RIACS
Annual Report
1994

Contract NAS 2-13721
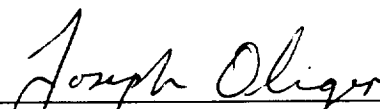
RIACS Principal Investigator
Joseph Oliger

NASA Technical Monitor
David Cooper

COTR
Ron Deiss

Joseph Oliger, Director

Table of Contents

I. INTRODUCTION

Joseph Oliger, Director

The Research Institute for Advanced Computer Science (RIACS) was established by the Universities Space Research Association (USRA) at the NASA Ames Research Center (ARC) on June 6, 1983. RIACS is privately operated by USRA, a consortium of universities with research programs in the aerospace sciences, under contract with NASA.

The primary mission of RIACS is to provide research and expertise in computer science and scientific computing to support the scientific missions of NASA ARC. The research carried out at RIACS must change its emphasis from year to year in response to NASA ARC's changing needs and technological opportunities. A flexible scientific staff is provided through a university faculty visitor program, a post doctoral program, and a student visitor program. Not only does this provide appropriate expertise but it also introduces scientists outside of NASA to NASA problems. A small group of core RIACS staff provides continuity and interacts with an ARC technical monitor and scientific advisory group to determine the RIACS mission. RIACS activities are reviewed and monitored by a USRA advisory council and ARC technical monitor.

Research at RIACS is currently being done in the following areas:
Parallel Computing
Advanced Methods for Scientific Computing
High Performance Networks
Learning Systems

Parallel compiler techniques, adaptive numerical methods for flows in complicated geometries and optimization have been identified as important problems to investigate for ARC's involvement in the Computational Grand Challenges of the next decade.

During the past year Professor Bertil Gustafsson of the University of Uppsala, Professor Antony Jameson of Princeton University, and Professor Wai-Pai Tang of the University of Waterloo have been visiting RIACS.

We had 8 visiting graduate students this year.

RIACS technical reports are usually preprints of manuscripts that have been submitted to research journals or conference proceedings. A list of these reports for the period January 1, 1994 through December 31, 1994 is in the Reports and Abstracts section of this report.

II. RESEARCH IN PROGRESS

A. PARALLEL COMPUTING

HIGH PERFORMANCE FORTRAN
Robert S. Schreiber

The HPF effort during 1994 concentrated on defining an agenda for an HPF 2 language specification, which would occur in 1995. There is considerable user interest in extending HPF so that it can carry more scientific parallel applications. To that end, a number of potential generalizations of HPF features and new HPF features have been discussed. A working document, the "HPF2 Scope of Activities," coauthored by Schreiber, discusses these potential new features. The areas covered include irregular and more powerful data mapping, task parallelism, computation/iteration mapping, and input/output.

An effort to correct, clarify, and interpret the HPF 1.0 standard was also pursued, and the result is an updated HPF standard, version 1.1, published in October.

RESEARCH IN HIGHLY PARALLEL MATRIX COMPUTATION
Robert S. Schreiber

Joint work with Petter Bjorstad (Univ. of Bergen) on routing in grid architectures for unstructured mesh applications was finished and the results presented at the SHPCC conference.

With Ed Rothberg (Intel Scientific Computers) an improved implementation of distributed memory sparse Cholesky was developed. Rothberg had discovered that block methods with two-dimensional mapping suffer significant loss of efficiency due to imperfectly balanced computational load. By altering the mapping, and in particular allowing independent mapping of block rows and block columns, Rothberg and Schreiber have improved the performance of these methods by as much as 20 percent. They have recently also improved the heuristic task scheduler and obtained a further improvement of about 10 percent. These results were presented at the Supercomputing 94 conference.

Fred Chong (MIT) undertook an implementation of the highly parallel algorithms of Alvarado, Pothen, and Schreiber for repeated solution of a sparse triangular linear system while visiting RIACS. He used the DSC algorithms of Gerasoulis and Yang for mapping and scheduling the computation graph produced by this algorithm, then executed it on the CM-5 using a lightweight tasking package from MIT. The encouraging results showed that the partitioned inverse method of Alvarado et al. is far better for highly parallel machines than the standard substitution method.

Finally, work with Xiaoye Li of Berkeley and John Gilbert of Xerox has resulted in a new formulation of sequential sparse LU factorization with partial pivoting that is better adapted to memory hierarchies that previous codes. Our experiments show the new software to be the fastest available. Li is now pursuing a parallel implementation.

5

COMPILING ARRAY PARALLELISM FOR DISTRIBUTED-MEMORY MACHINES

Robert S. Schreiber, Thomas J. Sheffler
Siddhartha Chatterjee, Leonid Oliker (U. of Colorado at Boulder)
John Gilbert (Xerox PARC)
Patty Hough (Cornell University; summer intern)

The goal of this project is to develop compiler algorithms to automate the task of data layout when compiling array-parallel languages such as Fortran 90 on distributed-memory parallel computers. Data layout is accomplished by alignment of arrays to a template and distribution of the template to the processors. The past year has seen much progress in the theory of data layout. We now have a good model of the distribution problem and a plausible proposed heuristic algorithm for it. We have also gained understanding of the data layout requirements of real programs by implementing our proposed algorithms and performing experiments to evaluate their effectiveness.

Our theory for the alignment phase of data layout was largely completed last year. This year, we implemented the proposed algorithms. Using a large body of test programs, we have shown that our linear programming based algorithms for alignment are robust and correct. We can successfully generate programs with nontrivial mobile offset alignments, for example.

The distribution phase of data layout is a difficult problem requiring a detailed model of the target parallel computer. (Alignment can be effectively handled using abstract models). One of our projects for the year was to develop a model of data parallel communication that is appropriate for distribution analysis. With Patty Hough, a visiting summer intern, we tuned our model to one particular parallel computer: the CM-5. We plan to gather performance data for other machines in the future.

Early this year, we gained enough confidence about our proposed distribution algorithm to begin to implement it. This allowed us to experiment with fine points of the algorithm at an early stage. The resulting algorithm uses a divide-and-conquer approach to successively dissect a data-parallel program into small pieces that are separately analyzed and then re-combined. Most of our time has been spent experimentally evaluating partioning algorithms that perform the dissection, and we have developed a partitioner that works well in practice. We also found that a set of graph-contraction tools that we developed significantly simplifies the problem. We are now at a stage where we are evaluating the data layouts generated by our suite of analysis tools.

B. ADVANCED METHODS FOR SCIENTIFIC COMPUTING

DYNAMIC ADAPTATION AND PARALLEL IMPLEMENTATION OF 3-D UNSTRUCTURED GRIDS TO HELICOPTER AERODYNAMICS AND ACOUSTICS
Rupak Biswas

Work on this project, in collaboration with Roger Strawn, of NASA Code AAC/YF, has been completed with the development of an option to control the quality of the resulting mesh after several levels of refinement. The method (code is called 3D_TAG) has been successfully implemented on the C-90 and is currently being used to compute large problems in rotor aerodynamics and acoustics. It is an efficient solution-adaptive procedure that anisotropically adds and removes mesh points in order to optimize their distribution. so that the flowfield is accurately modeled with a minimum of computational resources. An innovative data structure, that uses a combination of dynamically-allocated arrays and linked lists, allows points to be rapidly added and/or deleted.

One problem with the anisotropic subdivision of tetrahedral elements is that repeated refinement can lead to poor mesh quality. Poor mesh quality is defined as a grid deficiency that leads to inaccurate flowfield solutions. Our algorithm controls the quality of the mesh by never further subdividing anisotropically-refined elements. This effectively provides an upper bound on element face angles and controls the growth of the maximum vertex degree.

DYNAMIC MESH ADAPTATION OF UNSTRUCTURED HEXAHEDRAL GRIDS
Rupak Biswas

Anisotropic mesh adaptation is a powerful tool for computing steady and unsteady three-dimensional problems that require local grid modifications in order to efficiently resolve solution features. However, repeated anisotropic subdivision can significantly deteriorate the quality of a tetrahedral mesh. Previous work has demonstrated that isotropic refinement is required if mesh quality is to be controlled effectively for arbitrary refinement levels. This is a serious limitation when directional flowfield features are present, leading to an inefficient distribution of grid points in the final mesh. In addition, truly anisotropic subdivision is almost impossible to realize for real problems on tetrahedral meshes.

A remedy for this drawback is to use hexahedral elements which can be subdivided anisotropically in any of the three directions without mesh quality problems. Hexahedral meshes also yield more accurate flowfield solutions than their tetrahedral counterparts for the same number of edges. Our adaptation procedure uses an edge data structure that facilitates efficient subdivision by allowing individual edges to be marked for refinement or coarsening.

Hexahedral adaptation schemes generate hanging vertices when a hexahedron cannot be split into smaller hexahedra without continuously propagating the mesh refinement into regions where it is not desired. We solve this problem by using pyramids and prisms as buffers between refined and unrefined elements. These buffer elements are never subdivided however. If an edge of a pyramid or a prism is marked for subdivision, then this element and its siblings coarsen back to their parent hexahedron and further refinement is performed directly on the hexahedral element. We expect this hexahedral mesh adaptation scheme to yield superior solutions with far fewer mesh points than the tetrahedral scheme.

SOLUTION-ADAPTIVE STRUCTURED/UNSTRUCTURED OVERSET GRID SCHEME FOR HELICOPTER AERODYNAMICS
Rupak Biswas, Earl P. N. Duque (US Army AFDD), and
Roger C. Strawn (US Army AFDD)

We are developing a hybrid method that solves both the three-dimensional thin-layer Navier-Stokes equations and the Euler equations using overset structured and solution-adaptive unstructured grids with applications to helicopter rotor flowfields.

Structured grid methods work well in viscous boundary layer regions near body surfaces. They employ implicit solution methods that allow larger time steps compared to explicit schemes. However, they are not well-suited for grid adaptation required to capture a vortex wake. In contrast, efficient dynamic grid adaptation is a major advantage of unstructured grids. In the wake of a helicopter blade, the solution methodology needs to resolve the convecting tip vortex and the vortex sheet. Unstructured grid methods work well in these flow regions. Moreover, the grids in the rotor wake need not be as fine as those in the boundary layers; therefore, the CFL limit of an explicit method does not constrain the time step.

HELICOPTER NOISE PREDICTION USING COMPUTATIONAL AEROACOUSTICS
Rupak Biswas and Roger C. Strawn (US Army AFDD)

High-performance helicopters and tiltrotors generate excessive noise both in forward flight and during takeoffs and landings. Accurate prediction of helicopter noise is essential to its control. Traditional acoustic analogy approaches cannot model the near-field nonlinear phenomena from high-speed rotor blades. CFD techniques are much better suited for these near-field nonlinearities but cannot efficiently propagate acoustic signals over long distances. The goal of this work is the accurate and efficient prediction of helicopter noise for a wide variety of flight regimes and rotor blade shapes using a hybrid CFD/Kirchhoff scheme.

We have used a solution-adaptive unstructured-grid Euler scheme with a stationary Kirchhoff method to model high-speed impulsive (HSI) noise for a hovering rotor. We have also developed a non-rotating Kirchhoff formulation to model HSI noise in forward flight. We are currently implementing a unified Kirchhoff surface formulation for HSI and blade-vortex interaction (BVI) noise to be combined with a CFD method for rotors in unsteady forward flight.

TENSOR METHODS FOR CONSTRAINED OPTIMIZATION
Dan Feng and Robert Schnabel (CU Boulder)

This research investigates tensor methods for nonlinear constrained optimization problems. These are general purpose methods especially intended for problems where the constraint gradient matrix is rank deficient or ill-conditioned at the solution. They are adapted from the standard successive quadratic programming (SQP) method by augmenting the linear model of the constraints with a simple second order term. The second order term is selected so that the model of the constraints interpolates constraint function values from several previous iterations, as well as the current constraint function value and gradients. Similar to tensor methods for nonlinear equations, the tensor methods for constrained optimization require no more function and-derivative evaluations, and hardly more storage or arithmetic per iteration, than the standard SQP methods. It is shown that the tensor methods are very efficient computationally on singular and well-conditioned nonlinear equality constrained optimization problems. A complete presentation of these methods can be found in CU technical report CU-CS-729-94.

A paper in preparation analyzes the local convergence properties of a version of tensor method on problems where the constraint gradient matrix at the solution has rank deficiency one. The tensor model uses the same quadratic model of the Lagrangian as in the objective of the standard SQP model, and augments the linear model of the constraints in the SQP model by a rank one quadratic term. We show under certain conditions that the sequence of iterates generated by the tensor method based upon the tensor model converges locally and two-step Q-superlinearly to the solution with Q-order $3/2$ on an interesting class of singular problems. In a similar situation, we show that the standard SQP method converges only linearly with constant converging to $1/2$. To the best of our knowledge, the result of this kind is the first for the SQP method. Hence, tensor methods have a theoretical advantage over the standard SQP method. Our analysis also confirms that the tensor method converges at least at the same rate as the standard SQP method on problems where the constraint gradient matrix at the solution has full rank.

OPTIMALITY CONDITIONS FOR SINGULAR CONSTRAINED OPTIMIZATION
Richard Byrd (CU Boulder), Dan Feng and Robert Schnabel (CU Boulder)

Optimality conditions for singular constrained optimization are not well understood. When the constraint gradient matrix has rank deficiency at the solution, the Kuhn-Tucker conditions are usually not satisfied. We introduce several optimality conditions for singular constrained optimization problems for the case when the constraint gradient matrix at the solution is rank deficient. These conditions are a generalization of the traditional Kuhn-Tucker conditions for constrained optimization. These conditions have useful implications both for constructing algorithms and for analyzing methods for singular constrained optimization problems. This research will be presented in a forthcoming paper.

TENSOR-KRYLOV METHODS FOR LARGE SPARSE SYSTEMS OF NONLINEAR EQUATIONS
Dan Feng and Thomas H. Pulliam (NASA Ames)

Tensor-Krylov methods for nonlinear equations are a coupling of tensor model formation and solution techniques for nonlinear equations with Krylov subspace projection techniques for unsymmetric systems of linear equations. A new tensor-Krylov algorithm, the tensor-GMRES algorithm, is developed and is shown to have significant computational advantage over the analogous Newton-GMRES method. A complete presentation of the tensor-GMRES method can be found in RIACS Technical Report 94.12. A software package that implements the tensor-GMRES algorithm is under development. A paper that introduces this package is in preparation.

One disadvantage of GMRES is its relatively large memory requirement compared to other Krylov subspace methods (e.g. BiCG). This difficulty is overcome by hybrid schemes that retain the advantages of the tensor-GMRES algorithm and provide the freedom of making use of other Krylov subspace methods.

Applications of tensor-Krylov methods to NASA ARC2D and INS2D codes are currently under investigation. These methods also have potential applications to aerodynamic analysis codes in industry (e.g. Boeing TRANAIR).

ENERGY ESTIMATES FOR NONLINEAR CONSERVATION LAWS
Pelle Olsson and Joseph Oliger

The purpose of this project is to analyze under what circumstances one can obtain energy estimates for nonlinear conservation laws. The basic idea behind this work is to establish certain principles that are valid for scalar equations as well as systems in one or more space dimensions, where the domains can be irregular. Furthermore, we have aimed for methods that to a large extent can be applied also to the discrete problem. Contrary to the pure initial value problem, it turns out that one must impose rather stringent conditions on the fluxes in order to obtain an energy estimate for the initial-boundary value problem, even in the scalar case. It should be noted that these conditions are trivially satisfied in the linear case. Symmetrizability, i.e., existence of an entropy function, and homogeneity of the flux vectors are very natural and important properties if one is to obtain an energy estimate for nonlinear systems. As a by-product we get an explicit expression for the entropy function and the corresponding entropy flux. Furthermore, for scalar equations it is possible to arrive at a maximum principle for the initial-boundary value problem. The detailed results can be found in RIACS TR 94.01.

HIGH-ORDER DIFFERENCE METHODS FOR HYPERBOLIC SYSTEMS
Bertil Gustafsson (Upsalla University) and Pelle Olsson

Implicit finite difference approximations of hyperbolic initial-boundary value problems are attractive from the accuracy point of view, because the leading coefficient in the error term is very small. Little is known about the stability properties. The major difficulty lies in combining the analytic boundary conditions with a numerical boundary closure such that strong stability is obtained. We prove strong stability for a fourth-order accurate implicit operator.

The computational efficiency of the fourth-order implicit operator is compared with that of a fourth-order explicit operator. Issues addressed include: number of grid points needed to reach a given tolerance level, number of arithmetic operations, sensitivity to discontinuities, parallelization/vectorization aspects etc. Having a filter interact with the difference operators, we look at issues of particular interest for real applications, such as shock resolution, the possibility of entropy violating shocks, and smearing of contact discontinuities. The results are presented in RIACS TR 94.04.

HIGH-ORDER DIFFERENCE METHODS WITH SHARP SHOCK RESOLUTION
Bertil Gustafsson and Pelle Olsson

In this research project we have analyzed how to modify centered high-order finite difference methods so as to yield accurate solution when solving nonlinear conservation laws. Centered finite difference methods, especially high-order ones, are often computationally efficient when the solution of the underlying problem is smooth. For non-smooth solutions, however, these methods produce excessive oscillations, which ultimately will ruin the solutions completely. One way to overcome the spurious phenomena is to introduce numerical viscosity which will smooth the numerical solution. But there are caveats; viscosity must not be used such that unnecessary smoothing occurs. As the computational mesh is refined the viscous effects should decrease while still damping the oscillations. We have developed a general theory on how to achieve sharp shock resolution for high-order finite difference approximations of systems of conservation laws by adding artificial viscosity. The addition of artificial viscosity to the discrete scheme is shown to be equivalent to the Lax k-shock condition. We refer to RIACS TR 94.07 for details.

ENERGY ESTIMATES FOR DIFFERENCE APPROXIMATIONS OF NONLINEAR CONSERVATION LAWS
Pelle Olsson

This project is a direct continuation of previous research projects carried out by the same author while at RIACS ("Summation by Parts, Projections, and Stability" and "Energy and Maximum Norm Estimates for Nonlinear Conservation Laws"). The ideas of TR 93.04 and TR 94.01 are transferred to semi-discrete systems. It turns out that the notion of entropy introduced by Lax is the natural extension of the linear L_2 stability concept to difference approximations of nonlinear conservation laws. We give several different characterizations of an entropy condition and show that they are all equivalent. The skew-symmetric form - which is an entropy condition - discussed in TR 94.01 is used to define a difference approximation of arbitrary order of accuracy. Using projections and summation by parts the resulting scheme is shown to satisfy an entropy condition, which in turn implies an energy estimate. The theory applies to initial-boundary value problems for symmetrizable systems of conservation laws in several space dimensions. A RIACS Technical Report is in progress.

HIGH PERFORMANCE PRECONDITIONER FOR LARGE SPARSE SYSTEMS
Wei-Pai Tang

An effective linear solver is an essential part of many scientific programs. Commonly, most of the CPU time of the computation is spent on the linear solver. When the size of the problems grows and the difficulty of the simulation increases, the performance of the linear solver becomes crucial.

The use of preconditioned Krylov space methods has been proven to be a competitive solution technique for wide ranges of large sparse matrix problems. It is acknowledged now that a high quality preconditioner is the key to the success. The objective of this project is to investigate several issues which are important for the performance of iterative methods. In particular, our emphasis will be on construction of an effective preconditioner. Both ILU (incomplete LU) preconditioner and approximate inverse preconditioner are considered. Our new techniques have improved the quality of the preconditioner.

RESEARCH ON COMPUTATIONAL FLUID DYNAMICS
Antony Jameson

Antony Jameson worked as a visiting researcher from Princeton University at RIACS from september 1993 until september 1994. A primary focus of his research during this time has been to pursue the underlying mathematical aspects of computational fluid dynamics.

His work on numerical methods has concentrated on the development of a unified theory for non-oscillatory shock capturing discretization schemes for both structured and unstructured meshes. An initial development of this theory was presented in AIAA Paper 93-3359. Since then the research has matured and has been presented in a series of works highlighting its various aspects (RIACS TR's 94.15, and 94.16). The theory, has led to new family of schemes (Symmetric Limited Positive - SLIP and Convective Upwind Split Pressure - CUSP) that exhibit superior numerical results, for both shock capturing and boundary layer modeling (AIAA paper 94-0647 and AIAA paper 95-0466).

Over the past decade the principles underlying the design of non-oscillatory discretization schemes for compressible flows have been quite well established. A very large number of variations of artificial diffusion, upwind biasing and flux splitting have been proposed and tested. In the same period multigrid acceleration schemes have also been the subject of widespread investigation, and have proved effective, particularly for subsonic and transonic flow. The use of limiters to enforce monotonic solutions, has proved, however, to have an adverse effect on multigrid convergence. In fact, it is not uncommon for schemes with limiters to become trapped in limit cycles. Limiters also tend to reduce the accuracy of solutions, particularly in regions containing smooth extrema.

In Jameson's current work a systematic procedure for the analysis and design of a broad class of schemes which satisfy monotonicity constraints on both structured and unstructured grids is developed. Secondly he shows ways in which these schemes can be modified to improve both their accuracy and their rate of convergence to a steady state.

Two main issues arise in the design of non-oscillatory discrete schemes. First there is the issue of how to construct an approximation to a scalar convection or convection-diffusion equation which is non-oscillatory, captures discontinuities with high resolution, and is sufficiently accurate. Second there is the issue of how to construct a numerical flux for a system of equations with waves traveling at different speeds, and sometimes in opposite directions. These two issues can be treated essentially independently, and by combining alternative non-oscillatory formulations with different constructions of the numerical flux one arrives at a matrix of candidate high resolution schemes, all of which may have acceptable characteristics. The performance of such a matrix of schemes for viscous boundary layers is explored in the AIAA paper 94-0647.

It is suggested that the principle of non-increasing maxima and non-decreasing minima provides a convenient criterion for the design of non-oscillatory schemes. This principle contains the concept of total variation diminishing (TVD) schemes for one-dimensional problems, but can readily be applied to multi-dimensional problems with both structured and unstructured grids. Such local extremum diminishing (LED) schemes can be realized by making sure that the coefficients of the discrete approximation are non-negative. First order accurate schemes satisfying this principle are easily constructed, but are too diffusive. It is well known that schemes which strictly satisfy the LED principle fall back to first order accuracy at extrema even when they realize higher order accuracy elsewhere. This difficulty can be circumvented by relaxing the LED requirement. Therefore the concept of essentially local extremum diminishing (ELED) schemes is introduced.

One approach to the construction of high resolution schemes which combine monotonicity and higher order accuracy is to blend low and high order diffusive terms as, for example, in the Jameson-Schmidt-Turkel (JST) scheme. Jameson has proved that with appropriately chosen coefficients the JST scheme is LED. Moreover, the coefficients can be chosen so that the scheme is both second order accurate at smooth extrema and ELED.

Another approach to the construction of higher order schemes, which has been adopted by several authors is to add limited anti-diffusive terms to a lower order scheme. This procedure is used to derive a general family of symmetric limited positive (SLIP) schemes for both structured and unstructured meshes. Moreover, the switch between low and high order terms in the JST scheme can be formulated in such a way that the JST scheme becomes a special case of the SLIP scheme. A slight modification of the SLIP formulation produces a corresponding family of upstream limited positive (USLIP) schemes are also derived, and resemble some well known upwind schemes. The limiters in the SLIP and USLIP schemes can be relaxed so that the schemes are second order accurate at smooth extrema and ELED.

Jameson has also performed an analysis of the conditions under which a discrete stationary shock can contain a single interior point. It emerges that a characteristic decomposition is not necessary to meet these conditions. Perfect single point discrete shocks completely free of oscillations can be produced by simpler flux splittings belonging to the class of convective upwind and split pressure (CUSP) schemes, in which scalar diffusion is augmented by pressure differences. It is actually possible to obtain high resolution with almost no oscillation by introducing the right amount of scalar diffusion, though this seems to result in a scheme which is less robust than the CUSP scheme.

13

RESEARCH IN AERODYNAMIC SHAPE OPTIMIZATION
James Reuther and Antony Jameson

Since the inception of CFD, researchers have sought not only accurate aerodynamic prediction methods for given configurations, but also design methods capable of creating new optimum configurations. Yet, while flow analysis can now be carried out over quite complex configurations using the Navier-Stokes equations with a high degree of confidence, direct CFD based design is still limited to very simple two-dimensional and three-dimensional configurations, usually without the inclusion of viscous effects. The CFD-based aerodynamic design methods that do exist can be broken down into two basic categories: inverse methods, and numerical optimization methods.

Inverse methods derive their name from the fact that they invert the goal of the flow analysis algorithm. Instead of obtaining the surface distribution of an aerodynamic quantity, such as pressure, for a given shape, they calculate the shape for a given surface distribution of an aerodynamic quantity. Most of these methods are based on potential flow techniques, and few of them have been extended to three-dimensions. The common trait of all inverse methods is their computational efficiency. Typically, transonic inverse methods require the equivalent of 2-10 complete flow solutions in order to render a complete design. Since obtaining a few solutions for simple two-dimensional and three-dimensional designs can be done in at most a few hours on modern computers systems, the computational cost of most inverse methods is considered to be minimal. Unfortunately, they suffer from many limitations and difficulties. Their most glaring limitation is that the objective is built directly into the design process and thus cannot be changed to an arbitrary or more appropriate objective function.

An alternative approach, which avoids some of the difficulties of inverse methods, but only at the price of heavy computational expense, is to use numerical optimization methods. The essence of these methods is very simple: a numerical optimization procedure is coupled directly to an existing CFD analysis algorithm. The numerical optimization procedure attempts to extremize a chosen aerodynamic measure of merit which is evaluated by the chosen CFD code. Most of these optimization procedures require gradient information in addition to evaluations of the objective function. Here, the gradient refers to changes in the objective function with respect to changes in the design variables. The simplest method of obtaining gradient information is by ``brute-force'' finite differences. In this technique, the gradient components are estimated by independently perturbing each design variable with a finite step, calculating the corresponding value of the objective function using CFD analysis, and forming the ratio of the differences. These methods are very versatile, allowing any reasonable aerodynamic quantity to be used as the objective function. They can be used to mimic an inverse method by minimizing the difference between target and actual pressure distributions, or may instead be used to maximize other aerodynamic quantities of merit such as L/D. Unfortunately, these brute-force numerical optimization methods, unlike the inverse methods, are computationally expensive because of the large number of flow solutions needed to determine the gradient information for a useful number of design variables. Tens of thousands of flow analyses would be required for a complete design. In our research, a new method that avoids the limitation and difficulties of traditional inverse methods while retaining their inherent computational efficiency is developed.

Considerable effort as been focused in the last year to develop control theory based aerodynamic shape optimization methods. At the beginning of the this year methods were already in place which showed that control theory could be used in conjuction with

numerical optimization and computational fluid dynamics to create efficient design tools for flows governed by the potential flow equation. AIAA Paper 94-0499

During the course of the year the method has been extended to treat the Euler equations. In our paper at the Multi-Disciplinary Optimization conference this summer (AIAA paper 94-4272) results were shown that demonstrated that control theory could be used to design airfoils under transonics by employing both an analytic mapping and a general mesh approach. Various objective functions were demonstrated showing the versatility of the new method. In the work presented at VKI, the first examples of three-dimensional wing design using control theory were presented. Finally, in a paper to be presented at the January 1995 Aerospaces Sciences Meeting (AIAA paper 95-0123) results will be presented for the design of wing and wing-body configurations over general meshes.

The method dramatically reduces the cost of aerodynamic optimization by replacing the expensive finite-difference method of calculating the required gradients by an adjoint variable formulation. After deriving the differential form of the adjoint equations and posing the correct boundary conditions based the objective functions, the resulting system is discretized and solved on the same mesh as that used for the flow solution. A significant economization is thus achieved by recycling the same subroutines that were used for the flow solution or the solution of the adjoint equations. The resulting design process requires only one flow calculation and one adjoint calculation per gradient evaluation, as opposed to the hundreds required for a finite-difference method. In practice the computational cost of the new method is two orders of magnitude less then a conventional approach.

Thus the new methods retain all of the versatility and power of numerical optimization while simultaneously reducing their computational cost to be equivalent to inverse methods.

ADAPTIVE REFINEMENT OF COMPOSITE CURVILINEAR GRIDS
Steven Suhr

A software system is being designed and implemented to manage the adaptive refinement of composite curvilinear grids for the approximate solution of time-dependent partial differential equations. With the simplifying assumption that the spatial domain has fixed boundaries, an initial grid is constructed using a fixed set of overlapping grids, which collectively conform to the boundaries and cover the domain. Refinement grids, aligned with the original base grids, are added to maintain accuracy as the solution evolves. This approach organizes the grids into a geometrically nested tree of connected components, and it explores the use of curvilinear stairstep refinement grids.

The programming language Vorpal, currently under development, will be used in the implementation of this system, taking advantage of Vorpal's support for data structures, abstract data types, structured external files, and modular program structure. An important milestone in the transformation of Vorpal from a collection of useful concepts into a unified preliminary design will be the translation of adaptive grid code being implemented at Stanford by others for model problems in two space dimensions, from C into Vorpal. When an implementation of Vorpal exists, the adaptive grid system will be extended from two to three space dimensions, and a version which can be applied more readily to realistic and diverse problems will be created. As the adaptive grid system evolves and grows, the anticipated future support in Vorpal for concurrency should also be useful.

C. HIGH PERFORMANCE NETWORKS

BAY AREA GIGABIT NETWORK TESTBED
Marjory J. Johnson

The basic infrastructure for the Bay Area Gigabit Testbed (BAGNet) has been established, and application development has begun.

The fifteen organizations participating in BAGNet are connected by OC-3c (155 Mbps) links. Each site may have up to four hosts directly connected to the testbed. We have established a mesh of permanent virtual channels to provide connections between all possible pairs of hosts. We will use switched virtual channels as soon as the technology is ready. We are using AAL5 as the adaptation layer, and we are running IP over ATM.

The featured application for BAGNet is the teleseminar application. Our initial goal is to implement simple multicast video transmissions. We have set up point-to-multipoint permanent virtual channels, so that each host on BAGNet can broadcast to all BAGNet sites. This enables us to simulate multicast until a better solution is developed. Various sites are currently running experiments over these point-to multipoint virtual channels.

We are using the World Wide Web (WWW) to coordinate testbed activities and to post information about BAGNet. The URL of the general BAGNet home page on WWW, which is maintained by Lawrence Berkeley Laboratory (LBL), is http://george.lbl.gov/BAGNet.html. Site-specific BAGNet home pages, which contain information about the individual testbed sites, can be accessed from the LBL home page. The NASA Ames BAGNet home page can be accessed directly at http://www.nas.nasa.gov/NAS/groups/RND/netdev/bagnet/bagnet.html.

BAGNet is sponsored by Pacific Bell's CalREN (California Research and Education) program. M. Johnson, Bill Johnston of LBL, and Dan Swinehart of Xerox PARC are jointly coordinating the testbed project.

COLLABORATIVE SCIENCE
Marjory J. Johnson

The objective of this project is to develop new collaborative working paradigms for earth and space scientists by combining data-analysis tools with state-of-the art networking tools. These new paradigms will address the use of massive data sets that are stored remotely, collaboration between geographically separated colleagues, and joint data analysis.

Project collaborators include earth scientists at Ames and at the University of Arizona, space scientists at Lockheed Palo Alto Research Laboratory (PARL), and networking researchers at Sandia National Laboratories - Livermore.

The project infrastructure is almost complete. We are developing a collaborative work environment using a workstation which is directly connected to BAGNet (called the BAGNet workstation in the remainder of this text). Scientific data-analysis tools we are using include ANA, a data-analysis package developed at Lockheed PARL, and IDL, a commercial data-analysis package that is widely used by scientists in multiple disciplines. Software packages that we are installing to enable the collaborative work include DAVE

(Distributed Audio/Video Environment), a package developed at Sandia National Laboratories - Livermore for audio/video teleconferencing, and X/Telescreen, a commercial package that enables the sharing of X window applications. The BAGNet workstation currently has audio capabilities; we are in the process of adding video capabilities.

This project will involve both experimental work and analytic work to model and evaluate the new paradigms that we develop. At this time we have conducted preliminary collaboration experiments using only audio conferencing. We have conducted testing locally between two workstations at RIACS, between a workstation at RIACS and a workstation at the University of Arizona accessed over the Internet, and between the BAGNet workstation at RIACS and a workstation at Sandia accessed over BAGNet. Once the BAGNet workstation has video capabilities, we will be able to conduct more meaningful experiments and begin the modeling and evaluation aspects of the project.

Summer visitors who contributed to this project include Marc Bumble, a graduate student from Pennsylvania State University, Justin Paola, a graduate student from the University of Arizona, and Professor Robert Schowengerdt, also from the University of Arizona. Bumble, a computer scientist, assisted in development of the collaboration environment for my workstation, and assisted in the experimental work. Paola and Schowengerdt, specialists in image processing, are developing techniques for using neural networks to browse remote databases of earth science imagery.

D. LEARNING SYSTEMS

MODEL-BASED LEARNING
Wray Buntine

This research involves developing methods and tools for doing supervised and unsupervised learning. In 1994 a technique was developed for constructing a "learning compiler," a system that takes as input a specification for a data analysis problem and compiles source code for addressing the problem. The technique has been presented at conferences and graduate seminars. An early prototype developed by summer student Scott Roy, a Ph.D. student at Stanford University, enabled a rapid implementation of several algorithms including Cheeseman et al.'s Autoclass III. A second implementation is currently being planned.

The application with Marshall Space Flight Center progressed on the automatic analysis of spectral data from the Space Shuttle Main Engine (SSME). A suite of software developed in C and the IDL language (from Research Systems, Inc.) was installed at Marshall for their use. This will now be the main software tools for their analysis of the spectral data. Methods are currently being developed to increase the accuracy of the estimated produced by the system and improve anomaly detection.

SUPER-RESOLUTION/SURFACE MODELING
Peter Cheeseman

Proposed Capability: An algorithm that takes multiple images of the same area/object and combines the information to give a much higher resolution image, and identification of ground truth in the image.

Benefits: More science from images from Viking, Voyager, Galileo, Landsat, Tiros etc., through improved resolution. Also, lossless data compression and change detection for Earth observation data. An 8:1 increase in resolution has already been demonstrated.

User Interest: Preliminary discussions with LORAL (EOS-DIS contractor) have identified potential areas for collaboration. A demonstration project with University of Colorado at Boulder [Bill Emory] will develop software for earth science applications. Presentation to planetary scientists has resulted in collaborations to produce higher resolution images of sites of considerable scientific interest.

Industry and Academic Involvement: Collaboration and co-funding of development of our algorithms for earth science applications at the University of Colorado at Boulder.

Commercial Use: We are working with Research Triangle Institute to identify and interest potential commercial applications. Several short-term and long-term possibilities have been identified.

Additional Supporting Program: Currently no additional support, but there is potential for support from an EOS NRA in collaborating with other EOS investigators.

Beginning in late 1991, the Bayes group at ARC began a project to develop the theory and practice of multiple image data combination. Multiple images taken from similar loca-

tions and under similar lighting conditions contain similar - but not identical - information. Slight differences in instrument orientation and position produce mismatches between the pixel grids of different images. These mismatches ensure that any point on the ground is sampled differently in each image. The surface modeling project is designed to exploit these differences to build a super-resolved composite image that uses all the information from the separate images.

The basic theory behind our approach is that of inverse graphics. That is, if we knew what the ground is like, the lighting conditions, and the camera orientation and characteristics, etc., then we could predict what the camera would see (an image). This is the standard computer graphics problem. However, we have the inverse problem—we know what the images are, and we want to find the most probable ground truth (surface) that would have generated them, assuming we know the lighting conditions and camera characteristics. The most important (and difficult) part of this process is recovering the camera orientation and position for each image. To do this, we must register all the images with respect to each other to an accuracy of a small fraction of a pixel; this registration tells us how an image maps onto the ground truth model we are building. Our initial ground model is formed by letting each pixel "vote" on what the corresponding ground position should be depending on how much that ground position contributed to that pixel. This initial ground model is then used to project what each image should be (i.e., predict each pixel value). The differences between the predicted pixel value and the observed value is used to update the ground model until it cannot be further improved. This procedure produces an increase in both spatial resolution and gray-scale resolution.

The above procedure for super-resolution surface reconstruction from multiple images could be implemented with standard statistical (maximum likelihood) methods. The distinctly Bayesian contribution is to include neighbor correlation into the ground pixels model. In other words, in the surface reconstruction, a ground pixel is not just a function of the corresponding projected pixels, but also depends on its neighbor values. This correlation addition makes our approach much less sensitive to noise in the pixel values, and produces more realistic reconstruction. Test results with 8:1 resolution enhancement, and Viking Orbiter imagery with 2:1 to 4:1 enhancements are attached.

Fiscal Year 1994 Progress
In FY94 we have refined our registration algorithm so that we can achieve registration to at least within 1/10th. of a pixel. This very accurate registration has allowed us to obtain 8:1 improvement in spatial resolution, as shown in the figure. In the process of developing this accurate registration process, we had to invent a new procedure for function minimization from generated function values that maximized the use of information in these very expensive function values. This new procedure has applications well beyond our particular problem. We have also increased the speed and reliability of the code, and tested it on scientifically important test cases on Mars (also shown). We are also exploring potential commercial applications of this technology in collaboration with Research Triangle Park Institute.

Our test cases to date have effectively been 2-D surfaces (small patches of Mars or Earth). However, the Voyager imagery of the moons of Jupiter and Saturn, required us to extend the theory to projections onto spherical surfaces. We are currently testing this extended theory, and expect to get improved resolution images of some of these moons that exceeds the best currently available. We are also developing more extensive 3-D super-resolution surface reconstruction for irregular shaped objects such as Gaspra (an asteroid) from images obtained by the Galileo fly-by.

In 1994, we also began a collaborative research arrangement with the University of Colorado at Boulder to apply the super-resolution techniques to earth observing satellite images. This collaboration involves developing a 3-D surface model of selected earth test sites, using a USGS elevation model as the starting point. We have developed techniques for taking satellite earth observation data, particularly AVHRR data available daily, and integrating this information into the high resolution 3-D grid. Even though the AVHRR pixel size is 1 kilometer square, at best, we expect to get roughly an order of magnitude improvement in resolution by integrating many images. This approach allows us to use other satellite data, such as GOES, OLS, SPOT and LandSat, depending on availability.

Fiscal Year 1995 Plan
Our current algorithm is limited to flat areas with nearly identical viewing angle and lighting conditions. We plan to extend the theory and implementation surface reconstruction for spherical planetary surfaces, and apply the resulting algorithms to all Voyager planetary images of scientific interest. We will also continue the development of the theory and implementation of arbitrary 3-D surface models, with any combination of lighting and viewing conditions. This extension will require modeling of shadows and atmospheric conditions. These extensions will allow us to produce super-resolved images (3-D models) of objects such as asteroids and close-ups of earth's surface.

These extensions require that we extend the theory to integrate information across different spectral bands. We expect to be able to do this by allowing correlations between the bands that vary slowly across the image. For earth images we must also develop atmospheric models to remove the effects of clouds and haze from individual images. We expect that information from the half-hourly GOES images will help separate the rapidly changing cloud effects from the relatively constant earth background. We also plan to develop "classification models" where each ground grid point is assigned a percentage membership of all the known classes of ground cover (e.g., soil, grass, trees, etc.). The spectral characteristics of each class is known, so it is possible to calculate, for each grid point, what the expected spectral characteristics of that point should be. The observed characteristics allow us to update the classification vector at each point, including the assumption that the classification vectors for adjacent points are highly correlated. This kind of classification model will be extremely useful for quantitative earth science studies.

Primary Users: EOS-DIS, earth scientists, planetary scientists.

Potential Users: Anyone trying to get more information out of a set of images—e.g., high resolution television; improved microscope and telescope resolution; higher fax resolution; etc.

Delivery Year of Major Capability: 1995

KNOWLEDGE COMPILATION AND LEARNING FOR EFFICIENT GENERAL PROBLEM-SOLVING
Barney Pell

This research investigates methods to increase the generality of problem-solving systems by combining general-purpose knowledge with details of a specific problem to be solved in order to produce an efficient special-purpose system to solve that problem efficiently. This separation of general knowledge from specifics of a given problem provides several advantages. It facilitates transfer of knowledge from one problem domain to another, enables a user to specify the problem in high-level language without confusing the problem with its implementation, and the concise representation resulting from such a problem-formulation enables the machine to learn from experience across a variety of domains. Current work centers on extending the author's Ph.D. thesis research, which focused on games (like chess and checkers) to problems of industrial importance. Such problems include information retrieval and automatic scheduling.

INDUCTION FOR HEURISTIC INFORMATION RETRIEVAL
Barney Pell, Catherine Baudin (NASA/Recom) and
Smadar Kedar (ILS, Northwestern University)

This research involves developing methods for using supervised learning techniques to aid information retrieval (IR). A general problem within IR is that a user's QUERY may have no direct match to any indexed document, but it would be desirable for the system to return something related to that query anyway. The framework we are using involves having a knowledge engineer design a set of RETRIEVAL HEURISTICS. When the indexing system has no direct match to a user's query, the heuristics operate on the query to generate an extended query set, and the cases matching these extended queries are offered to the user.

For example, the user may ask a query about the function of a widget. If the system has no documents indexed under the function of widgets, some simple strategies are to return all documents about widgets (of which there may be thousands), or to provide nothing at all. The heuristic retrieval approach offers more sophisticated response strategies. For example, the heuristic "function-part" might determine that widgets are a part of a larger system, say gadgets, and would then produce an extended query for documents which provide detailed descriptions of the function of gadgets.

One limitation of this approach is that the heuristics may generate a large number of cases which are IRRELEVANT to the user. One reason for this is that the heuristics are OVER-GENERAL; another is that some heuristics may be more useful than others depending on CONTEXT.

One solution to this problem, which we have recently started exploring, is to use FEED-BACK from users to specialize these heuristics automatically. Whenever the user is presented with a list of extended cases, the user can designate a number of these cases as relevant or irrelevant (perhaps with varying degrees of relevance). The result is thus a set of EXAMPLES, consisting of properties of both the query and the extended case, each CLASSIFIED as to its relevance. The induction problem is to use these examples to learn a classification function, which will return for each new extended case a prediction of the relevance to the user. Such a function could be used to order extended cases by predicted relevance, which would greatly simplify the user's task by offering him the most relevant choices first. Details of this approach are presented as RIACS Technical Report 94.03.

III. TECHNICAL REPORTS

TR 94.01

ENERGY AND MAXIMUM NOR ESTIMATES FOR NONLINEAR CONSERVATION LAWS

Pelle Olsson and Joseph Oliger
January 1994, 26 pages
Submitted for publication.

We have devised a technique that makes it possible to obtain energy estimates for initial-boundary value problems for nonlinear conservation laws. The two major tools to achieve the energy estimates are a certain splitting of the flux vector derivative f(u)_x, and a structural hypothesis, referred to as a cone condition, on the flux vector f(u). These hypotheses are fulfilled for many equations that occur in practice, such as the Euler equations of gas dynamics. It should be noted that the energy estimates are obtained without any assumptions on the gradient of the solution u. The results extend to weak solutions that are obtained as pointwise limits of vanishing viscosity solutions. As a by-product we obtain explicit expressions for the entropy function and the entropy flux of symmetrizable systems of conservation laws. Under certain circumstances the proposed technique can be applied repeatedly so as to yield estimates in the maximum norm.

TR 94.02

ON THE DYNAMICS OF SOME GRID ADAPTION SCHEMES

P.K.Sweby (University of Reading, UK) and H.C. Yee (NASA Ames Research Center)
February 1994, 14 pages

Preprint for Proceedings of the 4th International Conference on Numerical Grid Generation in Computational Fluid Dynamics and Related Fields, Swansea, UK, April 6-8, 1994.

The dynamics of a one-parameter family of mesh equidistribution schemes coupled with finite difference discretisations of linear and non-linear convection-diffusion model equations is studied numerically. It is shown that, when time marched to steady state, the grid adaption not only influences the stability and convergence rate of the overall scheme, but can also introduce spurious dynamics to the numerical solution procedure.

TR 94.03

INCREASING LEVELS OF ASSISTANCE IN REFINEMENT OF KNOWLEDGE-BASED RETRIEVAL SYSTEMS

Catherine Baudin, Smadar Kedar, and Barney Pell

This paper is concerned with the task of incrementally acquiring and refining the knowledge and algorithms of a knowledge-based system in order to improve its performance over time. In particular, we present the design of DE-KART, a tool whose goal is to provide increasing levels of assistance in acquiring and refining indexing and retrieval knowledge for a knowledge-based retrieval system. DE-KART starts with knowledge that has been entered manually, and increases its level of assistance in acquiring and refining that knowledge, both in terms of the increased level of automation in interacting with users, and in terms of the increased generality of the knowledge. DE-KART is at the intersection of machine learning and knowledge acquisition: it is a first step towards a system which moves along a continuum from interactive knowledge acquisition to increasingly automated machine learning as it acquires more knowledge and experience. Keywords: Conceptual indexing and retrieval, increasing levels of assistance, inductive learning.

TR94.04

FOURTH ORDER DIFFERENCE METHHODS FOR HYPERBOLIC IBVP

Bertil Gustafsson and Pelle Olsson

In this paper we consider fourth order difference approximations of initial-boundary value problems for hyperbolic partial differential equations. We use the method of lines approach with both explicit and compact implicit difference operators in space. The explicit operator satisfies an energy estimate leading to strict stability. For the implicit operator we develop boundary conditions and give a complete proof of strong stability using the Laplace transform technique.

We also present numerical experiments for the linear advection equation and Burgers' equation with discontinuities in the solution or in its derivative. The first equation is used for modeling contact discontinuities in fluid dynamics, the second one for modeling shocks and rarefaction waves. The time discretization is done with a third order Runge-Kutta TVD method. For solutions with discontinuities in the solution itself we add a filter based on second order viscosity.

In case of the non-linear Burgers' equation we use a flux splitting technique that results in an energy estimate for certain difference approximations, in which case also an entropy condition is fulfilled. In particular we shall demonstrate that the unsplit conservative form produces a non-physical shock instead of the physically correct rarefaction wave.

In the numerical experiments we compare our fourth order methods with a standard second order one and with a third order TVD-method. The results show that the fourth order methods are superior.

TR 94.05

·UNSTRUCTURED GRIDS ON SIMD TORUS MACHINES

Petter E. Bjorstad (University of Bergen), and Robert Schreiber
March 1994, 8 pages

Unstructured grids lead to unstructured communication on distributed memory parallel computers, a problem that has been considered difficult. Here, we consider adaptive, off-line communication routing for a SIMD processor grid. Our approach is empirical. We use large data sets drawn from supercomputing applications instead of an analytic model of communication load. The chief contribution of this paper is an experimental demonstration of the effectiveness of certain routing heuristics. Our routing algorithm is adaptive, nonminimal, and is generally designed to exploit locality. We have a parallel implementation of the router, and we report on its performance.

TR 94.06

COMPUTATION OF HELICOPTER ROTOR ACOUSTICS IN FORWARD FLIGHT

Roger Strawn (US Army AFDD) and Rupak Biswas
March 1994, 10 pages

This paper presents a new method for computing acoustic signals from helicopter rotors in forward flight. The aerodynamic and acoustic solutions in the near field are computed with a finite-difference solver for the Euler equations. A nonrotating cylindrical Kirchhoff surface is then placed around the entire rotor system. This Kirchhoff surface moves subsonically with the rotor in forward flight. The finite-difference solution is interpolated onto this cylindrical surface at each time step and a Kirchhoff integration is used to carry the acoustic signal to the far field. Computed values for high-speed impulsive noise show excellent agreement with model-rotor and flight-test experimental data. Results from the new method offer high accuracy with reasonable computer resource requirements.

TR 94.07

HIGH-ORDER CENTERED DIFFERENCE METHODS WITH SHARP SHOCK RESOLUTION

Bertil Gustafsson and Pelle Olsson
July 1994, 24 pages

In this paper we consider high-order centered finite difference approximations of hyperbolic conservation laws. We propose different ways of adding artificial viscosity to obtain sharp shock resolution. For the Riemann problem we give simple explicit formulas for obtaining stationary one- and two-point shocks. This can be done for any order of accuracy. It is shown that the addition of artificial viscosity is equivalent to ensuring the Lax k-shock condition. We also show numerical experiments that verify the theoretical results.

TR 94.08

COMPARISONS OF NEURAL NETWORKS TO STANDARD TECHNIQUES FOR IMAGE CLASSIFICATIONAND CORRELATION

Justin D. Paola and Robert A. Schowengerdt (University of Arizona, Tucson)
July 1994, 11 pages
To be presented at IEEE Geoscience and Remote Sensing Symposium, August 10th, 1994, Pasadena, CA.

Neural network techniques for multispectral image classification and spatial pattern detection are compared to the standard techniques of maximum-likelihood classification and spatial correlation. The neural network produced a more accurate classification than maximum-likelihood of a Landsat scene of Tucson, Arizona. Some of the errors in the maximum-likelihood classification are illustrated using decision region and class probability density plots. As expected, the main drawback to the neural network method is the long time required for the training stage. The network was trained using several different hidden layer sizes to optimize both the classification accuracy and training speed, and it was found that one node per class was optimal. The performance improved when 3x3 local windows of image data were entered into the net. This modification introduces texture into the classification without explicit calculation of a texture measure. Larger windows were successfully used for the detection of spatial features in Landsat and Magellan synthetic aperture radar imagery.

TR 94.09

ARRAY DISTRIBUTION IN DATA-PARALLEL PROGRAMS

Siddhartha Chatterjee, John R. Gilbert, Robert Schreiber and Thomas J. Sheffler
July 1994, 17 pages

We consider distribution at compile time of the array data in a distributed-memory implementation of a data-parallel program written in a language like Fortran 90. We allow dynamic redistribution of data and define a heuristic algorithmic framework that chooses distribution parameters to minimize an estimate of program completion time. We represent the program as an alignment-distribution graph. We propose a divide-and-conquer algorithm for distribution that initially assigns a common distribution to each node of the graph and successively refines this assignment, taking computation, realignment, and redistribution costs into account. We explain how to estimate the effect of distribution on computation cost and how to choose a candidate set of distributions. We present the results of an implementation of our algorithms on several test problems.

TR 94.10

ALIGNING PARALLEL ARRAYS TO REDUCE COMMUNICATION

Thomas J. Sheffler, Robert Schreiber, John R. Gilbert, and Siddhartha Chatterjee.
July 1994, 18 pages

Axis and stride alignment is an important optimization in compiling data-parallel programs for distributed-memory machines. We previously developed an optimal algorithm for aligning array expressions. Here, we examine alignment for more general program graphs. We show that optimal alignment is NP-complete in this setting, so we study heuristic methods.

This paper makes two contributions. First, we show how local graph transformations can reduce the size of the problem significantly without changing the best solution. This allows more complex and effective heuristics to be used. Second, we give a heuristic that can explore the space of possible solutions in a number of ways. We show that some of these strategies can give better solutions than a simple greedy approach proposed earlier. Our algorithms have been implemented; we present experimental results showing their effect on the performance of some example programs running on the CM-5.

TR 94.11

AN OBJECT-ORIENTED APPROACH TO NESTED DATA PARALLELISM

Thomas J. Sheffler and Siddhartha Chatterjee
July 1994, 14 pages

This paper describes an implementation technique for integrating nested data parallelism into an object-oriented language. Data-parallel programming employs sets of data called "collections" and expresses parallelism as operations performed over the elements of a collection. When the elements of a collection are also collections, then there is the possibility for"nested data parallelism." Few current programming languages support nested data parallelism however.

In an object-oriented framework, a collection is a single object. Its type defines the parallel operations that may be applied to it. Our goal is to design and build an object-oriented data-parallel programming environment supporting nested data parallelism. Our initial approach is built upon three fundamental additions to C++. We add new parallel base types by implementing them as classes, and add a new parallel collection type called a "vector" that is implemented as a template. Only one new language feature is introduced: the "foreach" construct, which is the basis for exploiting elementwise parallelism over collections.

The strength of the method lies in the compilation strategy, which translates nested data-parallel C++ into ordinary C++. Extracting the potential parallelism in nested "foreach" constructs is called "flattening" nested parallelism. We show how to flatten "foreach" constructs using a simple program transformation. Our prototype system produces vector code which has been successfully run on workstations, a CM-2 and a CM-5.

TR 94.12

TENSOR-GMRES METHOD FOR LARGE SPARSE SYSTEMS OF NONLINEAR EQUATIONS

Dan Feng, Thomas H. Pulliam (NASA Ames Research Center)
August 1994, 38 pages

This paper introduces a tensor-Krylov method, the tensor-GMRES method, for large sparse systems of nonlinear equations. This method is a coupling of tensor model formation and solution techniques for nonlinear equations with Krylov subspace projection methods for nonsymmetric systems of linear equations. A major disadvantage of traditional tensor methods is that the solution of the tensor model requires the factorization of the Jacobian matrix, which may not be suitable for problems where the Jacobian matrix is large and has a "bad" sparsity structure for an efficient factorization. We overcome this difficulty by forming and solving the tensor model using an extension of a Newton-GMRES scheme. Like traditional tensor methods, we show that the new tensor method has significant computational advantages over the analogous Newton counterpart. Consistent with Krylov subspace based methods, the new tensor method does not depend on the factorization of the Jacobian matrix. As a matter of fact, the Jacobian matrix is never needed explicitly.

TR 94.13

IMPROVED LOAD DISTRIBUTION IN PARALLEL SPARSE CHOLESKY FACTORIZATION

Edward Rothberg (Intel Supercomputer Systems), and Robert Schreiber
August 1994, 10 pages
To appear: Proceedings of Supercomputing 94

Compared to the customary column-oriented approaches, block-oriented, distributed-memory sparse Cholesky factorization benefits from an asymptotic reduction in interprocessor communication volume and an asymptotic increase in the amount of concurrency that is exposed in the problem. Unfortunately, block-oriented approaches (specifically, the block fan-out method) have suffered from poor balance of the computational load. As a result, achieved performance can be quite low. This paper investigates the reasons for this load imbalance and proposes simple block mapping heuristics that dramatically improve it. The result is a roughly 20% increase in realized parallel factorization performance, as demonstrated by performance results from an Intel Paragon system. We have achieved performance of nearly 3.2 billion floating point operations per second with this technique on a 196-node Paragon system.

TR 94.14

STABILITY AND ERROR ESTIMATION FOR COMPONENT ADAPTIVE GRID METHODS

Joseph Oliger and Xiaolei Zhu
August 1994, 29 pages
Publication: accepted for publication by Journal of Comp. and Appl. Math.

Component adaptive grid (CAG) methods for solving hyperbolic partial differential equations (PDE's) are discussed in this paper. Applying recent stability results for a class of numerical methods on uniform grids, the convergence of these methods for linear problems on component adaptive grids is established here. Furthermore, the computational error can be estimated on CAG's using the stability results. Using this estimation, the error can be controlled on CAG's. Thus, the solution can be computed efficiently on CAG's within a given error tolerance. Computational results of linear problems in 1-D and 2-D are presented.

TR 94.15

ANALYSIS AND DESIGN OF NUMERICAL SCHEMES FOR GAS DYNAMICS ARTIFICIAL DIFFUSION, UPWIND BIASIS, LIMITERS AND THEIR EFFECT ON ACCURACY AND MULTIGRID CONVERGENCE

Antony Jameson (Princeton University)
August 1994, 40 pages

The theory of non-oscillatory scalar schemes is developed in this paper in terms of the local extremum diminishing (LED) principle that maxima should not increase and minima should not decrease. This principle can be used for multi-dimensional problems on both structured and unstructured meshes, while it is equivalent to the total variation diminishing (TVD) principle for one-dimensional problems. A new formulation of symmetric limited positive (SLIP) schemes is presented, which can be generalized to produce schemes with arbitrary high order of accuracy in regions where the solution contains no extrema, and which can also be implemented on multi-dimensional unstructured meshes. Systems of equations lead to waves traveling with distinct speeds and possibly in opposite directions. Alternative treatments using characteristic splitting and scalar diffusive fluxes are examined, together with a modification of the scalar diffusion through the addition of pressure differences to the momentum equations to produce full upwinding in supersonic flow. This convective upwind and split pressure (CUSP) scheme exhibits very rapid convergence in multigrid calculations of transonic flow, and provides excellent shock resolution at very high Mach numbers.

TR 94.16

ANALYSIS AND DESIGN OF NUMERICAL SCHEMES FOR GAS DYNAMICS 2 ARTIFICIAL DIFFUSION AND DISCRETE SHOCK STRUCTURE

Antony Jameson (Princeton University)
August 1994, 29 pages

The effect of artificial diffusion on discrete shock structures is examined for a family of schemes which includes scalar diffusion, convective upwind and split pressure (CUSP) schemes, and upwind schemes with characteristic splitting. The analysis leads to conditions on the diffusive flux such that stationary discrete shocks can contain a single interior point. The simplest formulation which meets these conditions is a CUSP scheme in which the coefficients of the pressure differences is fully determined by the coefficient of convective diffusion. It is also shown how both the characteristic and CUSP schemes can be modified to preserve constant stagnation enthalpy in steady flow, leading to four variants, the E and H-characteristic schemes, and the E and H-CUSP schemes. Numerical results are presented which confirm the properties of these schemes.

TR 94.17

OPTIMUM AERDYNAMIC DESIGN VIA BOUNDARY CONTROL

Antony Jameson (Princeton University)
August 1994, 33 pages
In proceedings of the AGARD-VKI Lecture Series on Optimum Design Methods in Aero-
dynamics, von Karman Institute for fluid Dynamics, 1994.

These lectures describe the implementation of optimization techniques based on control
theory for airfoil and wing design. In previous studies by the author it was shown that con-
trol theory could be used to devise an effective optimization procedure for two-dimension-
al profiles in which the shape is determined by a conformal transformation from a unit
circle, and the control is the mapping function. Recently the method has been implement-
ed in an alternative formulation which does not depend on conformal mapping, so that it
can more easily be extended to treat general configurations. The method has also been ex-
tended to treat the Euler equations, and results are presented for both two and three di-
mensional cases, including the optimization of a swept wing.

TR 94.18

CONTROL THEORY BASED AIRFOIL DESIGN USING THE EULER EQUATIONS

Antony Jameson, James Reuther
September 1994, 17 pages
Publication: 5th AIAA/NASA/USAF/ISSMO Symposium on Multidisciplinary Analysis
and Optimization AIAA paper 94-4272

This paper describes the implementation of optimization techniques based on control the-
ory for airfoil design. In our previous work it was shown that control theory could be em-
ployed to devise effective optimization procedures for two-dimensional profiles by using
the potential flow equation with either a conformal mapping or a general coordinate sys-
tem. The goal of our present work is to extend the development to treat the Euler equations
in two-dimensions by procedures that can readily be generalized to treat complex shapes
in three-dimensions. Therefore, we have developed methods which can address airfoil de-
sign through either an analytic mapping or an arbitrary grid perturbation method applied
to a finite volume discretization of the Euler equations. Here the control law serves to pro-
vide computationally inexpensive gradient information to a standard numerical optimiza-
tion method. Results are presented for both the inverse problem and drag minimization
problem.

TR 94.19

MESH QUALITY CONTROL FOR MULTIPLY-REFINED TETRAHEDRAL GRIDS

Rupak Biswas and Roger C. Strawn (US Army AFDD)
November, 14 pages

A new algorithm for controlling the quality of multiply-refined tetrahedral meshes is presented in this paper. The basic dynamic mesh adaption procedure allows localized grid refinement and coarsening to efficiently capture aerodynamic flow features in computational fluid dynamics problems; however, repeated application of the procedure may significantly deteriorate the quality of the mesh. Results presented show the effectiveness of this mesh quality algorithm and its potential in the area of helicopter aerodynamics and acoustics.

TR 94.20

WEIGHTED GRAPH BASED ORDERING TECHNIQUES FOR PRECONDITIONED CONJUGATE GRADIENT METHODS

Simon S. Clift (Queens University Canada) and Wei-Pai Tang (University of Waterloo, Waterloo, Ontario, Canada)
December 1994, 23 pages

We describe the basis of a matrix ordering heuristic for improving the incomplete factorization used in preconditioned conjugate gradient techniques applied to anisotropic PDE's. Several new matrix ordering techniques, derived from well-known algorithms in combinatorial graph theory, which attempt to implement this heuristic, are described. These ordering techniques are tested against a number of matrices arising from linear anisotropic PDE's, and compared with other matrix ordering techniques. A variation of RCM is shown to generally improve the quality of incomplete factorization preconditioners.

TR 94.21

PERFORMANCE ISSUES FOR ITERATIVE SOLVERS IN DEVICE SIMULATION

Qing Fan (P. A. Forsyth Univ. of Waterloo), J. R. F. McMacken (Goal Electronics Inc.), Wei-Pai Tang (University of Waterloo, Waterloo, Ontario, Canada)
December 1994, 19 pages

Due to memory limitations, iterative methods have become the method of choice for large scale semiconductor device simulation. However, it is well known that these methods still suffer from reliability problems. The linear systems which appear in numerical simulation of semiconductor devices are notoriously ill-conditioned. In order to produce robust algorithms for practical problems, careful attention must be given to many implementation issues. This paper concentrates on strategies for developing robust preconditioners. In addition, effective data structures and convergence check issues are also discussed. These algorithms are compared with a standard direct sparse matrix solver on a variety of problems.

TR UPDATES:

The following TR's have been *accepted* for publication:
- Summation by Parts, Projections, and Stability I, TR-93-04, to appear in <u>Mathematics of Computation</u>.
- Summation by Parts, Projections, and Stability II, TR-93-04, to appear in <u>Mathematics of Computation</u>.
- On the Superconvergence of Galerkin Methods for Hyperbolic IBVP, TR-93-07, to appear in <u>SIAM Journal of Numerical Analysis</u>.

The following TR's have been *submitted* for publication:
- Energy and Maximum Norm Estimates for Nonlinear Conservation Laws, TR-94-01, submitted to <u>Journal of Mathematical Analysis and Applications</u>.
- Fourth Order Difference Methods for Hyperbolic IBVP's, TR-94-04, submitted to <u>Journal of Computational Physics</u>.
- High-Order Centered Difference Methods with Sharp Shock Resolution, TR-94-07, submitted to <u>SIAM Journal of Numerical Analysis</u>.

IV. PUBLICATIONS

RUPAK BISWAS

Biswas, R., and Strawn, R. C., (US Army AFDD), "A New Procedure for Dynamic Adaption of Three-Dimensional Unstructured Grids," <u>Applied Numerical Mathematics</u> Volume 13, pages 437-452, February 1994.

Biswas, R., Devine, K. D. (RPI), and Flaherty, J. E. (RPI), "Parallel, Adaptive Finite Element Methods for Conservation Laws," <u>Applied Numerical Mathematics</u> Volume 14, pages 255-283, April 1994.

Strawn, R. C., (US Army AFDD) and *Biswas, R.*, "Computation of Helicopter Rotor Acoustics in Forward Flight," <u>Proceedings of the 19th Army Science Conference</u>, Orlando, FL, June 1994.

Biswas, R. and Strawn, R. C. (US Army AFDD), "Mesh Quality Assessment for Multiply-Refined Tetrahedral Grids," <u>Proceedings of the 14th IMACS World Congress</u>, Atlanta, GA, pages 580-583, July 1994.

Biswas, R. and Strawn, R. C. (US Army AFDD), "Dynamic Mesh Adaption for Tetrahedral Grids," <u>Proceedings of the 14th International Conference on Numerical Methods in Fluid Dynamics</u>, Bangalore, India, July 1994.

Strawn, R. C. (US Army AFDD) and *Biswas, R.*, "Numerical Simulations of Helicopter Aerodynamics and Acoustics," <u>Proceedings of the 6th International Congress on Computational and Applied Mathematics</u>, Leuven, Belgium, July 1994.

Strawn, R. C. (US Army AFDD), *Biswas, R.*, and Garceau, M. (Stanford), "Unstructured Adaptive Mesh Computations of Rotorcraft High-Speed Impulsive Noise, " to appear in <u>AIAA Journal of Aircraft</u>, 1994.

Strawn, R. C. (US Army AFDD), *Biswas, R.*, Lyrintzis, A. S. (Purdue), "Computation of Helicopter Blade-Vortex Interaction Noise with Kirchhoff Methods," accepted to 51st AHS Annual Forum, Fort Worth, TX.

Biswas, R. and Strawn, R. C. (US Army AFDD), "Anisotropic h-Refinement for Unstructured Hexahedral Grids," submitted to 3rd U.S. National Congress on Computational Mechanics, Dallas, TX.

Biswas, R.s and Strawn, R. C. (US Army AFDD), "Dynamic Mesh Adaption for Unstructured Hexahedral Grids, submitted to 12th AIAA Computational Fluid Dynamics Conference, San Diego, CA.

Duque, E. P. N. (US Army AFDD), Strawn, R. C. (US Army AFDD), and *Biswas, R.*, "A Solution Adaptive Structured/Unstructured Overset Grid Flow Solver with Applications to Helicopter Rotor Flows," submitted to 13th AIAA Applied Aerodynamics Conference, San Diego, CA.

SIDDHARTHA CHATTERJEE

Blelloch, G. E. (CMU), Hardwick, J. (CMU), Sipelstein, J. (CMU), M. Zagha(CMU), and *Chatterjee, S.*, "Implementation of a Portable Nested Data-Parallel Language," Journal of Parallel and Distributed Computing 21(1), pages 4-14, April 1994.

Chatterjee, S., "Programming Models, Compilers, and Algorithms for Irregular Data-Parallel Computations, "International Journal of High Speed Computing" 6(2), pages 183-222, June 1994.

Chatterjee, S., Gilbert, J. R. (Xerox), *Schreiber, R*, and *Sheffler, T. J.*, "Modeling Data-Parallel Programs with the Alignment-Distribution Graph," Journal of Programming Languages, Special issue on compiling and run-time issues for distributed address space machines, to appear.

Chatterjee, S., Gilbert, J. R. (Xerox PARC), Long, F. J. E. (UCSC), *Schreiber, R.*, and Teng, S.-H. (Xerox PARC), "Generating Local Addresses and Communication Sets for Data-Parallel Programs," Journal of Parallel and Distributed Computing, to appear, 1995.

Chatterjee, S., Gilbert, J. R. (Xerox PARC), *Schreiber, R.*, and Teng, S.-H. (Xerox PARC), "Optimal Evaluation of Array Expressions on Massively Parallel Machines," ACM Transactions on Programming Languages and Systems, to appear, 1995.

Chatterjee, S., Gilbert, J. R. (Xerox PARC), *Schreiber, R.*, and *Sheffler, T. J.*, "Array Distribution in Data-Parallel Programs," Languages and Compilers for Parallel Computing, Lecture Notes in Computer Science Series, Springer-Verlag, to appear 1995.

Chatterjee, S., "Locality, Communication, and Code Generation for Array-Parallel Languages," Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing (invited paper), San Francisco, CA, to appear February 1995.

Sheffler, T. J.,Schreiber, R., Gilbert, J. R. (Xerox PARC), and *Chatterjee, S.*, "Aligning Parallel Arrays to Reduce Communication," Proceedings of Frontiers'95, McLean, VA, to appear February 1995.

Sheffler, T. J. and *Chatterjee, S.*, "An Object-Oriented Approach to Nested Data Parallelism," Proceedings of Frontiers'95, McLean, VA, to appear February 1995.

JOSEPH OLIGER

Gustafsson, B., Kreiss, H.-O., Oliger, J.,"Time Dependent Problems and Difference Methods, to be published by J. Wiley and Sons in 1995.

THOMAS J. SHEFFLER

Chatterjee, S., Gilbert, J. R. (Xerox), *Schreiber, R*, and *Sheffler, T. J.*, "Modeling Data-Parallel Programs with the Alignment-Distribution Graph," Journal of Programming Languages, Special issue on compiling and run-time issues for distributed address space machines, to appear, 1994.

THOMAS J. SHEFFLER(cont'd)

Chatterjee, S., Gilbert, J. R. (Xerox), *Schreiber, R,* and *Sheffler, T. J.,* "Array Distribution in Data-Parallel Programs," <u>Proceedings of the Seventh Annual Workshop on Languages and Compilers for Parallelism,</u> pages 6.1–6.17, August, 1994.

Sheffler, T. J. and *Chatterjee, S.,* "Aligning Parallel Arrays to Reduce Communication, <u>Proceedings of Frontiers '95," the Fifth Symposium on the Frontiers of Massively Parallel Computation,</u> McLean, VA, February 1995.

Sheffler, T. J. and *Chatterjee, S.,* "An Object-Oriented Approach to Nested Data Parallelism," <u>Proceedings of Frontiers '95, the Fifth Symposium on the Frontiers of Massively Parallel Computation,</u> McLean, VA, February 1995.

WRAY BUNTINE

Buntine, W.L., "Operations for Learning with Graphical Models," <u>Journal of Artificial Intelligence Research</u> 2:159–225, December 1994.

Buntine, W.L., and Burckert, H.-J., "On Solving Equations and Disequations," <u>Journal of the ACM,</u> 1994.

Buntine, W.L., and Weigend, G. "Computing Second Derivatives in Feed-Forward Networks: a Review," <u>IEEE Transactions on Neural Networks,</u> 5(3), 9 pages, May, 1994.

Buntine, W.L., "Data Analysis With Graphical Models: Software Tools," <u>Computing Science and Statistics,Proceedings of the 26th Symposium at the Interface,</u> Raleigh, NC, 1994.

WEI-PAI TANG

Tang, W.-P. (Univ. of Waterloo), George, J.A., (Univ. of Waterloo), Ikramov, K and Tchugunov, V.N., (Univ. of Moscow), "On Doubly Symmetric Tridiagonal Forms for Complex Matrices," submitted to <u>SIAM J. Matrix Analysis and Applications,</u> 11 pages.

Tang, W.-P. (Univ. of Waterloo), George, J. A. (Univ. of Waterloo), Ikramov, K. and L. Matushkina, (Univ. of Moscow), "On a QR-like Algorithm for Some Structured Eigenvalue Problems," <u>SIAM J. Matrix Analysis and Applications,</u> 30 pages, accepted.

Tang, W.-P. (Univ. of Waterloo), George, A. Ikramov, Kh., Tchugunov, V., "On Doubly Symmetric Tridiagonal Forms for Complex Matrices,'" submitted to SIAM J. Matrix Analysis and Applications, 11 pages.

Fan, Q., Forsyth,P. A., Univ. of Waterloo, McMacken, J. R. F. (Goal Electronics Inc.), *Tang, W.-P.* (Univ. of Waterloo), "Performance Issues for Iterative Solvers in Device Simulation," RIACS 94.21, submitted to <u>SIAM J. Scientific Computing,</u> 19 pages.

Tang, W.-P. (Univ. of Waterloo), S. Clift, Queens University, "Weighted Graph Based Ordering Techniques for Preconditioned Conjugate Gradient Methods" RIACS TR94.20, accepted <u>BIT,</u> 18 pages.

DAN FENG

Byrd, R.H., *Feng, D.* and Schnabel, R.B., "On Optimality Conditions for Singular Constrained Optimization," to be submitted, 1994.

Feng, D. and Schnabel, R.B., "Tensor Methods for Equality Constrained Optimization," Report CU-CS-729-94, Department of Computer Science, University of Colorado, Boulder, Colorado, 1994.

MARJORY J. JOHNSON

Johnson, M.J., Johnston, W., Swinehart, D., "Bay Area Gigabit Testbed (BAGNet)," WWW document, http://george.lbl.gov/BAGNet.html

Johnson, M., Johnston, W., Swinehart, D., ``Bay Area Gigabit Testbed (BAGNet) - Project Descriptions," WWW document, http://george.lbl.gov/BAGNet.html

ROBERT SCHREIBER

Chatterjee, S., Gilbert, J. R. (Xerox), *Schreiber, R.,* and Sheffler, T. J., "Modeling Data-Parallel Programs with the Alignment-Distribution Graph," Journal of Programming Languages, Special issue on compiling and run-time issues for distributed address space machines, to appear.

Chatterjee, S., Gilbert, J. R. (Xerox), *Schreiber, R.,* and Sheffler, T. J., "Array Distribution in Data-Parallel Programs," Proceedings of the Seventh Annual Workshop on Languages and Compilers for Parallelism, pages 6.1–6.17, August, 1994.

Bjorstad, P.E, and *Schreiber, R.,* "Unstructured Grids on SIMD Torus Machines," Proceedings of the 1994 Scalable High Performance Computing Conference, pages 658-665.

Chatterjee, S., Gilbert, J. R. (Xerox), *Schreiber, R,* "The Alignment-Distribution Graph," Proceedings of the Sixth Annual Workshop on Languages and Compilers for Parallelism, Springer-Verlag, 1994.

Rothberg, E. and *Schreiber, R.,* "Improved Load Distribution in Parallel Sparse Cholesky Factorization," Proceedings of Supercomputing 94, pages 783–792.

Sheffler, T. J., Oliker, L., Gilbert, J. R. (Xerox), Chatterjee, S., Bau, D., and *Schreiber, R.,* "Data Parallel Array Allocation for Reducing Communication Cost," submitted to the 1994 ACM PLDI Symposium, 1994.

Koelbel, C.H., Loveman, D.B., Steele, Jr., G. L., Zosel M.E., and *Schreiber, R.,* "The High Performance Fortran Handbook," The MIT Press, 1994.

Chatterjee, S., Gilbert, J. R. (Xerox), *Sheffler, T. J.,* and *Schreiber, R.,* "Automatic Distribution in HPF," Proceedings of the Second Workshop on Environments and Tools for Parallel Scientific Computing, SIAM, 1994.

ROBERT SCHREIBER (cont'd)

Rothberg, E. and *Schreiber, R.*, "Efficient Parallel Sparse Cholesky Factorization,"Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing, SIAM, 1995.

PELLE OLSSON

Olsson, P., "High-Order Finite Difference Methods on Non-Smooth Domains," Computer Methods in Applied Mechanics and Engineering 116, 265-272, 1994.

Olsson, P., "The Numerical Behavior of High-Order Finite Difference Methods," Journal of Scientific Computing 9, 445-466, 1994.

Olsson, P., "Summation by Parts, Projections, and Stability I, TR 93.04," Mathematics of Computations, 31 pages, to appear in Mathematics of Computation, October 1995.

Olsson, P., "Summation by Parts, Projections, and Stability II, TR 93.04," Mathematics of Computations, 20 pages, to appear in Mathematics of Computation, July 1995.

Gottlieb, D. (Brown University), Gustafsson, B. (Uppsala University), *Olsson, P.*, Strand, B. (Uppsala University)," On the Superconvergence of Galerkin Methods for Hyperbolic IBVP,"TR 93.07, 20 pages, to appear in SIAM Journal of Numerical Analysis, October 1996.

Gustafsson, B. (Uppsala University), *Olsson, P.*, "Fourth-Order Difference Methods for Hyperbolic IBVP, TR 94.04," to appear in Journal of Computational Physics, 35 pages.

Olsson, P., Oliger, J., "Energy and Maximum Norm Estimates for Nonlinear Conservation Laws, TR 94.01," submitted to SIAM Review, 16 pages.

Gustafsson, B.(Uppsala University), *Olsson, P.*, "High-Order Centered Difference Methods with Sharp Shock Resolution, TR 94.07," submitted to SIAM Journal on Numerical Analysis, 25 pages.

BARNEY PELL

Pell, B., "A Strategic Metagame Player for General Chess-Like Games," Proceedings of the Twelfth National Conference on Artificial Intelligence, Volume 2, AAAI/MIT Press, pages 1378-1385, 1994.

Pell, B., "A Strategic Metagame Player for General Chess-Like Games, accepted for publication in Computational Intelligence Journal, 1995.

Matsubara, H. (ETL, Japan) and *Pell, B.* "Applying Metagamer to Shogi," in H. Matsubara, Editor, Game Programming Workshop in Japan'94, Japanese Computer Shogi Association, Kanagawa, Japan, October, 1994. (In Japanese).

Pell, B., "Logic Programming for Efficient General Game-Playing," submitted to the 1995 International Joint Conference on Artificial Intelligence, 1994.

BARNEY PELL (cont'd)

N. Muscettola (RECOM) and *Pell, B.,* "Toward Real-World Science Mission Planning," presented at the AAAI Fall Symposium on Planning and Learning: On to Real Applications, New Orleans, LA, Nov 4-6, 1994, in preparation as an IC technical report.

*Pell, B.*and S. Kedar (ILS), "Using Induction to Refine Information Retrieval Strategies," Baudin, C. (Recom), Proceedings of the Twelfth National Conference on Artificial Intelligence, Volume 1, AAAI/MIT Press, pages 553-559, 1994.

Baudin, C. (Recom), Kedar, S. (ILS), and *Pell, B.,* "Increasing Levels of Assistance in Refinement of Knowledge-Based Retrieval Systems," Knowledge Acquisition Journal Volume 6, Number 2, pages 179-196, 1994. Also appears as RIACS Technical Report 94.03, also in preparation as a book chapter, 1994.

C. Baudin (Recom), *Pell, B.* and S. Kedar (ILS)," Incremental Acquisition of Conceptual Indices for Multimedia Design Documentation," presented at the Workshop on Indexing and Reuse in Multimedia Systems, at the Twelfth National Conference on Artificial Intelligence, Seattle, WA, pages 157-164, 1994.

JAMES REUTHER/ANTONY JAMESON

Reuther, J. and Jameson, A., "Control Theory Based Airfoil Design for Potential Flow and a Finite Volume Discretization," AIAA paper, 94-0499, 32nd Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 1994.

Jameson, A. and *Reuther, J.,* "Control Theory Based Airfoil Design Using the Euler Equations," AIAA paper, 94-4272, 5th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization," Panama City Beach, FL, September 1994.

Jameson, A., "Optimum Aerodynamic Design via Boundary Control," Proceedings of the AGARD-VKI Lecture Series on Optimum Design Methods in Aerodynamics, von Karman Institute for Fluid Dynamics, 1994.

Jameson, A., "Aerodynamic Design Methods," Proceedings of the International Workshop on Solution Techniques for Large-Scale CFD Problems, CERCA, Montreal, September 1994.

Reuther, J. and Jameson, A., "Aerodynamic Shape Optimization of Wing and Wing-Body Configurations Using Control Theory," 33rd Aerospace Sciences Meeting and Exhibit, Reno, Nevada, to appear AIAA paper, 95-0123, January 1995.

V. SEMINARS AND COLLOQUIA

RUPAK BISWAS

Mesh Quality Assessment for Multiply-Refined Tetrahedral Grids, *R. Biswas*, 14th IMACS World Congress, Atlanta, GA, July 11, 1994.

Numerical Simulations of Helicopter Aerodynamics and Acoustics, *R. Biswas*, 6th International Congress on Computational and Applied Mathematics, Leuven, Belgium, July 27, 1994.

WRAY BUNTINE

Compiling Learning Algorithms, Seminar given at Seimens Corporate Research, *Buntine, W.*, Princeton, January 1994.

Learning Rules: Automating the Process, *Buntine, W.*, invited talk at ML-94 tutorial, State of the Art in Learning DNF Rules, July, 1994.

Compiling Learning Algorithms, *Buntine, W.*, Seminar given at The Real World Computing Center, Tsukuba Center, Japan, September, 1994.

Tools for Prototyping Networks, *Buntine, W.*, presentation at a special software session in Neural Networks for the Capital Markets, Pasadena, CA, November, 1994.

SIDDHARTHA CHATTERJEE

Optimizing Data Layout on Distributed-Memory Parallel Machines, *Chatterjee, S.*, PS Seminar, Carnegie Mellon University, Pittsburgh, PA, 17 January 1994;

Seminar, IBM T. J. Watson Research Center, Hawthorne, NY, 3 Februrary 1994.

Seminar, Digital Equipment Corporation, Maynard, MA, 7 February 1994.

Colloquium, Department of Computer Science, University of Minnesota, Minneapolis, MN, 25 February 1994.

Seminar, Department of Computer Science and Engineering, Southern Methodist University, Dallas, TX, 4 March 1994.

Seminar, Department of Computer Science, The Johns Hopkins University, Baltimore, MD, 8 March 1994

Seminar, Department of Computer Science, George Mason University, Fairfax, VA, 9 March 1994.

SIDDHARTHA CHATTERJEE (cont'd)

Seminar, Department of Computer Science, University of Maryland, College Park, MD, 10 March 1994.

Seminar, Department of Computer Science, Purdue University, West Lafayette, IN, 28 March 1994.

Colloquium, Department of Mathematics and Computer Science, DartmouthCollege, Hanover, NH, 29 March 1994.

Seminar, Department of Computer Science, The University of North Carolina, Chapel Hill, NC, 31 March 1994.

Seminar, Department of Electrical Engineering, Santa Clara University, Santa Clara, CA, 7 April 1994.

DAN FENG

Tensor Methods for Nonlinear Equality Constrained Optimization, Feng, D., 15th International Symposium on Mathematical Programming, Ann Arbor, Michigan, August 17, 1994.

Tensor-GMRES Method for Large Sparse Systems of Nonlinear Equations, Feng, D., Computational Algorithms and Applications Branch (RFC Branch), NASA Ames Research Center, Moffett Field, California, July 20, 1994.

JOSEPH OLIGER

Stability and Error Estimation for Component Adaptive Grid Methods, Oliger, J., in a special session on Adaptive Mesh Refinement for CFD Applications, 14th IMACS World Congress, Georgia Tech, July 11-15, 1994.

PELLE OLSSON

A Survey of High-Order Finite Difference Methods for Hyperbolic and Parabolic PDE's, Olsson, P., Bell Laboratories, AT&T, Murray Hill, NJ, (Mar. 1994).

Energy Estimates for Nonlinear Conservation Laws, RFC Seminar, Olsson, P., NASA Ames Research Center, Moffett Field, CA, (May 1994).

Energy and Maximum Norm Estimates for Nonlinear Conservation Laws, Olsson, P., 5th International Conference on Hyperbolic Problems, Stony Brook, NY, (Jun. 1994).

BARNEY PELL

Toward Real-World Science Mission Planning, AAAI Fall Symposium on Planning and Learning: On to Real Applications, New Orleans, LA, Nov 4-6, 1994.

A Strategic Metagame-Player for General Chess-Like Games, Twelfth National Conference on Artificial Intelligence, Seattle, WA, July 31-Aug 4, 1994; Computer Science Departmental Seminar, U.C. Davis, Nov. 17, 1994; Computer Science Department, U.C. Santa Cruz, Nov. 18, 1994.

Computer Game-Playing, B. Pell, Guest Lecture for a graduate AI course, Computer Science Dept.; Berkeley, Feb. 8, 1994.

Using Induction to Refine Information Retrieval Strategies, Twelfth National Conference on Artificial Intelligence, Seattle, WA, July 31-Aug 4, 1994.

Incremental Acquisition of Conceptual Indices for Multimedia Design Documentation, AAAI-94 Workshop on Indexing and Reuse in Multimedia Systems, Seattle, WA, July 31-Aug 4, 1994.

How Symbolic Systems prepared me for a Career in AI Research, B. Pell, Symbolic Systems Forum, Stanford University, February 3, 1994.

Modeling Aggregation Behavior in Autonomous Agents, B. Pell, Jan 27, 1994, NASA Ames; BOTS seminar, Computer Science Department, Stanford University, Feb., 1994.

ROBERT S. SCHREIBER

Language and Compiler Grand Challenges in Data-Parallel Programming, Schreiber, R., Mardi-Gras Conference on Teraflop Computing, Baton Rouge, LA, February, 1994.

Alignment and Distribution of Arrays in Data-parallel Programs, Schreiber, R.,Workshop on Environments and Tools For Parallel Scientific Computing, Blackberry Farm, Tennessee, May, 1994.

Parallel Matrix Computation, Invited lecture, Schreiber, R.,Fifth SIAM Conference on Applied Linear Algebra, June, 1994.

Scalable Sparse Cholesky Factorization." CERFACS "Sparse Days, Schreiber, R., Conference, St. Girons, France, July, 1994.

High Performance Fortran, Past, Present, Future, Schreiber, R., Intel Paragon Workshop, University of Bergen, October, 1994.

THOMAS J. SHEFFLER

Aligning Parallel Arrays to Reduce Communication, *Sheffler, T.J.*, in a poster session at Supercomputing '94 in Washington DC, December 1994.

Presented "Array Distribution in Data-Parallel Programs" *Sheffler, T.J.*,at the Seventh Annual Workshop on Languages and Compilers for Parallelism at Cornell, August 1994.

Presented a Seminar on Array Distribution, *Sheffler, T.J.*, Stanford University, September 1994.

ROBERT B. SCHNABLE

An Informal Discussion of Optimization Research, Robert B. Schnabel (Department of Computer Science University of Colorado at Boulder), January 18, 1994

WEIJIA SHANG

Two Fundamental Problems in Programming Distributed Memory Parallel Computers, Weijia Shang (Santa Clara University), March 15, 1994

IAIN DUFF

The Solution of Sparse Equations on High Performance Computers, Iain Duff (Rutherford Appleton Laboratories, Oxon, United Kingdom, April 13, 1994.

BRIAN TOTTY

Tunable Shared-Memory Abstractions for Distributed Data Structures, Brian Totty (University of Illinois), April 15, 1994.

VADIM MASLOV

Exact Array Dataflow Dependence Analysis: Why and How, Vadim Maslov (University of Maryland), April 18, 1994.

DAN SCALES

An Efficient Shared Memory Layer for Distributed Memory Machines, Dan Scales (Stanford University), May 3, 1994.

JUSTIN PAOLA
Comparisons of Neural Networks to Standard Techniques for Image Classification and Correlation, Justin Paola (University of Arizona), July 26, 1994.

VI.   OTHER ACTIVITIES

RUPAK BISWAS

Chaired session on CFD Algorithms for High-Performance Computer Architectures, 25th AIAA Fluid Dynamics Conference, Colorado Springs, CO, June 20-23, 1994.

Organized and chaired session on Adaptive Mesh Refinement Methods for CFD Applications, 14th IMACS World Congress, Atlanta, GA, July 11-15, 1994.

Attended and presented a paper at the 14th IMACS World Congress, Atlanta, GA, July 11-15, 1994.

Attended and presented a paper at the 6th International Congress on Computational and Applied Mathematics, Leuven, Belgium, July 25-29, 1994.

Interview article in USRA Quarterly, Winter-Spring, 1994.

WRAY BUNTINE

Learning rules: automating the process, Invited talk at ML-94 tutorial, State of the Art in Learning DNF Rules, July, 1994.

Learning with Graphical Models, Invited tutorial given at, Probabilities in Artificial Intelligence Summer School, Oregon, July, 1994.

Learning From Data: A Probabilistic Prospective, Invited tutorial co-taught with Padhriac Smyth (JPL) and given at the, National Conference on Artificial Intelligence (AAAI-94)}, Seattle, July, 1994.

Learning with Probabilities, Invited tutorial given at the 1994 {\it European Machine Learning Summer School, Paris, September, 1994.

Joined the editorial board for the Machine Learning Journal and the Journal of Artificial Intelligence Research.

Program Chair for SPIE's Applications of AI Workshop, SPIE-AAI-94; held in Orlando, April, 1994.

On program committee for European Conference on Machine Learning, held in Bonn, Germany, April, 1994.

On program committee and attended 11th International Machine Learning Conference, held in New Brunswick, New Jersey, July 10-13, 1994.

On program committee for 4th International Workshop on Inductive Logic Programming, held in Bonn, Germany, 12-14 September, 1994.

WRAY BUNTINE (cont'd)

On program committee and attended 10th Uncertainty in Artificial Intelligence Conference, held in Seattle, July 19-21, 1994.

Reviewer for 8th Annual Conference on Neural Information Processing Systems, held in Denver, 28 November to 1 December, 1994.

SIDDHARTHA CHATTERJEE

Joint USRA/Xerox patent application no. 08/217,404, title Generating Local Addresses and Communication Sets for Data-Parallel Programs, filed March 24, 1994.

Taught full-day tutorial titled Programming Languages for High Performance Computing at SHPCC'94, Knoxville, TN, May 1994.

DAN FENG

Visitor, ROBERT SCHNABEL (visiting Dan Feng). An Informal Discussion of Optimization Research, NASA Ames Research Center, Moffett Field, California, January 18, 1994.

MARJORY J. JOHNSON

Chaired session at 5th IFIP International Conference on High Performance Networking, June, 1994, France.

Led discussion on applications for high-performance networks at 4th IFIP International Workshop on Protocols for High-Speed Networks, August, 1994, Canada.

Served on program committee for 4th IFIP International Workshop on Protocols for High-Speed Networks, August, 1994, Canada.

Served on program committee for Hot Interconnects II, August, 1994, Palo Alto, CA.

Attended two gigabit testbed workshops.

Served on review committee for a DoE collaboratory program.

JOSEPH OLIGER

Professor of Computer Science, Stanford University.

Member, Board of Trustees of the Society for Industrial and Applied Mathematics.

Associate Editor, Journal of Computational and Applied Mathematics.

Editorial Advisory Committee, SIAM Review.

BARNEY PELL

Guest Editor, Special Issue on Learning and Planning in Games, Computational Intelligence Journal, to appear in 1995.

With Jens Krause (Princeton University), Pell is developing a computational model of emergent aggregation behavior in animals (e.g., schooling, flocking, herding).

Visited Cambridge University, UK for 3 weeks as visiting scientist, June 1994.

Attended the AAAI-94 Workshop on Comparative Analysis of Planning Systems, Seattle, WA, July 31-Aug 4, 1994.

Attended the 1994 Summer Institute on Probabilistic Reasoning in AI, Corvallis, Oregon July 22-27, 1994.

Attended the AAAI Spring Symposium on (a) Software Agents and (b) Believable Agents, Stanford CA, March 16-19, 1994.

JAMES REUTHER

Continued work in the development of the NASA Ames Optimization Group which I formed last year. Planned seminars and led discussions on optimization topics at Ames.

Worked with Code AAH in the High Speed Research Project to apply optimization techniques in the practical development of the next generation High Speed Civil Transport.

Participated in three wind tunnel test to support the HSR project that validated optimized designs.

Participated in dozens of tele-conferences with the Configuration Aerodynamics team for the HSR project to discuss various aspects of optimized configurations and determine future technology development.

ROBERT S. SCHREIBER

Conference organizer, Seventh SIAM Meeting on Parallel Processing for Scientific Computing, 1995.

Organizer, User-interface Workshop on Parallel Programming, Santa Barbara, CA, 1995

SIGNUM Vice-chair.

Program Committee, Scalable High Performance Computing Conference.

Program Committee, ACM Symposium on Parallel Algorithms and Architectures (SPAA).

Organizing Committee of the IMA Year on High-Performance Computing, IMA, Univ. of Minnesota.

Co-instructor in tutorial of Scalable Parallel Algorithms for Multicomputers, Scalable High Performance Computing Conference.

THOMAS J. SHEFFLER

Attended the Symposium on Parallel Algorithms and Architectures, Sheffler, T.J.,Cape NJ, June 1994.

STEVEN SUHR

Contributed a new scrolling algorithm which has been accepted for incorporation into a future version of the text editor GNU emacs, improving the visual stability of the approach which emacs uses in updating the displayed text on the screen.

VII. RIACS STAFF

ADMINISTRATIVE STAFF

Joseph Oliger, Director - Ph.D., Computer Science, University of Uppsala, Sweden, 1973. Numerical Methods for Partial Differential Equations (03/25/91 - present).

Frances B. Abel, Office and Financial Manager (5/5/88 - present).

Deanna M. Gearhart, Administrative Assistant II (5/9/88 - present).

Rufus White Jr., Systems Administrator (5/17/93 - present).

RIACS SCIENCE COUNCIL

Dr. Robert B. Schnabel (Convener), Department of Computer Science, University of Colorado, Boulder, Colorado 80309-0430.

Dr. James W. Demmel, Computer Science Division, 571 Evans Hall, University of California, Berkeley, California 94720.

Dr. Joseph Oliger (Ex-Officio), Director, Research Institute for Advanced Computer Science, NASA Ames Research Center, Moffett Field, California 94035-1000.

Dr. Kenneth W. Neves, Boeing Company, P.O. Box 24346, MS 7L-25, Seattle, Washington 98124.

Dr. Daniel A. Reed, Department of Computer Science, 1304 West Springfield Avenue, University of Illinois, Urbana, Illinois 61801-2987.

Dr. Marc Snir, IBM, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598.

Dr. David Cummings, Executive Director, Universities Space Research Association, Columbia, Maryland 21044.

SCIENTIFIC STAFF

Rupak Biswas, Ph.D., Computer Science, Rensselaer Polytechnic Institute, 1991, Large scale scientific computation using parallel and adaptive methods (9/16/91 - present).

Wray Buntine, Ph.D., Computer Science, University of Technology, Sydney, Australia, 1992, Mathematical and probabilistic modeling of problems in intelligent systems (2/5/90 - present).

Peter Cheeseman, Ph.D., Artificial Intelligence, Monash University, Australia 1979, Artificial intelligence and automatic control, induction of models under uncertainty, Bayesian inference, expert systems and robotics (5/6/85 - present).

Dave Gehrt, JD Law, University of Washington, 1972, UNIX system administration, security, and network based tools (1/84 - 7/85, 2/1/88 - present).

Marjory J. Johnson, Ph.D., Mathematics, University of Iowa, 1970, High-performance networking for both space and ground applications (1/9/84 -present).

Michael R. Raugh, Ph.D. - Mathematics, Stanford University, 1977, Mathematics and computers for modeling physical and biological systems (1/28/85 - 12/30/94).

Robert S. Schreiber, Ph.D. - Computer Science, Yale University, 1977, Parallel numerical algorithms and parallel computer architectures (8/29/88 - present).

VISITING SCIENTISTS

Remi Abgrall, Ph.D. - Research Scientist, University of California, Los Angeles, Description of the multi-resolution analysis of A. Harten (7/18/94 - 7/22/94).

Marsha Berger, Ph.D. - New York University, Computational fluid dynamics; parallel computing (6/27/94 - 8/31/94).

Peter Bradshaw, Ph.D. - Professor, Stanford University, Analysis of experimental data for turbulent flow over a backward-facing step (7/1/94 - 7/31/94).

Bertil Gustafsson, Ph.D. - Professor, Upsalla University, Sweden, Numerical methods for partial differential equations. Computational fluid dynamics (7/1/93 - 6/30/94).

Ami Harten, Ph.D. - Professor, Tel-Aviv University, Numerical solutions of PDES/multi-resolution analysis (6/27/94 - 7/31/94).

Antony Jameson, Ph.D. - McDonnell Professor of Aerospace Engineering, Princeton University, Numerical Methods, computational fluid dynamics, computational sciences (9/7/93 - 8/31/94).

Ki Dong Lee, Ph.D. - Associate Professor, University of Illinois, Computational aerodynamics, design optimization, grid generation and adaptation (7/1/94 - 7/23/94).

VISITING SCIENTISTS (cont'd)

Luigi Matinelli, Ph.D. - Assistant Professor, Princeton University, Computational methods for Navier-Stokes EQ Free Surface flows aerodynamic Des (8/1/94 - 8/27/94).

Kenneth G. Powell, Ph.D - Associate Professor, University of Michigan, Computational Aerodynamics (5/30/94 - 7/29/94).

Romanujain Jaganatham, Ph.D. - Assistant Professor, Louisiana State University, Compiling for parallel machines; parallel computing; programming environment (7/18/94 - 8/1/94).

Mohammad Saleem, Ph.D. - Professor, San Jose State University, Numerical Analysis: Acceleration-Convergence for finite difference formulation of Navier Stokes equations, Numerical linear Algebra: Krylov suspace methods (5/30/94 - 7/26/94).

Robert Schowengerdt, Ph.D. - Professor, University of Arizona, Image processing for earth remote sensing/Pattern recognition in large image databases/remote sensor modeling and simulation (6/1/94 -6/30/94).

Peter Sweby, Ph.D. - Lecturer in Mathematics, University of Reading, England, The dynamics of discretisations of convection reaction diffusion equations was investigated over a range of physical and numerical parameters (7/11/94 - 8/22/94).

Eitan Tadmor, Ph.D. - Visiting Professor, University of California, Los Angeles, Numerical solution of non-linear time dependent problems (8/15/94 - 8/16/94)

David Zingg, Ph.D. - Associate Professor, University of Toronto, Canada, Development and analysis of high-accuracy numerical methods applicable to simulations of fluid Flows, Acoustic Waves, and Electromagnetic Waves (6/27/94 - 9/2/94).

POST-DOCTORAL SCIENTISTS

Siddhartha Chatterjee, Ph.D. - Computer Science, Carnegie Mellon, Compilation for distributed-memory parallel machines; parallel algorithms and applications (11/1/91 - 10/7/94).

Dan Feng, Ph.D. - Computer Science, University of Colorado, Boulder, Numerical computation including optimization, solving systems of nonlinear equations, solving PDE's and their applications in aeronautics and astronautics; parallel numerical computation (9/1/93 - present).

Pelle Olsson, Ph.D. - Uppsala University, Sweden, Initial-boundary value problems for hyperbolic and parabolic PDEs, numerical methods for PDEs on parallel computers (11/2/92 - present).

POST-DOCTORAL SCIENTISTS(cont'd)

Barney Pell, Ph.D. - University of Cambridge, England, Artificial intelligence, machine lLearning, strategic reasoning, logic programming, automatic scheduling, information retrieval, multiple autonomous agents, artificial life (9/27/93 - present).

Thomas Scheffler, Ph.D. - Carnegie Mellon, Computer engineering, Parallel compiler techniques (6/11/93 - present).

RESEARCH ASSOCIATES

Leonid Oliker - Computer Science, University of Colorado, compilation of data parallel programs (9/1/94 - present).

James Reuther, Ph.D. - University of California, Davis, numerical optimization aerodynamic shape optimization numerical analysis CFD (9/6/94 - present).

Steven Suhr, - Research Associate, Computer Science, Stanford University, programming languages (7/1/92 - present).

STUDENTS

Marc Bumble - Computer Engineering, Penn State University, cellular satellite communications routing schemes, algorithms, and models (5/23/94 - 8/19/94).

Frederic Chong - Computer Science, Massachusetts Institute of Technology, parallel systems, irregular applications (8/29/94 - 9/30/94).

Patricia Hough - Applied Math, Cornell University, numerical linear algebra scientific computing (5/23/94 - 8/19/94).

Justin Paola - Electrical Engineering & Computer Science, UC Berkeley, (1990), image processing/remote sensing, University of Arizona (5/31/94 - 8/3/94).

Sabine Vermeersch - Scientific Computing & Computational Mathematics, Stanford University, Tensor -Krylov methods for systems of non-linear equations and applications for PSE's arising aerodynamics (6/20/94 - 8/26/94).

Xiaolei Zhu - Scientific Computing & Computational Mathematics, Stanford University, error analysis for adaptive grids (6/1/94 - 8/30/94).

CONSULTANTS

Tony F. Chan - Professor of Mathematics, University of California, Los Angeles, Efficient algorithms in large-scale scientific computing, parallel algorithms and computational fluid dynamics (10/01/86 - present).

John Gilbert - Research Scientist, Xerox Palo Alto Research Center, Parallel computing and theoretical computer science (5/1/92 - present).

Richard G. Johnson, Ph.D. - Physics, Indiana University, 1956, Global environmental problems and issues (11/1/92 - present).

Niel K. Madsen - Lawrence Livermore National Laboratory. Numerical solutions of partial differential equations, with specific interests in method of lines techniques, PDE soft ware, matrix algorithms for vector and parallel computers (10/18/90 - 12/31/94).

Robert Schnabel, Ph.D. - Professor, University of Colorado, Boulder, Numerical computation especially optimization nonlinear equations, parallel computation (1/1/94 - present).

Wei-Pai Tang, Ph.D. - Professor, University of Waterloo, Canada, numerical solution of partial differential equations, numerical linear algebra, parallel computations (7/1/94 - present).