Table of Contents Author Index Home

Get Acrobat® Reader Get Postscript Version Get PDF Version

First Parallel Results for a Finite-Element MHD Code on the T3D and T3E

Alice E. Koniges, X. Xu, and P. Covello, Lawrence Livermore National Laboratory,
Livermore, CA 94550
Steve Plimpton, Sandia National Laboratory, MS 1111 Albuquerque, NM 87185-1111
and the NIMROD TEAM

koniges@llnl.gov
www.llnl.gov

ABSTRACT:
NIMROD is a finite-element based magnetic fusion energy code. The code achieves parallelism by spatially decomposing the problem into an unstructured collection of structured blocks. The algorithm yields excellent scaling on both the T3D and T3E with the latter yielding a performance improvement factor of approximately 4.

KEYWORDS: Applications, finite-element, T3D, T3E, MHD

Introduction

NIMROD is a finite-element based magnetohydrodynamic (MHD) fusion energy code. The code is being developed by a team of scientists who are physically located at several different laboratories. Information on the team and the team process is contained in the NIMROD web pages . The focus of parallelization of the NIMROD code is based on the computationally intensive physics kernel. This kernel is written so as to run without modification on single processors or any platform that supports a MPI message-passing style of programming. This includes workstations, traditional vector supercomputers, and essentially all current-generation massively parallel machines. In this paper, we describe how the NIMROD kernel is structured to enable efficient parallelization. We present timings on several parallel machines including the T3D at LLNL, T3E's at NERSC and UT Austin, and the C90 at NERSC.

Parallel Decomposition

NIMROD represents the poloidal simulation plane as a collection of adjoining grid blocks; the toroidal discretization is pseudo-spectral. Within a single poloidal block the grid is topologically regular to enable the usual 2-D stencil operations to be performed efficiently. Blocks join each other in such a way that individual grid lines are continuous across block boundaries. Within these constraints, quite general geometries can be gridded, and parallelization is achieved by assigning one or more blocks (with their associated toroidal modes) to each processor. In parallel, the only interprocessor communication that is then required is to exchange values for block-edge or block-corner grid points shared by other processors. For general block connectivity, this operation requires irregular, unstructured interprocessor communication. We pre-compute the communication pattern and then exchange values asynchronously. This enables the block-connection operation to execute efficiently and scalably on any number of processors. NIMROD uses implicit time stepping to model long-timescale events and thus requires a robust iterative solver. To date, an explicit time-stepping routine and an iterative solver using conjugate gradient techniques have been implemented and tested in parallel for NIMROD. The iterative method uses simple diagonal (Jacobi) scaling as a matrix preconditioner. A second method (currently under parallel development) directly inverts the portion of the matrix residing on each block as a preconditioning step. Both iterative solvers perform their computations on a block-wise basis and thus work in parallel using the block-connection formalism described above. We present timings that illustrate the performance and convergence of both techniques as a function of (1) the number of blocks used to grid the poloidal plane and (2) the number of processors used.

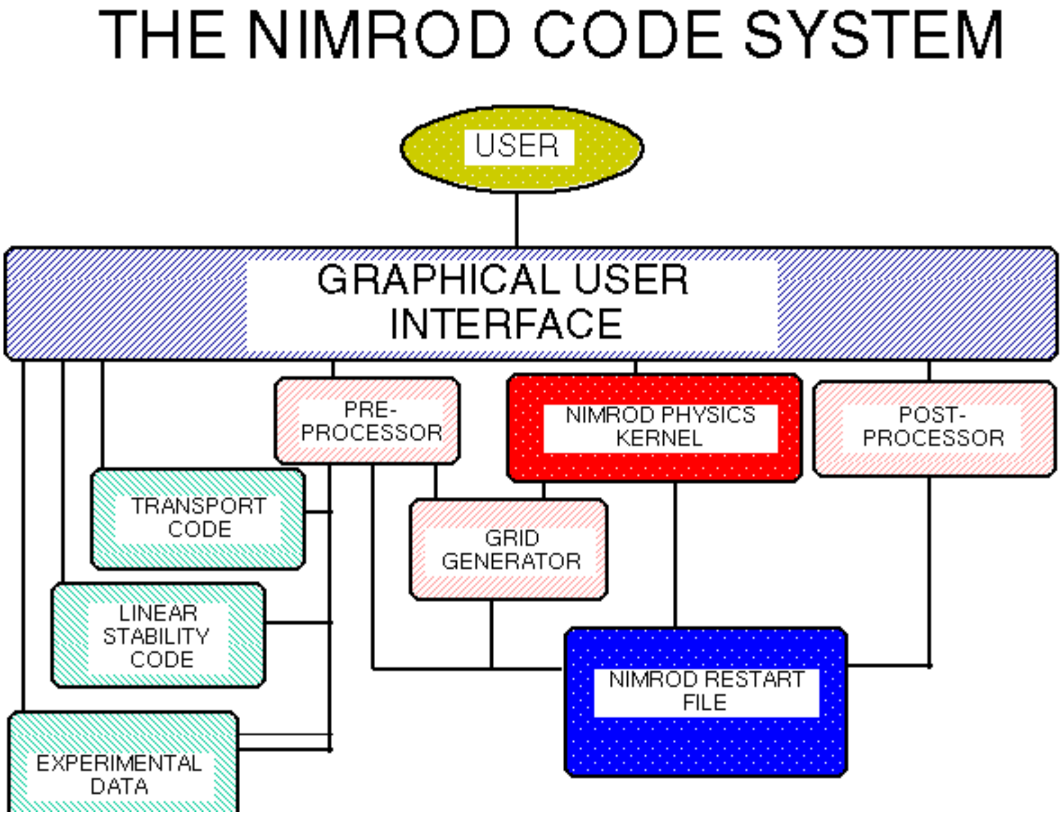

Figure 1 shows the NIMROD Code system.

The NIMROD Physics Kernel

The Nimrod physics kernel is based on solving the nonlinear, two-fluid dynamical equations for electrons and ions simultaneosly with Maxwell's equations. For this approximation, we assume that the ions and electrons have Maxwellian distribution functions and can thus be described as interacting fluids. This approach misses purely kinetic effects such as those related to Landau damping and finite Larmor radius, but includes all the cold plasma wave branches such as Langmuir oscillations, cyclotron, hybrid, and whistler waves, as well as low frequency Alfvén and sound waves[1]. Collisionless, as well as resistive, reconnection is also described by this model. The NIMROD code is presently being developed to solve the equations of such a model in three-dimensional toroidal geomery and time.

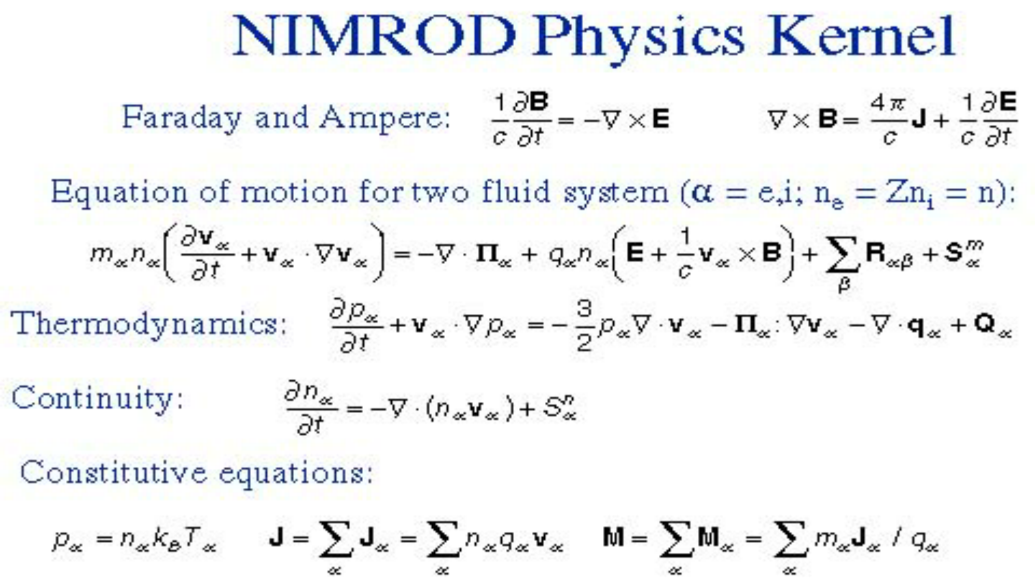

NIMROD Physics Kernel

Faraday and Ampere: $\frac{1}{c}\frac{\partial \mathbf{B}}{\partial t} = -\nabla\times\mathbf{E}$ $\quad\nabla\times\mathbf{B} = \frac{4\pi}{c}\mathbf{J} + \frac{1}{c}\frac{\partial \mathbf{E}}{\partial t}$

Equation of motion for two fluid system ($\alpha = e,i;\ n_e = Zn_i = n$):

$$m_\alpha n_\alpha\left(\frac{\partial \mathbf{v}_\alpha}{\partial t} + \mathbf{v}_\alpha\cdot\nabla\mathbf{v}_\alpha\right) = -\nabla\cdot\mathbf{\Pi}_\alpha + q_\alpha n_\alpha\left(\mathbf{E} + \frac{1}{c}\mathbf{v}_\alpha\times\mathbf{B}\right) + \sum_\beta \mathbf{R}_{\alpha\beta} + \mathbf{S}_\alpha^m$$

Thermodynamics: $\frac{\partial p_\alpha}{\partial t} + \mathbf{v}_\alpha\cdot\nabla p_\alpha = -\frac{3}{2}p_\alpha\nabla\cdot\mathbf{v}_\alpha - \mathbf{\Pi}_\alpha:\nabla\mathbf{v}_\alpha - \nabla\cdot\mathbf{q}_\alpha + \mathbf{Q}_\alpha$

Continuity: $\frac{\partial n_\alpha}{\partial t} = -\nabla\cdot(n_\alpha\mathbf{v}_\alpha) + S_\alpha^n$

Constitutive equations:

$$p_\alpha = n_\alpha k_B T_\alpha \qquad \mathbf{J} = \sum_\alpha \mathbf{J}_\alpha = \sum_\alpha n_\alpha q_\alpha \mathbf{v}_\alpha \qquad \mathbf{M} = \sum_\alpha \mathbf{M}_\alpha = \sum_\alpha m_\alpha \mathbf{J}_\alpha / q_\alpha$$

Spatial Discretization

Spatial discretization in the poloidal plane is accomplished with finite elements and a non-orthogonal Eulerian grid. The toroidal coordinate is treated with a dealiased, pseudospectral representation and Fast Fourier transforms (FFTs). The non-orthogonal poloidal grid consists primarily of unstructured blocks of structured quadrilaterals (RBLOCKs). In the poloidal plane using the RBLOCK representation, the dependent variables are approximated with bi-linear elements, and metric elements are treated as bi-cubic splines. In the core of the plasma the grid is chosen to be closely aligned with the initial axisymmetric ($n = 0$) flux surfaces. (Occasional regridding may be required as the axisymmetric component evolves.) Non-conforming RBLOCKs can be joined with blocks of unstructured triangular elements (TBLOCKs). A TBLOCK is also introduced to represent the region around the magnetic axis, thus avoiding the coordinate singularity inherent in pure flux coordinates. The initial choice to avoid the singularity while avoiding bottle necks in the parallel decomposition was to overlay a patch of RBLOCKS at the axis. This is shown in Figure 3. This led to numerical inaccuracies, however, and is being replaced by a pie-shaped collection of triangle elements, a TBLOCK, as shown in Figure 4.

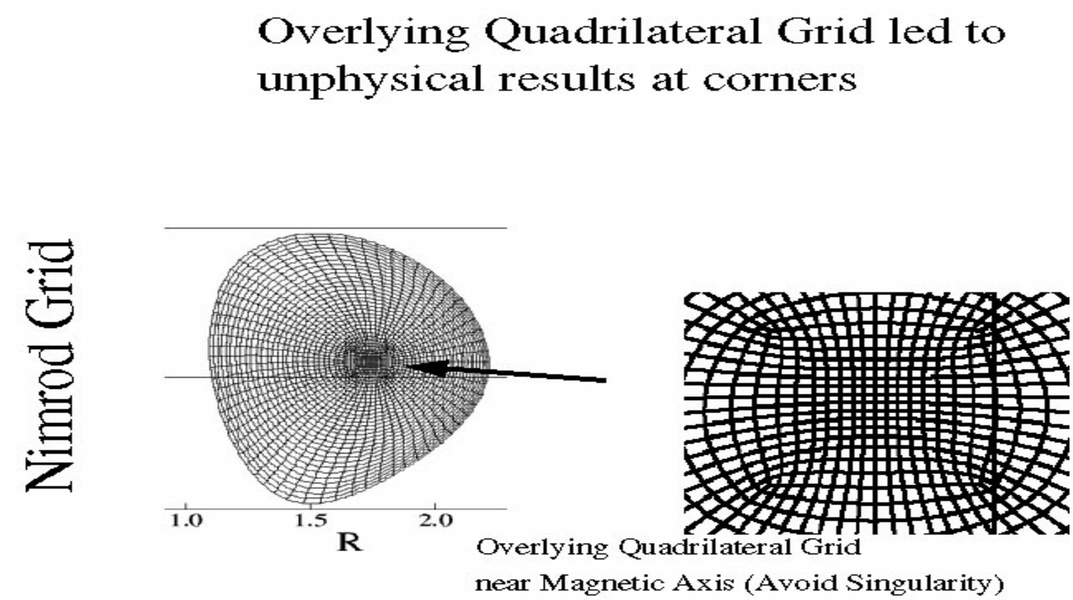

Overlying Quadrilateral Grid led to unphysical results at corners

Overlying Quadrilateral Grid near Magnetic Axis (Avoid Singularity)

Figure 3 shows initial rectangular patch.

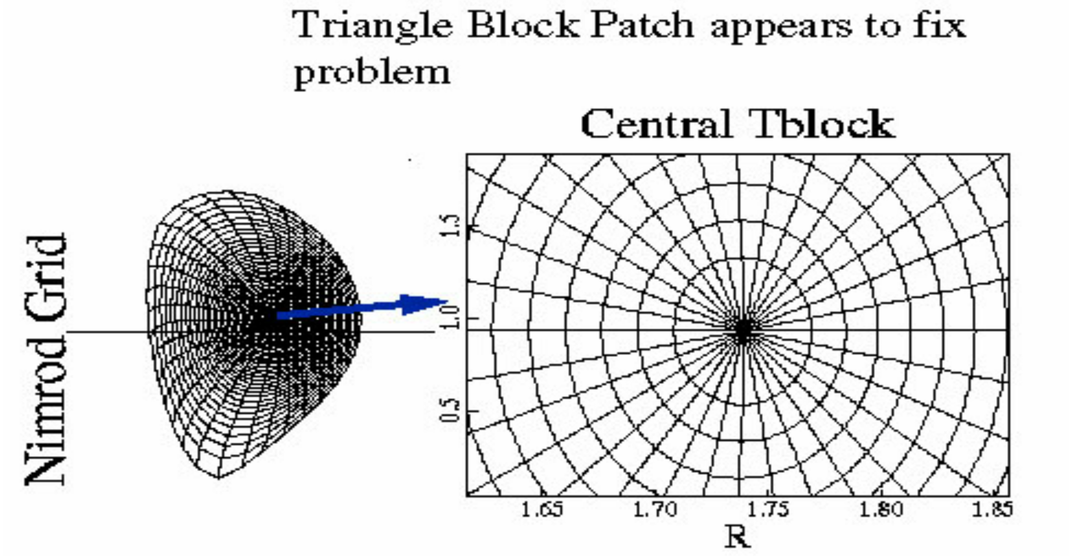

Triangle Block Patch appears to fix problem

Central Tblock

Figure 4 shows latest version with pie-shaped collection of triangular elements at the magnetic axis. Extensive numerical analysis of the spatial and temporal discretization described here has shown that this formulation is free of spectral pollution and the point-to-point "red-black" mode. Details of this analysis are given in Ref. 2.

Parallel Performance

NIMROD has been designed from the beginning to be used on state-of-the-art massively parallel computing platforms. This is significantly different from many other application codes which were ported to the MPP after years of use on a serial (often with vector processors) machine. Collections of RBLOCK or TBLOCK elements can be assigned to different CPUs. Communication between blocks is required primarily for the matrix multiply and the dot product phases of the CG solver. The pre-conditioner for the CG solver uses diagonal scaling on each RBLOCK or TBLOCK. Plans are underway to introduce more sophisticated preconditioning methods for nonlinear problems. Nearly ideal performance scalings with increases in problem size and number of CPUs have been obtained on both the T3D and T3E computers. More details of these scalings are given in the viewgraph pages. A comparison of T3D/E performance with the Cray C90 is given in Figures 5 and 6.

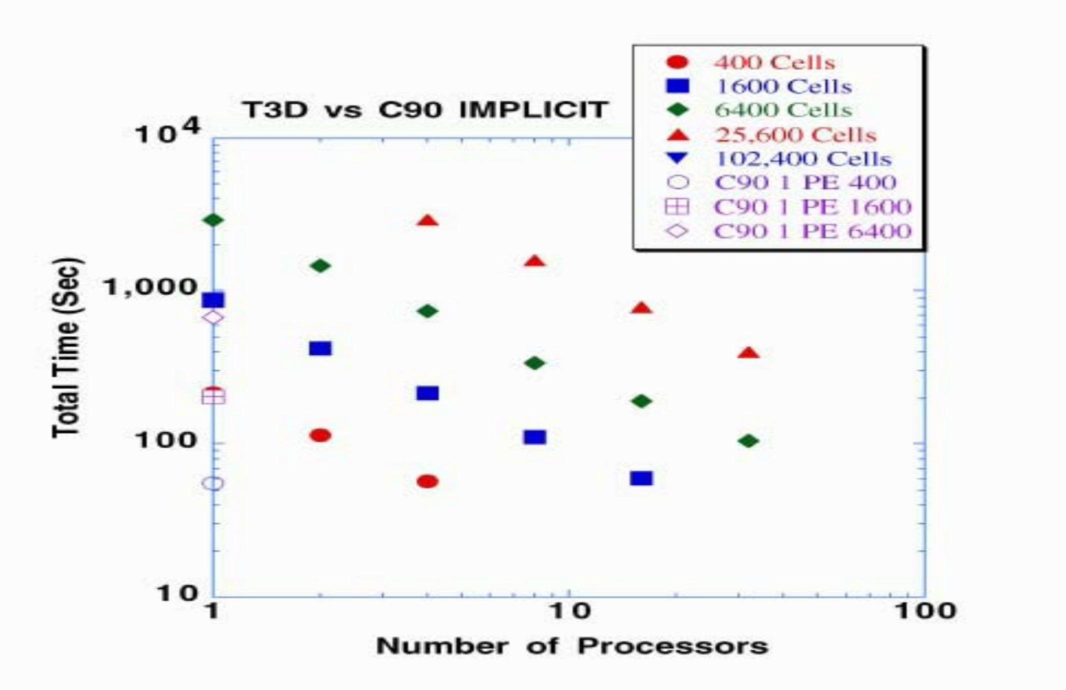

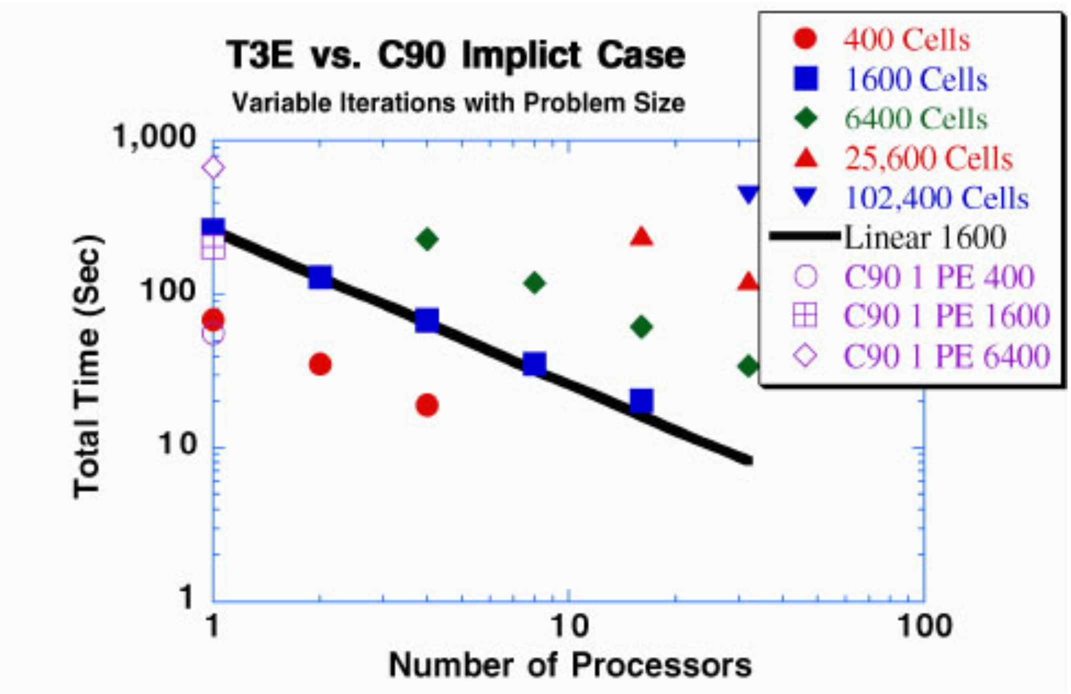

In all cases shown here, the code has not been optimized for a given machine. Optimization for the MPP machines will be performed after a majority of the non-linear physics effects are included. Reference 4 describes some general procedures for optimization on T3D/E-like architectures. The final figures give a scaled speed-up comparison of the T3D and the T3E.

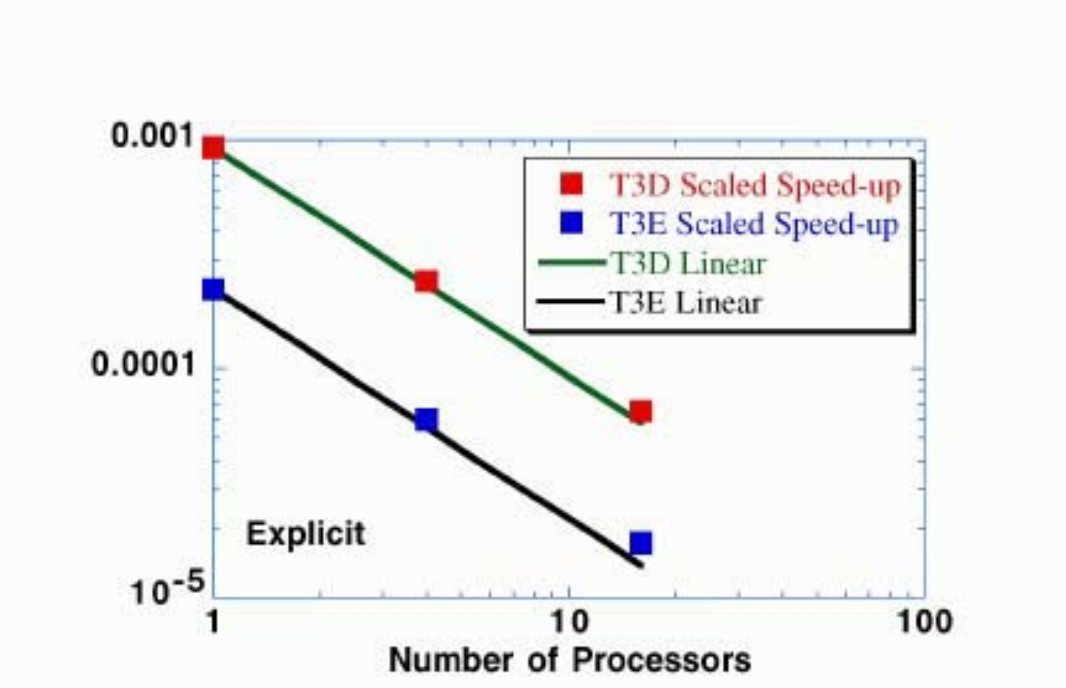

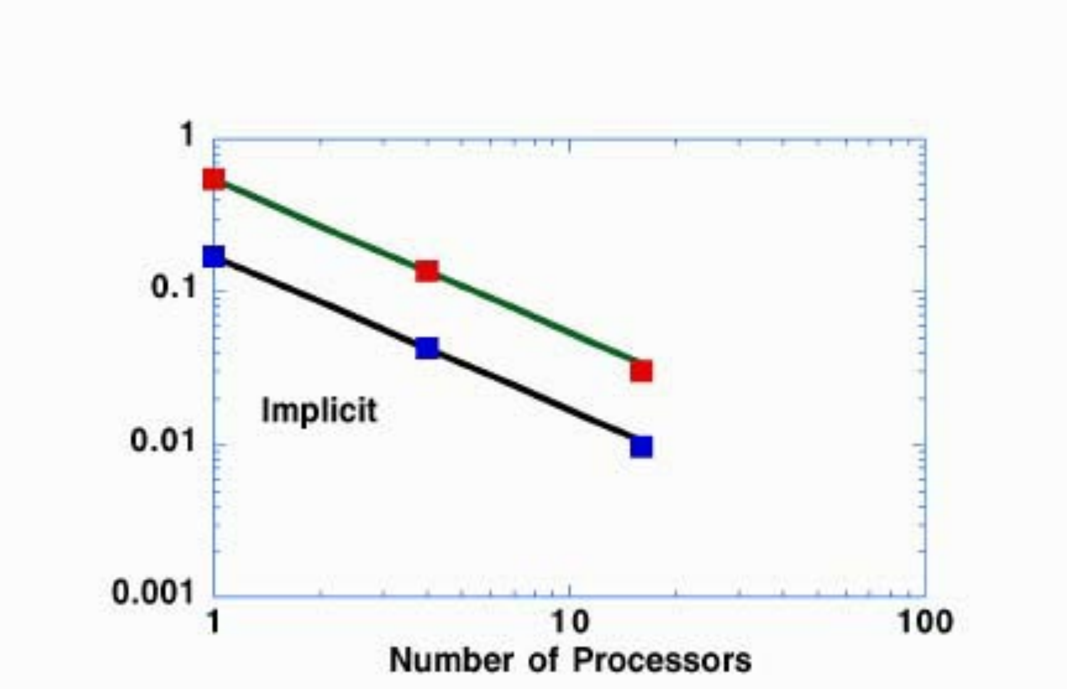

Here, scaled speed-up is speedup divided by problem size. We find that the T3E is roughly a factor of 4 faster which is identical to the performance improvement predicted by the specifications of the T3E which has a faster processor and allows for chaining of operations. The larger cache on the T3E may also contribute to the performance improvement. Scalability is virtually linear for both machines. Additional details can be found in the viewgraphs for the oral version of this presentation.

Acknowledgments

- Work at LLNL for DOE under Contract W7405-ENG-48.
- The DoE MICS Office supported SNL work on solvers and parallelization for the NIMROD project.
- Computer time provided by NERSC, LLNL, and Univ. Texas, Austin.

References

1. D. Schnack, et al., An Overview of the NIMROD Code, Proc. International Workshop on Nonlinear and Extended MHD, Madison WI, April 30-May 2, 1997.

3. A. H. Glasser, et al., "Numerical Analysis of the NIMROD Formulation", Proc. International Workshop on Nonlinear and Extended MHD, Madison WI, April 30-May 2, 1997.

4. A. E. Koniges and K.R. Lind, "Parallelizing Code for Real Applications on the T3D", Computers in Physics 9, 39 (1995).

Author Biography

Alice Koniges is Leader of Multiprogrammatic and Institutional Computing Research at the Lawrence Livermore National Laboratory in Livermore, Ca.

Lawrence Livermore National Laboratory
For information about this page, please contact Alice Koniges, koniges@llnl.gov.

Disclaimers
Last modified May 22, 1997.
UCRL-MI-127382