did not ruin components within the computers.

An NPAC monitoring program written by the systems staff allowed operations staff to do daily routine task and function checks. For example, this tool monitored load averages on computers, disk usage, printer status, and queues and the completion of required/scheduled backups.

The systems and operations staff created a Sun-based program called montool that monitored the NPAC computers via network connection. This tool allowed an operator to query and control many computer functions.

Operations staff also provided periods of stand-alone time for users who needed accurate timing for programs run on particular computers. Stand-alone time was provided routinely, and this also allowed the staff to run routine maintenance checks on the computers.

Research Support

Consulting

During this period the research consulting group grew. NPAC added a manager, as well as consultants with PhDs in fields other than computer science.

Research consulting became available via telephone and on a walk-in basis 9 to 5 weekdays. NPAC was also able to provide consulting through electronic mail. Using documentation sets, and in some cases, training sessions provided by vendors, consultants quickly became educated in using all of NPAC's computer architectures.

Training and Outreach

Since parallel computing is such a new field, and tends to create a natural apprehension in those familiar with serial computing, NPAC focused on training. NPAC presented training workshops and institutes on the use of all of its architectures, languages, operating systems, and in some cases, programming techniques. Workshop attendees were asked to evaluate each training session and often wrote of the high quality of the instructional content and the skills of instructors. NPAC staff also used comments from evaluation forms to continually improve the quality of training at the center.

Support and research staff developed training materials and workshops that ranged in length from a single afternoon to several full days. The Connection Machine Summer Institutes in 1989 and 1990 lasted for two weeks at a time. NPAC staff also traveled to NCSA to give on-site Connection Machine training in October 1990.

Workshops

Workshop topics included those on:

- The Alliant FX/80, which covered UNIX, FX/Ada, Fortran 8x and C. The Alliant workshop also covered using the ALLUS library.
- The Encore Multimax, which covered UNIX, Mach, Fortran 77, Parallel Ada, Franz Lisp, Multilisp, and C. Using the Encore Parallel Threads library was also covered.

- The Connection Machine, which covered UNIX, CM Fortran, *Lisp, C*, C/Paris and X Windows. Using the DataVaults and frame buffers was also covered, as well as programming for a wide variety of applications.

Many workshops concentrated on specific operating systems and languages. For example, the Mach operating system drew a lot of interest, independent of the Multimax, and so NPAC offered several workshops on Mach alone.

Institutes

The Connection Machine Summer Institutes drew participants from around the world. Each Institute featured selected guest lecturers from Thinking Machines Corporation. The first week covered the CM operating system and all the languages available on the CM. The second, special uses of each programming language and specific applications, including:

- special geometries;
- FFT;
- using the CM scientific library;
- N-body problems; and
- CM Fortran compiler optimization.

At the end of each institute, the participants presented some aspect of their research to the rest of the attendees. (See the August 1990 issue of *Parallel Computing News* in Appendix D.)

The Invited Lecture Series

The Invited Lecture Series became part of NPAC's outreach program. Each academic semester NPAC brought leaders in parallel computing research to Syracuse. These included:

- Daniel Hillis of Thinking Machines Corporation;
- Marvin Kalos, Director of the Cornell Theory Center;
- Keith Clarke
- Andre Van Tilborg;
- Harry Jordan;
- Constantine Polychronopolous; and
- Gary Sabot.

Geoffrey Fox also came to Syracuse as an Invited Lecturer while he was still at Caltech.

The NPAC User Group

The research consultants coordinated the NPAC User Group. This group met once or twice each academic semester. Sometimes these meetings were informal discussions and presentations. At times NPAC brought in representatives from companies involved in various aspects of parallel computing. Speakers included Walter Kleinfelder and Herbert Liberman of IBM's Thomas J. Watson Research Center, who presented a talk on the RP3 architecture. Shortly after their talk, RP3 accounts were made available to interested NPAC researchers. Kelly Murray of Top Level, Inc. also gave a talk on Top Level Common Lisp, which NPAC offered for a while on one of its computers.

The Visiting Researcher Program

NPAC's Visiting Researcher Program gave researchers the opportunity to work in residence at the NPAC facility with the staff, often on a one-on-one basis. Visitors included Stephen Vavasis of Cornell University, who was working on a fluid flow problem for Dr. Meng Lean of Xerox. Others included David Richards from the University of Edinburgh, several scientists from Rutgers University working on statistical mechanical problems, and several scholars from Brazil. (See the July 1990 *Parallel Computing News* in Appendix D.) Donald Pflug from Atlantic Research Corporation also visited NPAC for training on the Alliant FX/80.

Educational Presentations

NPAC staff routinely made presentations on various aspects of parallel computing to those in the educational community. These included:

- an NSF-sponsored undergraduate faculty enhancement program organized by Colgate University;
- a presentation to high school teachers at SUNY Oswego; and
- a talk at Supercomputing by University Persons in Education and Research (SUPER!).

NPAC staff also presented papers and/or were represented at conferences throughout the country including Supercomputing '89 and '90. (See the NPAC bibliography in the User Research Projects section and the SCCS bibliography and list of abstracts Appendix B.)

Course Work

NPAC facilities were used in several college courses in the Northeast. Gary Craig, an SU professor in electrical and computer engineering, used the Multimax (Mach machine) for his course in objected-oriented design. Ernest Sibert, an SU professor in the School of Computer and Information Science, taught a course on the Connection Machine every fall semester. SUNY Binghamton connected to NPAC facilities through the Internet for a course in comparative parallel architectures. The New Jersey Institute of Technology taught a course on the theory of parallel and VLSI computing using NPAC facilities through the internet.

Documentation

In addition to the reference manuals and user guides provided by vendors, which the center made available to researchers in the lending and reference libraries, NPAC also created its own documentation. These included:

- **Training Manuals**

NPAC staff wrote these training pieces for each workshop, continually updating them to meet the needs and interests of researchers.

- **Technical Note Series**

NPAC Technical Notes are timely updates on the center's software, hardware, and various research applications. These are issued as new developments occur to keep researchers informed and up-to-date.

- *Parallel Computing News*

Parallel Computing News, NPAC's newsletter, publishes timely news in the field of parallel computing, a training column, scientific papers, and profiles of work in progress every month.

- a Set of NPAC User Guides Under Revision

NPAC completed the first of three user guides on NPAC's computers with a generic "Introduction to UNIX" manual at the beginning.

- *An Overview of NPAC*

This document is an introduction to NPAC with instructions for how to apply for resources. It was mailed out to new and potential researchers as part of the application process. It was also intended as the first document in the user guide series.

- Brochures

- The NPAC brochure is a four-color glossy trifold detailing the services and facilities of NPAC.

- The ACTION brochure is a glossy eight-page document, which introduces ACTION, the center's technology transfer program.

Technology Transfer

When Geoffrey Fox joined NPAC in July of 1990, he and NPAC staff created the ACTION technology transfer program. ACTION is a high-performance outreach program that puts parallel computing to work in industry. ACTION transfers parallel computing technology to industry by evaluating parallel computing potential in commercial applications; identifying applications especially suitable for parallel computation; providing parallel computing training and tools; and assisting in converting and developing existing and new applications to run on parallel computers. Technology transfer partners now include:

- Niagara Mohawk;
- Digital Equipment Corporation;
- IBM;
- General Electric;
- TRW Corporation
- Shell Development Corporation; and
- Xerox.

CM Software Tools

Tools developed by NPAC staff include:

- a utility to write data from the CM to the DataVault;
- a special subprogram to access the FFT routine from the cm-lights software library;

- a CM frame buffer program to produce higher resolution color images; and
- diagnostic utilities for the CM and DataVault.

Alliant software library (ALLUS)

Although many of the free software routines in the Allus software library were available, the NPAC consulting staff customized these routines to better serve the user community.

Research

In addition to the ongoing research conducted by scientists throughout the country on the center's parallel computers, NPAC/SCCS staff participated in several ongoing research projects. They worked on evaluating several versions of parallel Fortran and a project in which they ported algorithms from LAPACK, a mathematical solver in the Alliant ALLUS library, to the Connection Machine. They evaluated these algorithms for accuracy and performance.

In a Fortran 90 project supported by the NSF, Geoffrey Fox with NPAC and SCCS researchers began work on a Fortran 90 project supported by the NSF. They have been exploring elements in CM Fortan and Fortran 90 that comprise portable Fortran precursors. Geoffrey Fox and his colleagues have since written several papers on portable Fortran. (These are detailed in the SCCS list of abstracts in Appendix E.)

Conclusions

During the third and fourth years of operation, NPAC became a national high performance computing center. As described in the Scope of Work, NPAC focused on service and learned an enormous amount about serving not only a local community, but also a community of remote users. As researchers became more well versed in UNIX and widespread use of the Internet became more common, NPAC was able to serve an increasing number of users. NPAC developed effective methods for supporting this community with resource allocation, training, documentation, consulting, and various outreach programs.

Because of its focus on service, NPAC did not concentrate on tool development and technology transfer during this period. In mid-1990, however, it became clear that researchers not only needed training, but they wanted tools to ease the application of parallel computing in their projects. When Geoffrey Fox joined the center, and the ACTION program was created, NPAC began to focus on research and development. These were two strengths needed in effective industrial technology transfer. Fox brought with him the knowledge and experience in parallel computing architecture and tool development that ultimately would interest the corporate world in parallel computing. He also brought with him a wealth of talented computer and application scientists. Six months after Fox became Director of NPAC, the center began to acquire an increasing number of corporate partners.

The Future

NPAC has now redirected its mission to focus on computational science education and technology transfer as defined by this new mission statement:

NPAC and SCCS support computational science research and education using large-scale parallel computing technology with an interdisciplinary collaboration among university, corporate and government researchers.

With Syracuse University, SCCS is now developing both graduate and undergraduate computational science curricula, which link scientific applications to computer science. SCCS plans to offer degrees in this field and continuing education courses for people in industry.

As part of the Center for Research on Parallel Computation (CRPC), an NSF parallel computing consortium, NPAC/SCCS will be collaborating with researchers in all fields. The goal is to advance computer technology, computer science, and computational science education through NPAC/SCCS's membership in CRPC, and to offer the advantages of membership to corporate participants in ACTION.

NPAC/SCCS will also continue to work on projects with industries involved in the ACTION program, developing industrial parallel computing applications, while at the same time, creating tools to ease the transition from serial to parallel computing. Current projects include those in:

- data fusion;
- signal processing;
- energy;
- benchmarking;
- distributed computing; and
- compiler development.

Appendix A - User Research Profiles

NPAC has had a growing group of researchers, scientists, engineers, students, and faculty using the parallel computing resources the center provides. The resource allocation committee grants computer time to researchers pursuing applications in richly diverse scientific areas including:

- computational fluid dynamics;
- neural nets;
- high energy physics;
- vision and image synthesis;
- medical systems modeling and
- molecular dynamics.

Scientists have been able to do computer modeling and simulations in these and other applications. The development of new parallel algorithms coupled with the speed of parallel computers has allowed them to run some simulations that would have been impossible on serial computers.

Brief summaries of each research project granted time on NPAC computers are compiled by the staff. Each project is categorized and described in terms of the field of study, computer(s) used, algorithms, and software used. These project profiles are stored in a database accessible to NPAC research support and administrative staff to facilitate interchanges among current and potential users. A complete set of User Research Profiles follows.

TITLE: Design, Evaluation, and Tuning of Application Programs on Parallel Architectures.

Principal Investigator (s):
Hanny H. Ammar, Clarkson University

Contact Person:
Hanny H. Ammar
336 Clarkson Hall
Potsdam, NY 13676
315-268-3805
FT00@clutx.clarkson.edu

Abstract of Research

The decreasing cost of Supermini multiprocessor systems has made parallel processing technology available to a wide range of institutions, companies, and government agencies. Tuning application programs to the underlying parallel architecture is essential to exploit the full advantage of parallel processing. In this project, we propose the development of a trace driven simulator that facilitates the tuning process of application programs. It also helps in evaluating the performance of the architecture and the detection of system bottlenecks which aids in upgrading the system by adding more memory or more CPU's. It is planned apply this technique to image processing algorithms on the Connection Machine and to numerical algorithms on the Encore Multimax.

Scientific Field of Study:	Tools for Parallel Program Development. (NSF: 345; CR: D1)
Architectures/Languages:	Encore Multimax and Connection Machine
Algorithm(s) used:	Trace driven simulator of a BUS oriented multiprocessor architecture.
Computational Techniques:	Parallel programming, graphics interface.
Goal of the Research:	To provide a nice user interface for programming that is also efficient.

TITLE: Analysis of Small Radomes of Arbitrary Shape.

Principal Investigator (s):
Ercument Arvas, Syracuse University

Contact Person:
Ercument Arvas
111 Link Hall
Syracuse , NY 13244
315-443-4430
earvas@sunrise

Abstract of Research

The purpose of this research project is to study the electromagnetic behavior of a small radome of arbitrary shape. The radome is composed of one or more layers of dielectric shells. A small radome has its maximum linear dimension smaller than the wavelength of the operation. The radome may have bends and edges such that the radius of curvature may be much smaller than the wavelength. The radome contains an antenna of arbitrary shape. The antenna is modelled as a metallic body. The radome may also be supported by a perfectly conducting body of arbitrary shape. This body is called "the reflector". Given the excitation, we are computing the field transmitted through the radome which may be backed by the reflector. The problem is being solved using the surface equivalence principle to obtain a set of coupled integral equations for the unknown surface equivalent electric and/or magnetic currents. These currents are the sources of the secondary field which exists because of the presence of the radome and/or the reflector. The integral equations are being numerically solved using the method of moments.

Scientific Field of Study:	Electromagnetics (NSF 631; CR J.2)
Architectures/Languages:	Fortran on the Alliant.
Algorithm(s) used:	Solving integral equations using the method of moments.
Computational Techniques:	Using MIMD concurrency to find subsolutions to quadrants of 2500×2500 matrices, using vector operations to speed up standard array operations like inner product.
Goal of the Research:	To make better weather protectors for radar antennas.

TITLE: Modeling a Small Neural Network

Principal Investigator (s):
Robert B. Barlow, Jr., Syracuse University

Contact Person:
Robert B. Barlow Jr.
Merrill Lane
Syracuse, NY 13244-5290
315-443-4164

Abstract of Research

This project is focused on applying the capabilities of parallel architectures toward understanding how the brain works. In particular, we are studying the function of the relative simple neural network of the retina of the horseshoe crab. This network has about 1000 receptor units that integrate visual information with a combination of excitatory and inhibitory mechanisms. It is the only neural network for which there exists a quantitative formulation based on physiological measurements. The Connection Machine permits us to model the time-dependent properties of the eye in real or near real time and thus compare the model results with direct physiological recordings. Computer simulation allows us to test the consequences of the brain's modulation of the retinal function. The retina changes state from day to night, and we are anxious to learn the relative importance of each component of the state change. Our progress to date has been substantial. We have a working model written in *Lisp, and have made many successful simulation runs. We are now in the process of attempting to improve performance of the model.

Scientific Field of Study:	Vision (NSF 449; CR J.3)
Architectures/Languages:	Connection Machine, *Lisp.
Algorithm(s) used:	Modeling a physical neural network with linear equations
Computational Techniques:	Mapping receptor units to processors, using NEWS communication to model interaction in the inhibitory field
Goal of the Research:	Understanding how the brain controls vision, both in horseshoe crabs and in humans.

TITLE: Implementations of Mathematical Algorithms on the Connection Machine.

Principal Investigator (s):
Marjory Baruch, Syracuse University
Nancy McCracken, Syracuse University

Contact Person:
Dr. Nancy McCracken
111 College Place
Rm 3-201 CST
Syracuse, NY 13244
315-443-1722
njm@nova.npac.syr.edu

Abstract of Research

Applied mathematicians are redesigning algorithms for classical problems to take advantage of multiple processors. Experiments in probability that have been too time consuming on serial machines appear readily adaptable to parallel machines. We plan to implement these algorithms on the Connection Machine.

The first class of probability problems that we have worked on can be cast as a cyclic cell automaton, where each cell uses the values of its eight neighbors to determine its own value at the next cycle. One of these algorithms has been implemented on the Connection Machine and its results shown on the frame buffer. Visualizing the results as the automaton converges has been quite helpful to researchers in probability theory. They are formulating other automata, which will also be implemented.

Scientific Field of Study:	Probability theory and automata theory (NSF 118; CR I.4)
Architectures/Languages:	Connection Machine/ C*, *Lisp
Algorithm(s) used:	Cyclic cell automata
Computational Techniques:	Grid communication
Goal of the Research:	Using parallel computing and visualization to assist in theoretical mathematics.

TITLE: Current Languages for Parallel Processing

Principal Investigator (s):
Milissa M. Benincasa

Contact Person:
Milissa M. Benincasa
RADC/COEE
Griffiss AFB, NY 13441
315-330-4476
benincasam@ionex.radc.af.mil

Abstract of Research

This research effort, entitled "Current Languages for Parallel Processing," will investigate the strengths and weaknesses of programming languages available for developing applications on parallel architectures. In order to evaluate these languages a small Command and Control application will be implemented. The chosen problem will calculate the electromagnetic spectrum of a battlefield using a terrain masking algorithm. This application will be written in each of the following languages: C, Ada, Fortran, and Joyce using the Encore Multimax as the target parallel architecture.

The purpose of this effort is to analyze the constructs for expressing parallelism in each of the above identified languages and to determine programming problems associated with implementing Command and Control algorithms. The result will be incorporated into a RADC technical report which will present a brief overview of each of the languages, the parallel constructs available in each language, and analysis of the constructs and their suitability to Command and Control applications,

Scientific Field of Study: Software Systems: Programming Languages (NSF: 314)

Architectures/Languages: Encore Multimax 520; Fortran, C, Ada, and Joyce

Algorithm(s) used: Terrain masking

Computational Techniques: Use the terrain masking algorithm to calculate the electromagnetic spectrum of a battlefield using new parallel constructs in C, Ada, Fortran, and Joyce on a MIMD (Encore Multimax 520) computer; shared memory and message passing techniques.

Goal of the Research:

Analysis of new parallel constructs implemented in C, Ada, Fortran and Joyce to determine programming problems associated with implementing Command and Control algorithms. Results of this study will be used to explore their utility for present and future work on Command and Control applications.

TITLE: Three-dimensional reconstruction, manipulation, and display of the brain to study the brain's control of blood pressure

Principal Investigator (s):
Scott B. Berger, Cornell University

Contact Person:
Scott B. Berger
411 E. 69th Street
New York, NY 10021
212-472-5594
BERGER@ROCKVAX.BITNET

Abstract of Research

For over 30 years, we have studied the brain's control of blood pressure through a wide variety of experimental approaches, including anatomical, physiological and biochemical methods. One of our image analysis tools is a facility for three-dimensional reconstruction of the brain, from serial histological sections, or from any other imaging modality which yields 2D consecutive sections (CT, MRI, PET, etc.). Our projects require a true volumetric representation, which is a formidable problem for computation and storage, and for which we designed a system based on a hierarchical data structure. This design is suitable for large scale parallel machines, and we have written a set of *Lisp algorithms which perform 3D reconstruction, manipulation and display of ultra-high resolution data sets.

Scientific Field of Study:	Image Processing (NSF 343; CR I.4)
Architectures/Languages:	Connection Machine/ *Lisp
Algorithm(s) used:	Constructing 3D images from 2D consecutive sections
Computational Techniques:	Hierarchical data structure on a massively parallel machine.
Goal of the Research:	Improve understanding of the functioning of the human brain.

TITLE: Parallel Algorithms for Image Processing and Image Synthesis/Parallel Algorithm for the Network Flow Problem

Principal Investigator (s):
George A. Betzos, Syracuse University
Ernest Sibert, Syracuse University

Contact Person:
George A. Betzos
4-116 Center for Science and
Technology
Syracuse, NY 13244
315-443-1026
betzos@amax.npac.syr.edu

Abstract of Research

In this project, we are considering parallel algorithms for realistic image synthesis. Various data parallel algorithms for hidden surface elimination and ray tracing have been implemented on the Connection Machine. These include intersection calculations between rays and various objects, realistic shadow generation, anti-aliasing based on oversampling and 5-stage ray tracing. We have also implemented image resolution algorithms for images up to 1024 by 1024 pixels. All these programs exhibit a significant speedup against sequential implementations. We are currently developing algorithms to handle transparency effects and we are adding a new acceleration scheme which should greatly improve the performance of the existing programs. Furthermore, the combination of ray tracing and the radiosity approach for the generation of highly realistic images will be implemented using MIMD techniques on the Multimax and FX/80.

Scientific Field of Study:	Computer Graphics (NSF:343; CR:I.3).
Architectures/Languages:	C* on the CM, C on the Multimax, C on the FX/80.
Algorithm(s) used:	Image enhancement algorithms, hidden surface elimination, ray tracing, shadow generation.
Computational Techniques:	SIMD representation of images, MIMD implementation of ray tracing.
Goal of the Research:	To find efficient methods for realistic image synthesis.

TITLE: Evaluation of Algorithms for Parallel Computation of Aerodynamic Flows.

Principal Investigator (s):
Edward A. Bogucz, Syracuse University
Frederic A. Lyman, Syracuse University

Contact Person:
MEdward A. Bogucz
139 Link Hall
Syracuse, NY 13244
315-443-3177
bogucz@suvvm.bitnet
eabogucz@sunrise

Abstract of Research

External aerodynamic flows at large Reynolds numbers often are modeled in terms of two separate regions: an inviscid, irrotational flow away from the body, and a thin viscous boundary layer adjacent to the body surface. In this subdivision of the flow field, a different set of equations governs the flow in each region. The solution in each flow region must be matched through appropriate boundary conditions to obtain a composite representation of the entire flow field. Practical calculation procedures based on this approach determine a solution of each flow field separately, and iterate until a consistent solution in both regions is obtained. The first part of this research is to obtain a more accurate description of inviscid flow fields with discrete vortices by developing a parallel algorithm that use Lagrangian vortex methods rather than the usual Eulerian grid. The second part of this research will exploit concurrent solutions to the free-stream and boundary layer regions, and develop new coupling techniques that will improve the overall solution for flows in which the boundary layer has a strong effect.

Scientific Field of Study:	Computational Fluid Dynamics. (NSF: 614; CR: J2)
Architectures/Languages:	Alliant/ Fortran, Connection Machine/CM Fortran.
Algorithm(s) used:	Algorithms for Lagrangian vortex methods, iteration of solutions for flow-field equations.
Computational Techniques:	Solution of different regions of flow equations in parallel.
Goal of the Research:	Solve aerodynamic problems for devices with external vortex flows, like helicopters, turbomachines, and downstream blade rows.

TITLE: Parallel VLSI Circuit Verification and Simulation

Principal Investigator (s):

Randal E. Bryant, Carnegie Mellon
Saul Kravitz, Carnegie Mellon
Derek Beatty, Carnegie Mellon

Contact Person:

Dr. Randal E. Bryant
Computer Science Department
Pittsburgh, PA 15213

Abstract of Research

The COSMOS project seeks to achieve new levels of performance in simulating and formally verifying VLSI circuits designed in MOS technology. One of our aims is to exploit parallelism to achieve higher simulation performance to allow full chip simulations of state of the art VLSI systems. We have implemented two versions of conventional simulation, using both data parallelism, by simulating circuit operations over a number of input sequences independently, and circuit parallelism, by evaluating circuit elements in parallel while simulating a single input sequence. Both of these algorithms obtain speeds up to 2 orders of magnitude faster than on conventional machines.

In addition, our project has developed symbolic simulation methods by representing input patterns with Boolean variables and nodes as Boolean functions of these variables. Our Boolean manipulation algorithms create and manipulate complex graph data structures much like connectionist networks, representing the Boolean functions as acyclic graphs. During the course of a symbolic simulation run on the Connection Machine, we may construct thousands of these graphs, each ranging up to 10,000 or more vertices. Further research is required to determine the most effective ways to exploit the many sources of potential parallelism in these algorithms. Ultimately, we would like to integrate the Boolean manipulation code with our existing simulation code to implement a symbolic simulator.

Scientific Field of Study: VLSI Design (NSF 632; CR B.7)

Architectures/Languages: Connection Machine / C/Paris, C*, and *Lisp

Algorithm(s) used: Circuit simulation, circuit verification.

Computational Techniques: Data parallelism, graph data structures.

Goal of the Research: To improve the design process and the reliability of VLSI circuits.

TITLE: The SISAL Project: Machine Access Proposal

Principal Investigator (s):
David Cann

Contact Person:
David Cann
Lawrence Livermore Nat'l Lab
P.O. Box 808
Livermore, CA 94550
415-423-7875
cann@antonio.llnl.gov

Abstract of Research

This proposal requests access to the parallel computers at the Northeast Parallel Architectures Center. We wish to use these machines to implement, test, and evaluate SISAL 2.0, a functional language for parallel numerical computation. The goal of the SISAL project is to explore methods for automatically exploiting parallel computer systems. The project itself is a collaborative effort between Lawrence Livermore National Laboratory, Colorado State University, and the University of Manchester.

Scientific Field of Study:	NSF: 314; CR: D3
Architectures/Languages:	Alliant FX/80 and the Encore Multimax
Algorithm(s) used:	Write an automatically parallelizing compiler for a purely functional language
Computational Techniques:	Scheduling parallel expression evaluation concurrently
Goal of the Research:	To make parallel programming easy and natural.

TITLE: Solution of Partial Differential Equations via the Finite Element Method - Connection Machine

Principal Investigator (s):
Steven P. Castillo, New Mexico State
Edward Hensel, New Mexico State

Contact Person:
Dr. Steven P. Castillo
Box 30001
Las Cruces, NM 88003-0001
505-646-3115

Abstract of Research

Currently, research is being conducted into the finite element solution of Laplace's equation, Poisson's equation, the wave equation, and the reduced wave equation. The solution of these differential equations and other large practical 2-D and 3-D problems is limited in size by time and memory constraints of current conventional supercomputers such as the CRAY XMP. We are implementing existing finite element code used for the solution of partial differential equations on the Connection Machine, and addressing the issues:

1. How easy is it to convert or write code for the Connection Machine?
2. How fast is the Connection Machine for solving large problems in relation to conventional supercomputers such as the CRAY XMP?
3. How fast is the Connection Machine in comparison to other parallel architectures for solving large problems?

Scientific Field of Study: Computational Engineering.
(NSF: 342; CR: J2)

Architectures/Languages: Connection Machine.

Algorithm(s) used: Finite elements method for solving linear partial differential equations.

Computational Techniques: Large matrices and operations on a SIMD machine.

Goal of the Research: To improve the performance of foundational algorithms used in many scientific disciplines.

TITLE: Evaluation of Matrix Solution Methods Using the Connection Machine

Principal Investigator (s):
Roger Chen, Syracuse University
Kirk Friedman, Syracuse University

Contact Person:
Kirk Friedman
118 North Street
Syracuse, NY 13206

Abstract of Research

A system of linear equations can take the form $Ax = b$ can be solved by a computer by several different methods. Two of the methods commonly used are Gauss-Seidel and Jacobian iterative methods. Each relies on iteration in which an element in the matrix is updated based on the elements surrounding it. Parameters which may effect the performance of each of these methods include the size of the matrix, the amount the matrix is changed from its original state before the solution method is applied, the bandwidth of the matrix and the sparsity of the matrix (how many elements it contains). Currently there is no data to form a set of tradeoff points where it would be beneficial for a researcher to select one method over another for solving a large series of equations.

In this project, we are implementing the two solution methods on the Connection Machine. Very large matrices on the order of 1000 by 1000 (and hundreds of thousands of elements) will be generated. These matrices will be changed by some percentage and each method will be used to arrive at the new matrix solution. The tradeoff points between each algorithm's performance will be determined initially based on matrix size and the percentage of the change. Repetitions will be performed in order to guarantee accurate results.

Scientific Field of Study: Computational Mathematics
(NSF 118; CR G.2)

Architectures/Languages: Connection Machine/ CM Fortran

Algorithm(s) used: Gauss-Seidel, Jacobian iteration.

Computational Techniques: Grid representation of arrays, nearest neighbor communication.

Goal of the Research: To provide mathematical tools and a method for deciding which tools will be most effective for a particular scientific problem.

TITLE: Forest's Fires, Lattice Dynamical Systems, Spatio-temporal Chaos and the Heart

Principal Investigator (s):
Dante R. Chialvo

Contact Person:
Dante R. Chialvo
766 Irving Ave.
Syracuse, NY 13210
315-464-7979
dchialvo@suvvm.bitnet

Abstract of Research

The mechanisms of life-threatening cardiac arrhythmias like ventricular tachycardia and ventricular fibrillation are poorly understood. The unrealistic simplicity of cellular automata's or the extreme complexity of the partial differential equation's models currently available of the heart's electrical dynamics is a major drawback to studying the problem theoretically. Using a forest's fire description as a caricature of cardiac cell dynamics, in the context of coupled map lattices, propagation of normal and abnormal waves in the heart will be studied. Numerical experiments in 2 and 3-dimensional excitable media simulated using that technique will be conducted in an attempt to understand the role of known rate-dependent electrical properties of heart tissue in inducing spatial disorder. Specifically, the relevance of the rate-related changes in conduction velocity and in the duration of the propagated activity will be investigated.

Scientific Field of Study:	Medical science (NSF 422; CR J.3)
Architectures/Languages:	Connection Machine 2
Algorithm(s) used:	Cardiac cell simulation in lattice dynamical systems
Computational Techniques:	Grid techniques
Goal of the Research:	To understand heart arrhythmias

TITLE: Parallel Computer Vision

Principal Investigator (s):
Alok N. Choudhary

Contact Person:
Alok N. Choudhary
CST, Rm 2-127
Syracuse University
Syracuse, NY 13244
315-443-4280
choudhar@fruit.ece.syr.edu

Abstract of Research

Computer Vision has been regarded as one of the most complex and computationally intensive problems. The algorithms involved employ a very broad spectrum of techniques from several areas, such as signal and image processing, advanced mathematics, graph theory, and artificial intelligence. Computer vision algorithms are normally divided into three levels: low level, intermediate level and high level. Low level algorithms involve pixel data, are structurally regular, and perform data independent and local computations. Intermediate level algorithms involve more complex data structures, are data dependent, and perform symbolic processing. High level algorithms possess all the properties of intermediate level algorithms and involve using of knowledge bases for interpretation. These algorithms use top down processing which involves accessing databases, and performing artificial intelligence operations. A Computer Vision System (CVS) involves algorithms from all levels, and several algorithms with vastly different characteristics need to be employed in a sequence where input of an algorithm is the output of the previous algorithm in the sequence.

Parallel processing has been accepted as the approach to providing necessary computational power to implement CVSs. However, success in exploiting parallelism for CVSs has been limited due to several reasons. First, most multiprocessor architectures are either very rigid (Such as an array of SIMD machines) or are designed for a specific set of applications, and therefore, are not suitable for complete CVSs. Second, mapping techniques for intermediate and high level algorithms are not well developed. Finally, techniques to develop parallel programs have lagged far behind other advances. In this research, we propose to investigate all three issues because they are most critical to make any significant advances in exploiting parallel processing for CVSs.

Scientific Field of Study: Computer Vision (NSF:322; CR: I4)

Architectures/Languages: CM-2, Encore Multimax 520 & 330, Alliant FX/80

Algorithm(s) used:

Convolutions, Edge detection

**Computational
Techniques:**

Divide & Conquer

Goal of the Research:

To develop computer systems that can see
and understand scenes.

TITLE: Parallel Algorithms for VLSI Testability Analysis.

Principal Investigator (s):
Gary L. Craig, Syracuse University

Contact Person:
Dr. Gary L. Craig
111 Link Hall
Syracuse, NY 13244-1240
315-443-4389
craig@sutcase

Abstract of Research

One of the testing challenges of static CMOS logic networks is the detection of stuck-open faults which cause combinational circuits to exhibit sequential behavior. Testing schemes for these faults require multiple-pattern tests. In an attempt to solve the problems of previous test generation algorithms, a minimum length robust multiple-pattern test generation/ testability analysis procedure has been developed. This process involves creating a database for the network of representative cubes for each circuit node and then generating a test pattern for every "fault" in the network. Both of these steps are well suited for parallel computation.

The algorithm for test pattern generation for single Path Delay faults in primitive gate combinational logic is similarly well suited for parallel computation. We have implemented these algorithms on a shared memory MIMD machine and achieved speedups of 4 to 5 times with 11 processors. It is believed that a slightly better speedup is achievable. In the case of stuck-open faults, there are faster sequential algorithms. However, the single path delay algorithm is comparable to competitive algorithms and shows more promise.

Scientific Field of Study:	VLSI Design (NSF 632; CR B.7)
Architectures/Languages:	Encore Multimax/ C, EPT
Algorithm(s) used:	Network database generator, test pattern generator.
Computational Techniques:	Data parallel decomposition.
Goal of the Research:	Improve the reliability of computer circuits.

TITLE: Evaluation of Distributed Operating System Features on MACH

Principal Investigator (s):
Gary L Craig, Syracuse University

Contact Person:
Dr. Gary L Craig
111 Link Hall
Syracuse, NY 13244-1240
315-443-4389
craig@sutcase

Abstract of Research

This project was established to investigate a variety of distributed operating system issues on the MACH operating system. Our initial work has focused on performance measures in the multiprocessor environment and has been evaluated on a number of parallel and distributed applications. We would also like to evaluate MACH to see if the various schemes suggested for load balancing and process migration in a heterogeneous network environment can be efficiently implemented.

During the course of our investigation, we intend to produce appropriate distributed software tools. We have written the basic user-level IPC (InterProcess Communication) instrumentation package necessary for debugging and performance evaluation. We have also written an animation package (graphical user interface) which interfaces with the instrumentation package.

Scientific Field of Study:	Distributed systems (NSF 314; CR D.4)
Architectures/Languages:	Encore Multimax/ MACH, C, Cthreads
Algorithm(s) used:	Performance evaluation, load balancing
Computational Techniques:	Interprocess Communication in a distributed system
Goal of the Research:	To improve distributed computing as a parallel programming tool.

TITLE: Parallel Graph Reduction Experiments

Principal Investigator (s):
James A. Davis

Contact Person:
James A. Davis
201 Coover Hall
Ames, Iowa 50011
515-294-0659
davis@isuee1.ee.iastate.edu

Abstract of Research

We are proposing to develop software to study parallel combinator graph reduction for the support of functional programming languages. Specifically, we will implement a multiprocessor version of SASL using parallel combinator graph reduction. Functional languages look to be a good platform for studying parallel processing because they provide referential transparency which allows us to reorder computations at run-time without introducing timing or synchronization problems. Another advantage of using functional languages, especially when implemented through combinator reduction, is that parallel expressions are identified automatically (i.e., without user augmentation of the source program).

We are proposing to build the parallel reduction software from a family of multiprocessor simulators that we have developed for past research projects. The target system will be a shared memory multiprocessor, where one processor is responsible for translating SASL to combinatory code, and the other processors concurrently evaluate the combinatory code in shared memory. We expect that the project will allow us to develop a software tool for studying: parallel combinators, the effectiveness of strictness augmentation in a lazy functional language, a scheduling scheme that performs dynamic load balancing through control of task placement, a scheme to vary the degree and granularity of parallelism in relation to the availability of physical resources, and to collect data on memory references so that we can look at cache and bus structures for reduction architectures in general.

Scientific Field of Study:	Software Systems (NSF 314) Programming Languages (CR D3)
Architectures/Languages:	Alliant FX/80, Encore Multimax 520 & 330
Algorithm(s) used:	Combinator graph reduction
Computational Techniques:	Master interpreter "farms" out pieces of computation to reduction slaves; Communicating through shared memory.

Goal of the Research:

To provide parallel implementations of functional languages.

TITLE: Parallel Computational Geometry

Principal Investigator (s):
Frank Dehne

Contact Person:
Frank Dehne
School of Computer Science
Herzberg Building
Carleton University
Ottawa Canada K1S 5B6
DEHNE@CARLETON.CA

Abstract of Research

Computational Geometry (CG) methods are powerful algorithm design tools for the efficient solution to a variety of problems occurring in CAD, graphics, image processing, and robotics. To support the design of efficient high performance parallel systems for such applications, I am studying parallel algorithms for CG and closely related problems on several models of computation; mainly hypercube architectures and processor arrays. During the recent years, I have presented a variety of parallel solutions for geometric problems including d-dimensional ECDF searching, 3-dimensional convex hulls, visibility from a point, parallel visibility, separability problems, Voronoi diagrams, computation of the configuration space of a robot, etc. I also studied closely related searching and dictionary problems for processor networks. During the last year, our project has focused on the testing and implementation of two of our recent theoretical results on hypercube algorithms. The first set of results deals with efficiently implementing parallel branch and bound algorithms on a fine grained hypercube (the Connection Machine). The second set of results deals with the problem of efficiently implementing data structures on a fine grained hypercube (the CM again). Our methods make extensive use of several data routing methods.

Scientific Field of Study:	Computational Geometry (NSF 343; CR I.4)
Architectures/Languages:	Connection Machine
Algorithm(s) used:	Parallel branch and bound
Computational Techniques:	Hypercube routing
Goal of the Research:	Faster image processing, graphics, computer aided design, etc.

TITLE: Manipulation and Implementation of the Hamming Net for Pattern Classification, Signal Processing - Using the Elements of Prediction and System Modeling

Principal Investigator (s):
M. A. Saleh Elmohammed
Ernest Sibert, Syracuse University

Contact Person:
M. A. S. Elmohammed
P. O. Box 397
University Station
Syracuse, NY 13210
315-474-2380

Abstract of Research

This project deals with artificial neural network models which are made of highly parallel rotational elements connected in a pattern that is reminiscent of biological nets. These are of interest primarily because they may be able to emulate the speed and performance of real neural nets using many simple "slow" computational elements operating in parallel. Thus, they offer one possible solution to the problem of obtaining the massive parallelism and computational requirements that are presumed to be required for such problems as signal processing and speech recognition. We are in the process of implementing the Hopfield model. There has been a great deal of work about the ability of Hopfield's model to serve as a CAM and as a pattern classifier for binary bit patterns (note: the analog side has not been much dealt with). A content-addressable memory retrieves one of (M) stored patterns given an input pattern which is a "noisy version" of the stored pattern. Basically, they're focusing on the classification problem for two reasons: 1) A CAM is essentially a classifier which outputs the exemplars for selected class instead of an index to that class. In other words, it determines which of (M) exemplar patterns is "most similar" to a noisy input pattern. 2) Classification is a very fundamental operation that is essential to the important problems of speech, image recognition and signal processing, whether achieved by biological or artificial means. The second part of this study focuses on non-linear signal processing using artificial neural nets with an added parameter of prediction and system modeling.

Scientific Field of Study: Neural Nets (NSF 322; CR I.5)

Architectures/Languages: *Lisp on the Connection Machine.

Algorithm(s) used: Using back propagation in neural nets to match a noisy signal to the best example.

Computational Techniques: Using fine-grained parallelism and a general communication network to model a neural net.

Goal of the Research:

To provide better models to improve signal processing.

TITLE: **Petroleum Reservoir Simulation**

Principal Investigator (s):
Philip Emeagwali

Contact Person:
Philip Emeagwali
University of Michigan
1781-7 Cram Circle
Ann Arbor, Michigan 48105
313-936-3127
emeagwal@caen.engin.umich.edu

Abstract of Research

Petroleum reservoir simulations are computation-intensive, and therefore require the use of supercomputers with a GFLOPS level of performance. This level of performance can be attained by implementing efficient and parallel algorithms on massively parallel supercomputers, such as the iPSC/2, the NCUBE/2, and the Connection Machine Model 2. Our goal is to study those algorithms that are suitable for these machines. Also, lattice gas automaton simulation will be studied as a viable alternative to traditional computational methods.

Scientific Field of Study:	NSF: 526; CR: J2
Architectures/Languages:	Connection Machine 2 and the Encore Multimax
Algorithm(s) used:	Modeling; Lattice gas automaton simulation
Computational Techniques:	Parallel simulation, Data parallelization
Goal of the Research:	Recovery of oil from petroleum resevoirs

TITLE: Electromagnetic Phenomenology Studies

Principal Investigator (s):
Paul Fanelli

Contact Person:
Paul Fanelli
RADC/RBCT
Griffiss AFB, NY 13441-5700
315-330-2465

Abstract of Research

This effort will install the General Electromagnetic Model for the Analysis of Complex Systems (GEMACS) on the Alliant FX/80 computer. GEMACS is an electromagnetic analysis tool which will be used to support RADC programs, such as the Air Defensive Initiative (ADI). The eventual goal is to port the program to the CM-2 machine.

Scientific Field of Study:	Electromagnetics (NSF 631; CR J.2)
Architectures/Languages:	Alliant FX/80
Algorithm(s) used:	Solving linear equations
Computational Techniques:	Matrix techniques
Goal of the Research:	Improve the modeling of realistic electromagnetic situations

TITLE: X Windows Graphics on the Connection Machine

Principal Investigator (s):

Wayne Fordyce
Robert A. Zeh

Contact Person:

Wayne Fordyce
Academic Computing Services
Hinds Hall
Syracuse University
Syracuse, NY 13244

Abstract of Research

This project will involve the development of a data visualization scheme utilizing the Connection Machine graphics primitives for X Windows. The animations will be driven directly by computations running on the Connection Machine. We have chosen X Windows, as a machine independent windowing system, to be able to run Connection Machine graphics in the Advanced Graphics Research Lab.

In this way, we expect to reach the goal of "interactive" data visualization (perhaps even real-time), eliminating the need for "batch-style" intermediate pre-processing of datasets, in preparation for visualization on an extreme computing output source.

Scientific Field of Study: Computer graphics (NSF 343; CR I.3)

Architectures/Languages: Connection Machine 2

Algorithm(s) used: Software tool development

**Computational
Techniques:**

Goal of the Research: Provide interactive data visualization.

TITLE: QCD with Dynamical Fermions on the Connection Machine

Principal Investigator (s):
Geoffrey Fox, California Institute of
G. Stuart Pawley, University of Edinburgh

Contact Person:
Clive Baillie
Mailstop 356-48, 1201 E.
California Blvd.
Pasadena, CA 91125
818-356-6791
clive@hamlet.caltech.edu

Abstract of Research

Quantum Chromo-dynamics (QCD) is the gauge theory of the strong interaction that is responsible for binding quarks and gluons into hadrons (the constituents of nuclear matter). Although QCD can be investigated analytically using perturbation theory at high energies, numerical study by computer simulation is necessary for studying QCD at lower (experimentally attainable) energies. In order to do this, the four-dimensional space-time continuum is discretized onto a four-dimensional hypercubic periodic lattice, with the quarks living on the lattice sites and the gluons living on the links connecting the lattice sites. The first three dimensions of the lattice are the spatial extent and the last dimension is the temporal extent. Because including virtual quark anti-quark pairs in the simulation is very costly in terms of computer time, many simulations are done without these fermions, yielding "quenched" or "pure gauge" QCD. However, with the availability of very high performance machines, we can include the virtual quarks and still run a simulation on a relatively large lattice. It is this "dynamical fermion" system that our code simulates. The gluons are represented by 3×3 complex SU(3) matrices associated with each link in the lattice, while the Wilson representation of the fermions is used for the quarks. These consist of 3×4 complex matrices associated with each site on the lattice. During a QCD simulation, updating these link matrices consumes most of the computer time, with the fermion part comprising the vast majority. This is because all updating algorithms currently require inversion of a very large matrix at each step.

We have developed this code in *Lisp and are running it on four different Connection Machines. The code sustains about 1 Giga flop and peaks at about 1.9 Giga flops. Using these calculations, we are getting close to the goal of 10% errors in the proton to rho mass ratio, which would validate QCD as the correct theory of strong interactions.

Scientific Field of Study: Theoretical Physics (NSF 134; CR J.2)

Architectures/Languages: Connection Machine/ *Lisp

Algorithm(s) used:	Lattice gauge theory with the Feynman Path Integral, Hybrid Monte Carlo Algorithm.
Computational Techniques:	Grid representation, nearest neighbor communication.
Goal of the Research:	To study and quantify the low-level interactions of quarks.

TITLE: Efficient Parallel Geometric Operations on Large Databases.

Principal Investigator (s):
Wm R. Franklin, Rensselaer

Contact Person:
Wm. R Franklin
6026 JEC
Troy, NY 12180
518-276-6077
franklin@turing.cs.rpi.edu

Abstract of Research

Efficient massively parallel geometric and computer aided design algorithms are the focus of this research. They are often the major compute-bound step in several diverse applications, such as visualization of complex scenes in computer aided design, interference detection in automated assembly, overlay of different maps in cartography, and design rule checking in integrated electronics. Since our techniques are also comparatively easy to implement, they are well suited for such applications. Our previous results on a sequential machine include processing a complete chip design with 1,819,064 (almost 2 million) edges. We found all 6,941,110 pairs of intersecting edges in 178 seconds of CPU time. The techniques are extendible to higher level operations. We are now implementing a polyhedron intersection program and a polyhedron visualization program in parallel.

Scientific Field of Study:	Computational Geometry (NSF 343; CR I.4)
Architectures/Languages:	Connection Machine
Algorithm(s) used:	Various geometric and CAD algorithms
Computational Techniques:	Data Parallelism
Goal of the Research:	To improve geometric and computer aided design algorithms.

TITLE: An extended theory of structural character recognition

Principal Investigator (s):
James Geller, New Jersey Institute of
Technology

Contact Person:
Dr. James Geller
Newark, NJ 07102
geller@vienna.njit.edu

Abstract of Research

The subject matter of the proposed research is real time character recognition. Work in the field of character recognition by computers has followed along three general directions: (1) use of low level features, (2) template matching, and (3) density-quadtrees representation. It is the third technique that we propose to use in this project.

The basic idea of density-quadtrees representation is that characters are learned by storing them in a quadtree. A density-quadtrees contains additional information not ordinarily found in a quadtree. The most important item of additional information is the relative pixel density of an area compared to its parent area. In this project we intend to develop an improved density-quadtrees algorithm that adds a normalization step to previous implementations. The nature of the density-quadtrees algorithm allows it to be very gainfully adapted to a Connection Machine implementation.

We have implemented the Teachable Letter Recognizer in *Lisp, doing the two main steps of centering the pattern to be recognized and analyzing its pixel densities on the Connection Machine. Although we have achieved improvements in the identification rates of letters, this version of the program still runs very slowly since the quadtree of already learned letters is still implemented on the frontend machine. We are now working on a new algorithm design to represent the quadtree on the Connection Machine in order to improve the speed of the algorithm.

Scientific Field of Study: Pattern Recognition. (NSF: 343; CR: I.5)

Architectures/Languages: Connection Machine/ *Lisp.

Algorithm(s) used: Character recognition.

Computational Techniques: Density-quadtrees algorithm

Goal of the Research: To develop a character recognizer for arbitrary alphabets.

TITLE: Scheduling of Distributed Computing Systems.

Principal Investigator (s) :
Arif Ghafoor, Syracuse University

Contact Person:
Ishfaq Ahmad
Graduate Lounge CST Building
Syracuse, NY 13244-4100

Abstract of Research

This research is related to the design of task-scheduling algorithms for distributed computing systems. We are analyzing different dynamic scheduling algorithms for distributed homogeneous computer networks which consist of large number of independent computing nodes connected with each other according to a certain topology. The performance of some existing algorithms are being compared with the newly proposed schemes. The comparisons are being done by extensively simulating the algorithms and collecting data for various performance indices such as system's throughput, response time, etc.

Scientific Field of Study: Computer Operating Systems
(NSF 314; CR D.4)

Architectures/Languages: Encore Multimax/ C; Connection Machine/
C*

Algorithm(s) used: Task scheduling algorithms.

**Computational
Techniques:** Extensive simulation

Goal of the Research: To improve the performance of networks of
computers.

TITLE: Parallel Computational Fluid Dynamics in 3-D

Principal Investigator (s):
James Glimm

Contact Person:
James Glimm
Applied Math Dept.
SUNY at Stony Brook
Stony Brook, NY 11794-3600

Abstract of Research

The purpose of this proposal is to conduct a feasibility study for the use of parallel computers to perform fluid dynamics computations using interface methods, i.e. front tracking. The importance of this study rests on several facts. Many of the currently intractable problems of fluid dynamics require the greatly increased power which parallel computers will be able to offer. Here we include multiple length scales, chaotic flows and especially three dimensional computations. In some of these cases, front tracking offers considerable advantages; for example it appears to be the only method to give correctly validated and experimentally confirmed mixing rates for the important problem of late time Rayleigh-Taylor mixing. However these computations and the agreement with experiment are limited by a restriction to two dimensions, so cost effective and powerful parallel computers are a fundamental requirement for continued scientific progress, in order to carry the computations into three dimensions.

Scientific Field of Study:	CFD (NSF 614; CR J.2)
Architectures/Languages:	Connection Machine 2 and the Alliant FX/80.
Algorithm(s) used:	Front tracking
Computational Techniques:	Matrix numerical techniques
Goal of the Research:	Improve the accuracy of modeling complex flow systems.

TITLE: Parallel compiler design using pipeline and data partitioning

Principal Investigator (s):
Amrit L Goel, Syracuse University

Contact Person:
Sanjay Khanna
908 Lancaster Ave.
Syracuse, NY 13210

Abstract of Research

We propose to develop a parallel compiler based on a new data partitioning scheme based on partitioning the grammar of the language being compiled. The scheme effectively decomposes the language specification into multiple subsets. The scheme requires definition and development of specialized subcompilers which will be implemented by a simple and practical method using Parser Generators. The proposed methodology has a potential speed-up factor larger than that obtained by the functional decomposition method. The overall structure of the proposed compiler would also achieve parallelism by pipelining the lexical analysis, syntax analysis and code generation stages. These compiler design ideas are being used to implement a parallel compiler for Pascal on the Multimax and C on the Connection Machine.

Scientific Field of Study:	Compiler Design (NSF:314; CR:D.3)
Architectures/Languages:	C on Encore Multimax, C* on Connection Machine
Algorithm(s) used:	Lexical analyzer and parser generator (Lex and YACC)
Computational Techniques:	Pipelined compiler, compiler based on a data partitioned grammar
Goal of the Research:	To enhance programming language usability by speeding up the compiling process.

TITLE: Performance Analysis of LU Factorization on Multiprocessor System.

Principal Investigator (s):
Amrit L. Goel, Syracuse University
Arif Ghafoor, Syracuse University

Contact Person:
Mallick Shameen Ahsan
111 Link Hall
Syracuse, NY 13244
315-443-4448

Abstract of Research

The solution of a system of linear equations is an important computation found in many engineering and scientific applications. One commonly used approach is to solve such a system is the LU factorization algorithm. Here, we propose a parallel decomposition of this algorithms and analyze its performance for hypercube multiprocessor systems and Connection Machines. Specifically, we have implemented these algorithms and are analyzing the effects of problem decomposition, communication overhead and load balancing on the speedup.

Scientific Field of Study:	Analysis of Parallel Algorithms. (NSF: 118; CR: G1.3)
Architectures/Languages:	Multimax and Connection Machine.
Algorithm(s) used:	LU factorization algorithm.
Computational Techniques:	Data partitioning and hypercube simulator.
Goal of the Research:	Improve the performance of an important algorithm in many scientific disciplines.

TITLE: Testing of Parallel Software

Principal Investigator (s):
Amrit L Goel, Syracuse University

Contact Person:
Nashaat Mansour
111 Link Hall
Syracuse, NY 13244

Abstract of Research

With the increasing use of parallel computers in critical applications, the problem of parallel software reliability is becoming more important. Testing is used to gain confidence about the correctness of programs and hence to ensure their reliability. Testing of parallel programs is inherently more difficult than testing of serial programs and much more work is needed in this field. Such work is expected to critically assess the applicability of some known testing approaches and to develop and implement an approach based on a test selection criterion leading to the development of a testing tool. These research objectives are intended to cover programs for MIMD machines.

Scientific Field of Study: Software Engineering (NSF 315; CR D.2)

Architectures/Languages: Encore Multimax

Algorithm(s) used: Standard testing algorithms

**Computational
Techniques:** Test selection

Goal of the Research: Develop tools for testing parallel programs

TITLE: Massively Parallel Algorithms for Computer Graphics

Principal Investigator (s):
Donald P. Greenberg, Cornell University
Peter Kochevar, Cornell University

Contact Person:
Peter Kochevar
120 Rand Hall
Ithaca, NY 14853
607-255-6705

Abstract of Research

This work proposes to investigate the possibilities that massively parallel architectures may provide for the generation of realistic image synthesis. This research entails developing highly distributive algorithms for modeling shape and form, and for simulating interaction of light with matter. These algorithms will be implemented on the Connection Machine by assigning processors to regions of space and then arraying them into a three-dimensional grid with the stipulation that adjacent processors in the lattice correspond to proximate regions of space. Light paths will be constrained to follow lattice links and the sum over all paths from light sources to each lattice site will be computed inductively by all processors in parallel. The memory contents of certain designated measurement sites will collectively constitute the desired image at the end of a simulation. These algorithms have been successfully implemented on the Connection Machine.

Scientific Field of Study: Graphics and Image Processing
(NSF:343; CR:I.3)

Architectures/Languages: Connection Machine/ C/Paris

Algorithm(s) used: Feynman's path integral formulation of quantum electrodynamics to simulate behaviour of light.

Computational Techniques: 3-D lattice simulator, design of a high level graphics language and a translator to C/Paris.

Goal of the Research: To explore the possibilities that massively parallel architectures provide for the generation of realistic image synthesis.

TITLE: Parallel Algorithms in Joyce

Principal Investigator (s):
Per Brinch Hansen, Syracuse University
Nawal Coptý, Syracuse University

Contact Person:
Nawal Coptý
CIS
Syracuse NY 13244

Abstract of Research

The purpose of this research is to design, implement and analyze parallel algorithms in the Joyce programming language. Joyce is a programming language designed by Prof. Per. Brinch Hansen for distributed systems. It is based on Pascal and CSP, and it permits recursive activation of communicating agents that exchange messages through synchronous channels. Joyce at present is implemented on the Encore Multimax multiprocessor at the Northeast Parallel Architectures Center. The algorithms under investigation fall into the following areas: computational geometry, pattern matching, mathematical computation, fault detection in digital circuits. Presently, the research is focusing on various algorithms in the area of matrix computation. Several parallel paradigms are being investigated, and their relative performance compared.

Scientific Field of Study:	Parallel Languages and Algorithms (NSF 314; CR D.3)
Architectures/Languages:	Encore Multimax.
Algorithm(s) used:	Various parallel algorithms
Computational Techniques:	Message passing
Goal of the Research:	To better understand how to program parallel computers

TITLE: Molecular Dynamic Simulation of Polymer Chain Motions.

Principal Investigator (s):
Barry Hardy, Syracuse University
Philip Rice, Syracuse University

Contact Person:
Barry Hardy
320 Hinds Hall
Syracuse, NY 13244
315-422-3832
bjhardy@suvvm.bitnet

Abstract of Research

This research focuses on the use of molecular dynamics simulation model for understanding conformational behaviour and mechanisms of motion of cellulose chains. Sophiscated computer simulation techniques have been developed on the Alliant computer at the NPAC. The model has been applied for investigating the influence of solute-solvent interactions on the conformational energetics and dynamics of a saccharide, cellobiose, in water. The results are being prepared for publications and have also been presented at the ACS Meeting in Boston, April 1990. We plan to extend this molecular dynamics work to simulate periodic models of crystalline and non-crystalline regions of the cellulose polymer.

Scientific Field of Study:	Molecular Dynamics in Liquids. (NSF:144; CR:J.2)
Architectures/Languages:	Alliant/ Fortran
Algorithm(s) used:	Molecular Dynamics and Molecular Brownian Dynamics simulation methods; energy minimization techniques.
Computational Techniques:	Parallel processing; vector-concurrent evaluation of sums.
Goal of the Research:	Further understand the mechanisms underlying polymer chain motion.

TITLE: Computation of Passive Microwave and Millimeter-wave integrated circuits in Parallel Computers.

Principal Investigator (s):
Roger F. Harrington, Syracuse University
Chung-I Hsu, Syracuse University

Contact Person:
Roger F. Harrington
121 Link Hall
Syracuse, NY 13244-1260
315-443-4391

Abstract of Research

This research project deals with electromagnetic analysis and numerical computation of printed circuits in microwave and millimeter-wave regimes. The method of moments is employed as a numerical scheme, which is well known to be simple, accurate and efficient.

Consider an arbitrarily shaped planar circuit printed on a grounded dielectric slab. Due to the high operation frequency, i.e. microwave and millimeter-wave regimes, quasi-static analysis loses its accuracy. A more accurate full wave analysis is sought to solve the problem. This is done by modeling the conducting surface with rectangular patches and computing the moment matrix elements with Sommerfeld-type integrals, which can be evaluated concurrently for each element. Then the matrix equation is solved using Gaussian elimination or LU decomposition algorithm, yielding the desired quantities. We have successfully implemented some extremely complicated numerical integration schemes for the Sommerfeld-type integrals. The input and mutual impedances of two printed wire or strip antennas computed agree very well with the data published in the literature. Our next step is to characterize the printed circuits, which is more complicated and time consuming than the antenna problem.

Scientific Field of Study:	Electromagnetic fields (NSF 631; CR J.2)
Architectures/Languages:	Alliant/ Fortran
Algorithm(s) used:	Method of moments, Sommerfeld-type integrals, solving linear equations.
Computational Techniques:	Computing matrix elements in parallel.
Goal of the Research:	Understand and improve extremely high-frequency printed circuit boards.

TITLE: Research on Alopex Optimization Algorithms.

Principal Investigator (s):
Erich M. Harth, Syracuse University
Theodore Kalogeropoulos, SU
Volker Weiss, SU

Contact Person:
Erich M. Harth
317 Physics Building
Syracuse, NY 13244
315-443-2565

Abstract of Research

Optimization of a physical situation is generally performed by determining the "best" value of the cost function, within a specified domain of adjustable variables. The Alopex Method of optimization solves this problem for cost functions which may be non-linear in a large number of variables and provides solutions which avoid local extrema by the stochastic introduction of noise. Optimization by Alopex algorithms has been successfully applied to a great variety of problems, including visual receptive fields, models of sensory perception, curve fitting, the traveling salesman problem, and crystal formation. Since the calculation of the algorithm for some of these problems involves large numbers of identical operations, the use of parallel computers could provide substantial speedups as well as providing the capability of solving problems involving large numbers of variables, for example, many-particle systems.

Scientific Field of Study: Numerical Algorithm Development
(NSF:342; CR:G.1)

Architectures/Languages: Alliant / Fortran; Connection Machine/ C* .

Algorithm(s) used: The Alopex method of optimization.

Computational Techniques: Iterative operations on large arrays.

Goal of the Research: Design an optimization algorithm applicable to many particle systems in diverse scientific disciplines.

TITLE: Analyzing the Performance of Large Parallel Computer Systems Based on a Parallel Petri Net Model Which is Already Mapped to the Connection Machine.

Principal Investigator (s):
Fengjia Hu, Syracuse University
Dan Pease, Syracuse University

Contact Person:
Fengjia Hu
839 Livingston Avenue
Syracuse, NY 13210
315-472-3212

Abstract of Research

Petri Nets are event/condition graphs that are excellent tools for modeling asynchronous operations. They have been expanded to stochastic and timed forms; in such form, they are excellent tools for modeling parallel computer systems. This work is aimed at:

1. Expanding the complexity of the systems that can be modeled from those having under 100 elements to those having over 1000 elements.
2. Using parallel machines to implement Petri net tools.
3. Extending Petri Net concepts to mix synchronous and asynchronous operations in a single model.

Scientific Field of Study:	Petri Nets (NSF 311; CR G.2).
Architectures/Languages:	Connection Machine
Algorithm(s) used:	Algorithms for analyzing Petri Nets
Computational Techniques:	SIMD computation
Goal of the Research:	Developing tools for designing parallel machines.

TITLE: To Develop and Implement Parallel Algorithms for Computer Vision

Principal Investigator (s):
Hussein A. Ibrahim, Columbia Univ.
Lisa G. Brown, Columbia University

Contact Person:
Hussein A. Ibrahim
450 Computer Science Building
New York, NY 10027
212-280-2736

Abstract of Research

The goal of this project is to develop parallel algorithms in both low-level and medium-level vision. In the first year, several parallel stereo, depth interpolation, and texture algorithms for highly parallel computer architectures have been developed. An environment has also been developed to program pyramid and multi-resolution algorithms on the Connection Machine. Several of the parallel algorithms for stereo and texture have been implemented in this environment, including a new auto correlation-based texture algorithm. In the pyramidal environment, an image is represented as an image pyramid consisting of a set of decreasing resolution versions of the image. Several classic pyramid operations have been used to test the environment such as stereo matching, edge refinement, pyramid search and similar logarithmically-improved pyramid operations. We propose to continue developing other parallel computer vision algorithms that use the multi-resolution approach to vision and that fuse information from cooperative processes to increase the certainty of computing surface parameters in the image.

Scientific Field of Study:	Computer Vision (NSF 343; CR I.4.8)
Architectures/Languages:	Connection Machine/ *Lisp
Algorithm(s) used:	Stereo and texture vision algorithms on a pyramid of images of increasing resolution.
Computational Techniques:	Fine-grained parallel representation of a pyramid of vision images.
Goal of the Research:	To increase the capabilities of machine vision.

TITLE: Application of Artificial Neural Networks for Learning
 Control of Robot Motion.

Principal Investigator (s):
Can Isik, Syracuse University

Contact Person:
M. K. Ciliz
111 Link Hall
Syracuse, NY 13244
315-443-1125

Abstract of Research

Autonomous control of robots requires the utilization of sensory information at various levels of complexity that are, in general, descriptive, rather than purely numeric. Hierarchical controllers have been proposed to solve the path planning, navigation, obstacle avoidance, and motion control problems. Rule-based methods have been predominantly utilized since they allow symbolic processing of descriptive information. Since the correctness of the rules depends heavily on the environment, changes in the environment should produce changes in the control rules. In this project, the utilization of neural net architectures for adaptive-learning control is considered. The conceptual design of a two-level learning system is being investigated, and the developed system will be implemented on a parallel machine.

Scientific Field of Study: Robotics (NSF:322; CR:I.2).

Architectures/Languages: Connection Machine.

Algorithm(s) used: Neural net two-level learning system.

**Computational
Techniques:** Representing a neural net on a fine-grained
 parallel machine.

Goal of the Research: To improve automated devices.

TITLE: Multiple Scattering Effects from Complex Radar Targets

Principal Investigator (s):
Mr Paul R Jedrezejewski, Syracuse

Contact Person:
Mr Paul R Jedrezejewski
Merrill Lane
Syracuse, NY 13210-4080
315-425-5100

Abstract of Research

The aim of this research is to develop a practical radar cross section prediction scheme including multiple scattering effects using physical optics by taking advantage of parallel computer architecture. Physical optics is a high frequency scheme for predicting the radar cross section of a body. Depending on the direction of the incident field, the body is divided into an illuminated and a shadowed portion. The currents on the illuminated portion are approximated by twice the component of the incident magnetic field intensity tangential to the surface. The currents on the shadowed portion are assumed to be zero. The currents are integrated over the surface of the scatterer resulting in the radiated far field. When the fields of interest lie in a specular direction (direction whose angle relative to the surface normal at the point of reflection is equal to the angle of incidence relative to the same normal), physical optics works very well. Multiple reflected fields may be observed at points of observation which are not specular to the incident radiation. The multiple scattering effects are what is of interest in this work. To date, many runs on NPAC's Alliant computer have been made, resulting in the successful completion of radar imaging experiments which were not feasible on previously available (serial) computers.

Scientific Field of Study:	Radar Imaging (NSF 631; CR J.2)
Architectures/Languages:	Fortran on the Alliant.
Algorithm(s) used:	Inverse Synthetic Aperture Radar Imaging
Computational Techniques:	Automatic parallelization and vectorization on the Alliant FX/80
Goal of the Research:	Better radar imaging systems

TITLE: Parallel Processing of Dynamic Design Sensitivity Analysis for Controlled Multibody Mechanical Systems

Principal Investigator (s):
Woon-Sung Lee

Contact Person:
Woon-Sung Lee
Mechanical & Aerospace Eng.
Truman Campus
Univ. of Missouri at Kansas City
600 W. Mechanic
Independence, MO 64050
816-235-1282
lee@vaxa.cuep.umkc.edu

Abstract of Research

The objective of the proposed research is to develop parallel processing strategies for very efficient dynamic design sensitivity analysis of controlled multibody mechanical systems such as robots, tractors, and passenger cars. The efficient dynamic design sensitivity analysis will provide a designer with valuable information about quantitative effects of design variations on system performance, and will lead to a truly interactive design optimization environment for large scale engineering problems such as design of vehicle systems for passenger safety and comfort.

The proposed research will develop graph theoretic methods to exploit low level parallelism in system and sensitivity equations. Additional lower level parallelism will be exploited to reduce critical path in recurrence relations of the recursive Newton-Euler formulation. The multi-rate integration methods will be developed to exploit high level parallelism by a separation of time scales. The proposed research will use the Alliant FX/80 to develop the parallel processing strategies because heterogeneous computational threads in design sensitivity computation require frequent synchronization and exchange of data.

Scientific Field of Study: Dynamic Systems and Control
(NSF: 621; CR: J2)

Architectures/Languages: Alliant FX/80 and Fortran

Algorithm(s) used: Multi-rate integration method with an error control scheme to monitor the step size

Computational Techniques: Exploitation of inherent parallelism identified by the number of independent kinematic chains and use of data recurrence

Goal of the Research:

To provide information about quantitative effects of design variations on multibody mechanical system performance.

TITLE: Fluid Simulation on the Connection Machine

Principal Investigator (s):
Meng Lean, Xerox Corporation
Stephen A. Vavasis, Cornell University

Contact Person:
Stephen A. Vavasis
Dept. of Computer Science
Cornell University
Ithaca, NY 14853

Abstract of Research

Fluid simulation is among the most computationally intensive tasks in engineering research. The problem under consideration is the simulation of a pool of liquid to which forces are applied. The fluid is modeled as a potential flow, i.e., the velocity is the gradient of a potential function, and the potential satisfies Laplace's equation. The surface of the fluid is treated as a free surface satisfying nonlinear equations, including a term for surface tension. This problem has applications to design and engineering problems involving ink flows in printer devices.

A program written in CM-Fortran has already been implemented to carry out part of this task in conjunction with Xerox Corp. In particular, the boundary element calculations and the free surface calculations have been implemented and have passed some preliminary tests. In addition to CM-Fortran, part of the code has been written in C/Paris and Fortran/Paris.

Scientific Field of Study: CFD Parallel Algorithm.
(NSF: 342; CR: J2)

Architectures/Languages: Connection Machine/Fortran, C/Paris.

Algorithm(s) used: The boundary element method and free stream calculations.

Computational Techniques: Numerical integration and iterative techniques on the CM.

Goal of the Research: Algorithm development for solving ink flows in printer devices.

TITLE: Computational Electromagnetic Fields in Parallel Computers.

Principal Investigator (s):
Meng Lean, Xerox Corporation
Roger F. Harrington, Syracuse University
Taoyun Wang, Syracuse University

Contact Person:
Roger F. Harrington
121 Link Hall
Syracuse, NY 13244-1260
315-443-4391

Abstract of Research

The method of moments in electromagnetic field theory was developed by our research group in the nineteen sixties when high speed computers were finding many practical applications. Because of its simplicity, efficiency, and accuracy, this method has been widely used in electromagnetics engineering and has established itself as one of the standard methods in numerical analysis of electromagnetic fields. Since the matrix structure of the method of moments should allow us to develop parallel algorithms for this technique, we should be able to apply this technique to physical problems of increasing complexity, which could require hundreds or even thousand of expansion functions. Our current research is to consider the electromagnetic penetration into a conducting body of arbitrary shape through an arbitrary aperture. We wish to determine the aperture field distribution, the electric current in the conducting surface, the radar cross section, and the field inside the conducting body. We plan to calculate these quantities by modeling the conducting surface and the aperture with planar triangular patches and computing the impedance of the aperture-closed conducting surface and other related matrices.

We have implemented a system of one conducting body and one aperture in it. We have learned how to structure data dependencies and concurrency to make computing the moment matrices 2 to 4 times faster than before. Now we have the possibility of determining quantities for a multiple conducting and dielectric body and aperture system.

Scientific Field of Study: Electromagnetic fields (NSF 631; CR J.2).

Architectures/Languages: Fortran on the Alliant

Algorithm(s) used: Method of moments, matrix inversion, Gaussian elimination.

Computational Techniques: Computing matrix elements in parallel.

Goal of the Research: Discover fundamental properties of electromagnetic fields.

TITLE: Parallelism in Delay-Constrained Logic Synthesis

Principal Investigator (s):
Yiu M Leung

Contact Person:
Yiu M Leung
608 Leon Drive
Endicott, NY 13760
607-754-5394
ymleung@suvvm

Abstract of Research

One of the most difficult, error-prone and time consuming tasks in VLSI circuit design is the logic synthesis of random combinational logic. Automation of the synthesis process can result in shorter design time and circuit improvements in terms of area and performance. Timing analysis is an integral part of the synthesis process, but is computationally very intensive.

Some of the synthesis phases seem to be ideal candidates for parallel processing. It is felt that marked speed up can be achieved for delay-constrained logic synthesis by designing efficient algorithms which utilize parallelism.

Scientific Field of Study: VLSI Design (NSF 332; CR B.7).

Architectures/Languages: Connection Machine, Alliant, Mach

Algorithm(s) used: To be developed

Goal of the Research: Develop fast parallel algorithms for delay-constrained logic synthesis

TITLE: Intelligent Spectral and Image Analysis

Principal Investigator (s):
George C. Levy, Syracuse University

Contact Person:
George C. Levy
1-104 CST
Syracuse, NY 13244-4100

Abstract of Research

Analysis of nuclear magnetic resonance (NMR) of infrared spectra often requires reduction of spectral data near the limits of detection and is almost always accompanied by non-idealities such as undetermined baselines, incomplete feature separation, and inaccurate representation of feature lineshapes. Research is currently under way on three fronts: optimization of algorithms used for data processing in NMR and infrared spectroscopy, development of new algorithms for magnetic resonance imaging (MRI), and the combination of expert system and statistical analysis of spectra. Three programs for NMR and a program for infrared data processing have been developed: NMR1 for analysis of 1D NMR spectra; NMR2 for 2D and 3D spectra and MRI; MEM for maximum entropy method processing and linear prediction; and SPECTIR for infrared data processing. These programs have several computationally intensive algorithms which can be speeded up by vectorization or by and by parallel processing of subregions of the spectral data.

Scientific Field of Study: Nuclear Magnetic Resonance.
(NSF:345; CR:J.2).

Architectures/Languages: Alliant/ Fortran; Connection Machine .

Algorithm(s) used: Fast Fourier transform; large matrix operations on up to 3D spectral data; 2D peak deconvolution, curve fitting and integration; single value decomposition; principle component analysis and projection of latent structures; rule-based systems.

Computational Techniques: Vectorization of large matrix operations data decomposition.

Goal of the Research: Optimization of algorithms used for NMR data processing.

TITLE: Modelling of Turbulent Shear Flows

Principal Investigator (s):
Jacques Lewalle

Contact Person:
Jacques Lewalle
157 Link Hall
Syracuse, NY 13244
315-443-4211
jlewalle@sunrise.acs.syr.edu

Abstract of Research

To this day, the understanding of the spreading angles of turbulent jets remains elusive. A new spectral closure of the turbulence equations has been developed and calibrated, and preliminary runs on the VAX-8800 points at determining factors for the spreading angles. Thorough calculations of plane and circular jets might shed a new light on this long standing problem. In addition, the model will be used for the calculation of turbulent flows in a turbulent environment (such as pipe jets, coflowing jets and wakes, etc.) and two-phase turbulent shear flows, for which no alternative models are currently available.

By comparison, with standard turbulence models such as k-e, the model developed at Syracuse University requires 6 to 10 times as many variables (depending on the Reynolds number), and progresses in steps two orders of magnitude smaller. Production runs on the VAX would be prohibitively long, but the highly parallel structure of the code should yield considerable speedup on NPAC machines. On the basis of year-long testing on the VAX, production runs on the Alliant FX/80 could begin within days.

Scientific Field of Study:	CFD (NSF:614; CR:J2)
Architectures/Languages:	FX/80 and Fortran
Algorithm(s) used:	Thomas algorithm(3 point finite difference tridiagonal structure) is used for both diffusion and spectral transfer of eddy size & location.
Computational Techniques:	Compiler optimization, the parallel processing of diffusion and transfer is naturally coded into the outer loops. Calculation of matrix elements is vectorized.

Goal of the Research:

Clearer understanding of the physical factors governing the spreading angles of turbulent jets.

TITLE: A Parallel Language Interpreter and Garbage Collector.

Principal Investigator (s):

Xue Li, Syracuse University

Nancy J. McCracken, Syracuse University

Contact Person:

Ms. Xue Li

313 Link Hall

Syracuse, NY 13244

315-443-2163

Abstract of Research

One of the more promising sources of parallelism in programming language implementation is in the implementation of functional languages. In this project, an implementation of a parallel functional language interpreter that can also run in parallel with one or more garbage collectors is explored. The interpreter that was implemented will be a simple lambda-calculus graph-reduction interpreter as described by Simon Peyton-Jones. The garbage collector was a variation of the on-the-fly garbage collector originally proposed by Dijkstra. This kind of garbage collector can run in parallel with the interpreter using shared memory without any need for explicit locking. The synchronization is all handled by the marking of the garbage collector.

The parallel garbage collector that was implemented was Lamport's extensions to Dijkstra's collector. Parallel speedups were measured on the Encore Multimax 520, but the results were negative. The algorithm has such overhead in introducing multiple collectors that it actually runs slower on more processors.

Scientific Field of Study:

Parallel language implementation
(NSF 314; CR D.3)

Architectures/Languages:

Encore Multimax/ C

Algorithm(s) used:

Graph-reduction programming language
interpreter, on-the-fly garbage collector.

**Computational
Techniques:**

Coarse-grained parallel processes using
shared memory without explicit locking.

Goal of the Research:

To provide more efficient programming
languages.

TITLE: Monte Carlo and Iterative Solutions of Partial Differential Equations on Massively Parallel Computers

Principal Investigator (s):
Dr. Michael Mascagni

Contact Person:
Dr. Michael Mascagni
NIDDK
Bethesda, MD 20892
301-496-4325
mascagni@ncifcrf.gov

Abstract of Research

Scientific Field of Study:	Numerical analysis (NSF 312; CR 6.1)
Architectures/Languages:	CM Fortran/CM2
Algorithm(s) used:	Monte Carlo techniques, iterative solvers
Computational Techniques:	Grid computations
Goal of the Research:	To improve fundamental numerical techniques on a massively parallel computer.

TITLE: Three Dimensional Wave Propagation on the Connection Machine

Principal Investigator (s):
Brian J. McCartin

Contact Person:
Brian J. McCartin
275 Windsor St.
Hartford CT, 06120
203-548-2497
mccartin@mstr.hgc.edu

Abstract of Research

Three dimensional wave propagation is central to acoustics and electromagnetics. This research proposal centers about the numerical solution of the three dimensional Helmholtz equation using Artificial Time. In this method, previously applied in two dimensions by Kriegsmann and Morawetz, the steady state wave propagation problem is imbedded in an artificial transient problem. This transient problem is then marched to a steady state. The Connection Machine will be used to explore the rate of convergence of various imbedding schemes and the appropriateness of various time marching schemes. The successful completion of this project would lead to further studies with potential practical applications.

Scientific Field of Study: Theoretical Physics (NSF: 134)
Algorithm Development (NSF: 341)

Architectures/Languages: Connection Machine 2

Algorithm(s) used: Flexible finite difference scheme (Control Region Approximation) to solve numerically three dimensional Helmholtz equation using Artificial Time

Computational Techniques: The algorithm will be applied to a model problem consisting of a metallic sphere with dielectric/magnetic layers. The Connection Machine 2 will be used to explore this massively parallel problem with over 4 million grid points. Each grid point will be assigned to a single processor and will use only local interconnections

Goal of the Research:

Demonstrate the feasibility of using the CM-2 to perform large-scale three dimensional frequency domain scattering computations using the Artificial Time concept and study relative merits of various time marching schemes. The results would lead to advance studies of optimal absorbing boundary conditions.

TITLE: Droplet Motion in Numerically Simulated Turbulence

Principal Investigator (s):
John B. McLaughlin

Contact Person:
John B. McLaughlin
Dept. of Chemical Eng.
Clarkson University
Potsdam, NY 13699-5705
315-268-6663
jmclau@sun.soe.clarkson.edu

Abstract of Research

The goal of this project is to determine the values of the lift and drag forces acting on a rigid sphere in a linear shear flow of viscous fluid. The results would be useful in determining the trajectories of aerosols in turbulent flows over solid surfaces. The computations would involve the use of finite volume techniques to solve the Navier-Stokes equation for the flow around the sphere. A code has been developed on a Cray-2 computer, but the code has not been optimized to take advantage of concurrency and vectorization. The investigators would expect to make use of NPAC consultants for advice in optimizing the code for the Alliant FX/80 computer.

Scientific Field of Study:	CFD (NSF:614; CR:J2)
Architectures/Languages:	Alliant FX/80 and Fortran
Algorithm(s) used:	Finite volume method to solve the Navier-Stokes equations for a flow around a rigid sphere in a linear shear
Computational Techniques:	Use of the vector and concurrent facility and efficient memory utilization through the use of PARALIN to solve the linear system at the core of the calculations
Goal of the Research:	To study the deposition of aerosols onto solid surfaces from turbulent shear flow.

TITLE: Synchronization of Pacemaker Activity Within the Sinoatrial Node

Principal Investigator (s):
Donald C. Michaels, SUNY Health
Science Center

Contact Person:
Dr. Donald C. Michaels
750 Adams Street
Syracuse, NY 13210
315-473-4681

Abstract of Research

In our current research, mathematical models are being used to investigate synchronization of pacemaker activity within the sinoatrial node (the natural pacemaker of the heart) and propagation of electrical activity within and between various cardiac tissues. Cells are coupled through ohmic resistances to form one-, two-, and three-dimensional arrays. We have begun applying the tools of non-linear dynamics and chaos theory to the analysis of the patterns of activity obtained in our simulations. These tools allow us to investigate the complex and irregular dynamics observed under various conditions. We have converted an existing Fortran code into C* and parallelized it. In particular, the two-dimensional arrays of cells mapped nicely to the architecture of the Connection Machine, since each cell can be assigned to a single processor, and the interactions between cells can be modeled with the transfer of information between processors. We are revising the integration scheme to use a finite element method. That will lead to major improvements in computational accuracy and hopefully in computational speed.

Scientific Field of Study: Human Physiology (NSF 422; J.3)

Architectures/Languages: C* on the Connection Machine.

Algorithm(s) used: Non-linear dynamics and chaos theory.

Computational Techniques: Mathematically modeling cells on a grid with nearest neighbor communication.

Goal of the Research: Understanding pacemaker and other electrical activity in the heart.

TITLE: The Phase Problem of X-Ray Crystallography

Principal Investigator (s):
Russ Miller, SUNY Buffalo

Contact Person:
Noushin Bashir
State University of New York at
Buffalo
223 Bell Hall
Buffalo, NY 14260

Abstract of Research

The diffraction patterns of X-rays can be used to determine the structural information of crystals. Since X-rays have short wave lengths, in addition to intensity, they also have phases. Therefore, measuring the intensities and the phases of X-rays scattered by crystals can be used to determine the structure of a crystal. The phases can be deduced from the available information on scattered X-rays intensities. The methods used to find the phases are called direct methods, and the goal of this research is to implement and extend the existing techniques to solve "the phase problem".

A new function has been proposed, call it R, that when minimized can be used to determine the phases of the given crystal. We propose using parallel machines to minimize R. The goal of this research is to use the Connection Machine for minimizing R in parallel. Initially, we have implemented a simulated annealing algorithm to solve the problem. Additional probabilistic approaches will also be considered.

Scientific Field of Study:	Parallel algorithms (NSF 345; CR D.1)
Architectures/Languages:	Connection Machine/ CM Fortran
Algorithm(s) used:	Simulated Annealing
Computational Techniques:	Data parallelism
Goal of the Research:	Cost effective ways of solving the phase problem.

TITLE: Cartographic Generalization for Parallel Processing

Principal Investigator (s):
James E. Mower

Contact Person:
James E. Mower
Department of Geography
Social Sciences 147
University at Albany
Albany NY 12222
JEM97@ALBNY1VX

Abstract of Research

The goal of this project is the implementation of parallel processing algorithms for the automation of cartographic line simplification. We are expanding our previous work on the simulation of connectionist solutions for automated line simplification on serial machines. Now we are using true parallel implementations. Our principal technique is to apply Butterfield's procedure for assigning lines to feature categories. The focus of our research is to achieve fast execution times by mapping the procedure to the massively parallel Connection Machine.

Scientific Field of Study:	Cartography (NSF 452; CR J.2)
Architectures/Languages:	Connection Machine
Algorithm(s) used:	Butterfield's procedure
Computational Techniques:	Data parallelism
Goal of the Research:	Computer assisted map making

TITLE: An Interactive Information Retrieval System Based on a Connectionist Model.

Principal Investigator (s):
Robert N. Oddy, Syracuse University

Contact Person:
Robert N Oddy
4-206 CST
Syracuse, NY 13244
315-443-5616

Abstract of Research

The objective of this work is to explore a type of automatic information retrieval system which interacts with its user by building, and continuously adjusting, a stored image of his/her information need. This contrasts with the usual approach of matching a clearly expressed query with portions of structured records. We are resuming earlier work on a prototype information retrieval system, called Thomas, which was able to find documents relevant to the user on the basis of an extremely sketchy initial query and evaluative reactions given by the user as the dialogue proceeded. Thomas' processes are intrinsically highly parallel, akin to spreading activation in large networks. The C* programs we have developed manipulate networks of words, names, and documents. The texts are handled in conventional ways, in indexed files on the VAX host computer, and the Connection Machine is used for the operations on the network structure.

Scientific Field of Study: Informational Retrieval
(NSF:321; CR:H.3).

Architectures/Languages: C* on the Connection Machine.

Algorithm(s) used: Connectionist nets.

Computational Techniques: Data parallel representation of connectionist nets.

Goal of the Research: To provide successful information retrieval for users without clearly expressed queries.

TITLE: Information Retrieval and Discourse Linguistics, Based on Graph Processing and Connectionism.

Principal Investigator (s):
Robert N. Oddy, Syracuse University
Elizabeth Liddy, Syracuse University

Contact Person:
Robert N. Oddy
4-206 CST
Syracuse, NY 13244
315-443-5616

Abstract of Research

The purpose of this project is to explore the application of discourse linguistics techniques to information retrieval from textual databases. We argue that information retrieval takes place in the context of problem-solving activities in the life of the user, and that keyword-based systems, even with semantic and statistical capabilities, have no way of capturing this context. The discourse-level structures of two important text types - user's problem statements and document abstracts - are seen to be closely related to the working situations of the users, and thus form suitable vehicles for an interactive retrieval system to negotiate with the user. We are working on a system which will automatically analyze empirical abstracts into discourse-level components, corresponding to aspects of empirical research methods. This information is incorporated into a connectionist model, where clue words and components are represented by nodes connected by weighted links (some inhibitory). Part of the network is fixed and represents general knowledge about the structure of abstracts, and part is specific to the features of the user's problem statements.

Scientific Field of Study: Information Retrieval
(NSF:321; CR:H.3).

Architectures/Languages: C* on the CM2.

Algorithm(s) used: Connectionist nets.

Computational Techniques: Data parallel representation of a connectionist net.

Goal of the Research: To provide successful information retrieval of text documents for users without clearly expressed queries, corresponding to real-life situations of scientific researchers.

TITLE: Parallel Algorithms in Computational Geometry

Principal Investigator (s):
Seoung-Jun Oh, Syracuse University
Minsoo Suk, Syracuse University

Contact Person:
Seoung-Jun Oh
B18-3 Slocum Heights
Syracuse, NY 13210
315-445-2115

Abstract of Research

Computational Geometry is concerned with the design and analysis of computational algorithms for solving geometry problems. There are basically five categories in computational geometry according to the nature of geometric objects involved: convexity, intersection, geometric searching, proximity, and optimization. The framework for parallelism in computational geometry is efficient for both synthesis as well as analysis, since 1) all geometry problems can be described as collections of primitive geometry functions for which we can easily analyze the computational complexities, and 2) they are connected very systematically based on parallel problem solving techniques. Several algorithms in each primary area of computational geometry are being developed.

Scientific Field of Study: Computational Geometry
(NSF:343; CR:I.4).

Architectures/Languages: Connection Machine.

Algorithm(s) used: Various algorithms for finding convex hulls, performing geometric searching, and for solving problems of proximity and geometric intersection.

Computational Techniques: SIMD parallelism.

Goal of the Research: Investigating parallel techniques for basic algorithms of computational geometry.

TITLE: Simulation and Analysis of Fine-Grained Parallel Architectures.

Principal Investigator (s):

John V. Oldfield, Syracuse University
Charles D. Stormon, Coherent Research

Contact Person:

Charles D. Stormon
100 East Washington Street
Syracuse, NY 13202
315-426-0929

Abstract of Research

The SUMAC (Syracuse University Machines for Associative Computation) Project at the Syracuse University CASE Center is a long-term, broad-scope venture into fine-grained parallel processing computer architectures and applications. We are currently developing applications and implementations of associative architectures. Some implementations use several prototypes of our VLSI Content Addressable Memory (CAM) chip, while others are simulated using the Alliant and the Connection Machine. Current applications work is centered on Logic Programming, Image Processing, Expert Systems for Automated Test, and Knowledge-base Management Systems.

Scientific Field of Study:

Computer architectures and applications (NSF:313; CR:C.3).

Architectures/Languages:

Alliant, Connection Machine

Algorithm(s) used:

Associative memory, logic programming, expert systems.

Computational Techniques:

Fine-grained implementations of associative memory.

Goal of the Research:

To investigate new computer architecture designs.

TITLE: Application of Parallel Processing and Neural Networks to the Analysis of Remotely Sensed and Medical Imagery

Principal Investigator (s):
Dr. Jo Ann Parikh

Contact Person:
Dr. Jo Ann Parikh
103 MapleWood Dr.
Ithaca, NY 14850
607-257-2684
parikh@msson7.msi.cornell.edu

Abstract of Research

The research to be performed on the Connection Machine consists of the development and implementation of neural networks using parallel algorithms for detection of lineaments in remotely sensed Landsat data of the Canadian Shield and for detection of abnormalities in ultrasonic medical imagery. The primary objectives of the research are to evaluate different neural network models for solving the selected image pattern recognition problems and to evaluate different computer architectures for implementation of neural network paradigms. Conventional serial computers will be used to extract features from the image data and to perform statistical classification and image preprocessing tasks. The performance characteristics of algorithms developed on the Connection Machine will be compared with the performance characteristics of similar algorithms developed on the IBM 3090 and on the Intel Hypecube.

Scientific Field of Study: Pattern matching (NSF 343; CR I.5)

Architectures/Languages: Connection Machine 2

Algorithm(s) used: Neural networks

Computational Techniques: General communication; data parallel nodes

Goal of the Research: Improve image analysis in remote sensing applications.

TITLE: Performance Evaluation and Modeling of Multiprocessor Architectures.

Principal Investigator (s):
Daniel J. Pease, Syracuse University

Contact Person:
Daniel J. Pease
271 Link Hall
Syracuse, NY 13244-1240

Abstract of Research

The design of a machine would be considerably aided by being able to simulate the running of the machine under a variety of assumptions about the machine's use and the algorithms used in the machine. By using the isomorphism between Petri nets and dataflow graphs, we approach this problem by having the architect describe the machine using dataflow graphs and actually running their architectural designs on the Petri nets. Timed Petri nets have already been implemented on the Connection Machine, and preliminary results show that machine designs with 1000 processors and 2000 memory modules can be run in less than an hour. (Previous net implementations on sequential machines ran 16 processor designs in 9 hours.) We have also validated the Petri net implementation by modeling some existing architectures; namely, the Connection Machine, Alliant, and the Multimax. The results are consistent with the actual running of the machines.

Scientific Field of Study: Computer Architecture
(NSF:313; CR:C.3).

Architectures/Languages: C* on the Connection Machine.

Algorithm(s) used: Modeling multiprocessors with dataflow graphs, timed Petri nets.

Computational Techniques: Data parallel implementation of timed Petri nets.

Goal of the Research: To help design the next generation of parallel architectures.

TITLE: Parallel Algorithms for Random Pattern Testability Computation

Principal Investigator (s):
Vijay Pitchumani, Syracuse University

Contact Person:
Pankaj Mayor
CST
Syracuse University
Syracuse, NY 13244

Abstract of Research

SIMD and MIMD algorithms for computing random pattern testability in computer logic networks will be developed. The complexity of the algorithms will be determined analytically. The algorithms are being implemented on the Connection Machine and the Multimax. Performance of this algorithms will be compared with those that run on serial machines. Considerable speedup is expected. The results from this work are expected to make a significant contribution to the field of computer logic testing.

Scientific Field of Study:	VLSI Design (NSF 632; CR J.2)
Architectures/Languages:	Connection Machine, Encore Multimax
Algorithm(s) used:	Algorithms for random pattern testability analysis, including those of Savir et al, and Jain et al (STAFAN algorithm)
Computational Techniques:	Data parallelism on a SIMD and a MIMD machine.
Goal of the Research:	To make logic network fault detection easier easier and faster, and to provide a fast and reliable mechanism for identifying hard to detect faults.

TITLE: Development of Parallel Algorithms for Testing and Simulation of Logic Circuits.

Principal Investigator (s):
Vijay Pitchumani, Syracuse University

Contact Person:
Vinod Narayanan
111 Link Hall
Syracuse, NY 13244
315-443-4410

Abstract of Research

Parallel algorithms for testing and simulation of logic circuits will be developed. These algorithms will be intended primarily for SIMD architectures, and will be implemented on the Connection Machine. The specific problems addressed are 1) weighted random pattern testing based on test counts; 2) logic and fault simulation for gate level networks; 3) switch level simulation; 4) deterministic test generation. Solutions developed for these problems on non-parallel architectures are proving to be inadequate for the current complexity of VLSI designs. These problems have a considerable amount of data-level parallelism. The Connection Machine can result in significantly faster solutions to these problems. Such speed-ups will help ease the testing and design verification problems for large VLSI circuits.

Scientific Field of Study:	VLSI Design (NSF 332; CR B.7).
Architectures/Languages:	C* and *Lisp on the Connection Machine.
Algorithm(s) used:	Weighted random pattern testing, logic simulation, switch-level simulation, and deterministic test generation.
Computational Techniques:	SIMD processing of test patterns and logic simulation.
Goal of the Research:	To develop a self-contained parallel test and design verification environment for VSLI design.

TITLE: Parallel Algorithms for VLSI Layout Compaction.

Principal Investigator (s):
Vijay Pitchumani, Syracuse University

Contact Person:
Sachin Sapatnekar
514 Clarendon Street
Syracuse, NY 13210
315-443-1894
sachin@cmx.npac.syr.edu

Abstract of Research

The idea of VLSI is to pack as much as possible into as little space as possible to obtain the best possible space efficiency. At the same time, the physical layouts of VLSI systems must obey certain design rules that dictate physical constraints in the spacing between adjacent wires or devices. With the advent of more and more complex circuits, the need for a computer aided design tool that handles problems that were previously handled by human designers has been underlined. Parallel architectures play an important role in the computer aided design of such tools. The basic approach to parallel layout compaction is to divide the layout region into a number of parts that may be allotted to different processors. Each of these sub-regions is separately compacted using conventional serial compaction algorithms. Next, these sub-regions must be integrated into a whole that satisfies the boundary conditions that prevail between regions. This is an iterative process that must be repeated several times before a final result is obtained.

Scientific Field of Study: VLSI Design (NSF 332; CR B.7).

Architectures/Languages: Connection Machine

Algorithm(s) used: Parallel decomposition of layout regions.

Computational Techniques: MIMD data parallelism.

Goal of the Research: To provide computer aided design tools for VLSI designers.

TITLE: Explicit Methods of Structural Mechanics on Parallel Machines

Principal Investigator (s):
Edward J. Plaskacz,
Argonne National Lab

Contact Person:
Edward J Plaskacz
9700 S Cass Building
221-E101
Argonne, IL 60439
312-972-7975
plaskacz@antares.mcs.anl.gov

Abstract of Research

To study the feasibility of finite element analysis on the Connection Machine as well as learning the idiosyncracies of programming on the Connection Machine, a short and simple code in FORTRAN was used as a tool. In deference to the vast volume of FORTRAN code which has been developed over the past 30 years and the infeasibility of "translating" these codes into a specific language dialect for each innovative computer architecture being developed, the approach I have taken in the execution of my research is to translate the computationally intensive protions of a FORTRAN solid mechanics code into C* and linking it to a FORTRAN driving program which manages I/O. This approach will undoubtedly capture the interest of organizations using the Connection Machine on a time sharing basis in the future for whom extensive code renovation is infeasible.

Scientific Field of Study:	Structural Mechanics. (NSF: 622; CR: J.2)
Architectures/Languages:	Connection Machine/C*.
Algorithm(s) used:	Finite Element Analysis.
Computational Techniques:	Translation of an existing FEM code to C*.
Goal of the Research:	Improving performance of existing FEM applications.

TITLE: Two-dimensional Kalman Filtering Using Parallel Processors.

Principal Investigator (s):
David J. Potter, Clarkson University

Contact Person:
David J. Potter
Potsdam, NY 13676

Abstract of Research

The project involves the implementation of two-dimensional Kalman filtering on different parallel architectures. Currently, Sequent and CLIP machine implementations have been studied, but it would be useful to study implementation on the architectures available at NPAC. Two-dimensional Kalman filtering is computationally intensive. However, unlike most image processing algorithms, the parallelism is not simply on the data at each pixel, which would give machines such as CLIP (processor arrays) the best advantage for optimization. The Kalman filter is based on the scan-line model such that a pixel in the middle of the image is a function of pixels above and to the left of that pixel. For all three systems, the Encore, the Alliant and the Connection Machine, the same basic wavefront technique is used.

Scientific Field of Study: Image Processing (NSF 343; CR I.4)

Architectures/Languages: Alliant, Encore Multimax, Connection Machine

Algorithm(s) used: Two-dimensional Kalman Filtering

Computational Techniques: MIMD Wavefront techniques, and some SIMD techniques adapted from the CLIP

Goal of the Research: Complete a study of parallelization of Kalman Filtering, to date conducted on a Sequent, a CLIP machine

TITLE: Concurrent Computations in General Relativistic Astrophysics

Principal Investigator (s):
Jorge Pullin

Contact Person:
Jorge Pullin
Physics Dept.
Syracuse University
Syracuse, NY 13244-1130
315-443-1821
pullin@suhep.phy.syr.edu

Abstract of Research

We propose to use NPAC's CM-2 machine to run a one dimensional time dependent Relativistic Hydrodynamics code which is already tested and running on vector machines. This research will serve as a test bed for the potential use of parallel machines in more realistic two and even three dimensional situations in a near future. Among several other applications, the two and three dimensional codes could be successfully used in the study of gravitational radiation in gravitational collapse of stars, processes of accretion of matter onto stars and compact stars coalescence and therefore provide a theoretical framework from which to interpret experimental data that will soon be available from interferometric gravitational wave detectors.

Scientific Field of Study:	Gravitational Physics (NSF: 135; CR: J2)
Architectures/Languages:	Connection Machine and CM-Fortran
Algorithm(s) used:	Numerical solution of a system of partial differential equations
Computational Techniques:	Data parallel processing on the Connection Machine
Goal of the Research:	Provide theoretical framework to interpret experimental data soon to be available from interferometric gravitational wave detectors.

TITLE: Parallel Algorithms for Image Processing and Pattern Recognition

Principal Investigator (s):
Sanjay Ranka, Syracuse University

Contact Person:
Sanjay Ranka
313 Link Hall
Syracuse, NY 13244
315-443-4457
ranka@top.cis.syr.edu

Abstract of Research

Many applications in real-time image processing and pattern recognition require computational capability far beyond what can be obtained from the fastest single processor computer available. Hence the usefulness of these applications is limited. Advanced in VLSI (Very Large Scale Integration) have provided a cost-effective means of obtaining increased computational power by using parallel computers. The primary thrust of the proposed research is to continue our investigations into the design and implementation of algorithms for computationally intensive image processing and pattern recognition tasks on parallel computers.

We will carry this out in a number of different methods with the overall goal of improving the speed and performance of image processing and pattern recognition systems. We propose to develop efficient computational techniques for some important primitive tasks on different parallel models of computation. We have implemented several new squared-error clustering algorithms, which we developed with Sartaj Sahni. These clustering algorithms perform pattern matching by partitioning a set of patterns into groups and are applicable to machine vision.

Scientific Field of Study:	Image Processing. (NSF: 343; CR: I.4.0)
Architectures/Languages:	Connection Machine/*Lisp
Algorithm(s) used:	Squared-error clustering.
Computational Techniques:	Implementing basic hypercube operations on the hypercube of the Connection Machine.
Goal of the Research:	To develop an efficient computational technique for image processing.

TITLE: Physical Synthesis and Verification of VLSI Integrated Circuits

Principal Investigator (s):
Rob A Rutenbar, Carnegie Mellon

Contact Person:
Rob A Rutenbar
Pittsburgh, PA 15213
412-268-3334
rutenbar@gauss.ece.cmu.edu

Abstract of Research

As the size of Very Large Scale Integrated Circuits (VLSI ICs) increases, the computational demands on the algorithms that synthesize and verify the geometric layout of the required IC masks increases. Despite relentless optimization of sequential algorithms with good serial complexity, the new generation of application-specific ICs with 100,000 to 1,000,000 devices requires enormous amounts of CPU time to design. This project focuses on two parts of this problem: layout synthesis problems based on divide-and-conquer schemes such as graph bipartitioning, and layout verification problems based on computational geometry methods. Both problems are well-suited to massively parallel machines: they require very large data sets, need modest computations on each object in the data set, and exhibit some form of locality of computation. At this writing, our work on layout verification on the CM2 is close to completion. We successfully extended our earlier work to incorporate new methods for flattening mask layouts, new methods for resolving connectivity, and methods for transforming the identified wires and devices into a complete netlist representing the circuit. In regard to layout synthesis, substantial progress was made. We have completed first versions of CM algorithms for direct placement methods based on optimizing a quadratic wirelength metric.

Scientific Field of Study: Design of Very Large Scale Integrated Circuits (NSF 632; CR J.6)

Architectures/Languages: C/Paris on CM2 Connection Machine

Algorithm(s) used: Divide-and-Conquer Methods, Graph Partitioning, Computational Geometry

Computational Techniques: Fine-grain data-parallel problem decompositions

Goal of the Research: Reduce design/verification times, and greatly improve design quality of very large custom integrated circuits.

TITLE: Digital enhancement of remotely sensed data for mineral exploration

Principal Investigator (s):
Joe Robinson, Syracuse University
Ashok K. Joshi, Syracuse University

Contact Person:
Ashok K. Joshi
211 Link Hall
Syracuse NY 13244

Abstract of Research

Remotely sensed data from satellite platforms have been used extensively for geological exploration since the beginning of the last decade. The blurring of the image due to modulation of the sensor's optical system, detector and atmosphere often causes degradation of the image. Due to this blurring effect the sharp contacts and edges are not readily visible in the imageries. Photo-interpretation of these satellite images is the more commonly used technique and not much emphasis has been paid to digital image processing of remotely sensed data. Available software and hardware have been used to improve the quality of image for manual interpretation but efforts have not been made for quantitative analysis. Also most of the processing for improving the spatial quality of the image has been done in spatial domains using convolution kernels that are not efficient enough to enhance the edges selectively. This type of processing often increases the noise in the image at the expense of resolution improvement. Convolution operators degrade the image quality considerably and result in processed images that cannot be used for spectral analysis.

Scientific Field of Study:	Image Processing (NSF 343; CR I.4)
Architectures/Languages:	Connection Machine, Encore Multimax
Algorithm(s) used:	Fast Fourier Transform
Computational Techniques:	Grid representation of image points, nearest neighbor communication.
Goal of the Research:	To identify more detailed information from the image in order to improve the exploration strategy in mineral exploration.

TITLE: A Connection Machine Implementation of SUPER

Principal Investigator (s):
J. Alan Robinson, Syracuse University
Ernest Sibert, Syracuse University

Contact Person:
Dr. Ernest Sibert
313 Link Hall
Syracuse, NY 13244
315-443-3110
sibert@cmx.npac.syr.edu

Abstract of Research

Robinson and Greene proposed a high-level programming system known as SUPER (for Syracuse University Parallel Expression Reduction) which combines both functional and relational constructs. The language is defined by means of an abstract graph-reduction machine based on techniques pioneered by David Turner, but considerably extended to encompass the relational features of SUPER. Such a system naturally admits parallel reductions, without difficult synchronization requirements. The result is, at least in principle, a very expressive, generalized logic-programming system which can be given a parallel implementation, yet does not require programmers to specify parallel computations explicitly.

We have written a first implementation of SUPER on the Connection Machine, using a refined definition of the system. This effort has exposed a number of interesting questions concerning the abstract definition as well those raised by the limit of finite resources and those specific to the CM architecture. We see already much better ways to implement such languages on the Connection Machine.

Scientific Field of Study:	Parallel Programming languages (NSF:314; CR:D.3).
Architectures/Languages:	Connection Machine/ *Lisp.
Algorithm(s) used:	Parallel graph reductions.
Computational Techniques:	SIMD Combinator graph reduction, data parallel representation of graphs.
Goal of the Research:	To provide very high-level programming languages with a natural programming style.

TITLE: A Computational Approach to the Assessment of the Visibilities of Exit Signs Under Adverse Conditions

Principal Investigator (s):
Badrinath Roysam

Contact Person:
Badrinath Roysam
Dept. of ECSE
R.P.I.
Troy, NY 12180
518-276-8067
roysam@ecse.rpi.edu

Abstract of Research

This project is concerned with quantitative assessment of the visibility of directly or indirectly-lighted exit signs in buildings under adverse conditions such as smoke, dust and fumes. A related problem is the visibility of highway signs under adverse weather conditions. Our approach to this study is based on the synthesis of computer graphic images of the exit signs under various lighting and visibility conditions using quantitative models for the exit signs as well as the light transmission medium. This will allow us to study the influence of specific design choices such as lighting arrangements and colors under a wide range of carefully-controlled visibility conditions. Our initial efforts will involve static monochrome imagery, which will be used to validate the underlying mathematical models and computational methods. These results will then be extended in stages to full true-color static imagery, and then to dynamic imagery via the visualization of image sequences. The latter development will allow the quantitative visualization of exit signs under rapidly changing visibility conditions.

Scientific Field of Study:	Vision modeling
Architectures/Languages:	Connection Machine Model CM-2
Algorithm(s) used:	Ray tracing, radiosity method
Computational Techniques:	Data parallel across photon bundles
Goal of the Research:	To improve traffic sign visibility.

TITLE: Time-Parallel Circuit/Switch Simulation on Massively Parallel Processors

Principal Investigator (s):

Karem A Sakallah, University of Michigan

Contact Person:

Karem A Sakallah
2213 EECS Building
Ann Arbor, MI 48109-2122
313-936-1350
karem@caen.engin.umich.edu

Abstract of Research

Detailed simulation of VLSI circuits at the electrical and switch levels continues to be a major computational bottleneck in the design verification process. They are, thus, prime candidates for implementation on parallel machines. In particular, we feel that the maximum performance gain can only be realized with fine-grain parallel machines such as the Connection Machine. This requires "regularizing" the simulation algorithms so that they minimize random access patterns. In addition, the availability of thousands of processing elements makes it possible to exploit parallelism along the time axis by replicating the circuit equations over time. At this time we have completed the implementation on the CM of a test generation algorithm for digital combinational circuits at the gate level. The test generation method used is known as threshold value simulation method. We are also continuing our research into circuit simulation on the Connection Machine. Specifically, we are seeking a method to enhance the performance of the CM for the transient simulation of small circuits.

Scientific Field of Study: Circuit Simulation (NSF 634; CR I.6)

Architectures/Languages: Connection Machine

Algorithm(s) used: Threshold value simulation

Computational Techniques: Data parallelism (mapping gates and connections to processors)

Goal of the Research: More efficient and faster circuit design and testing

TTITLE: Implementation and Performance Evaluation of a
Reconfigurable Parallel Computer Vision System

Principal Investigator (s):
Ashok Samal

Contact Person:
Ashok Samal
Dept. of Computer Science
115 Ferguson Hall
Univ. of Nebraska - Lincoln
Lincoln, NE 68588-0115
402-472-2217
samal@fergvax.unl.edu

Abstract of Research

Multiprocessors have been proposed to reduce the time taken for object recognition. Building a parallel computer vision system is a two step process: (a) design and implement parallel algorithms for individual vision operations, and (b) integrate them into a coherent system. The emphasis so far has been on the first step and the second step has not received enough attention.

We have proposed a framework for a parallel computer vision system. the necessary and desirable features of such a system have been identified. Such features include support for large number of operators, dynamic reconfigurability, system efficiency, easy user interaction, architecture independence, etc. We have also given an initial design which incorporates most of these features. Our goal now is to implement the design on different types of multiprocessors, especially large ones, test the performance of the system and modify the design if necessary.

Scientific Field of Study:	Computer vision (NSF 322; CR I.4)
Architectures/Languages:	Connection Machine 2, Encore Multimax 330, and the Alliant FX/80
Algorithm(s) used:	Incorporates parallel versions of traditional vision algorithms
Computational Techniques:	Functional parallel system decomposition; data parallel vision algorithms
Goal of the Research:	Improved parallel vision system design.

TITLE: Parallel Computation of EM Fields for Applications in High Speed Integrated Circuit Design.

Principal Investigator (s):
Tapan K. Sarkar, Syracuse University

Contact Person:
Madhavan Swaminathan
111 Link Hall
Syracuse, NY 13244-1240
315-443-4450

Abstract of Research

Electromagnetic scattering from surfaces of arbitrary shape involve the solution of complex integro-differential equations. The method of moments provides us with a technique to convert these equations into a matrix equation, which can be solved for the unknowns. This technique uses a two step process: 1) generation of the complex matrix using a numerical procedure. 2) solution of this matrix involving thousands of unknowns. Parallel computing can reduce the time needed to generate and solve the matrix.

Scientific Field of Study: Electromagnetic Field
Computation. (NSF:631; CR: J.2).

Architectures/Languages: Alliant/ Fortran.

Algorithm(s) used: The method of moments.

Computational Techniques: Generating and solving large matrices in parallel.

Goal of the Research: To solve electromagnetic scattering and radiation from arbitrarily shaped conducting bodies of sufficient size to use with real objects like aircraft, ships, etc.

TITLE: **Massively Parallel Simulation of Queueing Networks on SIMD Machines**

Principal Investigator (s):
Robert G. Sargent, Syracuse University

Contact Person:
Mr. Bruce Cota
313 Link Hall
Syracuse, NY 13244
315-472-8273
cota@amax.npac.syr.edu

Abstract of Research

Parallel discrete event simulation has been the subject of a great deal of recent work but nearly all of this work has dealt with the MIMD model of computation. Queueing network models are a common subject of discrete event simulation and we believe that the structure of these models will make the use of SIMD computation feasible. We plan to experiment with methods of simulating queueing networks on the connection machine. We intend to use combinations of different methods of parallel simulation in an attempt to exploit the full power of the connection machine.

Scientific Field of Study: Queueing Networks (NSF 342; CR I.6)

Architectures/Languages: Connection Machine

Algorithm(s) used: Parallel discrete event simulation

**Computational
Techniques:** SIMD computation

Goal of the Research: Develop methods for simulating queueing networks efficiently on SIMD architectures

TITLE: Parallel Algorithms for Discrete Event Simulation.

Principal Investigator (s):
Robert G. Sargent, Syracuse University

Contact Person:
Robert G. Sargent
431 Link Hall
Syracuse, NY 13244-1240
315-443-4348

Abstract of Research

Our research addresses two problems in the area of large-scale discrete event simulations: they typically have long computer run times, and they usually are labor-intensive with respect to the development of the simulation models. We are developing an approach for parallel discrete event simulation on a shared-memory multiprocessor computer regarding these two problems. Our approach differs in three ways from most other research: 1) It is targeted for shared memory parallel computers; 2) It identifies when events can "safely" be executed in parallel; 3) We plan to evaluate it experimentally on an actual parallel computer. Our approach will include a parallel implementation of a small modeling language that we are developing for this purpose in C.

Scientific Field of Study: Simulation and Modeling
(NSF:345; CR:I.6)

Architectures/Languages: C/EPT on the Encore Multimax.

Algorithm(s) used: Queueing network modeling.

Computational Techniques: Using a modelling language to model a system as a number of interacting entities.

Goal of the Research: Speeding up large-scale discrete event simulations.

TITLE: Calculations of the 3-D Structures of Proteins.

Principal Investigator (s):
Harold Scheraga, Cornell University

Contact Person:
Harold Scheraga
Baker Lab
Ithaca, NY 14853

Abstract of Research

The goal of this project is to predict the three-dimensional structure of relatively small proteins (50-60 amino acid residue range) by using algorithms to locate the minimum of potential energy that corresponds to their native conformations. Based on the hypothesis that the native conformation of a protein is determined completely by its constituent sequence of amino acids, statistical mechanical arguments show that the native conformation must correspond to a minimum in the potential energy surface that governs the motion of the protein in its physiological environment. Computational models will be generated which should closely predict the most physiologically stable and therefore biologically consistent protein structure. The first of two algorithms to be used is based on building an amino acid chain by combining small pieces for which all minima are located within a fixed energy of the global minima. These are in turn combined on the basis of locating the low-energy minima of the new fragment, and so on in tandem. The second algorithm uses an adaptive Monte Carlo method.

Scientific Field of Study:

Molecular Biochemistry
(NSF:411; CR:J.3).

Architectures/Languages:

C* on Connection Machine,
Fortran on Alliant FX/80.

Algorithm(s) used:

Finding the minimum of a potential energy
function, adaptive Monte Carlo method.

**Computational
Techniques:**

Automatic and manual parallelizing existing
algorithms, data parallelism.

Goal of the Research:

To understand protein structure in order to
solve one of the great unsolved problems,
protein folding.

TITLE: Calculations of the 3-D Structures of Proteins.

Principal Investigator (s):
Harold Scheraga, Cornell University

Contact Person:
Harold Scheraga
Baker Lab
Cornell University
Ithaca, NY 14853

Abstract of Research

The goal of this project is to predict the three-dimensional structure of relatively small proteins (50-60 amino acid residue range) by using algorithms to locate the minimum of potential energy that corresponds to their native conformations. Based on the hypothesis that the native conformation of a protein is determined completely by its constituent sequence of amino acids, statistical mechanical arguments show that the native conformation must correspond to a minimum in the potential energy surface that governs the motion of the protein in its physiological environment. Computational models will be generated which should closely predict the most physiologically stable and therefore biologically consistent protein structure.

Scientific Field of Study: Molecular Bioscience.

Architectures/Languages: C* on CM, Fortran on Alliant.

Algorithm(s) used: Finding the minimum of a potential energy function, adaptive Monte Carlo method.

Techniques: Automatic and manual parallelizing existing algorithms, data parallelism.

Goal of the Research: To understand protein structure in order to solve one of the great unsolved problems, protein folding.

TITLE: Efficient Methods for Numerical Analysis.

Principal Investigator (s):
Fred H. Schlereth, Syracuse University

Contact Person:
Fred H. Schlereth
227 Link Hall
Syracuse, NY 13244-1240

Abstract of Research

The purpose of this research is to experiment with efficient methods for numerical analysis on the FX/80 and the Connection Machine (CM2). Currently, algorithms and architectures for the solution of very large matrices (10,000 by 10,000) are being investigated. Data flow between the processing elements and memory is the main problem in taking advantage of parallelism for scientific computing problems, especially as the size of the problems gets large. One aim of the project is to apply computing strategies developed for ASIC parallel computers to these matrix computations.

Scientific Field of Study: Numerical Analysis
(NSF:312; CR:G.1).

Architectures/Languages: C/PARIS on the CM2 and
Fortran on the Alliant FX/80.

Algorithm(s) used: Various matrix operations.

Computational Techniques: Minimizing data flow between
processors and memory.

Goal of the Research: To analyze the performance of the
Connection Machine and the FX/80 on
typical matrix problems, and propose
alternative architectures if appropriate.

TITLE: Implement Fast Parallel Graph Reduction on the Connection Machine.

Principal Investigator (s):
Dr. Ernest Sibert, Syracuse University

Contact Person:
Feng Yang
313 Link Hall
Syracuse, NY 13244

Abstract of Research

Our research is to implement a functional language interpreter on the Connection Machine by the technique of graph reduction, using an interpreter with a very small instruction set so that the same interpreter can be run in parallel at all nodes of reduction. For this project we are working on the following tasks: 1) develop an algorithm to transform lambda-expressions into SKI-combinator expressions for an efficient implementation of graph reduction; 2) investigate the possibility of parsing lambda-expressions in parallel; 3) develop a graph reduction scheme to implement an extension of SKI-combinators in parallel; 4) develop a combination of reference-counter and mark-collecting garbage collection methods and implement it on the Connection Machine.

Scientific Field of Study: Functional Programming (NSF 312; CR D.3).

Architectures/Languages: *Lisp on the CM.

Algorithm(s) used: Transforming lambda-expressions into SKI-combinator expressions, parsing of lambda-expressions, graph reduction, reference-count and mark/collect garbage collection.

Computational Techniques: SIMD graph reduction, data parallel representation of graphs.

Goal of the Research: Developing schemes for parallel implementation of functional programming.

TITLE: Parallel Graph Algorithms for the CM/PRAM Simulation on the CM.

Principal Investigator (s):
Ernest Sibert, Syracuse University
Arif Ghafoor, Syracuse University

Contact Person:
Dr. Ernest Sibert
CIS Department
3-171 CST
Syracuse, NY 13244-4100

Abstract of Research

This research consists of the investigation and implementation of published parallel graph algorithms which match a SIMD architecture well. The goal is to provide a collection of efficient parallel graph algorithms which will be useful for some common computationally intensive subproblems found in many applications. Effort shall also be made to improve upon the known bounds of problems. We are now implementing an efficient simulation of the Parallel Random Access Machine (PRAM) on the Connection Machine. The PRAM model can be used to express many parallel algorithms in an abstract way, without having to write details of communication, etc.

Scientific Field of Study: Parallel Algorithms
(NSF:345; CR:D.1).

Architectures/Languages: Connection Machine/ *Lisp

Algorithm(s) used: A collection of well-known parallel graph algorithms.

Computational Techniques: A fine-grained representation of graphs, simulation of the PRAM model.

Goal of the Research: To find an efficient SIMD implementation of algorithms which are common to many applications.

TTILE: Connectionist Algorithms for Hard Combinatorial Problems

Principal Investigator (s):
Jadranka Skorin-Kapov

Contact Person:
Jadranka Skorin-Kapov
Dept of Applied Mathematics
SUNY@Stony Brook
Stony Brook, NY 11794

Abstract of Research

A massively parallel algorithm to perform tabu search for the Quadratic Assignment Problem (QAP) is proposed. The algorithm is based on a connectionist model. For a QAP of size n , the algorithm requires $O(n^2)$ processors. The resultant speedup is also $O(n^2)$. The algorithm is implemented on a Connection Machine and its performance is evaluated.

Scientific Field of Study:	Parallel algorithms (NSF 311; F.2)
Architectures/Languages:	Connection Machine 2
Algorithm(s) used:	Boltzmann machines
Computational Techniques:	Data parallel nodes with optimum communication strategies
Goal of the Research:	To improve general optimizing algorithms.

TITLE: SIMD Massively Parallel Exhaustive Search of Small State Spaces.

Principal Investigator (s):
Lewis Stiller, Boston University

Contact Person:
Lewis Stiller
111 Cummington Street
Boston, MA 02215
617-353-1543

Abstract of Research

Exhaustive state space search can yield optimizations which are qualitatively different than heuristic search would be likely to yield. This technique has already been used to find new solutions to chess end games with 5 pieces and to optimize assembly language programs. This research proposal is to implement an exhaustive search of a game space by computing all possible moves from a game configuration. This method achieves greater functional simplicity at a cost of considerable computational redundancy. The CM-2 will be used for the implementation, which will permit heretofore infeasible real-time domain interaction and experimentation. This may help use endgames to explore induction learning and expert systems. This project will also explore the canonical exponentially time complete game Peek, a game played on propositional formulae.

Scientific Field of Study: Artificial Intelligence (NSF 322; CR I.2).

Architectures/Languages: C/Paris on the Connection Machine

Algorithm(s) used: Game playing by searching the game space.

Computational Techniques: Massively parallel And/Or tree backup algorithm, using operators on each node carefully coded for bit-serial architecture.

Goal of the Research: Increase our knowledge of human intelligence and learning.

TITLE: Development of a neural net simulator

Principal Investigator (s):
Minsoo Suk, Syracuse University

Contact Person:
Ganesh Ramamoorthy
CST
Syracuse, NY 13244
315-443-2218
ramamoor@cmx.npac.syr.edu

Abstract of Research

Connectionism refers to a class of massively parallel architectures that consists of several simple processing units, each connected to several other similar units. The knowledge in such systems is stored in the pattern of connections between the units and the corresponding weights. Since knowledge is stored by a pattern of activation over several processing units, it provides a distributed representation. Some connectionist models that are used as neural net classifiers are Hopfield net, Hamming net, Carpenter/Grossberg classifier, Kohonen's Self Organizing Feature Maps, and the Perceptron approach. Since a good deal of time is spent writing programs to emulate the networks under study, it would be immensely useful to have a framework which one could use for simulating all neural nets of this sort. We are now implementing a general neural net simulator on the Connection Machine.

Scientific Field of Study: Neural Nets (NSF 322; CR I.2)

Architectures/Languages: Connection Machine

Algorithm(s) used: Back-propagation, Hopfield nets, Hamming nets, Carpenter/Grossberg classifiers, Kohonen's self organizing feature maps, perceptrons

Computational Techniques: Data parallel representation of neural nets

Goal of the Research: To provide a basic tool in the study of neural nets.

TITLE: Parallel Processing in Knowledge-based Computer Vision Systems.

Principal Investigator (s):
Minsoo Suk, Syracuse University

Contact Person:
Dr. Minsoo Suk
113 Link Hall
Syracuse, NY 13244
315-443-4405

Abstract of Research

Our research is concerned with the development of robust techniques for 3-D object recognition procedures dealing with multiple-object scenes with partial occlusion and background clutter. Our problem at present is limited to rigid objects with six degrees of freedom. We will extend our results to rigid objects with parameterized variations, rigid objects in motion and deformable objects. The ultimate goal of our research is the better understanding of the following theoretical and practical issues: development of suitable indexing schemes for indexing into a library of object models represented at multiple levels of abstraction, insight into how symbolic and numerical constraint satisfaction techniques could be integrated into a problem-solving mechanism, and a more robust and general framework for 3-D object recognition.

Scientific Field of Study: Artificial Intelligence. (NSF:322; CR: I.2).

Architectures/Languages: Connection Machine (SIMD), starlisp; Alliant (MIMD), Encore Multimax (MIMD).

Algorithm(s) used: Hough Transform

Computational Techniques: SIMD image processing on the Connection Machine, rule-based shared memory systems like a blackboard on the Multimax, neural net architectures on the Alliant.

Goal of the Research: To improve techniques for 3-D object recognition.

TITLE: Specification and Testing of Parallel Programs

Principal Investigator (s):
Andrew H. Sung

Contact Person:
Andrew H. Sung
Dept. of Computer Science
New Mexico Tech
Workman 25 - Leroy St.
Socorro, NM 87801
505-835-5126
sung@nmt.edu

Abstract of Research

In software engineering, a description of the intended behavior of a program is termed as its specification. It is an accepted fact that formalization of this aspect evidently leads to a more precise end product and ambitiously, even to its automation. Research in software specification, however, appears to be heading toward a very precise mathematical framework, making it difficult for the average programmers to apply most of the proposed methods.

The present state of parallel processing is characterized by the proliferation of parallel machine architectures and relatively slower developments of software technologies. The introduction of generic logic specifications is intended to redress this imbalance. We propose to introduce generic concepts written in PARLOG constructs which are to be translated into a parallel programming language.

The proposed thesis is to use first-order predicate logic for the derivation of program specifications, write these specifications into computable Horn clause logic programs in declarative-style using PARLOG and do transformations on these programs into an imperative language specific to an intended machine or architecture. For the initial stage, the work will be concentrated on two versions of Fortran: that which run on a SIMD machine and on an MIMD machine (in particular the Connection Machine, the FX/80, and the Multimax 520).

The objective of this research is three-fold:

- 1) To enhance the clarity of program specifications via stepwise refinements;
- 2) To introduce generic parallel software specifications that would be applicable to a wide variety of architectures; and
- 3) To facilitate software validation through test data generation using the executable specification in PARLOG and the target code in Fortran.

Scientific Field of Study: Software engineering (NSF 315; CR D.2)

Architectures/Languages:	CM-2, Encore Multimax 520, Alliant FX/80, and PARLOG and Fortran
Algorithm(s) used:	Automatic program translation
Computational Techniques:	Generation of MIMD parallel programs from formal logic specifications
Goal of the Research:	To provide a high-level software environment for parallel programming.

TITLE: Parallel Constraint Satisfaction

Principal Investigator (s):

Michael J. Swain, University of Rochester

Paul R. Cooper

Christopher M. Brown

Contact Person:

Michael J. Swain

University of Rochester

Rochester, NY

Abstract of Research

Many problems in Artificial Intelligence can be posed as constraint satisfaction problems (CSP's), in which the possible labels for a set of variables must be determined subject to a set of constraints. One of these problems is the graph isomorphism problem, which we are interested in for object recognition. Sequential algorithms have been developed and analyzed for arc consistency, a method for eliminating local inconsistencies that cannot possibly take part in the global solution to the constraint satisfaction problem. We have developed a parallel version of this algorithm, a SIMD message-passing algorithm for arc consistency that obtains linear worst case speed-up over the best sequential algorithm. We are implementing this on the CM-2 and plan to test its performance on large constraint satisfaction problems.

Scientific Field of Study:

Machine Vision (NSF 343; CR I.4)

Architectures/Languages:

Connection Machine/ C*

Algorithm(s) used:

Arc consistency in constraint satisfaction problems

**Computational
Techniques:**

Data parallelism, network communication

Goal of the Research:

Real-time object recognition.

TITLE: Coupled Atmospheric-Ocean-Sea Ice Modelling of the Tertiary Cooling Trend

Principal Investigator (s):
Jozef Syktus and L. A. Frakes

Contact Person:
Jozef Syktus
Geology & Geophysics
Adelaide University
Adelaide 5001
South Australia

Abstract of Research

We propose to study numerically climate changes using a mathematical model. The influence of the various forcing parameters such as carbon dioxide concentration changes, continental movement through time, etc. will be investigated on the past and future climates. The climate model will be tuned and made to work properly in the area of coupling between ocean and atmospheric circulation.

Scientific Field of Study: Climate Dynamics (NSF:514; CR: J2)

Architectures/Languages: Alliant FX/80 and Fortran

Algorithm(s) used: Numerical Solution of Coupled Differential Equations

Computational Techniques: Parallel Processing techniques on the Alliant

Goal of the Research: Reproduce past climates from the last several centuries and future predictions on a global scale.

TITLE: Exploring the use of Parallelism in the Development of Constructive Computational Theories of Social Phenomenon.

Principal Investigator (s):
Stuart J. Thorson, Syracuse University
James P. Bennett, Syracuse University
Robert Wolfson, Syracuse University

Contact Person:
Stuart J. Thorson
313 Maxwell Hall
Syracuse, NY 13244
315-443-2306

Abstract of Research

We are continuing work on exploring the use of parallelism in the development of constructive computational theories of social phenomenon, along two fronts. The first involves the use of the Connection Machine to investigate properties of various social choice mechanisms. An example of a social choice problem is the election of a president. The basic idea behind these problems is that some sort of aggregation function is applied to a collection of individual preference orders to identify some "best" social order. These sort of problems are typically investigated formally by showing that certain aggregation functions do or do not exist which meet certain conditions. The proofs are in general existence proofs (and not constructive proofs) perhaps in part because of the computational complexity involved. Thus we are using the CM to "reduce" the computational complexity by letting each processor contain the preferences of one agent. Our second line of research has been to use a theorem proving program (ITM) to try to prove theorems about problems such as the one discussed above.

Scientific Field of Study: Political Science
(NSF:455; CR:J.4).

Architectures/Languages: *Lisp on the CM, Lisp on the Multimax.

Algorithm(s) used: Showing that aggregation functions (which meet certain conditions) do or do not exist, proving theorems about social choice problems.

Computational Techniques: Massively parallel computation (letting each processor represent one agent), parallel theorem proving.

Goal of the Research: Investigation of constructive computational theories of social phenomenon.

TITLE: Computational Algorithms for Solving Unsteady Nonlinear Problems in Fluid Dynamics.

Principal Investigator (s):
Daniel T. Valentine, Clarkson
A. Gaber Mohamed, NPAC

Contact Person:
Dr. Daniel T. Valentine
Clarkson University
Potsdam, NY 13699
315-268-2204

Abstract of Research

The aims of the proposed project are: 1) to develop high-order, vectorized computational algorithms that take advantage of parallel architectures to solve unsteady nonlinear problems in fluid dynamic and, in particular, flow problems that are modeled by the Navier-Stokes equations. 2) To apply the computational methods to predict the development of flow separation over asymmetric, afterbody geometries. 3) To extend the method to handle afterbody flows with shear-layer control (i.e. with suction and/or blowing at the flow boundaries). 4) To pay particular attention to post processing time-dependent solutions in order to study the development of flow separation, the transients associated with an instantaneous imposition of a shear-layer control device, among other unsteady flow problems).

Scientific Field of Study: CFD Algorithm Development.
(NSF: 342; CR: J2)

Architectures/Languages: Alliant FX/80 Fortran.

Algorithm(s) used: Navier-Stokes equations.

Computational Techniques: Simulation, where at each iteration ELLPACK is used to solve linear elliptic partial differential equations. (ELLPACK was modified to run on the Alliant FX/80 by the investigators to use FX/BLAS and PARALIN.)

Goal of the Research: Development of a computational algorithm to solve time-dependent nonlinear applied mechanics problems.

TITLE: Parallel Sorting with Serial Comparators.

Principal Investigator (s):

Pramod K. Varshney,
Syracuse University

Contact Person:

Pramod K. Varshney,
ECE Department
Syracuse University

Abstract of Research

Traditionally, parallel sorting of n m -bit keys consider m -bit comparisons of keys as the basic unit of time. Thus, if we look at the input to a parallel sorting algorithm as an $n \times m$ array of bits (representing n m -bit numbers), traditional approaches operate on entire rows of the input at a time. This does not take advantage of the fine-grained nature of the input. We propose an approach which processes the input, column by column. The advantage of this approach is that the sorting time is not constrained by the size of the comparator in the processors. Two parallel algorithms based on this approach have been implemented on the Connection Machine. This first uses n processors to sort n m -bit numbers in $O(m \log n)$ time. The second algorithm requires n^2 processors and uses $O(\log m + \log n)$ time.

Scientific Field of Study: Parallel Algorithms.
(NSF: 345; CR: F2.1)

Architectures/Languages: *Lisp and Paris on the Connection Machine.

Algorithm(s) used: Parallel bit-representation sorting.

**Computational
Techniques:**

Storing n m -bit keys in a bit-representation on a fine-grained parallel machine, sorting by parallel one bit comparisons.

Goal of the Research: To design and analyze novel parallel data decomposition approaches to traditional problems.

TITLE: Fluid Simulation on the NPAC Connection Machine 2

Principal Investigator (s):
Stephen Vavasis

Contact Person:
Stephen Vavasis
Computer Science Dept.
4105C Upson Hall
Cornell University
Ithaca, NY 14853
607-255-9213
vavasis@cs.cornell.edu

Abstract of Research

A program was written in CM-Fortran to solve the potential flow equation for a single drop of fluid ejected from a pan of liquid. In particular, the boundary element calculations and the free surface calculations have been implemented and have passed some preliminary test. Mesh generation will be added and animation of fluid motion will be produced.

Scientific Field of Study: CFD and Parallel Algorithms

Architectures/Languages: CM-2, CM Fortran, CM C/Paris and Fortran/Paris

Algorithm(s) used: Boundary element method and free surface calculations.

Computational Techniques: Numerical integrations and solution of dense systems of linear equations using iterative methods that fit the CM architecture.

Goal of the Research: Solve engineering problems involving ink flows in printer devices.

TITLE: Parallel Execution of Some Power System Programs

Principal Investigator (s):

Y Wallach

Contact Person:

Y Wallach

CST, Rm 4-191

Syracuse University

Syracuse, NY 13244

315-443-4679

yehuda@top

Abstract of Research

A number of programs will be run on the Alliant FX/80 and the Multimax 520. These programs apply the numerical solution of very large, extremely sparse non-linear algebraic and of non-linear differential equations to problems of electric power systems (and some electronics circuits).

Scientific Field of Study:

Engineering Systems
(NSF:634; CR: J2)

Architectures/Languages:

Alliant FX/80, Multimax 520 and the C language

Algorithm(s) used:

Various standard numerical algorithms for sparse non-linear algebraic and differential equations

Computational Techniques:

Vectorization, Multiprocessing, ASP processing

Goal of the Research:

To use parallel computing to make possible the analysis and optimization of larger electrical power systems than is now possible.

TITLE: Radar Image Clutter Cancellation

Principal Investigator (s):
Taoyun Wang

Contact Person:
Frank Frazekas
Adaptive Technology Inc.
309 Curtis St.
Syracuse, NY 13208
315-475-1121

Abstract of Research

Under a contract with RADC, Adaptive Technology, Inc. is developing high performance radar jammer and clutter cancellation algorithms using the so-called Space-Time method. Part of the computer code has been developed in a DEC MicroVax but the computational capacity of the MicroVax virtually hampers further development of the code. We wish to move our code development to NPAC parallel machines, primarily the Alliant FX/80 and possibly the Connection Machine 2.

Scientific Field of Study: NSF: 343; CR: I4

Architectures/Languages: Alliant FX/80 and Fortran, and the Connection Machine 2 and CM Fortran

Algorithm(s) used: Space-Time method

**Computational
Techniques:**

Goal of the Research:

TITLE: Modifications to the APL2 Language to Support Expression and Parallel Execution of Generalized Algorithms.

Principal Investigator (s):
Mr. Robert G. Willhoft, IBM Corporation
Daniel J. Pease, Syracuse University

Contact Person:
Mr Robert G Willhoft
1320 Main St.
Vestal, NY 13850
607-748-8386
willhoft@suzeus or
nacs791@sutcase

Abstract of Research

The APL language has received attention in the last several years as a language that can be used to express parallel algorithms. The primary focus of interest has been its ability to express algorithms on vector or array arguments directly without the need for the programmer to convert them into sequential loops as in most languages. Despite this very significant processing advantage, both APL and APL2 do not handle a number of other significant parallel structures that are necessary for efficient, general parallel behavior. APL tends to be very good at expressing fine grain parallelism, but has few facilities to express course grain parallelism or data flow structures. This research focuses on some of the enhancements that would be necessary to the APL2 language to provide some of these missing features. We are working on implementing critical subsections of the new language on various parallel architectures as a feasibility study.

Scientific Field of Study:	Parallel Languages (NSF 314; CR D.3)
Architectures/Languages:	CM, Multimax, Alliant.
Algorithm(s) used:	Interpreter for a MIMD language with array operations
Computational Techniques:	Data flow analysis, vectorization, data parallelism
Goal of the Research:	To design a parallel language that enables the programmer to better express both MIMD and SIMD parallelism.

TITLE: Extracting Program Structure Knowledge from Fortran Programs
for Use in Converting to Parallel Algorithms

Principal Investigator (s):
Mr Robert G Willhoft, IBM Corporation
Daniel J Pease, Syracuse University

Contact Person:
Mr Robert G Willhoft
1320 Main St.
Vestal, NY 13850
607-748-8386
willhoft@suzeus or
nacs791@sutcase

Abstract of Research

There has been a great deal of research done on the automatic conversion of FORTRAN programs to run on parallel machines. This work includes a broad spectrum of work ranging from vectorizing compilers to the extraction of complete data flow algorithms from sequential algorithms. Much of this work focuses on the conversion of FORTRAN to a specific machine. However, most of these methods concentrate on extracting static data from the program. In doing this, much of the possible parallelization is lost because of data dependencies that are unknown to the static program analyzer. This means that in any case the worst possible assumptions must be made. On the other hand, dynamic analysis of a program is time consuming and subject to misleading information due to the problem of selecting the correct input data. We claim, however, that by extracting knowledge from the FORTRAN code it should be possible to gain most of the advantages of dynamic analysis without actually having to run the code.

Scientific Field of Study: Languages for parallel computation
(NSF 314; CR D.3)

Architectures/Languages: CM, Alliant, Multimax.

Algorithm(s) used: Data flow analysis

**Computational
Techniques:** Serial compilation for parallel execution

Goal of the Research: To make parallel computers easier to
program

TITLE: Implementation of the Image Algebra on the Connection Machine.

Principal Investigator (s):
Joseph N. Wilson, University of Florida

Contact Person:
Joseph N. Wilson
Gainsville, FL 32611

Abstract of Research

The Image Algebra Project is an algebraic system developed to support the specification and implementation of computer vision algorithms. The Image Algebra is high-level notation supporting image and generalized template operands, and providing a variety of operations that directly manipulate images and templates. The first phase of the work has been the implementation of the Image Algebra in *LISP. The goal of this system is that the user will be able to utilize the power of the Connection Machine's massively parallel architecture without having to understand the underlying implementation of Image Algebra operations. We are currently looking into different approaches to implementing image algebra operations not considered in the previous work. We are considering efficient schemes for representation of images over sparse point sets. We are also working on the implementation of an algorithm to perform stereo segmentation.

Scientific Field of Study: Machine Vision (NSF:343; CR:I.4).

Architectures/Languages: *Lisp on the Connection Machine.

Algorithm(s) used: Macros and functions for generalized template operations .

Computational Techniques: Data parallelism where each processor is mapped to an individual pixel, nearest-neighbor communication.

Goal of the Research: Provide general image software on the Connection Machine.

TITLE: An Experimental & Theoretical Study of the Dynamics of a Simple Neural System

Principal Investigator (s):
Raimond L Winslow, University of
Minnesota

Contact Person:
Raimond L Winslow
Room 6-255 Millard Hall
Minneapolis, MN 55455
612-626-5475
rai@neurr.med.umn.edu

Abstract of Research

A fundamental problem in neuroscience is to understand the ways in which non-linear, voltage-dependent conductances determine the dynamics of both single neurons as well as networks composed of such neurons. We are using the horizontal cell network of the turtle retina as a model system for studying this problem. Data obtained from physiological studies of voltage-dependent conductances in isolated turtle horizontal cells are being used to construct time-domain models of individual cells. Large arrays of model horizontal cells are being studied by mapping each cell in an $N \times N$ network model onto a single CM2 processor. The functional role of different voltage-dependent membrane conductances are determined by computing model responses to various light stimuli, and comparing these light-evoked responses to those measured in the intact retina. Substantial progress has been made on this project (due in part to the support provided by NPAC), and we have been able to use our model on the CM2 to make specific predictions concerning the ways in which receptive field properties change as a function of light and dark adaptation.

Scientific Field of Study:	Neurobiology (NSF 447; CR J.3)
Architectures/Languages:	Connection Machine
Algorithm(s) used:	Direct simulation of neurons
Computational Techniques:	Data parallelism, NEWS network for local communication
Goal of the Research:	To better understand neural systems

TITLE: Parallel and Vector Computing for Nonlinear Network Optimization

Principal Investigator (s):
Stavros A. Zenios, The Wharton School

Contact Person:
Stavros A. Zenios
University of Pennsylvania
Room 1310, Steinberg Hall-
Dietrich Hall
Philadelphia, PA 19104-6366
215-898-6727

Abstract of Research

Nonlinear programming problems with network constraints arise in a wide range of engineering, management, statistical, economic and other applications. Such problems are typically characterized by their large size that may extend to thousands of variables and constraints. Our goal is to increase the size of solvable models and their domain of application via the design and implementation of suitable parallel algorithms and their testing with real data on parallel computers. Our research is to design and implement distributed relaxation algorithms for a massively parallel computer. This involves a mapping of the network topology to the interconnection pattern of the computer and resolving issues like algorithm partitioning and synchronization. We have placed emphasis on communications, as well as computational, efficiency.

We have implemented the matrix balancing algorithms on the Connection Machine and have achieved some impressive performance statistics. For example, in updating a specific 500 x 500 matrix, solved to a high degree of accuracy, a Vax-8700 took 1 hour and 20 minutes, a Cray XM/P took 1.5 minutes, and the CM-2 took only 12 seconds. We are now looking into problems of simulating complex financial situations.

Scientific Field of Study: Operations Research (NSF 641; CR G.1.6)

Architectures/Languages: Connection Machine/ C/Paris

Algorithm(s) used: The relaxation algorithm for nonlinear network problems.

Computational Techniques: Mapping nodes of a network to Connection Machine processors so as to minimize communication times, parallel linesearch involving bracketing procedures.

Goal of the Research:

Optimizing nonlinear networks with millions of nodes; for example, air traffic control systems, telecommunications networks, or cashflow management.

TITLE: Sparse Matrix Algorithms for Multiprocessors

Principal Investigator (s):
Earl Zmijewski

Contact Person:
Earl Zmijewski
6626 Picasso Rd. #23
Goleta, CA 93117
805-968-2283
linchic@cornu.ucsb.edu

Abstract of Research

I am going to develop and implement SIMD parallel matrix algorithms for the Connection Machine 2. I plan to work on single-value decomposition and Shur decomposition using a Joccobi Iterative algorithm. I also am going to develop an algorithm for the Cholesky factorization of Space matrices. This algorithm utilizes elimination trees.

Scientific Field of Study: Parallel mathematical Algorithms (NSF: 342; CR: G4)

Architectures/Languages: Connection Machine 2

Algorithm(s) used: Joccobi iteration and Elimination tree

Computational Techniques: Data parallel, Grid communication

Goal of the Research: Increase the class of SIMD numerical techniques.

Appendix B - Selected User Research Bibliography

This is a bibliography of many of the technical papers published by NPAC users.

Selected User Research Bibliography

Baillie, C. F., Brickner, F.G., L. Johnson, R. Gupta, "QCD with dynamical fermions on the Connection Machine," *Proceedings of Supercomputing 89*, Reno, Nevada, New York, 1989.

Borer, P. N., Bishop, K.D., Heffron, Hoffman, R. E., Levy, G.C., and Wang, S.S., "Evaluation of Contrast Enhancement Methods for 2D NMR Spectroscopy," Abstract in *30th ENC, Asilomar, CA, USA*, 4/1989.

Bryant, R. E., "Formal Verification of Memory Circuits by Switch-Level Simulation," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 1990.

Carlson, E., Rutenbar, R., "Design and Performance Evaluation of New Massively Parallel VLSI Mask Verification Algorithms in JIGSAW," *Proc. ACM/IEEE Design Automation Conference*, June 1990.

Criske, D. M., "Development and Performance of a Two-Dimensional Panel Program on a Massively Parallel Computer," *AIAA Northeast Regional Student Conference*, West Point, NY, April 1990.

Elmohammed, M.A. Saleh, *The Effect of Category Boundary and Non-boundary Input on the Learning Time and Size of a Neural Net Classifier*. Phd thesis (in progress), Syracuse University, CIS, Syracuse, New York, August 1989.

Gupta, R., "Hadron Spectrum from the Lattice," invited talk at workshop, *Lattice 89*, Capri, Sept, 1989, to appear in *Nucl. Phys. B (Proc Suppl.)*, 1990.

Ganguly, Shantanu and Pitchumani, Vijay, *Compaction of a Routed Channel on the Connection Machine*. Paper presented at the 26th ACM/IEEE Design Automation Conference, 1989.

Hachem, N. I., *Dynamic Key-Ordered File Structures for the Random and Sequential Access to Very Large Data/Knowledge Bases*. Phd Dissertation, Syracuse University, ECE, Syracuse, New York, December 1988.

Hachem, N.I., and Berra, P.B., "Key Sequential Access Methods for Very Large Files Derived from Linear Hashing," *Proceedings of the Fifth International Conference on Data Engineering*, pp. 305-312, February 1989.

Hardy, Barry, "Utilizing Parallelism in Molecular Simulations of Flexible Polymer Models," *Parallel Computing News* Vol. 3 (1) PP.4-8, February 1990.

Hardy, Barry, a movie on molecular simulation work presented at the American Chemical Society (Division of Carbohydrate Chemistry) National Meeting, Dallas, Texas, 4/1989.

Hariri, Salim, Mohamed, Gaber, and Mutlu, Hasan, "Modeling Availability of Parallel Computers," to be published in a special issue of *IEEE on Experimental Reliability*.

Hoffman, R.E., Kumar, A., Bishop, K.D., Borer, P. N., Levy, G.C., "Application of the Maximum Likelihood Method to a Large 2D-NMR Spectrum Using a Parallel Computer," *J. Magn. Reson.* (in press).

Hoffman, R.E., "Impact of Parallel Computers on Magnetic Resonance Spectroscopy and Imaging," *NMRINK, New Methods Research Inc.*, Vol. 1 (2), 5, 1988.

Hoffman, R. E. and Levy, G. C., "Optimization of Magnetic Resonance Imaging and NMR Data Processing with the Alliant FX/80," *Parallel Computing News*, Vol. 1 (5), 3 1988.

Hoffman, R. E., "Phasing of Spectra Past and Future," *NMRINK, New Methods Research Inc.*, Vol 2 (2), 3 1989.

Hoffman, R.E. and Levy, G.C., "Spectral Deconvolution by Simulated Annealing," *J. Magn. Reson.* (in press).

Hoffman, R. E. and Levy, G. C., "Advanced NMR Data Processing Using Graphics Supercomputer Workstations," Abstract in *Ninth International Meeting on NMR Spectroscopy, Warwick, U.K., 7/1989*.

* Kayssi, A.I., Sakallah, K.A., "Automatic Test Pattern Generation on the Connection Machine."

Kravitz, S.A., Bryant, R.E., Rutenbar, R.A., "Logic Simulation on Massively Parallel Architectures," *Proc. ACM/IEEE International Symposium on Computer Architecture*, May 1989.

Kravitz, S.A., Bryant, R.E., Rutenbar, R.A., "Massively Parallel Switch Level Simulation: A Feasibility Study," *Proc. ACM/IEEE Design Automation Conference*, June 1989.

Kravitz, S.A., "Massively Parallel Switch Level Simulation: A Feasibility Study," Ph.D. thesis, Department of Electrical and Computer Engineering, 1989.

* Lui, E.M., W-P Zhang, "A Parallel Frontal Solver for Large Scale Structural Analysis."

Levy, G. C., Macur, A., Begemann, J., Crowther, M. and Delaglio, "Off-line NMJR Data Processing for Nicolet and GE NMR spectrometers," *Relaxation Times*, Vol. 5, 4 1988.

Levy, G. C., "How Changes in Computer Technology are Revolutionizing the Practice of Chemistry," *J. Chem. Info. Comp. Sci.*, Vol. 28, 167 1988.

Levy, G. C., "A Comparison of Workstation Computers Used in Chemical Design," *Chem. Design Auto News*, Vol. 3 (11), 6 1988.

Levy, G. C., "Hierarchical Computer Networks for the Spectroscopy Laboratory," *Spectroscopy*, Vol. 3, 30 1988.

Levy, G. C., "Computers in Magnetic Resonance Spectroscopy," *Computers in Physics*, 32 1988.

Levy, G. C., Delaglio, F., Sole, P., Grahn, H., Begemann, J. and Hoffman, R. E., "Data Processing in NMR Spectroscopy: a Look to the Future," Abstract in UCLA Symposia, Park City, Utah, USA, 1/1989.

Levy, G. C., Borer, P.N., and Wang, S.S., "Significant Challenges in NMR Data Processing Protein and Other Biopolymer Structure Determination," Abstract in Ninth International Meeting on NMR Spectroscopy, Warwick, UK, 7/1989.

Mackie, Terry L., *An Implementation of the Teachable Letter Recognizer on the Connection Machine*. MS thesis, New Jersey Institute Of Technology, CIS, New Jersey, May 1989.

Mayor, P., Pitchumari, V. and Narayanan, V., "Parallel Automated Test Pattern Generation on the Connection Machine," Presented at the 20th International Test Conference, Washington, D.C., August, 1989.

Michaud, M. C., M.A. Schroeder, "Parallel Processing Performance Evaluation," *Supercomputing '89*, Reno, Nevada, November 1988.

Mohamed, A.G., Valentine, D.T. and Hessel, R.E., Users Guide to Some Useful Mathematical Software. *Report Number MIE-185*, Clarkson University, Potsdam, New York, May 1989.

Mohamed, A. G., Valentine, D. T., and Hessel, R. E., "Numerical Study of Laminar Separation Over an Annular Backstep," in press at *Computers and Fluids*, June 1988.

Mohamed, A.G., and Valentine, D.T., "Direct Simulation of Navier-Stokes Flow," The 6th Iroquois Conference, Moravia, New York, March 17-19, 1989.

J.E. Mower, "Implementing Automated Cartographic Name Placement on the Connection Machine," Association of American Geographers National Conference, Toronto, 1990.

J.E. Mower, "Line Feature Classification through Competitive Learning," submitted to *The American Cartographer (Cartography and GIS)* for publication in a special issue on analytical cartography, scheduled for publication January 1991.

J.E. Mower, "A Comparison of Statistical and Adaptive Pattern Recognition Approaches to Line Feature Identification," *proc. 4th International Symposium on Spatial Data Handling*, Zurich, 1989.

J.E. Mower, "Automated Cartographic Name Placement Using Connectionist Computing Environments," *Proc. 14th International Cartographic Conference*, Budapest, 1989.

Narayanan, V. and Pitchumani, V., "A Massively Parallel Algorithm for Fault Simulation on the Connection Machine," Paper 41.8, 26th ACM/IEEE Design Automation Conference.

* Potter, D. J., M.P. Cline, "Parallel Algorithms for 2D Kalman Filtering," submitted to *Frontiers of Massively Parallel Computation '90*.

* Sarkar, T. K., T.R. Arabi, R.F. Harrington, A.R. Djordjevic, "Analysis of Arbitrarily Oriented Microstriplines Utilizing a Quasi-Dynamic Approach."

Shin, D., Berra, P. B., "Surrogate File Approach to Managing First Order Terms in Secondary Storage," *Proceedings of the SPIE Conference on Applications of Artificial Intelligence VII*, pp. 1051-1062, April 1989.

Vaidyanathan, R., Hartmann, C.R.P., and Varshney, P.K., "The Bit-PRAM: A New Model for Parallel Computation," Submitted for presentation at the 1989 IEEE Symposium on Foundations of Computer Science.


Valentine, D.T., and Mohamed, A.G., "High-order Navier-Stokes Solver and a New Test Problem," *Bull. Am. Phys. Soc.*, Vol. 33, No. 10, 2265, 1988.

Valentine, D.T., and Mohamed, A.G., "Taylor's Vortex Array: A New Test Problem for Navier-Stokes Codes," (submitted for publication).

Appendix C - An Overview of NPAC

This document, written for potential NPAC researchers, also serves as a brief summary of all of the center's activities during its third and fourth years.

An Overview of NPAC



Northeast Parallel Architectures Center

C-iii/C-iv

An Overview Of NPAC

November, 1989

Copyright © 1989, NPAC. All rights reserved.
Northeast Parallel Architectures Center
111 College Place
Syracuse, NY 13244
(315) 443-1722
npac@nova.npac.syr.edu

Ada is a registered trademark of the U.S. Government, Ada Joint Program Office (AJPO).

Alliant, Concentrix, FX/Ada, FX/FORTRAN, and FX/80 are trademarks of Alliant Computer Systems Corporation.

Connection Machine is a registered trademark of Thinking Machines Corporation. C*, *Lisp, Paris, CM-1, CM-2, and DataVault are trademarks of Thinking Machines Corporation.

Franz Lisp is a trademark of Franz, Incorporated.

Multimax, UMAX, and Encore Parallel Threads (EPT) are trademarks of Encore Computer Corporation.

NeWS is a trademark of Sun Microsystems, Incorporated.

Symbolics is a trademark of Symbolics, Inc.

ULTRIX, VAX, and VMS are registered trademarks of Digital Equipment Corporation.

UNIX is a trademark of AT&T Bell Laboratories.

Preface

Audience

This overview was written for prospective users of the Northeast Parallel Architectures Center.

Purpose

We have provided this document to introduce the center and explain the services and facilities we offer.

Overview

Section 1 outlines NPAC's facilities and services, who uses them, and the types of research projects that best utilize NPAC's parallel architectures and resources.

Section 2 details NPAC's parallel computer architectures and describes applications for each.

Section 3 provides an overview of the software available at NPAC.

Section 4 describes NPAC's activities and support services.

Section 5 outlines NPAC's policies.

Section 6 explains the way in which resource units are determined.

Section 7 details the procedures for applying for NPAC resources.

Section 8 consolidates the essential NPAC names, telephone numbers and addresses given throughout this document into **Table 8-1, Communicating with NPAC**.

Conventions

In examples of characters you type, your entry is printed in **bold sans-serif** type.

Table of Contents

Section	Page
1/Getting Acquainted	1
2/Computers/Architectures/Applications	3
3/Software Available at NPAC	9
4/Activities/Support Services	11
5/Policies	15
6/Resource Units	17
7/Requesting Resources	19
8/Communicating With NPAC	26

List of Tables

Table	Page
3-1/Software	9
6-1/Resource Unit Equivalents	17
8-1/Communicating with NPAC	26

Section 1/Getting Acquainted

What is NPAC?

**a Center Offering
Parallel Computing
Resources
and Services**

NPAC is a parallel computing center that operates under contract to the Rome Air Development Center, Rome, New York, with funding from the Defense Advanced Research Projects Agency (DARPA) in Washington D.C. Established in 1987 at Syracuse University, NPAC supports research in the operation, application, and improvement of parallel computers. By providing the latest commercially available hardware and software, including a full set of common languages, NPAC assists researchers in many fields.

Who Uses NPAC?

**Academic,
Corporate, and
Government
Researchers**

The NPAC user community is presently comprised of research and teaching faculty, as well as graduate and undergraduate students. Scientists and engineers working in applied and theoretical sciences also make up a large part of this community.

Qualified researchers in academic, corporate, and government institutions across the nation may use these facilities. Researchers in academic and government laboratories (funded by the National Science Foundation and DARPA) may use them free of charge, subject to review. The NPAC Corporate Program offers corporate researchers access to these facilities and services by arrangement.

What Type of Research is Done at NPAC?

**Artificial
Intelligence,
Computational
Fluid Dynamics
and More**

Nuclear magnetic resonance, molecular dynamics, computational fluid dynamics, real-time imaging, vision and image synthesis are just a few areas in which scientists are currently using NPAC resources for applied research. NPAC has a particular interest in allocating resources to researchers in artificial intelligence, logic programming, signal processing, image processing, information retrieval, algorithm evaluation, VLSI design, and parallel architecture evaluation and measurement.

The NPAC Environment:

**SIMD Machines:
the
Connection
Machines,
MIMD
Machines:
the Encore
Multimax,
and the
Alliant/FX
Series**

The NPAC environment offers several distinct kinds of parallel computer architectures from several vendors. Each architecture (described in detail below) is suitable for a variety of applications in parallel computing. These architectures include massively parallel, single-instruction, multiple-data (SIMD) systems. The two Connection Machines represent the SIMD architecture, providing high-performance hosts, mass storage facilities, and graphic displays. Multiple instruction, multiple data (MIMD) architectures include bus-based, shared memory systems represented by the Encore Multimax and the Alliant/FX shared cache computers.

**UNIX
Environment**

All of NPAC's parallel computer architectures may be used in a UNIX-based or UNIX-compatible environment.

**Networking
Environment**

NPAC's networking environment allows users to work on any of our computers by using the Internet.

**Documentation
and Training**

NPAC provides software, documentation, training, and expert consultation to assist researchers in the use of advanced parallel computing techniques. NPAC also has an Invited Lecture Series and a Sponsored Seminar series, maintains a lending library, and supports user groups for various architectures and research interests.

Section 2/Computers/Architectures/Applications

Each of NPAC's architectures is a commercially available computer with a system of hardware and software that supports a distinct parallel computing environment. These architectures and the types of parallel computing they support are described below.

The Alliant FX/80

Alliant's FX/80 is a multiple-instruction, multiple-data (MIMD) minisupercomputer offering a variety of parallel programming techniques. The FX/80 can execute multiple jobs by scheduling one job per processor. Programs may also be executed concurrently on several processors which have been grouped together to perform as a single resource. In addition, the FX/80 features a FORTRAN compiler described below.

NPAC System Configuration: Hardware

The FX/80 has 128 megabytes of shared memory. It comes equipped with three types of computing resources: six interactive processors (IPs), eight advanced computational elements (ACEs) that have vector registers and a floating point arithmetic pipeline, and a parallel computational complex that groups a set of ACEs. Each computing resource processes jobs with a time-sliced, round-robin technique. It also has a concurrency control bus and two types of caches: an interactive processor cache (IPC) and a computational processor cache (CPC).

Software Languages

The FX/80 runs Ada and versions of C, and FORTRAN (FORTRAN 8x).

Operating Systems

It runs Concentrix, a proprietary operating system, based on 4.3 BSD UNIX.

**FORTTRAN
Compiler**

The FX/80 has an optimizing FORTRAN compiler that works with the concurrency control bus and automatically converts sequential FORTRAN to parallel and vectorized FORTRAN.

**Application
Examples**

The FX/80 can be used for a wide variety of applications. You can explore parallel programming in Ada, C, or FORTRAN on the FX/80, or if you have already written a FORTRAN program for a serial processor, you can use the FX/80's FORTRAN compiler to automatically make your program concurrent and vectorize it.

Documentation

Three manuals on the FX/80 are *FX/SERIES Architecture Manual*, *FX/FORTRAN Language Manual*, and *FX/FORTRAN Programmers' Handbook*.

**The Connection
Machines**

NPAC operates two single-instruction, multiple-data (SIMD) Connection Machines, models CM-1 and CM-2, from Thinking Machines Corporation. The CM-1 and CM-2 utilize low-cost LSI technology to provide massively parallel computational capabilities for applications ranging from large-scale scientific computation to artificial intelligence systems.

Each Connection Machine provides two paradigms for interprocessor communication: nearest-neighbor and general communication. Both increase computational throughput dramatically.

**NPAC System
Configuration:
Hardware**

The Connection Machines share a VAX 8800 front end. The CM-2 has a Symbolics Lisp front end that provides an efficient, single-user environment. The CM-1 has 32K processors with 4K bits of local memory per processor. The CM-2 has 32K processors with 64K bits of memory per processor plus 1024 floating-point accelerators, and is capable of an aggregate speed of 5 gigaflops. The two 5-gigabyte data storage units of the CM-2 provide data transfer at up to 40 megabytes per second. Two frame buffers offer a high resolution display from the Connection Machine of 1280 by 1024 pixels, each with a 24-bit color value.

**Software
Languages**

Both the CM-1 and the CM-2 run extended versions of C and Lisp (C* and *Lisp) and Parallel Instruction Set (Paris), the assembly language for the Connection Machine. The CM-2 also runs CM FORTRAN, a version of FORTRAN with many 8x and CM-specific extensions.

Libraries

NPAC currently supports libraries for common mathematical functions. Users can access these libraries using all of the languages mentioned above.

Operating Systems

The VAX front end runs ULTRIX 3.1, a 4.3 BSD-based system. The Symbolics runs Genera version 7.2.

Application Examples

Current applications on the Connection Machines include research projects in computational fluid dynamics, document retrieval, seismic processing, circuit optimization, and circuit simulation.

Documentation

Prospective users will find the following documentation helpful: *Connection Machine Parallel Instruction Set*, *Connection Machine I/O Programming*, *Connection Machine Programming in *Lisp*, *Connection Machine Programming in C**, *Connection Machine Graphics Programming*, and *Connection Machine Programming in FORTRAN*.

The Multimax

Encore Computer Corporation's Multimax is a shared-memory, bus-based, MIMD machine offering an environment for exploring new parallel programming methods. The Multimax 320 runs the Mach distributed operating system and the Multimax 520 runs the UMAX operating system. Each operating system provides access to a different programming environment and library (described below).

**NPAC System
Configuration:
Hardware**

The Multimax 520 has 16 XPC processors, 128 megabytes of shared memory and is rated at an aggregate speed of 32 MIPS. The Multimax 320 has 20 APC processors, 40 MIPS and 64 megabytes of shared memory.

**Software
Languages**

Programming languages for both Multimaxes include C, Ada, FORTRAN, and Lisp. In addition, the Multimax 520 runs Multilisp. Its programming languages include Encore Parallel FORTRAN, Encore's version of FORTRAN with built-in library calls.

**Operating
Systems**

The Multimax 520 runs a version of UMAX and, under a special agreement, the Multimax 320 runs the Mach distributed operating system. Mach was initially designed at Carnegie-Mellon University. It is now being refined with a particular focus on increasing its parallel capabilities by Encore. Both UMAX and Mach have UNIX-like user interfaces.

**C Threads
Library**

The Mach distributed operating system has a C Threads Library that provides a C language interface for manipulating threads of control and shared memory. The constructs of the C Threads Library are used for creating and using tasks and threads, protection of critical regions of code with mutual exclusion mutex variables, and synchronization of tasks.

Encore Parallel Threads

Like Mach, the UMAX operating system has a threads library called Encore Parallel Threads (EPT). It offers a complete parallel programming environment. EPT allows you to create and use threads, and to synchronize the use of shared memory.

Application Examples

Although the Multimax is well-suited for most MIMD applications, it is particularly useful for exploring new parallel programming methods. Currently, applications on the Multimax include language implementation and practical mathematical studies.

Documentation

Documentation that will help you in programming the Multimaxes are *Encore Parallel Threads Manual*, *Multimax Technical Summary*, *Encore Parallel FORTRAN Manual*, *A Programmer's Guide to the Mach System Calls*, *A Programmer's Guide to the Mach User Environment*, and *C Threads*.

Section 3/Software Available at NPAC

Table 3-1 below lists the software available at NPAC on each of the machines.

SOFTWARE	MACHINES			
	Alliant FX/80	Encore Multimax (UMAX)	Encore Multimax (Mach)	Thinking Machines' CM (VAX Front end)
Operating Systems	Concentrix	UMAX	Mach	ULTRIX
Languages	FX/FORTRAN (FORTRAN 8x) FX/Ada	Encore Parallel FORTRAN FORTRAN 77 Parallel Ada Franz Lisp Multilisp	FORTTRAN 77 Parallel Ada	CM FORTRAN *Lisp Simulator *Lisp C* C/Paris
	C	C	C	
Windowing Systems	X11 NeWS			X11 NeWS
Editors	GNU Emacs vi	GNU Emacs vi	GNU Emacs vi	GNU Emacs vi
Libraries	EISPACK ELLPACK LINPACK QUADPACK ODEPACK	Encore Parallel Threads	C Threads	cm-lights, National CM Users' Library

Table 3-1 Software

Section 4/Activities/Support Services

NPAC has a full suite of activities and support services. They are listed and described below.

Consulting 9 to 5 Weekdays

NPAC's consulting staff is available every weekday from 9 to 5 (Eastern Time). You can reach them by electronic mail (preferred method) at consult@nova.npac.syr.edu or by telephone at (315) 443-1722.

Visiting Researcher Program

The Visiting Researcher Program provides an opportunity for researchers to come to NPAC and work in residence for periods of time. Visiting researchers have access to all NPAC facilities, local resources, services, staff, and colleagues.

Training Workshops and Training Institutes

NPAC has an ongoing education program tailored to your training needs. Classes are offered in specific architectures, languages, and applications. These classes range in length from one-day workshops to longer training institutes. Contact Betty LaPlante for more information on training workshops and institutes by sending electronic mail (preferred method) to training@nova.npac.syr.edu, or by calling her at (315) 443-1722.

Invited Lecture Series

NPAC hosts an Invited Lecture Series each academic semester that brings world-renowned leaders in parallel computing research to Syracuse. Lecturers in the past have included Geoffrey Fox of the California Institute of Technology whose talk was entitled "The Physical Structure of Concurrent Computers and Parallel Algorithms" and Daniel Hillis of Thinking Machines Corporation whose topic was "Using the Connection Machine to Study Evolution."

Sponsored Seminar Series

NPAC offers a Sponsored Seminar Series in which vendors, researchers, and developers present subjects of interest to the NPAC community. Speakers in the past have included Steve Morton of Oxford Computers who presented "Intelligent Memory Chips Multiply a Matrix Times Vector With Unprecedented Power, Flexibility, and Economy" and Frank McCabe of Imperial College, London, England, whose topic was "Parallel Architectures for Symbolic Computing."

User Groups

The NPAC User Group currently meets to discuss different architectures and research interests. NPAC also distributes Alliant User Society software. For more information about user groups contact consulting at consult@nova.npac.syr.edu.

Publications Training Manuals, Newsletter, User Guide

NPAC provides training manuals for those who participate in workshops and upon request. You can receive a free subscription to *Parallel Computing News*, the monthly newsletter, by sending your name and mailing address to editor@nova.npac.syr.edu (preferred method) or by calling (315) 443-1722. NPAC researchers automatically receive the *NPAC User Guide*, a guide for new users, and *Parallel Computing News*.

Technical Report Series

NPAC plans to make researchers' papers available in the future in the Technical Report Series. If you are interested in this series, send electronic mail to npac@nova.npac.syr.edu (preferred method) or call (315) 443-1722.

**Technical
Notes**

NPAC users periodically receive new Technical Notes on software, hardware, and application updates. These will eventually be incorporated into the next version of the *NPAC User Guide*. Users should keep these Technical Notes until they receive a revised version of the *NPAC User Guide*.

**Documentation
Manuals and
References**

NPAC has a lending library of manuals from vendors as well as a reference library (see "Libraries" for more information).

**Libraries
Lending**

NPAC maintains a library of technical documentation produced and distributed by computer vendors, available for short-term loan to researchers. This collection includes multiple copies of the most popular and widely used volumes, and pre-release notes for new features.

Reference

NPAC also keeps a reference library of current subjects and titles. General topic areas include architectures, languages, artificial intelligence, and parallel and distributed computing.

Backups

NPAC currently backs up all researchers' files on a regular schedule. Our consultants will help you restore any files that may have been accidentally deleted. Please contact us as soon as you need to recover your files at consult@nova.syr.edu. Include your name, userid, the name of the computer your files were on, the full pathname of the files/directories you need to restore, the date the files were created, and the date they were lost.

**NPAC
Corporate
Program**

The NPAC Corporate Program offers companies membership benefits that include low-cost use of NPAC parallel architectures, consulting, documentation, and participation in other activities supported by NPAC. For more information contact:

Dr. Elizabeth Schermerhorn
Deputy Director
NPAC
111 College Place
Syracuse, New York 13244-4100;

by electronic mail at ecs@nova.npac.syr.edu;
by telephone at (315) 443-1723

Section 5/Policies

NPAC distributes computer time according to the policies and approval of the Resource Allocation Board. Scientists whose research is funded by DARPA should apply for computer time directly through DARPA. Other researchers can apply for NPAC resources by completing a Request for Resources packet described in the "Requesting Resources" section below.

Policy Guidelines

1. NPAC resources are available to university, corporate and government researchers upon submission and evaluation of a Request for Resources.
2. A Request for Resources submitted by university researchers (either faculty, staff, or graduate students) consists of three copies of a completed Request for Account application for each member of a research team, an abstract of the research project, and a research proposal with optional supporting materials and references.
3. A Request for Resources for class instruction accounts submitted by a faculty member requires Instructional Request for Account applications for each student, faculty member, and teaching assistant. A letter describing the course, and demonstrating the requirements for the computing resources should accompany the rest of the materials submitted.
4. NPAC operates a testbed environment. Stability and longevity of computing resources is not guaranteed. However, NPAC expects to provide, where reasonable, compatible products and operating systems. NPAC will make practical accommodations to support existing research on discontinued hardware.

Policy Guidelines (continued)

5. NPAC does not charge researchers working under DARPA or NSF grants, or their institutions, for time on NPAC machines.
6. NPAC will charge corporations for CPU time on NPAC machines based upon the Corporate Membership Program guidelines.
7. The NPAC Resource Allocation Board reviews all university research proposals to determine the appropriateness of the request.
8. NPAC immediately awards a small, gratis start-up allocation upon submission of a completed Request for Resources.
9. The Resource Allocation Board allocates all NPAC resources in quantities of computer time, in amounts described as resource units (see "Resource Units" below.)
10. Researchers must acknowledge use of NPAC resources in all research-related publications and presentations. The wording of this acknowledgement should read:

"This work was conducted using the computational resources of the the Northeast Parallel Architectures Center (NPAC) at Syracuse University, which is funded by and operates under contract to DARPA and the Air Force Systems Command, Rome Air Development Center (RADC), Griffiss Air Force Base, NY, under contract # F306002-88-C-0031."

NOTE

As resources become scarce, NPAC reserves the right to change this policy.

Section 6 /Resource Units

NPAC awards computer time by resource units. A resource unit (r.u.) is approximately equivalent to one CPU hour on a Cray X-MP single processor. Table 6-1 below gives NPAC computer r.u. equivalents. As a reference, note that the VAX 8800 is rated at 0.11 r.u./system hour.

Resource Units	Computer	r.u./system hour
	CM-1	1.0
	CM-2	2.0
	Multimax 320	0.33
	Multimax 520	1.4
	Alliant FX/80	0.33

Table 6-1 resource unit equivalents

(For the convenience of those familiar with the National Science Foundation's system, one resource unit is equal to one service unit [1.0 r.u. = 1.0 s.u.])

NOTE

NPAC reserves the right to re-evaluate and change the r.u./system hour ratings for its computers following significant system upgrades.

Section 7/Requesting Resources

Procedures

The procedures for University researchers, those affiliated with Federal agencies, and Corporate researchers differ. Refer to the appropriate sections below.

University Researchers

University researchers, please complete all of the steps listed below. If you are a graduate student or a faculty member applying for class accounts, refer to the appropriate sections on the following pages.

NOTE

The Principal Investigator or Co-Principal Investigator must be a researcher with a faculty- or staff-level appointment.

University Researchers/Faculty and Staff

Step 1:
**Obtain Request
for Resources**

Obtain a Request for Resources packet from NPAC by contacting Lisa Deyo at (315) 443-1723, or by sending electronic mail to acct-req@nova.npac.syr.edu. This packet contains the forms and guidelines for completing the three parts of the Request for Resources: the Request for Account; the Project Abstract; the Research Proposal.

Step 2:
**Complete
Requests for
Accounts**

Complete a Request for Account application for each member of the research team. NPAC assigns a separate userid to each person working on a research team. You are also assigned a separate userid for each project you are working on. This is for accounting purposes.

Step 3:
**Write Project
Abstract**

Write an abstract of your project. This abstract should be no more than two paragraphs.

Step 4:
**Write Research
Proposal**

This proposal can be up to three pages long. It should describe the work you intend to do and demonstrate the need for parallel computing resources. Please include a description of the research objective and the computational and/or algorithmic strategy to be employed. You can include any applicable support documentation. It will accompany your proposal for review by the Resource Allocation Board.

Step 5:
Copy

Make three copies of each completed section of your Request for Resources.

**Step 6:
Mail Completed
Packet**

Mail three copies of your Request for
Resources to:

Lisa Deyo
NPAC
111 College Place
Syracuse, New York 13244-4100

For Help

If you have any questions, contact Lisa Deyo
at (315) 443-1723 or by electronic mail at
acct-req@nova.npac.syr.edu.

**Decision
Notification**

If the Resource Allocation Board awards a
start-up allocation, NPAC issues userids and
passwords to the researchers involved in the
project. Later Principal Investigators are notified of
the final award by U.S. mail.

University Researchers/Graduate Students

Graduate Students please complete all the steps listed below.

NOTE

The Principal Investigator or Co-Principal Investigator must be a faculty- or staff-level appointment.

Step 1:
**Obtain Request
for Resources**

Obtain a Request for Resources packet from NPAC by contacting Lisa Deyo at (315) 443-1723, or by sending electronic mail to acct-req@nova.npac.syr.edu. This packet contains the forms and guidelines for completing the three parts of the Request for Resources packet: the Request for Account; the Project Abstract; the Research Proposal.

Step 2:
**Complete
Requests for
Accounts**

Complete a Request for Account application for each member of the research team. NPAC assigns a separate userid to each person working on a research team. You are also assigned a separate userid for each project on which you are working. This is for accounting purposes.

Step 3:
**Obtain
Advisor's
Signature**

Obtain your advisor's signature on your own Request for Account application(s). NPAC will issue separate userids and passwords to you and your advisor for the research project(s). You will both be considered Co-Principal Investigators.

Step 4:
**Write Project
Abstract**

Write an abstract of your project. This abstract should be no more than two paragraphs.

**Step 5:
Write Research
Proposal**

This proposal can be up to three pages long. It should describe the work you intend to do and demonstrate the requirements for parallel computing resources. Please include a description of the research objective and the computational and/or algorithmic strategy to be employed. You can include any applicable support documentation. It will accompany your proposal for review by the Resource Allocation Board.

**Step 6:
Copy**

Make three copies of each completed section of your Request for Resources.

**Step 7:
Mail Completed
Packet**

Mail three copies of your Request for Resources to:

Lisa Deyo
NPAC
111 College Place
Syracuse, New York 13244-4100

For Help

If you have any questions, contact Lisa Deyo at (315) 443-1723 or by electronic mail at acct-req@nova.npac.syr.edu.

**Decision
Notification**

If the Resource Allocation Board awards a start-up allocation, NPAC issues userids and passwords to the researchers involved in the project. Later, Principal Investigators are notified of the final award by U.S. mail.

University Researchers/Class Accounts

Step 1:

Submit Letter and Course Description

Submit a letter describing your need to use NPAC's facilities and a course description to:

Steven Sather
Manager, Support Services
NPAC
111 College Place
Syracuse, New York 13244-4100

or by electronic mail to steve@nova.npac.syr.edu.

Decision Notification

After this material has been reviewed, you will be notified of the class instruction award. If your request has been approved, you will receive Request for Account forms for each student in the class.

Step 2:

Submit Request for Account Forms

Submit the Request for Account forms to the NPAC office.

NPAC informs the instructor of the userids and passwords for the entire class when they have been issued. It is the responsibility of the instructor to notify the class members of their userids and passwords, and to provide copies of documentation.

Special Notes

Each instructor should designate a teaching assistant who will be responsible for answering students' questions and for communicating with NPAC's consulting staff.

Course accounts are valid only for the academic semester in which the course is given.

At the end of the semester, instructors will be asked to send in an evaluation.

**Corporate
Researchers**

Contact:

Dr. Elizabeth Schermerhorn
Deputy Director
NPAC
111 College Place
Syracuse, New York 13421-4100;

by electronic mail at ecs@nova.npac.syr.edu;
by telephone at (315) 443-1723.

**Federal Agency
Researchers**

You may be eligible for resources at no charge
under NPAC contract with DARPA. For more
information contact:

DARPA/ISTO
1400 Wilson Boulevard
Arlington, VA 22209-2308;

by electronic mail at pllcomp@vax.darpa.mil.

Section 8/Communicating with NPAC

Table 8-1 lists the essential people to contact, telephone numbers, and e-mail addresses given in this Overview. Address U.S. mail to the appropriate NPAC staff members (listed below) at NPAC, 111 College Place, Syracuse, NY 13244-4100.

Information Needed	Person/E-mail Address/Number
consulting	consult@nova.npac.syr.edu (315) 443-1722
training	Betty LaPlante training@nova.npac.syr.edu (315) 443-1722
help with Request for Resources	Lisa Deyo lsd@nova.npac.syr.edu (315) 443-1723
help with class accounts	Steven Sather steve@nova.npac.syr.edu (315) 443-1722
newsletter subscriptions	editor@nova.npac.syr.edu (315) 443-1722
Technical Report Series	npac@nova.npac.syr.edu (315) 443-1722
corporate researchers (how to apply)	Dr. Elizabeth Schermerhorn ecs@nova.npac.syr.edu (315) 443-1723
federally funded researchers (how to apply)	DARPA/ISTO 1400 Wilson Boulevard Arlington, VA 22209-2308 pllcomp@vax.darpa.mil
general questions	Lisa Deyo lsd@nova.npac.syr.edu (315) 443-1723

Table 8-1 Communicating with NPAC

Appendix D - *Parallel Computing News*

These are the issues of NPAC's monthly newsletter cited in this report.

PARALLEL COMPUTING NEWS

NORTHEAST PARALLEL ARCHITECTURES CENTER AT SYRACUSE UNIVERSITY

VOLUME 3
August

NUMBER 8
1990

NPAC's 1990 Connection Machine Summer Institute

By Diane Purser

"The level of instruction was exemplary," said George Mesina, one of twenty researchers who attended NPAC's second annual Connection Machine Summer Institute (CMSI). Their interests and backgrounds varied widely. Eddie Paramore, a professor at Tuskegee University explored teaching parallel processing at the undergraduate level. Jo Ann Parikh, a computer science professor at Southern Connecticut State University, worked on an algorithm in *Lisp to be used in a program for detecting abnormalities in ultrasonic images. Mesina, a mathematician at the Idaho National Engineering Lab, came to the CMSI to develop parallel programs for solving computational fluid dynamics problems.

"We're very pleased with the broad range of research interests among the CMSI participants," said Steven Sather, manager of NPAC Support Services.

Topics covered during the first week included an overview of the Connection Machine (CM) architecture, and introductions to the three CM programming languages, CM Fortran, C*, and *Lisp. The second week concentrated on specific programming techniques and special features, including matrix multiplication, using the DataVault, and Connection Machine graphics. Lennart Johnsson, director of Computational Science at Thinking Machines Corporation

Continued P. 1



INSIDE THIS ISSUE

TMC SPEAKERS	2
FROM THE TRAINING NOTEBOOK	3
CMSI PARTICIPANT PROFILES	6

*The CMSI participants
in front of the Science and
Technology building*

NPAC director focuses on education, the future of parallel computing

By Diane Purser

"Computation has the power to change many fields," said NPAC's director Geoffrey Fox at the closing CMSI luncheon. Founder of the Syracuse Center for Computational Science, Fox plans to integrate computer science and parallel computing with a variety of application fields at Syracuse University. "We need to develop interdisciplinary people, not teams," he said.

He predicted that the development of the teraflop supercomputer, capable of trillions of floating-point operations per second, will be crucial to the future of parallel computing. "The trend in software development," he added, "will be to create software suited for particular computing problems, rather than according to hardware constraints."

Fox stressed his commitment to education in computational science for students at the undergraduate, master's, and PhD levels. "Computing is not a technical trade that you pick up on the job," he added. "This argument is valid at all levels of education." □

Lennart
Johnsson chats
with researchers
after his talk.



TMC speakers highlight special Connection Machine features and applications

By Diane Purser

The CMSI featured three guest speakers from Thinking Machines Corporation (TMC), Lennart Johnsson, Pietro Rossi, and James Frankel.

Johnsson, director of Computational Sciences, spoke on the Connection Machine Scientific Software Library (CMSSL). The library, which he described as "multilingual," can be called from all supported languages on the Connection Machine; CM Fortran, *Lisp, Paris, Lisp/Paris, C/Paris, and F77/Paris. Routines in the library include:

- matrix multiply;
- fft;

- matrix inversion;
- matrix inversion and solve;
- qr factorization;
- tridiagonal system solver;
- matrix vector multiply;
- random number generator;
- histogramming;
- and triangular system solver.

TMC scientist Rossi described several computationally intensive scientific applications that are particularly well-suited to the Connection Machine (CM). Currently running on the CM, these applications include problems in quantum chromodynamics; quantum chemistry; computational fluid dynamics; and solving large matrices.

Frankel, the leader of TMC's C* project, gave an in-depth presentation of the latest version of C*, a language based on ANSI C, which provides facilities to program data parallel computers efficiently. "This language is suitable for any computer you want to program as a SIMD machine," he said. Future goals include providing a more efficient way of storing and manipulating data (slice-wise storage), a variety of language extensions, and developing a standard C* with other computer vendors. 

The beta version of the CMSSL Release 2.0 will soon be available on NPAC's CM-2.



PARALLEL COMPUTING NEWS

Parallel Computing News is published by the Northeast Parallel Architectures Center (NPAC) at Syracuse University. The Northeast Parallel Architectures Center promotes and explores advanced computing technology by providing parallel architectures and research support to university, corporate, and government researchers nationwide.

To subscribe, send name and mailing address to: *Parallel Computing News*

NPAC
111 College Place
Syracuse, NY 13244

Editor

Diane W. Purser

e-mail address

editor@nova.npac.syr.edu

Phone

315-443-1722

Data parallel processing in CM Fortran at NPAC

Over the last two decades, FORTRAN has earned overwhelming acceptance as one of the essential tools for most engineering and scientific applications. CM Fortran is an implementation of FORTRAN 77 supplemented with array-processing extensions from FORTRAN 90. The vector constructs of FORTRAN 90 map naturally onto the data parallel architecture of the Connection Machine (CM) system, which is well-suited for computations on large data sets. These notes summarize the basic features of CM Fortran and other utilities that are made available by the NPAC Research Consulting group.

Components of CM Fortran

CM Fortran is designed to combine the familiarity of FORTRAN 77, the expressive power of FORTRAN 90, and the computational power of the CM system. CM Fortran thus combines :

- FORTRAN 77 supplemented with the VAX FORTRAN extensions;
- FORTRAN 90 array extensions and intrinsic functions for manipulating arrays;
- CM Fortran compiler directives for efficient memory utilizations and use of the virtual processing mechanism, whereby each physical processor simulates some number of virtual processors by subdividing its memory.

Summary of Features

The current version of CM Fortran (version 0.7) provides a number of language and implementation enhancements, as well as corrections to previously reported problems. Implementation of CM Fortran is not yet complete. However, by supplementing the features supported by the current version with other features supplied by the NPAC Research Consulting group, programming of real applications is made possible

on the CM-2 at NPAC. This section lists the language features supported by Thinking Machines Corporation (TMC). These are:

•Declarations

The choice of declaration syntax is entirely a matter of style and convenience; it does not affect program behavior. Programmers can choose between FORTRAN 77-style declaration or FORTRAN 90 subscript triplet style (e.g., **REAL, ARRAY(lower bound:upper bound:stride)::A,B,C)**). Note that declaration syntax does not determine the home of the array; that is, the machine on which an array is allocated, either the front end or the CM. The home of an array is determined at compile time separately for each program unit. It is the use of the array, within a program unit (in expressions, assignment, and intrinsic function calls) that determines the array's home. Care must be exercised to ensure that the actual and dummy array homes match. Version 0.7 offers an interface block construct, which checks for type and shape mismatches, and the **-list** option, which checks on the home of arrays at compile time. In brief, the work in an application program is divided as follows:

- the front end stores and processes all scalar data, including subscripted arrays;
- the CM stores and processes all arrays that are referenced as array objects. (The **FORALL** construct is the only statement that runs on the CM with subscripted arrays.)

•Array Constructors

Array constructors provide a means of denoting a rank-one expression as an arbitrary sequence of scalar values, as a regular sequence of scalar values, or as a combination of the two enclosed by

square brackets. The current CM Fortran version supports nested array constructors and multiple array constructors forming an array constructor.

- Array Expressions

Whole arrays and array sections can be used with the ordinary set of FORTRAN operators, provided that all arrays within one expression are conformable. Two arrays are conformable if they have the same shape. Operations on arrays are performed on an element-by-element basis.

A fundamental problem in parallel computing is defining algorithms that have "locality of reference." This can be done on the CM by using arrays or array sections of conformable parent arrays that are allocated in the same set of virtual processors. This generates in-processor operations that do not require more expensive interprocessor communication.

- Conditional Operations

Operations that are conditional on the element values or the subscript values can be done by using the **WHERE** and **FORALL** statements respectively.

- Subroutines

Under CM Fortran, any dummy array argument can be of explicit shape, adjustable (dynamic) or non-adjustable, assumed size (front end arrays only), assumed-shape (CM arrays only), or locally defined arrays (automatic). Program performance can be increased by using the **INTENT** statement, which prevents needless copying when passing an array as an argument to a subroutine. CM Fortran intrinsic functions, elemental functions, inquiry functions, and array-reduction functions are supported under CM Fortran version 0.7. All the intrinsic functions supported under CM

Fortran version 0.7 are listed in the *CM Fortran reference manual*.

The following example demonstrates the simplicity of programming in CM Fortran with the features described above. More examples are in /u1/cmsi/demo.

Assume a CM programmer wants to evaluate the integral

$$\pi = \int_0^1 \frac{4}{(1+x^2)} dx$$

using a simple rule; for example, the rectangular rule, which breaks the area under the curve into a certain number of slices where each slice is a rectangle. On the CM, each processor can be assigned a specific rectangle, and the area of all the rectangles can be evaluated simultaneously. Due to the fact that performance improves by using a higher virtual processor ratio, the programmer is going to evaluate the area of 128K slices on 8K actual processors. Each processor will simulate 16 virtual processors by subdividing its memory. The program is organized as follows:

```

1  program pi
2  parameter (intrvls=131072)
3  integer i(intrvls)
4  real sumall, width, x(intrvls),
+   slice(intrvls)
5  sumall = 0.0
6  width = 1.0/intrvls
7  i = [1:intrvls]
8  x = (i - 0.5)*width
9  slice = width*(4.0/(1.0+x*x))
10 sumall = sum(slice)
11 write(6,*) 'Sum = ',sumall
12 write(6,*) 'Error = ',sumall-
+   3.14159265358979323846
13 stop
14 end
```

Lines five and six show the work on scalars that will be done on the front end. Line seven is an array constructor (CM). Line eight is an array expression (CM). Line nine is the area of slice calculation, which is an in-processor operation (CM). Line 10 is a call to the intrinsic sum reduction function,

which is a CM operation that ends up with a scalar passed to the front end.

Features Supported by NPAC

The NPAC Research Consulting group, primarily through the effort of Roy Heimbach, offers the following CM Fortran features:

- An interface to the CM Scientific Subroutine Library (CMSL) Fast-Fourier-Transform (**FFT**) routine, which can be called from CM Fortran, is available in /u1/cmsi/lib/cmf/fft. The next release of CMSL will offer many new features, among them an **FFT** routine that can be called from CM Fortran. As soon as we have the new release, NPAC's interface should not be used.
- An interface to several of the basic parallel file system routines is also available in /u1/cmsi/lib/cmf/datavault. It includes a simple interface for transposing data in CM Fortran arrays from parallel format, suitable for parallel files on the datavault, to serial format, suitable for front end files, and vice versa.
- A procedure to use the intrinsic function **MATMUL** to multiply complex matrices is available in /u1/cmsi/lib/cmf/complex-matmul.
- An additional utility for matrix-sorting, which can be called from CM Fortran, is also provided in /u1/cmsi/lib/cmf/sort.

A README file is provided in each of these directories as a guide to using these features.

Optimizing CM Array Directives

To improve performance on the CM, an application program must maximize processor use and streamline or eliminate interprocessor communications. CM Fortran provides two compiler directives, **LAYOUT** and **ALIGN**, which can improve performance in the following ways:

- The **CMF\$ LAYOUT** directive allows the

user to set the orientation, processor ordering, and communications weight attributes of the virtual-processor set in which the array is allocated. In some cases, the **LAYOUT** directive will help eliminate interprocessor communications that would occur when using the default canonical array layout (one element per processor). When interprocessor communications are unavoidable, the **LAYOUT** directive can optimize the speed of communication along specified dimensions.

- The **CMF\$ ALIGN** directive causes specified sections of two possibly nonconformable arrays to be aligned in the same virtual processor set. This way, elemental operations on their corresponding elements do not involve interprocessor communications.

Conclusion

CM Fortran on NPAC's CM-2 provides a number of language and implementation enhancements that make data parallel processing of real applications possible. Future versions of CM Fortran are even more promising due to the fact that CM Fortran programs will be compiled into microcodes that will run faster than hand-written assembly (Paris) codes. 

—A. Gaber Mohamed
NPAC research consultant

References

- "Introduction to CM Fortran," *CMSI-1990 Notes*, NPAC, Syracuse University.
- Getting Started in CM Fortran*, Version 5.2-0.6, TMC (Feb. 1990).
- CM Fortran Reference Manual*, Version 5.2-0.6, TMC (Sept. 1989).
- CM Fortran Release Notes*, Version 5.2-0.7, TMC (March 1990).

"I've found CM Fortran to be very user-friendly....I won't have to spend a lot of time working on programming tricks."



George Mesina

Computational fluid dynamics

By Diane Purser

Many engineers confront problems in fluid dynamics, such as designing aerodynamic bodies for planes, trains and automobiles. The equations that govern fluid flow can be solved analytically for a few simple problems, but not for real-world problems. In computational fluid dynamics (CFD), approximate answers can be found with the aid of a computer.

George Mesina, a mathematician at the Idaho National Engineering Laboratory (INEL), wants to make designers' lives easier by developing new and more efficient CFD algorithms for parallel computers. Mesina and his colleagues in the INEL CFD effort, have created several interrelated finite difference CFD algorithms designed for vector and parallel processing.

These algorithms are based on approximating the governing differential equations by discrete analogues. The derivatives are approximated by finite differences defined with respect to a mesh or grid that overlays the flow region. In each cell of the mesh there is a finite difference equation for each

governing equation. Every time the entire set of finite difference equations is solved, the approximate solution is advanced in time by one time-step. By watching this solution change over many time-steps, one can see the evolution of the flow field.

"We're trying to find more efficient methods for solving these equations," Mesina said, "either by finding the answer more quickly, or by arriving at a more accurate answer in the same amount of computer time."

This is why he became interested in the greater speed and capacity of the Connection Machine, and decided to attend the CMSI. The massively parallel architecture of the CM is well-suited for explicit methods (those in which the solution of any one finite difference equation depends entirely on information from the previous time-step). Each cell of the grid is associated with a separate processor on the CM, so that an entire time-step requires only the amount of time needed to evaluate the finite difference equations in any one cell. The amount of time required to solve the entire equation set is the same, regardless of the number of equations in the set, unless it exceeds the number of processors in the computer. Semi-implicit and fully implicit codes can be similarly adapted to the CM using iterative methods.

Mesina plans to rewrite the INEL programs in CM Fortran and run them on NPAC's CM-2. "I've found CM Fortran to be very user-friendly," said Mesina. "I won't have to spend a lot of time working on programming tricks." 

Mesina received his PhD in applied mathematics from the University of Pittsburgh in 1988. His dissertation was titled Iterative solutions to Navier-Stokes difference equations.

"The Connection Machine is great for this application....I can assign one gene to each processor, and since all the calculations for each iteration are the same for all the genes, no communication is needed between the processors."

Molecular models

By Paul H. Hebner

"I wouldn't have known my model is so well-suited to the Connection Machine if I hadn't attended the CMSI," said Sophia Wang, a Syracuse University graduate student in chemistry. Working with George Levy, professor of chemistry, Wang is developing an unconventional molecular modeling method that mimics the natural process of evolution at the gene level.

The purpose of Wang's model is to explore the molecular structure of DNA and RNA using computational methods. The model assumes a specific, structured environment, and tracks the evolution of a limited population of genes over a period of time. "What we hope to see," explained Wang, "is the population evolving in a direction that closely matches our experimental data."

The experimental data Wang uses is taken directly from the results of tests performed on small pieces of synthetic DNA and RNA using nuclear magnetic resonance (NMR) technology. These tests measure the distance between the atoms within each molecule and the angle of the chemical bonds between those atoms. "The genes in my model are only symbolic representations

of a specific set of characteristics," said Wang, "and the range of characteristics, which are molecular angles, is determined by the experimental data."

Fitness to the environment is determined by the molecular angle and atomic distances represented in the gene. Those genes whose calculated molecular angles and distances have a greater similarity to the experimental values and are, therefore, more fit, tend to survive and dominate the population over time. Genes that are less fit tend to disappear from the population.

"The Connection Machine is great for this application," said Wang. "The model can be mapped like a grid. I can assign one gene to each processor, and since all the calculations for each iteration are the same for all the genes, no communication is needed between the processors." This allows her model to run very fast. "This is a perfect SIMD model," she added. 

Wang is a native of Taiwan, where she completed her undergraduate work. She earned a master's degree in physics from the University of Oregon in 1986, and since then, has been working with Levy at Syracuse University.

(manufacturers of the CM) spoke on the CM Scientific Subroutine Library (CMSSL). Other TMC speakers included physicist Pietro Rossi, who discussed scientific applications currently running on the CM, and James Frankel, project leader of the C* group, who presented a talk on the new C* language.

"The CMSI is a wonderful forum for participants to share their research interests and needs with Thinking Machines Corporation," said Sather who organized the workshop with NPAC support coordinator Betty LaPlante. TMC staff members sat in on


classes throughout the two-week session, and were available to answer researchers' questions.

The NPAC Research Consulting group taught classes and assisted researchers in developing parallel programs. On the last day of the workshop, participants presented short talks on their areas of interest.

The CMSI concluded with a luncheon in an atrium of the Science and Technology building where NPAC's new director, Geoffrey Fox, spoke on the future of parallel computing. [See accompanying article.] 

"CMSI" from P. 1

"I look forward to working with the NPAC research consulting staff in the future."



In a semantic network, objects and their relationships are represented on a directed graph. The objects themselves are the nodes of the graph, and the relationships between the objects are the arcs between the nodes. The larger the network is (the more objects and relationships), the more information it contains and

Semantic networks

By Paul H. Hebner

Jamshid Goshtasbi came to NPAC's CMSI to continue his research in artificial intelligence-oriented computer architectures using the Connection Machine Model CM-2.

His current research project involves using NPAC's CM-2 to simulate a special-purpose architecture developed at Oxford University for processing semantic networks. This architecture is called Wafer-Scale Integrated Semantic Network Architecture, and was originally developed by Jose G. Delgado-Frias, currently at SUNY Binghamton, and William R. Moore of Oxford.

"The problem in artificial intelligence applications is how to represent the knowledge base being used, and then how to manipulate that representation," explained Goshtasbi. Semantic networks are a very popular way of representing knowledge in AI applications. However, different computer architectures are suitable for different representations. "The Connection Machine was originally designed with semantic network applications in mind," said Goshtasbi.

the greater knowledge potential it has. With a suitable parallel architecture, the network manipulation can be done very fast.

There are two types of knowledge that can be derived from semantic networks. The first, explicit knowledge, comes from the objects on the graph and is limited by the number of objects on the graph. The second, implicit knowledge, is what can be inferred by the relationships between the objects, and depends upon the complexity of the graph and how it is set up.

Goshtasbi also plans to use NPAC's CM-2 directly in semantic network manipulation. "I look forward to working with the NPAC research consulting staff in the future," he said.

A doctoral student in advanced technology at SUNY Binghamton, Goshtasbi earned a BS in physics and an MS in nuclear engineering in his native Iran. Since coming to the United States in 1978, he earned an MS in electrical engineering from Penn State University, where he is now an assistant professor of general engineering at the Scranton campus.

"Centers like NPAC can play a big role in helping scientists share information and experience."

Finite element mechanics

By Paul H. Hebner

"Your thinking will be limited by the computer you have access to," said M. Fouad Ahmad, an assistant professor of engineering mechanics at the University of Wisconsin in Milwaukee.

Ahmad teaches finite element mechanics at both the graduate and undergraduate levels. He came to the CMSI because he wanted to gain experience on a highly parallel architecture and apply that experience to his teaching and research.

"Most codes for finite element analysis were designed in the mid-sixties for serial machines," explained Ahmad. His goal is to exploit the power of parallel computation in the classroom and in finite element research.

Ahmad started using the NCSA Finite Element Analysis library in class to teach finite element analysis and to do comparisons between different computers. He is now interested in performance issues.

Traditional textbook problems in finite element analysis can now be done on a PC. "I have introduced real-life problems as part of the course-work," said Ahmad. The students start out on classic problems using their PCs and then move to real-life applications where some problems can take 1.5 hours to complete. I then show them how the same problems take only 16 seconds to run on a supercomputer. I then ask the students to look for real-life problems to analyze for their final projects."

"What supercomputers allowed us to do was tackle larger problems," explained Ahmad. A structural analysis of brake shoes from an automobile is one example of the type of analysis actually done by his students using the Cray X-MP supercomputer at the NCSA. "Symmetrical plates with one circular hole do not exist in real-life," he



Fouad Ahmad is interested in the role of supercomputers in teaching mechanical engineering.

said, referring to one of the traditional finite analysis problems, "nor can you go to the hardware store and buy a frictionless pulley. We live in a nonlinear, 3-D world."

"Every new machine creates a new debate," said Ahmad. "Scientists need to think about what problems they will be solving in two to three years. Computation is cheap, and machines are getting much more powerful. Simplified methods [quantitative] can be very illusive. The question is: What can you gain by using the Cray or any other supercomputer architecture?"

"Centers like NPAC can play a big role in helping scientists share information and experience. They are very valuable in terms of training, conferences and providing access to high-performance computational environments that not every academic institution can afford." □

Ahmad earned his PhD in theoretical and applied mechanics from the University of New Hampshire, and worked as an industry consultant for several years before joining the University of Wisconsin. He is currently a visiting research scientist at the National Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign.

"With the Connection Machine, I can push the algorithm to the next level of performance."

Data compression

By Paul H. Hebner

Sandy Shaw is founder of Compact Data, Inc., a company formed in late 1988 to create and market a prototype of a new data-compression technique he calls fractal formed data compression, or FFDC. This technique uses fractal geometry to form data into compressible patterns.

The compression algorithm Shaw has developed is highly parallel, and the degree of compression is dependent upon the number of operations performed. "Our biggest limitations now are file size and computational power," said Shaw. This makes NPAC's Connection Machine a

storage, either locally or on a network.

"One of our goals is to put the decompression functions of FFDC on a single chip," explained Shaw.

A very high degree of data compression is possible using FFDC. While Shaw expects compression ratios of 50:1 to 100:1, the computing power of the massively parallel Connection Machine is needed to achieve that level of performance.

Unlike other high-performance compression methods, however, FFDC is completely noiseless. No data or details are lost as a result of the process, making it ideal for text files, executable files, medical images, and other data that cannot have errors.

Shaw's technique has the additional advantage of being able to take data compressed by conventional methods and compact it even further. "My technique starts where others leave off," he said. "With the Connection Machine, I can push the algorithm to the next level of performance. The next step is to take data that has been already compressed once and run it through the algorithm again for even greater compression."

Shaw was very pleased with his experience at the CMSI. "The lectures were very well structured," he said. "The NPAC consulting staff was sharp and very helpful." As a result of his work in the CMSI, and as part of his effort to develop a commercial prototype of FFDC, Shaw is now considering requesting research time on NPAC's CM-2 through the Corporate Partners Program. □

Originally from Texas, Shaw earned a BS in physics from the University of Houston in 1976 and an MS in physics from SUNY Stony Brook in 1978.



Sandy Shaw

natural choice for the further development of FFDC.

FFDC is an extremely asymmetric technique, requiring approximately ten million times more computation for compression than is needed for decompression. Although supercomputer performance is necessary for data compression, decompression can be done on a PC or workstation with a specialized add-on board. This makes it well-suited for read-only applications, such as CD-ROM and archival

"We came to the Summer Institute to see what we could learn about massively parallel computation....I believe that the Connection Machine has real potential in our areas of research."

Fluid mechanics

By Paul H. Hebner

"As problems become more complicated, the average computer time needed to solve them tends to become longer," explained Csaba Zoltani, a CMSI participant and researcher from the U.S. Army Ballistic Research Laboratory at the Aberdeen Proving Ground, Maryland. "We need to find ways to reduce run-times."

Zoltani attended the CMSI with his colleague, Monte Coleman, a senior computational scientist at the Laboratory. "We came to the Summer Institute to see what we could learn about massively parallel computation, and what problems are best suited to computers like the Connection Machine," said Zoltani. "I believe that the Connection Machine has real potential in our areas of research."

Zoltani's primary area of research is fluid mechanics. His current focus is the investigation of hypersonic flow fields through computational methods. These flows are extremely complex and involve the interaction of many discontinuities, shock waves, and may include chemical reactions as well. "I want to be able to

speed-up the computation for these problems," he said. Even limited problems of this type can take hundreds of hours of supercomputer time to solve.

The immediate goal is to take the codes used for these problems and run them on the Connection Machine. They will then compare its performance to that of the Cray X-MP supercomputer currently being used at that laboratory.

Zoltani is also interested in the economics of massively parallel computation. "Even a ten percent reduction in run-time could create a significant savings," he explained. "What we need to do is learn how to design efficient algorithms for the Connection Machine." □

Zoltani's background is in physics and mechanical engineering. He earned his undergraduate degree from MIT and did his graduate work in Switzerland. He considers himself both a theoretician and an experimentalist. "Computations are a complement to the experimental work we do here," Zoltani commented. "We verify our codes by comparisons with actual flow measurements."

techniques. She plans to modify this algorithm in the future and use larger neural network models on a more comprehensive data set.

In addition, she hopes to gain some speed-up in these algorithms by rewriting them to run on large, massively parallel machines. Because neural networks have a structure that is naturally parallel, she plans to implement them on a number of parallel computer architectures. This work will be done jointly with with Dr. Meledath Damodaran from the University of Bridgeport. They will then compare the performances of each architecture.

While attending the CMSI, with the help of several NPAC consultants, Parikh developed a program in *Lisp based on the backpropagation learning algorithm. She plans to refine this program and evaluate possible speed-ups. "I am impressed with how knowledgeable NPAC's consulting group is," she said, "and I look forward to working with them in the future." □

Parikh teaches courses in artificial intelligence and programming languages at Southern Connecticut State University in New Haven, CT. She received her PhD in computer science from the University of Maryland in 1977.

Computer Vision from P 12

"Computers don't get tired when performing tedious tasks, and they can be programmed to analyze numerical data objectively."

*Jo Ann Parikh
worked on
developing an
algorithm in
"Lisp."*



Computer vision and neural networks

By Diane Purser

Physicians study ultrasonic images to determine abnormalities in liver tissue. Geologists analyze remotely-sensed Landsat images to find fracture patterns on the earth's surface. The problem for computer scientists is to provide these specialists with interpretive/diagnostic tools that maintain accuracy over long periods of time with 100% consistency.

CMSI participant, Jo Ann Parikh, a professor of computer science at Southern Connecticut State University, wants to address this problem. She designs neural network models for the detection of geologi-

cal lineaments (features on the earth's surface, which may represent fractures), and for finding abnormalities in ultrasonic medical imagery.

"Computers don't get tired when performing tedious tasks, and they can be programmed to analyze numerical data objectively," said Parikh, now a visiting professor in the Mathematics Department at Cornell University. The implication is that computers can be used to improve the accuracy of a physician's diagnosis, and to aid the geological photointerpreter in locating geological structures of interest.

Parikh has implemented an algorithm based on Hopfield neural networks to detect geological lineaments. Written in vectorized FORTRAN, the image pattern recognition algorithm processes a 512 X 512 matrix on the IBM 3090-600J (at the Cornell National Supercomputing Facility) in about half an hour. In collaboration with John DaPonte from Southern Connecticut State University, she has developed a program that uses the backpropagation learning algorithm to detect abnormalities in ultrasonic images. The results* of this algorithm have been compared to those derived from conventional statistical

Continued P. 11



**Syracuse University
Northeast Parallel Architectures Center
111 College Place
Syracuse, NY 13244**

**Non-Profit Org.
U.S. Postage
Paid
Syracuse University
Syracuse, New York**

PARALLEL COMPUTING NEWS

NORTHEAST PARALLEL ARCHITECTURES CENTER AT SYRACUSE UNIVERSITY

VOLUME 3
September

NUMBER 9
1990

Solving the mystery of arrhythmias

By Diane Purser

"A good idea in one area of nature is almost always mimicked in another," said Don Michaels, a pharmacology professor at the SUNY Health Science Center. He and his colleagues have applied this principle to the functioning of the sino-atrial (SA) node, the natural pacemaker of the heart. They are developing computer models of the SA node in order to gain an understanding of arrhythmias, which are disorders in the regular beating of the heart.

The natural phenomenon to which Michaels referred was the tendency of certain fireflies in Malaysia and New Guinea to emit light in unison. They gather by the thousands on bushes and trees, lighting them up simultaneously, much like flashing Christmas tree lights. The question is, how do they coordinate their activity? An outside stimulus (like a flashlight) can affect the pacemakers in their brains, and change their individual rates. Thus, a group of this species of firefly acts as an outside stimulus to its neighbors, and the tail flashing of these neighbors adjusts to a new common frequency. Michaels and his colleague, Jose Jalife, thought that the pacemaker of the heart might work in a similar way, and applied this theory to a computerized model of the SA node.

The SA node is a group of interconnected cells located in the muscle tissue of the upper right chamber of the heart. These

cells supply the electrical impulses that keep the heart beating. Each of the cells is capable of firing spontaneously at its own frequency, and they all fire at slightly different frequencies. The overall activity of the heart's pacemaker is established by the coordination of the firing of the individual cells. The hypothesis Michaels and Jalife wanted to test is that a rhythmic heartbeat results from a dynamic process in which thousands of pacemaker cells communicate electrically. Eventually they reach a kind of "democratic consensus," and discharge rhythmically and harmonically to initiate each beat.

The heart can originate its own electri-



Don Michaels and systems analyst Ed Matyas model heart arrhythmias.

cal activity; however, the frequency of these cells is also normally affected by outside stimuli. For example, the brain often sends chemical signals (like the neurotransmitter acetylcholine, which usually slows the heart rate and rhythm) to the SA node. Other chemicals (such as adrenalin) can increase heart rate.

"We wanted to model an SA node to simulate the transmission of electrical activity between cells. We also wanted to

Continued P 2

INSIDE

THIS	ISSUE
CALENDAR	3
FROM THE TRAINING NOTEBOOK	4
NEW CM FRONT END	6
NEW RESEARCH AT NPAC	7
CM UPDATE	7
NEW STAFF	8

"We believe the capabilities of the CM-2 will bring us closer to understanding how arrhythmias happen."

"Arrhythmias" from P.1

simulate the response of the SA node to outside stimuli," Michaels said.

They used a mathematical model based on a set of differential equations. This model represented channels in the cell membrane through which different chemical electrolytes (such as sodium, potassium and calcium ions) can pass. Michaels and Jalife were then able to reconstruct electrical discharges using this mathematical description. Their model successfully reproduced the repetitive activity of a single cell, and eventually they were able to simulate the interaction of two cells.

Because the SA node is made up of numerous cells, however, Michaels and Jalife wanted to develop a model to simulate many cells interacting. Working with systems analyst Ed Matyas, they devised a two-dimensional simulation of large numbers of cells, and developed a color scheme to simulate real-world data. This scheme uses a matrix in which each cell is represented by one block of the matrix. Each color represents a relative firing time. The visual result is a matrix with blocks of color in which the "dominant" pacemaker region (that of the cells firing earliest) appears to conduct to other regions of the matrix represented by different blocks of color.

Having developed this sophisticated model, they could then explore more subtle complexities of the heart's physiology. "We're really interested in discovering what starts arrhythmias," said Michaels. "We simply don't know why some arrhythmias begin. Since they can lead to sudden cardiac death, we want to find an answer." Collaborating with Dante Chialvo, Michaels and Jalife began investigating chaotic activity in SA nodes affected by stimulation of the vagus nerve (the nerve which runs from the brain to the heart). When vagal stimulation occurs, acetylcholine is released in the SA node, causing a phase shift in which the cells either delay or speed up the formation of subsequent electrical impulses for a brief period. When these vagal stimuli are applied repetitively, the SA node can go into an arrhythmia and beat very erratically. Normally the cells of the SA node quickly regulate themselves again, and the heart returns to its regular rhythm. If this doesn't happen, a type of arrhythmia may occur which can be life-threatening.

To understand arrhythmias more completely, Michaels and his colleagues developed a mathematical model of the effects of acetylcholine pulses on the electrical activity of the SA node. They used the color coded matrix to simulate chaotic



PARALLEL COMPUTING NEWS

Parallel Computing News is published by the Northeast Parallel Architectures Center (NPAC) at Syracuse University. The Northeast Parallel Architectures Center promotes and explores advanced computing technology by providing parallel architectures and research support to university, corporate, and government researchers nationwide.

To subscribe, send name and mailing address to: *Parallel Computing News*

NPAC
111 College Place
Syracuse, NY 13244

Editor

Paul H. Hebner

e-mail address

editor@nova.npac.syr.edu

Phone

315-443-1722

pacemaker activity on the IBM 3090-600J at the Cornell National Supercomputer Facility. They were then able to plot the irregular dynamics of the cells' electrical impulses according to variations in the concentration of acetylcholine and frequency and duration of the pulses. This study gave them a useful framework in which to examine real-life arrhythmias.

"We really needed a larger matrix to conduct more realistic studies of arrhythmias," said Michaels. This is why he decided to implement another model of the SA node on the Connection Machine Model CM-2.

Working with biomathematician Alain Vinet, Michaels and Matyas developed a mathematical model for the CM-2. "The Connection Machine is a natural tool for this kind of model," said Michaels. "We can map each of the cells in the SA node to separate processors on the CM-2." Matyas

wrote the program for this model in C*. They were then able to run simulations two to four times faster than was possible using the IBM 3090.

In the future, Michaels and his colleagues plan to refine their model by developing larger matrices, and exploring the possibilities of getting better performance using CM Fortran. They also want to create visual images of the SA node using the CM-2's graphics display facilities. "It's a big jump from studying fireflies to modeling arrhythmias on a computer," said Michaels, "But we believe the capabilities of the CM-2 will bring us closer to understanding how arrhythmias happen." □

Michaels is an associate professor in the Department of Pharmacology at the SUNY Health Science Center at Syracuse. He received his PhD in pharmacology from the SUNY Health Science Center in 1974.

Calendar

October 16, 17	Alliant Training Workshop at NPAC
October 23	Mach Training Workshop at NPAC
November 28-30	Connection Machine Training Workshop at NPAC

For more information, please send electronic mail to training@nova.npac.syr.edu, or call Betty LaPlante at (315)443-1722.

NPAC discontinues CM-1 service.

NPAC announces the discontinuation of CM-1 service, effective October 1, 1990. Thinking Machines Corporation has made significant enhancements to this early implementation of their SIMD computing architecture, and has declared the era of the CM-1 to be ended. It is not without some nostalgia and sadness that the plug will be pulled. □

- Elizabeth C. Schermerhorn
NPAC Deputy Director

PCN
NPAC CALENDAR

Optimizing Programs with the Alliant FX/80 Fortran Compiler

The Alliant FX/80 Fortran compiler provides scientists and engineers the means of converting FORTRAN 77 computer codes, often called "dusty decks," into highly efficient concurrent and vector machine codes using automatic compiler options. The features of the Alliant compiler are:

- ease of selecting compiler options at compilation;
- programmer freedom from establishing interprocessor communications;
- ease of interpreting compiler messages for aiding optimizations.

The following improvements in execution speed can be observed when optimizing programs with the Alliant Fortran compiler [1]:

- for vectorization, two to four times;
- for concurrency, a factor approaching the number of processors (ACEs) used.

Program conversion and optimization

The following set of steps is a useful strategy for routine migration of Fortran codes to the Alliant.

1. Execute code with sample data before migration

It is important to establish correct program behavior and results before converting code to insure that subsequent versions are computing correctly. For the rest of these steps, **program.f** is the program that will be optimized.

2. Compiling with no optimization

sfx% fortran program.f

This is a very important step, and provides the programmer with the opportunity to check for any programming errors or language incompatibilities, since the original code may include language specific to other architectures.

3. Compiling with global optimizations

sfx% fortran -O program.f

Invoking the global optimization (**-O** compiler option) causes a code analysis to be done. Next, the compiler applies a standard set of optimizations that remove redundancy and inefficient references [1]. For example, it will move a constant assignment from within a DO loop. It then vectorizes eligible statements. Finally, in the same invocation, the program is made concurrent.

4. Compiling with global and **-autoinline** optimizations

sfx% fortran -O -autoinline program.f

The **-autoinline** option can select suitable subroutine calls and function references for in-line expansion. This option often makes other optimizations (such as vectorization and concurrency) possible once the expansion has been performed.

5. Compiling with global, **-autoinline** and **-DAS** optimizations

sfx% fortran -O -autoinline -DAS program.f

The **-DAS** option allows the compiler to generate code that makes valid associative arithmetic transformations. It also allows differently ordered transformations to be made, depending on the size of the complex of processors. Beware, however, that inconsistent results may occur if the number of physical processors varies from one runtime to another.

Interpretation of Compiler Messages

At each compilation step, the compiler produces a listing file if the **-l** compiler option is specified. The listing consists of informational messages that a programmer can use to guide the choice of automatic compiler options.

Generally, program listings contain optimization messages, loop summaries, expansion summaries, and explanations of optimization messages. Message in a listing identify each section of code that contains:

- each loop that was optimized;
- any loop or array that could not be optimized fully for concurrency or vectorization;
- any loop or array operation that cannot be optimized because associativity was not allowed;
- each subprogram and whether or not it could be expanded in-line.

The listing shows the source code, line numbers, and any messages.

A successful conversion and optimization of a scientific application on the Alliant will serve as an example. The program, written in FORTRAN 77, computes the classical deflection function of a particle scattered by a central Coulomb potential [2,3]. The main program declares a function, **FUNT**, external and this is called through another function, **CGI**, in the main program to compute the integral for the deflection using Gauss-Chebyshev Quadrature [4]. The program has two nested DO loops.

The program execution and system time (User Time) and the compiler options used during the optimization of this program are summarized in the Table. Details are given in the NPAC Alliant Training Workshop Notebook . Copies may be requested from NPAC, training@nova.npac.syr.edu.

For this example, the best automatically optimized version of the code runs about twice as fast as the unoptimized code, and

nearly three times faster than the dusty deck version run on the VAX 8800 computer.

Manual tuning techniques

In general, a programmer can select areas for further optimization by timing loops and using the profiler. In addition, alterations in the programming style may be considered for improving code efficiency, such as data manipulation, logic, loops (use of **cvd\$** directives), loop counts, loop lengths, conditional loop processing, loop branching, subprograms, and input/output statements.

As a first approach in manual tuning, the functions and/or subroutines can be expanded in-line manually where feasible. Subsequently the code can be tested and re-compiled, as described in steps 3 through 5.

The Table shows the improvement in performance (4* and 5*) following optimization, after the code was tuned manually.

Timing comparisons for example code optimization

Step #	Compiler Options	User Time (sec.)
1		18.7 (VAX 8800)
2		10.6
3	-O	5.7
4	-O -autoinline	5.7
5	-O -autoinline -DAS	5.7
4*	-O -autoinline	1.05
5*	-O -DAS -autoinline	0.95

The function, **FUNT**, is expanded in-line in the function **GCI** because **-autoinline** did not work in Step 5.

The process of adapting code to the Alliant is made easier by the power and

Continued P. 6

flexibility of the Fortran compiler. Most applications will show marked improvement in execution speed through use of standard compiler options. Additional performance may be gained through careful manual tuning. 

-Louis C. Vaz
NPAC Research Consultant

References

1. Alliant Computer Systems Corporation, *FX/FORTRAN Language Manual*,
2. W. Norenberg and H. A. Weidenmuller, "Introduction to the Theory of Heavy-Ion Collisions," Lecture Notes in *Physics* 51, Springer Verlag, NY, 1976.
3. H. Goldstein, *Classical Mechanics*, Second Edition, Edison Wesley Publishing Company, Reading 1980.
4. B. Carnahan, H. A. Luther and J. O. Wilkes, *Applied Numerical Methods*, John Wiley & Sons, Inc., NY 1969.

NPAC announces new SUN 4 front end for the CM

This October, NPAC will offer a SUN front end to the Connection Machine Model CM-2. The new computer will replace the Symbolics 3650 and will allow us to provide better service due to the rich, fully supported software and libraries available on the SUN. It will have CM Fortran, C*, *LISP, and CMSSL. This new front end will be accessible from the Internet.

The initial configuration will be a SUN 4 with one interface to the CM-2. In the near future, we will announce both the name of the new computer and the policies and procedures for the migration of the CM user community to this new platform. Regular CM users should watch the message of the day (MOTD) on cmx for announcements concerning this change. 

NPAC director's publications address parallel computing problems.

NPAC's director, Geoffrey C. Fox, is closely involved with several important publications that deal directly with problems and new developments in parallel computing. NPAC would like to make the readers of the *PCN* aware of their availability.

Solving Problems on Concurrent Processors, volumes I and II, combine natural and computer science to explain how major scientific algorithms can be implemented on large parallel machines. Both volumes are co-authored by Geoffrey C. Fox, and are available from the publisher, Prentice Hall, Englewood Cliffs, NJ, 07632.

Concurrency: Practice and Experience, edited by Geoffrey C. Fox, is a quarterly journal that serves as a forum for practical experience on concurrent machines. It is available from John Wiley & Sons, Ltd., Baffins Lane, Chichester, Sussex PO19 1UD, England.

International Journal of Modern Physics C - Physics and Computers is published quarterly by World Scientific Publishing Company and is available c/o Publications Expediting Services, 200 Meacham Avenue, Elmont, NY 11003. Geoffrey C. Fox is co-managing editor of this journal, which publishes review and research articles on the use of computers in the physical sciences. 

New research projects at NPAC

NPAC is pleased to announce this listing of new research as a regular feature of *Parallel Computing News*. We believe this will serve the research community by encouraging the exchange of ideas and experience.

The following new research projects have been awarded time at NPAC:

- ***Massively parallel accelerated artificial compressibility computations***, George L. Mesina, Idaho National Engineering Labs.
- ***Application of simulated annealing to a neural network model - the Boltzman Machine***, C. Mohan and Swarupa N. Sahoo, CIS, Syracuse University.
- ***Parallel computational fluid dynamics in 3-D***, James Glimm and Yuefan Deng, SUNY Stonybrook.
- ***Optimization of electrical energy generation within constraints using dynamic programming techniques on a shared-memory multi-processor***, Yehuda Wallach, CIS, Syracuse University, and David Koester, RADC/TOFM.
- ***Dynamic state estimation on multisensor network systems***, Michael A. Palis, University of Pennsylvania.
- ***Hydrodynamics, shocks and instabilities***, Joel L. Lebowitz, Rutgers University. 

Connection Machine update

NPAC's Connection Machine Model CM-2 now has the following new resources:

- **The Connection Machine Scientific Software Library Version 2.0 Beta II (CMSSL)**, a set of numerical routines designed to exploit the massive parallelism of the CM-2 that can be called from all supported languages, including CM Fortran, C*, *Lisp, Paris, Lisp/Paris, C/Paris, F77/Paris;
- **C* Version 6.0 Beta 1**, an initial release of a new version of the C* data-parallel programming language that supports data-parallel programming idioms, allows efficient access to all user-visible components of the CM system;
- **cmdbx**, a tool for the debugging and execution of CM Fortran programs at the source level, using VAX or Sun-4 front ends;
- **Multidrop CM I/O bus**, which allows either DataVault to be accessed from any of the four sequencers on NPAC's CM-2. 

Inquiries about any of these new CM resources can be directed to the NPAC consulting staff at consult@nova.npac.syr.edu.

PCN

NEW RESEARCH AT
NPAC

Kim Mills joins NPAC research consulting staff

NPAC welcomes Kim Mills, our newest research consultant. Mills holds a PhD in environmental science from the SUNY College of Environmental Science and Forestry (ESF), and was a technical consultant at the Cornell National Supercomputing

Facility (CNSF) prior to joining NPAC in July.

Mills became involved with computational science while in graduate school at ESF. There he used a computer model to describe and sample temporal patterns in climate for use in environmental analysis.

"My particular interest is in developing environmental applications for the Connection Machine," said Mills. "More and more disciplines are using supercomputers as standard tools," he explained. "What I would like to do at NPAC is work toward that as a standard for environmental scientists as well."

Mills was attracted to NPAC by more than the many opportunities for computational science. "I was very impressed by NPAC's presentations and the quality of the staff," he commented. "That's what really attracted me. Also, parallel computation is a fascinating way of doing science." □

- Paul H. Hebner



*Research consultant
Kim Mills looks forward to
developing environmental
applications.*



**Syracuse University
Northeast Parallel Architectures Center
111 College Place
Syracuse, NY 13244**

**Non-Profit Org.
U.S. Postage
Paid
Syracuse University
Syracuse, New York**

Appendix E - SCCS Bibliography and Abstracts

This is a bibliography of technical papers recently published by SCCS researchers and selected abstracts.

Syracuse Center for Computational Science Documentation Log

- SCCS 1: C3P-939 Case Conference Talk. Summer 1990
- SCCS 2: Fox, Geoffrey C., "Physical Computation," Invited talk at International Conference on Parallel Computing: Achievements, Problems and Prospects, Anacapri, Italy, June 3-9, 1990 C3P-928, CRPC-TR90090.
- SCCS 2b: Fox, Geoffrey C., "Physical Computation," Invited talk at International Conference on Parallel Computing: Achievements, Problems and Prospects, Anacapri, Italy, June 3-9, 1990 C3P-928, CRPC-TR90090 (Final version).
- SCCS 3: "NPAC and SCCS", G. Fox, Slide presentation to IBM July 26, 1990.
- SCCS 4: "Something to do with CRPC at Syracuse", G. Fox, Talk given at CRPC Annual Conference, Rice University, Houston, Texas, August 27, 1990.
- SCCS 5: "Static Performance Estimation in A Data Partitioning Tool", V. Balasundaram, Talk given at CRPC Annual Conference, Rice University, Houston, Texas, August 27, 1990.
- SCCS 6: "Migrating FORTRAN Codes To Parallel Computers", G. Fox, presentation to G. Weigand at Cornell University, September 27, 1990.
- SCCS 7: "Physical Computation and Software Systems: 1. New Approaches to Optimization based on physical analogs; 2. "Computing" is "just" an optimization problem", Invited Keynote Speech by G. Fox, at KBSA Conference, Sheraton Hotel, Liverpool, NY, September 26, 1990.
- SCCS 8: "Prospects for Parallel Computers", slide presentation given by G. Fox at Rome Air Development Center, Griffiss AFB, NY, September 21, 1990 and Cornell University, September 25, 1990.
- SCCS 9: Gandhi, Amar and Fox, Geoffrey, "Solving Problems in Navigation", September, 1990.
- SCCS 10: Levy, George and Fox, Geoffrey, "National Research Resource for Computational Methods in NMR", Proposal to the NIH

Biomedical Research Technology Program, 1991-1996, Submitted September 1990.

- SCCS 11: "Portable Parallel Programs", talk given by G. Fox to North American Transputer User Group Meeting, Cornell University, October 11, 1990.
- SCCS 12: Baillie, Clive F. and Coddington, Paul D., "Comparison of Cluster Algorithms for 2-d Potts Models", Caltech Concurrent Computation Program, California Institute of Technology, Pasadena, CA 91125 (C³P-945).
- SCCS 13: "Parallel Computing Architecture and Software", Invited talk by G. Fox at 63rd Semiannual Symposium at New York State Section of American Physical, IBM, Poughkeepsie, New York, October 19, 1990.
- SCCS 14: Balasundaram, Vasanth, Fox, Geoffrey, Kennedy, Ken, Kremer, Ulrich, "A Static Performance Estimator to Guide Data Partitioning Decisions", October 12, 1990.
- SCCS 14b: "Static Performance Estimation in A Data Partitioning Tool", Balasundaram, Vasanth; Fox, Geoffrey; Kennedy, Ken; and Kremer, Ulrich, *Third ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, April 21-24, 1991, Williamsburgh, VA.
- SCCS 15: Fox, Geoffrey, Choudhary, Alok, Ranka, Sanjay, Wu, Min-you, "Scalable FORTRAN: A Software Development Environment for SIMD/MIMD Distributed Memory Parallel Computers", BAA# 90-21 High Performance Computing (HPC) Software Science and Technology, Proposal Abstract, October, 1990.
- SCCS 16: Fox, Geoffrey, Biham, Ofer, Proposal to National Science Foundation REU NSF 90-79, "Undergraduate Research in Computer Science and Computational Physics", 1991-1994, November 1990.
- SCCS 17: "High Level Software Tools and the Astrophysical N Body Problem", Overheads of talk by G. Fox at ICASE Workshop on Unstructured Computation on Multiprocessors, Nags Head, North Carolina, October 29-31, 1990.
- SCCS 18: "Computer Science Software Tools and Languages to Support Computational Physics", G. Fox, O. Biham, Research proposal to

IBM Research Initiation Grant Program, October 30, 1990. (See #49)

- SCCS 19: Fox, Geoffrey, Goldberg, Marvin, "Development of Advanced Computer Methods for SSC Data Analysis", proposal submitted to Texas National Research Laboratory Commission, October 30, 1990.
- SCCS 20: Rose, Kenneth; Gurewitz, Eitan; Fox, Geoffrey C., "Statistical Mechanics and Phase Transitions in Clustering", Physical Review Letter, Volume 65, Number 8, Caltech Concurrent Computation Program, Pasadena, California 91125, May 2, 1990.
- SCCS 21: Rose, K.; Gurewitz, E.; Fox, G. C., "Constrained Clustering as an Optimization Method*", Caltech Concurrent Computation Program, California Institute of Technology, Pasadena, CA, September 17, 1990.
- SCCS 22: Keppenne, G. L.; Ghil, M.; Fox, G. C.; Flower, J. W.; Kowala, A.; Papaccio, P. N.; Rosati, J. J.; Shepanski, J. F.; Spadaro, F. G.; Dickey, J. O.; *Parallel Processing Applied to Climate Modeling*, November 1990
- SCCS 23: Fox, G. C., "Applications of Parallel Supercomputers: Scientific Results and Computer Science Lessons", 4, 47-90, *Natural and Artificial Parallel Computation*, ed. M. Arbib, J. A. Robinson, The MIT Press, Cambridge, Massachusetts, 1990
- SCCS 24: "Undergraduate Laboratory in Computational Physics and Parallel Computation", G. C. Fox, O. Biham, S. Ranka, A. Choudary, Proposal to The National Science Foundation, November 1990
- SCCS 25: Fox, G. C., Mansour, N., "An Evolutionary Approach to Load Balancing Parallel Computations," submitted to *Sixth Distributed Memory Computing Conference*, April 1991.
- SCCS 25b: Fox, G. C., Mansour, N., "An Evolutionary Approach to Load Balancing Parallel Computations," submitted to *Sixth Distributed Memory Computing Conference*, May 1991.
- SCCS 26: Fox, G. C. "The Northeast Distributed Center for Parallel Computing," an unsolicited proposal to the National Science Foundation by the Northeast Parallel Architectures Center at

Syracuse University in collaboration with the Cornell Theory Center, June 20, 1990.

- SCCS 27: Fox, G. C., Choudhary, A., Ranka, S., "Developing Large Industrial Applications on NCUBE-2", December 3, 1990.
- SCCS 28: "Compiling Fortran for the Intel Touchstone", Kennedy, Kenneth, Center for Research on Parallel Computation, Rice University, talk given December 14, 1990, joint CIS/SCCS lecture.
- SCCS 29: Fox, G. C., "Achievements and Prospects for Parallel Computing", Invited Talk at International Conference on Parallel Computing: Achievements, Problems and Prospects, Anacapri, Italy, June 3-9, 1990. C³P-927
- SCCS 30: Fox, G. C., "Undergraduate Computational Science Summer Internships", Proposal to the National Science Foundation REU Supplement to the Center for Research in Parallel Computation, January 1991.
- SCCS 31: Fox, Geoffrey C., "Parallelization of Scientific Codes for IBM3090", A proposal from NPAC at Syracuse University to IBM Corporation, December 12, 1990.
- SCCS 32: Rose, Kenneth, "Deterministic Annealing, Clustering, and Optimization", Thesis in partial fulfillment of the requirements for the Degree of Doctor of Philosophy, California Institute of Technology, Defended December 6, 1990, 1991. (C3P-950, CRPC-TR91114)
- SCCS 33: Mansour, N., Fox, G. C., "A Hybrid Genetic Algorithm for Task Allocation in Multicomputers," paper December 1990 (Old See #36)
- SCCS 34: Fox, G. C., Ranka, Sanjay, Choudhary, Alok, "Development of Software and Applications for Parallel Processors", A Joint Study with IBM Yorktown, December 1990
- SCCS 35: Fox, Geoffrey C., Choudhary, Alok, Ranka, Sanjay, Wu, Min-you, and Kennedy, Kenneth, "Scalable FORTRAN: A Software Development Environment for SIMD/MIMD Distributed Memory Parallel Computers", BAA# 90-21, High Performance Computing (HPC) Software Science and Technology, White Paper, October 26, 1990.

- SCCS 36: Mansour, Nashaat, "A Hybrid Genetic Algorithm for Task Allocation in Multicomputers", April 1991, will appear in *International Conference on Genetic Algorithms and Applications*, July 1991.
- SCCS 37: Fox, G. C., "Nearterm Examples of the Possibilities of Parallel Computing for the Financial Industry", narrative put together for Carolyn May, State Relations, Syracuse University, for presentation to members of the State Legislature, December 1990.
- SCCS 38: Chiu, T. W., "Guage theories on the random-block lattice*", Physics Dept. Caltech and Department of Physics, National Taiwan University, 22 April 1990. (C3P-955, CRPC-TR91113)
- SCCS 39: Fox, Geoffrey C., "Parallel Software Engineering", A Proposal to CASE, Syracuse University, January, 1991.
- SCCS 40: Min-You Wu, G. Fox. "Test Suite and Performance for Fortran90 Compilers," December 1990.
- SCCS 41: "An Outline of Fortran90 Compiler for Distributed Memory Systems," December 1990.
- SCCS 42: Fox, Geoffrey, Hiranadani, Seema, Kennedy, Ken, Koelbel, Charles, Kremer, Uli, Tseng, Chau-Wen, Wu, Min-you, "Fortran D Language Specifications", Rice COMP TR900141, December 1990.
- SCCS 42b: Fox, Geoffrey, Hiranadani, Seema, Kennedy, Ken, Koelbel, Charles, Kremer, Uli, Tseng, Chau-Wen, Wu, Min-you, "Fortran D Language Specifications", Rice COMP TR900141, December 1990, *Revised, February 1991*.
- SCCS 42c: Fox, Geoffrey, Hiranadani, Seema, Kennedy, Ken, Koelbel, Charles, Kremer, Uli, Tseng, Chau-Wen, Wu, Min-you, "Fortran D Language Specifications", Rice COMP TR90079, December 1990, *Revised, April 1991*.
- SCCS 43: Fox, Ghandi, "Physical Optimization for Navigation and Robot Manipulation" (From Terri)
- SCCS 44: Fox, Geoffrey C., Choudhary, Alok, Ranka, Sanjay, Wu, Min-you, Kennedy, Kenneth, "Scalable FORTRAN: A Software Development Environment for SIMD/MIMD Distributed Memory Parallel Computers", BAA# 90-21, High Performance

Computing (HPC) Software Science and Technology, Proposal,
December 1990. (REVISED)

- SCCS 45: "Coarse Grain Parallelization of Fortran-77 Programs", talk given by Josef Stein, Visiting Researcher, to SCCS/NPAC Meeting, January 22, 1991.
- SCCS 46: Mohamed, A. Gaber, Valentine, Daniel T., and Hessel, Richard E., "Numerical Study of Laminar Separation Over an Annular Backstep", paper submitted to *Computers & Fluids: An International Journal*, Pergamon Press, New York, January 9, 1990
- SCCS 47: "MOVIE", Wojtek Furmanski, Talk slides.
- SCCS 48: A. Moore, G. Fox, J. Alman and R. Goodman 1991 "A VLSI Neural Network for Color Constancy" In Touretzky, D. S., Lippman, R., (eds.) *Advances in Neural Information Processing Systems 3*. San Mateo, CA: Morgan Kaufmann.
- SCCS 49: Fox, Geoffrey C., "Computer Science Software Tools and Languages to Support Computational Physics", a research proposal to the IBM Research Initiation Grant Program, The Syracuse Center for Computational Science, October 30, 1990. (ACIS Proposal) (same as #18)
- SCCS 50: Fox, Geoffrey C., "ACTION Industrial Technology Transfer for NYS Industry for Parallel Computing", Talk at Thursday Morning Roundtable, January 31, 1991.
- SCCS 51: Fox, Geoffrey C., "Computer Hardware, Advanced Mathematics, and Model Physics (CHAMMP) Program", Proposal submitted to Department of Energy with TRW, January 30, 1991.
- SCCS 52: Salmon, John K., "Parallel Hierarchical N-Body Methods", Caltech Ph.D. Thesis, December 6, 1990. (CRPC-90115)
- SCCS 53: Fox, Geoffrey, C.; Furmanski, Wojtek; "JTF Quarterly Report and Project Summary", February 21, 1991
- SCCS 54: Fox, Geoffrey, C., "JTF Center of Excellence", talk given at NPAC as part of the JTF Quarterly Report, February 21, 1991.
- SCCS 55: Furmanski, Wojtek, "MOVIE: M Reference Manual - DRAFT", version dated February 21, 1991.

- SCCS 56: Fox, Geoffrey C., "Resource Material on Computational Science", information packet containing SCCS 23, 29, 67, 69, PCN Vol. 3 #5 used for presentation and general information at conference March 15, 1991, Chicago IL.
- SCCS 57: Fox, Geoffrey C.; Furmanski, Wojtek; "Syracuse University Joint Tactical Fusion Center of Excellence", Cost and Technical Proposal, RFP#L07-6744, June 21, 1990.
- SCCS 58: Stormon, Charles, "Coherent Research, Inc.", presentation of facility at JTF Quarterly Report, NPAC, Syracuse University, February 21, 1991.
- SCCS 59: Fox, Geoffrey C.; Furmanski, Wojtek, "JTF Quarterly Report Syracuse Center for Computational Science: Report on MOVIE Project for July through September 1990", September, 1990.
- SCCS 60: Fox, Geoffrey, C.; Furmanski, Wojtek, "JTF Quarterly Report Package", containing SCCS 53, SCCS 54, SCCS 55, SCCS 57, SCCS 58, SCCS 59; February 21, 1991.
- SCCS 61: "Presentation to the Joint Tactical Fusion Program", multiple contributors, JTFPMO Techbase Coord. Meeting at RSRE, Malvern, England, September 11-13, 1990.
- SCCS 63: "Computer Hardware, Advanced Mathematics, and Model Physics (CHAMMP) Program Research Grant", a proposal in response to notice 91-3, Joint Proposal with TRW, February 1991.
- SCCS 64: Fox, Geoffrey, "Parallel Computing: Software Applications and Their Linkage", presentation slides "B", February 28, 1991.
- SCCS 65: Apostolakis, John; Baillie, Clive F.; Fox, Geoffrey C.; "Investigation of the 2d O(3) Model Using the Overrelaxation Algorithm", *Physical Review D*, Volume 43, Number 8, pp. 2687-2693, April 15, 1991. C³P-943, COLO-HEP-231, CRPC-TR90078.
- SCCS 66: Zenios, Stavros A., "Adventures with a Million Processors: Financial Simulations and Large Scale Optimization on a Massively Parallel Connection Machine", presentation materials, NPAC at Syracuse University, March 7, 1991.
- SCCS 67: Fox, Geoffrey, C., "Computational Science Education", notes and excerpt from "Achievements and Prospects for Parallel Computing" Invited talk at International Conference on Parallel Computing: Achievements, Problems and Prospects Anacapri,

Italy June 3-9, 1990. SCCS-29; C³P-927; CRPC-TR90083, from meeting on January 23, 1991.

- SCCS 68: Leung, Alvin, "Experience of LU Factorization on Concurrent Computers: Results from SURF 89 Research on Parallel Pivoting from Summer 90; Findings of LU Factorization on CM", slides from NPAC/SCCS Wednesday Lunchtime Presentation Program, March 13, 1991.
- SCCS 69: Fox, Geoffrey, C., "Confidential Documentation of the Caltech Proposals in Computational Science", to be used as portion of information packet on Computational Science Program at Syracuse University for Exemplary Parallel Repository Advisory Committee Meeting, Chicago, Illinois, March 15, 1991.
- SCCS 70: O'Farrell, William G., and McCracken, Nancy Jean, "Association Graphs and Their Implementation on the Connection Machine", Technical Paper, March 18, 1991.
- SCCS 71: Fox, Geoffrey C., "Slide Presentation D", March 1991.
- SCCS 72: "White Paper for ONR Neural Supercomputer DoD University Research Initiative", Joint study with JPL, Yale University, and Syracuse University, Geoffrey Fox, PI, March 20, 1991.
- SCCS 73: "An Evaluation of High Performance and Parallel Computing Hardware and Software needs for DoD Applications", An extension of a taxonomy of problem and software architectures, March 20, 1991.
- SCCS 74: Mansour, N., Fox, G. "Parallel Genetic Algorithms with Application to Load Balancing for Parallel Computation", Submitted to *Supercomputing '91* in Albuquerque, NM, March 1991.
- SCCS 75: Saltz, Joel, "Thoughts of a Great Man", Slide Presentation to NPAC, March 21, 1991.
- SCCS 76: Fox, Geoffrey C., "NIH Site Visit Presentation", Overheads, March 21, 1991.
- SCCS 77: Ahmad, Ishfaq and Wu, Min-You "Express versus iPSC/2 Primitives: A Performance Comparison," SCCS Technical Report, April 1991.

- SCCS 78: Fox, Geoffrey C., "Parallel Problem Architectures and Their Implications for Parallel Software Systems", Presentation at DARPA Workshop, Providence, Rhode Island, Feb 28, 1991, White paper, April 1991.
- SCCS 78b: Fox, Geoffrey C., "The Architecture of Problems and Portable Parallel Software Systems," Submitted to *Supercomputing '91*, Albuquerque, NM, November 18-22, 1991
- SCCS 79: Fox, Geoffrey C. "Remarks on Fortran 90 and Distributed Memory LAPACK," BLACS Workshop, Houston, Texas, March 28, 1991.
- SCCS 80: Fox, Geoffrey C. "Physical Computation ASAS Program Office TECHBASE Group at Syracuse University," Presentation APO TECHBASE Meeting, Syracuse NY April 9-10, 1991.
- SCCS 81: Furmanski, Wojtek "MOVIE and Map Separates," View Foils, JTF Semiannual Coordination Meeting, Syracuse NY April 9-10, 1991
- SCCS 82: Furmanski, Wojtek "MOVIE and Map Separates," JTF Semiannual Coordination Meeting, Syracuse NY April 9-10, 1991
- SCCS 83: Fox, Geoffrey C. "Parallel Computing and Education," submitted to Daedalus, Journal of the American Academy of Arts and Sciences, April 1991. C3P-958, CRPC-TR91123.
- SCCS 84: Fox, Geoffrey C., Frieder, Gideon, "Computational Science Education for Undergraduates", a proposal submitted to the National Science Foundation, April 12, 1991.
- SCCS 85: Fox, Geoffrey C., "Preliminary Analysis Enhanced Clustered Fortran Language Specification", Dated December 7, 1990, appears March 29, 1991.
- SCCS 86: Fox, Geoffrey C., "Resource Documents for Computational Science Education," A compilation of materials, March 1991.
- SCCS 87: Fox, G. C., "Joint Tactical Fusion Concurrent Processing Program, Contract #958847", proposal to the Jet Propulsion Lab, Caltech, April 1991.
- SCCS 88: Wu, Min-You and Fox, Geoffrey C. "Compiling Fortran90 Programs for Distributed Memory MIMD Parallel Computers," April 1991. CRPC-TR91126

- SCCS 89: Copies of Slide Presentations from APO Techbase Semiannual Coordination Meeting, April 9-10, 1991
- SCCS 90: Furmanski, Wojtek, "MOVIE--a software system for computational science," Slides from SCCS presentation, April 17, 1991.
- SCCS 91: Fox, Geoffrey C. "FortranD as a Portable Software System for Parallel Computers," Presentation at Supercomputing USA/Pacific '91 Conference, Santa Clara, CA, June 1991. CRPC - TR91128
- SCCS 92: Fox, Geoffrey C. "Approaches to Physical Optimization," Submitted to SIAM J. Sci. Stat. Comp, April 1991. C3P-959, CRPC-TR91124
- SCCS 93: Li, Yong "TITLE," Report for CIS 90, Spring 1991, April 26, 1991.
- SCCS 94: Mohamed, A. Gabor, Fox, Geoffrey C. "Note: Blocked LU Factorization in Engineering Applications on a Minisupercomputer," submitted to *Microcomputers in Civil Engineering*, An International Journal from Elsevier Applied Science, ASCE 1991 Engineering Mechanics Conference, Mechanics Computing in the 1990's and Beyond. April 1991. CRPC-TR91131.
- SCCS 95: O'Farrell, William G. "Garbage Collection Algorithms for the Connection Machine," Ph.D. dissertation, Defended April 19, 1991.
- SCCS 96: Baillie, Clive "Monte Carlo Simulation of Dynamically Triangulated Random Surfaces as Models of String Theory," presentation to SCCS/NPAC, April 29, 1991.
- SCCS 97: Marinari, Enzo "Random Self-Interacting Chains, and Connections to Protein Folding", presentation of research to NPAC/SCCS members, May 1991.
- SCCS 98: Fox, Geoffrey C., "Evaluating Fortran D for NASA Distributed Supercomputing Applications, 1991 President's Fund", for JPL at California Institute of Technology, Syracuse Portion, May 2, 1991.
- SCCS 99: Fox, Geoffrey C. "Evaluation of Parallel Computing for Over the Horizon Radat (OTH)," Proposal to General Electric Company, May 5, 1991.

- SCCS 100: Fox, Geoffrey C. "Syracuse University School of Management Project", May 1991.
- SCCS 101: Apostolakis, John; Baillie, Clive F.; Fox, Geoffrey C.; "The 2-d O(3) model by overrelaxation", *Lattice 90*, Proceedings of the International Symposium, Tallahassee, Florida, October 8-12, 1990.
- SCCS 102: Coddington, P. D., Baillie, C. F. "Parallel Cluster Algorithms", Proceedings of the Conference "Lattice '90", Tallahassee, October 1990, to be published in Nucl. Phys. B. (Proc Suppl.)
- SCCS 103: Apostolakis, J., Coddington, P., Marinari, E. "A Multi-Grid Cluster Labeling Scheme", Technical article, publication pending, May 1991.
- SCCS 104: Kennett, R. G. "The MOVIE Software Environment - Description and Applications," CSE6882 Term Paper, May 1991.
- SCCS 105: "A Proposal to the GE Foundation for Support of The Faculty for the Future Program", submitted by Syracuse University, by invitation from Dr. Clifford V. Smith, Jr. to Melvin A. Eggers, April 1991.
- SCCS 106: Fox, Geoffrey C., "CRPC Application Review - February 11-14, 1991", Quarterly report of the progress of the CRPC projects, February 25, 1991. CRPC-TR91151.
- SCCS 107: "Parallel Computing for Signal Processing and Ocean Environmental Modeling," Geoffrey C. Fox, PI, Proposal to General Electric Company, May 15, 1991.
- SCCS 108: Apostolotkis, John "Thermal Rounding of the Charge Density Wave Depinning Transition" submitted to Physical Review Letters, May 1991.
- SCCS 109: Wu, Min-You, Shu, Wei, "Scattering Schedule for Problems with Unpredictable Structures," April 1981
- SCCS 110: Fox, Geoffrey C. and Balasundaram, Vasanth "Parallel Programming as an Optimization Problem," Invited Paper at Second Internatinal Conference on Expert Systems for Numerical Computing, Purdue University, April 25, 1990.

- SCCS 111: Fox, Geoffrey C. "Hardware and Software Architectures for Irregular Problem Architectues," Invited talk at ICASE Workshop on Unstructured Scientific Computation on Scalable Microprocessors, Nagshead, NC, October 29-31, 1990.
- SCCS 112: Edelsohn, D. J. and Fox, G. C. "Hierarchical Tree-Structures for Particle Simulations," hardcopy of materials for poster presented at SCCS Poster Session, May 23, 1991.
- SCCS 113: Fox, Geoffrey C. "Parallel Computing: Software, Applications and their linkage," Slide Talk A, Rev. April 16, 1991.
- SCCS 113b: Fox, Geoffrey C. "Parallel Computing: Software, Applications and their linkage," Slide Talk A, Rev. June 14, 1991.
- SCCS 114: Fox, Geoffrey C., "Future of Big Computing: A Triumph for Lilliputions," Slide of article from New York Times, November 25, 1990.
- SCCS 115: Fox, Geoffrey, C., "Parallel Computing: Some HPC software implications that we have learnt from the use of parallel computers," Slide Talk B, Rev. April 5, 1991.
- SCCS 115b: "Portable Software for Parallel Computers," Invited Talk by G. Fox at SC91 USA/Pacific Supercomputing Conference, Santa Clara, CA, June 20, 1991.
- SCCS 116: Fox, Geoffrey, C., "Opportunity for a New World-Class Industry; The Supercomputing Corridor," Slide Talk C, Rev. April 16, 1991.
- SCCS 117: Fox, Geoffrey, C., "The Supercomputing Corridor the Central New York Advantage," Slide Talk, Rev. May 28, 1991.
- SCCS 118: Fox, Geoffrey, C., "New Computer Services Company," Slide Talk, Rev. May 28, 1991.
- SCCS 119: Fox, Geoffrey, C., "Computational Science," Slide Talk, Rev. May 28, 1991.
- SCCS 120: Fox, Geoffrey, C., "Examples of Physical Computation," Slide Talk D, Rev. April 5, 1991.
- SCCS 120b: Fox, Geoffrey, C., "Examples of Physical Computation," Slide Talk D, Rev. June 26, 1991.

- SCCS 121: Shapiro, Richard, "UTRC Experience with the Connection Machine," Slide presentation United Technologies Corporation Invited Speaker, April 1991.
- SCCS 122: Mansour, N. and Fox, G. C., "Physical Optimication Methods for Allocating Data to Multicomputer Nodes," INTERNAL ONLY DOCUMENT, June, 1991.
- SCCS 123: "Problems and Software Engineering Solutions to C3I Applications," Rome Labs/NPAC Software Technology for High Performance Architectures Workshop, June 12-14, 1991, Slide Presentations from Presenters, June 15, 1991. (*Very Large Document*)
- SCCS 124: Harblin, T. "Preliminary Proposal (6/91) to the Nason Foundation", for support of The Alexander G. Nason Prize Postdoctoral Fellowship Program in Computational Science at Syracuse University, June 24, 1991.
- SCCS 125: Stein, J.; Salzmann, D. "A Percolation Theory Approach to the Lowering of Ionization Potential in Hot and Dense Plasmas," submitted to Phys. Rev. Letters, June 1991.
- SCCS 126: Stein, J.; Salzmann, D. "The Effect of Percolaton on the Ionization Potential in Hot and Dense Plasmas," Presented to the 8th APS Topical Conference on Atomic Processes in Plasmas, Portland, ME, August 1991.
- SCCS 127: "ACTION-NYS Evaluation of the Use of Express in Future Computer Systems," Fox, G. C., Pending Proposal with IBM, August 1991.
- SCCS 128: "Northeast Parallel Architectures Center (NPAC) Proposal to Lawrence Livermore Laboratory," Proposal 8/1/91-9/30/91, July 1991.
- SCCS 129: "Software Engineering for High Performance C3I Systems," Proposal to Rome Laboratory with Donald Leskiw 12/1/91-11/30/92, July 1991.
- SCCS 130: Mills, K.; Vinson, M. "Financial Modeling on the Connection Machine," CASE Center Summer Conference, Minnowbrook Lodge on Blue Mountain Lake, July 9-11, 1991.
- SCCS 131: Shelley, M. J. and Vinson, M. "Coherent Structures on a Boundary Layer in Rayleigh-Bénard Turbulence," June, 1991.

- SCCS 132: Stein, J., Paul, M., Fox, G. "On Outer-Loop Parallelization of Existing, Real-Life Fortran-77 Programs," slide presentation, July, 1991.
- SCCS 133: Fox, Geoffrey C., Furmanski, Wojtek, "ASAS Quarterly Report for April-June 1991", Contract #353-6216, July, 1991.

"The NPAC/SCCS Technical Report Series is a way of sharing the lessons we've learned in computational science and parallel computing with others."

Geoffrey Fox

The NPAC/SCCS Technical Report Series

This series offers a wide variety of scientific articles on all aspects of parallel computing and computational science. For your convenience, we have compiled a selected list of abstracts, and Geoffrey Fox has written an introduction to each. These reports illustrate the broad interdisciplinary work now underway in our Computational Science group.

Algorithms and Applications

The first two papers are continuations of physics simulations started at Caltech; John Apostolakis and Paul Coddington moved to Syracuse. Clive Baille is now at the University of Colorado at Boulder.

SCCS-12 addresses the difficulty of parallelization of clustering algorithms. We intend to focus on this in the future, here at Syracuse University.

SCCS-12

Baillie, C.F., and Coddington, P.D., "Comparison of Cluster Algorithms for 2-d Potts Models."

Abstract: We have measured the dynamical critical exponent z for the Swendsen-Wang and the Wolff cluster update algorithms, as well as a number of variants of these algorithms, for the $q = 2$ and $q = 3$ Potts models in two dimensions. We find that although the autocorrelation times differ considerably

between algorithms, the critical exponents are the same. The results for $q = 2$ favor a logarithmic increase in the autocorrelation time with lattice size, implying that z is zero.

SCCS-65

Apostolakis, J., Baille, C., Fox, G., "Investigation of the 2d 0(3) Model Using the Overrelaxation Algorithm," February 1991.

Abstract: We investigate the 2d 0(3) model with the standard action by Monte Carlo simulation at couplings b up to 2.05. We measure the energy density, mass gap and susceptibility of the model, and gather high statistics on lattices of size $L \leq 1024$ using the FPS T-series vector hypercube. Asymptotic scaling does not appear to set in for this action, even at $b = 2.05$, where the correlation length is 304. We observe a 20% difference between our estimate $m/L \frac{1}{ms} = 3.52(6)$ at this b and the recent exact analytical result. We use the overrelaxation algorithm interleaved with Metropolis

updates and show that decorrelation time scales with the correlation length and the number of overrelaxation steps per sweep. We determine its effective dynamical critical exponent to be

$z' = 1.079(10)$; thus critical slowing down is reduced significantly for this local algorithm that is vectorizable and parallelizable.

SCCS-22 is the first paper in a continuing collaboration involving Syracuse University, UCLA, Parasoft, TRW, and the Jet Propulsion Laboratory. This work is important to us in demonstrating the utility of FORTRAN 90 as a portable software environment.

SCCS-22

Keppenne, G. L., Ghil, M., Fox, G.C., Glower, J. W., Kowala, A., Papaccio, P. N., Rosati, J.J., Shepanski, J. F., Spadaro, F. G., Dickey, J.O., "Parallel Processing applied to Climate Modeling," November 1990.

PCN

Special Section

Abstract: A climate model of intermediate size and complexity is used to investigate the relative advantages and disadvantages of current and future computer architectures for Global Change studies. The model is based on the primitive-equation system used in atmospheric general circulation models (GCMs) and implements an efficient transform method to switch back and forth between the model's state spectral and grid-point representations.

The target architectures for the implementation of the model are a multiple instruction, multiple data (MIMD) NCUBE hypercube, a single instruction, multiple data (SIMD) Connection Machine Model CM-2 and a CRAY Y-MP vector-supercomputer operating in single-processor mode. Details of the implementation on each architecture are presented, along with a discussion of programming techniques and human considerations, as well as parallelization and vectorization issues. FORTRAN 90 has proven an efficient programming language for all three architectures involved, either as is or after direct translation.

SCCS-23 and 29 review parallel computing with an emphasis on applications.

SCCS-23

Fox, G. C., "Applications of Parallel Supercomputers: Scientific Results and Computer Science Lessons," 4, 47-90, *Natural and Artificial Computation*, ed. M. Arbib, J. A. Robinson, The MIT Press, Cambridge, Massachusetts, 1990.

Abstract: This paper describes a number of major computations that were carried out in 1988-1990 at Caltech on a variety of parallel machines: hypercubes, transputer arrays, the Connection Machine, and the AMT DAP. We compare their performance with that of several advanced-architecture serial machines, including the conventional CRAY and ETA-10 supercomputers, on the same problems. From our varied use of parallel machines we derived a number of lessons concerning hardware, software, and performance, including the idea of an appropriate match between problem category and type of parallel architecture. It is our hope that these lessons will both encourage scientific users to use parallel machines properly and help computer scientists to design better hardware and software. We also discuss the emergence of a new academic discipline—computational

science—motivated by the larger and more important role played in science by computation in general, but especially by parallel computation.

SCCS-29

Fox, G.C., "Achievements and Prospects for Parallel Computing," Invited talk at the International Conference on Parallel Computing: Achievements, Problems and Prospects, Anacapri, Italy, June 3-9, 1990.

Abstract: Parallel computing works for the majority of large-scale computations. The development of parallel hardware designs has been largely transferred to industry, while universities continue major research efforts into better software environments. We describe a classification of problems about how different software models are needed for portable user-friendly high performance implementations on parallel machines. The education of a new generation of computational scientists will be a major challenge to our universities.

Physical Computation

SCCS-2 and SCCS-92 review the use of physical analogies to provide new approaches to computation.

SCCS-2

Fox, G. C., "Physical Computation," Invited talk at the International Conference on Parallel Computing: Achievements, Problems, and Prospects, Anacapri, Italy, June 3-9, 1990.

Abstract: Physical Computation embraces a variety of physical analogies used to tackle nontraditional problems. We describe Monte Carlo and deterministic methods, including simulated annealing and neural networks. Applications include economic change in Eastern Europe, the traveling salesman problem, vehicle navigation, track finding, and parallel computer load balancing.

SCCS-32

Fox, G. C., "Approaches to Physical Optimization," Submitted to *SIAM J. Sci. Stat. Comp.*, April 1991.

Abstract: Physical optimization is the use of analogies from nature to solve optimization problems. This approach leads to approximate solutions with time complexities that are much lower than traditional exact methods. The algorithms parallelize straightforwardly and promise the solution of many practical problems on the large-scale parallel computers expected in the future. We describe and

contrast four related physical optimization methods; the well-known simulated annealing and neural network approaches, as well as the important but less well understood deterministic annealing approach.

SCCS-21 is the last of a series of papers written with Rose and Gurewitz. We had a fruitful collaboration at Caltech on the deterministic annealing approach to clustering and tracking. SCCS-21 generalizes to the TSP. Further details can be found in Rose's PhD thesis.

Ken Rose is now on the Electrical Engineering faculty at the University of California/Santa Barbara. E. Gurewitz has returned to Israel.

SCCS-21

Rose, K., Gurewitz, E., Fox, G.C., "Constrained Clustering as an Optimization Method," Caltech Concurrent Computation Program, California Institute of Technology, Pasadena, CA, September 17, 1990.

Abstract: Our deterministic annealing approach to clustering, derived on the basis of the principle of maximum entropy, is independent of the initial state, and produces natural hierarchical clustering solutions by going through a sequence of phase transitions. This approach is

modified here for a larger class of optimization problems by adding constraints to the free energy. The concept of constrained clustering is explained, and three examples are given where it is used as a means to introduce deterministic annealing. First, the previous clustering method is improved by adding cluster mass variables and a total mass constraint. Secondly, the Traveling Salesman Problem is reformulated as constrained clustering, yielding the Elastic Net approach to the problem. More insight is gained by identifying a second Language Multiplier, which is related to the tour length, and can also be used to control the annealing process. Finally, the "open path" constraint formulation is shown to relate to self-organization and dimensionality reduction in unsupervised learning. A similar annealing procedure is applicable in this case, as well.

SCCS-9 applies the deterministic approach to navigation, a very promising area, which we continue to investigate at Syracuse University. Traditional methods are computationally prohibitive for large problems, such as the simultaneous movement of many (≥ 20) robot links

or vehicles. We show that deterministic annealing provides an alternative approach, which scales well to large problems. In this preliminary study we present the general theory and a very simple example.

SCCS-9

Gandhi, A. and Fox, G. C., "Solving Problems in Navigation," September, 1990.

Abstract: The problem of trying to find the minimal-time path from point A to point B through a field of obstacles is both an interesting and a practical one. To solve the navigation problem, one can construct a graph representing all the possible ways of going from A to B. In this way, the navigation problem is mapped onto a multidimensional combinatorial optimization problem.

Andy Moore was my student at Caltech, and he continues his research in the Computation and Neural Systems program at Caltech.

SCCS-48

Moore, A., Fox, G., Alman, J., and Goodman, R., "A VLSI Neural Network for Color Constancy," in Touretzky, D. S., Lippman, R., eds. *Advances in Neural Information Processing Systems 3*, San Mateo, CA:

Morgan Kaufmann, 1991.

Abstract: A system for color correction has been designed, built, and tested successfully; the essential components are three custom chips built using subthreshold analog CMOS VLSI. The system, based on Land's Retinex theory of color constancy, produces colors similar in many respects to those produced by the visual system. Resistive grids implemented in analog VLSI perform the smoothing operation central to the algorithm at video rates. With the electronics system, the strengths and weaknesses of the algorithm are explored.

The preliminary results show the advantage of the parallel architectures for problems of this type and comparable or larger sizes. The results can be generalized to other disciplines, given the similarities of the model with others used in fluid dynamics and engineering.

Genetic Algorithms have been explored for several optimization problems at Syracuse University. We used some of the load balancing problems originally tackled by annealing in SCCS-25.

Fox, G.C., Mansour, N. "An Evolutionary Approach to Load Balancing Parallel Computations," *Proceedings of the Sixth Distributed Memory Conference, April, 1991.*

Abstract: We present a new approach to balancing the work load in a multicomputer when the problem is decomposed into subproblems mapped to the processors. It is based on a hybrid genetic algorithm. A number of design choices for genetic algorithms are combined in order to ameliorate the problem of premature convergence that is often encountered in the implementation of classical genetic algorithms. The algorithm is hybridized by including a hill-climbing procedure, which significantly improves the efficiency of the evolution. Moreover, it makes use of problem-specific information to evade some computational costs and to reinforce favorable aspects of the genetic search at some appropriate points. The experimental results show that the hybrid genetic algorithm can find solutions within 3% of the optimum in a reasonable time. They also suggest that this approach is not biased towards particular problem structures.

Software

In SCCS-78 and SCCS-91 we apply the concept of problem architectures to define the characteristics of software systems that will be scalable to future parallel systems.

SCCS-78

Fox, G. C., "Parallel Problem Architectures and Their Implications for Parallel Software Systems," Presentation at DARPA Workshop, Providence, Rhode Island, Feb. 28, 1991, White Paper, April 1991.

Abstract: We show how the structure or architecture of applications suggest the nature of parallel software systems that will run portably on a variety of parallel machines, both those available now, and those expected during the coming decade. The discussion is illustrated by lessons learned from real applications implemented on current MIMD and SIMD machines. These are mainly academic problems and the extrapolation to complex industrial and government applications is unproven, but we believe our methodology will still be applicable.

SCCS-91

Fox, G. C., "FORTRAN D as a Portable Software System for Parallel Computers," Presentation at Supercomputing USA/Pacific '91 Conference, Santa Clara, CA, June 1991.

Abstract: We discuss how extensions of FORTRAN in a distributed computing environment may be the basis of a portable software system for a heterogeneous computing network consisting of SIMD and MIMD parallel machines connected with conventional (super) computers.

In SCCS-42c, we propose extensions to either FORTRAN 77 or FORTRAN 90, which are the basis of a portable parallel FORTRAN (FORTRAN D) for distributed-memory machines. We expect a language like FORTRAN D to become an industry standard. This is a major project for us, involving Rice University and other collaborators. SCCS's component of this research focuses on FORTRAN 90, as described in SCCS-88.

SCCS-42c

Fox, G. C., Hiranadani, S., Kennedy, K., Koelbel, Charles, K., Uli, T., Chau-Wen, Wu, M., "FORTRAN D Language Specifications," *Rice Comp* TR90-141,

December 1990. (Revised February 1991.)

Abstract: This paper presents FORTRAN D, a version of FORTRAN 77 enhanced with data decomposition specifications. It is designed to support two fundamental stages of writing a data-parallel program: problem mapping using sophisticated array alignments, and machine mapping through a rich set of data distribution functions. We believe that FORTRAN D provides a simple machine-independent programming model for most numerical computations. We intend to evaluate its usefulness for both programmers and advanced compilers for a variety of parallel architectures.

SCCS-88

Wu, M., Fox, G., "Compiling FORTRAN 90 Programs for Distributed-memory MIMD Parallel Computers," April 1991.

Abstract: This paper describes the design and motivation for a FORTRAN 90 compiler, a source-to-source parallelizing compiler for distributed-memory systems. We discuss the methodology of parallelizing FORTRAN programs and the principle of compiler design. Then we describe compiler directives, data partitioning and sequentialization,

communication insertion, and implementation of intrinsic functions. Some basic optimization techniques are also presented. We use an introductory example of Gaussian elimination to explain the compiling techniques. Other sample programs in our test suite, such as FFT and the N-body problem, are briefly discussed with their performance.

Any parallel software environment needs many software tools. In SCCS-5b and SCCS-14, we describe a data partitioning tool. This will be combined with physical computation methods later in our research.

SCCS-5b

Balasundaram, V., Fox, G. C., Kennedy, K., Kremer, U., "Static Performance Estimation in a Data Partitioning Tool," Third ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), Williamsburgh, VA, April 21-24, 1991.

Abstract: The choice of the data domain partitioning scheme is an important factor in determining the available parallelism and hence, the performance of an application on a distributed memory multiprocessor. In this paper, we

present a performance estimator for statically evaluating the relative efficiency of different data partitioning schemes for any given program on any given distributed memory multiprocessor. Our method is not based on a theoretical machine model, but instead uses a set of kernel routines to "train" the estimator for each target machine. We also describe a prototype implementation of this technique and discuss an experimental evaluation of its accuracy.

SCCS-14

Balasundaram, V., Fox, G. C., Kennedy, K., Kremer, U., "A Static Performance Estimator to Guide Data Partitioning Decisions," October 12, 1990.

Abstract: An approach to distributed memory parallel programming that has recently become popular is one where the programmer explicitly specifies the data decomposition using language extensions, and a compiler generates all necessary communication. While this frees the programmer from the tedium of thinking about message-passing, no assistance is provided in determining the data decomposition scheme that gives the best performance on the

target machine. In this paper, we discuss performance estimation as part of an interactive software tool that provides assistance for this very task. We describe a new approach to performance estimation that is based on experiments rather than on a fixed machine model. Preliminary experiments indicate that the proposed techniques will work well for a large class of scientific programs.

How to Order NPAC/SCCS Technical Reports

Check off the report numbers on the business reply card enclosed in this issue, or send a request to:

Lisa Deyo,
Executive Secretary
111 College Place
Syracuse University
Syracuse, NY
13244-4100


Email: [scs-report](mailto:scs-report@nova.npac.syr.edu)

[@nova.npac.syr.edu](mailto:scs-report@nova.npac.syr.edu)



Appendix F - ACTION Brochure

This is a brochure detailing NPAC/SCCS's technology transfer program.



ACTION, a high performance computing outreach program, puts parallel computing to work in industry. **ACTION** combines the facilities of the Northeast Parallel Architectures Center, **NPAC**, and the services of the Syracuse Center for Computational Science, **SCCS**, targeting innovative commercial opportunities for parallel computing.

ACTION:

- ☐ evaluates the use of parallel computing technology on commercial applications
- ☐ identifies applications especially suitable for parallel computation
- ☐ assists in converting and developing existing and new applications to run on parallel computers
- ☐ provides parallel computing training and tools.

ACTION bridges the gap between industry and academia, offering:

- ☐ a variety of high performance parallel computers
- ☐ visualization capabilities
- ☐ tools to ease the transition to parallel computing
- ☐ assistance selecting appropriate parallel computer architectures and transferring application codes from sequential to parallel computers
- ☐ code migration, and optimization
- ☐ training workshops and institutes
- ☐ continuing education, apprenticeships, and co-op programs
- ☐ PhD-level consultants, user groups, publications
- ☐ service, including 24-hour computer operations coverage.

ACTION members enjoy these benefits:

- ☐ immediate access to research and development results
- ☐ royalty-free, nonexclusive license to technology
- ☐ a team approach to projects, incorporating the expertise of highly trained staff, faculty, graduate students, and researchers
- ☐ access to talented graduate and undergraduate students in a wide variety of scientific fields
- ☐ reduced tuition fees in computational science continuing education courses, and training workshops and institutes
- ☐ the opportunity to work in residence at Syracuse University and participate in all of the parallel computing center's activities.

**...we must apply
parallel computers to
industrial production,
thereby serving the
needs of our nation
and strengthening
the U.S. economy.**

Cover Photo

NPAC's Connection Machine Model CM-2 from Thinking Machines Corporation provides up to 5 gigaflops of supercomputing performance. The CM-2 offers 10 gigabytes of high-speed disk storage on the connected DataVaults and supports visualization on its graphics display system.


And Advances In Parallel Computing

By the year 2000, high performance parallel computers will have a sustained performance of a teraflop (a trillion operations per second). They will lead the growth in computation and simulation as a major force in universities, industry, and society. In academia, computation will evolve as a third approach to science, joining the traditional methodologies of theory and experiment. In industry, a foundation of modern high performance parallel computing will be key to economic competitiveness and leadership.

This growth in computer power will accelerate developments in manufacturing, simulation, and intelligent systems central to tomorrow's social and economic needs. Increasing research in languages, software tools, and programming environments will lead to major improvements in these and other areas of science, research, and industry. The challenge to people working in these areas will be how to use these systems. This is where the field of computational science will best be able to serve them.

Computational science integrates computer science with a wide variety of application areas that use computers. These include such diverse fields as biology, economics, electronics, materials science, and physics. To train people in this field, I have set up at Syracuse University a new interdisciplinary computational science and research program dedicated to teaching people from many varied disciplines to use computers, and to develop tools and methodologies to use them well. I have provided this academic enterprise with Syracuse University's parallel computing facilities, services, programs, and an expert staff. In an effort to foster a close partnership between industry and academia, these services, facilities, and programs are now available to industry through the ACTION program.

Applying parallel codes developed in academia to industrial applications is not a simple extrapolation, and represents the greatest challenge in computational science. Therefore, I intend to complement ACTION's industrial outreach with a major research effort into the appropriate methods for code migration and development of software.




Geoffrey Fox at Syracuse University's new Center for Science and Technology.

Dedicated research will strengthen the partnership between academia and industry, allowing us to put parallel computing into practical, widespread utilization in industry. Having proven that "parallel computing works" in academia, we must apply parallel computers to industrial production, thereby serving the needs of our nation and strengthening the U.S. economy. Giving corporations a competitive edge with high performance computing will create new jobs, save existing ones, and foster a robust computing service industry. This is the promise of 21st century technology, and the mission of ACTION.

Geoffrey C. Fox

Director, NPAC & SCCS
Professor, Computer Science & Physics

Transferring the Technology to You:



The Alliant FX/80's FORTRAN compiler automatically converts sequential FORTRAN to parallel and vectorized FORTRAN.

Computers

ACTION offers you the latest in parallel computing technology, providing computers, facilities, and services

NPAC operates a variety of parallel computers, which represent each of the important parallel architectures available today. All of them are connected to the Internet, which makes them available from almost anywhere in the United States. Each supports a distinct computing environment with a full suite of programming languages, including FORTRAN, Ada, C, and Lisp.


NPAC's Connection Machine Model CM-2 from Thinking Machines Corporation features 32,768 processors and 1024 floating-point accelerator chips, and is capable of a performance of more than 5 gigaflops. In addition, it has 10 gigabytes of high-speed storage, and two graphics display systems for visualization.

The Multimax 520 and 320 from Encore Computer Corporation are shared-memory, bus-based computers.

The Multimax 520 has 16 XPC processors, 128 megabytes of shared memory, and is rated at a speed of 170 MIPS. The Multimax 320 has 20 APC processors, is rated at 40 MIPS, has 64 megabytes of shared memory, and runs the Mach distributed operating system.

The FX/80 from Alliant Computer Systems Corporation is a minisupercomputer offering 8 advanced computational elements, 6 interactive processors, and 128 megabytes of shared memory. Its FORTRAN compiler automatically converts sequential FORTRAN to parallel and vectorized FORTRAN.


NPAC also offers a 32-node NCUBE 2 hypercube system from NCUBE Corporation, founded on the basis of work done by Geoffrey Fox and other researchers at Caltech. This multiple-instruction, multiple-data computer is a general-purpose, distributed-memory machine. It features a back-end parallel disk system, 4 megabytes of memory on each node, runs Express, a portable programming tool, and will soon run the ORACLE database information management system.



The center provides office space and workstations for visiting researchers.

Facilities

Training Facility



ACTION offers year-round training in the center's spacious new training facility.

Workstations

NPAC provides a training facility with dedicated space for up to 20 students, with 20 workstations connected directly to the center's parallel computers.

Additional Sun workstations are available for visitors. These workstations, connected directly to the center's local area network, support a wide variety of parallel computing software.

The Resources of NPAC and the Services of SCCS

SCCS provides a staff of consultants with PhDs in computing and other applied sciences who are parallel computing experts. They are available to assist users through individual or small-group consultation, or by telephone or electronic mail.

The SCCS software engineering team consists of a group of experts in the field of parallel computing who can assist with the migration or creation of large, commercial codes for parallel computers.


NPAC offers parallel computer cycles for testing and production use in industry.

Services

Consulting

Software Engineering

Computer Cycles



The center provides tape storage, and does routine backups, tape mounts, and file restores.

Collaborative Projects

SCCS and Cornell University have collaborated in developing the Supercomputing Corridor, a partnership among the two universities, New York State industry, the local Metropolitan Development Association, and New York State. The Corridor is expected to accelerate economic development and technology transfer through the incorporation of parallel computing. NPAC's parallel computing facilities and SCCS's highly trained software engineering staff combined with Cornell's broad user community provide a unique testbed for innovative parallel computing software projects. In addition, both SCCS and the Cornell Theory Center have substantial experience in developing, evaluating, and implementing parallel programming tools.

ACTION

SCCS also collaborates with the Center for Research in Parallel Computing (CRPC), an NSF-funded Science and Technology Center. As members of CRPC's Grand Challenge Consortium, an outreach program to integrate parallel computing into specific application areas, SCCS is associated with numerous industrial, government, and academic projects nationwide. The Consortium project involves SCCS and ACTION members in a parallel computing information exchange through annual workshops, parallel software development and engineering consulting, training, and visits to the sites of various participants.

ACTION

Individual SCCS researchers often collaborate with industry. For example, working with General Electric, SAIC, and Knowledge Systems Company, Syracuse University electrical engineering professor Dan Pease developed the PAAS Parallel Assessment and Analysis System. This software assesses the potential of parallelism in applications, and analyzes the parallel computer best suited to each application. It provides an interactive environment, allowing users to prototype different approaches to making their applications parallel and to simulate running them on different parallel computers.

ACTION

The work of Auburn University professor Cherri M. Pancake is another example. She collaborates with Syracuse University, IBM, the Cornell Theory Center, and the National Center for Supercomputing Applications. Together, they plan to create a universal user interface for remote access to supercomputing sites. Using this menu-driven interface, it will no longer be necessary to learn several different operating systems to program a variety of remote supercomputers.

ACTION

Geoffrey Fox and a team of Syracuse University computer science faculty collaborate with a company called Parasoft. Their goal is to make existing sequential software immediately useful on highly parallel computers by developing a FORTRAN preprocessor. This preprocessor will map the FORTRAN 90 standard to several different parallel computing architectures, reducing the time it takes to prepare sequential FORTRAN codes to run on parallel computers.

In Academia


Understanding how the brain works has always been one of science's greatest challenges. Dr. Robert Barlow hopes to discover more about the neural networks within the brain by studying the eyes of the horseshoe crab (*Limulus polyphemus*). Using a massively parallel computer, the Connection Machine Model CM-2, Barlow analyzes nerve signals the *Limulus* eye sends to the brain, and those the brain transmits back to the eye.

Barlow, a neuroscience professor at the Institute for Sensory Research at Syracuse University, discovered a connection between the crab's eye and its brain early in his career. He found the *Limulus* eye changes dramatically at night. Its photoreceptors become a million times more sensitive to light due to a clock in the brain, which signals the eyes to adjust their photosensitivity after sunset.

Understanding the connection between the crab's eye and its brain, he hopes, will bring scientists one step closer to understanding similar, but more complex connections in humans. Running simulations of the *Limulus* eye's response to light patterns on the CM-2 allows him to explore this link between physiology and behavior. The eye's nearly 1,000 photoreceptors are mapped to an equal number of processors. The CM-2's data storage device stores video simulations of images to which crabs respond in their natural environment. He then uses these simulations to test responses in the model eye. To ensure accurate results, he compares the CM-2 responses to those measured from live crabs, using sonar signals to transmit the electrical activity recorded from their eyes underwater.

Parallel computing, Barlow feels, is crucial to furthering his understanding of the neural basis of behavior.

"What once took us weeks to simulate on a serial computer now takes us a matter of seconds on the CM-2," said Barlow. "This brings us a lot closer to understanding the neural networks within the *Limulus* visual system, and someday, we'll be able to extend this knowledge to more complicated systems."



Syracuse University professor Dr. Robert Barlow feels that parallel computing is crucial to understanding the neural basis of behavior.

And Industry

Tracking the formation and motion of an ink drop no larger than a speck of dust presents many engineering design problems. Xerox's Dr. Meng Lean has found a solution—real-time simulations that save both time and money, and enhance the quality of the final product—in this case, printing devices.


"Our primary aim in doing real-time simulation is electronic prototyping," said Lean. "If we can accelerate product development, we minimize prototype cycle time, ensure quality, and lower cost."

Only parallel processors and supercomputers can provide the computing power needed for real-time simulation. This is why Xerox became a member of the Corporate Partners Program, and decided to provide support for several university researchers also interested in exploring high performance computing. "There are a lot of advantages to combining basic academic research with high performance computing and the real-life constraints of industry," said Lean.

For example, the design challenge of predicting the formation and motion of ink drops in printing devices has given rise to a large, interdisciplinary project. The goal is to develop a robust algorithm for simulating a three-dimensional free-surface drop ejection. Several Cornell university researchers have implemented experimental algorithms for this simulation on Syracuse University's Connection Machine Model CM-2 and on the IBM 3090 at the Cornell National Supercomputing Facility.

Lean combines strong simulation algorithms like these with supercomputing hardware and computer graphics. This is part of a concept he calls "synamation." The aim is to construct dynamic computing and rendering environments, which further accelerate the pace of product development. Engineers can now sit down at computers with real-time graphics capabilities and interactively design and test prototypes.

According to Lean, as an increasing number of researchers and business people become interested in collaboration, parallel processing and supercomputers will play a much larger role in the design process. "This is why we at Xerox want to establish strong ties between industry and academia," said Lean. "Combining these resources opens a wide gamut of possibilities for prototyping, saving design time, and ultimately, lowering costs for consumers while maintaining higher quality standards."



Putting Parallel Computing to Work

Healthcare


In an effort to understand arrhythmic heart disorders more completely, pharmacology professor Don Michaels models the heart's natural pacemaker on a parallel computer. He then simulates its reactions to changing body conditions. "Parallel computing has allowed us to conduct more realistic studies of arrhythmias," said Michaels.

Earth and Space Resources

Hoping to recover more oil from petroleum reservoirs, computational scientist Philip Emeagwali models them on a massively parallel computer, and simulates the efficient management of groups of oil wells. "Many engineering problems are very important socially," said Emeagwali. "And only a massively parallel computer has the computational speed required to accurately simulate petroleum reservoirs."

Manufacturing

Supported by the Westinghouse Foundation, Syracuse University engineer Can Isik simulated the motion of robots by mapping the dynamic model of the robot to a neural network on a parallel computer. "It's important to be able to predict and control the motion of robots, especially, if you want to put them to work in industry," said Isik. "Parallel computers are a natural for this kind of problem, and they save a lot of time."



Researchers are able to control the motion of industrial robots using parallel computers.

Cornell computer scientist Stephen Vavasis uses parallel computing for modeling fluid motion. Vavasis has developed new algorithms to solve the dense system of linear equations that arise in fluid problems. This work is a collaborative effort with Xerox; the goal of the project is to simulate drop ejection from a free surface. "Some of these computationally intensive problems can only be solved by harnessing the power of thousands of processors simultaneously," he said.

Wharton School professor Stavros Zenios simulates the effect of interest rates in contingent cash flows using massively parallel computers. In a project with the World Bank, he has simulated estimations of input/output tables for developing countries. He has also used parallel computing for running financial simulations of mortgage-backed securities. "The challenges of developing your applications on a massively parallel computer can be great," he said. "But then, so can the rewards."

Computer and Information Science


Finance

The visibility of highway signs is especially important in adverse weather, such as fog and rain. Supported in part by Niagara Mohawk Power Corporation, RPI professor Badrinath Roysam and graduate student Philip Getto are developing a parallel computation method to assess the visibility of these signs under different weather conditions. This will allow them to compare quickly different designs for signs and lighting arrangements.

As a student at Syracuse University, Sergio Rene Ramirez Chavez designed a more powerful memory chip for parallel processors. He then ran simulations of the chip on two different parallel computers. Now Ramirez collaborates with Syracuse University researchers in an effort to develop higher speed parallel computers to be used in database retrieval applications, such as library searches. He is currently employed at Bell Northern Research.

Transportation

Electronics



Parallel computers assist scientists in designing more visible highway signs.

Technology Transfer In **ACTION**

Corporate Partners Program

ACTION's Corporate Partners Program encourages researchers from commercial firms to experiment with parallel processing in a low-cost, low-risk environment. The basic program includes an introductory number of computing cycles, and participation in many of Syracuse University's parallel computing center's ongoing support programs.

These include:

- ☐ user groups, which provide a forum for information exchange
- ☐ the Visiting Researcher Program, an opportunity for corporate researchers to work in residence at the parallel computing center
- ☐ publications, including training manuals, user guides, and a monthly newsletter; also included are technical articles, finished compilations of work, and a yearly book of abstracts, which describes ongoing research at the center, and includes a bibliography of published papers
- ☐ an annual conference in which corporate members present reports on their work in parallel computing, participate in round-table discussions and seminars on important parallel computing issues, and get updates from colleagues.

Additional Programs

The Corporate Partners Program can be augmented with further "add-on" programs to meet the needs of individual companies. Corporations may select any number of services from each of these programs.

Computational Science Research and Development

In this program, a team of Syracuse University parallel computing specialists work closely with individual corporations assisting in:

- ☐ selecting the appropriate parallel computer for each application
- ☐ transferring application codes from sequential to parallel computers
 - ☐ writing new parallel codes for applications
 - ☐ developing parallel software for each application
- ☐ providing portable programming environments and software tools for use on a broad range of parallel computers
- ☐ developing performance modeling techniques to measure and optimize the performance of each selected parallel computer.

Education and Technology Transfer

ACTION offers a number of parallel computing education and technology transfer options for individuals and corporations. These include:

- ☐ apprenticeships, co-ops, and university courses to match the interests and skills of graduate and undergraduate students with the needs of corporate ACTION members
- ☐ continuing education for professionals in the work force, with a concentration on state-of-the-art computing techniques for dealing with real-life industrial applications
- ☐ training workshops and institutes, where people from corporations around the world have the opportunity to meet with the staff of major parallel computing vendors, and the center's consultants, staff, and researchers
 - ☐ advisory services for those who wish to create and manage new parallel computing centers
 - ☐ outreach to professional organizations, encompassing a broad range of scientific and engineering fields
- ☐ support for start-up companies, providing assistance with the design and/or installation of parallel computing environments, temporary file storage, and interim use of computing resources
- ☐ increased visibility for parallel computing vendors wishing to highlight new products by donating or loaning them to the parallel computing center; the center then learns to use these products, and supports and promotes them.

DISTRIBUTION LIST

addresses	number of copies
RL/C3CA ATTN: Raymond A. Liuzzi Griffiss AFB NY 13441-5700	10
Syracuse University Northeast Parallel Architecture Ctr ATTN: Dr. Geoffrey Fox Syracuse NY 13441-5700	5
RL/SUL Technical Library Griffiss AFB NY 13441-5700	1
Administrator Defense Technical Info Center DTIC-FDAC Cameron Station Building 5 Alexandria VA 22304-6145	2
Defense Advanced Research Projects Agency 1400 Wilson Blvd Arlington VA 22209-2308	2
RL/C3AB Griffiss AFB NY 13441-5700	1
HQ USAF/SCTT Washington DC 20330-5190	1
SAF/AQSC Pentagon Rm 4D 269 Wash DC 20330	1

Naval Warfare Assessment Center
GIDEP Operations Center/Code QA-50
ATTN: E Richards
Corona CA 91718-5000

1

HQ AFSC/XTH
Andrews AFB MD 20334-5000

1

HQ TAC/DRIY
ATTN: Maj. Divine
Langley AFB VA 23665-5575

1

HQ TAC/XP-JSG
Langley AFB VA 23665-5520

1

ASD/ENEMS
Wright-Patterson AFB OH 45433-6503

1

WRDC/AAAI-4
Wright-Patterson AFB OH 45433-6543

1

WRDC/AAAI-2
ATTN: Mr Franklin Hutson
WPAFB OH 45433-6543

1

AFIT/LDEE
Building 642, Area B
Wright-Patterson AFB OH 45433-6583

1

WRDC/MTEL
Wright-Patterson AFB OH 45433

1

AAMRL/HE
Wright-Patterson AFB OH 45433-6573

1

Air Force Human Resources Lab
Technical Documents Center
AFHRL/LRS-TDC
Wright-Patterson AFB OH 45433

1

AUL/LSE
Bldg 1405
Maxwell AFB AL 36112-5564

1

HQ ATC/TTOI
ATTN: Lt Col Killian
Randolph AFB TX 78150-5001

1

US Army Strategic Def
CSSD-IM-PA
PO Box 1500
Huntsville AL 35807-3801

1

Commanding Officer
Naval Avionics Center
Library D/765
Indianapolis IN 46219-2189

1

Commanding Officer
Naval Ocean Systems Center
Technical Library
Code 9642B
San Diego CA 92152-5000

1

Cmdr Naval Weapons Center Technical Library/C3431 China Lake CA 93555-6001	1
Superintendent Code 524 Naval Postgraduate School Monterey CA 93943-5000	1
Space & Naval Warfare Systems Comm Washington DC 20363-5100	1
CDR, U.S. Army Missile Command Redstone Scientific Info Center AMSMI-RD-CS-R/ILL Documents Redstone Arsenal AL 35898-5241	2
Advisory Group on Electron Devices Attn: Documents 2011 Crystal Drive, Suite 307 Arlington VA 22202	2
Los Alamos National Laboratory Report Library MS 5000 Los Alamos NM 87544	1
AEDC Library Tech Files/MS-100 Arnold AFB TN 37389	1
Commander/USAISC ATTN: ASOP-DO-TL Bldg 61801 Ft Huachuca AZ 85613-5000	1
1839 EIG/EIT Keesler AFB MS 39534-6348	1

AFWC/ESRI
San Antonio TX 78243-5000

1

SEI JPO
ATTN: Major Charles J. Ryan
Carnegie Mellon University
Pittsburgh PA 15213-3890

1

Director NSA/CSS
W157
9800 Savage Road
Fort Meade MD 21055-6000

1

NSA
ATTN: D. Alley
Div X911
9800 Savage Road
Ft Meade MD 20755-6000

1

Director
NSA/CSS R12
ATTN: Mr. Dennis Heinbuch
9800 Savage Road
Fort George G. Meade MD 20755-6000

1

DoD
R31
9800 Savage Road
Ft. Meade MD 20755-6000

1

DIRNSA
R509
9800 Savage Road
Ft Meade MD 20775

1

Director
NSA/CSS
R08/R & E BLDG
Fort George G. Meade MD 20755-6000

1

DOD Computer Center
C/TIC
9800 Savage Road
Fort George G. Meade MD 20755-6000

1

ESD/AV
HANSCOM AFB MA 01731-5000

1

DCMAO/GWE
ATTN: JOHN CHENG
US COURTHOUSE/SUITE B-34
401 N MARKET
WICHITA KS 67202-2095

1

ESD/IC
HANSCOM AFB MA 01731-5000

1

FL 2807/RESEARCH LIBRARY
OL AA/SULL
HANSCOM AFB MA 01731-5000

1

Technical Reports Center
Mail Drop D130
Burlington Road
Bedford MA 01731

1

Defense Technology Sec Admin (DTSA)
ATTN: STTD/Patrick Sullivan
400 Army Navy Drive
Suite 300
Arlington VA 22202

1

DARPA/SYSTO
ATTN: Dr Barry Boehm
3701 North Fairfax Drive
Arlington VA 22203-1714

1

DARPA/SYSTO
ATTN: Dr William Scherlis
3701 North Fairfax Drive
Arlington VA 22203-1714

1

DARPA/SYSTO
ATTN: Stephen Cross, Lt Col, USAF
3701 North Fairfax Drive
Arlington VA 22203-1714

1

DARPA/CSTO
ATTN: Mr Brian Boesh
3701 North Fairfax Drive
Arlington VA 22203-1714

1

DARPA/CSTO
ATTN: Dr G. Weigand
3701 North Fairfax Drive
Arlington VA 22203-1714

1

RL/IRRA
ATTN: Mr M. Hinman
Griffiss AFB NY 13441-5700

1

DARPA/SYSTO
ATTN: Gio Wiederhold
3701 North Fairfax Drive
Arlington VA 22203-1714

1

**MISSION
OF
ROME LABORATORY**

Rome Laboratory plans and executes an interdisciplinary program in research, development, test, and technology transition in support of Air Force Command, Control, Communications and Intelligence (C³I) activities for all Air Force platforms. It also executes selected acquisition programs in several areas of expertise. Technical and engineering support within areas of competence is provided to ESD Program Offices (POs) and other ESD elements to perform effective acquisition of C³I systems. In addition, Rome Laboratory's technology supports other AFSC Product Divisions, the Air Force user community, and other DOD and non-DOD agencies. Rome Laboratory maintains technical competence and research programs in areas including, but not limited to, communications, command and control, battle management, intelligence information processing, computational sciences and software producibility, wide area surveillance/sensors, signal processing, solid state sciences, photonics, electromagnetic technology, superconductivity, and electronic reliability/maintainability and testability.