NASA-CR-195117

RIACS

Research Institute for Advanced Computer Science
NASA Ames Research Center

NASA
IN-99-CR
204352
40P

Annual Report

1993

(NASA-CR-195117) RIACS/USRA Annual Report, 1993 (Research Inst. for Advanced Computer Science) 40 p

N94-25003

Unclas

G3/99 0204352

RIACS / USRA

RIACS
Annual Report
1993

Contract NAS 2-13721

Submitted to:

Contracting Officer: Dr. Michael J. Hutnik
Code RN, MS 241-1

CO, Code ASR, MS 241-1

External Relations Office
Technology Utilization, Code DXR, MS 223-3

Ames Research Patent Counsel
Code DP, MS 200-11A

Submitted by:
Research Institute for Advanced Computer Science (RIACS)

An Institute of:
Universities Space Research Association (USRA)

RIACS Principal Investigator
Joseph Oliger

NASA Technical Monitor
Frank Ron Bailey

NASA Alternate Technical Monitor
Masayuki Omura

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 and Technology
- 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 six months Professor Bertil Gustafsson of the University of Uppsala and Professor Antony Jameson of Princeton University have begun visits to RIACS. Dr. Dan Feng from the University of Coloradohas joined RIACS as a post doctoral scientist. He is working on numerical methods for optimization. Dr. Thomas Sheffler from Carnegie Mellon University has joined RIACS as a post doctoral scientist working on parallel compiler techniques.

We had 5 visiting graduate students this year that worked on projects over the summer.

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, 1993 through December 31, 1993 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 Language Definition was completed in May, 1993. The High Performance Fortran Handbook, by Koelbel, Loveman, Schreiber, Steele, and Zosel was published by MIT Press in the fall. HPF is a language based on Fortran 90 but with significant extensions for data parallel computation, for programming parallel, distributed memory machines. It is expected to be adopted widely in the parallel computing industry.

RESEARCH IN HIGHLY PARALLEL MATRIX COMPUTATION
Robert S. Schreiber

With Petter Bjorstad, a software message routing system for grid and toroidal machines has been developed in order to assess their usability for solving PDEs with unstructured meshes. By June, the effort had resulted in experimental data that showed clearly that fine-grained, highly parallel meshes of processors could be efficient for sparse matrix computation. This work has been presented at a number of meetings and seminars, and a paper is in progress.

With Yan Huo, (Princeton, Dept. of Applied Physics) an efficient, massively parallel implementation of some well known methods for the symmetric eigenvalue problem was completed. What was not known was whether these methods, tridiagonal reduction, multisection, and inverse iteration, would be efficient on such machines. These results indicate that this is the method of choice for parallel computation of the symmetric, dense eigen problem.

With Ed Rothberg (Intel Scientific Computers) an improved implementation of distributed memory sparse Cholesky was developed. Like Rothberg's earlier codes, it uses a two-dimensional mapping of matrix blocks. (Schreiber's earlier work with Gilbert on scalability of sparse solvers shows that this approach is theoretically efficient while column-mapped methods cannot be). Rothberg had discovered that his methods suffered 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.

Taken together with recent work of Robert Lucas (Supercomputing Research Center), who has developed an efficient, massively parallel SIMD sparse Cholesky along lines mapped earlier by Gilbert and Schreiber, this work has shown for the first time that highly parallel, distributed memory machines are useful for sparse direct solvers.

RESEARCH IN COMPILER OPTIMIZATION

Robert S. Schreiber

A paper is in preparation with Jaganatham Ramanujam (LSU, Dept. of Computer Science) on a new loop transformation technique that allows considerable simplification of the theory of Wolfe of optimizing loop transformations and also, for the first time, allows compilers to generate blocked versions of codes as complicated as gaussian elimination with partial pivoting.

COMPILING ARRAY PARALLELISM FOR DISTRIBUTED-MEMORY MACHINES

Robert S. Schreiber, Siddhartha Chatterjee, Thomas J. Sheffler
John R. Gilbert (Xerox PARC), Leonid Oliker (U. of Colorado at Boulder; summer intern), Jingke Li (Portland State University; summer visitor), David Bau (Cornell; Xerox PARC summer intern), Shang-Hua Teng (MIT; Xerox PARC summer visitor), Erik Boman (Stanford)

Work continues on developing compiler algorithms to automate the task of data layout in the context of compiling array-parallel languages such as Fortran 90 on distributed-memory parallel computers. We have made significant progress in both the theoretical and practical sides of the problem.

The various components of our theory of array alignment have been unified by developing the Alignment-Distribution Graph (ADG), an intermediate representation of array-parallel programs. The ADG closely resembles other program representations such as Static Single Assignment form and the Program Dependence Web, but is tailored for the analysis of data layout. It uses novel techniques to represent the effects of control flow and mobile alignments dependent on loop induction variables. A joint USRA/Xerox patent on the representation has been filed.

We have developed a new algorithm for axis and stride alignment based on the constraint graph method. This algorithm is more robust than our previous technique of compact dynamic programming. We have proved the NP-completeness of several related problems, and are investigating possible extensions of this technique to offset alignment. We have also developed an algorithm based on network flow to determine array replication. Finally, we have made some progress in determining program phases - portions of a program that should be distributed in the same way.

We began the work of implementing the data structures and algorithms this summer. We have implemented the ADG data structures and algorithms for axis, stride, offset, and replication alignment. We have a partial front end to the system and are experimenting with benchmark codes to quantify the effectiveness of our techniques.

NESTED PARALLELISM FOR UNSTRUCTURED PROBLEMS

Thomas J. Sheffler, Siddhartha Chatterjee

While array-oriented languages are suitable for problems with regular structure, there is little consensus on what is appropriate programming language support for parallel algorithms on problems with sparse or irregular structure. We believe that many of these types of applications can be effectively handled in a data-parallel language supporting nested parallelism. These ideas originated in Blelloch's work on the NESL programming language and in the dissertations of Sheffler and Chatterjee. Our approach adds a single parallel collection type (the vector) and one new language construct (a parallel FOREACH loop) to an object-oriented language (C++). The result is an object-oriented language that supports parallel operations on distributed collections of arbitrary data types.

The viability of the approach has been demonstrated in a proof-of-concept implementation that properly executes programs exhibiting nested parallelism. Rather than requiring development of a complex compiler, we have pushed the complexity of the language implementation into a sophisticated class library and runtime system. These make use of several advanced features of C++, including inheritance, operator and function overloading, and class templates. This approach will allow us to construct an initial compilation system consisting of a translator from our language into standard C++ that can be compiled using any of a number of commercially available compilers.

PARALLEL COPUTERS AND THREE-DIMENSIONAL COMPUTATIONAL ELECTROMAGNETICS

Niel K. Madsen

We have continued to enhance our ability to use new massively parallel processing (MPP) computers to solve time-domain electromagnetics problems. New vectorization techniques have improved the performance of our code DSI3D by factors of 5 to 15 depending upon the computer used. New radiation boundary conditions and far-field transformations now allow us to compute radar cross section (RCS) values for complex objects. A new parallel data extraction code has been developed which allows the extraction of data subsets from large problems which have been run on parallel computers for subsequent post-processing on workstations with enhanced graphics capabilities.

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 RFG and of U.S. Army AFDD/ATCOM, continues on developing efficient and robust techniques for the fast, dynamic refinement and coarsening of three-dimensional unstructured grids. The method has been successfully implemented on the Cray C-90 and used to solve steady-state problems in helicopter aerodynamics and acoustics. We are currently in the process of porting the procedure to distributed-memory parallel computers like the CM-5 that will enable us to solve realistic forward-flight problems.

This work aims to solve problems in rotary-wing aerodynamics where wake systems and acoustic waves are extremely complex. Unlike a fixed wing, where the wake is rapidly convected downstream, a rotor operates in close proximity to its own wake. The key to accurately modeling these flow features with computational fluid dynamics techniques is to minimize any numerical dissipation that can artificially diffuse the rotor wakes and acoustic signals. Our solution-adaptive approach adds and removes points locally to provide higher resolution for moving flow features such as rotor-wake vortices and acoustic waves. This local increase in grid resolution provides an excellent mechanism for minimizing numerical dissipation as long as the mesh points can be efficiently added and deleted from the mesh.

An efficient procedure has been developed for the simultaneous coarsening and refinement of three-dimensional tetrahedral meshes. An innovative data structure, that uses a combination of dynamically-allocated arrays and linked lists, allows the mesh connectivity to be rapidly reconstructed after individual points are added and/or deleted. The data structure is based on edges of the mesh rather than the tetrahedral elements that not only enhance the efficiency but also facilitates anisotropic mesh adaption. Error indicators are used to identify regions of the mesh that require adaption. The overall objective is to optimize the distribution of mesh points so that the flowfield is accurately modeled with a minimum of computational resources.

The key to success for a dynamic grid adaption procedure is the ability to efficiently add and delete points from the mesh. For an unsteady flow, this coarsening/refinement step must be completed every few time steps, so its efficiency must be comparable to that of the flow solver. It must also have a reasonable memory requirement. Our mesh adaption scheme has all of these attributes. It can also be applied to a variety of important NASA problems in addition to rotorcraft aerodynamics.

TENSOR METHODS FOR CONSTRAINED OPTIMIZATION

Dan Feng, 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. Similarly 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.

An analysis of the local convergence properties of a version of tensor methods has recently been done 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 similar situation, we show that the standard SQP method converges only linearly with constant converging to 1/2. (As we are aware of, 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 that characterize the singular constrained minimizers are also studied. This is one of the fundamental unsolved problems in constrained optimization. Feng, together with Robert Schnabel and Richard Byrd (CU Boulder) give a group of conditions that characterizes an interesting class of singular constrained minimizers. This gives us a better understanding of the current tensor methods as well as the SQP method for constrained optimization, and possible new ways of devising more robust and reliable methods for constrained optimization.

The work summarized above will be described in details in forthcoming separate papers.

TENSOR-KRYLOV METHODS FOR LARGE SPARSE SYSTEMS OF NONLINEAR EQUATIONS

Dan Feng

This research investigates tensor methods for system of nonlinear equations whose Jacobian matrix is large and sparse. Factorization-based tensor methods for nonlinear equations are not suitable for this type of large problems, since the fill-ins produced in the factorization process will usually soon run out the memory space. Previous tensor-Krylov methods that use two calls of a linear iterate solver (such as GMRES) have seen respectable improvement in both efficiency and robustness over analogous Newton-Krylov methods on a set of standard test problems. The margin of improvement is significantly larger on problems where the Jacobian matrix is ill-conditioned or rank deficient at the solution. However, the expense of the extra call of the linear iterate solver at each tensor-Krylov iteration could be formidable in many situations. A tensor-Krylov method that uses an amount of work equivalent to that of a single call of the linear iterate solver at each iteration would be desirable. This goal is achieved by a GMRES version of tensor-Krylov algorithms. This algorithm has shown competitive with its tensor counterpart in preliminary tests. Tensor-Krylov methods have potential applications to large sparse nonlinear systems arisen from the implicit scheme for partial differential equations. An interesting instance is the implicit scheme for the Euler and Navier-Stokes equations. These partial differential equations are playing a growing important role in computational fluid dynamics.

HIGH-ORDER DIFFERENCE METHODS FOR HYPERBOLIC SYSTEMS

Bertil Gustafsson, 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 obtains. 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.

SUPERCONVERGENCE OF GALERKIN METHODS AND COMPACT DIFFERENCE METHODS FOR HYPERBOLIC IBVPS

David Gottlieb (Brown University), Bertil Gustafsson, Pelle Olsson, Bo Strand (Upsalla University)

Error estimates for finite element approximations based on Galerkin methods for du/dt=Pu are determined by the approximation property of the finite element space. However, for problems with periodic solutions and uniform grids there is superconvergence at the grid points. If the approximating subspace is of order r, the error is of order 2r for first order systems. The question is whether this superconvergence property remains in the presence of boundary conditions for finite computational domains.

In this project it was shown that the superconvergence property is lost, no matter how the approximating subspace is chosen, if there are characteristics leaving the domain.

The Galerkin approximation can be viewed as an implicit compact difference method. By leaving the Galerkin formulation, there is greater flexibility when modifying the approximation near the boundary. However, it was shown that the higher order accuracy cannot be retained if the norm is modified only locally near the boundary.

STABILITY THEORY FOR COMPACT IMPLICIT DIFFERENCE APPROXIMATIONS

Bertil Gustafsson

Implicit compact difference approximations of Pade' type have a much smaller error constant than the corresponding explicit approximations. Therefore the extra work required for solving the linear systems at each time-step may well pay off. However, there are no general stability results for these methods when computational boundaries are present. In the project "Superconvergence of Galerkin methods and compact difference approximations for hyperbolic IBVP," we have shown that energy estimates impose severe restrictions on the approximations. The more general stability theory based on the Laplace transform technique does not apply in its present form, since the implicit form introduces certain singularities in the transformed equations."

In this project Gustafsson is working on the generalization of the stability theory. The goal is to prove that the Kreiss condition implies strong stability for a general class of approximations. For the special case of the fourth order accurate operator and scalar equations, strong stability has been proved in the project "High-order difference methods for hyperbolic systems."

RESEARCH ON COMPUTATIONAL FLUID DYNAMICS

Antony Jameson

While at RIACS since Sept 1993, Antony Jameson has been pursuing research in both the underlying mathematical aspects of computational fluid dynamics, and in related interdisciplinary projects.

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. This theory is built on stability in the L_∞ norm, and on the more stringent requirement of local extremum diminishing (LED) schemes, which also guarantee positivily of solutions evolving from initially positive data. An initial development of this theory was presented in AIAA Paper 93-3359. While at RIACS Antony Jameson has augmented the theory with a new analysis of the structure of discrete shock waves, which establishes general conditions for stationary stocks with a single interior point. The theory, which was presented in seminars at Stanford University and NASA Ames, has led to a new family of shock capturing schemes which exhibit superior numerical results, and also show promise for the accurant treatment of boundary layers (AIAA Paper 94-0647).

Antony Jameson's interdisciplinary research has been focussed on two areas. The first project is the treatment of aeroelastic problems by coupling a structures model with a new multigrid implicit time stepping scheme. Results of this work were presented in AIAA Paper 94-0056. The second project is the development of aerodynamic optimization techniques based on control theory, in collaboration with NASA Ames Advanced Aerodynamics Concepts Branch. The underlying idea is to find the gradient of the objective function with respect to an arbitrarily large (possibly infinite) number of design parameters by solving an adjoint equation. This can lead to tremendous savings in computational costs. Results of the application of this method to transonic airfoil design are given in AIAA Paper 94-0499

ITERATIVE METHODS FOR NON-HERMITIAN LINEAR SYSTEMS

Noël M. Nachtigal

During the period since the last report, we continued the work on the quasi-minimal residual method (QMR), in collaboration with Dr. Roland Freund at AT&T Bell Labs. We were interested in developing an alternate implementation of the QMR method, based on a coupled two-term recurrence version of the underlying nonsymmetric Lanczos algorithm, rather than the classical three-term Lanczos recurrence. We had already developed the theory behind the new implementation, and work was in progress on writing a software package that implements the new version of the QMR algorithm. The package has been completed, and in fact, it has grown beyond our original anticipations. It now contains the FORTRAN-77 implementation of six algorithms from the QMR family: the coupled QMR algorithm, with and without look-ahead, the original QMR algorithm, again with and without look-ahead, as well as the QMR-from-BCG algorithm proposed by Freund and Szeto, and the transpose-free QMR algorithm proposed by Freund. It also contains an eigenvalue solver based on the three-term recurrence version of the Lanczos algorithm with look-ahead. The new package supersedes a previous package we had made available on netlib, as it covers all the algorithms in the older package. We have rewritten the original QMR codes, so as to make them consistent with the new setup, and also to update them to benefit from the new strategies developed during the work for the coupled QMR algorithm. As before, the new package is made available to the scientific community through netlib.

SUMMATION BY PARTS, PROJECTIONS, AND STABILITY

Pelle Olsson

When solving a partial differential equation numerically it is necessary to have some bound of the growth rate of the solution, since otherwise round-off errors could grow arbitrarily fast. This upper bound can be established by ensuring some kind of stability. We have elected to use the energy method, because it can be applied to the continuous as well as the discrete model. Furthermore, it can be applied to general domains, which is important when studying multidimensional problems.

Stability of the continuous problem is established by means of an integration-by-parts procedure introducing boundary terms, some of which must be eliminated to ensure stability. For the finite difference model integration by parts is replaced by summation by parts. This amounts to designing the discrete difference operator ensuring that, in addition to the accuracy requirements, certain conditions of antisymmetry are met. As a consequence, the common problem of finding proper "numerical" boundary conditions will be eliminated; they will be built in the discrete difference operator.

The analytic boundary conditions are yet to be incorporated. We propose a certain projection operator, which interacts with the difference operator so as to generate boundary terms that are completely analogous to those of the continuous problem. This can be done for any type of linear boundary conditions. Thus, an energy estimate is obtained for the discrete problem, provided there is one for the analytic model. This conclusion remains true for domains in several space dimensions, even if the boundary is non-smooth. Furthermore, using this projection operator allows us to derive stability results for a larger class of finite difference operators than those considered by Kreiss and Scherer in their pioneering work. Stability is proved for explicit high-order finite difference approximations of mixed hyperbolic-parabolic variable coefficient systems subject to general inhomogeneous boundary conditions. A complete presentation can be found in RIACS Technical Report 93.04.

ENERGY ESTIMATES FOR NONLINEAR CONSERVATION LAWS

Pelle Olsson, 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. A RIACS Technical Report describing the research is in progress.

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 adaptive grid system will initially be applied to model problems in two space dimensions, but it can be extended compatibly to three dimensions and to more realistic problems.

The programming language Vorpal, currently under development, is expected to 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. Recent work in this project has focused on the design of the composite adaptive grid system and on improving the design of Vorpal to make the implementation of abstract data types such as grids more practical. As the adaptive grid system evolves and grows, it should also be able to use the future support in Vorpal for concurrency and interactivity.

C. HIGH PERFORMANCE NETWORKS & TECHNOLOGY

HIGH-PERFORMANCE COMMUNICATIONS AND NETWORKING

Marjory J. Johnson

The objective of this project is to support the High Performance Computing and Communications Program.

A major activity during the past year has been continuation of the effort to establish a gigabit testbed within the Bay Area. As reported earlier, M. Johnson has been collaborating with persons from the Numerical Aerodynamic Simulation Systems Division (NAS) and the NASA Science Internet Project Office (NSIPO) to represent NASA in this effort. Fourteen organizations within the Bay Area are involved, including major technology companies, research organizations, universities, and government laboratories.

After three years of planning and writing proposals, physical deployment of the testbed has finally begun. Our testbed is sponsored by Pacific Bell's CalREN (California Research and Education) program, a program initiated this year to stimulate the development and dissemination of high-speed communications applications within the state of California. Under CalREN's sponsorship our Bay Area Gigabit Testbed will receive access to Pacific Bell's ATM network services. On December 30, 1993, the first two testbed sites, NASA Ames and Xerox Palo Alto Research Center (PARC), were connected to the Pacific Bell ATM switch via OC-3c links. We are currently conducting experiments between these two sites. We anticipate that all fourteen sites will be connected to the testbed by the end of February, 1994. M. Johnson, Bill Johnston of Lawrence Berkeley Laboratory, and Dan Swinehart of Xerox PARC are jointly coordinating the testbed project.

Now that the testbed is a physical reality, testbed participants are focusing on development of the network infrastructure to support a diverse set of distributed applications. Applications of particular interest to NASA include networked mass storage, teleseminars, and collaborative work environments.

In other HPCCP-related activities, M. Johnson directed the work of two summer students. Justin Paola, a graduate student at the University of Arizona, returned for a second summer to continue the RIACS/UA collaboration on "Content-based Query and Browse of Earth Science Imagery Databases using High Performance Computers and Networks," a project funded by NASA under the HPCC/ESS (Earth System Science) Program. This summer Justin completed a survey paper analyzing the use of neural networks for the classification of remotely sensed multispectral imagery. The paper has been submitted for publication. During the coming year we plan to investigate networking issues related to the project.

The second summer student working with M. Johnson was Michael Kumbera, a graduate student from the University of Wisconsin - Milwaukee. Michael implemented the NAS multigrid benchmark on the CM-5. His accomplishments provide a foundation for future work to analyze the performance of the interprocessor communication architecture of the CM-5 and other parallel computers.

Finally, M. Johnson also participated in a joint NASA/DoD effort to develop interoperable data communications standards for use in both civil and military space projects. A program plan has been prepared and presented to both agencies.

HPCCP SOFTWARE EXCHANGE: THE RIACS CONTRIBUTION

Michael Raugh

Overall objectives of the HPCC Software Exchange are to utilize Internet and database technologies to facilitate software sharing and reuse for the HPCC Program, coordinate related efforts at several agencies, and develop a common interface to locator services and access mechanisms for software and related information at various locations on the Internet.

Envision the Internet as a large "library" of separate, individually administered repositories, each with its own collection of special holdings, standards for submission and rules for access, and local taxonomy and indexes. Develop a client/server software system that can deal with such heterogeneity and provides reference tools such as catalogues, and indexes for major resources to speed resource discovery, much like in an ordinary library. Use existing resources such as Netlib as a model for software archiving and GAMS as a model for cross-indexing of software modules. Use emerging standards for cataloguing, such as USMARC and RIG profiling. Use existing Internet access tools such as anonymous ftp, Archie, Gopher, telnet, WAIS, WorldWideWeb, where appropriate, to access resources. Demonstrate Logical Books as a consistent interface for representative databases in the Library, such as the Software Union Catalogue, the Catalogue of Repositories, and the Book of Standards, as well as Netlib and GAMS. Extend the HPCC Software Exchange system to software, data, documentation, bulletin boards and other resources that are useful to HPCC investigators.

RIACS is the "librarian" for the Software Exchange. A Catalogue-of-Repositories Book (repcat), has been developed with a subject index and prototype browser that permits access to over 200 repositories on the Internet by means of a few quick mouse clicks. A series of perl programs have been written to simplify and automate inclusion of new acquisitions into the library and updating of the various catalogue indexes. The Catalogue of Repositories and the Logical Library were demonstrated at a full-dress review at GSFC on September 1. The Software Union Catalogue, the mechanism favored by the review committee for software searching, is based upon repcat. Fourteen corporations have been invited to participate in a "corporation" subsection of the Logical Library by providing a statement of corporate relevance to (or interest in) the HPCC Program, as well as providing a further option of a telnet address for bulletin boards, software, or documentation of corporate or user materials relevant to the HPCC Program. In ongoing dialogues, six corporations have responded thus far by providing a statement, and one has begun preparing a workstation with corporate info for HPCCP Internet access. In addition, to show the suitability of the Books model for presenting combined graphics and text in a normal book-like presentation, an edition of the current CAS Project Plan is being prepared in which text and scanned graphics are combined using html as markup language and Mosaic as the access tool within the book. Thus icons can be clicked on to view graphics or to navigate between consecutive pages, while at the same time the Books' contents and pages mechanisms, as well as the Books' indexes, can be used to jump to chosen locations within the document. This demonstrates the relative strengths of the Books and Mosaic approaches to representing databases in a particularly striking way. Mosaic is excellent for setting up hypertext links within and between documents while also allowing for the mixing of imagery (and other constructions, such as audio presentation) on a "page," whereas Books provides convenient means for navigating and locating specific information within a large database. Books also can easily incorporate any access mechanism, such as Mosaic, as desired.

The project has shown that it is possible to make access to a very broad spectrum of heterogeneous multi-media information relevant to software development, data, and bibliographics, accessible in a convenient Library and Books paradigm. The client/server software utilizes existing repositories, search and access mechanisms and enhances the value of these by making them accessible through a single library with refined indexes and browsing facilities of its own.

Library and Books software is under continual refinement. Release to the HPC community awaits approval from HQ. Meanwhile, the development team continues to publish Books of various kinds, some featuring very large databases such as the EOS Master Directory and Data, and others demonstrating the capability of Books to incorporate markup documents, to demonstrate the power and versatitility of the basic architecture.

D. LEARNING SYSTEMS

MODEL-BASED LEARNING

Wray Buntine

This research involves developing methods and tools for doing supervised learning and regression. Supervised learning is the process of analyzing data to construct a "classifier," which is used for predicting one of the discrete variables in a record from the other variables.

In October 1992 a new applied project was initiated with Marshall Space Flight Center. This involves developing and applying data analysis methods to diagnosis Space Shuttle Main Engines via high resolution spectral readings. A pilot data analysis study on the Aviation Safety Reporting System completed in October 1992 yielded promising results, and a proposal is being prepared for more work.

The IND tree learning package was submitted to COSMIC in October 1992 as Version 2.0. This fixed numerous bugs in the early release and improved the interface. A new beta test version was released in December 1992 that included rule learning and decision graphs.

INDUCTION FOR HEURISTIC INFORMATION RETRIEVAL

Barney Pell, Catherine Baudin (NASA/Recom)

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 OVERGENERAL; 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 FEEDBACK 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.

DATA-BAYESIAN SUFACE MODELLING FROM MULTIPLE IMAGES

Peter Cheeseman (RIACS), Bob Kanefsky (RECOM),
Rich Kraft (RECOM), John Stutz (NASA)

Beginning in late 1991, the Bayes group (FIA) began a project to develop the theory and practice of multiple image data combination. Multiple images taken from similar locations and under similar lighting conditions contain similar - but not identical - information.

Slight differences in instrument orientation and position ensure mismatches between the pixel grids of different images. These mismatches ensure that any point on the ground is sampled differently in each image. We register all the images with respect to each other to an accuracy of a small fraction of a pixel, then combine the information from the multiple images to increase both spatial resolution and grey-scale resolution. In FY94 we have refined our registration algorithm so that we can achieve registration to at least within 1/10th. of a pixel. Our current algorithm is limited to flat areas with nearly identical viewing angle and lighting conditions. We plan to extend the theory and implementation to spherical planetary surfaces, and later to include arbitrary 3-D surface models, with any combination of lighting and viewing conditions. This extension will require modeling of shadows and atmospheric conditions. In addition, we plan to integrate information across different spectral bands. Target areas will be choosen to aid specific planetary scientist's needs.

(Note: See graph next pg.)

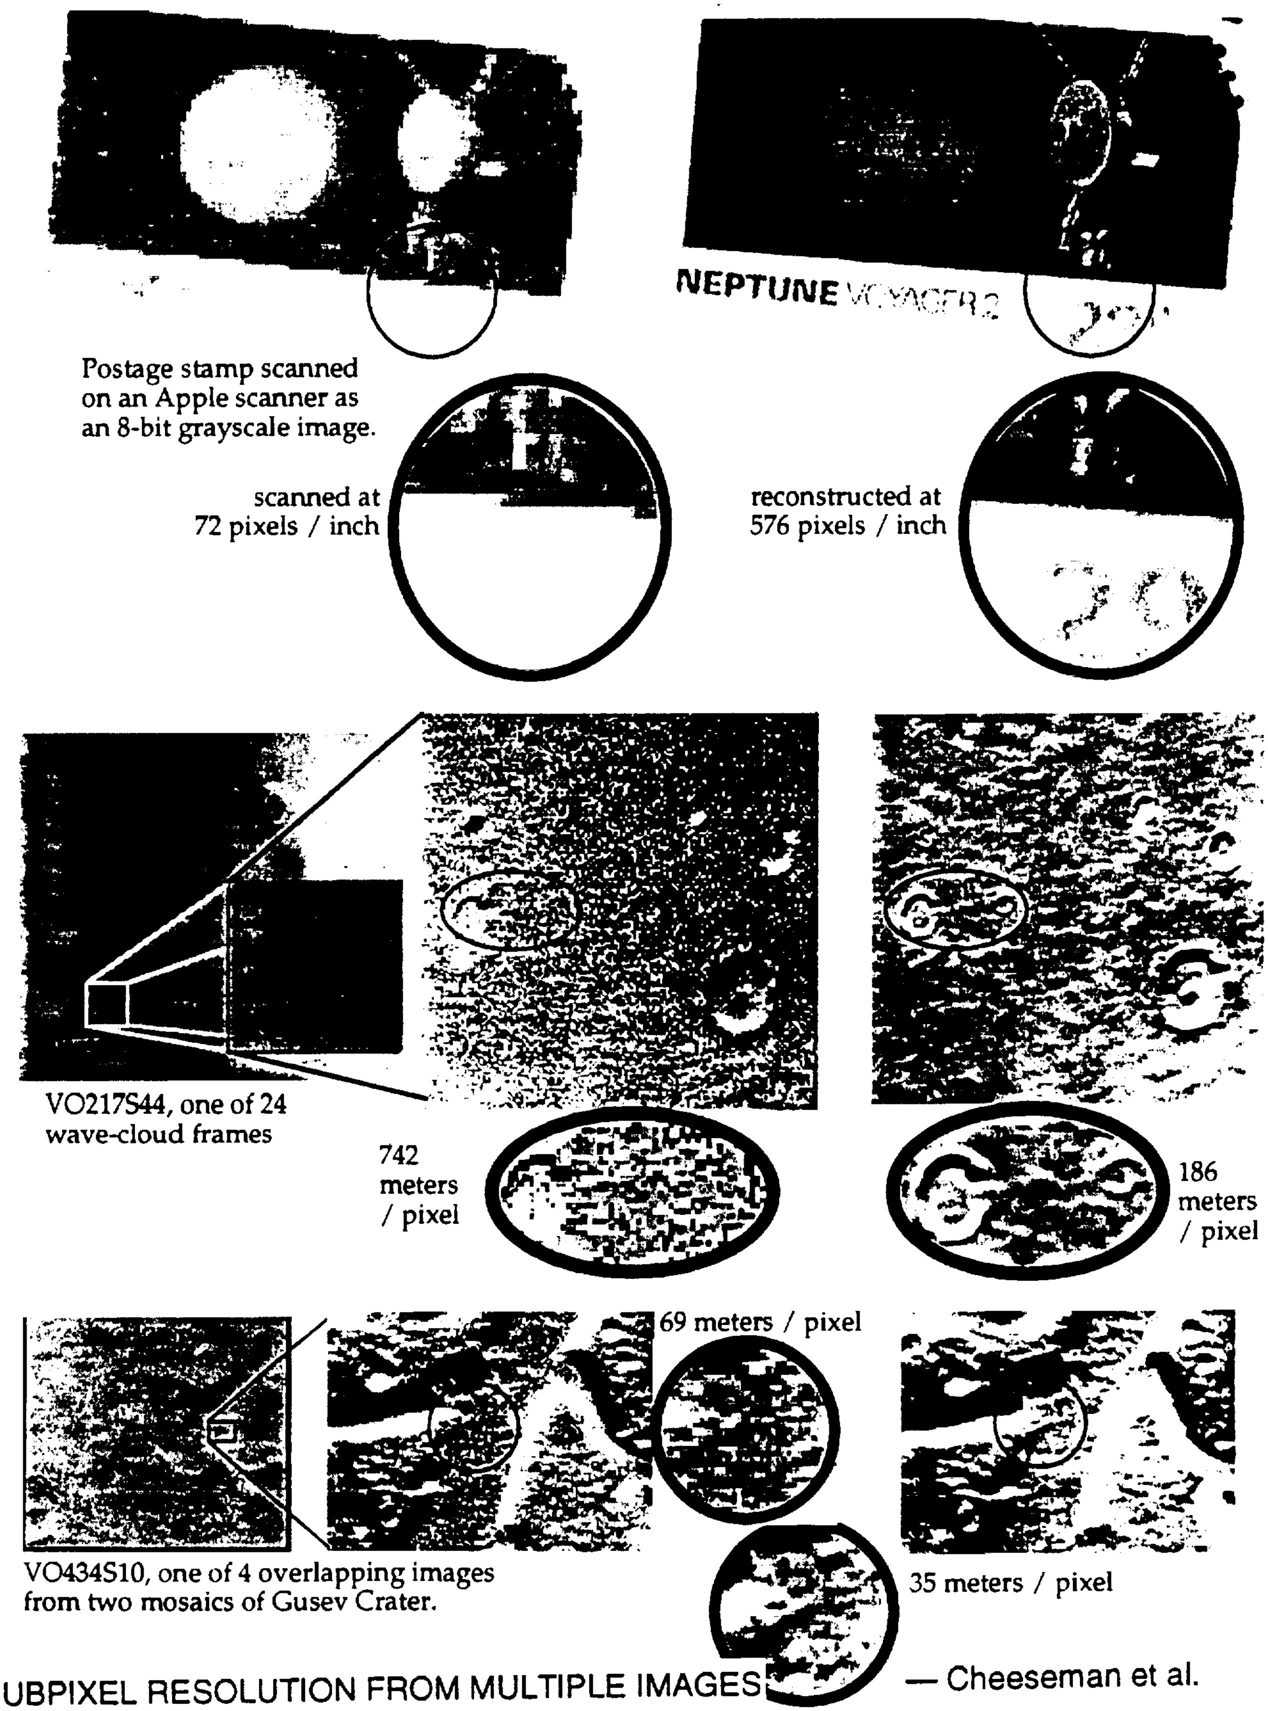

SUBPIXEL RESOLUTION FROM MULTIPLE IMAGES — Cheeseman et al.

III. TECHNICAL REPORTS

A NEW PROCEDURE FOR DYNAMIC ADAPTION OF THREE-DIMENSIONAL UNSTRUCTURED GRIDS
Rupak Biswas and Roger Strawn (US Army Aeroflightdynamics Directorate)
RIACS TR 93.01, January 1993, 12 pages

To appear in Applied Numerical Mathematics, 31st AIAA Aerospace Sciences Meeting & Exhibit, Reno, Nevada, Paper Number AIAA-93-0672.

A new procedure is presented for the simultaneous coarsening and refinement of three-dimensional unstructured tetrahedral meshes. This algorithm allows for localized grid adaption that is used to capture aerodynamic flow features such as vortices and shock waves in helicopter flowfield simulations. The mesh-adaption algorithm is implemented in the C programming language and uses a data structure consisting of a series of dynamically-allocated linked lists. These lists allow the mesh connectivity to be rapidly reconstructed when individual mesh points are added and/or deleted. The algorithm allows the mesh to change in an anisotropic manner in order to efficiently resolve directional flow features. The procedure has been successfully implemented on a single processor of a Cray Y-MP computer. Two sample cases are presented involving three-dimensional transonic flow. Computed results show good agreement with conventional structured-grid solutions for the Euler equations.

EFFICIENT, MASSIVELY PARALLEL, EIGENVALUE COMPUTATION
Yan Huo (Princeton University) and Robert S. Schreiber
RIACS TR 93.02, January 1993, 24 pages

In numerical simulations of disordered electronic systems, one of the most common approaches is to diagonalize random Hamiltonian matrices and to study the eigenvalues and eigenfunctions of a single electron in the presence of a random potential. In this paper, we describe an effort to implement a matrix diagonalization routine for real symmetric dense matrices on massively parallel SIMD computers, the Maspar MP-1 and MP-2 systems. Results of numerical tests and timings are also presented.

GENERATING LOCAL ADDRESSES AND COMMUNICATION SETS FOR DATA-PARALLEL PROGRAMS
Siddhartha Chatterjee, John R. Gilbert (Xerox PARC), Fred J. E. Long (UCSC), Robert S. Schreiber, and Shang-Hua Teng (MIT)
RIACS TR 93.03, April 1993, 10 pages

Appears in the Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, San Diego, CA, 19-22 May 1993. To appear in Journal of Parallel and Distributed Computing.

Generating local addresses and communication sets is an important issue in distributed-memory implementations of data-parallel languages such as High Performance Fortran. We show that for an array A affinely aligned to a template that is distributed over p processors with a cyclic(k) distribution, and a computation involving the regular section A(l:h:s), the local memory access sequence for any processor is characterized by a finite state machine of at most k states. We present fast algorithms for computing the essential information about these state machines, and extend the framework to handle multidimensional arrays. We also show how to generate communication sets using the state machine approach. Performance results show that this solution requires very little runtime overhead and acceptable preprocessing time.

SUMMATION BY PARTS, PROJECTS, AND STABILITY
Pelle Olsson
RIACS TR 93.04, June 1993, 61 pages

Presented at the 1993 SIAM Annual Meeting, Philadelphia, PA, July 12-16,1993

We have derived stability results for high-order finite difference approximations of mixed hyperbolic-parabolic initial-boundary value problems (IBVP). The results are obtained using summation by parts and a new way of representing general linear boundary conditions as an orthogonal projection. By rearranging the analytic equations slightly we can prove strict stability for hyperbolic-parabolic IBVP. Furthermore, we generalize our technique so as to yield strict stability on curvilinear non-smooth domains in two space dimensions. Finally, we show how to incorporate inhomogeneous boundary data while retaining strict stability. Using the same procedure one can prove strict stability in higher dimensions as well.

A REVIEW AND ANALYSIS OF NEURAL NETWORKS FOR CLASSIFICATION OF REMOTELY SENSED MULTISPECTRAL IMAGERY
Justin D.Paola(University of Arizona) and Robert A. Schowengerdt (University of Arizona)
RIACS TR 93.05, July 1993, 47 pages

Submitted to Remote Sensing of Environment

A literature survey and analysis of the use of neural networks for the classification of remotely sensed multispectral imagery is presented. As part of a brief mathematical review, the backpropagation algorithm, which is the most common method of training multi-layer networks, is discussed with an emphasis on its application to pattern recognition. The analysis is divided into five aspects of neural network classification: 1) input data preprocessing, structure, and encoding, 2) output encoding and extraction of classes, 3) network architecture, 4) training algorithms, and 5) comparisons to conventional classifiers. The advantages of the neural network method over traditional classifiers are its non-parametric nature, arbitrary decision boundary capabilities, easy adaptation to different types of data and input structures, fuzzy output values that can enhance classification, and good generalization for use with multiple images. The disadvantages of the method are slow training time, inconsistent results due to random initial weights, and the requirement of obscure initialization values (e.g., learning rate and hidden layer size). Possible techniques for ameliorating these problems are discussed. It is concluded that, although the neural network method has several unique capabilities, it will become a useful tool in remote sensing only if it is made faster, more predictable, and easier to use.

THE ALIGNMENT-DISTRIBUTION GRAPH
Siddhartha Chatterjee, John R. Gilbert (Xerox PARC), and Robert S. Schreiber
RIACS TR 93.06, August 1993, 15 pages

Appears in the Proceedings of the Sixth Annual Languages and Compilers for Parallelism Workshop, Portland, OR, 12-14 August 1993.

Implementing a data-parallel language such as Fortran 90 on a distributed-memory parallel computer requires distributing aggregate data objects (such as arrays) among the memory modules attached to the processors. The mapping of objects to the machine determines the amount of residual communication needed to bring operands of parallel operations into alignment with each other. We present a program representation called the alignment-distribution graph that makes these communication requirements explicit. We describe the details of the representation, show how to model communication cost in this framework, and outline several algorithms for determining object mappings that approximately minimize residual communication.

ON THE SUPERCONVERGENCE OF GALEKIN METHODS FOR HYPERBOLIC IBVP
David Gottlieb (Brown University), Bertil Gustafsson, Pelle Olsson, and Bo Strand (Uppsala University)
TR 93.07, August 1993, 19 pages

Submitted to SIAM Journal on Numerical Analysis

Finite element Galerkin methods for periodic first order hyperbolic equations exhibit superconvergence on uniform grids at the nodes, i.e., there is an error estimate of order 2r instead of the expected approximation order r. In this paper it will be shown that no matter how the approximating subspace is chosen, the superconvergence property is lost if there are characteristics leaving the domain. We shall also discuss the implications of this result when constructing compact implicit difference schemes.

MOBILE AND REPLICATED ALIGNMENT OF ARRAYS IN DATA-PARALLEL PROGRAMS
Siddhartha Chatterjee, John R. Gilbert (Xerox PARC), Robert S. Schreiber
RIACS TR 93.08, September 1993, 10 pages

Appears in the Proceedings of Supercomputing'93, Portland, OR, 15-19 November 1993.

When a data-parallel language like Fortran 90 is compiled for a distributed-memory machine, aggregate data objects (such as arrays) are distributed across the processor memories. The mapping determines the amount of residual communication needed to bring operands of parallel operations into alignment with each other. A common approach is to break the mapping into two stages: first, an alignment that maps all the objects to an abstract template, and then a distribution that maps the template to the processors.

We solve two facets of the problem of finding alignments that reduce residual communication: we determine alignments that vary in loops, and objects that should have replicated alignments. We show that loop-dependent mobile alignment is sometimes necessary for optimum performance, and we provide algorithms with which a compiler can determine good mobile alignments for objects within DO loops. We also identify situations in which replicated alignment is either required by the program itself (via SPREAD operations) or can be used to improve performance. We propose an algorithm based on network flow that determines which objects to replicate so as to minimize the total amount of broadcast communication in replication. This work on mobile and replicated alignment extends our earlier work on determining static alignment.

A BROWSING TOOL FOR THE INTERNET LOGICAL LIBRARY OF THE HPCCP SOFTWARE EXCHANGE
Ross Biro (Stanford University)
TR 93.09, September 1993, 20 pages

As the quantity of information available on the Internet grows, locating a particular piece of information becomes more difficult. One possible solution is for a database of pointers to all available information to be maintained at a central site. Subject classifications for all the information could also be maintained in order to make searching possible. This paper describes one possible method of searching such an index. In particular a prototype browsing tool has been created using TCL/TK to demonstrate several possible features: rapidly scanning at any rank of the index, narrowing the index to any scope, regular-expression searching, and creation of a list of pointers answering to any set of index terms. The prototype browser is an easy-to-use independent X application designed for use in the Catalog of Repositories of the HPCC Software Exchange.

UNSTRUCTURED ADAPTIVE MESH COMPUTATIONS OF ROTORCRAFT HIGH-SPEED IMPULSIVE NOISE
Roger Strawn (US Army Aeroflightdynamics Dir.), Michael Garceau (Stanford University), and Rupak Biswas
RIACS TR 93.10, October 1993, 12 pages

AIAA 15th Aeroacoustics Conference, Long Beach, California, Paper No. AIAA-93-4359, Submitted to AIAA Journal of Aircraft,

A new method is developed for modeling helicopter high-speed impulsive (HSI) noise. The aerodynamics and acoustics near the rotor blade tip are computed by solving the Euler equations on an unstructured grid. A stationary Kirchhoff surface integral is then used to propagate these acoustic signals to the far field. The near-field Euler solver uses a solution-adaptive grid scheme to improve the resolution of the acoustic signal. Grid points are locally added and/or deleted from the mesh at each adaptive step. An important part of this procedure is the choice of an appropriate error indicator. The error indicator is computed from the flowfield solution and determines the regions for mesh coarsening and refinement. Computed results for HSI noise compare favorably with experimental data for three different hovering rotor cases.

GLOBAL ASYMPTOTIC BEHAVIOR OF ITERATIVE IMPLICIT SCHEMES
H.C. Yee (Fluid Dynamics Division NASA), P.K. Sweby (University of Reading, Whiteknights, England)
RIACS TR 93.11, December 1993, 57 pages

Will be submitted to J. of Bifurcation and Chaos

The global asymptotic nonlinear behavior of some standard iterative procedures in solving nonlinear systems of algebraic equations arising from four implicit linear mul three models of 2 X 2 systems of first-order autonomous nonlinear ordinary differential equations (ODEs) is analyzed using the theory of dynamical systems. The iterative procedures include simple iteration and full and modified Newton iterations. The results are compared with standard Runge-Kutta explicit methods, a non-iterative implicit procedure, and the Newton method of solving the steady part of the ODEs. Studies showed that aside from exhibiting spurious asymptotes, all of the four implicit LMMs can change the type and stability of the steady states of the differential equations (DEs). They also exhibit a drastic distortion but less shrinkage of the basin of attraction of the true solution than standard explicit methods. The simple iteration procedure exhibits behavior which is similar to standard explicit methods except spurious steady-state numerical solutio attraction of the non-iterative implicit procedure mimic more closely the basins of attraction of the DEs than the three iterative implicit procedures for the four implicit LMMs. The results of the study can be used as an explanation for possible causes and cures of slow convergence and non-convergence of steady-state numerical solutions when using an implicit LMM time-dependent approach in computational fluid dynamics.

SUPER-RESOLVED SURFACE RECONSTRUCTION FROM MULTPLE IMAGES
P. Cheeseman, B. Kanefsky (RECOM), R. Hanson (RECOM), J. Stutz (NASA)
Tech Report: FIA-93-02, Feb. 1993

This paper describes a Bayesian method for the combining information from multiple images of the same surface to construct a super-resolved surface model. We develop the theory and algorithms in detail for the 2-D surface reconstruction problem, and show the results on actual images. These results show dramatic improvements in resolution. The Bayesian approach uses neighbor correlation information as well as combining information from multiple images. The reconstructed surfaces have both significantly higher spatial resolution and grey scale resolution. We show how this theory can be extended to super-resolved 3-D surface reconstruction from multiple images, such as a series of frames from a moving TV camera. We also show that this approach can be applied to diffraction blurred images, offering the possibility of resolving images below the wavelength of light.

[Also in prep. as book chapter].

A STRATEGIC METAGAME PLAYER FOR GENERAL CHESS-LIKE GAMES
Barney Pell
Tech Report FIA-93-32

Appeared in Proceedings of the AAAI Fall Symposium on Games: Learning and Planning, Raleigh, NC, October 1993, November 1993.

This paper reviews the concept of Metagame and discusses the implementation of \metagamer, a program which plays Metagame in the class of {...m symmetric chess-like games}, which includes chess, Chinese-chess, checkers, draughts, and Shogi. The program takes as input the {...m rules} of a specific game and analyses those rules to construct for that game an efficient representation and an evaluation function, both for use with a generic search engine. The strategic analysis performed by the program relates a set of general knowledge sources to the details of the particular game. Among other properties, this analysis determines the relative value of the different pieces in a given game. Although \metagamer does not learn from experience, the values resulting from its analysis are qualitatively similar to values used by experts on known games, and are sufficient to produce competitive performance the first time the program actually plays each game it is given. This appears to be the first program to have derived useful piece values directly from analysis of the rules of different games.

IV. PUBLICATIONS

RUPAK BISWAS

An Adaptive Mesh-moving and Refinement Procedure for One-dimensional Conservation Laws, R. Biswas, J. E. Flaherty (Rensselaer Polytechnic Institute), D. C. Arney (United States Military Academy). Applied Numerical Mathematics, Volume 11, Number 4, February 1993, pages 259-282.

Dynamic Mesh Adaption for Triangular and Tetrahedral Grids, R. Biswas, R. Strawn (US Army Aeroflightdynamics Directorate). NASA Unstructured Grid Generation Techniques and Software Workshop, Hampton, Virginia, NASA CP 10119, September 1993, pages 181-192.

Parallel, Adaptive Finite Element Methods for Conservation Laws, R. Biswas, K. D. Devine (Rensselaer Polytechnic Institute), J. E. Flaherty (Rensselaer Polytechnic Institute). To appear in Applied Numerical Mathematics.

Numerical Simulations of Helicopter Aerodynamics and Acoustics, R. Strawn (US Army Aeroflightdynamics Directorate), R. Biswas. Submitted to 19th Army Science Conference, Orlando, Florida.

Dynamic Mesh Adaption for Tetrahedral Grids, R. Biswas, R. Strawn (US Army Aeroflightdynamics Directorate). Submitted to 14th International Conference on Numerical Methods in Fluid, Dynamics, Bangalore, India.

PETTER BJORSTAD

Two Different Data-Parallel Implementations of the BLAS, P. Bjorstad and T. Sorevik (University of Bergen). To appear in NATO proceedings from a meeting on parallel computing, Cetraro, Italy, June 1994.

Dataparallel BLAS As a Basis for LAPACK on Massively Parallel Computers, P. Bjorstad and T. Sorevik (University of Bergen). In NATO Advanced Study Institute, published by Kleuver 1993.

Large Scale Structural Analysis on Massively Parallel Computers, P. Bjorstad and J. Cook, (University of Bergen). In NATO Advanced Study Institute, published by Kleuver 1993.

WRAY BUNTINE

On Solving Equations and Disequations, W.L. Buntine and H. J. Burckert, Journal of the ACM, to appear.

Computing Second Derivatives in Feed-Forward Networks: a Review, W.L. Buntine and A. Weigend, IEEE Transactions on Neural Networks, to appear.

WRAY BUNTINE(cont'd)

Classification and Beyond: Exploring the ASRS Database, R. Kraft and W.L. Buntine in Proceedings of Fourth International Conference on Human-Machine Interaction and Artificial Intelligence in Aerospace, Toulouse, France, 1993.

Initial Exploration of the ASRS Database, R. Kraft and W.L. Buntine in proceedings of Seventh International Symposium on Avaition Psychology, Columbus, Ohio, 1993.

SIDDHARTHA CHATTERJEE

Implementation of a Portable Nested Data-Parallel Language, G. E. Blelloch (CMU), S. Chatterjee, J. Hardwick (CMU), J. Sipelstein (CMU), and M. Zagha (CMU). Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pages 102-111. San Diego, CA, May 1993. Accepted for publication in Journal of Parallel and Distributed Computing, 1994.

Compiling Nested Data-Parallel Programs for Shared-Memory Multiprocessors, S. Chatterjee. ACM Transactions on Programming Languages and Systems, 15(3), pages 400-462, July 1993.

Programming Models, Compilers, and Algorithms for Irregular Data-Parallel Computations, S. Chatterjee, 1993. Accepted for publication in International Journal of High Speed Computing, 1994.

Solving Linear Recurrences with Loop Raking, G.E. Blelloch (CMU), S. Chatterjee, and M. Zagha (CMU), 1993. Accepted for publication in Journal of Parallel and Distributed Computing, 1994.

Automatic Array Alignment in Data-Parallel Programs, S. Chatterjee, J. R. Gilbert, R. Schreiber, and S.-H. Teng. Proceedings of the Twentieth Annual SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Charleston, SC, January 1993, pages 16--28.

Generating Local Addresses and Communication Sets for Data-Parallel Programs, S. Chatterjee, J. R. Gilbert, F. J. E. Long, R. Schreiber, and S.-H. Teng. Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, San Diego, CA, May 1993, pages 149-158.

JOSEPH OLIGER

Special Year in Numerical Analysis at the Namkai Institute, R. B. Kellogg and J. Oliger, SIAM News, Volume 26, Number 6, October 1993, page 22.

ROBERT S. SCHREIBER

Optimal Parallel Solution of Sparse Triangular Systems, F. L. Alvarado (University of Wisconsin) and R. Schreiber. SIAM Journal on Scientific and Statistical Computing, 14 (1993) pp. 446-460.

Block "LU" Factorization, J. E. Demmel (University of California), N. J. Higham (Manchester University), R. Schreiber. Journal of Numerical Linear Algebra and Applications, to appear.

Scalability of Sparse Direct Solvers, R. Schreiber, In A. George, J. R. Gilbert, J.W.H. Liu, editors, Graph Theory and Sparse Matrix Computation, pp. 191--209. The IMA Volumes in Mathematics and Its Applications, Volume 56, Springer-Verlag, New York, 1993.

Highly Parallel Sparse Triangular Solution, F. L. Alvarado (University of Wisconsin), A. Poghen (University of Waterloo), R. Schreiber. In A. George, J. R. Gilbert, and J.W.H. Liu, editors, Graph Theory and Sparse Matrix Computation, pp. 141--157. The IMA Volumes in Mathematics and Its Applications, Volume 56, Springer-Verlag, New York, 1993.

Efficient, Massively Parallel Eigenvalue Computation, Y. Huo (Princeton University), R. Schreiber. International Journal of Supercomputer Applications, 7:4 (1993).

Efficient Eigenvalue Computation on the Maspar, Y. Huo (Princeton University), R. Schreiber, In F. T. Luk, editor, Advanced Signal Processing Algorithms, Architectures, and Implementations IV, SPIE Vol. 2027, SPIE, Bellingham, Washington, 1993.

Subway: A Communication Compiler for the Maspar. P. E. Bjorstad (University of Bergen), R. Schreiber. Submitted to the The 1994 Scalable High Performance Computing Conference.

Book Review: "Very Large Scale Computation in the 21st Century," edited by J. P. Mesirov, R. Schreiber. SIAM Review 35 (1993).

Improved load distribution in distributed memory sparse Cholesky, E. Rothberg (Intel Corporation), Robert S. Schreiber. In preparation.

The High Performance Fortran Handbook, C.H. Koebel (Rice University), D. B. Loveman (Digital Equipment Corporation), R. Schreiber, G. L. Steele, Jr. (Thinking Machines Corporation), M. E. Zosel (Lawrence Livermore National Laboratory). The MIT Press, Cambridge, Massachusetts, 1994.

THOMAS SHEFFLER

Implementing the Multiprefix Operation on Parallel and Vector Computers, T. Sheffler. Presented at the "5th Annual Symposium on Parallel Algorithms and Architectures", June 30 - July 2, 1993, Velen, Germany.

An Analysis of Hashing on Parallel and Vector Computers, T. Sheffler. Presented at the 1993 International Conference on Parallel Processing, St. Charles, IL, August 16 - 20, 1993.

STEVEN SUHR

3D Composite Grids Using Bezier Curves and Surfaces in Component Adaptive Methods, R. G. Venkata, Stanford University, S. C. Suhr, J. Oliger, J. H. Ferziger, Stanford University. Stanford CLaSSiC Manuscript 93-32, August 1993, 17 pages.

V. SEMINARS AND COLLOQUIA

RUPAK BISWAS

A New Procedure for Dynamic Adaption of Three-Dimensional Unstructured Grids, R. Biswas, R. Strawn (US Army Aeroflightdynamics Directorate), at 31st AIAA Aerospace Sciences Meeting & Exhibit, Reno, Nevada, January 13, 1993.

Dynamic Mesh Adaption for Triangular and Tetrahedral Grids, R. Biswas, R. Strawn (US Army Aeroflightdynamics Directorate), at NASA Unstructured Grid Generation Techniques and Software Workshop, Hampton, Virginia, April 27, 1993.

WRAY BUNTINE

Prior Probabilities, W. Buntine, MIT, Lab. for CS, May 3rd, Stanford AI SIG Lunch, KSL, April 30; IBM Almaden Res. Center, March 17; Int. Computer Sceinec Inst., Berkeley, February 25, 1993.

Learning with Graphical Models, W. Buntine, MIT, Dept. Behavior and Brain Science, May 4, 1993.

OPAD Data Analysis at AIAA/SAE/ASME/ASEE, W. Buntine, 29th Joint Propulsion Conference, Monterey, California, 1993, April, 1993.

Learning with Probabilistic Networks W. Buntine, W. Buntine, at the Workshop on Bayesian Analysis in Expert Systems, Pavia, Italy, June 4th, 1993, and at Neural Networks for Computing, Snowbird, Utah, April 7th 1993.

Invited speaker at the Neural Information Processing Systems Workshop, *Putting it All Together: Methods for Combining Neural Networks*, W. Buntine, December 4th, 1993.

Compiling Learning Algorithms, W. Buntine, at University of Colorado, Boulder, December 7th, 1993.

Tutorial given at *Uncertainty in Artificial Intelligence Conference on Learning: a probabilistic approach*, W. Buntine, Washington, DC, July 9, 1993.

SIDDHARTHA CHATTERJEE

Implementing Array Parallelism on Distributed-Memory Machines, S. Chatterjee, NASA Ames Research Center, June 29, 1993.

Automatic Alignment of Fortran 90 Arrays, S. Chatterjee, UMIACS Seminar on High Performance Computing, University of Maryland, College Park, MD, March 16, 1993.

SIDDHARTHA CHATTERJEE(cont'd

Compiling Array Parallelism for Distributed Memory, S. Chatterjee, The Portland Group, Portland, OR. 16 August 1993.

Automatic Array Alignment in Data-Parallel Program, S. Chatterjee, Twentieth Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages, Charleston, SC, 10 January 1993.

Generating Local Addresses and Communication Sets for Data-Parallel Programs, S. Chatterjee, Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, San Diego, CA, 21 June 1993.

PETER CHEESEMAN

Super-resolution Work, P. Cheeseman as an invited speaker at, the 11th Maximum Entropy and Bayesian Inference Conference, Santa Barbara, October 1993.

DAN FENG

Parallel Methods for Block Bordered Systems of Nonlinear Equations and Tensor Methods for Constrained Optimization, D. Feng, NASA Ames Research Center, February 11, 1993.

Tensor Methods for Equality Constrained Optimization, D. Feng, NASA Ames Research Center, September 23, 1993.

BERTIL GUSTAFSSON

The Boundary Treatment for High-order Difference and Finite Element Methods, B. Gustafsson, 1993 SIAM Annual Meeting, Philadelphia, PA, July 12-16, 1993.

Open Boundary Conditions for the Navier-Stokes Equations, B. Gustafsson, Wright-Patterson Air Force Base, Dayton, OH, Aug 16-18, 1993.

High Order Methods for Hyperbolic Systems, B. Gustafsson, University of New Mexico, Albuquerque, NM, Nov 8-9, 1993.

High Order Methods for Hyperbolic Systems, B. Gustafsson, Los Alamos National Laboratory, Los Alamos, NM, Nov 10, 1993.

PAUL HAVLAK

Efficient Analysis of Symbolic Values, P. Havlak, Rice University, NASA Ames Research Center, March 4, 1993.

NICHOLAS J. HIGHAM

A Parallel Algorithm for Computing the Polar Decomposition, N. J. Higham, University of Manchester, NASA Ames Research Center, June 22, 1993.

MARLIS HOCHBRUCK

Krylov Subspace Methods for Nonsymmetric, and Symmetric Inconsistent Systems On Massively Parallel Computer, M. Hochbruck, Universitaet Karlsruhe, Germany, NASA Ames Research Center, June 8, 1993.

MICHAEL KUMBERA

Developing the NAS Multi Grid (MG) benchmark on the CM-5, M. Kumbera, University of Wisconsin-Milwaukee, NASA Ames Research Center, August 20, 1993.

NOËL NACHTIGAL

Iterative Methods in Linear Algebra, N. Nachtigal, Stanford University, Stanford, CA, July 12, 1993.

Iterative Methods in Linear Algebra, N. Nachtigal, Argonne National Laboratory, Argonne, IL, March 26, 1993.

An Implementation of the QMR Method Based on Coupled Two-term Recurrences, N. Nachtigal, Oak Ridge National Laboratory, Oak Ridge, TN, February 22, 1993.

JOSEPH OLIGER

Presented two weeks of lectures on *Composite Adaptive Grid Methods for Partial Differential Equations* as an invited participant in Namkai Institute's "Special Year in Numerical Analysis," Tianjin, China, May 1993.

Presented lecture series on *Adaptive Component Grid Methods for Time Dependent PDE's* as a keynote lecturer at the Institute for Mathematics and its Applications Summer Program, University of Minnesota, July 1993.

PELLE OLSSON

Stable Boundary Conditions for High Order Finite Difference Methods on Domains with Non-smooth Boundary, P. Olsson, Uppsala University, Sweden, NASA Ames Research Center January 26, 1993.

Energy Estimates for High-Order Finite Difference Methods on Non-smooth Domains, P. Olsson, 1993 SIAM Annual Meeting, Philadelphia, PA, July 12-16, 1993.

Summation by Parts, Projections, and Stability, P. Olsson, Stanford University, April 19, 1993.

Summation by Parts, Projections, and Stability, P. Olsson, UCLA, November 9, 1993.

BARNEY PELL

A Strategic Metagame Player for General Chess-Like Games, B. Pell, AAAI Fall Symposium on Games: Learning and Planning, held in Raleigh, NC, October 22-24, 1993.

J. RAMANUJAM

Beyond Unimodular Transformations, J. Ramanujam, Louisiana State University, NASA Ames Research Center, January 12, 1993.

ABHIJIT SAHAY

Log P: Towards a Realistic Model of Parallel Computation, A. Sahay, UC Berkeley, NASA Ames Research Center, March 18, 1993.

ROBERT S. SCHREIBER

Alignment and Distribution of Arrays in Data-parallel Programs, R. Schreiber, DIMACS Workshop on Models, Architectures and Technologies for Parallel Computation, September, 1993; COMPCON, San Francisco, February, 1993.

Subway: A Communication Compiler for the Maspar, R. Schreiber, DIMACS Workshop on Parallel Algorithms for Unstructured and Dynamic Problems, June, 1993; MIT Parallel Computing Seminar, October 1993; Cornell Theory Center, October, 1993; Conference on Parallel High-Performance Applications, Royal Institute of Technology, December, 1993.

L. RIDGWAY SCOTT

Scalable Molecular Dynamics Computations, L.R. Scott, University of Houston, NASA Ames Research Center, January 28, 1993.

BJÖRN SJÖGREEN

Accuracy in the Use of Finite Difference and Finite Volume Methods for Compressible Navier-Stokes Computation, B. Sjögreen, March 25, 1993

PETER SU

Efficient Parallel Algorithms for Proximity Problems, P. Su, Duke University, NASA Ames Research Center, January 6, 1993.

EDRISS S. TITI

4 Part Lecture Series: *Methods for the Reduction of PDEs to ODEs and Spurious Numerical Solutions for Dissipative PDEs (Parts I - III), and Existence of Global solutions for 3D Navier-Stokes Equations in the Presence of Symmetry (Part IV)*, E. Titi, University of CA, Irvine, NASA Ames Research Center, April 29, 1993

VI. OTHER ACTIVITIES

RUPAK BISWAS

R. Biswas was selected to receive the NASA Ames Contractor Council Excellence Award in recognition of his outstanding contribution to the mission of Ames Research Center October 14, 1993.

WRAY BUNTINE

Taught a tutorial on *An Introduction to Statistics and Decision Theory,* at the Fourth International Workshop on Artificial Intelligence and Statistics, Ft. Lauderdale, January 5th, 1993.

Co-taught a tutorial on *Learning from Data: A probabilistic framework,* W. Buntine, with P. Smyth (JPL) at the Uncertainty in Artificial Intelligence Conference, Washington, July 9, 1993. Also member of the program committee.

Member of the program committee, *National Artificial Intelligence Conference,* Washington, D.C., July 11-15, 1993.

Member of the program committee, *6th Annual ACM Conference on Computational Learning Theory,* Santa Cruz, CA, July 26-28, 1993.

Co-taught a graduate class *Computer-based Probabilistic Reasoning (CS354)* at Stanford University, Dept. of Computer Science, W. Buntine, with P. Cheeseman, Fall 1993.

SIDDHARTHA CHATTERJEE

Attended and presented a paper at the Twentieth ACM SIGPLAN-SIGACT Symposium on *Principles of Programming Languages,* Charleston, SC, January 10-13, 1993.

Attended and presented a paper at the Fourth ACM SIGPLAN Symposium on *Principles and Practice of Parallel Programming,* San Diego, CA, May 19-22, 1993.

Co-taught a course on the *Programming Language NESL, DAGS'93 School on Parallel Programming,* S. Chatterjee with G. Blelloch, Dartmouth College, Hanover, NH June 24-27, 1993.

Joint USRA/Xerox patent application no. 08/104,755 title, *Mobile and Replicated Alignment of Arrays in Data-Parallel Programs,* filed, August 11, 1993.

Attended and presented paper at the *Sixth Annual Languages and Compilers for Parallelism Workshop,* Portland, OR, August 12-14, 1993.

Developed and co-taught (with K. Yelick) CS 294-2, A graduate Seminar on Parallel Programming Languages, University of California, Berkeley, CA, Fall semester 1993.

SIDDHARTHA CHATTERJEE (cont'd)

Participated in *IBM HPF Customer Partnership Meeting,* IBM Santa Teresa Laboratory, San Jose, CA, November 11-12, 1993.

Attended *Supercomputing'93,* Portland, OR, November 15-19, 1993.

PETER CHEESEMAN

P. Cheeseman, in conjuction with W. Buntine (RIACS) taught a course, *CS354 Probabilistic Methods in Computing,* at Stanford, Fall quarter, 1993.

DAN FENG

Participate in discussions of a newly-formed group, *NASA Ames Optimization Group,* and also presented a seminar. The group has regular meetings that attract dozens of people who are interested in optimization from many Ames branches and contractors.

MARJORY J. JOHNSON

Organized and led panel discussion on *Applications for Global Data Networks,* 1st Global Data Networking Conference, Cairo, Egypt, December, 1993.

Member, program committee for *1st Global Data Networking Conference,* Cairo, Egypt, December, 1993.

Member, program committee for *14th International Conference on Distributed Computing Systems,* to be held in Poland, June, 1994.

Member, program committee for *5th IFIP Conference on High Performance Networking,* to be held in France, June, 1994.

NOËL M. NACHTIGAL

Attended *XII Householder Meeting,* Lake Arrowhead, Ca, July 13-18, 1993.

JOSEPH OLIGER

Professor of Computer Science, Stanford University.

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

Chairman, Computing Division Advisory Committee, National Center for Atmospheric Research.

BARNEY PELL

Member, *Organizing Committee for, AAAI Fall Symposium on Games: Learning and Planning,* held in Raleigh, NC, October 22-24, 1993.

Presented, *A Strategic Metagame Player for General Chess-Like Games,* AAAI Fall Symposium on Games: Learning and Planning, held in Raleigh, NC. October 22-24, 1993.

With Jens Krause (Mount Alison University, New Brunswick, Canada), Pell is developing a computational model of *Emergent Aggregation Behavior in Animals* (e.g., schooling, flocking, herding).

ROBERT S. SCHREIBER

Began working as Chair of the *Seventh SIAM Meeting on Parallel Processing for Scientific Computing,* 1993.

Finished term as Chair of the *SIAM Activity Group on Supercomputing.*

Elected *SIGNUM Vice-chair.*

Program Committee of the *International Conference on Supercomputing.*

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

Co-instructor in tutorial of *Scalable Parallel Algorithms for Multicomputers,* Supercomputing '93.

Co-instructor in tutorial of *High Performance Fortran for Programmers and Implementors,* Supercomputing 93.

Attended Sixth SIAM Meeting on *Parallel Processing for Scientific Computing.*

Attended Salishan Conference on *High Performance Computing.*

Attended *Federated Computing Research Conference.*

Organized session on *Parallel Eigenvalue Computation,* Householder Meeting on Numerical Algebra.

THOMAS SHEFFLER

Attended the *NASA Distributed Computing Workshop,* NASA Ames Research Center, October 18-19, 1993.

Represented RIACS at the *IBM High Performance FORTRAN Partnership Meeting* IBM Santa Teresa Laboratory, November 11-12, 1993.

Participated in a weekly graduate seminar on *Parallel Programming Languages* at University of California, Berkeley.

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).

Evangeline Tanner, Administrative Assistant I (4/5/90 - 5/7/93).

Jessica Casillas, Systems Administrator (11/25/92 - 3/23/93).

RIACS SCIENCE COUNCIL

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

Dr. Dennis B. Gannon, (Covener), Director, Center for Innovative Computer Applications (CICA), Indiana University, Bloomington, Indiana 47405.

Dr. Joseph Flaherty, Chairman Department of Computer Science, Rensselaer Polytechnic Institute, Troy, New York 12080.

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

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

Dr. David Gottlieb, Division of Applied Mathematics, Brown University, Providence, Rhode Island 02912.

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

Dr. Thomas H. Pulliam, NASA Ames Research Center, MS 202A-I, Moffett Field, California 94035.

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

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

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

SCIENTIFIC STAFF

Wray Buntine, Ph.D., Computer Science, University of Technology, Sydney, Australia, 1992. Mathematical and probabilistic modeling of problems in intelligent systems (1/2/90 - 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 - present).

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

Steven Suhr, - Research Associate, Computer Science, Stanford University, programming languages (5/1/93 - present).

VISITING SCIENTISTS

Fernando Alvarado, Ph.D. - Professor, University of Wisconsin, Madison. Sparse Matrices/Parallel Computation/Power Systems (7/6/93 - 7/20/93).

Marsha Berger, Ph.D. - New York University. Computational fluid dynamics; parallel computing (7/1/93- 8/31/93).

Petter Bjorstad, Ph.D. - Professor of Computer Science, University of Bergen, Norway. Parallel Computing, Domain Decomposition Algorithms for PDE's (7/1/92 - 6/10/93.)

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

Antony Jameson, Ph.D. - McDonnell Professor of Aerospace Engineering, Princeton University. Numerical Methods, Computational Fluid Dynamics, Computational Sciences (9/7/93 - present).

Sir James Lighthill, Ph.D. - Research Fellow, University College of London. Work on aeroacoustics in general, and also on the noise field of jets at high supersonic speeds in particular (8/2/93 - 8/6/93).

Yu-Tai Lee, Ph.D. - Naval Architect, Dept of the Navy, David Taylor Research Center. Conduct flow prediction through two flapping airfoils interacting with a downstream stationary airfoil (8/9/93 - 8/22/93).

VISITING SCIENTISTS (cont'd)

Ki Dong Lee, Ph.D. - Associate Professor, University of Illinois. Computational Aerodynamics, Design Optimization, Grid Generation and Adaptation (7/21/93 - 8/13/93).

Jingke Li, Ph.D. - Assistant Professor, Portland State University. Compiler techniques for Parallel Computers (6/15/93 - 9/15/93).

Kenneth G. Powell, Ph.D - Associate Professor, University of Michigan. Computational Aerodynamics (5/3/93 - 5/14/93).

Peter Sweby, Ph.D. - Lecturer in Mathematics, University of Reading, England. The dynamics of discretizations of convection reaction diffusion equations was investigated over a range of physical and numerical parameters (8/9/93 - 8/27/93).

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 (6/24/93 - 9/3/93).

Romesh Jain, Ph.D - Self-employed Consultant. Incorporation of a multigrid technique in the GRAPE-3D code to accelerate the convergence of its solution for higher convergence tolerances as well as for finer grids (5/17/93 - 11/10/93).

POST-DOCTORAL SCIENTISTS

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

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

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).

Noël M. Nachtigal, Ph.D. - Massachusetts Institute of Technology. Iterative methods for large, sparse, non-Hermitian linear systems (9/2 /91 - 8/31/93).

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).

Barney Pell, Ph.D. - University of Cambridge, England. Computer biology, artificial intelligence, machine learning, automatic scheduling, multiple autonomous agents, and strategic reasoning (9/27/93 - present).

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

STUDENTS

Ross Biro - Mathematics, Stanford University. Algebraic topology, computational homotopy theory, homological algebra (6/14/93 - 12/31/93).

Michael Kumbera - Computer Science, University of Wisconsin, Milwaukee. Parallel architectures and algorithms (5/26/93 - 8/20/93).

Justin Paola - Electrical Engineering & Computer Science, UC Berkeley, (1990). Image processing/remote sensing, University of Arizona (5/26/93 - 7/27/93).

Leonid Oliker - Computer Science, University of Colorado. Compilation of data parallel programs (6/1/93 - 8/6/93).

Brian Rogoff - Scientific Computing/Computational Mathematics, Stanford University. Scientific visualization (6/14/93 - 9/15/93).

Steven Suhr - Computer Science, Stanford University, Programming Languages (7/1/92 - 4/30/93).

CONSULTANTS

Roland Freund - Technical Staff, AT&T Bell Laboratories. Iterative methods for matrix computations. Numerical solution of partial differential equations (5/26/93 - 6/2/93).

Rose Gamble - Assistant Professor, University of Tulsa. Artificial intelligence (knowledge-based systems). Formal methods, software engineering, concurrent program design (7/26/93 - 8/25/93).

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).

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

Jeffrey D. McDonald - Applications Specialist, MASPAR Computer Corporation. Computational fluid dynamics (5/1/92 - 12/31/93).

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