ELSEVIER

Advances in Engineering Software Vol. 29, No. 3–6, pp. 289–296, 1998

0965-9978/98/$19.00+0.00

Scalable, finite element analysis of electromagnetic scattering and radiation. Error estimation and H-adaptivity

Tom Cwik*, John Lou & Daniel S. Katz

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

In this paper a method for simulating electromagnetic fields scattered from complex objects is reviewed; namely, an unstructured finite element code that does not use traditional mesh partitioning algorithms. The complete software package is implemented on the Cray T3D massively parallel processor using both Cray Adaptive FORTRAN (CRAFT) compiler constructs to simplify portions of the code that operate on the irregular data, and optimized message passing constructs on portions of the code that operate on regular data and require optimum machine performance. The above finite element solution package is then integrated into an error estimation and adaptive mesh refinement algorithm. An error for the fields over the mesh is estimated and used to drive an adaptive mesh refinement algorithm that refines the mesh where errors are above a given tolerance. This refined mesh is then used once again for a finite element solution of the fields. After estimating the error a second time, the mesh is again refined as needed. This process continues until the error is reduced to an allowable level. © 1998 Published by Elsevier Science Ltd. All rights reserved.

Key words: Electromagnetic fields, finite element analysis, simulation, adaptive mesh refinement.

1 PARALLEL FINITE ELEMENT ANALYSIS FOR ELECTROMAGNETICS

Large scale parallel computation can be an enabling resource in many areas of engineering and science. The available memory capacity and computational speed on large distributed memory machines can allow the simulation of complicated engineering components if the simulation algorithm attains an appreciable fraction of the machine peak performance, and if undue cost in porting the code or in developing the code for the parallel machine is not incurred. The issue of code parallelization is especially significant when considering unstructured mesh simulations. The unstructured mesh models considered in this paper result from a finite element simulation of electromagnetic fields scattered from geometrically complex objects (either penetrable or impenetrable.) The finite element model is used to capture the complex materials involved in the simulation, and to maintain fidelity of the structure's geometry. The unstructured mesh must be distributed among the processors, as must the resultant sparse system of linear equations. Since a distributed memory architecture does not allow direct access to the irregularly distributed unstructured mesh and sparse matrix data, partitioning algorithms which are not needed in the sequential software have traditionally been used to efficiently spread the data among the processors. In this paper, an alternate method for simulating electromagnetic fields scattered from complex objects is reviewed: namely, an unstructured finite element code that does not use traditional mesh partitioning algorithms. The complete software package is implemented on the Cray T3D massively parallel processor using both Cray Adaptive FORTRAN (CRAFT) compiler constructs to simplify portions of the code that operate on the irregular data, and optimized message passing constructs on portions of the code that operate on regular data and require optimum machine performance.

The finite element modeling software begins with mesh data constructed on a workstation using a commercially available CAD meshing package. Because the electromagnetic scattering simulation is an open region problem (scattered fields exist in all space to infinity), the mesh must be truncated at a surface that maintains accuracy in the modeled fields and limits the volume of free space that is

*Author to whom all correspondence should be addressed.

meshed. Local, absorbing boundary conditions can be used to truncate the mesh, but these may be problematic because they become more accurate as the truncating surface is removed from the scatterer, requiring greater computational expense, and they may be problem dependent. The approach outlined in this paper solves either the two- or three-dimensional vector Helmholtz wave equations using a coupled finite element–integral equation method. A specific integral equation (boundary element) formulation that efficiently and accurately truncates the computational domain is used. A partitioned system of equations results from the combination of discretizing the volume in and around the scatterer using the finite element method and discretizing the surface using the integral equation method. The system of equations is solved using a two-step solution, combining a sparse iterative (or direct) solver and a dense factorization method. The matrix equation assembly, the solution and the calculation of observable quantities are all computed in parallel, utilizing varying number of processors for each stage of the calculation.

This finite element solution package is then integrated into an error estimation and adaptive mesh refinement algorithm. An error for the fields over the mesh is estimated and used to drive an adaptive mesh refinement algorithm that refines the mesh where errors are above a given tolerance. This refined mesh is then used once again for a finite element solution of the fields. After estimating the error a second time, the mesh is again refined as needed. This process continues until the error is reduced to an allowable level.

2 THE COUPLED FINITE ELEMENT–INTEGRAL EQUATION MODEL

To practically compute a solution to exterior electromagnetic scattering problems, the domain must be truncated at some finite surface where the Sommerfeld radiation condition is enforced, either approximately or exactly. Approximate methods attempt to truncate the mesh using only local field information at each grid point, whereas exact methods are global, needing information from the entire mesh boundary.[1] The global method used here couples a finite element solution interior to the bounding surface with an efficient integral equation solution that exactly enforces the Sommerfeld radiation condition. The problem domain is divided into interior and exterior regions, separated at the mesh boundary. The unknown sources in the integral equation are directly related to the tangential fields on the mesh boundary, and the radiation condition is implicitly enforced exactly through the use of the free-space Green's function. Fields in the two regions are coupled by enforcing boundary conditions on tangential field components at the mesh boundary, thereby producing a unique and exact solution to Maxwell's equations in both regions.

2.1 Two-dimensional formulation—a conformal surface

The two-dimensional scattering geometry is shown in Fig. 1.

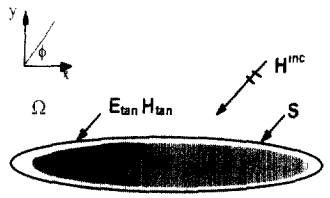

Fig. 1. Geometry of the two-dimensional scatterer for TE_z polarization.

The truncating surface S, taken to lie on or close to the boundary of the scatterer, divides the geometry into interior region V and exterior region Ω. Volume V can be inhomogeneous, containing lossy material and perfect conductors which can lie partly on the surface.[2] For the TE_Z polarization and for wavenumber k, the incident field is

$$\mathbf{H}^{\text{inc}}(\rho) = \hat{\mathbf{z}} e^{jk(x\cos(\phi) + y\sin(\phi))} e^{jwt} \tag{1}$$

which is due a source in the exterior region.

To model the inhomogeneous region V, the wave equation for $\mathbf{H} = \hat{\mathbf{z}} H_z$

$$\nabla \times \frac{1}{\varepsilon_r(\rho)} \nabla \times \mathbf{H}(\rho) - \mu_r(\rho) k^2 \mathbf{H}(\rho) = 0 \quad \rho \in \mathbf{V} \tag{2}$$

is used; ρ is the two-dimensional position vector. Both ϵ and μ are functions of position and may be complex, i.e. $\varepsilon = \varepsilon' - \mathbf{j}\varepsilon''$ and $\mu = \mu' - \mathbf{j}\mu''$, to account for dielectric and magnetic loss. To model the homogenous exterior region Ω, the surface magnetic field integral equation (eqn (3))

$$\begin{aligned} \mathbf{H}(\rho) = 2\mathbf{H}^{\text{inc}}(\rho) + \frac{\mathbf{jk}}{2} \oint_{\partial u} \mathbf{H}_{\tan}(\rho') H_1^{(2)}(\mathbf{k}|\rho - \rho'|) \\ \times \cos(\hat{\mathbf{n}}', \rho - \rho')\, \mathbf{dl}' + \frac{\omega\varepsilon}{2} \oint_{\partial u} \mathbf{E}_{\tan}(\rho') H_0^{(2)} \\ \times (\mathbf{k}|\rho - \rho'|)\, \mathrm{dl}' \quad \rho, \rho' \in \mathbf{S} \end{aligned} \tag{3}$$

is used; $\mathbf{H}_{\tan} = \hat{\mathbf{n}} \times \mathbf{H}^+$ and $\mathbf{E}_{\tan} = \hat{\mathbf{n}} \times \mathbf{E}^+$ are the tangential field components just exterior to the surface S; $H_0^{(2)}$ and $H_1^{(2)}$ are the outgoing first- and second-order Hankel functions; $(\hat{\mathbf{n}}', \rho - \rho')$ denotes the angle between source and observation point, and principal value integrations are implied.

Boundary conditions must be enforced at all material boundaries, at infinity, and at the boundary surface of the two solutions (S). At material boundaries, which only reside within V, boundary conditions are enforced in the finite element solution by a proper choice of elements, or by explicitly zeroing tangential electric field components to model perfectly conducting surfaces. In the integral equation (eqn (3)), the Sommerfeld radiation condition is naturally enforced by the outgoing Green's function. The only boundary conditions remaining need to be enforced at the surface S. Equating tangential electric field components gives

$$\mathbf{E}_{\tan}(\rho) + \frac{1}{\mathbf{j}\omega\varepsilon} \hat{\mathbf{n}} \times \nabla \times \mathbf{H}(\rho) = 0 \quad \rho \in \mathbf{S} \tag{4}$$

on the surface S. If part of the surface encloses a perfect conductor, $\mathbf{H}(\mathbf{r})$ within the conductor is identically zero, and the usual boundary condition $\mathbf{E}_{\tan} = 0$ applies on the surface. Similarly, equating tangential magnetic field components gives

$$\mathbf{H}_{\tan}(\rho) - \hat{\mathbf{n}} \times \mathbf{H}(\rho) = 0 \quad \rho \in \mathbf{S} \tag{5}$$

on the surface S.

The wave equation (eqn (2)) and integral equation (eqn (3)), as well as the boundary conditions of eqns (4) and (5), are combined into three equations. This system of equations is discretized and solved numerically yielding the magnetic fields in the interior region and tangential magnetic and electric fields on the surface S. Far fields are then easily found from an integration of the surface tangential fields.

For a finite element solution of the wave equation, eqn (2) is dotted against a testing function, integrated over the volume, and integrated again by parts yielding the 'weak form'

$$\int_V \left[\frac{1}{\varepsilon_r}\nabla T_z \cdot \nabla H_z - \mu_r k^2 T_z \cdot H_z\right] \mathbf{dv} + \int_{\partial V} \frac{1}{\varepsilon_r} T \cdot [\hat{\mathbf{n}} \times \nabla \times \mathbf{H}]\, \mathbf{ds} \tag{6}$$

The boundary condition on tangential components of the electric field (eqn (4)) is combined into this equation by noting that in the contour integral in eqn (6), $\hat{\mathbf{n}} \times \nabla \times \mathbf{H} = \mathbf{j}\omega\varepsilon\mathbf{E}_{\tan}$; therefore, substituting eqn (4) into this integral gives

$$\int_V \left[\frac{1}{\varepsilon_r}\nabla T_z \cdot \nabla H_z - \mu_r \mathbf{k}^2 T_z \cdot H_z\right] \mathbf{dv} - \mathbf{j}\omega\varepsilon_0 \int_{\partial V} T_z E_{\tan}\, \mathbf{ds} \tag{7}$$

for the first equation of the system ($\mathbf{E}_{\tan} = \hat{\mathbf{t}} E_{\tan} \hat{t}$ is the surface tangent). The second equation is found by enforcing eqn (5) in a weak sense, i.e. dotting eqn (5) against a testing function $\mathbf{U}$ and integrating

$$\int_{\partial V} \mathbf{U} \cdot [\hat{\mathbf{n}} \times \mathbf{H}(\rho) - \mathbf{H}_{\tan}]\, \mathbf{ds} = 0 \tag{8}$$

The third and final equation of the system is found from a standard moment solution to eqn (3), dotting the equation against a testing function $\mathbf{W}$ and integrating along the contour S. This equation will not be written explicitly, the reader is referred to the method of moments literature (for example, Ref. [3]) for further details.

2.2 Three-dimensional formulation — a surface of revolution

In the three-dimensional formulation, the bounding surface chosen is the minimal surface of revolution that fits around the scatterer (Fig. 2). The integral equation is discretized using sub-domain basis functions along the surface of revolution generator, and Fourier harmonics azimuthally,

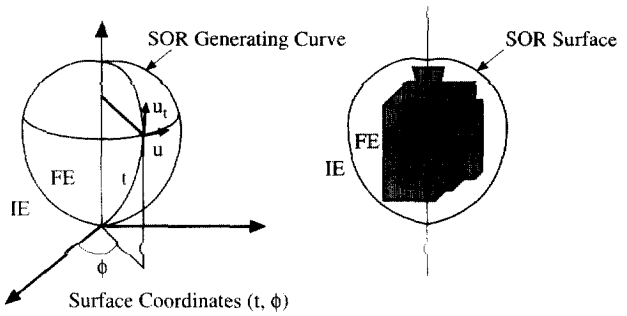

Fig. 2. Geometry of the scattering problem showing interior and exterior regions of model.

to greatly limit the storage necessary in the integral equation component of the model.[4,5]

As in the two-dimensional case, a finite element discretization is used in the interior region. The analogous weak form of the wave equation in eqn (6) for the three-dimensional case is

$$\frac{\eta_0}{jk_0}\int_V\int\int \frac{1}{\varepsilon_r}(\nabla \times \bar{H}) \cdot (\nabla \times \bar{W}^*) - k^2 \mu_r \bar{H} \cdot \bar{W}^*\, \mathrm{dv} - \int_{\partial V}\int \bar{E} \times \hat{n} \cdot \bar{W}^*\, \mathrm{ds} = 0 \tag{9}$$

$\bar{H}$ is the magnetic field (the $\bar{H}$ equation is used in this paper; a dual $\bar{E}$ equation can also be written), $\bar{W}$ is a testing function, the asterisk denotes conjugation, and $\bar{E} \times \hat{n}$ is the tangential component of $\bar{E}$ on the surface of revolution S (∂V). Eqn (9) represents the fields internal to and on the surface of S. These fields will be modeled using a set of properly chosen finite element basis functions. In eqn (9), ε_r and μ_r are the relative permittivity and permeability, respectively, and k_0 and η_0 are free-space wavenumber and impedance, respectively.

A set of tetrahedral, vector edge elements (Whitney elements)[6] are used to discretize eqn (9),

$$\bar{W}_{mn}(r) = \lambda_m(r)\nabla\lambda_n(r) - \lambda_n(r)\nabla\lambda_m(r) \tag{10}$$

in which $\lambda(r)$ are the tetrahedral shape functions and indices (m,n) refer to the two nodal points of each edge. These elements are used for both expansion and testing (Galerkin's method) in the finite element domain.

In the formulation of the integral equation, fictitious electric ($\bar{J} = \hat{n} \times \bar{H}$) and magnetic ($\bar{M} = -\hat{n} \times \bar{E}$) surface currents, equivalent to the tangential magnetic and electric fields just on the exterior of the boundary surface, are defined on the boundary. These currents produce the scattered fields in the exterior region. A linear combination of the electric field integral equation and the magnetic field integral equation is used in this formulation, and it can be succinctly expressed as

$$Z_M[\bar{M}/\eta_0] + Z_J[\bar{J}] = V_i \tag{11}$$

where Z_M and Z_J are the integro-differential operators[5], and V_i represents the incident field.

The integral equation on the surface of revolution is discretized by a set of basis functions with piecewise linear variation along the surface of revolution generator, and with an azimuthal Fourier modal variation. Applying Galerkin's method, both expansion and testing functions are given as

$$\begin{bmatrix} \bar{U}^t \\ \bar{U}^\phi \end{bmatrix} = \begin{bmatrix} \hat{t} \\ \hat{\phi} \end{bmatrix} \frac{T_k(t)}{\rho(t)} e^{jn\phi} \tag{12}$$

in which $T_k(t)$ is a triangle function spanning the kth annulus on the surface of the revolution surface. The variables t and ϕ refer to the local surface of revolution coordinates, and ρ is the distance from the z axis to a point on the surface of revolution. Each annulus spans two segments along the generator, each referred to as a strip. Adjacent triangles overlap on one segment.

At the artificial surface of revolution separating the interior and exterior regions, boundary conditions on the continuity of tangential field components must be enforced. Three equations are written for the three unknown field quantities of interest, the magnetic field $\bar{H}$ internal to the volume V and the electric and magnetic surface currents, $\bar{J}$ and $\bar{M}$, on the boundary. Continuity of the magnetic field across the boundary is enforced in a weak sense

$$\int\int_{\partial V} (\hat{n} \times \bar{H} - \bar{J}) \cdot (\hat{n} \times \bar{U}^*) \, ds = 0 \tag{13}$$

where $\bar{U}$ is a testing function. Continuity of the electric field across the boundary is made implicit in the finite element equation in the surface integral term $\hat{n} \times \bar{E}$ by replacing this term with M.

The surface integral in eqn (9) and the first component of the integral in eqn (13) are termed the coupling integrals, since with a convenient choice of the unknown in the first and of the testing function in the second, they are made to couple interior and exterior field representations. To evaluate these terms, the finite element basis function $\bar{W}$ is evaluated approximately on the portion of surface of revolution projected from the triangular facet of the tetrahedron onto a strip. Such projections are curved triangles, curved quadrilaterals, or curved pentagons. The evaluations of the integrals are performed numerically. These coupling integrals, as well as the discretization of the second surface integral in eqn (13), complete the discretization of the problem.

2.3 Numerical solution of the linear system

After discretization, a coupled linear set of equations results. In the two-dimensional formulation, first-order simplex elements are used for discretization of the finite element equation (eqn (1)), and piecewise triangular basis functions are used in the integral equation. Having introduced the basis and testing functions for the volume as well as the surface unknowns, substitution into the complete set of equations yields

$$\begin{vmatrix} \mathbf{K} & \mathbf{C} & \mathbf{0} \\ \mathbf{C}^\dagger & \mathbf{0} & \mathbf{Z}_0 \\ \mathbf{0} & \mathbf{Z}_\mathbf{M} & \mathbf{Z}_\mathbf{J} \end{vmatrix} \begin{vmatrix} \mathbf{H} \\ \mathbf{M} \\ \mathbf{J} \end{vmatrix} = \begin{vmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{V}_\mathbf{i} \end{vmatrix} \tag{14}$$

where

$$\begin{aligned} \mathbf{K} &= \langle K_p[\bar{W}_p] \cdot \bar{W}_q \rangle \\ \mathbf{C} &= -\eta_0 \langle \bar{U}_m \cdot \bar{W}_q \rangle \\ \mathbf{Z}_0 &= \eta_0 \langle \bar{U}_m \cdot [\bar{n} \times \bar{U}_n] \rangle \\ \mathbf{Z}_\mathbf{M} &= \langle Z_{Mm}[\bar{U}_m] \cdot \bar{U}_n \rangle \\ \mathbf{Z}_\mathbf{J} &= \langle Z_{Jm}[\bar{U}_m] \cdot \bar{U}_n \rangle \end{aligned} \tag{15}$$

The superscript $\dagger$ indicates the adjoint of a matrix. Note that both $\mathbf{K}$ and $\mathbf{C}$ are sparse, $\mathbf{Z}_0$ is tri-diagonal, and $\mathbf{Z}_\mathbf{M}$ and $\mathbf{Z}_\mathbf{J}$ are banded. In particular, the system is complex, non-symmetric and non-Hermitian. This set of equations has the same form in both the two-dimensional and three-dimensional formulations, though the density of the $\mathbf{Z}_\mathbf{M}$ and $\mathbf{Z}_\mathbf{J}$ blocks differ in the two cases.

The parallel solution to this matrix equation system is completed in two steps. Initially $\mathbf{H}$ in the first equation in eqn (14) is written as $\mathbf{H} = -\mathbf{K}^{-1}\mathbf{C}\mathbf{M}$ and substituted into the second equation resulting in

$$\begin{vmatrix} \mathbf{Z}_\mathbf{K} & \mathbf{Z}_0 \\ \mathbf{Z}_\mathbf{M} & \mathbf{Z}_\mathbf{J} \end{vmatrix} \begin{vmatrix} \mathbf{M} \\ \mathbf{J} \end{vmatrix} = \begin{vmatrix} \mathbf{0} \\ \mathbf{V}_\mathbf{i} \end{vmatrix} \tag{16}$$

where $\mathbf{Z}_\mathbf{K} = -\mathbf{C}^\dagger\mathbf{K}^{-1}\mathbf{C}$. This relatively small system is then solved directly for $\mathbf{M}$ and $\mathbf{J}$. By solving the system in two steps, the interior solution is decoupled from the incident field V_i, allowing for efficient solutions when many excitation fields are present as in monostatic radar cross-section simulations.

The relative numbers of unknowns in $\mathbf{H}$ and $\mathbf{M}$ (or $\mathbf{J}$) makes the calculation of $\mathbf{K}^{-1}\mathbf{C}$ the major computational expense. This operation is the solution of a system of equations, $\mathbf{KX} = \mathbf{C}$, where $\mathbf{C}$ is a rectangular matrix with a potentially large number of columns in the case of electrically large scatterers. The solution is accomplished by using a symmetric variant of the quasi-minimum residual iterative algorithm. The resulting overall matrix (eqn (16)) is treated as being dense, and the solution of this second system is accomplished via a direct dense LU decomposition, since its size is relatively small.

3 ADAPTIVE MESH REFINEMENT

An adaptive mesh refinement package that constructs a set of adaptively refined meshes based on an error estimate is added to the above formulation. These error estimates follow from the Cauchy convergence properties of finite element mathematical analysis and only require a small percentage of the time needed to solve the complete system

over the previous step's mesh. Based on the error estimate generated, the mesh is adaptive refined — independent of the user — where needed and the fields are calculated over the new mesh by solution of the sparse matrix equation. This step is repeated until the error in the fields over the mesh is uniform and reduced below a user prescribed level.

3.1 Error estimation

Error estimates have been well developed in finite element structural and fluid analysis[7]. Estimates of error in the magnetic or electric field have been less developed in electromagnetic analysis. This work follows from the error estimates in Meyer and Davidson[8] where p-enrichment was implemented. The error estimate is a local estimate of the error in the fields over each element. The estimate is computed element by element rather than over the entire mesh for computational efficiency. By defining u to be the exact solution to the wave equation (either E or H), and $\hat{u}$ to be the computed solution over a given mesh, the estimate can be derived. The error in the field for a given solution is then $e = u - \hat{u}$. The derivation of the estimate e follows that of standard finite element analysis.[7] The result is an equation that estimates the error that would exist over the next refined mesh, based on the global calculated error over the coarse mesh. The equation for this error is

$$B_k(e,t) = L_k(t) + \hat{B}(\hat{u},t) + \tfrac{1}{2}\int_{\partial V_k/\partial V} \hat{n}_k\cdot(a_k\nabla\hat{u}_k - a_j\nabla\hat{u}_j)t \,\mathrm{d}s \tag{17}$$

where

$$B_k(e,t) = \int_k (a\nabla e\cdot\nabla t - k_0^2 bet)\,\mathrm{d}V \tag{18}$$

and the surface integral is along edges of adjoining elements k and j, and the term $L_k(t)$ is a surface integral at the computational boundary where a radiation boundary condition is imposed. $\hat{B}$ is the evaluation of eqn (18) using the known solution $\hat{u}$, resulting in a vector of known entries. Eqn (17) is therefore an equation for e over the kth element, given known right-hand-side terms found from a coarse mesh solution. The testing functions t are the fine mesh elements, thereby producing a solution for an estimate of the error in the coarse grid projected onto a refined mesh. If this error is above a tolerance value, the element will be refined; if the error is below the given tolerance the element need not be refined. This refinement decision based on the error estimate is the adaptive h-refinement process.

3.2 Adaptive mesh refinement

The h-refinement procedure consists of two main components, an adaptive mesh adjustment step and an adaptive mesh refinement step. The input to our adaptive mesh refinement procedure is a coarse mesh, an array containing local error estimates for the coarse mesh, and a prescribed error tolerance. The adaptive mesh adjustment step implements a scheme that, given a mesh with a subset of elements indicated to be refined, generates information necessary for producing a consistently refined global mesh. The adaptive mesh refinement step performs the actual refinement operations on the input mesh based on the refinement information stored in each coarse element.

The adaptive mesh refinement algorithm is implemented for triangular meshes and will be extended to tetrahedral meshes. The software is developed in FORTRAN 90 using many of its advanced features such as modules, user-defined structures, pointers, dynamic memory management and array operations. The software is therefore highly modular and robust. It has interfaces to both FORTRAN 77 and FORTRAN 90.

4 PARALLEL IMPLEMENTATION AND RESULTS

A parallel implementation of the above finite element formulation consists of an initial decomposition of the sparse matrix data, solution of the sparse matrix equation for the unknown magnetic field over the mesh and the surface tangential electric and magnetic fields, and finally the calculation of observable quantities such as the radar cross-section. This implementation is then connected to the error estimation algorithm and adaptive mesh refinement algorithms. A description of the parallel three-dimensional finite element implementation will be given with results, followed by results from the two-dimensional adaptive mesh refinement algorithm.

The matrix decomposition code, termed P_SLICE, consists of a number of subroutines. Initially, the potentially large mesh files are read (READ). Then the connectivity structure of the sparse matrix is generated and reordered (CONNECT), followed the generation of the complex-valued entries of **K** (FEM), building the connectivity structure and filling the **C** matrix (COUPLING). Finally, the individual files containing the row slabs of **K** and the row slabs of **C** must be written on disk (WRITE). For each processor that will be used in the matrix equation solver, one file containing the appropriate parts of both the **K** and **C** matrices is written.

Cray Research Adaptive FORTRAN (CRAFT) is used for the matrix decomposition stage of the simulation. All large arrays are declared using *CDIR$* directives to be shared in either a block manner or a cyclic manner for the leading dimension, with the non-leading dimension distributed degenerately. Using a block distribution of a matrix of size 256 on four processors leads to the first 64 elements residing on processor 0, the next 64 elements on processor 1, etc. A cyclic distribution would lead to processor 0 having elements (1, 5, 9, ···), processor 1 having elements (2, 6, 10, ···), etc. A two-dimensional array with a degenerate distribution of the second dimension leads to all elements of the array having a given index in the first dimension being on the same processor, regardless of the index in the second

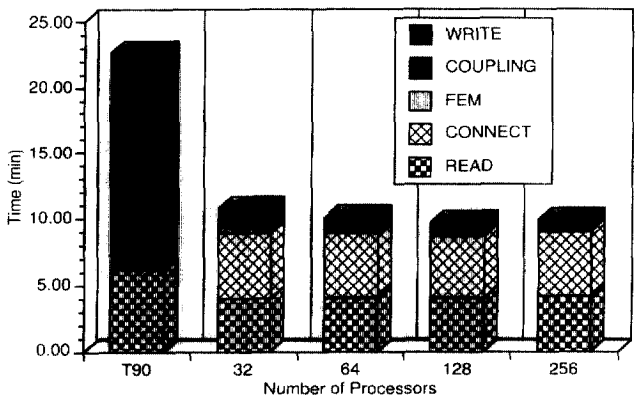

Fig. 3. Computation time and scaling for a relatively large simulation (dielectric cylinder with 579 993 edges, radius = 1 cm, height = 10 cm, permittivity = 4.0 at 2.5 GHz). First column shows the time for single processor T90. Times on T90 for CONNECT and FEM have been combined.

dimension. For example, a two-dimensional array of size (256,10) distributed degenerately over the second dimension will have elements $(i,1)$, $(i,2),\cdots,(i,10))$ all located on the same processor. Which processor this will be is dependent on the value of i, and the method of distribution over the first dimension.

Routines which could be easily parallelized by CRAFT directives were FEM and part of COUPLING. The directive *CDIR$ DO SHARED* was added to the parallelizable loops to automatically distribute the work over all the processors. Other routines that could be executed in parallel with a combination of CRAFT and message passing included the READ and WRITE routines. The remaining routines (CONNECT, and a second part of COUPLING) are basically sequential routines, where only one processor is doing the majority of work, while using data spread across many (usually all) processors.

Two files are read in the READ routine, one containing finite element data, and the other containing integral equation data. The finite element file is at least an order of magnitude larger than the integral equation file, and is read by four processors. By using these four processors, the time of the READ routine is reduced roughly by a factor of 3 as compared to reading the file with one processor.

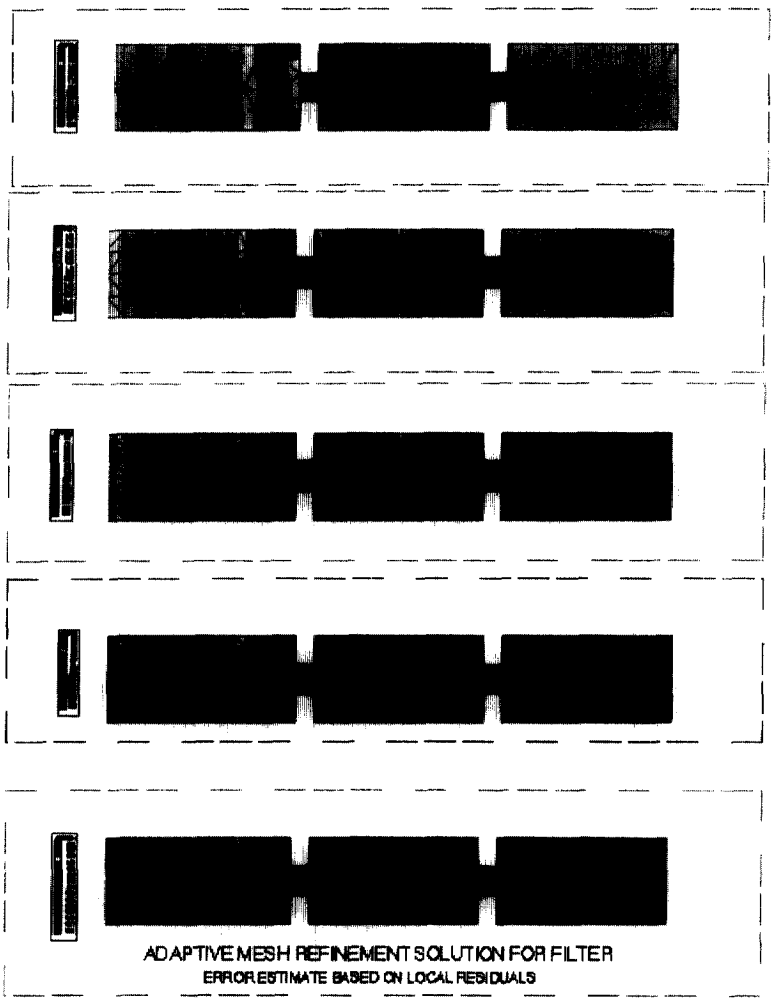

Fig. 4. Two-dimensional waveguide filter structure. Dimensions are 0.45 cm height main section; 0.05 cm height reduced height section, centered in guide; 0.9 cm length of each main section; 0.1 cm length of reduced height section; Frequency = 15 GHz; magnetic field amplitude plotted.

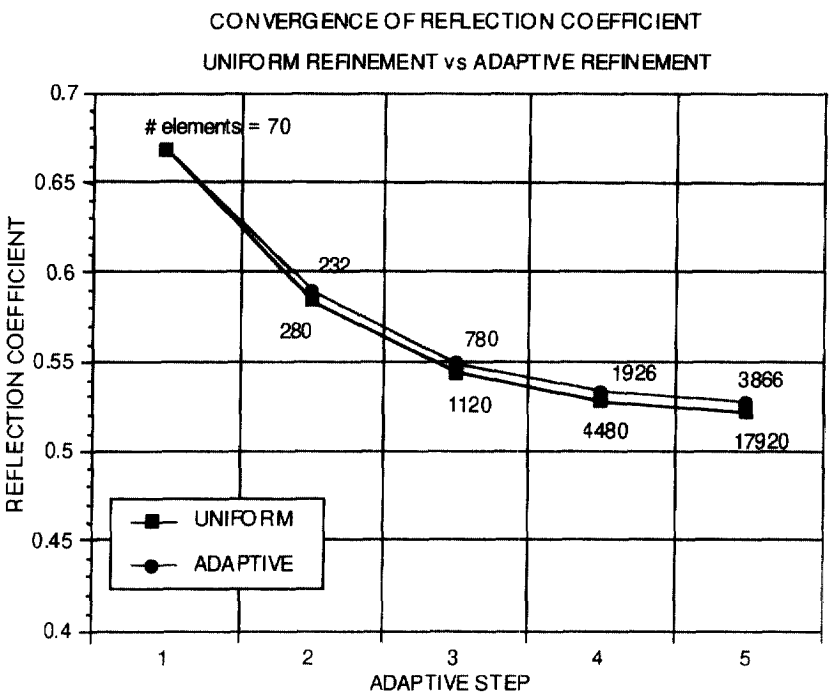

Fig. 5. Convergence of reflection coefficient for the waveguide filter of Fig. 5. Frequency = 15 GHz. The number of triangular finite elements is listed for each refinement step.

Further reduction in this time may be possible; however, this factor of 3 is currently sufficient. In the WRITE algorithm, data is assembled on each processing element and written to disk. On the T3D, it is faster to assemble a local array and write out that data than to write out a distributed array directly, since as the number of processors increases, more writes of smaller amounts of data are being performed, and disk and network contention develops. Scaling beyond this point quickly leads to diminishing returns from each processor.

Fig. 3 shows the performance of P_SLICE over varying numbers of processors for a simulation of scattering from a dielectric cylinder.[9] The number of edges is the number of finite element unknowns in the problem. It may be observed that for the routines that have been parallelized, doubling the number of processors reduces the amount of time by a factor of approximately two. For routines that are sequential, where only one processor is doing the work using the other processors' data, the time goes up very slightly as the number of processors for the overall code are increased. This is due strictly to communication latency. As the number of processors increases, the percentage of array elements which are not local increases, and the time to load or store these elements is longer than the time to load or store local elements. The I/O time should have roughly the same behavior, but for practical tests the I/O time is more dependent on the I/O load of the other T3D processors and the load on the front-end YMP that is between the T3D and the disks than the number of T3D processors being used in P_SLICE. It is clear that the routines that benefit most from the parallel implementation on the T3D are COUPLING and WRITE. Results for the solver portion of the code can be found in Ref. [9].

An example of results from the adaptive mesh refinement algorithm is shown in Fig. 4. The problem considered is a two-dimensional waveguide filter structure. An initial coarse mesh of the geometry is generated. An H-polarized (H perpendicular to the paper) lowest order mode is incident from the left. Absorbing boundary conditions are used at the left and right ports. Shown colored on the plots is the magnetic field magnitude for solutions over the initial coarse mesh and through four refinements. The refinement is driven by the error estimate calculated over the mesh at each step of the calculation. It is seen that the fields converged to constant values as the refinement process continues.

Shown in Fig. 5 is the convergence of the reflection coefficient for the waveguide filter for each adaptive step. Plots are shown for uniform refinement (dividing each triangular element uniformly independent of an error estimate) and the adaptive refinement. It is seen that nearly exact convergence of the reflection coefficient can be found using nearly four times fewer elements using the adaptive mesh refinement process.

ACKNOWLEDGEMENTS

The three-dimensional developments presented in this paper were performed in collaboration with Cinzia Zuffada and Vahraz Jamnejad at Jet Propulsion Laboratory (JPL). The JPL/Caltech Supercomputer used in this investigation

was provided by the NASA Offices of Mission to Planet Earth, Aeronautics, and Space Science. The Cray T90 used in this investigation was provided by the Information Services Department of Cray Research. The research described in this paper was performed at the JPL, California Institute of Technology, under contract to the National Aeronautics and Space Administration (NASA).

REFERENCES

1. Jin, J.-J., *The finite Element Method in Electromagnetics*. Wiley, New York, 1993.
2. Cwik, T. Coupling finite element and integral equation solutions using decoupled boundary meshes. *IEEE Trans. Antennas Propagat.*, 1992, **AP-40**(12), 1496–1504.
3. Poggio, A. and Miller, E., Integral equation solutions of three-dimensional scattering problems. In *Computer Techniques for Electromagnetics*, Chapter 4. Pergamon Press, London, 1973.
4. Boyse, W. and Seidl, A. A hybrid finite element method for near bodies of revolution. *IEEE Trans. Mag*, 1991, **MAG-27**, 3833–3836.
5. Cwik, T., Zuffada, C. and Jamnejad, V. Efficient coupling of finite element and integral equation representations to model scattering from three-dimensional objects. *IEEE Trans. Antennas Propag.*, 1996, **AP40**.
6. Lee, J. and Mittra, R. A note on the application of edge-elements for modeling 3-dimensional inhomogeneously-filled cavities. *IEEE Transactions on Microwave Theory and Techniques*, 1992, **40**(9), 1767–1773.
7. Ainsworth, M. and Oden, T. A procedure for a posteriori error estimation for h–p finite element methods. *Comp. Meths. in Applied Mechs. and Eng.*, 1992, **101**, 73–96.
8. Meyer, F. and Davidson, D. Adaptive-mesh refinement of finite element solutions for two-dimensional electromagnetic problems. *IEEE Antennas and Prop. Magazine*, 1996, **37**(5), 77–83.
9. Cwik, T., Katz, D., Zuffada, C. and Jamnejad, V., The application of scalable distributed memory computers to the finite element modeling of electromagnetic scattering. *Intl. Jrnl. Numer. Methods. Eng.*, 1998, **41**, 759–776.