DOE/ER/25158--T1

FINAL REPORT

August, 1994

High Performance Computing Equipment for Environmental Remediation Modeling and First Principles Simulation of Materials Properties

co-Principle Investigators:
James Glimm, W. Brent Lindquist

State University of New York at Stony Brook

Grant #DE-FG02-93ER25158
Scientific Computing Staff
Office of Basic Energy Research
Department of Energy

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Contents

1 Summary

A 56-node Intel Paragon parallel computer was purchased with major support provided by this grant, and installed in July, 1993, in the Center for Scientific Computing, Department of Applied Mathematics and Statistics, SUNY - Stony Brook. The targeted research funded by this proposal consists of work to support the Stony Brook and Brookhaven National Laboratory contributions to the Partnership in Computational Science (PICS) program; namely environmental remediation modeling of ground water transport, Car-Parrinello first principles molecular dynamics calculations, and the supporting development of the parallelized VolVis graphics package. Research accomplishments to date for this targeted research is discussed in §2. This computer has also enabled or enhanced many other projects conducted both by the Center for Scientific Computing and by the Department of Applied Mathematics and Statistics. These other projects include two- and three-dimensional gas dynamics using front tracking, other molecular dynamics applications, kidney modeling, and global optimization techniques applied to DNA-protein interactions. Technical summaries of these additional projects are presented in §3.

The targeted research includes users from the Departments of Applied Mathematics and Computer Science at SUNY - Stony Brook, as well as staff scientists from the Departments of Physics and Applied Sciences at Brookhaven National Laboratory. The additional projects involve university faculty from the above departments as well as the Departments of Physics and Chemistry. Regular users of this machine currently include 10 faculty members, 8 postdoctoral fellows, more that 12 PhD students and approximately 8 staff members from BNL.

2 Targeted Research Accomplishments to Date

2.1 Environmental Remediation Modeling of Ground Water Transport

The contribution of Stony Brook researchers in this area of the PICS effort consists of three broad tasks: encorporation of front tracking capabilities into the GCT code being developed; development, selection and encorporation of effective parameter estimation technologies; testing of the GCT code on field sites.

Encorporation of Front Tracking Capability

Two dimensional front tracking capability was encorporated into version 1.0 of the PICS GCT code [2]. It supports the ability to run both in scalar and parallel modes. Encorporation of the front

tracking options involves addition to the transport part of the GCT code. Only minor changes to the GCT code were necessary, though further enhancement will be achieved in later versions. In recognition of structural changes in future versions of the developing GCT code, encorporation of the front tracking option was achieved by establishing communication between the flow component of GCT 1.0 and the front tracking routines. Where possible, communication was restricted through function calls. This maintained a separation between the flexible data structures required for front tracking and the more rigid FORTRAN vectors employed in the GCT code. The resulting code consists of GCT with a run time front tracking option. The I/O remains that of GCT unless the front tracking option is picked, in which case an additional input file is required. The front tracking routines produce their own output files, and do not interfere with the GCT I/O. Choosing the front tracking option places constraints on the type of problem that can be run. The front tracking option supports a two dimensional domain embeddable in a rectangle. The corresponding GCT domain should be a three dimensional block with homogeneous physics in the third direction. The fluid formulation of the front tracking and GCT code are compatible only in the saturated regime with the GCT diffusion and reaction coefficients set to zero. Unsaturated domains and non trivial reaction coefficients will be implemented in the next upgrade of this implementation. Front tracking is not defined for non zero diffusion coefficients.

The resolution enhancing capabilities of the front tracking code were dramatically demonstrated in calculation of tracer flow along a channel defined by neighboring impermeable zones. In addition to eliminating any numerical diffusion of the tracer front, front tracking enhancements also prevent 'leakage' of the flow into the impermeable zones. These results have been presented in a briefing to the Office of Management and Budget by John Cavalini on behalf of the DOE Office of Scientific Computing and by the ORNL Center for Computational Science at Supercomputing '93 in Portland, OR.

Continued Front Tracking Development

Development of the front tracking algorithms continues in two major areas, increased robustness and refinement of the two dimensional parallel algorithms and the development of three dimensional front tracking algorithms, which of necessity, are parallel algorithms.

Though initially limited to $n \times 1$ domain decomposition strategies, continued refinement of the two dimensional parallel algorithms has resulted in code that is now capable of handling $n \times m$ domain decompositions with increasingly complex wave interactions across subdomain boundaries. This robustness is allowing us to explore more complex flow behaviors. We are currently building

a parallel version of the point redistribution algorithm which redistributes the points along the one dimensional moving front tracking grids. This version will have improved robustness, and guarantee that the parallel redistribution produces the same results as serial redistribution.

We are continuing development of three dimensional front tracking, as part of a GCT implementation planned for 1995. A working prototype version of the three dimensional front tracking code is now in place. This prototype allows the testing of the geometry and computational libraries used to extend the front tracking code to three dimensional flows. The implementation is also fully parallelized for distributed memory parallel computers (ie. the Paragon). The transition from a prototype code into a working robust scientific code is now the highest priority in the program development work.

In this study, we continue our work on the development of front tracking algorithms for solving nonlinear hyperbolic systems of conservation laws in three space dimensions.

Extending ideas from our two-dimensional algorithm, we use an underlying uniform grid with additional lower-dimensional grids, ie. surfaces, introduced at appropriate locations for tracking discontinuities. A time step consists of a normal direction sweep procedure to update both the discontinuity location and the state variables to the left and right of the discontinuity. A tangential sweep procedure is then applied to further modify the tracked front states to complete the integration of the dynamical equations at the interface. An operator splitting finite difference method is then used to update states on grid points away from the tracked fronts. The parallelization of this three dimensional code (as in our two dimensional code) uses standard domain decomposition ideas with overlapping domains. The code has been tested on Intel iPSC/860 and Paragon computers. As the code is still in a development stage, we have not run performance studies of the algorithms, rather we have concentrated on validating the output of the algorithms.

Parameter Estimation Technology

Mathematical parameter estimation and conditioning (kriging) techniques have traditionally produced maximum-likelihood model geological realizations that are recognized as having spatial variation that is too smooth, and not necessarily predictive of flow and transport properties, even for average behavior. We have developed a substantially new statistical theory for permeability spatial distribution [12,13.14,15,20,21,22]. This theory produces permeability realizations with controllable random spatial variation yet allows for specification of the permeability at given measurement locations. The theory allows understanding of the scaling properties of the permeability with change of length scale. We have analyzed the predictions of this theory for the resultant dispersion

in single phase flow. The theory gives consistent flow dispersivity behavior in low order perturbation theory, and in renormalized perturbation theory, both in the Lagrangian and Eulerian pictures. The predicted dispersivities have been checked numerically and seem to be valid for moderately large strengths of heterogeneity. The theory is consistent with observed field dispersion data [11].

This theory will be employed practically for conditional simulations, in which sparsely known geological data will serve to restrict the ensemble from which the random geologies are drawn [12]. We have developed numerical code to produce two and three dimensional conditioned ensembles of permeability data having the spatial distribution properties developed in the above theory but which match measurements at a set of given locations. This kriging code is implemented on the Paragon and can be called as subroutines to produce permeability realizations in a format directly readable as input by, for example, the GCT code. Additionally this code can be called dynamically by the Texas A&M history matching code which attempts to produce an optimum permeability field based upon minimization procedures.

We also mention another use of this permeability theory, that of guiding the acquisition of data by emphasizing the value of data containing multiple length scales. The analysis of field data from a spatial distribution point of view is hampered by lack of multiscale measurements, and by measurement inconsistencies between different instruments which have different intrinsic measurement length scales [11].

Field Testing of the GCT Code

We have completed the modification of a data set describing the Peconic River Watershed on Long Island, NY to GCT input form. This field site is designed to test the current capabilities of GCT 1.2, namely single and two phase flow with no chemistry or radionuclide decay processes. After initial computation on the 56 node paragon at Stony Brook, large grid runs will be done on the Paragon at ORNL and the results critically compared to industrial-standard runs that have been done on this site utilizing the Modflow code on a 486-processor based PC. The field site covers a region of Long Island approximately 25 kilometers by 16 kilometers, running east–west from Brookhaven National Laboratory to Flanders Bay, and from the north–south from Long Island Sound to the Atlantic Ocean. In terms of human impact, the region is relatively pristine, containing the Long Island Pine Barrens, one of the remaining undeveloped regions on the Island, and subject to intense development pressure. The unique dependence of Long Islands drinking supply solely upon its underlying system of three major aquifers, and the continued preservation of a supply of potable water to serve the entire population (currently approximately 2.5 million) of Suffolk and

Nassau counties requires extremely farsighted decisions concerning the impact of development upon this large watershed. In conjunction with colleagues in Earth Sciences we plan to investigate the mass balance of several chemical species (ie. nitrogen, sodium) through this watershed region in order to achieve an understanding of the current distribution of chemical loading on this water system.

2.2 Car-Parrinello First Principles Molecular Dynamics Calculations

This project is directed at developing new scalable parallel techniques and algorithms for performing very large scale simulations of key material properties. The goal is to apply these techniques to solving problems that are well beyond the limits of present methods in order to guide the discovery and the development of advanced materials. Enabling software being developed allows large scale quantum simulations using Car-Parrinello (or Ab Initio Pseudopotential) first principles molecular dynamics. The immediate scientific applications are to predict the properties of metals and alloys especially in complex systems containing more than 100 atoms per unit cell. It grew out of previous studies at Brookhaven which utilized the Car-Parrinello method to calculate properties of liquid metals. The codes used previously are being rewritten in a more easily parallelizable form so that they can be run efficiently on the Paragon. Simulations on technologically relevant metallic alloys, such as aluminum-magnesium, and the extension of the calculations to non-zero temperature are being planned.

In the plane wave pseudopotential method employed here the wave functions are expanded in Fourier components. Typically, 100 plane waves are required per atom in the unit cell of a periodically repeated crystal. This means that for interesting alloy systems which might contain 100 atoms, solving the Schrödinger equation amounts to diagonalizing a matrix which is $10,000 \times 10,000$, beyond the range of standard methods on serial computers. However, the number of eigenvalues required is of the same order as the number of atoms – in this case 100. This suggests the use of iterative methods for solving the Schrödinger equation. These methods require a set of wave functions equal to the number of occupied energy levels, which vastly reduces the storage requirements. They also require evaluation of the action of the Schrödinger matrix on the individual eigenvectors (termed an "update"). In the plane wave pseudopotential approach this can be done with fast Fourier transforms.

Our strategy for parallelizing these codes is as follows. We assign a group of states to each processor. In general there would be one state per processor, but more are possible. We assume the eigenvector for that state will fit in memory. For the typical case of 10,000 (complex) Fourier

components, this is only 1.28 Mbytes. The iterative update of the wavefunctions can then proceed completely in parallel. There is no problem with latency because all states will require the same amount of time to update. Further, there is no need for data transfer among processors for this step. As this is the CPU intensive part of the calculation, the overall performance should scale with the number of processors. This will be true provided the memory requirements for each node are not exceeded.

The major activity of this project has been the development of a new pseudopotential code [9] based on iterative diagonalization which is especially designed for the treatment of metals [7], requiring the extension of the original Car-Parrinello method. Fermi statistics had to be incorporated, and multiple k-point capability for Brillouin zone averages included. In addition, significant time was needed to test the convergence of the iterative diagonalization scheme and to obtain a preliminary parallel version which was run on the Paragon .

We have also developed [3] a new iterative minimization technique which resembles the preconditioned conjugate gradient method. In this simple method, the "update" is followed by a subspace diagonalization and it is possible to prove that it always converges.

In addition to the first principles calculations, which are the main thrust of this work, we have performed simulations on large silicon clusters using a tight binding molecular dynamics approach [19]. This code is a simple example of an electronic structure method which has been ported to the Paragon already.

2.3 Volume Based Visualization

At Stony Brook, Ari Kaufman and his students have been developing the underlying algorithms of a 3-D volume visualization system [1,16], employing both regular and irregular grids. They have developed and implemented a set of fast volume viewing algorithms, as well as a host of other visualization tools. Specifically, they have developed a hierarchical data structure which is exploited by the viewing algorithms to produce the highest quality output for a given interaction speed and a given machine. The algorithms include a volumetric navigation projection algorithm, the PARC (Polygon Assisted Ray Casting) algorithm, and a volumetric ray tracing algorithm.

These algorithms have been incorporated within VolVis [1], a versatile volume visualization software system, which includes a variety of manipulation and rendering tools, flexible I/O, a volumetric navigation facility, key-frame animation generator, quantitative analysis tools, and a generalized protocol for communicating with 3-D input devices (e.g., a Spaceball). The current version of VolVis, version 1.0, which includes over 75,000 lines of C code, is running on a variety of

workstation platforms, including Silicon Graphics, Sun Sparcstation, and HP 400. Although VolVis has been tested most extensively on the Silicon Graphics family (five different machines: a personal Iris, 420, 240, Elan, and Extreme), it can be compiled on any machine that has X windows and Motif with or without a graphics package (such as GL, Starbase). We plan to port it during 1994 to IBM RS6000 and IBM PC.

The output visualization algorithms have also been parallelized to work on a distributed heterogeneous network of workstations, on the Intel iPSC/860 and Paragon architectures, and on a shared memory multiple-processor SGI. Specifically, the implementation of our ray casting algorithm on the iPSC/860 yielded a sub-second projection time for a moderate size grid. During the past year the code has been ported to the Paragon; the version 2.0 implementation of VolVis on the Paragon will be available by 12/31/94. Additionally the Paragon has been used for experimenting with parallel versions of the navigation projection algorithm.

The case studies for testing VolVis algorithms have included a variety of data sets, with a focus on output visualization of the materials properties and 3-D volume visualization of the groundwater front tracking data. The groundwater group at Stony Brook has used several version of VolVis and provided some valuable feedback. We have also communicated with the visualization group at Rice to coordinate data formats, so it will be possible to link our visualization systems. We anticipate a link of VolVis with Rice surface-based visualization system by 12/31/94. At Stony Brook, VolVis has been tested successfully with data sets from Ames Labs (periodic minimal function and molecular dynamic data sets of Carbon 60 - Che Ting Chan), from Solid State Division ORNL (3D mesh of charge densities of a silicon crystal - Victor Milman), and from SUNY-BNL (sandstone - Brent Lindquist).

3 Technical Summaries - Additional Projects

3.1 Gas Dynamics - Front Tracking

The development of a two-dimensional parallel interface method for multi-phase flow problems has been completed. Parallel efficiency exceeding 80% has been achieved on the Paragon. The increased problem size accessible through parallelization has enabled the reaffirmation of the consistency of two-phase flow mixing rate with experiments for incompressible flow and, for the first time, allowed determination of the mixing rate for compressible flow [5,10]. In addition, a comparative statistical analysis of a two phase flow mixing problem [6] is close to completion, and an extensive study of Richtmyer-Meshkov instability has been conducted.

Further progress has been made to enhance the two dimensional front tracking method, which has completed its evolution from a prototype numerical implementation into a working scientific tool. All significant simulations of gas dynamic flows conducted by a group involving James Glimm, J. Grove, R. Holmes, and B. Boston now use the parallel version of this code. The increased resolution made possible by this parallelization has led to important discoveries in the Richtmyer-Meshkov problem of a shock accelerated interface. The simulations clearly show the effects of compressibility on the initial behavior of the shock accelerated interface and explain why theoretical models that ignore this effect do not agree with experiments. These computations require fine grid resolution in the region between the reflected and transmitted waves produced by the shock refraction. At early times, the narrowness of this region puts a large burden on the computation. The resolution required to isolate the nonlinear wave structures affecting the interface growth rate was only made possible by the combination of front tracking and parallelism.

3.2 Flow in Porous Media

We have implemented a Glowinski-Wheeler domain decomposition for mixed finite element (Raviart Thomas basis functions) solution of elliptic equations. We use a Bramble-Pasciak type preconditioner for the resultant system of equations [17]. This elliptic solver is designed to work with our two dimensional front tracking, domain decomposition package for solving an associated hyperbolic system of equations as part of the equations for flow in porous media. Currently this project is in the stage of robustness testing of the parallelized algorithms. The resultant parallelized code will later be used for the computation of physically interesting porous media flow problems.

3.3 Porous Media Microgeometry Characterization

Fundamental theories of rock structure are limited by the absence of high resolution, pore level, three dimensional images which could be used for statistical analysis. The ability to produce such images in a non-destructive manner would also allow for repeated measurements of dynamic processes such as fluid motion which could be correlated to the medium properties. We are concerned with the geometrical and statistical analysis of the pore structure of such high resolution images produced by collaborators at the National Synchrotron Light Source at BNL. As the three dimensional data sets are large and the geometric structure of the pore complicated, the computational power provided by parallel architectures is crucial to this work. This work has been done using both the Intel i/860 and Paragon parallel computers.

Our main tool is the computation of the medial axis (intuitively, the 'spine') of the void space

in the rock sample. Computation of the medial axis also provides, simultaneously, a measure of a minimum cross sectional area of the void space at any point along its length. We are engaged in the statistical analysis of the resultant pore medial axes computed for three rock samples, Berea sandstone - a 'standard' rock type recognized by the oil industry, Danish chalk - an rock type that is extremely difficult to work with using invasive imaging techniques such as microsanding combined with visual microscopy, and a packed glass bead sample consisting of 100 micron diameter beads - an intensively studied theoretical approximation to rock. To date we have completed analyses of the statistical distributions of the size of disconnected void components. We have also achieved rigorous measurements of the distribution of tortuosities for these real rock samples. We have computed two point correlation functions for both the void spaces and for the grains of the rock samples. We have used Dijstra's algorithm on the medial axis network to compute shortest void space paths through the sample. Currently we are coding network flow algorithms to provide estimates of maximal flow fields through the void network.

3.4 Linear Algebra

In [8] a method for matrix multiplication on distributed-memory MIMD architectures based on Strassen's method is presented. Timing tests, performed on the Paragon, demonstrate that the method realizes the full potential of Strassen's method with a complexity of $4.7M^{2.807}$ on the system-level rather than node-level. The parallel efficiency is nearly 100% when the processor number is divisible by 7. Otherwise, the performance falls back to that of the traditional matrix multiplication method whose complexity is $2M^3$ coupled with the BMR method, a scalable method on distributed-memory machines.

3.5 Renal Modeling

A parallel algorithm for solving the multinephron model of the renal inner medulla has been developed [18]. The intrinsic nature of this problem supplies sufficient symmetry for high-level parallelism on distributed-memory parallel machines. A nearly perfect speedup is achieved by evenly distributing the load and minimizing the cost of communication. Parallelization makes it feasible to study realistic models such as the rat kidney with $30,000$ nephrons. On the 56 node Paragon one can handle more than 2000 nephrons, compared with the 100 nephrons systems that can be studied on a high-end work station.

3.6 Molecular Dynamics for Thin Film Deposition Studies

A portable parallel algorithm has been developed that can study 10^6 particles with multiple interaction ranges. On the 56-node Paragon, a rate of 1 step/minute has been achieved. The code, built upon PVM, is independent of architecture; it can be ported to all distributed-memory MIMD parallel computers without any changes. The communication and load imbalance overhead is minimized; a parallel efficiency of 95% is recorded for up to 56 nodes. The code is currently utilized to study thin film depositions in which the particles interact in a combined form of a Leonard-Jones potential for the incident species, and a Molière potential for particles near the surface of the substrate after bombardment. The algorithm has a special advantage in treating the load imbalance caused by the multiple interaction ranges in this application. Moderate progress has been made for sputtering deposition, while plasma deposition is in an initial stage of development.

3.7 Quantum Molecular Dynamics

The study of quantum molecular dynamics, involving the study of the rotational and vibrational properties and the electronic structure of molecules such as HCN and H_2O_2, requires immense computational power, to solve large eigenvalue problems. Only a dozen major institutions and national laboratories can afford to perform such computations. This project is conducted by F. Webster and G. C. Lo from the Stony Brook's Chemistry Department.

3.8 Protein Folding

This Paragon has been used in the study of a sub-problem of protein folding. An efficient global minimization algorithm has been developed for problems with multiple local minima, which lie at the heart of these biological problems. A study of analysis and prediction of protein binding on DNA, based on pattern recognition of hydrogen bonds [4] is essentially complete. This study has identified the roles of hydrogen bonds in the binding of proteins on DNA. The data set is so large that completion of the work without access to parallel architectures is infeasible.

3.9 Operating Systems

Our site has been selected as a beta site for testing the operating system SUNMOS on the Paragon.

4 References

[1] R. S. AVILA, L. M. SOBIERAJSKI AND A. E. KAUFMAN, *Towards a comprehensive volume visualization system*, in Visualization '92 Proceedings, IEEE Computer Society Press, October, 1992, pp. 13–20.

[2] B. BIELEFELD, W. B. LINDQUIST AND F. TANGERMAN, *Encorporation of Two Dimensional Front–Tracking into PICS GCT 1.0*, State University of New York at Stony Brook, Report No. SUNYSB-AMS-93-07, 1993.

[3] J. H. BRAMBLE, A. V. KNYAZEV AND J. E. PASCIAK,, *A Subspace Preconditioning Algorithm for Eigenvector/Eigenvalue Computation*, In preparation.

[4] G. CAMPBELL, Y. DENG, J. GLIMM, Y. WANG, Q. YU, M. EISENBERG AND A. GROLLMAN, *Analysis and prediction of protein binding on dna: pattern recognition of hydrogen bonds*, Journal of Molecular Biology.

[5] Y. CHEN, Y. DENG, J. GLIMM, G. LI, D. H. SHARP AND Q. ZHANG, *A Renormalization Group Scaling Analysis For Compressible Two-Phase Flow*, Phys. Fluids A, 5 (1993), pp. 2929–2937.

[6] Y. CHEN, J. GLIMM, D. H. SHARP AND Q. ZHANG, *A Two Phase Flow Model of the Rayleigh-Taylor Mixing Zone*, In preparation .

[7] N. CHETTY, M. WEINERT, T. S. RAHMAN AND J. W. DAVENPORT,, *Vacancies and Impurities in Aluminum and Magnesium Using Large Supercell Total Energy Calculations*, In preparation.

[8] C-C. CHOU, Y. DENG, G. LI AND Y. WANG, *Parallelizing Strassen's Method for Matrix Multiplication on Distributed-Memory MIMD Architectures*, Stony Brook Preprint: SUNY-AMS-93-17, Submitted in November 1993.

[9] J. W. DAVENPORT, N. CHETTY, R. B. MARR, S. NARASIMHAN, J. E. PASCIAK, R. F. PEIERLS, T. S. RAHMAN AND M. WEINERT,, *First Principles Pseudopotential Calculations on Aluminum and Aluminum Alloys*, in Proceedings of the TMS Symposium on Alloy Phase Stability and Design, Pittsburgh, PA, October, 1993.

[10] Y. DENG, J. GLIMM AND D. H. SHARP, *Mixing And Chaotic Microstructure*, Los Alamos Science, 21 (1993), pp. 124–132.

[11] L. FONSECA, J. GLIMM AND W. B. LINDQUIST, *An Analysis of Field Data Permeability*, State Univ. of New York at Stony Brook, Report No. SUNYSB-AMS-92-16, 1992.

[12] J. GLIMM AND W. B. LINDQUIST, *Scaling Laws for Macrodispersion*, in Computational Methods in Water Resources IX, vol. 2: Mathematical Modeling in Water Resources, T. F. Russell, R. E. Ewing, C. A. Brebbia, W. G. Gray and G. F. Pinder, eds., Computational Mechanics Publications, Southampton, 1992, pp. 35–49.

[13] ———, *A Theory of Fluid Mixing*, in Nonliner Hyperbolic Problems, A. Donato and F. Oliveri, eds., Lecture Notes on Numerical Fluid Dynamics, 43, Vieweg, Weisbaden, 1993, pp. 291–297.

[14] J. GLIMM, W. B. LINDQUIST, F. PEREIRA AND R. PEIERLS, *The Multi-fractal Hypothesis and Anomalous Diffusion*, Revista Brasileira de Matematica Aplicada e Computacional, 11 (1992), pp. 189–207.

[15] J. GLIMM, W. B. LINDQUIST, F. PEREIRA AND Q. ZHANG, *A Theory of Macrodispersion for the Scaleup Problem*, Transport in Porous Media, 13 (1993), pp. 97–122.

[16] A. KAUFMAN, D. COHEN AND R. YAGEL, *Volume Graphics*, IEEE Computer, 26 (1993), pp. 51–64.

[17] D. MIRKOVIC, B. LINDQUIST AND G. LI, *Domain decomposition for the front tracking reservoir simulator*, in Proceedings of the Seventh International Conference on Domain Decomposition Methods, State College, October 1993.

[18] H. WANG, J. L. STEPHENSON, Y. DENG AND R. P. TEWARSON, *An Efficient Parallel Algorithm for Solving n-Nephron Models of the Renal Inner Medulla*, Computers and Mathematics with Applications, Accepted for publication.

[19] M-L. YEH, J. KIM AND F. KHAN, *Parallel Decomposition of the Tight Binding Fictitious Lagrangian Algorithm for Molecular Dynamics Simulations of Semiconductors*, submitted to Computers in Physics..

[20] Q. ZHANG, *The Transient Behavior of Mixing Induced by a Random Velocity Field*, Stony Brook Preprint SUNYSB-AMS-92-17, 1992.

[21] ———, *The Asymptotic Scaling Behavior of Mixing Induced by a Random Velocity Field*, Adv. in Appl. Math. (1994), to appear.

[22] Q. ZHANG AND J. GLIMM, *Inertial Range Scaling of Laminar Shear Flow as a Model of Turbulent Transport*, Comm. Math. Phys., 146 (1992), pp. 217–229.