A Parallel Finite-Element Tearing and Interconnecting Algorithm for Solution of the Vector Wave Equation with PML Absorbing Medium

C. T. Wolfe, *Member, IEEE,* U. Navsariwala, *Member, IEEE,* and Stephen D. Gedney, *Senior Member, IEEE*

***Abstract*—A domain decomposition method based on the finite-element tearing and interconnecting (FETI) algorithm is presented for the solution of the large sparse matrices associated with the finite-element method (FEM) solution of the vector wave equation. The FETI algorithm is based on the method of Lagrange multipliers and leads to a reduced-order system, which is solved using the biconjugate gradient method (BiCGM). It is shown that this method is highly scalable and is more efficient on parallel platforms when solving large matrices than traditional iterative methods such as a preconditioned conjugate gradient algorithm. This is especially true when a perfectly matched layer (PML) absorbing medium is used to terminate the problem domain.**

***Index Terms*—Finite-element methods, PML's, vector wave equation.**

I. Introduction

THE finite-element method (FEM) is an effective means for analyzing a plethora of electromagnetic problems. The FEM's principal attribute is that it efficiently models highly irregular geometries as well as penetrable and inhomogeneous material media. The linear system of equations that results from a FEM discretization is highly sparse and can be solved using efficient solution techniques for sparse matrices based on either direct methods [1] or iterative methods [2], [3]. Direct methods have the advantage that multiple right-hand sides can be treated efficiently. However, storing the factorized matrix is memory intensive for large matrices. Iterative methods, such as the conjugate gradient (CG) method, are much less memory intensive. However, an iterative solution must typically be performed for each right-hand side.

Over the past decade, high-performance computing has been achieved via multiprocessing. In a massively parallel environment, traditional sequential algorithms will not necessarily scale and can lead to a very poor utilization of the multiprocessor's architecture. As a result, specialized algorithms that directly exploit the parallel architecture must be developed. For the solution of sparse matrices, parallel algorithms based on domain decomposition methods (DDM's) have been the most successful [2], [3]. DDM's essentially partition the global mesh discretizing the problem domain into several nonoverlapping contiguous subdomains. Subsequently, a sparse matrix derived from a variational formulation can be introduced for each subdomain. The subregions are then coupled together through some boundary constraint. The advantage of this approach is that if the mesh is equally partitioned, then this approach can lead to a highly scalable algorithm. Therefore, the efficiency of the parallel algorithm can be dependent upon both the algorithm devised for the matrix solution as well as the partitioning algorithm. There has been extensive effort in the areas of effective and efficient mesh partitioning algorithms, including the recursive inertia partitioning algorithm [4], spectral bisection methods [5], [6], the METIS algorithm [7], and the greedy algorithm [8], [9].

The focus of this paper is on the development of the parallel algorithm for the matrix solution. An early approach to this problem was a divide-and-conquer technique developed by Patterson *et al.* [10], [11]. This approach consisted of partitioning the global matrix using an automatic partitioning scheme. Subsequently, a global iterative solver based on the biconjugate gradient (BiCG) method was used to solve the distributed sparse matrix. Each matrix vector multiply of the BiCG algorithm was done in parallel. Interprocessor communication is required to concatenate the resultant vector, as well as to perform the global dot-product operations. One of the difficulties with the divide-and-conquer scheme is that the number of iterations required for convergence is dependent upon the condition number of the global matrix.

Alternatively, a parallel direct solution method for two-dimensional (2-D) FEM analysis was introduced by Lee *et al.* [12], [13]. This technique coupled the subdomain solutions by enforcing tangential field continuity between adjacent subdomains leading to a global matrix representing only the tangential fields on the shared boundaries. The global matrix is much smaller than the original FEM matrix and can be solved using a direct method.

In [14], Deprés introduced a hybrid iterative DDM for the 2-D Helmholtz problem [14]. To this end, an iterative method was proposed for which each iteration consists of solving the fields interior to each subdomain and then constraining the field continuity at the interface of each subdomain by enforcing a Robin-type transmission condition on the boundary fields (this transmission condition essentially enforces the continuity of both the tangential electric and magnetic field intensities across the shared boundaries). Depres also introduced a relaxation scheme in [15] that greatly accelerated the iterative process. Later Stupfel [16] extended this method by prescribing a new ABC [17] at the exterior boundary and using an "onion-like"

Manuscript received June 2, 1997; revised October 7, 1999. This work was supported by the National Science Foundation under Grant ECS-9624628.

C. T. Wolfe is with Lexmark International, Inc., Lexington, KY 40550 USA.

U. Navsariwala and S. D. Gedney are with the Department of Electrical Engineering, University of Kentucky, Lexington, KY 40506-0046 USA.

Publisher Item Identifier S 0018-926X(00)01272-2.

partition of the computational domain improving the efficiency of the transmission condition and overall performance of this DDM.

The focus of this paper will be on the application of a hybrid iterative solution based on the method of Lagrange multipliers. This method is modeled after the finite-element tearing and interconnecting (FETI) method originally developed by Farhat and Roux [18]. Specifically, the FEM discretization of the weak form equation for each subdomain will be posed. The solutions of each subdomain will be constrained through the use of Lagrange multipliers by enforcing the continuity of the tangential fields across each boundary interface. A reduced system of equations representing the Lagrange multipliers is then derived and is solved using a preconditioned BiCG algorithm. The advantage of this method is that each subregion can be solved completely independently, leading to a scalable algorithm. Second, the number of iterations required to solve the global problem is dependent on the order of the matrix representing the Lagrange multipliers as opposed to the global matrix.

Another challenge of modeling electromagnetic wave phenomena in unbounded medium via the FEM is the accurate and efficient termination of the discrete volume using an absorbing boundary condition or an absorbing layer. The goal of such terminations is to be nonreflective, while minimizing the overall mesh dimensions and the computational overhead associated with the truncation operator. Recently, Berenger [19] proposed the perfectly matched layer (PML) as an absorbing medium for orthogonal finite-difference time-domain (FDTD) methods. Using the PML to truncate FDTD meshes, reflection errors due to the absorber have been shown [20]–[25] to be -80 dB or less. Berenger's PML is based on a split-field formulation, which is well suited for an orthogonal grid implementation.

A similar PML based on a uniaxial anisotropic medium was introduced by Sacks *et al.* [23] for finite-element frequency-domain methods, and by Gedney [24]–[26] for FDTD methods. The uniaxial PML (UPML) technique is better suited for applications such as the FEM, which utilize unstructured grids since it does not rely on field splitting based on orthogonal projections. Furthermore, it avoids much of the complexity introduced by a stretched coordinate form of Berenger's PML [27]. A number of recent papers have been published on the application of PML and UPML terminations of FEM meshes. Sacks *et al.* [23], provided an example of a simple dipole radiating in an open region. Gong and Volakis [28] used a PML to terminate a microstrip line. Rappaport introduced the stretched coordinate PML formulation for finite-element applications [27]. By adjusting the PML parameters, the layer was found to produce a reflection error of -30 dB or less for a line current radiating in an unbounded region. Likewise, Lyons *et al.* [29], used a PML to terminate an air-filled waveguide and found reflection errors of -30 to -55 dB. Kingsland *et al.* [30] investigated the use of PML for propagation and scattering problems. Finally, Chew and Jin [31] analyzed the PML in discretized space in order to optimize the layer with respect to its parameters.

Even though PML methods have shown promising performance, one difficulty that arises is that the matrix becomes poorly conditioned when PML absorbing layers are present [32]. As a result, iterative solvers suffer a substantial increase in the number of iterations needed for convergence. It will be shown herein that the number of iterations required by the FETI algorithm for large matrices is dramatically reduced, even with the presence of PML absorbing layers as compared to a direct iterative solution.

This paper is organized as follows. The FEM formulation with UPML absorbing layers is presented in Section II. The FETI algorithm is then presented in Section III. In Section IV, the efficiency and scalability of the FETI algorithm for the solution of the FEM matrices with UPML absorbing media is presented.

II. Finite-Element Implementation Using a Uniaxial PML

Assume that a lossy inhomogeneous half-space is interfaced with a uniaxial anisotropic medium in the $z = 0$ plane. In the anisotropic medium, Maxwell's equations are described in the frequency domain as

$$\nabla \times \vec{E} = -j\omega\mu_0\mu_r\bar{\bar{s}}\vec{H},\ \nabla \times \vec{H} = j\omega\varepsilon_0\varepsilon_r\bar{\bar{s}}\vec{E} \tag{1}$$

where $\hat{\varepsilon}_r$ is a complex and frequency dependent and $\bar{\bar{s}}$ is a 3×3 tensor. It was shown in [23]–[25] that an arbitrarily polarized plane wave propagating from a medium with material parameters $\hat{\varepsilon}_r$ and $\mu_r$is impinging on a planar half-space at the $z = 0$ interface described as a uniaxial medium with material parameters $\varepsilon_0\hat{\varepsilon}_r\bar{\bar{s}}$ and $\mu_0\mu_r\bar{\bar{s}}$, the interface will be reflectionless if

$$\bar{\bar{s}} = \begin{bmatrix} s_z & 0 & 0 \\ 0 & s_z & 0 \\ 0 & 0 & s_z^{-1} \end{bmatrix}. \tag{2}$$

This will hold true for arbitrary polarization, angle of incidence, and frequency spectrum. For application to the finite method, the anisotropic medium must be highly attenuative such that any wave entrant in the medium will attenuate rapidly. By choosing $s_z = \kappa_z + (\sigma_z/j\omega\varepsilon_0) = \alpha - j\beta$ [22]–[25], then κ_z amplifies the attenuation of the evanescent portion of the wave, and the imaginary term attenuates the propagating portion.

In the corner regions, where the UPML layers overlap, $\bar{\bar{s}}$ will be represented by the product of tensors. This can easily be derived by matching a UPML to a uniaxial medium [26]. Subsequently, $\bar{\bar{s}}$ is generalized as

$$\bar{\bar{s}} = \begin{bmatrix} \dfrac{s_y s_z}{s_x} & 0 & 0 \\ 0 & \dfrac{s_x s_z}{s_y} & 0 \\ 0 & 0 & \dfrac{s_x s_y}{s_z} \end{bmatrix}. \tag{3}$$

Secondly, due to discretization errors, the PML medium can suffer from numerical reflection error. To circumvent this problem, the material parameters have been scaled using a polynomial scaling [18]–[26].

The FEM is used to compute the fields in the volumetric space. The vector wave equation in the uniaxial medium is derived from (1)

$$\nabla \times \mu_r^{-1}\bar{\bar{s}}^{-1}\nabla \times \vec{E} - \omega^2\mu_0\varepsilon_0\hat{\varepsilon}_r\bar{\bar{s}}\vec{E} = 0 \tag{4}$$

where UPML is assumed throughout the volume. Performing the inner product with a testing function defined over the finite volume Ω and utilizing Green's first identity results in the weak form equation

$$\iiint_{\Omega}\left[\nabla\times\vec{T}\cdot\mu_r^{-1}\bar{\bar{s}}^{-1}\nabla\times\vec{E}-\omega^2\mu_0\varepsilon_0\hat{\varepsilon}_r\vec{T}\cdot\bar{\bar{s}}\vec{E}\right]d\Omega - \oiint_{\partial\Omega}\left[\vec{T}\cdot\hat{n}\times\mu_r^{-1}\bar{\bar{s}}^{-1}\nabla\times\vec{E}\right]dA=0. \tag{5}$$

The finite-element solution is performed by discretizing the volume into element domains and expanding the testing and trial vector functions using vector edge elements (in this paper first-order Whitney elements are employed). Then, the first variation of (5) is evaluated at the stationary point, leading to a sparse linear system of equations.

Inside the working volume, the medium is assumed to be isotropic. Specifically from (3), $\bar{\bar{s}}$ reduces to the identity tensor. Within the PML region, $\bar{\bar{s}}$ is anisotropic and the parameters are assumed to be spatially dependent as an mth-order polynomial along the normal axes. To more accurately represent the spatial variation, Gaussian quadrature numerical integration was used to perform the integrals in (5) [33]. Furthermore, since the PML interfaces are assumed planar the implementation of (5) is quite simple in the FEM routine and no special preprocessing is required. The exterior boundary backing the PML region is assumed to be a PEC wall. On this wall, a Dirichlet boundary condition is enforced.

III. Finite-Element Tearing and Interconnecting Method

The finite-element tearing and interconnecting method (FETI) is a domain decomposition technique based on a hybrid variational principle. For simplicity consider a domain Ω divided into nonoverlapping subdomains Ω_i. Adjacent subdomains will share common boundaries defined by $\Gamma_{i,j}$. The vector electric fields and testing functions within each subdomain are discretized separately into finite elements. Then evaluating the first variation of (5) at a stationary point yields

$$K_i e_i = f_i \tag{6}$$

where

- e_i vector field unknowns in region Ω_i;
- K_i stiffness matrix;
- f_i forcing vector in Ω_i.

The discrete fields in each subregion are constrained to enforce the continuity of the tangential electric fields across the shared boundary. Specifically, let $B_i e_i$ map e_i into the space of tangential fields on the shared boundaries $\Gamma_{i,j}$. Subsequently, using the method of Lagrange multipliers, it is desired to solve the linear system of equations

$$Ke = f \tag{7}$$

with the constraint

$$Be = 0 \tag{8}$$

where K is a block diagonal matrix defined by

$$K=\begin{bmatrix} K_1 & 0 & \cdots & 0 \\ 0 & K_2 & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & K_N \end{bmatrix}. \tag{9}$$

e, f, and B are described as

$$e=\begin{bmatrix} e_1 \\ e_2 \\ \vdots \\ e_N \end{bmatrix} \quad f=\begin{bmatrix} f_1 \\ f_2 \\ \vdots \\ f_N \end{bmatrix} \quad B=[B_1 \quad B_2 \quad \cdots \quad B_N]. \tag{10}$$

B is a matrix representing the continuity of tangential fields such that $B_i e_i + B_j e_j = 0$. It is noted that the matrix B is highly sparse and all nonzero entries are simply ± 1. From the theory of Lagrange multipliers, the solution of this constrained linear system is equivalent to finding the stationary point of the functional

$$f = \tfrac{1}{2}e^T K e - e^T f + \lambda^T B e \tag{11}$$

where λ is the vector of Lagrange multipliers. Taking the first variation of (11) and evaluating it at the stationary point leads to the symmetric linear system of equations

$$\begin{bmatrix} K & B^T \\ B & 0 \end{bmatrix}\begin{bmatrix} e \\ \lambda \end{bmatrix} = \begin{bmatrix} f \\ 0 \end{bmatrix}. \tag{12}$$

From the first row of (12), it follows that

$$e = K^{-1}(f - B^T\lambda) \tag{13}$$

combining this with the second row of (12) leads to the symmetric sparse linear system of equations

$$BK^{-1}B^T\lambda = BK^{-1}f. \tag{14}$$

The parallel algorithm then proceeds as follows.

1) K is factorized using a sparse matrix factorization method. Note that this is done completely in parallel due to the block diagonal characteristics of K defined in (9). As a result, each K_i is factorized independently and concurrently on each processor.
2) A preconditioned BiCG algorithm is used to compute the solution for the vector of Lagrange multipliers λ from the reduced-order matrix equation in (14). It is noted that the matrix vector multiplies are performed completely in parallel since the products of vectors with the matrix blocks K_i^{-1} and B_i can be performed explicitly in parallel.
3) Once the Lagrange multipliers are computed, the electric fields are computed in parallel from (13).

It is noted that the order of the matrix in (14) is greatly reduced as compared to the order of the global matrix. As a result, the iterative scheme is expected to converge rapidly. It is also noted that if there is symmetry or repetitiveness in the geometry resulting in blocks with identical K_i, then such blocks only need to be stored and factored once on one processor. (It is noted that proper load balance should still be monitored in such instances.) Furthermore, since the global problem is of reduced

order it would be possible to solve the problem with a direct solver. However from a scalability issue, an iterative solver was implemented.

The performance of a number of preconditioners for accelerating the BiCG solution of (14) was studied. Interestingly, diagonal preconditioning based on the full matrix $BK^{-1}B^T$ did not significantly reduce the number of iterations. Regardless, this preconditioner is prohibitively expensive to construct, as it requires the rigorous evaluation of $K^{-1}B^T$. Incomplete diagonal preconditioners based only on the diagonal of K also do not accelerate the BiCG solution.

The most successful preconditioner found that significantly accelerates the BiCG convergence without insurmountable computational overhead was BKB^T. While, this preconditioner is more expensive than a simple diagonal preconditioner, it does significantly accelerate the convergence of the BiCG algorithm and reduces the overall CPU time for the solution.

The condition number of the reduced-order system in (14) was found to be sensitive to the aspect ratio of the subdomains. This was also reported in [34]. More specifically, highly elongated domains or domains with highly irregular surfaces tended to lead to poorer conditioned matrices and, hence, slower convergence. It was found that regulating (or smoothing) the domain shape greatly improved this. If one were to decompose the mesh using an automated mesh partitioning scheme such as METIS, such smoothing can be done automatically using simulated annealing or other techniques as reported in [35].

IV. Results

In this section, the efficiency of the FETI solution will be presented for FEM models with UPML. To this end, the results for a microstrip line are presented. This problem was chosen since the computational domain was terminated on five sides by a UPML layer (see Fig. 1). The microstrip line was excited by a 4.0-GHz voltage. The relative permittivity and permeability of the dielectric under the microstrip line were chosen to be 3.2 and 1.0, respectively.

The microstrip line length was set to 0.08629 m and the depth of the PML on the end walls was effectively six-cell-radii thick, on the side walls five-cell-radii thick, and on the top wall five-cell-radii thick. First-order tetrahedral edge elements or Whitney elements were used in this study and there was no effort to make the tetrahedral elements symmetric about any axis, however, the PML interfaces were planar. No special meshing was done in the PML regions. From (4), we arbitrarily set the order of the spatial polynomial to $m = 0$. This reduces the complexity of the expressions for (3) with minimal increase in error. For the study, the values used were $s_x = s_y = (2 - j1.8)$ and $s_z = (1 - j1.8)$. These values were experimentally determined by minimizing the error of the electric field along the microstrip line.

To study the efficiency of the FETI algorithm, the parallel FETI solution outlined in Section III was implemented on a 32-processor subcomplex of an HP SPP2200. Scalability results were found for 8, 16, and 32 processors by keeping the (unknowns/processor) approximately equivalent. Three sets of problems were investigated in this manner. The first set assigns approximately 3000 unknowns per processor, the second approximately 6000, and the third set approximately 9000. In order to obtain the number of unknowns desired, the number of cells was scaled in the x-,y-, and z-directions. The geometry was kept constant, so the average cell size was decreased. The benchmark used to compare against the FETI algorithm was a divide-and-conquer-type preconditioned BiCGM. The preconditioner used was an inexpensive diagonal preconditioner. The FETI method used BKB^T as its preconditioner for the BiCGM as previously discussed. In addition, different mesh partitioning methods were used in conjunction with the FETI method. The ones presented are: 1) heuristic RIP (H-RIP), which recursively slices along the long axis of the geometry but only along a planar boundary leaving "cube-like" subdomains; 2) METIS; and 3) METIS with simulated annealing (METIS-SA), which recursively subdivides domains but uses simulated annealing after each partitioning in order to smooth the interface boundaries. Since the eigen spectrum of the divide-and-conquer method is a function of the overall geometry and not of individual subdomains, it is necessary to show results using only one of the partitioning methods.

Tables I–III summarize the number of degrees of freedoms for the global problems and their iteration counts (in parenthesis) for the various solution methods. For the divide-and-conquer method this is simply the number of field variables. For the FETI method this would be the number of Lagrange multipliers. It can

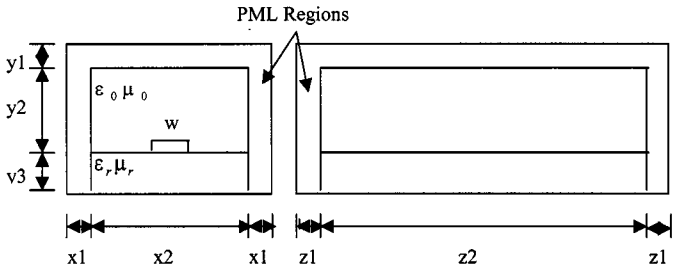

Fig. 1. Dimensions for microstrip line problem. ($x1 = 0.012$ m): PML depth in x-direction; ($x2 = 0.0238$ m): problem depth in x-direction; ($y1 = 0.012$ m): PML depth in y-direction; ($y2 = 0.0085$ m): free-space depth in y-direction; ($y3 = 0.0021$ m): dielectric depth in y-direction; ($z1 = 0.016$ m): PML depth in z-direction; ($z2 = 0.05492$ m): problem depth in z-direction; ($w = 0.00548$ m): width of microstrip. Note: the conductor thickness of the microstrip line was infinitesimal.

TABLE I
COMPARISON OF DEGREES OF FREEDOM (ITERATION COUNTS) FOR GLOBAL PROBLEM WITH 3000 UNKNOWNS PER PROCESSOR

	8 Domains	16 Domains	32 Domains
D&C / METIS	24336 (19474)	46572 (19116)	86943 (29564)
FETI / METIS	2199 (1789)	5510 (2634)	12323 (4955)
FETI / METIS-SA	2107 (1619)	5242 (2106)	11700 (3909)
FETI / H-RIP	1944 (991)	5016 (1375)	10842 (2015)

TABLE II
COMPARISON OF DEGREES OF FREEDOM (ITERATION COUNTS) FOR GLOBAL PROBLEM WITH 6000 UNKNOWNS PER PROCESSOR

	8 Domains	16 Domains	32 Domains
D&C / METIS	46572 (18869)	86943 (29720)	172766 (49309)
FETI / METIS	3355 (1489)	8111 (2287)	19342 (7912)
FETI / METIS-SA	3181 (1273)	7911 (2070)	18431 (5381)
FETI / H-RIP	2986 (956)	7439 (1530)	16884 (2540)

TABLE III
COMPARISON OF DEGREES OF FREEDOM (ITERATION COUNTS) FOR GLOBAL PROBLEM WITH 9000 UNKNOWNS PER PROCESSOR

	8 Domains	16 Domains	32 Domains
D&C / METIS	86943 (29659)	172766 (48235)	314591 (68304)
FETI / METIS	4819 (1554)	12790 (3352)	29101 (12372)
FETI / METIS-SA	4684 (1121)	12319 (2349)	27335 (6130)
FETI / H-RIP	4421 (1128)	12020 (2002)	24886 (3431)

TABLE IV
COMPARISON OF SOLUTION TIMES IN CPU-SECONDS AND (RELATIVE SPEEDUP) FOR A FIXED PROBLEM WITH 86 943 UNKNOWNS

	8 Domains	16 Domains	32 Domains
D&C / METIS	1072 (1.0)	569 (1.88)	406 (2.64)
FETI / METIS	420 (1.0)	507 (0.82)	418 (1.00)
FETI / METIS-SA	382 (1.0)	406 (0.94)	242 (1.57)
FETI / H-RIP	323 (1.0)	249 (1.29)	130 (2.48)

be seen for the different partitioning methods that the number of Lagrange multipliers as well as the associated iteration counts vary considerably.

Figs. 2–4 report the solution time of both methods. For the divide-and-conquer method this is the iteration time of the BiCGM. For the FETI method this includes the factorization of the K_i matrices and then finding the solution to global problem (15) using BiCGM and the local problem (14) using a direct solver. Both methods used 10^{-6} as the termination criteria for the BiCGM. Because of this low value the iteration counts are relatively high. Table IV reports solution times and relative speedup for a fixed problem of 86 943 unknowns. Speedup is defined as execution time of y processors divided by the execution time for one processor. However, a measure of *relative speedup* can be defined as the execution time of y processors divided by the execution time for any number of x processors. Therefore, relative speedup is in *relation* to x processors instead of one processor if the serial time is not available. A mathematical expression of this definition is

$$\text{Relative Speedup}_{x\ \text{processors}} = \frac{\text{Execution Time}_{x\ \text{processors}}}{\text{Execution Time}_{y\ \text{processors}}},\ y \geq x. \qquad (15)$$

The relative time used in Table IV is the solution time for eight processors (i.e., $x = 8$). Therefore, linear relative speedup would be 2.0 for 16 processors and 4.0 for 32 processors.

Analyzing the results, it is obviously apparent that the FETI method reduces both the number of unknowns for the global problem as well as the iterations to solve it. For all problems considered, FETI with the heuristic RIP partitioning method produced the best results. However, FETI with METIS did not always produce better results than the divide-and-conquer method, especially as the number of processors of processors grow. METIS does an excellent job keeping the number of cells per processor constant. However, as a tradeoff, irregular boundaries between subdomains are produced. FETI unfortunately is highly sensitive to nonsmooth boundaries. By using simulated annealing, results were improved dramatically as is shown in the tables. By contrast, FETI with the heuristic RIP produces the most unbalanced loads regarding cells per processor. However, it does produce the best results with its regular boundaries.

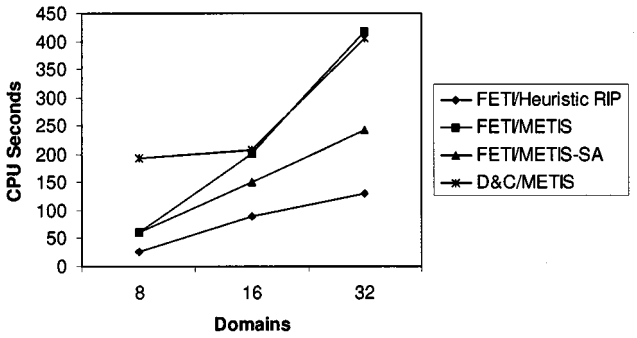

Fig. 2. Comparison of solution times in CPU-seconds for problem with 3000 unknowns per processor.

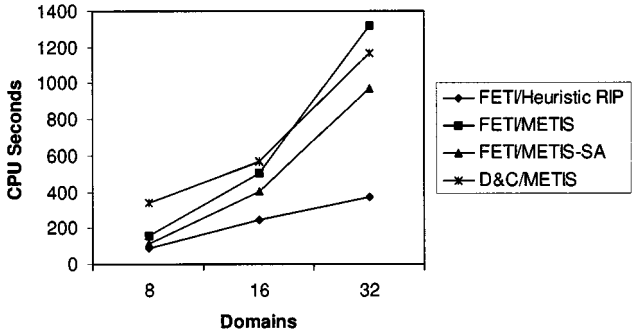

Fig. 3. Comparison of solution times in CPU-seconds for problem with 6000 unknowns per processor.

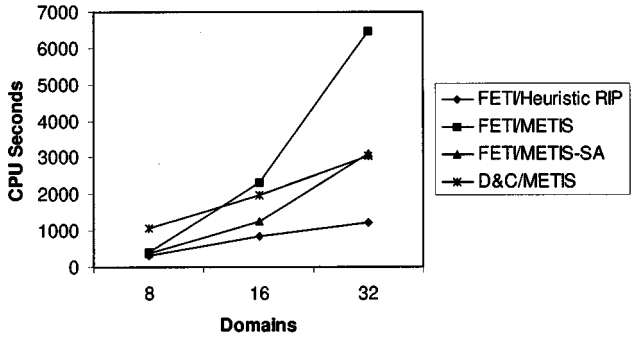

Fig. 4. Comparison of solution times in CPU-seconds for problem with 9000 unknowns per processor.

Also, it is noticeable that the time per iteration for the divide-and-conquer is considerably less. One of the main reasons is the preconditioner used. The diagonal preconditioner does a relatively good job and is very inexpensive. However, the preconditioner we used for the FETI method decreases the execution time but is relatively expensive. Other inexpensive preconditioners investigated had little or negative effect on the solution time.

V. CONCLUSIONS

A highly scalable domain decomposition algorithm referred to as the finite-element tearing and interconnecting (FETI) algorithm for solving electromagnetic problems with PML was presented. It was shown that this method can be more efficient for the analysis of general FEM problems on parallel computers

than a divider and conquer scheme. The FETI scheme is better suited for parallel computers with a moderate number of processors. One reason for this is that the algorithm is best suited when the Lagrange multipliers is kept small compared to the global problem size. Second, when using automatic partitioning schemes such as METIS, finely partitioned meshes tend to result in highly irregular domains. The condition of the reduced order matrix produced by the FETI algorithm can be sensitive to irregular domains. It was shown that this is circumvented through the use of an optimization method based on simulated annealing that smoothed out highly irregular domains [35]. The tradeoff is that the workload can become unbalanced. However, it was found that the FETI algorithm is much less sensitive to load imbalances than the shape of the mesh partitions.

It should also be noted that this method was investigated for use with PML. As was stated in the introduction, the addition of PML has a dramatic effect on the condition number of the global matrix. When the PML is removed the performance of the divide and conquer improves increases more dramatically than that of the FETI algorithm. This is due to, once again, the irregular boundaries.

Future work planned for maximizing the efficiency of the FETI algorithm is the development of an improved preconditioner that is computationally inexpensive. This will be necessary for the FETI algorithm to scale over a large numbers of processors.

ACKNOWLEDGMENT

The authors would like to thank C. Fisher and S. Sathaye at the University of Kentucky Computing Center for system support of the HP SPP2200.

REFERENCES

[1] I. S. Duff, A. M. Erisman, and J. K. Reid, *Direct Methods for Sparse Matrices*, New York: Oxford Univ. Press, 1986.

[2] W. Hackbusch, *Iterative Solution of Large Sparse Systems of Equations*, New York: Springer-Verlag, 1994.

[3] Y. Saad, *Iterative Methods for Sparse Linear Systems*. Boston, MA: PWS-Kent, 1996.

[4] B. Nour-Omid, A. Raefsky, and G. Lyzenga, "Solving finite element equations on concurrent computers," in *Symposium on Parallel Computation and their Impact on Mechanics*, Boston, MA, 1987.

[5] S.-H. Hsieh, G. Paulino, and J. Abel, "Recursive spectral algorithms for automatic domain partitioning in parallel finite element analysis," *Comput. Method. Appl. Mech. Eng.*, vol. 121, pp. 137–162, 1995.

[6] R. Van Driessche and D. Roose, "An improved spectral bisection algorithm and its application to dynamic load balancing," *Parallel Computing*, vol. 21, pp. 29–48, Jan. 1995.

[7] G. Karypis and V. Kumar, "A fast and high quality multilevel scheme for partitioning irregular graphs," Department of Computer Science, University of Minnesota, Tech. Rep. TR 95-035, 1995.

[8] C. Farhat and M. Lesoinne, "Automatic partitioning of unstructured meshes for the parallel solution of problems in computational mechanics," *Int. J. Numer. Methods Eng.*, vol. 36, pp. 745–764, 1993.

[9] C. Farhat, N. Maman, and W. Brown, "Mesh partitioning for implicit computations via iterative domain decomposition: Impact and optimization of the subdomain aspect ratio," *Int. J. Num. Meth. in Eng.*, vol. 38, pp. 989–1000, 1995.

[10] J. Patterson, T. Cwik, R. Ferraro, N. Jacobi, P. Liewer, T. Lockhart, G. Lyzenga, J. Parker, and D. Simoni, "Parallel computation applied to electromagnetic scattering and radiation analysis," *Electromagn.*, vol. 10, pp. 21–40, Jan.–June 1990.

[11] R. D. Ferraro, "Solving partial differential equations for electromagnetic scattering problems on coarse-grained concurrent computers," in *Computational Electromagnetics and Supercomputer Architecture*, T. Cwik and J. Patterson, Eds. Cambridge, MA: EMW, 1993, vol. 7, pp. 1111–154.

[12] R. Lee and V. Chupongstimun, "A partitioning technique for the finite element solution of electromagnetic scattering from electrically large dielectric cylinders," *IEEE Trans. Antennas Propagat.*, vol. 42, pp. 737–741, May 1994.

[13] Y. S. Choi-Grogan, K. Eswar, P. Sadayappan, and R. Lee, "Sequential and parallel implementations of the partitioning finite-element method," *IEEE Trans. Antennas Propagat.*, vol. 44, pp. 1609–1616, Dec. 1996.

[14] B. Després, "Domain decomposition method and boundaries the Helmholtz problem," in *Proc. Int. Symp. Math. Numer. Aspects Wave Propagat. Phenomena*. Strasbourg, France, 1992, pp. 44–52.

[15] ——, "Domain decomposition method and the Helmholtz problem," in *Proc. 2nd Int. Conf. Math. Aspects of Wave Propagat.*. Dover, DE, 1993, pp. 197–206.

[16] B. Stupfel, "A fast_domain decomposition method for the solution of electromagnetic scattering by large objects," *IEEE Trans. Antennas Propagat.*, vol. 44, pp. 1375–1385, Oct. 1996.

[17] ——, "Absorbing boundary conditions on arbitrary boundaries for the scalar and vector wave equations," *IEEE Trans. Antennas Propagat.*, vol. 42, pp. 773–780, 1994.

[18] C. Farhart and F. X. Roux, "A method of finite element tearing and interconnecting and its parallel solution algorithm," *Int. J. Numer. Method Eng.*, vol. 32, pp. 1205–1227, 1991.

[19] J. P. Berenger, "A PML for the absorption of electromagnetic waves," *J. Computat. Phys.*, vol. 114, no. 2, pp. 185–200, Oct. 1994.

[20] W. C. Chew and W. H. Weedon, "A 3-D perfectly matched medium form modified Maxwell's equations with stretched coordinates," *IEEE Microwave Guided Wave Lett.*, vol. 7, pp. 599–604, Sept. 1994.

[21] C. E. Reuter, R. M. Joseph, E. T. Thiele, D. S. Katz, and A. Taflove, "Ultrawideband absorbing boundary condition for termination of waveguiding structures in FD-TD simulations," *IEEE Microwave Guided Wave Lett.*, vol. 4, no. 10, pp. 344–346, Apr. 1994.

[22] D. S. Katz, E. T. Thiele, and A. Taflove, "Validation and extension to three-dimensions of the Berenger PML absorbing boundary condition for FD-TD meshes," *IEEE Microwave Guided Wave Lett.*, vol. 4, pp. 268–270, Aug. 1994.

[23] Z. S. Sacks, D. M. Kingsland, R. Lee, and J. F. Lee, "A perfectly matched anistropic absorber for use as an absorbing boundary condition," *IEEE Trans. Antennas Propagat.*, vol. 43, pp. 1460–1463, Dec. 1995.

[24] S. D. Gedney, "An anisotropic PML absorbing media for the truncation of FDTD Lattices," *IEEE Trans. Antennas Propagat.*, vol. 44, pp. 1630–1639, Dec. 1996.

[25] ——, "An anisotropic PML absorbing media for FDTD simulation for fields in lossy dispersive media," *Electromagn.*, vol. 16, pp. 399–416, July/Aug. 1996.

[26] ——, "The PML absorbing medium," in *Advances in Computational Electrodynamics: The Finite Difference Time Domain*, A. Taflove, Ed. Boston, MA: Artech House, 1998.

[27] C. M. Rappaport, "Perfectly matched absorbing boundary conditions based on anisotropic lossy mapping of space," *IEEE Microwave Guided Wave Lett.*, vol. 5, pp. 90–92, Mar. 1995.

[28] J. Gong and J. L. Volakis, "Optimal selection of uniaxial artificial absorber layer for truncating finite element meshes," *Electron. Lett.* , vol. 31, no. 18, pp. 1559–1561, Aug. 31, 1995.

[29] M. R. Lyons, A. C. Polycarpou, and C. A. Balanis, "On the accuracy of PMLs using a finite element formulation," in *IEEE Symp. Microwave Theory Tech.*, San Francisco, CA, June 1996, pp. 205–208.

[30] D. M. Kingsland, J. Gong, J. L. Volakis, and J. F. Lee, "Performance of an anisotropic artificial absorber for truncating finite-element meshes," *IEEE Trans. Antennas Propagat.*, vol. 44, pp. 975–982, July 1996.

[31] W. C. Chew and J. M. Jin, "PML's in the discretized space: An analysis and optimization," *Electromagn.*, vol. 16, pp. 325–340, July/Aug. 1996.

[32] Y. Y. Botros and J. L. Volakis, "Preconditioned generalized minimal residual (GMRES) solver for domains truncated by PML absorbers," in *1998 ACES Conf.*, Monterey, CA, Mar. 1998, pp. 639–646.

[33] P. C. Hammer, O. J. Marlowe, and A. H. Stroud, "Numerical integration over simplexes and cones," *Math. Tables Aids Comp.*, vol. 10, pp. 130–137, 1956.

[34] C. Farhat and F. X. Roux, "Implicit parallel processing in structural mechanics," in *Computat. Mech. Adv.*, 1994, vol. 2, pp. 1–124.

[35] D. Vanderstraeten and R. Keunings, "Optimized partitioning of unstructured finite element meshes," *Int. J. Numer. Methods in Eng.*, vol. 38, pp. 433–450, 1995.

C. T. Wolfe (M'90) was born in Danville, KY, in 1963. He received the B.S. degree in electrical engineering from the University of Kentucky, Lexington, and the M.S. degree in electrical engineering from the University of Missouri-Rolla, in 1985 and 1989, respectively. He is currently working toward the Ph.D. degree in electrical engineering at the University of Kentucky.

From 1985 to 1989, he worked in the Avionic Laboratories, McDonnell–Douglas, St. Louis, MO. Since 1989 he has worked at Lexmark International, Inc., Lexington, KY, designing networking products for their laser printers. His research interests include domain decomposition methods for use with finite elements and algorithmic specific computer architectures.

U. Navsariwala (S'94–M'98) received the Bachelor of Technology degree from the Indian Institute of Technology, Mumbai (Bombay), India, in 1991, and M. S. E. E. degree from University of Kentucky, Lexington, KY, in 1993.He is currently working toward the Ph.D. degree at the University of Kentucky.

From 1991 to 1998, he was a Research Assistant at the University of Kentucky, Lexington, where he worked on the development of finite-element time-domain (FDTD) algorithms, irregular unstructured FDTD and other computational electromagnetic algorithms. From 1998 to 1999 he was employed by T&M Antennas, where he was working as a Research Engineer, specializing in the development of finite-element method (FEM) software for antenna analysis and the design of small size antennas for use in cellular phone handsets. He is currently a Senior Engineer at Motorola Laboratories, where he is involved in the development of high performance FEM algorithms for three-dimensional electromagnetic fields solutions. His current research interest focus on conformal and internal antenna design, use of photonic bandgap structures for user terminal antennas, and model-order reduction algorithms.

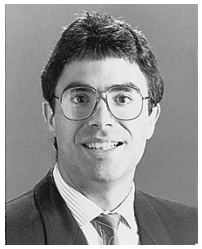

Stephen D. Gedney (S'84–M'91–SM'98) received the B.Eng. degree (honors) from McGill University, Montreal, PQ, Canada, in 1985, and the M.S. and Ph.D. degrees in electrical engineering from the University of Illinois, Urbana-Champaign, in 1987 and 1991, respectively.

He is an Associate Professor of electrical engineering at the University of Kentucky. From 1985 to 1987, he worked for the U.S. Army Corps of Engineers, Champaign, IL, where he was engaged in research in electromagnetic pulse simulation and propagation. Since 1991 he has been with the Department of Electrical Engineering at the University of Kentucky. In the summers of 1992 and 1993, he was a NASA/ASEE Faculty Fellow at the Jet Propulsion Laboratory, Pasadena, CA. In 1996, he was a visiting Professor at the Hughes Research Laboratories, Malibu, CA. His research includes computational electromagnetics, specifically in the development of rigorous solutions of Maxwell's equations for the analysis of microwave circuits and devices and electrical interconnects of electronic packages. He also specializes in the development of parallel algorithms for computational electromagnetics on high-performance parallel computers.