MULTIPOLE-ACCELERATED PRECONDITIONED ITERATIVE METHODS FOR SOLVING THREE-DIMENSIONAL MIXED FIRST AND SECOND KIND INTEGRAL EQUATIONS*

K. NABORS†, T. KORSMEYER‡ AND J. WHITE§

Abstract. First and second kind surface integral equations with $\frac{1}{r}$ and $\frac{\partial}{\partial n}\frac{1}{r}$ kernels are generated by a variety of three-dimensional engineering problems, and are typically solved using Galerkin or collocation schemes in which the singularities are integrated analytically. Such approaches lead to a dense $N \times N$ matrix problem, which may or may not be symmetric. In this paper, we present an overlapping-block preconditioned GMRES iterative algorithm for solving the dense matrix problem, where an adaptive multipole algorithm is used to compute the iterates. Experimental evidence is given to demonstrate that in practice, the combined algorithm is nearly order N.

1. Introduction. First and second kind surface integral equations with $\frac{1}{r}$ and $\frac{\partial}{\partial n}\frac{1}{r}$ kernels are generated by a variety of three-dimensional engineering problems. For such problems, Nystr-type algorithms can not be used directly, but an expansion for the unknown, rather than for the entire integrand, can be assumed and the product of singular kernel and the unknown integrated analytically. Combining such a substitution with a Galerkin or collocation scheme for computing the expansion coefficients gives a general discretization technique, but leads to dense matrix problems. In this paper, we present an overlapping-block preconditioned GMRES iterative algorithm for solving the dense matrix problem, where an adaptive multipole algorithm is used to compute the iterates. Note that this approach follows along lines originally suggested in [1].

The outline of the paper is as follows. The boundary-element formulation and the GMRES iterative algorithm for solving the generated matrix problem are reviewed in Section 2. A brief description of the hierarchical multipole algorithm is given in Section 3. The preconditioning strategy for accelerating GMRES convergence is presented in Section 4. Experimental results using our program, FASTCAP, to analyze a wide variety of structures are presented in Section 5. Finally, conclusions and acknowledgments are given in Section 6.

2. Formulation. For exposition in this short paper, we consider only the first-kind formulation used to compute electrostatic capacitances or forces in three-dimensional structures. Electrostatic capacitances are useful figures of merit for designers of electronic packaging [2], and microelectromechanical system designers are interested in electrostatic forces [3]. The exterior Dirichlet problem on the associated unbounded, multiply-connected domain corresponding to the conductor surfaces in an infinite homogeneous medium is formulated using a single layer charge density denoted σ [2, 4] (for a second kind formulation of this exterior problem see [5]). It then follows that σ must satisfy the integral equation

$$\psi(x) = \int_{surfaces} \sigma(x') \frac{1}{\|x - x'\|} da', \qquad x \in surfaces. \tag{1}$$

where $\psi(x)$ is the known conductor surface potential, da' is the incremental conductor surface area, x, $x' \in \mathbf{R}^3$, and $\|x\|$ is the usual Euclidean length of x given by $\sqrt{x_1^2 + x_2^2 + x_3^2}$.

A standard approach to numerically solving (1) for σ begins by expanding σ with piece-wise constant finite elements and then using point collocation. That is, the conductor surfaces are broken into N small panels or tiles, and it is assumed that on each panel i, a charge, q_i, is uniformly distributed. Then for each panel, an equation is written which relates the known potential at the center of that i-th panel, denoted

* This work was supported by the Defense Advanced Research Projects Agency contracts N00014-91-J-1698 and N00014-90-J-1085, the National Science Foundation contract (MIP-8858764 A02), F.B.I. contract J-FBI-88-067, and grants from Digital Equipment Corporation and I.B.M.

† Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139. (ksn@rle-vlsi.mit.edu)

‡ Dept. of Ocean Engineering, Massachusetts Institute of Technology Cambridge, MA 02139. (xmeyer@flying-cloud.mit.edu)

§ Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology Cambridge, MA 02139. (white@rle-vlsi.mit.edu)

p_i, to the sum of the contributions to that potential from the N charge distributions on all N panels [4]. The result is a dense linear system,

$$Pq = p \tag{2}$$

where $P \in \mathbf{R}^{N \times N}$, q is the vector of panel charges, $p \in \mathbf{R}^N$ is the vector of known panel potentials, and

$$P_{ij} = \frac{1}{a_j} \int_{panel_j} \frac{1}{\|x_i - x'\|} \, da', \tag{3}$$

where x_i is the center of the i-th panel and a_j is the area of the j-th panel. The integral in (3) can be evaluated analytically [6], and since the discretization uses point collocation, P can be unsymmetric.

The dense linear system of (2) can be solved to compute panel charges from a given set of panel potentials, and the electrostatic capacitances can be derived by summing the panel charges. If Gaussian elimination is used to solve (2), the number of operations is order N^3. Clearly, this approach becomes computationally intractable if the number of panels exceeds several hundred. Instead, consider solving the linear system (2) using a conjugate-residual style iterative method like GMRES [7]. Such methods have the form given below:

Algorithm 1: GMRES algorithm for solving (2)

Make an initial guess to the solution, q^0.
Set $k = 0$.
do {
  Compute the residual, $r^k = p - Pq^k$.
  if $\|r\| < tol$, return q^k as the solution.
  else {
    Choose α's and β in
      $q^{k+1} = \sum_{j=0}^{k} \alpha_j q^j + \beta r^k$
    to minimize $\|r^{k+1}\|$.
    Set $k = k + 1$.
  }
}

The dominant costs of Algorithm 1 are in calculating the N^2 entries of P using (3) before the iterations begin, and performing N^2 operations to compute Pq^k on each iteration. Described below is our adaptive hierarchical multipole algorithm which, through the use of carefully applied approximations, avoids forming most of P and reduces the cost of forming Pq^k to order N operations. This does not necessarily imply that each iteration of the GMRES algorithm can be computed with order N operations. If the number of GMRES iterations required to achieve convergence approaches N, then to perform the minimization in each GMRES iteration will require order N^2 operations. This problem is avoided through the use of a preconditioner, also described below, which reduces the number of GMRES iterations required to achieve convergence to well below N for large problems.

3. The Multipole Algorithm. Low order multipole and local expansions can be used to accurately compute the potentials at N evaluation points due to N panel charges in order N operations, regardless of the evaluation point and charge distributions. Below, we summarize the hierarchical multipole algorithm, in a somewhat modified form from that originally presented in [8, 9]. The multipole algorithm computes an approximation to the potential due to a charge distribution, but the approximation is still a linear function of the charges. It is possible to represent this linear transformation directly, in terms of transformation matrices whose entries are functions only of geometry. This is an important observation for two reasons: the transformation matrices need be computed only once, but can then be reused many times in the inner loop of the GMRES algorithm; and the formulation in terms of matrices will make transparent a simple adaptive multipole algorithm.

To begin, the smallest cube that contains the entire collection of panels for the problem of interest is determined. This cube will be referred to as the level 0, or root, cube. Then, the volume of the cube is broken into eight equally sized child cubes, referred to as level 1 cubes, and each has the level 0 cube as its parent. The panels are divided among the child cubes by associating a panel with a cube if the panel's center point is contained in the cube. Each of the level 1 cubes is then subdivided into eight level 2 child cubes and the panels are again distributed based on their center point locations. The result is a collection of 64 level 2 cubes and a 64-way partition of the panels. This process is repeated to produce l levels of cubes, and l partitions of panels starting with an 8-way partition and ending with an 8^l-way partition. The number of levels, l, is chosen so that the maximum number of panels in a finest, or l-th, level cube is less than the number of coefficients in a the multipole expansion (see below).

Remark *As the charges in this problem are not point charges, but are distributed on panels, it is necessary to ensure that each panel is mostly contained in a finest level cube.*

The following terms are used to concisely describe the multipole algorithm, Algorithm 2.

Definition 3.1. Evaluation Points of a Cube: The center points of the panels associated with the cube.

Definition 3.2. Nearest Neighbors of a Cube: Those cubes which have a corner in common with the given cube.

Definition 3.3. Second Nearest Neighbors of a Cube: Those cubes which are not nearest neighbors but have a corner in common with a nearest neighbor. Note that there are at most 124 nearest and second nearest neighbors of a cube, excluding the cube itself.

Definition 3.4. Interaction Cubes of a given cube: Those cubes which are either the second nearest neighbors of the given cube's parent, or are children of the given cube's parent's nearest neighbors, excluding nearest or second nearest neighbors of the given cube. There are a maximum of 189 interaction cubes for a given cube, roughly half are from a level one coarser than the level of the given cube, the rest are on the same level.

For the j-th cube on level i: its parent's index on level $i-1$ is denoted $F(i,j)$; its associated set of child-cube indices is denoted $C(i,j)$; its set of interaction-cube index pairs is denoted $I(i,j)$; the vector of multipole-expansion coefficients representing the panel charges in the cube is denoted $M_{i,j}$; the vector of local-expansion coefficients representing the evaluation point potentials in $I(i,j)$ is denoted $L_{i,j}$; and if the cube is a finest-level cube, $N(l,j)$ is its set of nearest and second-nearest neighbor cube indices, and $q_{l,j}$ and $p_{l,j}$ are its panel charge and panel potential vectors, respectively. The matrices $Q2M$, $M2M$, $M2L$, $L2L$, and $L2P$ represent charge to multipole, multipole to multipole, multipole to local, local to local, and local to potential transformation operators [10].

A general approach to making Algorithm 2 adaptive is to break up the problem domain nonuniformly, leading to finest-level cubes of different sizes, where the sizes are chosen so that each finest-level cube has roughly the same number of panels [11]. For the kinds of charge and evaluation point distributions common to boundary-element problems, such an approach can sometimes require more computation than a nonadaptive algorithm [12]. Instead, a more effective approach in this setting is to avoid transformation from panel charges to multipole expansions, and to avoid forming local expansions, whenever such representations are inefficient. For example, consider computing $M_{l,j}$ from $Q2M(l,j)q_{l,j}$. If the number of entries in vector $q_{l,j}$ is smaller than the number of multipole coefficients in $M_{l,j}$, which for a v-th order multipole expansion is $(v+1)^2$,then $q_{l,j}$ is a more efficient representation of the charge distribution, and that representation can be propagated through the algorithm instead of $M_{l,j}$. Note that using such an approach requires some rather straight-forward bookkeeping and a few more transformation operators.

4. Preconditioning the Iterative Method. In general, the GMRES iterative method applied to solving (2) can be significantly accelerated by *preconditioning* if there is an easily computed good approximation to the inverse of P. We denote the approximation to P^{-1} by $\tilde{C}$, in which case preconditioning the GMRES algorithm is equivalent to using GMRES to solve

$$P\tilde{C}x = p. \tag{4}$$

Algorithm 2: Multipole Algorithm for Computing $p = Pq$

/* THE UPWARD PASS: Computes Multipole expansions for every cube at every level. */
For each finest-level cube $j = 1$ to 8^l {
$M_{l,j} = Q2M(l,j)q_{l,j}$
}
For each level $i = l-1$ to 2 {
For each level i cube $j = 1$ to 8^i {
$M_{i,j} = \sum_{\tilde{j} \in C(i,j)} M2M(i,j,i+1,\tilde{j})M_{i+1,\tilde{j}}$
}
}
/* THE DOWNWARD PASS: Computes Local expansions for every cube. */
For each level $i = 2$ to l {
For each level i cube $j = 1$ to 8^i {
$L_{i,j} = L2L(i,j,i-1,F(i,j))L_{i-1,F(i,j)} + \sum_{\tilde{i},\tilde{j} \in I(i,j)} M2L(i,j,\tilde{i},\tilde{j})M_{\tilde{i},\tilde{j}}$
}
}
/* EVALUATION: Evaluates the potential in each finest-level cube. */
For each finest-level cube $j = 1$ to 8^l {
$p_{l,j} = L2P(l,j)L_{l,j} + \sum_{\tilde{j} \in N(l,j)} Q2P(l,j,l,\tilde{j})q_{l,j}$
}

for the unknown vector x, from which the charge density is computed by $q = \tilde{C}x$. Clearly, if $\tilde{C}$ is precisely P^{-1}, then (4) is trivial to solve, but then $\tilde{C}$ will be very expensive to compute.

In [13] and [14], it was suggested that a good approximation to P^{-1} can be derived by solving overlapping subproblems. Such an approach fits naturally with the hierarchical multipole algorithm because the preconditioner can be constructed and applied in a cube-by-cube fashion. First, $\tilde{C}$ is computed by inverting a sequence of reduced P matrices, one associated with each finest-level cube, as in Algorithm 3 below.

Algorithm 3: Forming $\tilde{C}$.

For each finest-level cube $i = 1$ to 8^l {
Form P^i, the potential coefficient matrix for the
reduced problem considering only the panels contained
in cube i and cube i's neighbors.
Compute $\tilde{C}^i = (P^i)^{-1}$.
For each panel k in cube i or cube i's neighbors {
if panel k is not in cube i {
delete row k from $\tilde{C}^i$.
}
}
}

The $\tilde{C}^i$ matrices taken together form a convenient representation of $\tilde{C}$ with each of the N rows in $\tilde{C}$ appearing as a row in one of the $\tilde{C}^i$ matrices. Comparing Algorithm 3 with Algorithm 2 indicates that the association of each $\tilde{C}^i = (P^i)^{-1}$ with a particular finest-level cube in Algorithm 3 guarantees the use of only those elements of the full P matrix which are already required in Algorithm 2, and therefore the

additional computation involved in calculating the preconditioner is only the inversions of the small P^i matrices.

The product $P\tilde{C}x^k$, which is used in a preconditioned GMRES algorithm applied to solving (4), is computed in two steps. First, the preconditioner is applied to form $q^k = \tilde{C}x^k$ using Algorithm 4 below. Then, Pq^k is computed using Algorithm 2 in the previous section.

Algorithm 4: Forming $q = \tilde{C}x$.

For each finest-level cube $i = 1$ to 8^l {
  For each panel j in finest-level cube i {
    For each panel k in cube i or its neighbors {
      Add $\tilde{C}^i_{jk}x_k$ to q_j.
    }
  }
}

5. Experimental Results. In this section, results from computational experiments using our program FASTCAP are presented to demonstrate that the preconditioned, adaptive, multipole-accelerated (PAMA) 3-D capacitance extraction algorithm described above really deserves a four-letter mnemonic [12]. The structures described below were created with the solid-modeling program PATRAN, or by computer program, and all capacitance calculations were performed using FASTCAP. The multipole-accelerated algorithms in FASTCAP use, by default, second-order multipole expansions and a GMRES convergence tolerance (see Algorithm 1) of 0.01.

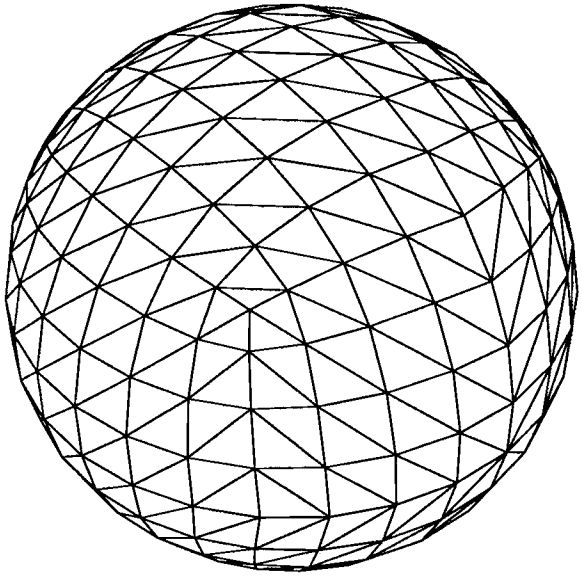

FIG. 1. *The Sphere1 discretization of the unit sphere.*

To demonstrate absolute accuracy, the FASTCAP program was used to compute the capacitance of a unit sphere, discretized as in Figure 1, and a unit cube, discretized as in Figure 2. In Table 1, the capacitances computed using the PAMA algorithm are compared with the capacitances computed using direct factorization of P in (2) (Direct), and with analytic results for the unit sphere and with reference results for the unit cube. As can be seen from the table, the results using the PAMA algorithm are easily within one percent of the analytic or reference results.

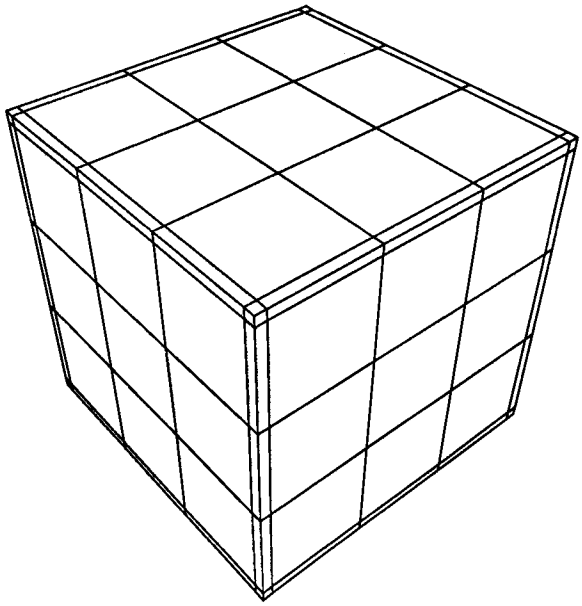

FIG. 2. *The Cube1 discretization of the unit cube.*

Method	Problem	
	Sphere1 768 panels	Cube1 150 panels
Direct	110.6	73.26
PAMA	110.5	73.28
Other	111†	73.5, 73.4

TABLE 1

Capacitance values (in pF) illustrating FASTCAP's accuracy. †By analytic calculation. ‡From [2], [15].

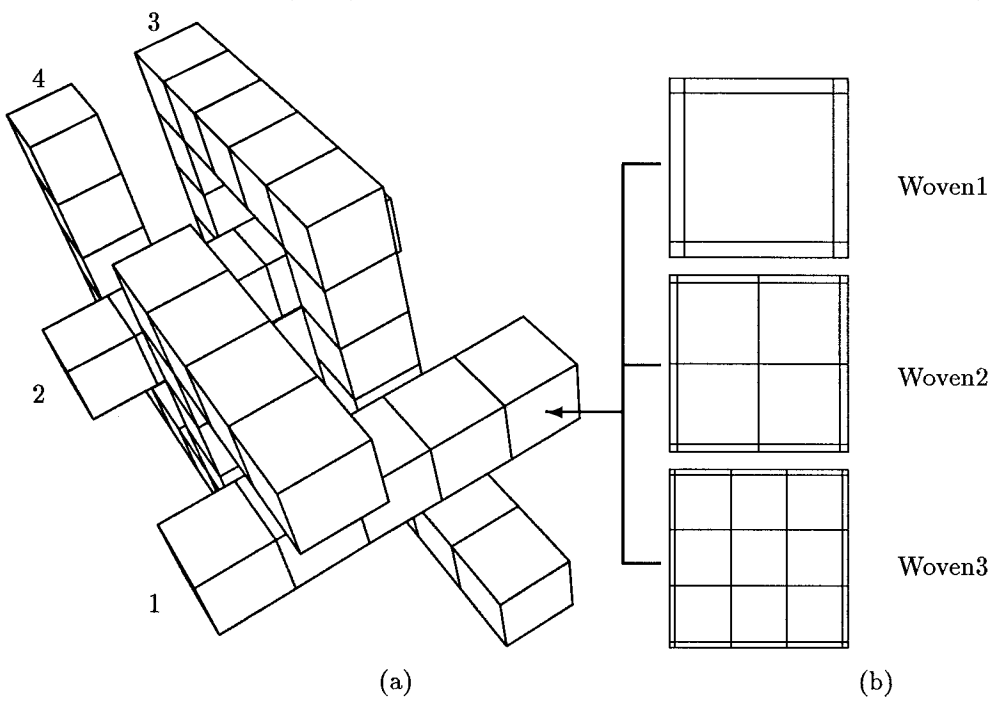

FIG. 3. *The 2 × 2 woven bus problem: bars have 1m×1m cross sections. The three discretizations are obtained by replacing each square face in (a) with the corresponding set of panels in (b).*

The PAMA algorithm is nearly as accurate as the direct factorization method even on more complex problems, such as the 2 × 2 woven bus structure in Figure 3. The capacitances computed using the two methods are compared in Table 2, using coarse, medium, and fine discretizations of the woven bus structure, also shown in Figure 3. Note that the coupling capacitance C_{12} between conductors one and two, which is forty-times smaller than the self-capacitance C_{11}, is computed nearly as accurately with the PAMA algorithm as with direct factorization.

Method	Problem					
	Woven1 1584 Panels		Woven2 2816 Panels		Woven3 4400 Panels	
	C_{11}	C_{12}	C_{11}	C_{12}	C_{11}	C_{12}
Direct	251.6	−6.353	253.2	−6.446	253.7	−6.467
PAMA	251.8	−6.246	253.3	−6.334	253.9	−6.377

TABLE 2

Capacitance values (in pF) illustrating FASTCAP's accuracy for the complicated geometry of Figure 3.

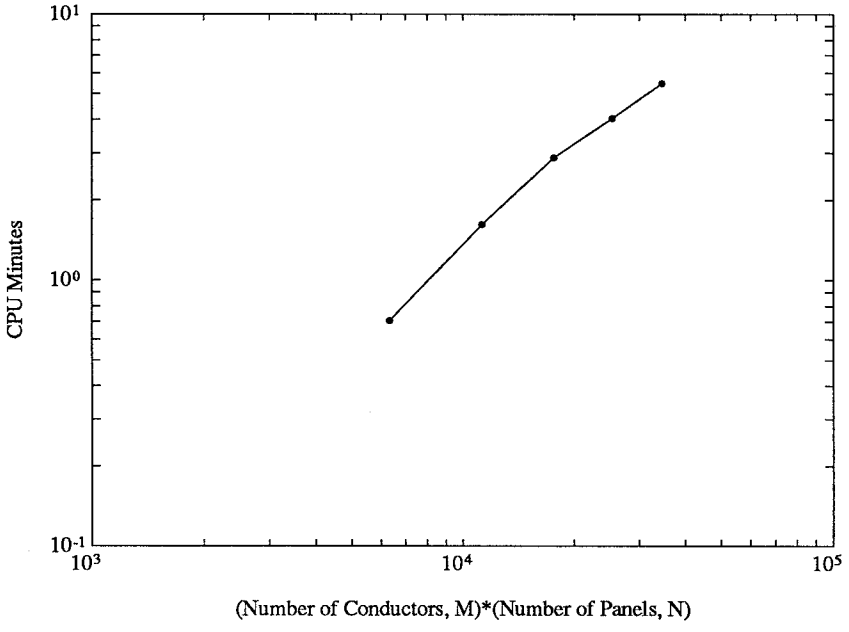

FIG. 4. *Executione time as a function of mn for the MA, AMA, and PAMA algorithms applied solving to progessively finer discretizations of the 2 × 2 woven bus problem.*

The computational cost of using the FASTCAP program is roughly proportional to the product of the number of conductors, M, and the number of panels N. This is experimentally verified by computing the capacitances of the 2 × 2 woven bus structure in Figure 3, with progressively finer discretizations. In Figure 4, the execution times required to compute these capacitances are plotted as a function of MN, and as the graph demonstrates, the execution time does grow nearly linearly.

To demonstrate the effectiveness of various aspects of the PAMA algorithm on a range of problems, in Table 3 the execution times required to compute the capacitances of six different examples using four different methods are given. The examples Cube2 and Sphere2 are finer discretizations of the unit cube and sphere in Figures 1 and 2; the 2 × 2 Woven Bus example is the Woven3 discretization of the Figure 3 problem; the 5 × 5 Woven Bus is the geometry of Figure 3(a) extended to five conductors crossing five conductors and discretized as for Woven1 (Figure 3(b)); the example Via, shown in Figure 5, models a pair of connections between integrated circuit pins and a chip-carrier; and the example Diaphragm, shown in

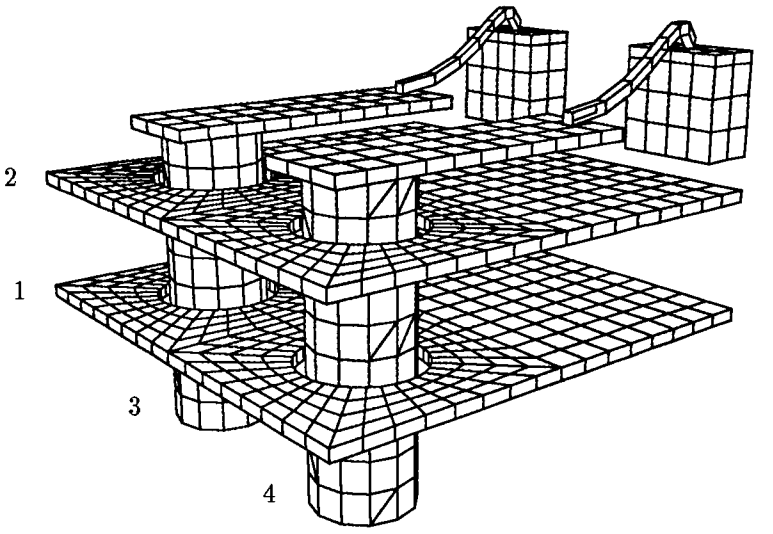

FIG. 5. *Two signal lines passing through conducting planes; via centers are 2mm appart.*

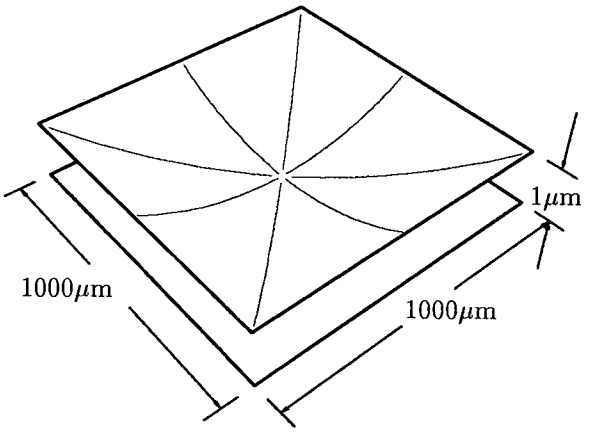

FIG. 6. *A schematic illustration of the diaphragm problem. The two plates are* $0.02\mu m$ *appart at the center.*

Figure 6, is a model for a microsensor [16].

Comparing the execution times for the larger, realistic examples in Table 3, it can be seen that using the adaptive multipole algorithm (AMA) improves FASTCAP's execution time by a factor of two over using the multipole algorithm (MA) alone, and that using the preconditioner can reduce the execution time by nearly a factor of five. The improvement due to the adaptive algorithm is small because it is being compared to the MA algorithm in [10], which ignores empty cubes, and is therefore already somewhat adaptive. Exploiting empty cubes is trivial to implement, and makes an enormous difference. For the largest problem, the 5×5 Woven Bus, more than 252,000 out of 262,000 cubes used to partition the problem domain are empty.

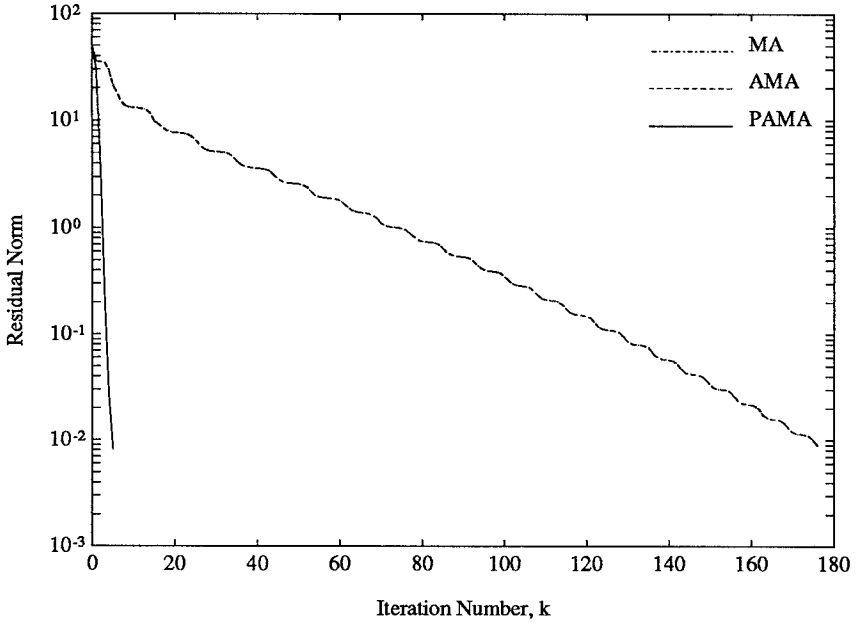

FIG. 7. *The GMRES residual norms for the linear system solution corresponding to the Diaphragm problem (Figure 6) with the top conductor at unit potential. As is evident here, the AMA method generally makes the iterate calculation more efficient, while the PAMA algorithm leads to fewer iterations.*

The reduction in execution time afforded by the adaptive algorithm (AMA) over the normal multipole algorithm (MA) stems from more efficient Pq product computation in each iteration of the GMRES algorithm. Using the preconditioner reduces execution time by decreasing the number of iterations required to achieve convergence. Because of various program overheads, comparing total execution times can hide the sometimes dramatic effect the preconditioner can have on GMRES convergence. To show the impact of the preconditioner more directly, in Figure 7 we plot the norm of the residual, $\|p - Pq^k\|$, as a function of iteration for the MA, AMA, and PAMA algorithms applied to solving the Diaphragm problem. As is evident in the figure, the MA and AMA algorithms converge nearly identically, as expected, but the residuals computed with the PAMA algorithm drop considerably faster. It is this rapid convergence that easily offsets the fact that one PAMA iterate requires slightly more computation than an AMA iterate.

6. Conclusions and Acknowledgments. In this short paper, a preconditioned, adaptive, multipole-accelerated approach to solving first-kind surface integral equations with $\frac{1}{r}$ and $\frac{\partial}{\partial n}\frac{1}{r}$ kernels was described, and shown to be effective for several engineering problems. Also, experimental evidence is given to demonstrate that in practice, the combined algorithm is nearly order N. Several analytical issues have not been addressed in this short paper, but will be considered more carefully in forthcoming work. The calculation of multipole expansion coefficients for the charge distribution over panels can be performed analytically, though the derivation is lengthy. The preconditioner has a stronger theoretical justification, and is in fact more robust in practice, than the presentation would suggest.

Method	Cube2 294 Panels	Sphere2 1200 Panels	2x2 Woven Bus 4400 Panels	Via 6185 Panels	Diaphragm 7488 Panels	5x5 Woven Bus 9630 Panels
Direct	0.11	3.2	185	(490)	(890)	(1920)
MA	0.06	0.3	6.0	11	8.7	42
AMA	0.05	0.2	3.3	4.7	5.9	23
PAMA	0.05	0.2	2.3	3.2	1.3	11

TABLE 3
CPU times in minutes on an IBM RS6000/540. Times in parentheses are extrapolated.

The authors would like to thank Prof. Senturia for many valuable discussions on various approaches to preconditioning, and for his and Brian Johnson's help in linking FASTCAP with the PATRAN solid-modeling program. The authors would also like to thank David Ling and Albert Ruehli of the I.B.M. T. J. Watson Research Center for their helpful suggestions about capacitance calculations.

REFERENCES

[1] V. Rohklin, "Rapid solution of integral equation of classical potential theory," *J. Comput. Phys.*, vol. 60, pp. 187–207, 1985.

[2] A. E. Ruehli and P. A. Brennan, "Efficient capacitance calculations for three-dimensional multiconductor systems," *IEEE Transactions on Microwave Theory and Techniques*, vol. 21, pp. 76–82, February 1973.

[3] S. D. Senturia, R. M. Harris, B. P. Johnson, S. Kim, K. Nabors, M. A. Shulman, and J. K. White, "A computer-aided design system for microelectromechanical systems (memcad)," *IEEE Journal of Microelectromechanical Systems*, vol. 1, pp. 3–13, March 1992.

[4] S. M. Rao, T. K. Sarkar, and R. F. Harrington, "The electrostatic field of conducting bodies in multiple dielectric media," *IEEE Transactions on Microwave Theory and Techniques*, vol. MTT-32, pp. 1441–1448, November 1984.

[5] A. Greenbaum, L. Greengard, and G. B. McFadden, "Laplace's equation and the dirchlet-neumann map in multiply connected domains," tech. rep., Courant Institute, New York, March 1991.

[6] J. N. Newman, "Distributions of sources and normal dipoles over a quadralateral panel," *Journal of Engineering Mathematics*, no. 20, pp. 113–126, 1986.

[7] Y. Saad and M. H. Schultz, "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM Journal on Scientific and Statistical Computing*, vol. 7, pp. 856–869, July 1986.

[8] L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *Journal of Computational Physics*, vol. 73, pp. 325–348, December 1987.

[9] L. Greengard, *The Rapid Evaluation of Potential Fields in Particle Systems.* Cambridge, Massachusetts: M.I.T. Press, 1988.

[10] K. Nabors and J. White, "Fastcap: A multipole accelerated 3-D capacitance extraction program," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 10, pp. 1447–1459, November 1991.

[11] J. Carrier, L. Greengard, and V. Rokhlin, "A fast adaptive multipole algorithm for particle simulations," *SIAM Journal on Scientific and Statistical Computing*, vol. 9, pp. 669–686, July 1988.

[12] K. Nabors and J. White, "Fast capacitance extraction of general three-dimensional structures," *IEEE Trans. on Microwave Theory and Techniques*, June 1992.

[13] S. A. Vavasis, "Preconditioning for boundary integral equations," in *Proceedings of the Copper. Conf. on Iterative Methods*, April 1991.

[14] K. Nabors, S. Kim, J. White, and S. D. Senturia, "An adaptive multipole algorithm for 3-d capacitance calculation," in *Proceedings of Int. Conf. on Comp. Design*, October 1991.

[15] M. A. Jaswon and G. T. Symm, *Integral Equation Methods in Potential Theory and Elastostatics.* London: Academic Press, 1977.

[16] B. P. Johnson, S. Kim, S. D. Senturia, and J. White, "MEMCAD capacitance calculations for mechanically deformed square diaphragm and beam microstructures," in *Proceedings of Transducers 91*, June 1991.