SCIENTIFIC COMPUTING
RESEARCH ENVIRONMENTS
FOR THE MATHEMATICAL SCIENCES

NSF 98-72009

Final Report

Matthias Heinkenschloss Petr Klouček Yin Zhang

August 30, 2001

Summary

This report describes the research projects and accomplishments made possible through the availability of the sixteen processor SGI Origin 2000, purchased in parts with the funds from NSF SCREMS grant NSF 98-72009. To date the SGI Origin 2000 has served as the main computing facility in many interdisciplinary projects involving 48 faculty, research scientists, postdocs, graduate and undergraduate students from six departments at Rice University, as well as several visiting scholars and collaborators from other universities. Computations performed on the SGI Origin 2000 have led to 44 journal articles, proceedings articles, and technical reports. Availability of the SGI Origin 2000 on campus has led to a significant increase in the complexity of the problems we are able to tackle and it has served as the catalyst for several of the research projects described in the report. The sixteen processor SGI Origin 2000 continues to be a widely used and important computing resource on campus.

1 Introduction

In July 1998 the three PIs were awarded SCREMS grant NSF 98-72009 which, together with cost sharing from the Office of the Vice Provost for Research and Graduate Studies, the Office of the Dean of the George R. Brown School of Engineering, and funds of the PIs was used to purchase an eight processor SGI Origin 2000 R12000 with 8 GB of RAM and approximately 100 GB of hard disk. Subsequently, a group of researchers in Computer Science and the Center for High Performance Software (HiPerSoft) Research at Rice University received funding for the purchase of an SGI Origin parallel computer. To leverage resources, these funds were used to enlarge the number of processors in the existing SGI Origin to sixteen and to increase memory appropriately.

The sixteen processor SGI Origin 2000 R12000, called MAPY, has been used for research on *Numerical Solution of Time–Dependent and Pseudo Time–Dependent Optimization Problems*, *Advanced Numerical Methods for Computation of Thermodynamics of Microstructures in Solids*, *Optimization Algorithms for the*

Phase Problem in Crystallography by the three PIs as outlined in the SCREMS research proposal and for *Compiler, Library, and Tools Research* conducted by research groups in CS and HiPerSoft. Availability of MAPY has promoted the interaction of research groups across campus on computing intensive projects, which has led to several new, funded, interdisciplinary research projects. MAPY was and still is one of the largest computers available on campus. From the beginning MAPY resources were made available to select research groups that required computing resources not available otherwise.

To date, MAPY has been used by over

- 21 faculty, faculty fellows, and research scientist,
- 5 visiting faculty and visiting scholars,
- 2 postdocs,
- 19 graduate students
- 6 undergraduate students

from the

- Bioengineering (BIOE),
- Computational and Applied Mathematics (CAAM),
- Computer Science (CS),
- Earth Science (ESCI),
- Mechanical Engineering and Materials Science (MEMS),
- Physics and Astronomy (PHYS)

departments at Rice University in interdisciplinary research projects, as well as by a few scientists from other universities collaborating on research projects with the PIs. In addition to the researchers associated with the projects detailed in the following section, MAPY has been made available to a number of faculty, postdocs, and graduate students in various departments at Rice University for software testing and similar purposes.

Short summaries of the main research projects conducted on MAPY are given in the next section together with a list of personnel involved in the respective projects and publications generated. To date, results generated with the help of MAPY have been used in over 6 journal articles, 27 proceedings articles, 11 technical reports, 2 PhD theses, 1 masters thesis, and 1 Habilitations thesis. Furthermore, MAPY has played a pivotal role in the development of several software packages that are made available though the research groups. Several PhD theses which use MAPY are still in progress. The technical reports are expected to lead to journal publications. Most projects outlined in the next section are ongoing and will lead to additional publications.

In summary, the SCREMS grant has been catalyst in fostering collaborative and interdisciplinary research across the Rice campus and beyond. The availability of this computing resource has enabled us to tackle problems otherwise beyond reach. The impact of the SCREMS grant has far exceeded the original vision in our proposal.

2 Projects

2.1 Optimization of Systems Governed by Time–Dependent Partial Differential Equations

Many physical phenomena can be modeled by linear or nonlinear time–dependent partial differential equations (PDEs). Optimization of such systems, in the context of optimal control, optimal design, or parameter estimation plays an important role in science and engineering. Algorithms for the solution of time–dependent PDEs often involve marching in time, starting from an initial condition. In optimization, however, the values of the solution of the PDE at later times feed into the optimization at early times. This coupling in time makes the practical solution of these very large-scale optimization problems challenging. Examples of such optimization problems are given in Sections 2.2 and 2.9 of this report. For specific problems, suboptimal control schemes, such as reduced bases techniques or instantaneous control, have been applied successfully. However, the analysis of these techniques is incomplete and the limits of their applicability are not clearly described.

This research builds upon the existing work in optimization methods and in control of time–dependent PDEs to construct a framework for iterative methods for this class of problems. Our framework aids the analysis of iterative methods and their implementation for a broad class of applications. To explore and evaluate this general framework, we have applied our methodology to some smaller linear quadratic optimal control problems. An extension to nonlinear optimal control problems related to problems in fluid mechanics and the application to the one of our target applications, the problem described in Section 2.9 are in progress.

Our target applications involve about $O(10^7)$ state variables and $O(10^5)$ control variables. Without a computing resource like MAPY testing and further development would only be possible on overly simplistic problems. Our iterative methods introduce parallelism into the problem at the optimization level, for example through time–decomposition, in addition to already exisiting parallelism on the PDE level. Thus MAPY is a good hardware platform for development and testing of our software.

This on-going research project is funded by NSF DMS-0075731. Descriptions of our framework and initial results are presented in [34, 19].

Personnel: S.S. Collis (Assistant Prof., MEMS), M. Heinkenschloss (Associate Prof., CAAM), A. Comas (Graduate student, CAAM).

2.2 Aeroacoustic Physics and Control

This ongoing project applies large-scale aeroacoustic simulation along with optimization to construct optimal control strategies that reduce the adverse effects of flow induced noise. A motivating application is the prediction and control of far-field sound due to blade-vortex interaction in rotorcraft. The physics for this problem are modeled by the compressible Navier–Stokes equations or Euler equations in the near-field which is coupled to a wave equation that models the acoustics in the far-field. This project has made use of MAPY for the following activities: (1) high accuracy fluid-structure induced aeroacoustic simulations; (2) multi-physics optimization: near-field uses compressible Navier-Stokes, far-field uses acoustic analogy; and (3) development of efficient, parallel optimization algorithms. Our research has addressed fundamental issues arising in the numerical solution of optimal control problems for unsteady compressible flows, such as choice of control spaces and regularization terms, adjoint and gradient computations, but we also have investigated specific control problems. Computing resources comparable to MAPY are necessary for these computations, since even modest discretizations of the unsteady flow physics generate 10^7 state variables and 10^5 control variables. Optimal control runs with modest discretizations take about 8 hrs using eight processors on MAPY.

Figure 1 shows the result of an optimal control calculation. The uncontrolled flow computations show pressure variations resulting from the interaction of a vortex and a point source. The goal is to reduce the average aeroacoustic noise in the rectangle Ω_0 near the bottom wall over time $[t_0, t_f]$ through suction and blowing of fluid on a portion of the bottom wall. The average aeroacoustic noise is quantified by $\int_{t_0}^{t_f} \int_{\Omega_0} (p - p_{\text{amb}})^2$, where p_{amb} is the ambient pressure and p is the flow pressure. Without control, $\int_{t_0}^{t_f} \int_{\Omega_0} (p - p_{\text{amb}})^2 = 0.49$. With optimal suction and blowing on the bottom wall, $\int_{t_0}^{t_f} \int_{\Omega_0} (p - p_{\text{amb}})^2 = 0.34$.

Figure 1: Optimal control of aeroacoustic noise through suction and blowing on the bottom wall. The plots in the top row show $p - p_{\text{amb}}$ for the no control case and the plots in the bottom row show $p - p_{\text{amb}}$ for the optimally controlled case. Optimal suction and blowing on the bottom wall reduced the objective function $\int_{t_0}^{t_f} \int_{\Omega_0} (p - p_{\text{amb}})^2$ to 0.34 compared to 0.49 in the no control case.

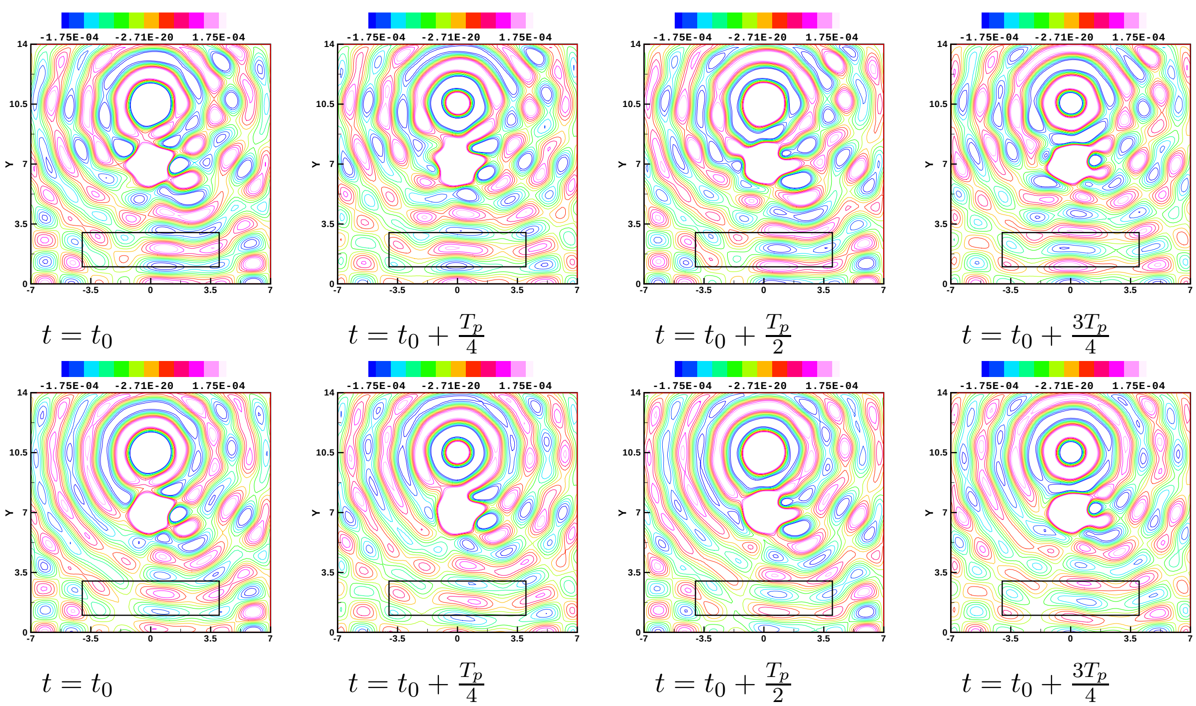

Most of this research is funded by TX-ATP (3604-001) and research on prediction and control of far-field sound due to blade-vortex interaction in rotorcraft is carried out in conjunction with Bell Helicopter, Textron. Present research findings are documented in [20, 21, 22, 58]. Papers on adjoint computations on the differential equation level for this problem and aeroacoustic control governed by the Euler equations are in preperation.

Personnel: S.S. Collis (Assistant Prof., MEMS), M. Heinkenschloss (Associate Prof., CAAM), M. Ulbrich and S. Ulbrich (Assistant Prof., TU Munich, Germany, visiting scholars CAAM 1999-2000), K. Ghayour (Postdoc., CAAM & MEMS), G. Chen, S. Ramakrishnan, Z. Smith (Graduate students, MEMS), N. Bou-Rabee, R. Purple and E. Rareshide (Undergraduate students, MEMS).

2.3 Modeling and active control of vibrations

In regular use, machinery is subjected to periodic stresses. This results in acoustic waves travelling through the material. Since these waves are small in comparison to the potential energy of the overall machine, conversion to heat is an effective method of noise reduction. Thus highly damping materials may be used either in part or in full to accomplish this task.

Shape memory alloys exhibit such significant damping properties. These are special alloys that change their microstructure from that of a stiff, rotationally symmetric phase (Austenite) to a ductile, less symmetric phase (Martensite) when cooled or put under stress. The desirable damping properties are a result of movement within the twin boundaries in the Martensite phase, as well as the motion of the incoherent Austenite-Martensite Interface, and these damping properties are significantly temperature dependent.

This research project focused on the development of active control of vibrations and noise reduction based on the improved models of shape memory alloys patches. We have completed the first phase of the development of computational modeling of microscopic-mesoscopic coupling necessary to model shape memory behavior. This resulted in the *subgrid projection method*. As a part of this research development we designed robust active control tools within the framework quasi-stochastic partial differential equations in order to control and stimulate motion of incoherent interfaces in active materials.

The results are published in [23], [41], [53].

Figure 2: Example of the active control of decrease of total energy as a function of variable ohmic heating on the boundary of shape memory alloy patch. This computational example indicates that 90% reduction of the energy is possible within 20 elastic wave vibrations (about 0.001s) in the 2-4kHz vibrations range, [41].

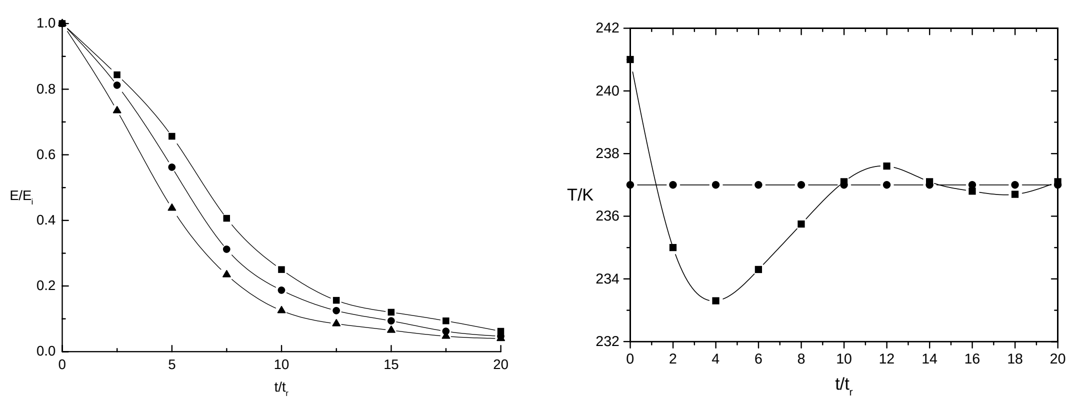

Personnel: D. Cox (Professor, Statistics Dept.), P. Klouček (Assistant Prof., CAAM), D. R. Reynolds (Graduate student, CAAM).
Collaborators: F. D. Fisher (Professor and Chair, Dept. of Machanical Engineering, U. of Leoben, Austria), T. I. Seidman (Professor, Dept. of Mathematics, UMBC), K. Tanaka (Dept. of Aerospace Engineering, Tokyo Metro. Inst. of Technology, Hino/Tokyo, Japan).

2.4 The Computational Modeling of Internal Surfaces in Crystals

Internal surfaces in crystals form domain walls in ferroelastic or co-elastic materials. These pattern boundaries are typically the result of a phase transition between two phases with different symmetries. The underlying phase transition thus often creates twin structures. Mechanical twinning is technologically important in superconductors, industrial shape memory alloys, and it is observed in mineralogically and petrologically relevant systems. Major mesoscale microstructures formed by twins include junctions where twins intersect, S-shaped domain walls, and needle twins.

From the mathematical point of view, the pattern morphology such as the needle twins corresponds to very low levels of the free energy. Various mesoscale microstructures are located at different local minima. The numerical approach to finding these patterns by minimization the free energy is rather difficult. These patterns are formed by a competition of events occurring on many different scales.

Figure 3: Basic comparison of the results. The first density plot corresponds to the minimization of of the free energy corresponding to a domain wall in a shape memory material in the domain $(0, 16) \times (0, 4)$ discretized by the regular mesh 121×121. The density plot exhibits characteristic "triangular regions" first observed by E. Lifschitz in 1944. The second density plot visualizes the result of the minimization of the same problem but computed with the adaptive scaling method. It seems that this minimizer corresponds to the more complicated domain structure observed by Privorotskij. The spatial resolution in the region $(0, 2) \times (0, 4)$ is the same for both examples. The white stripes running across the lower density plot correspond to the positions of the interfaces, [40].

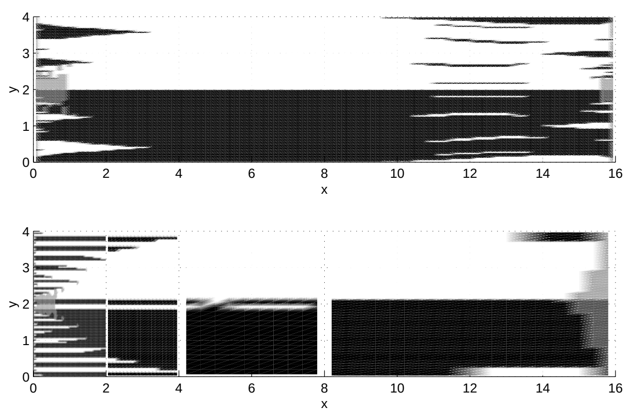

We have developed for this project a finite element method which is piecewise continuous on the microscopic scale of the spatial resolution h but that is discontinuous on the mesoscopic scale h^{δ}, $\delta \in (0, 1)$, with the scaling parameter δ determined by the equidistribution of the free energy. The method proved capable of capturing the morphology of mesoscopic needle twin structures frequently found in ferroic and co-elastic crystals. The approach is based on a domain decomposition method that interpolates between the scale on order of the size of crystal and the microscopic scale of finite element approximation h. The scale interpolation is enabled by incorporating a frequency adaptivity. The computational results yield geometrical structures observed in uniaxial ferromagnets and in the vicinity of the Austenite-Martensite interfaces.

The results are published in [40].

Figure 4: Comparison of the free energies. The plot depicts the value of $\mathcal{J}(u_h, (0, 1/2) \times (0, 4))$ against the Steepest Descent iteration number. The dashed line correspond to the minimization in the entire domain, while the solid line represents the value of the energy for the minimizers obtained by the method developed for this project, [40].

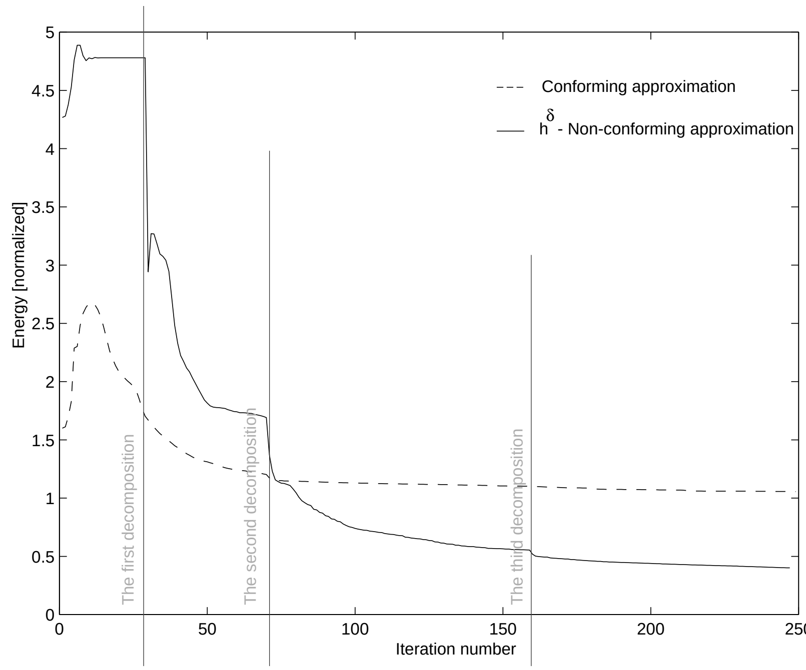

Personnel: P. Klouček (Assistant Prof., CAAM), L. A. Melara (Graduate student, CAAM).
Collaborators: A. Quarteroni (Professor, Dept. of Mathematics, EPF Lausanne, Switzerland).

2.5 The Detachment of Bubbles under a Porous Rigid Surface

The modeling of instabilities during the production of aluminium typically used is based on a Magneto-Hydrodynamic equations disregarding the gas production. In the absence of magneto-hydrodynamic effects, the gas production, i.e., its evolution under an anode and its contribution to the instability of the system, has been studied both experimentally and analytically.

The experiments show that there exist conditions at which bubbles, which result from a gas production, abruptly change shape during their growth while attached under a solid-liquid interface contributing thus to an instability of the fluid motion. They become suddenly flatter and movable. This implies existence of a thin liquid layer separating the bubbles from the rigid surface they were attached to prior to the shape change. This phenomenon is referred to as *detachment*. The detachment of gas bubbles under a surface and exposed to Archimedes drag was not studied before. This phenomenon is different from a gas bubble detachment frequently found in boiling liquids that are above the surfaces where the detachment takes place. We view detachment as a mechanism for changing the contact (wetting) angle, i.e., the angle between the tangent to the bubble and the flat rigid surface (anode) at the point of contact. The understanding of the origin of the detachment mechanism is a necessary step towards the description of the gas fluxes under anodes and of their effects on the stability.

We have developed a theory and a computational model for the detachment of bubbles under rigid surfaces. We view the detachment as a critical phenomenon, taking place on a molecular scale. Within the continuum approach, it manifests itself by the discontinuity of the derivative of the free energy. The free energy corresponds to a bubble attached to a rigid surface and the derivative is taken with respect to an order parameter. As such, the detachment can be viewed as the first order phase transition. We have formulated the conditions at which the detachment occurs as a constrained optimization problem. These conditions are determined by the critical pressure inside the bubble and by the unique width of the liquid layer separating the bubble from the surface. In order to determine these two parameters, we consider solutions of the Euler-Langrage equations corresponding to the free energy of the gas enclosed in a bubble. We parametrize these solutions by the pressure and the layer width. The parameters are then computed by requiring the conservation of the mass density and of the free energy before and after the detachment. The computations we made show extreme accuracy and correlation with the experimental results.

The results are published in [43], [42].

Personnel: P. Klouček (Assistant Prof., CAAM), M. V. Romerio (Research Professor, Dept. of Mathematics, EPF Lausanne, Switzerland).

Figure 5: Plot of the two and three dimensional surfaces of the bubbles having the same free energy and mass density at the critical point. The results were obtained by a constrained nonlinear optimization problem in [43].

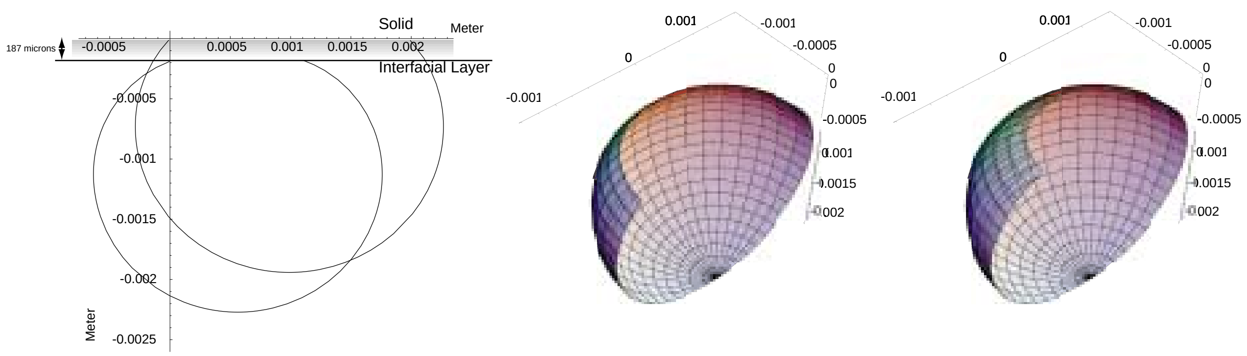

2.6 Computational Biology and X-Ray Crystallography

This research focuses on solving the so-called molecular replacement problem in X-ray Crystallography which is the primary technique for determining protein structures. The knowledges of such structures are crucial in understanding many biological processes, and in designing drugs for, say, cancers.

The problem can be formulated as a global optimization problem with a small number of variables but with highly expensive and difficult objective functions. The availability of MAPY has greatly sped up the turn-around time for our computational experiments, enabling us to run some of the experiments involving data sets of sizes close to 10 Gigabytes, and to run multiple experiments in parallel. For instances, using one processor, MAPY is about five time faster in generating two-dimensional slices of the objective functions like the pair given in Figure 6 than a Sun UltraSparc2 workstation. These two-dimensional slices help us understand how different formulations can affect the level of difficulty in solving the problem. A large number of such slices were needed in the process. With eight processors used, our turn-around time on MAPY was about forty times faster than on a workstation. Without the computing power of MAPY, this project would have been impossible to carry out in a reasonable amount of time on ordinary workstations. Two research papers has been written so far in this project [59, 37]. A third paper [60] is on a different topic of computational biology. In all these works, MAPY has played a pivotal role.

Personnel: Y. Zhang (Associate Prof., CAAM), R. Tapia (Prof., CAAM), G. Phillips (Prof., now Univ. Wisconsin), Z. Wu (Research scientist, CAAM, now Assistant Prof. at Iowa State), D. Jamrog (Graduate student, CAAM).

2.7 Algorithms for Large-Scale Semidefinite Programs

This is a joint project between Zhang and his collaborators, Monteiro and Burer, at Georgia Tech. Semidefinite programming has been one of the most active research areas in optimization with its many newly found applications. However, the capacity of solving large-scale semidefinite programs (SDPs) has been severely limited by the extraordinarily high demand on computational resources. With the aid of MAPY, We were able to test and experiment on a novel algorithm for large-scale SDPs, and able to solve some large-size problems that were never solved before. Our software package, called BMZ, is currently the only program that is capable of effectively solving certain SDPs with over ten thousand constraints. This NSF grant had a direct and important impact on the research reported in two recent papers [5, 6]. The following software package has been developed on MAPY:

BMZ: An ANSI C package for solving a certain class of large-scale semidefinite programs via nonlinear transformation and a first-order log-barrier algorithm.

Personnel: Y. Zhang (Associate Prof., CAAM), R. Monteiro (Prof., Georgia Tech.), S. Burer (Graduate student, Georgia Tech.).

2.8 Continuous Optimization Techniques for Discrete Optimization

This project also involves the collaboration between Zhang and his co-workers at Georgia Tech. Two software packages have been developed and tested on MAPY, and another one is to be released soon. MAPY has again been providing the vital computing resource in this research project and contributed significantly to the success of the research, leading to three papers [7, 4, 62] so far. The following software packages have been developed on MAPY:

CirCut: A fast and scalable, Goemans-Williamson quality heuristics for the maximum cut problem, maximum bisection problem, and other graph partitioning problems, code written in Fortran 90, available from http://www.caam.rice.edu/∼zhang/circut/; and

Max-AO: High-quality heuristics for the maximum stable set, maximum clique, and related problems based on continuous optimization formulations, written in ANSI C, electronically available from http://www.math.gatech.edu/∼burer/Max-AO/.

Personnel: Y. Zhang (Associate Prof., CAAM), R. Monteiro (Prof., Georgia Tech.), S. Burer (Graduate student, Georgia Tech.), L. Gao (Graduate student, CAAM).

Figure 6: The colors in each picture represent the level sets of a slice of a function arising in the molecular replacement problem in X-ray Crystallography. In each case only two variables are varied and the remaining variables are kept fixed. Such pictures help us to understand how difficult it is to locate the global optima and how closely the global optima correlate to the true solutions. In the above case, the picture on the right indicates a smoother, thus better function than the one on the left. A large number of such slices need to be studied. Generating them is very computation-intensive. MAPY greatly speeds up the turn-around time in generating pictures like these.

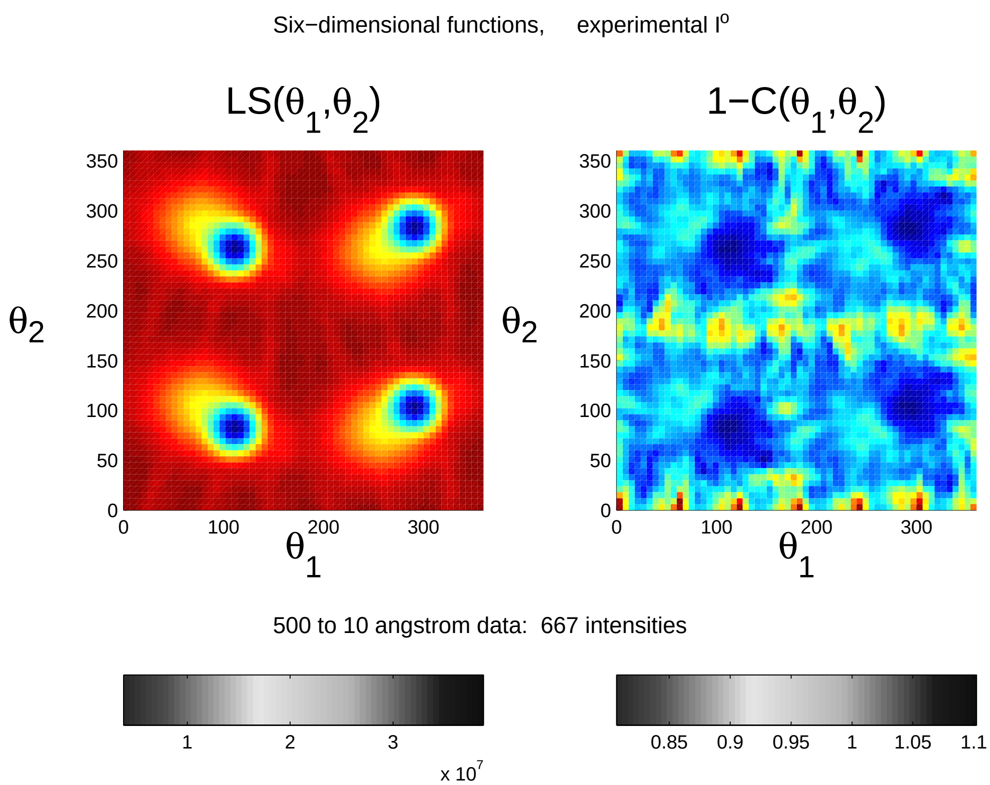

2.9 Turbulence Physics and Control

Flow control offers the potential for modifying complex fluid flows to achieve a variety of objectives including drag reduction, noise suppression, and heat transfer modification. Possible applications of this technology include reducing the drag on aircraft wings and increasing the mixing in internal combustion engines, both of which could lead to significant performance increases. Recent advances in micro-scale actuators and sensors, often called Microfabricated Electro-Mechanical Systems (MEMS), have made the application of flow control more practicable. However, efficient and reliable control strategies must be developed which can achieve the objective while minimizing the number of sensors and actuators.

Figure 7: Optimal control of turbulent channel flow: (a) Turbulent structures near the bottom wall of an uncontrolled flow, (b) Evolution of drag for uncontrolled and optimally controlled flows. As optimization interval ΔT increases, the initially turbulent flow becomes laminar resulting in a 60% reduction in drag.

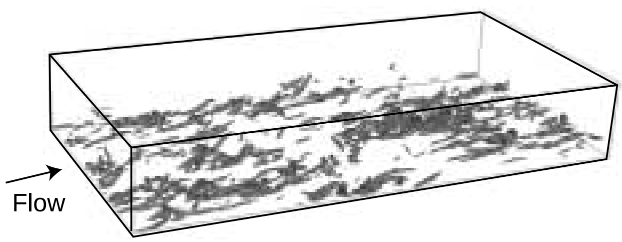

(a)

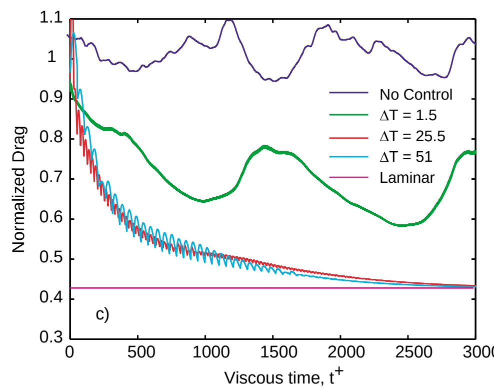

(b)

This project utilizes cutting-edge simulation tools to develop *optimal* control strategies to reduce the adverse effects of turbulence. Due to iterations inherent in nonlinear optimization, these simulations are often 10 to 20 times the expense of traditional single-run studies and therefore require extensive computational resources and the availability of the NSF-SCREMS supported SGI Origin 2000 has been critical to our success. To efficiently solve these problems, we have developed a unique parallel computational tool for simulating optimal control of turbulence using Large-Eddy Simulation (LES) based on the dynamic subgrid-scale model. Our work has pioneered the use of LES to investigate the control of turbulence which leads to dramatic improvements in speed over more traditional Direct Numerical Simulations (DNS). Our code is based on OpenMP and is highly tuned for the SGI Origin 2000 architecture.

This project has received funding from Texas Advanced Technology Program (TX-ATP) (3604-017) and NASA (NGT-1-52215) and extensions to this work are currently funded by NSF (DMS-0075731) in collaboration with M. Heinkenschloss (see Section 2.1) to explore novel approaches to solving large-scale optimization problems. This research is described in the papers [10, 12, 13, 16, 18, 52], the graduate theses [9, 39], and the conference abstracts [11, 17, 38, 50, 51]

Personnel: S.S. Collis (Assistant Prof., MEMS), R.D. Prabhu (Postdoc., MEMS), Y. Chang, S. Kellogg (Graduate students, MEMS), J. Templeton, A.Q. Nguyen (Undergraduate students, MEMS).

2.10 Computational Electrophysiology

The group composed of John Clark, Professor of Bioengineering, Steve Cox, Professor of Computational and Applied Mathematics, and Boyce Griffith, NYU Graduate Student have used MAPY as their basic engine in studies of conduction in both heart and nerve cells.

The heart study focuses on the role of calcium release at gap junctions in conduction block based on three dimensional extensions highly detailed models [54] of frog atrial cells. Figure 8 a. depicts a cross section of a bundle composed of seven strands of atrial fibers. A typical strand is 50 cells long. In order to draw biological conclusions each cell is partitioned into 20 sections and time is sampled every 100 microseconds over a simulation of 5 to 30 real seconds. Figure 8 b. is a snapshot of an action potential traveling to the right in a longitudinal slice.

The nerve study attempts to recover constitutive properties of dendritic neurons from single and dual potential recordings, see, e.g., [24]. Both the heart and nerve work came down to solving the equations of electrodiffusion (a hybrid ode/semilinear parabolic system) in realistic geometries and over long intervals of time. General C++ discretization, io, and graphics routines were developed as front and rear ends for PETSc [1] and ScaLAPACK [2].

Figure 8: a. Cross section of a bundle composed of seven strands of atrial fibers. b. Snapshot ofan action potential traveling to the right in a longitudinal slice.

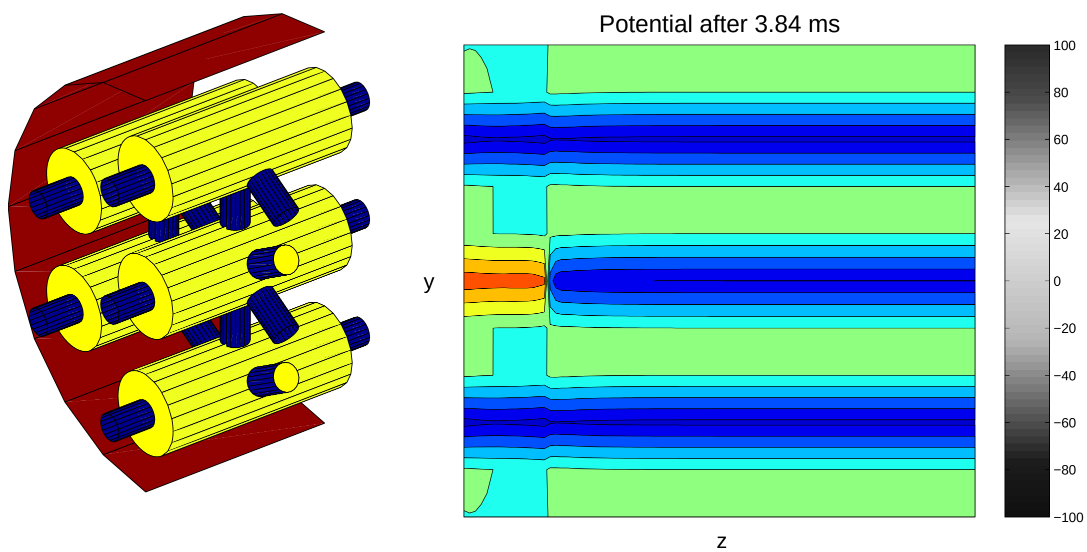

a b

Personnel: J. Clark (Prof. BIOE), S. Cox (Prof., CAAM), B. Griffith (Graduate Student, NYU, formerly undergrad CAAM).

2.11 Modeling the Physics of the Magnetospheric Substorm

We have implemented coupled fluid-equilibrium codes using the SGI Origin 2000 platform to calculate large scale structure and dynamics of the Earth's magnetosphere. Figure 9 shows the system generated currents that are formed in the system during simulation of a substorm. A substorm is an intrinsic instability of the magnetosphere. These are long calculations that are not possible without a machine such as the SGI Origin 2000. Its understanding has been a major unsolved problem in magnetosphere physics. Work done has made some important steps towards solving the substorm problem. This coupled code is the first attempt to model the full set of partial differential equations that describe the physics of the substorm process. This work has been done in collaboration with Dr. Joachim Birn of Los Alamos National Labs and Michael Hesse of the Goddard Space Flight center. This work is described in the papers [44, 45, 57, 56].
Personnel: P. Kloucek (Assistant Prof., CAAM) F. Toffoletto (Assistant Prof., PHYS), J. Wightman (Graduate Student, CAAM).

Figure 9: Computed field-aligned currents before (left) and after (right) Earth's magnetospheric substorm. This current has a dramatic effect on the motion of plasma in the Earth's ionosphere which can be compared to observations.

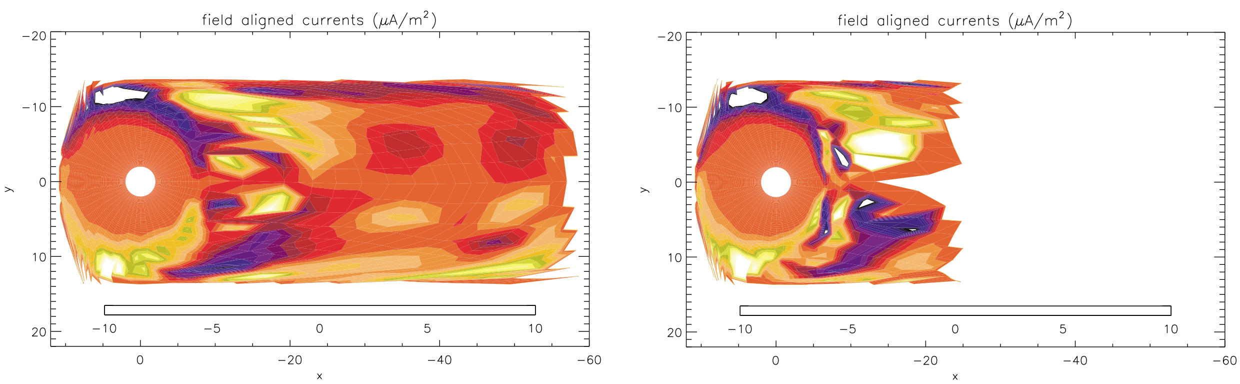

2.12 Subduction Zone Thermal Modeling Using Finite Elements

MAPY was used to compute mantle convection velocity in subduction zones. A subduction zone is where one oceanic plate subducts underneath another oceanic or continental plate. Many geology events happen in this zone, for example, earthquakes and volcanos. The main controlling factor for these geological events is the temperature. Our research objective is to model the temperature field for several subduction zones around the world. To model the subduction zone temperature, we need to know the mantle convection velocity field because this is one of most important parameter affecting the temperature. Previous subduction zone models used simplified slab (subducting cool plate) geometry. From earthquake observations, however, one can determine that a single angle slab geometry is not sufficient and that instead a curved slab should be used in most cases. We have developed a new algorithm using slab geometry fitting to earthquake hypocenter data. Our new curved slab model provides a more realistic treatment of subduction zones. This will affect our understanding of temperature control on the occurrence of intermediate depth earthquakes in a subduction zone. We have applied our model to 10 subduction zones around the world.

This work is described in [36].

Personnel: P. Kloucek (Assistant Prof., CAAM), William Leeman (Prof., ESCI), Virginia Sisson (Assistant Prof., ESCI), S. Huang (Graduate Student, ESCI).

Figure 10: Geographic map of Alaska subduction zone. The bight color dots represent earthquake distribution. The red arrow indicates the direction of movement of the Pacific plate. The white line is the location of our study.

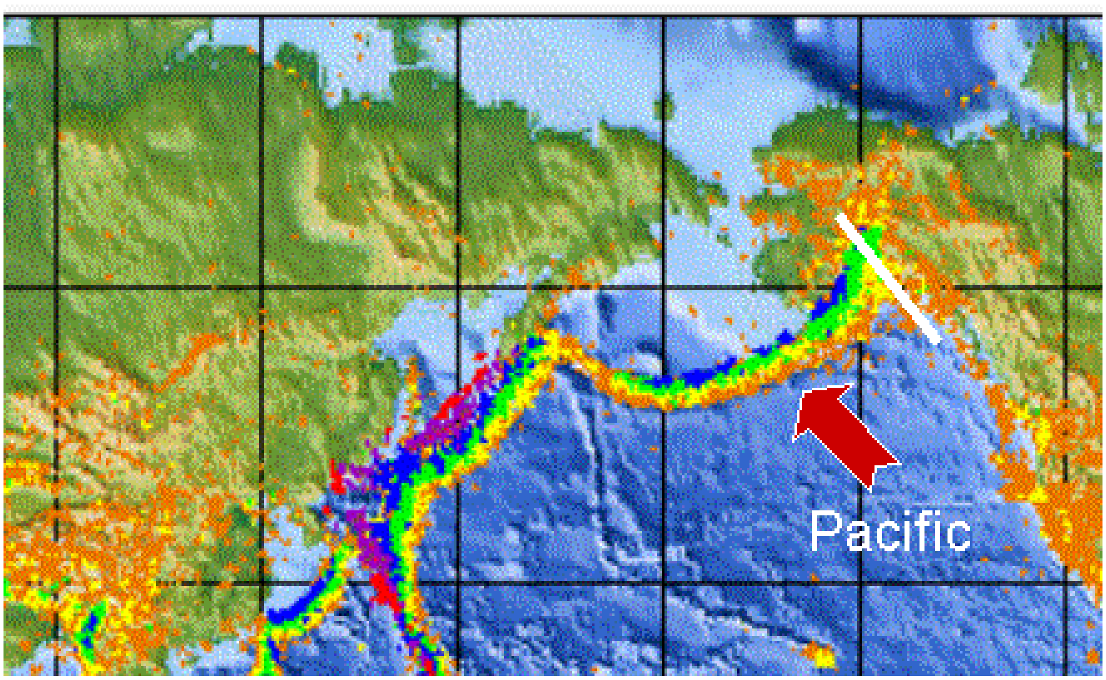

Figure 11: Temperature distribution of the Current Aleutiians subduction zone from our model computation. These results indicate 1) The coolest temperature is not on the slab surface, but inside the slab. 2) The wedge corner at about 100km has a large temperature gradient and is very hot; magma melting is possible in this corner.

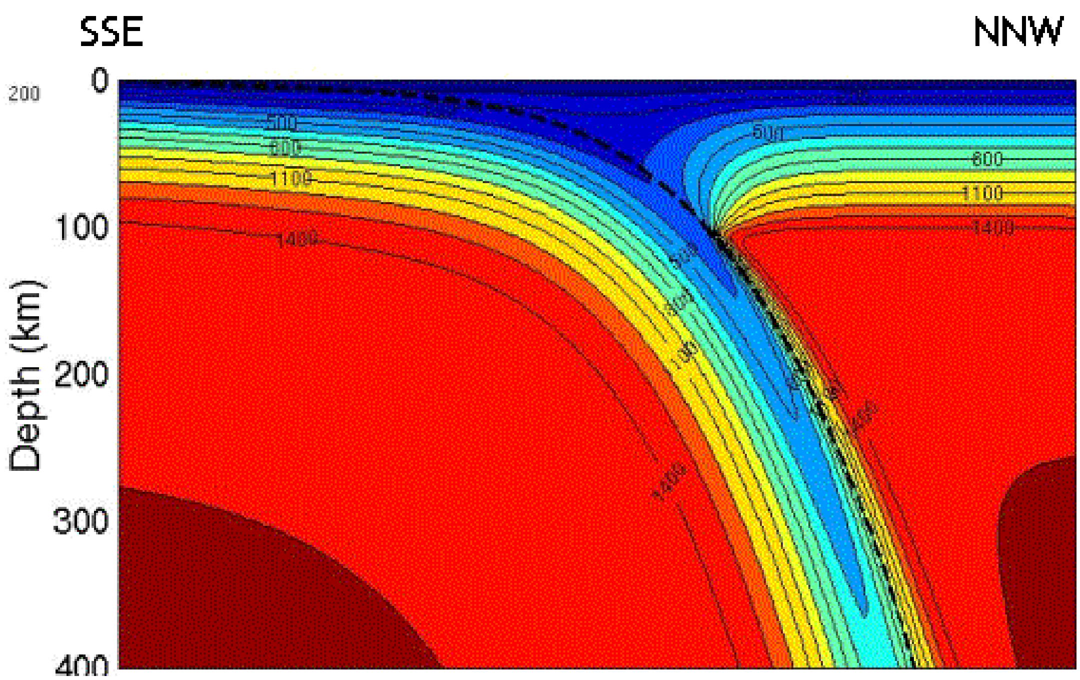

2.13 The Rice Automatic Differentiation Project

The purpose of the Rice Automatic Differentiation Project is to provide true analytic derivatives for mathematical models expressed as computer programs. The ability to obtain derivatives for computer models has proved useful in many diverse areas. A sample of user comments yields these representative applications:

- Sensitivity-based design
- Gradient-based optimization
- Solution of non-linear equations
- Stability studies for near-chaotic processes

The primary tool employed in the Rice Automatic Differentiation Project is a program augmentation tool called Adifor. MAPY has been, and continues to be, an important resource in Adifor development. The Rice Automatic Differentiation Project used MAPY in two ways. First, when extending Adifor 3.0 to differentiate MPI programs, we needed to test and verify our algorithms on many different MPI implementations. We used both the native MPI as well as the MPICH implementation on MAPY to develop and verify our MPI differentiation algorithms. Second, when applying our differentiation algorithms to a NASA CFD code, we were able to test and tune the differentiated code in it's native environment by using MAPY. Since our collaborators were using a (bigger) Origin 2000, we were able to collect realistic performance tuning data in a timely fashion by running our profilers on MAPY.

This work has led to the papers [8], [55].

Personnel: Michael Fagan (Research Scientist, CAAM), Alan Carle (Research Scientist, CAAM (on medical leave)).

2.14 Compiler, Library, and Tools Research

The SGI Origin 2000 provided an excellent environment for developing and testing compiler optimization techniques. The hardware performance counter support built in to the R12000 processor were crucial for understanding performance problems in applications, and for evaluating how well compiler techniques were able to ameliorate these problems.

The shared-memory hardware was vital for evaluating code produced by our data-parallel compiler, and was essential for the development of thread library technology.

The high-performance processors provided a capable cycle server for exploring large search spaces in the development of algorithms to compute multipartitionings.

2.14.1 Multipartitioning for Parallelizing Line Sweep Computations

Multipartitioning is a strategy for partitioning multi-dimensional arrays among a collection of processors so that line-sweep computations can be performed efficiently. The principal property of a multipartitioned array is that for a line sweep along any array dimension, all processors have the same number of tiles to compute at each step in the sweep. This property results in full, balanced parallelism. A secondary benefit of multipartitionings is that they induce only coarse-grain communication.

Previously, computing a d-dimensional multipartitioning required that $p^{\frac{1}{d-1}}$ be integral, where p is the number of processors. We developed an algorithm to compute a d-dimensional multipartitioning of an array of ρ dimensions for an arbitrary number of processors, for any d, $2 \leq d \leq \rho$. When using a multipartitioning to parallelize a line sweep computation, the best partitioning is the one that exploits all of the processors and has the smallest communication volume. To compute the best multipartitioning of a ρ-dimensional array, a cost model was developed for selecting d, the dimensionality of the best partitioning, and the number of cuts along each partitioned dimension. In practice, our technique will choose a 3-dimensional multipartitioning for a 3-dimensional line-sweep computation, except when p is a prime; previously, a 3-dimensional multipartitioning could be applied only when $\sqrt{p}$ is integral.

2.14.2 Data-Parallel Compilation

Strategies for partitioning an application's data determine both the range of suitable parallelizations and their potential efficiency. For multi-directional line-sweep computations, *multipartitioned* data distributions offer better parallel efficiency and scalability than block unipartitionings. We developed extensions to the program analysis and code generation techniques in Rice dHPF compiler for High Performance Fortran that enable it to support multipartitioned data distributions. In addition, optimizations were developed to enable dHPF to generate efficient multipartitioned code. Experiments applying dHPF's compiler support for multipartitioning to parallelize two serial versions of NASA's computational fluid dynamics application benchmarks (NAS SP and BT) show that the performance of the multipartitioned code generated by dHPF is within a few percent of that of hand-coded parallel implementations using multipartitioning.

2.14.3 Improving Memory Hierarchy Performance

Reusing data in cache is critical to achieving high performance on modern machines, because it reduces the impact of the latency and bandwidth limitations of direct memory access. To date, most studies of software memory hierarchy management have focused on the latency problem in loops. However, today's machines are increasingly limited by the insufficient memory bandwidth—latency-oriented techniques are inadequate because they do not seek to minimize the total amount of memory transfer of the whole program. To address the bandwidth limitation, this research explores the potential for global cache reuse—that is, reusing data across whole program and over the entire data. A two-step strategy for improving bandwidth utilization is proposed. The first step fuses computations on the same data to enable the caching of repeated accesses. The second step groups data used by the same computation to effect contiguous access to memory. While the first step reduces the frequency of memory access, the second step improves its efficiency. Experiments demonstrate the effectiveness of this strategy. This approach is suitable for automating in production compilers.

The performance of irregular applications on modern computer systems is particularly hurt by the wide gap between CPU and memory speeds because these applications typically underutilize multi-level memory hierarchies, which help hide this gap. Strategies for applying data and computation reorderings to improve memory hierarchy utilization for irregular applications were studied. This research evaluated the impact of reordering on data reuse at different levels in the memory hierarchy. This research focused on studying coordinated data and computation reordering based on space-filling curves and introduced a new architecture-independent multi-level blocking strategy for irregular applications. For two particle codes studied, the most effective reorderings reduced overall execution time by a factor of two and four, respectively. Preliminary experience with a scatter benchmark derived from a large unstructured mesh application showed that careful data and computation ordering reduced primary cache misses by a factor of two compared to a random ordering.

Recently, there have been several experimental and theoretical results showing significant performance benefits of recursive algorithms on both multi-level memory hierarchies and on shared-memory systems. In particular, such algorithms have the data reuse characteristics of a blocked algorithm that is simultaneously blocked at many different levels. Most existing applications, however, are written using ordinary loops. A new compiler transformation was developed that can be used to convert loop nests into recursive form automatically. This algorithm is fast and effective, handling loop nests with arbitrary nesting and control flow. The transformation achieves substantial performance improvements for several linear algebra codes even on a current system with a two level cache hierarchy. As a side-effect of this work, an improved algorithm for transitive dependence analysis (a powerful technique used in the recursion transformation and other loop transformations) was developed that is much faster than the best previously known algorithm in practice.

2.14.4 An API for POSIX Threads in Fortran Programs

Pthreads is the library of POSIX standard functions for concurrent, multithreaded programming. The POSIX standard only defines an application programming interface (API) to the C programming language, not to Fortran. Many scientific and engineering applications are written in Fortran. Also, many of these applications exhibit functional, or task-level, concurrency. They would benefit from multithreading, especially on symmetric multiprocessors (SMP). An interface to that part of the Pthreads library that is compatible with standard Fortran was developed. This interface consists of two primary source files: a Fortran module and a collection of C wrappers to Pthreads functions. The Fortran module defines the data structures, interface and initialization routines used to manage threads. The stability and portability of the Fortran API to Pthreads is demonstrated using common mathematical computations on three different systems.

2.14.5 Tools for Understanding Memory Hierarchy Performance

Application performance tuning is a complex process that requires assembling various types of information and correlating it with source code to pinpoint the causes of performance bottlenecks. Existing performance tools don't adequately support this process in one or more dimensions. To address this problem, two tools were developed to provide better diagnosis of performance problems. *MHsim* is a memory hierarchy simulator that produces source-level information not otherwise available about memory hierarchy utilization and the causes of cache conflicts. *HPCView* is a tool that combines data from arbitrary sets of instrumentation sources and correlates it with program source code. Both tools report their results in scope-hierarchy views of the corresponding source code and produce their output as HTML databases that can be analyzed portably and collaboratively using a commodity browser. In addition to receiving daily use within our group, these tools are being used successfully by several code development teams in DoD and DoE laboratories.

2.14.6 Personnel

John Mellor-Crummey, Senior Faculty Fellow
Robert Fowler, Research Scientist
Richard Hanson, Research Scientist
Vikram Adve, Research Scientist
David Whalley, Visiting Scholar
Alain Darte, Visiting Scholar
Chen Ding, Graduate Student
Yuan Zhao, Graduate Student
Gabriel Marin, Graduate Student

Qing Yi, Graduate Student
Daniel Chavarria-Miranda, Graduate Student
Trushar Sarang, Undergraduate
Nathan Tallent, Research Programmer
Monika Mevencamp, Research Programmer
Fengmei Zhao, Research Programmer
Li-Wei Peng, Research Assistant
Jie Song, Research Assistant

2.15 Improving Fine-Grained Irregular Shared-Memory Benchmarks by Data Reordering

We have demonstrated that data reordering can substantially improve the performance of fine-grained irregular shared-memory benchmarks, on both hardware and software shared-memory systems. In particular, we have evaluated two distinct data reordering techniques that seek to co-locate in memory objects that are in close proximity in the physical system modeled by the computation. The effects of these techniques are increased spatial locality and reduced false sharing.

We have evaluated the effectiveness of the data reordering techniques on a set of five irregular applications from SPLASH-2 and Chaos. Both techniques were implemented in a small library, allowing us to enable them in an application by adding less than 10 lines of code. Our results on one hardware and two software shared-memory systems show that, with data reordering during initialization, the performance of these applications is improved by 12%–99% on the SGI Origin 2000, 30%–366% on TreadMarks, and 14%–269% on HLRC.

In our study, the SGI Origin 2000, being the newest generation of hardware shared memory machine, served as our only hardware DSM platform. Thus it is crucial in demonstrating that our techniques are applicable on both software DSMs and hardware DSMs.

Details of this study are documented in [35].

Personnel: Alan Cox (Associate Prof., CS), Charlie Hu (Research Scientist, CS), Willy Zwaenepoel (Prof., CS).

3 Personnel

Faculty, Faculty Fellows, and Research Scientist

Vikram Adve, Research Scientist, CS
Alan Carle, Research Scientist, CAAM
John Clark, Prof. BIOE
Scott Collis, Assistant Prof, MEMS
Alan Cox, Associate Prof., CS
Steven Cox, Prof., CAAM
Michael Fagan, Research Scientist, CAAM
Robert Fowler, Research Scientist, CS
Richard Hanson, Research Scientist, CS
Matthias Heinkenschloss, Associate Prof., CAAM
Charlie Hu, Research Scientist, CS
Petr Kloucek, Assistant Prof., CAAM
William Leeman, Prof., ESCI
John Mellor-Crummey, Senior Faculty Fellow, CS
George Phillips, Prof. BIOE, now Univ. Wisconsin),
Virginia Sisson, Assistant Prof., ESCI
Richard Tapia, Prof., CAAM
Frank Toffoletto, Assistant Prof., PHYS
Zhijun Wu, Research scientist, CAAM, now Assistant Prof. at Iowa State
Yin Zhang, Associate Prof., CAAM
Willy Zwaenepoel, Prof., CS

Visiting Faculty and Scholar

Alain Darte
Renato Monteiro, Prof., Georgia Tech.
Michael Ulbrich, Assistant Prof., TU Munich, Germany
Stefen. Ulbrich, Assistant Prof., TU Munich, Germany
David Whalley

Postdocs

K. Ghayour, CAAM & MEMS
R.D. Prabhu, MEMS

Graduate Students

Sam Burer, Graduate student, Georgia Tech.
Yong Chang, MEMS
Daniel Chavarria-Miranda, CS
Guaquan Chen, MEMS
Agata Comas, CAAM

Chen Ding, CS
Liyan Gao, CAAM
Boyce Griffith, Graduate Student at NYU (formerly undergrad CAAM)
Diane Jamrog, CAAM
Steve Kellogg, MEMS
Saijin Huang, ESCI
Gabriel Marin, CS
Luis Melara, CAAM
Srinivas Ramakrishnan, MEMS
Dan Reynolds, CAAM
Zachary Smith, MEMS
Jennifer Wightman, CAAM
Qing Yi, CS
Yuan Zhao, CS

Undergraduate Students

N. Bou-Rabee, CAAM & MEMS
A.Q. Nguyen, MEMS
R. Purple, MEMS
E. Rareshide, MEMS
Trushar Sarang, CS
J. Templeton, MEMS

Research Programmers and Research Assistants

Nathan Tallent, CS, Research Programmer
Monika Mevencamp, CS, Research Programmer
Fengmei Zhao, CS, Research Programmer
Li-Wei Peng, CS, Research Assistant
Jie Song, CS, Research Assistant

References

[1] S. Balay, K. Buschelman, W. D. Gropp, D. Kaushik, L. Curfman McInnes, and B. F. Smith. PETSc home page, 2001. http://www-fp.mcs.anl.gov/petsc.

[2] L. S. Blackford, et. al. *ScaLAPACK Users' Guide*. SIAM Press, Philadelphia, 1997.

[3] B. Broom, D. Chavarría-Miranda, G. Jin, R. Fowler, K. Kennedy, and J. Mellor-Crummey. Overpartitioning with the Rice dHPF compiler. In *Proceedings of the 4th Annual HPF User Group meeting*, Tokyo, Japan, Oct. 2000.

[4] S. Burer, R. Monteiro, and Y. Zhang. Maximum stable set formulations and heuristics based on continuous optimization. Technical Report TR00-34, Department of Computational and Applied Mathematics, Rice University, Houston, TX 77005, 2000.

[5] S. Burer, R. Monteiro, and Y. Zhang. A computational study of a gradient-based log-barrier algorithm for a class of large-scale SDPs. Technical Report TR01-11, Department of Computational and Applied Mathematics, Rice University, Houston, TX 77005, 2001.

[6] S. Burer, R. Monteiro, and Y. Zhang. Interior-point algorithms for semidefinite programming based on a nonlinear programming formulation. *Computational Optimization and Applications*, To appear, 2001.

[7] S. Burer, R. Monteiro, and Y. Zhang. Rank-two relaxation heuristics for max-cut and other binary quadratic programs. *SIAM journal on Optimization*, To appear, 2001.

[8] A. Carle and M. Fagan. Automatically differentiating mpi-1: The complete story. In *3rd International Conference/Workshop on Automatic Differentiation*. Springer-Verlag, June 2000. to be published December 2001.

[9] Y. Chang. *Reduced Order Methods for Optimal Control of Turbulence*. PhD thesis, Rice University, 2000.

[10] Y. Chang and S. S. Collis. Active control of turbulent channel flows based on large eddy simulation. In *Proceedings of the 1999 ASME/JSME Joint Fluids Engineering Conference, FEDSM99-6929*, 1999.

[11] Y. Chang and S. S. Collis. Hybrid DNS/LES methods for turbulence control. *Bull. Am. Phys. Soc.*, 44(8):41, 1999.

[12] Y. Chang and S. S. Collis. Approximate models for optimal control of turbulence. *under preparation*, 2001.

[13] Y. Chang and S. S. Collis. Viscous effects in turbulence control. *under preparation*, 2001.

[14] D. Chavarría-Miranda and J. Mellor-Crummey. Towards compiler support for scalable parallelism. In *Proceedings of the Fifth Workshop on Languages, Compilers, and Runtime Systems for Scalable Computers*, Lecture Notes in Computer Science 1915, pages 272–284, Rochester, NY, May 2000. Springer-Verlag.

[15] D. Chavarría-Miranda, J. Mellor-Crummey, and T. Sarang. Data-parallel compiler support for multipartitioning. In *European Conference on Parallel Computing (Euro-Par)*, Manchester, United Kingdom, Aug. 2001.

[16] S. S. Collis and Y. Chang. On the use of LES with a dynamic subgrid scale model for the optimal control of wall bounded turbulence. In D. Knight and L. Sakell, editors, *Recent Advances in DNS and LES*, pages 99–110. Kluwer Academic Publishers, 1999.

[17] S. S. Collis and Y. Chang. Viscous effects in turbulence control simulations. *Bull. Am. Phys. Soc.*, 45(9):156, 2000.

[18] S. S. Collis, Y. Chang, S. Kellogg, and R. D. Prabhu. Large eddy simulation and turbulence control. In *AIAA paper 2000-2564*, 2000.

[19] S. S. Collis, A. Comas, and M. Heinkenschloss. Time–domain decomposition preconditioning of nonlinear optimal control problems. Technical report, Department of Computational and Applied Mathematics, Rice University, Houston, TX 77005–1892, 2001. in preparation.

[20] S. S. Collis, K. Ghayour, M. Heinkenschloss, M. Ulbrich, and S. Ulbrich. Numerical solution of optimal control problems governed by the compressible Navier–Stokes equations. In *Proceedings of the International Conference on Optimal Control of Complex Structures*. Birkhäuser Verlag, 2001.

[21] S. S. Collis, K. Ghayour, M. Heinkenschloss, M. Ulbrich, and S. Ulbrich. Optimal control of unsteady compressible viscous flows. Technical report, Department of Computational and Applied Mathematics, Rice University, Houston, TX 77005–1892, 2001. submitted for publication.

[22] S. S. Collis, K. Ghayour, M. Heinkenschloss, M. Ulbrich, and S. Ulbrich. Towards adjoint–based methods for aeroacoustic control. In *39th Aerospace Science Meeting & Exhibit, January 8–11, 2001, Reno, Nevada, AIAA Paper 2001–0821*, 2001.

[23] D. Cox, P. Klouček, and D. Reynolds. The non-local relaxation of non-attainable differential inclusions using subgrid projection method: One-dimensional theory and calculations. *submitted to: J. Nonlin. Sci.*, 2001. Technical Report, *École Polytechnbique Fédérale de Lausanne, Département de Mathématiques, Analyse and Analyse Numérique, #10, July 2001.*

[24] S. J. Cox and B. E. Griffith. Recovering quasi–active properties of dendrites from dual potential recordings. *J. of Computational Neuroscience*. to appear.

[25] A. Darte, J. Mellor-Crummey, R. Fowler, and D. Chavarría-Miranda. On efficient parallelization of line-sweep computations. In *9th Workshop on Compilers for Parallel Computers*, Edinburgh, Scotland, June 2001.

[26] C. Ding. *Improving Effective Bandwidth through Compiler Enhancement of Global and Dynamic Cache Reuse*. PhD thesis, Dept. of Computer Science, Rice University, Jan. 2000.

[27] C. Ding and K. Kennedy. Bandwidth-based performance tuning and prediction. In *Proceedings of IASTED International Conference on Parallel Computing and Distributed Systems*, November 1999.

[28] C. Ding and K. Kennedy. Improving cache performance in dynamic applications through data and computation reorganization at run time. In *Proceedings of the SIGPLAN '99 Conference on Programming Language Design and Implementation*, Atlanta, GA, May 1999.

[29] C. Ding and K. Kennedy. Inter-array data regrouping. In *Proceedings of The 12th International Workshop on Languages and Compilers for Parallel Computing*, August 1999.

[30] C. Ding and K. Kennedy. Memory bandwidth bottleneck and its amelioration by a compiler. In *Proceedings of the 2000 International Parallel and Distributed Processing Symposium*, Cancun, Mexico, May 2000.

[31] C. Ding and K. Kennedy. Improving effective bandwidth through compiler enhancement of global cache reuse. In *International Parallel and Distributed Processing Symposium*, San Francisco, CA, Apr. 2001. (Best Paper Award.).

[32] R. J. Hanson, C. R. Breshears, and H. A. Gabb. Algorithm xxx: A Fortran interface to POSIX threads. Submitted to *ACM Transactions on Mathematical Software* July, 2000.

[33] R. J. Hanson, C. R. Breshears, and H. A. Gabb. Using a Fortran interface to POSIX threads. In *The Architecture of Scientific Software*, The Netherlands, 2001. Kluwer Press, Inc.

[34] M. Heinkenschloss. Time–domain decomposition iterative methods for the solution of distributed linear quadratic optimal control problems. Technical Report TR00–31, Department of Computational and Applied Mathematics, Rice University, Houston, TX 77005–1892, 2000.

[35] Y. C. Hu, A. Cox, and W. Zwaenepoel. Improving fine-grained irregular shared-memory benchmarks by data reordering. In *Proceedings of SC 2000, Dallas, TX, November 2000*, 2000.

[36] S. Huang, P. W. Leeman, B. V. Sisson, and P. Kloucek. A numerical model for subduction zone thermal structure. *EOS Transactions, American Geophysical Union, 1999 Fall Meeting*, 80, 1999.

[37] D. Jamrog, G. Phillips, R. Tapia, and Y. Zhang. The effect of the separation of variables on the molecular replacement method. Technical Report TR99-18, Department of Computational and Applied Mathematics, Rice University, Houston, TX 77005, 1999.

[38] S. Kellogg and S. S. Collis. Model-based predictive control of turbulent channel flow. *Bull. Am. Phys. Soc.*, 44(8):41, 1999.

[39] S. M. Kellogg. Immersed boundary methods with applications to flow control. Master's thesis, Rice University, 2000.

[40] P. Klouček and L. A. Melara. The computational modeling of internal surfaces in crystals. *submitted to: J. Comp. Phys.*, 2001. Technical Report, *École Polytechnbique Fédérale de Lausanne, Département de Mathématiques, Analyse and Analyse Numérique, #12, August 2001.*

[41] P. Klouček, D. Reynolds, and T. I. Seidman. On the robust control of incoherent phase transformations. *Manuscript*, 2001.

[42] P. Klouček and M. V. Romerio. The detachment of bubbles under a porous rigid surface during aluminium electrolysis. Technical Report 11, École Polytechnbique Fédérale de Lausanne, Département de Mathématiques, Analyse et Analyse Numérique, 1015 Ecublens, Lausanne, Switzerland, July 2001. submitted to: Math. Models & Methods in Applied Sciences.

[43] P. Klouček and M. V. Romerio. The first order phase transition of bubbles at solid-liquid interface. Technical Report 1, École Polytechnbique Fédérale de Lausanne, Département de Mathématiques, Analyse et Analyse Numérique, 1015 Ecublens, Lausanne, Switzerland, January 2001. submitted to: Arch. Rat. Mech. Anal.

[44] P. Klouček and F. R. Toffoletto. The three dimensional non-conforming finite element solution of the chapman-ferraro problem. *J. Comp. Physics*, 150:549–560, 1999.

[45] P. Klouček and F. R. Toffoletto. Rough nonconforming finite element solver for the magnetosphere-ionosphere coupling equation. *ISSS-6 Proceedings*, 2001.

[46] J. Mellor-Crummey, R. Fowler, and D. Whalley. On providing useful information for analyzing and tuning applications. In *Joint International Conference on Measurement & Modeling of Computer Systems*, Cambridge, MA, June 2001.

[47] J. Mellor-Crummey, R. Fowler, and D. Whalley. Tools for application-oriented performance tuning. In *Proceedings of the 15th ACM International Conference on Supercomputing*, Sorrento, Italy, June 2001.

[48] J. Mellor-Crummey, D. Whalley, and K. Kennedy. Improving memory hierarchy performance for irregular applications. In *Proceedings of the 13th ACM International Conference on Supercomputing*, pages 425–433, Rhodes, Greece, June 1999.

[49] J. Mellor-Crummey, D. Whalley, and K. Kennedy. Improving memory hierarchy performance for irregular applications using data and computation reorderings. *International Journal of Parallel Programming*, 29(3), June 2001.

[50] R. D. Prabhu and S. S. Collis. Proper orthogonal decomposition of turbulent channel flow with and without control. *Bull. Am. Phys. Soc.*, 44(8):180, 1999.

[51] R. D. Prabhu and S. S. Collis. Assessing reduced order models for flow-control using POD. *Bull. Am. Phys. Soc.*, 45(9):156, 2000.

[52] R. D. Prabhu, S. S. Collis, and Y. Chang. The influence of control on proper orthogonal decomposition of wall-bounded turbulent flows. *Phys. Fluids*, 13(2):520–537, 2001.

[53] D. Reynolds. Vibration damping using martenstic phase transformation. *Ph. D. Dissertation*, December 2002. Department of Computational and Applied Mathematics, 6100 main Street, Houston, TX 77005.

[54] M. Shumaker, J. W. Clark, and W. R. Giles. Simulations of passive properties and action potential conduction in an idealized bullfrog atrial trabeculum. *Math. Biosci.*, 116:127–167, 1993.

[55] P. Sundaram, S. Agrawal, J. Hager, A.Carle, and M. Fagan. Viscous design optimization using adjifor - an hpccp perspective. In *HPCCP/CAS Workshop*. NASA Ames Research Center, February 2000.

[56] F. R. Toffoletto, J. Birn, M. Hesse, R. W. Spiro, and R. Wolf. Modeling inner magnetospheric electrodynamics. In P. Song, G. L. Siscoe, and H. Singer, editors, *Proceedings of the Chapman Conference on Space Weather*. Am. Geophys. Un., Washington, D. C., 2000. to appear.

[57] F. R. Toffoletto, R. W. Spiro, R. Wolf, M. Birn, and M. Hesse. Computer experiments on substorm growth and expansion. In A. Wilson, editor, *Proceedings of the Fifth International Conference on Substorms. St. Petersburg, Russia, 16-20 May 2000 (ESA SP-443)*, pages 351–355. ESA, Noordwijk, 2000.

[58] M. Ulbrich. Nonsmooth Newton-like methods for variational inequalities and constrained optimization problems in function spaces. Habilitation Thesis, Fakultät für Mathematik, Technische Universität München, 80290 München, Germany, 2001.

[59] Z. Wu, G. Phillip, R. Tapia, and Y. Zhang. A fast newton's method for entropy maximization in phase determination. *SIAM Review*, To appear, 2001.

[60] Z. Wu and Y. Zhang. Solving the double digestion problem as a mixed-integer linear program. Technical Report TR01-12, Department of Computational and Applied Mathematics, Rice University, Houston, TX 77005, 2001.

[61] Q. Yi, V. Adve, and K. Kennedy. Transforming loops to recursion for multi-level memory hierarchies. In *Proceedings of the SIGPLAN '00 Conference on Programming Language Design and Implementation*, Vancouver, British Colombia, June 2000.

[62] Y. Zhang and L. Gao. On numerical solution of the maximum volume ellipsoid problem. Technical Report TR01-15, Department of Computational and Applied Mathematics, Rice University, Houston, TX 77005, 2001.