Center for Research on Parallel Computation Application Activities

Geoffrey C. Fox

CRPC-TR92213
March, 1992

Center for Research on Parallel Computation
Rice University
P.O. Box 1892
Houston, TX 77251-1892

Center for Research on Parallel Computation
Application Activities

March 13, 1992

The CRPC application work plays an important and natural role in unifying CRPC research. Application scientists interact with software and algorithm researchers both within the CRPC and in the external academic and industrial community. The application projects, coordinated by Geoffrey Fox, serves three major functions within the CRPC: (I) as a testbed for the development of core parallel algorithms such as those described in the preceding sections on optimization, differential equations and linear algebra, (II) as a testbed for the software and other basic computer science research in the CRPC, and (III) as an outreach demonstrating the relevance of parallel computing in particular application areas.

The application activities were described in detail in CRPC-TR91151 which is currently updated to September 1991. This included seven Rice, seven Caltech, six Los Alamos and four Syracuse applications listed in Table 1. However, we have recently decided to both increase the amount and change the emphasis of applications within CRPC. The change of emphasis involves more involvement with application scientists outside CRPC. These will be the ultimate consumers of CRPC software and algorithms; further, we want to impact the very best computations which will be the first grand challenges running on teraflop machines. We must go outside CRPC to find these application grand challenges.

The application activities are not consisted independently but rather grouped into larger grand challenge consortia involving different mixes of CRPC and outside participation. Our model for a (new) grand challenge consortium would involve

1. An annual workshop bringing together appropriate computer science and application scientists. These workshops would be used to isolate key algorithm and software issues and monitor progress.

2. Parallelization of selected kernels and complete applications.

3. An exchange of visits, in which application scientists would visit CRPC sites and CRPC scientists would visit Consortium member sites.

4. Training (short courses) and access to parallel computing facilities.

5. Assistance by the CRPC in the application of its software and algorithm technologies to the parallelization of specific applications.

Each project will leverage existing funding from both the CRPC and the chosen project areas. However, we will need additional funding from industry or government to cover the essential interdisciplinary research and development, as well as project infrastructure such as workshops, training, travel, and other management expenses. The project research would typically fund two senior researchers (usually one in the application area and one in parallel software) and three graduate students in applications, algorithms, and software.

	Activity	Lead Scientist	Institution	Role	Comments
R1	Flow in Porous Media	Wheeler	Rice	I (PDE) III (Geoscience)	4 codes, work with UT-Austin
R2	Seismic Inversion	Symes	Rice	III (Geoscience)	2 codes
R3	Scheduling	Bixby	Rice	I (Optimization)	Airline Problem Linear Programming
R4	Pilot Plant Design	Dennis	Rice	I (Optimization) III (Geoscience)	Nonlinear Programming
R5	Optimal Well Placement	Dennis Lewis	Rice	I (Optimization) III (Geoscience)	Continuous Optimization
R6	Molecular Dynamics	Phillips	Rice	III (Biophysics)	Partial Parallelization of CHARMM
R7	Storm Modeling	Mellor-Crummey	Rice	II (FORTRAN)	Base code of Oklahoma NSF Center CAPS
C1	Spectral–CFD	Fischer	Caltech	III (CFD)	MIT Collaboration
C2	Vortex Flow–CFD	Harrar Taylor	Caltech	II (PCN)	800 Lines PCN 4000 Lines FTN
C3	Astrophysical Particle Dynamics	Salmon Edelsohn	Caltech Syracuse	II (FORTRAN) III (Physics)	Irregular
C4	Vortex Method–CFD	Leonard	Caltech	III (CFD)	Irregular
C5	Cellular Automata	Sturtevant	Caltech	III (CFD)	Only Integer Calculations
C6	Plasma Physics	Liewer Brackbill	Caltech Los Alamos	III (CFD)	Particle in Cell, 4 codes
C7	PCN Benchmark Suite	Chandy	Caltech	II (PCN)	3 codes
L1	QCD	Gupta Fox	Los Alamos Syracuse	III (Physics)	First Large Scale CM-2 Calculations
L2	Molecular Dynamics	Holian	Los Alamos	III (Physics)	Local Forces for Metals
L3	Ocean	Smith	Los Alamos	III (Climate)	NCAR, Naval PGS Collaboration
L4	Multifluid	Smith	Los Alamos	III (Physics)	24,000 Lines FTN 8,000 lines CMFTN
L5	Adaptive Mesh	Saltzman	Los Alamos	I (PDE)	SIMD
L6	Multigrid	Dendy	Los Alamos	I (PDE)	SIMD
S1	Climate	Fox Mills	Syracuse	II (FORTRAN) III (Climate)	Collaboration with TRW & UCLA SIMD+MIMD
S2	Statistical Physics	Coddington	Syracuse	II (FORTRAN) III (Physics)	"Irregular" Clustering Algorithm SIMD+MIMD
S3	FORTRAN D Benchmark	Wu Stein	Syracuse	II (FORTRAN)	8 codes
S4	Linear Programming	Wu	Syracuse	I (Optimization) II (FORTRAN)	Simplex

Table 1: CRPC Application Project Chart — September 1991

Name	CRPC Institutions	Status
Computational Fluids	Caltech	Ongoing (in-house CRPC) See Appendix and Algorithm section
Optimization	Rice	Ongoing (in-house CRPC) See Optimization section
Computational Biology	Rice	Keck funded Rice-Baylor Collaboration See Appendix
Molecular Modeling (CHARMM)	Rice-Syracuse	Proposed Collaboration with Karplus and Co-workers
Geoscience	Rice	Funded by CRPC and State of Texas See Appendix
Industry	Syracuse	Funded by New York State See Appendix
Statistical Physics and Optimization	Syracuse	Funded by CRPC and IBM See Appendix

Table 2: Grand Challenge Consortium

Currently, we divide our activities into seven such consortia listed in Table 2, and with more details in the appendix. We note that some individual applications in Table 1, such as optimization and benchmark sets for PCN and Fortran D, are now described in other sections of this report.

We believe that CRPC application activities will grow by increasing our collaborations with outside scientists. For this reason, we intend to start a major educational activity in collaboration with the four NSF supercomputer centers. We view this as an activity of the proposed NSF metacenter which was originally discussed as a nationwide link of the centers at Cornell, Illinois, Pittsburgh, and San Diego. However, CRPC provides an intellectual breadth to this concept. Initially our educational outreach will be a two week summer institute aimed at the NSF Supercomputer Center staff and a carefully selected set of application scientists who can be expected to be early "power users" of CRPC and other high performance parallel machines. Our goal is both to accelerate integration of CRPC parallel computing technology and conversely feedback to CRPC, lessons from the applications. As part of this plan, we will increase the fraction of CRPC facilities available to the general community to 75%.

Appendix: Grand Challenge Consortia

1. Computational Fluid Dynamics (Caltech)

C11: Software for the Numerical Analysis of Bifurcation Problems

Scientific Goal

Development of algorithms and general software for the numerical analysis and control of bifurcation problems in ordinary differential equations and in certain classes of partial differential equations. The aim is to develop a successor to the widely used AUTO86 software package which was written at Caltech. The new version is intended to be implemented on a number of parallel machines.

Parallelization Issues

For robustness, the linearized equations arising in Newton's method will be solved by direct methods (nested decomposition). This considerably complicates parallel implementation on distributed memory machines.

CRPC Facilities

Use of the Gamma and Delta Touchstone at Caltech

CRPC Collaborators

Eusebius J. Doedel, Herbert B. Keller

CRPC Funding

CRPC - 68%, DoE - 32%

Outreach

Users of the AUTO86 software (in many research areas; at universities and in industry.)

C12: Parallel Computation of Wavy-Vortex Flows

Scientific Goal

Simulate numerically the transition from Taylor-vortex to wavy-vortex flow in one configuration of the Taylor experiment. We seek to study phenomena such as period-doubling and would like to show that wavy-vortices develop from Hopf bifurcation.

Parallelization Issues

Develop scalable, highly portable implementation of large application program using PCN to compose routines written in Fortran and C.

CRPC Facilities

Use a large variety of parallel computers including the Intel Delta.

CRPC Collaborators

Dr. David Harrar II and Prof. Herbert B. Keller, Applied Mathematics, Caltech

Prof. Stephen Taylor, Computer Science, Caltech

CPRC Support
Postdoc.

Outreach/Technology Transfer outside CRPC
None

Publications

1. D. L. Harrar II, H. B. Keller, D. Lin, and S. Taylor, "Parallel Computation of Taylor-Vortex Flows," in *Proceedings of Parallel Computational Fluid Dynamics*, Stuttgart, Germany, 10–12 June 1991, (ed. K. G. Reinsch et al.) Elsevier.

C13: Algorithms for Numerical Linear Algebra

Scientific Goal
Develop, analyze and implement algorithms for numerical linear algebra, specifically those based on continuation and Krylov subspace methods. Analyze connections between discrete iterative method of numerical linear algebra and continuous dynamical systems.

Parallelization
Determine how the above methods can best be implemented on parallel machines.

CRPC Collaboration

C. T. Lenard
Some work was done with Gautam Shroff (now at IIT New Delhi) while he was with CRPC.

CRPC Support
Full support

Outreach/Technology Transfer
None as yet.

C14: Dynamics and Control of Bluff Body Flows

Scientific Goal
Study of unsteady separated flows using fast, viscous vortex methods.

Parallelization
Implementation of fast N-body techniques with additional complications of imbedded solid surfaces and close interaction viscous effects.

CRPC Facilities
Previous use of Mark III at Caltech. Will use Intel Delta in the future.

CRPC Collaborators

A. Leonard, Petros Koumoutsakos (student), Greg Winckelmans (postdoc). Also follow work of others with similar techniques, e. g., John Salmon at Caltech.

CRPC Support

1/2 graduate student

Outreach/Technology Transfer

No technology transfer yet.

C15: Fundamental Studies of Turbulence

Scientific Goal

Study of transport, mixing, and vortex structures in homogeneous turbulence using Fourier spectral methods on the Navier-Stokes equations.

Parallelization

Efficient transpose required to perform large-scale 5123 three-dimensional FFT's.

CRPC Facilities

INTEL Delta

CRPC Collaborators

A. Leonard, Dale Pullin and Mei-Jiau Huang (student) at Caltech

CRPC Support

None

Outreach/Technology Transfer

Work is being done in collaboration with colleagues at NASA Ames. Efficiencies developed for the homogeneous turbulence code will be applicable to other turbulence spectral codes at Ames.

C16: Parallel Particle Methods

Scientific Goal

A better understanding of systems and devices with large numbers of particles with long range interactions. This includes plasmas, gravitational systems and others. Many continuous PDE's can be looked at in this way as well.

Parallelization

Load balancing and minimizing communication are, as with most parallel problems, the main focus. The complications here are complex geometries that complicate the load balancing and long range interaction that complicate the communication.

CRPC Facilities

I've been using the Gamma at Caltech and will be migrating to the Delta. I will also be using the CM-5 at Los Alamos.

CRPC Collaborators

Steve Roy

Potentially Roy Williams for this work with distributed triangular meshes.

CRPC Support

My support comes partly from CRPC and partly from Caltech

Outreach/Technology Transfer

I'm collaborating with the John Reynders and Willey Lee at the Princeton Plasma Physics Lab.

C17: Application of Lattice Gas Methods to the Computation of Compressible Flow

Scientific Goal

Determine the dynamic and thermodynamic properties of multi-speed lattice gases, for application to the computation of compressible flow (Euler and Navier Stokes equations).

Parallelization

Lattice gases are amenable to straightforward domain decomposition and are self-balancing. Direct simulations are relatively inefficient because each site is treated at each time step, but equilibrium flux methods yielding finite-difference Euler solvers are very efficient. To properly resolve complex problems large amounts of memory in each node are required.

CRPC Facilities

Intel Gamma

CRPC Collaborators

G. Doolen, LANL

B. Sturtevant

CRPC Funding

One graduate student

Key Publications

B. Nadiga, "A Study of Multi-Speed Discrete-Velocity Gases", CIT Thesis, 1992 (In preparation).

C18: Euler Flow Simulations

Scientific Goal

A long unresolved problem is whether three-dimensional Euler flow can develop singularities in finite time. A recent proposed flow by Grauer-Sideris lets us simulate a fully developed three-dimensional flow with two-dimensional calculations. We plan to use the Delta to computer these flows with maximum possible resolution.

Parallelization

The main issue is that the I/O capabilities of the Delta are taxed to the limit. Obtaining the concurrent version of the computational parts of program was fairly straightforward.

CRPC Facilities

Use of the Gamma and Delta Touchstone

CRPC Collaborators

Eric F. Van de Velde, Daniel Meiron

CRPC Funding

1/3 Eric F. Van de Velde

Outreach/Technology Transfer

Mike Shelley, Russel Caflish

C19: Concurrent Adaptive Finite-Element Methods for Reaction-Diffusion Equations

Scientific Goal

Singularly perturbed reaction-diffusion equations exhibit multiple time and space scales and locally complex phenomena, like moving transition layers to steep fronts. Robust and accurate computations require very fine space-time grids. This is computationally feasible with irregular adaptive grids. We shall develop concurrent adaptive methods involving the finite-element method, strategies for refining and coarsening of the grids, and numerical solution methods.

Parallelization

Concurrency issues are addressed by the Associations software developed by Roy Williams.

CRPC Facilities

Use of Gamma and Delta Touchstone

CRPC Collaborators

Eric F. Van de Velde, Roy Williams

CRPC Funding

1/3 Eric F. Van de Velde, Roy Williams, Donald Estep

Outreach/Technology Transfer

Possibly some sites interested in specific applications.

2. Optimization (Rice)

See appropriate section of this report.

3. Computational Biology (Rice, Baylor)

Over the past six months, we have taken the first steps towards forming interdisciplinary research team with researchers in the Keck Center for Computational Biology to attack the most important biological applications via parallel computing. The teams will consist of computer scientists from the CRPC, scientists with high-impact computational problems in biology and new research staff members hired to produce parallel implementation that can be widely distributed and used. Based on the strengths of the Keck Center and the importance of the problems, we have launched grand challenge efforts in four major areas: molecular dynamics, medical imaging, genetics, and neuroscience. In each of these areas, the Keck Center will provide seed resources to get the effort started and, in conjunction with the CRPC, will seek complementary funding to bring the effort to fruition. The parallel CHARMM effort described in the following section is a prime example of CRPC outreach to problems in computational biology in which Keck resources are brought to bear on specific computational problems and applications that can have broad and dramatic impact on biology over the next decade.

4. Molecular Modeling (Rice, Syracuse, Harvard)

This is a collaboration with Karplus's group at Harvard and his co-workers, including C. Brooks (C.M.U.), B. Brooks (M.I.H.), and K. Schulter (Illinois). We intend that the computer science work will involve Rice, Syracuse and a collaboration with Joel Saltz. Currently, we are pursuing a two-pronged approach with the short goal of modifying the existing CHARMM code. In the longer term, we plan the development of a new code using Fortran D and new chemistry inputs. We will also explore new algorithms, such as those developed for astrophysical applications by Salmon and Edelsohn.

R11: Parallelization of CHARMM Molecular Dynamics Package

(K. Kennedy, A. Carle, J. Mellor-Crummey, G. Fox, S. Ranka, in collaboration with M. Karplus (Harvard), C. Brooks (CMU), J. Saltz (ICASE), G. Phillips (Keck), M. Pettitt (Keck/UH), B. Brooks (NIH), W. Celmaster (DEC), J. Graham (INTEL), J. Bailey (TMC))

Scientific Goal

Our efforts in parallelization of CHARMM will be focused on a number of key problems in biochemistry/biophysics for which the computational enhancements will greatly benefit the scientific applications. In addition, it is crucial that our proposed solutions are significant when applied to real problems. Therefore, we have selected three general topic areas in which applications will be supported and explored. These are in the general area of virial uncoating mechanism and its inhibition, the characterization of non-compact states of proteins for models of protein folding intermediates and the investigation of cooperativity in the binding of small ligands (oxygen) to hemoglobin. In each of these three areas computations being carried out today are pushing the technological and resource base for large-scale molecular simulations, and in that sense could benefit greatly by the software

developments proposed in this grant request. In addition, each problem is itself a problem of great interest within the community of biochemistry/biophysics.

Parallelization

Our approach to parallelization of CHARMM will involve two concurrent phases. Phase 1 will attempt to improve the preliminary implementations of parallel CHARMM for the iPSC/860 developed by Phillips at Rice and Bernie Brooks at NIH, and to port those codes to the Delta machine and the CM-5. It is expected that this effort will produce a useful full-feature version of CHARMM by the end of 1992. Phase 2, which would proceed in parallel, will pursue a major reimplementation of CHARMM using state-of-the-art algorithms and written in "high-performance Fortran."

Publications

1. S. L. Lin, J. Mellor-Crummey, B.M. Pettitt, G.N. Phillips, Jr, "The Parallelization of CHARMM for the iPSC/860: Molecular Dynamics on a Distributed-memory Multiprocessor"
2. S. Ranka, G. Fox, J. Saltz, R. Das "Parallelization of CHARMM Molecular Dynamics Code on Multicomputers"

S18: Adaptive Solvers for Astrophysical Problems

(D. Edelsohn, G. C. Fox, in collaboration with J. Salmon (Caltech), B. Elmegreen, M. Lemke (IBM, Yorktown), and D. Quinlan (Colorado))

Scientific Goal

The main focus during the past research year has been to study algorithmic properties of adaptive elliptic differential equation solvers such as the Barnes-Hut Tree code, the Fast Multipole Method, and adaptive multilevel (*multigrid*) methods. New underlying common features which allow techniques developed for one class of methods to be utilized in other classes have been found.

In the coming year more work will be done regarding the optimal sub-division depth for adaptive techniques. A multi-component N-body simulation will be developed using the *Asynchronous Fast Adaptive Composite* multi-level method with *Method of Local Corrections* enhancements based upon the C++/M++ programming system. Under the guidance of Bruce Elmegreen of IBM Corporation's Physical Sciences Division, this parallel computing research will be applied to the study of arms in spiral galaxies using 3D simulations.

Parallelization

Work has been also done on providing a parallel programming environment based upon C++ using array extensions similar to Fortran 90.

We are looking at providing extensions described in the Fortran D specification as user-controllable optimizations. We plan to re-implement the communication using ParaSoft's Express MIMD communication library to remove device-dependencies and improve portability.

Publications

1. D. J. Edelsohn and G. C. Fox, "Hierarchical Tree-Structures for Particle Simulations", Syracuse Center for Computational Science, Syracuse University, Presentation at Physics Computing '91, San Jose, California, May 1991

2. D. J. Edelsohn, "Hierarchical Methods for N-body Problems", Syracuse Center for Computational Science, Syracuse University, Presentation at Supercomputing '91, Albuquerque, New Mexico, November 1991

3. D. J. Edelsohn, "Hierarchical Tree-Structures as Adaptive Meshes", Syracuse Center for Computational Science, Syracuse University, CRPC-TR91186 / SCCS-192, November 1991, (submitted for publication)

5. Geoscience (Rice)

Below is a brief description of the technical results obtained by the "Flow in Porous Media Parallel Project" which has major funding from the State of Texas. 18 papers are in preparation.

I. Parallel Scaling Issues in Reservoir Simulation
(Todd Arbogast, Clint Dawson, Phil Keenan, Marcelo Rame, Mary Wheeler)

Extensive computational studies on two codes, RICE-UTCHEM (parallel INTEL version of UTCHEM) and PIERS (black oil code released to Rice by Exxon Production Research) have been carried out.

UTCHEM is a three-dimensional multicomponent, multiphase chemical flood reservoir simulator which involves 11 components and a pressure equation. It is used by more than 20 oil companies and has been highly vectorized. UT has benchmarked this code for a collection of machines and I believe they are willing to provide such a copy.

We are presently working with UT and computer vendors, Digital (Maspar) and Thinking Machines (CM-5) on parallel versions. Our work on the Delta has been disappointing so far. We are presently involved in speeding up the linear solver by domain decomposition preconditioners (see comments below). At the same time, we have formulated some new advection schemes which we plan to put in RICE-UTCHEM.

The work on RICE-UTCHEM was done primarily by Marcelo Rame with assistance from Todd Arbogast and Clint Dawson. Rame ported the code to the INTEL Delta. These preliminary results are disappointing but not surprising; there appears to be no speedup on the Delta over the INTEL/860. In addition the Delta environment is still not stable. ***** just saw Paul Messina and we will try to make some runs for larger problems....*******

We are now beginning to address certain subroutines where the scaleup is poor; in particular the linear solver. Dawson recently incorporated a domain decomposition procedure into the RICE-UTCHEM reservoir simulator. This procedure is based on a method defined by Glowinski and Wheeler and then extended to parabolic equations by Dawson and

Dupont. For speedups the domain decomposition procedure needs to be accelerated by applying a multilevel technique as defined by L. Cowsar and Wheeler or by developing a good preconditioner.

The work on PIERS was primarily done by Phil Keenan and Mary Wheeler. To facilitate studying scaling behavior on the IPSC/3 Hypercube several computer science type changes were made to the code. In particular the dimensioning program was revised to use command line arguments and to understand equations; the source files were split into individual subroutines. PIERS was originally written to run on only 16 processors. With help from Exxon researchers we modified the code to run on 256 processors. We have successfully tested the use on 64 processors. In addition the code has been modified to treat rectangular mesh orderings in addition to gray code nearest neighbor topologies. The I/O and graphics has been modified for use on Sparc stations rather than on IBM PC's. In addition Keenan wrote a graphics display program for viewing graphical results with X-11 and with PostScript.

Six major test cases have been studied. The bottleneck in parallel speedups is the coupled nonlinear system of equations that needs to be solved at each time step. Even so, this code has better efficiencies than any of the parallel results that have been reported to date. Our future work includes addressing the nonlinear solver.

II. Algorithms

(Todd Arbogast, Ashok Kumar Chilikapati, Clint Dawson, Mary Wheeler, Nai-ying Zhang)

Arbogast and Wheeler have defined and analyzed two mixed characteristic procedures for treating linear advection problems. These procedures are conservative. In addition maximum principles can be enforced by applying slope limiting techniques.

Code implementation of these schemes for three dimensional problems which arise in unstable miscible displacement and to waste management problems is being carried out by Ashok Kumar Chiliapati. Recent numerical experiments indicate that the methods are computationally more efficient than characteristic Galerkin methods.

In addition, Arbogast, Wheeler, and Zhang have also defined and analyzed a new method for nonlinear degenerate advection diffusion problems. This work applies to immiscible displacement problems.

III. Applications: Dual-porosity Simulation

(T. Arbogast)

Previous work included the completion of a code to study the dual-porosity model in the case of a single-phase fluid. This code assumes two space dimensions, and it has three main options:

1. microscopic simulation on the finest scale (i.e., this option can be used to solve for the "true" flow for small problems);

2. the dual-porosity approximation to the true flow;

3. unfractured porous media simulation.

Computational results demonstrated that the dual-porosity model of a single-phase fluid predicts very well the true microscopic solution for the flow of either an oil or a gas (CO_2). The use of some average permeability and porosity in an unfractured simulation cannot match the microscopic solution in general.

Work has progressed on a two-phase, incompressible, immiscible code which simulates, for example, a waterflood. It has the same three options as the first code, so a true description of the flow can be obtained for a very small reservoir (with a few matrix blocks) and compared to the dual-porosity approximation and an unfractured simulation.

The goal is to verify the accuracy of the dual-porosity model that is considered. This model is derived by homogenization of the microscopic, fine scale description of the flow. This model has one ambiguity related to the location of the oil-water contact line in the fractures. A combination of the computational results and theoretical concerns from the theory of homogenization should remove this ambiguity.

A modification of this code will be made to test the accuracy of a mathematically motivated simplification of the equations. This new model is much simpler than the original and should be reasonably accurate in practice. (It is similar to a model of deSwaan, SPEJ '78).

These two codes are completely parallelizable in the matrix rock computations. Future work includes running one or both codes on a parallel machine; this will require a domain decomposition approach for the fracture equations.

OUTREACH

1. Flow group visited UT Austin on 31 October 1991.
2. Affiliates meeting was held on 17 December 1991 with 11 companies attending.
3. Numerous technical presentations were made by the group at various meetings and colloquium. To name a few, Exxon, ACM Houston chapter, Lambrecht Germany, University of New Mexico, Center for Nonlinear Studies in Science at Los Alamos, Pacific Northwest Laboratory, American Geophysical Union National Meeting in San Francisco, University of Texas, and University of Houston.

6. Industrial Outreach (Syracuse

New York State has funded a broad outreach program to develop the industrial applications of parallel machines. Initial work has involved detailed projects with IBM and General Electric and a broad survey of 20 companies, summarized in Table 3. We expect to start soon two new significant projects. One involves transient stability analysis of Niagara Mohawk's electrical transmission network. The second involves a financial modeling problem with a Wall Street company which is based on a pilot project described below, with the management school at Syracuse. We also have a summary of our current survey of New York industries and the possibilities for parallel computing.

Company/Agency	Applications	Current Big Computer	Problem Class	Can One Use HPC?	Relevant Code Exist?
General Electric Company (Syracuse)	Acoustic beam forming	Special & iWarp	Synchronous	yes	no
General Electric Company (Syracuse)	Ocean Environmental Modeling	many	Loosely Synchronous	yes	maybe
Grumman Corporation United Technologies Corporation	CFD & Structures	Cray & CM-2	Synchronous Loosely Synchronous	yes	yes
Grumman Corporation PAR Technology Corporation	Avionics & Command and Control (JSTARS)	Special & VAX	Loosely Synchronous Complex	yes	probably not
Otis Elevator Company	Manage repair database	Not implemented yet	Loosely Synchronous Complex	yes	no
Lamson Corporation (small Syracuse Co.)	Design exhaust pumps	None	Loosely Synchronous	?	yes
Carrier Corporation	Design air conditioners	VAX	Loosely Synchronous	? (is physics of noise understood)	yes
Atlantic Research Corporation (ARC)	Antenna Fields	VAX	Embarrassingly Parallel	yes	yes (little change)

Table 3A. New York State Parallel Computing Corporate Survey

Company/Agency	Applications	Current Big Computer	Problem Class	Can One Use HPC?	Relevant Code Exist?
Niagara Mohawk Power Corporation	Control & plan electrical power grid	IBM 3090	Loosely Synchronous	yes	no
MONY Insurance Co.	Process insurance policies	IBM 3090	Embarrassingly Parallel	no	yes
Empire Blue Cross/ Blue Shield	Process Medicare claims	IBM 3090	Embarrassingly Parallel	no	yes
Prudential Securities Inc.	Predict values of securities	various, iPSC/860 & Hypercube	some Embarrassingly Parallel	?	yes (small)
Securities Industry Automation Corporation (SIAC)	Control trading on floor of NY and American Stock Exchange	Tandem & PC	Embarrassingly Parallel	?	yes
U.S. Air Force (ESD)	Command and control &battle management for SDI	many	Loosely Synchronous Complex	yes	probably not
U.S. Army	Data fusion	VAX	Loosely Synchronous Complex	yes	somewhat
PAR Technology Corporation	Infrared sensors in aircraft	Special-purpose	Synchronous	yes	no
Abrams/Gentile Entertainment Inc. (AGE)	Entertainment, toys, theme parks, virtual reality	Must be low-cost (toys) Supercomputer (theme park)	Loosely Synchronous (complex)	yes	no

Table 3B. New York State Parallel Computing Corporate Survey

S19: Financial Modeling

(G. Cheng, K. Mills, M. Vinson, G. Fox)

Scientific Goal

Investigate relevance of parallel computing to financial modeling in both academic and the commercial (Wall Street) community. Initially, we are using an academic option pricing model and comparing it with a data set consisting of three years sets of trades.

Parallelization

We have efficient programs in CMFortran and MPFortran for Maspar. We found serious problems with treatment of arrays in MPFortran which are a useful lesson for HPFortran.

Collaboration

T. Finucane in Syracuse University School of Management

Publications

G. Cheng, K. Mills, M. Vinson, and G. Fox, "Mapping a Two-Dimensional Problem Structure to a One-Dimensional Array for Improved Performance on the DECmpp", preprint

7. Statistical Physics and Optimization

S11: Random Surfaces and Quantum Gravity

(M. Bowick, P. Coddington, L. Han, G. Harris, E. Marinari, in collaboration with C. Baillie)

Scientific Goal

Numerical simulations of Triangulated Random Surfaces provide an invaluable tool to try to build and understand a sensible theory of quantum gravity. On the physics side we have followed both an analytical and a numerical approach, trying to clarify issues connected to 2D matrix models, and the existence of a phase transition where a continuum theory could be defined.

We plan to set up large scale simulations to understand more about this possible continuum theory. By means of analytic work we will try to decide which are the relevant questions one has to answer.

Parallelization

On the computing side we are at present porting our code (by means of *Express*) to the Delta machine. We also plan to run large scale simulations on the CM-5, writing a program that is message passing based for the inter node communication, and which use effectively the four vector chips that are on each node (exploiting the four parallelism, the presence of one adder and one multiplier in each floating point chip, and the pipelining features of the chip). The problem of the local parallelization of the code will be non-trivial because the *dynamical triangulations* are, in principle, highly non-homogeneous.

Publications

1. C. F. Baillie, D. A. Johnston and R. D. Williams, "Computational Aspects of Simulating Dynamically Triangulated Random Surfaces", *Comp. Phys. Comm.*, **58**, (1990) 105

2. C. F. Baillie, S. M. Catterall, D. A. Johnston and R. D. Williams, "Further Investigations of the Crumpling Transition in Dynamically Triangulated Random Surfaces", *Nucl. Phys. B*, **348**, (1991) 543

3. C. F. Baillie, and D. A. Johnston, "Crumpling Dynamically Triangulated Random Surfaces in Two Dimensions", *Phys. Lett. B*, **258**, (1991) 346

4. M. Bowick and E. Brezin, "Universal Scaling of the Tail of the Density of Eigenvalues in Random Matrix Models", *Phys. Lett. B*, **268**, (1991) 21

5. K. N. Anagnastopoulos, M. J. Bowick and N. Ishibashi, "An Operator Formalism for Unitary Matrix Models", Mod. Phys. Lett. A, **6**, (1991) 2727.

S12: Heteropolymers and Disordered Systems

(E. Marinari, P. Pliszka, in collaboration with G. Iori, G. Parisi, M. V. Struglia)

Scientific Goal

We are studying the behavior of heteropolymers with random self-coupling. This is a specific part (with potential application to the long standing problem of protein folding) of a more general research program investigating optimization algorithms for disordered systems.

Parallelization

Our preliminary studies have been done on workstations. Now we have a code running on the CM-2, and we plan to run large scale simulations. On the CM-2 we are also studying the different performances of different kind of parallelizations. We will also adapt the code to the DECmpp.

Publications

1. G. Iori, E. Marinari and G. Parisi, "Random Self-Interacting Chains: a Mechanism for Protein Folding", *J. Phys. A: Math. Gen.*, **24** (1991) 5349

2. G. Iori, E. Marinari and G. Parisi and M. V. Struglia, "Random Chains and Protein Folding", at the 1991 Trieste ICTP Conference on *Physics of Biological Systems*

S13: Optimization Methods, Cluster Reconstruction and Monte Carlo Techniques

(J. Apostolakis, P. Coddington, L. Han, E. Marinari, A. Su in collaboration with C. Baillie, A. Sokal and G. Parisi)

Scientific Goal

The problem of cluster labeling is becoming more and more crucial in different domains of physics. We have first encountered it as one aspect of percolation based non local Monte Carlo methods, crucial in order to fight critical slowing down in spin models. We plan to continue our study of how different cluster methods work on computers with different architecture and in different physical situations.

Parallelization

We have concentrated on two approaches to component labeling. One of these combines regular lattice communications (with the scan operation) and general communication (*get & send*) to quickly propagate labels over small and large distances respectively. The other is a simple multi-scale method which we tested for the critical Potts and pure percolation models and found to have good scaling properties, converging in a number of steps proportional to the logarithm of the length of the lattice.

We have also recently proposed a new Monte Carlo procedure which, allowing β to become a dynamical variable, works as an effective annealing for finite T models. We plan to implement this approach for different problems (we are writing a Spin Glass Three-Dimensional Ising code on the DECmpp), and to study its effectiveness.

Publications

1. C. F. Baillie and P. D. Coddington, "Cluster Identification Algorithms for Spin Models—Sequential and Parallel", *Concurrency: Practice and Experience*, **3**, (1991) 129

2. C. F. Baillie and P. D. Coddington, "Comparison of Cluster Algorithms for 2D Potts Models", *Phys. Rev. B*, **43**, (1991) 10617

3. C. F. Baillie and P. D. Coddington, "Parallel Cluster Algorithms", *Nucl. Phys. B (Proc. Suppl.)*, **20**, 76 (1991)

4. P. D. Coddington and C. F. Baillie, "Empirical Relations between Static and Dynamic Exponents for Ising Model Cluster Algorithms", to appear in *Phys. Rev. Lett.*

5. P. D. Coddington and C. F. Baillie, "Dynamical Exponents for Potts Model Cluster Algorithms", to appear in *Nucl. Phys. B (Proc. Suppl.)*, (1992)

6. J. Apostolakis, P. Coddington and E. Marinari, "A Multi-Grid Cluster Labeling Scheme", *Europhys. Lett.*, **17**, (1992) 189

7. J. Apostolakis, P. Coddington and E. Marinari, "New SIMD Algorithms for Cluster Labeling", NPAC and Syracuse University preprint, February 1991

8. E. Marinari and G. Parisi, "The Annealed Annealing: a New Efficient Monte Carlo Method", Roma *Tor Vergata* and SCCS preprint

S14: Monte Carlo Simulations of Spin Models and Lattice Gauge Theories: Improved Algorithms and the Scaling Regime

(J. Apostolakis, G. C. Fox, E. Marinari, in collaboration with C. Baillie, M. Guagnelli, R. Gupta and G. Parisi)

Scientific Goal

We are trying here to match the use of large resources in order to set up large scale simulations with the need of finding effective algorithms which will perform well on massively distributed parallel computers, when the space of configurations becomes very large.

A goal of ours is to test asymptotic scaling in the $O(3)$ model, the simplest system with similar scaling behavior to QCD, by simulating the model at large correlation lengths. In this, we followed up on our work using the overrelaxation algorithm by using the Swendsen-Wang-Wolff algorithm, in a simulation on a Connection Machine-2 with a program written in C*. The nature of this simulation algorithm breaks up the regular lattice into very irregular parts with fractal paths.

We intend to utilize our $O(3)$ simulation program and extend it to perform a Monte Carlo Renormalization Group calculation to determine the renormalization group flow starting from the standard action.

Parallelization

The most important challenge of this work was to investigate new algorithms for connected component labeling on an SIMD computer, and, at the same time, find an efficient method to evaluate an improved estimator.

Publications

1. J. Apostolakis, C. Baillie and G. C. Fox, "Investigation of the 2D $O(3)$ Model Using the Overrelaxation Algorithm", *Phys. Rev. D*, **43**, (1991) 2687

2. M. Guagnelli, M. P. Lombardo, E. Marinari, G. Parisi and G. Salina, "The Quenched Mass Spectrum in Lattice QCD on a One Gigaflops Computer", to appear in *Nucl. Phys. B*

3. M. Guagnelli, M. P. Lombardo, E. Marinari, G. Parisi and G. Salina, "The Quenched Mass Spectrum in Lattice QCD", to appear in *Nucl. Phys. B (Proc. Suppl.)*

S15: Dynamical Depinning Transitions

(A. Middleton, O. Biham, in collaboration with D. S. Fisher and P. B. Littlewood)

Scientific Goal

Many physical systems, including superconductors, interfaces between two fluids, and charge-density waves, can be described as elastic media subject to random pinning forces

and an external drive. A problem in these systems with collective transport, is that of the "depinning" transition, where the external force overcomes the pinning and the system "slides". This transition is a novel type of dynamic critical phenomenon.

Our research has focussed on understanding the critical exponents of the transition, mode-locking in an a.c. drive, and the rounding of the transition due to thermal effects. We have performed large-scale simulations of the equations of motion and use analytic bounds found by one of us to improve the numerical simulation. The disorder in the system requires the study of large $(128^3, 512^2)$ systems and averages over many realizations.

We have been able to compare the behavior of related models and have shown that an automaton can reproduce important aspects of the behavior of the full model, allowing for a relatively precise determination of critical exponents. We plan to extend these calculations to a model in the co-moving frame of the system (rather than the current lab frame) which will necessitate a novel numerical approach.

Parallelization

Much of the problem is straightforward to parallelize, by describing the elastic medium as a set of particles on a lattice, connected by springs. The extension to studies in the co-moving reference frame will require more careful communications between neighboring processors on a SIMD machine and enough memory for each particle to store a large part of the rest-frame pinning potential.

Publications

1. A. A. Middleton and D. S. Fisher, "Critical Behavior of Pinned Charge-Density Waves below the Threshold for Sliding," *Physical Review Letters*, **66**, 92 (1991)

2. A. A. Middleton, D. S. Fisher, and P. B. Littlewood, Reply to comment on "Critical Behavior of Pinned Charge-Density Waves below the Threshold for Sliding," *Physical Review Letters*, **67**, 3873 (1991)

3. A. A. Middleton, "Asymptotic Uniqueness of the Sliding State for Charge Density Waves," *Physical Review Letters*, **68**, 670 (1992)

4. A. A. Middleton, O. Biham, P. B. Littlewood and P. Sibani, "Complete Mode-locking in models of Charge-density Waves," to appear in *Physical Review Letters*

5. A. A. Middleton, "Thermal Rounding of the Charge Density Wave Depinning Transition," to appear in *Physical Review B (Rapid Communications)*

6. A. A. Middleton and D. S. Fisher, "Sub-threshold Critical Behavior of Charge-Density Waves," in preparation

S16: Cellular Automaton Models of Traffic Flow

(O. Biham, A. A. Middleton)

Scientific Goal

We are studying cellular automaton models of particle flow, or of the flow of "cars" on a grid pattern in two dimensions, in order to gain insight into such particle flow or traffic.

This simple model captures many essential features of traffic or particle flow, including segregation into bands (pattern formation) and the occurrence at high densities of traffic jams, or stuck configurations. The transition between flowing and stuck traffic is very sharp (first order in the density for several models).

We are currently studying simple grids, but will extend these simulations to more complex geometries to explore the effects of disorder (such as real traffic geometries or fluid flow in porous media).

Parallelization

We have performed these simulations on the DECmpp. We have packed configurations into 32-bit integers and used bit-shift operators in MPL for efficiency. Particle flow uses nearest-neighbor communications. We currently move 1.4 billion cars per second, on large grids (to avoid finite-size effects, which are quite large).

Publication

1. O. Biham, A. A. Middleton, and D. Levine, "Traffic Jams as a First-Order Dynamical Phase Transition", to be published

S17: High Temperature Superconductivity: Elastic String in a Random Potential

(A. A. Middleton, in collaboration with M. Dong, M. C. Marchetti, and V. Vinokur)

Scientific Goals

The discovery of high-temperature superconductivity has spurred a large effort to understand the properties of these new materials. In particular, the motion of magnetic flux has implications for the practical use of such materials and is also of great theoretical interest.

We model magnetic flux lines in the superconductor as elastic strings. These strings experience forces due to currents in the superconductor and impurity pinning forces. We calculate the motion of these flux lines as a function of current and pinning strength. Below the critical current, the flux lines do not move.

Parallelization

The numerical study of this model is based upon the numerical integration of the equations of motion for the elastic string. Using machines with many nodes, we are able to simulate many strings simultaneously, to quickly explore the possible parameter space. We have used both CM-Fortran and C* in the simulations.

Publication

1. M. Dong, M. C. Marchetti, A. A. Middleton and V. Vinokur, "Elastic String in a Random Potential," submitted for publication