Preconditioning for Multipole-Accelerated 3-D Inductance Extraction.

M. Kamon J. White
Research Laboratory of Electronics and the
Department of Electrical Engineering and Computer Science,
Massachusetts Institute of Technology, Cambridge, MA 02139.
Phone: (617) 253-2543 FAX: (617) 258-7859 *

1 Introduction

A wide range of integrated circuit and packaging design problems require accurate estimates of the coupling inductances of complicated three-dimensional structures. To perform the required magnetoquasistatic analysis, the most commonly used approach is to apply finite-difference or finite-element techniques to a differential problem formulation. However, finite-element techniques require that the entire 3-D volume be discretized, and generating such a volume discretization for complex structures can become cumbersome. Instead, volume-element methods can be applied to solving integral formulations of the problem, in which case only the interior of the conductors need be discretized. Unfortunately, volume-element methods generate dense matrix problems which, if solved directly, grow in computational cost like n^3, where n is the number of elements into which the problem is discretized.

Recently, solving the dense matrices associated with boundary or volume-element methods has been made substantially more efficient through the use of iterative solution techniques accelerated by "fast-multipole" algorithms. This combined approach reduces the computational (and storage) cost of using boundary and volume-element methods to nearly $O(n)$, provided the number of iterations required to achieve convergence does not increase with problem size. In order to reduce the number of iterations required, the matrix can be "pre-conditioned", that is the matrix can be pre- or post-multiplied by an easily computed approximation to its inverse. Below, we give a brief description of the multipole accelerated approach, followed by two approaches to preconditioning, one based on local inversion and one based on direct factorization of the sparse resistance matrix. Examples are given to demonstrate that the preconditioners substantially accelerate convergence.

2 Inductance Formulation

In the case of magnetoquasistatic analysis, used for extracting inductances, the conductor current density, J, satisfies $\nabla \cdot J = 0$ and for any point x in the conductor,

$$\frac{1}{\sigma}J(x) + \frac{j\omega\mu}{4\pi}\int_{V'}\frac{J(x')}{\|x-x'\|}dV' = -\nabla\psi(x), \tag{1}$$

where here ψ is a scalar potential, σ is the conductivity, μ is the magnetic permeability, ω is the frequency of interest, and V' is the conductor volume.

To numerically compute J, the conductor volume is discretized into b filaments, and in each filament the conductor current is assumed constant. A system of equations for the filament currents, which are denoted by the vector I_b, is then generated by insisting that at filament intersection points, the directed sum of currents associated with the intersecting filaments is zero. In addition, the filament currents must satisfy

$$ZI_b = (R + j\omega L)I_b = V_b, \tag{2}$$

where V_b, $I_b \in C^b$, b is the number of branches (number of current filaments), $Z \in C^{b\times b}$ is the complex impedance matrix, $R \in \Re^{b\times b}$ is the diagonal matrix whose elements are associated with the dc resistance of each current filament, and $L \in \Re^{b\times b}$ is the dense matrix of partial inductances [1]. Specifically,

$$L_{i,j} = \frac{\mu_0}{4\pi}\int_{filament_i}\int_{filament_j}\frac{l_i(X_i)\cdot l_j(X_j)}{|X_i - X_j|}d^3x_i d^3x_j, \tag{3}$$

where $X_i, X_j \in \Re^3$ are the positions in filament i and j respectively, and $l_i, l_j \in \Re^3$ are the unit vectors in the direction of current flow in filaments i and j. Using mesh analysis, it is possible to combine the current conservation constraint with (2) to yield

$$MZM^t I_m = V_s \tag{4}$$

where $I_m \in \Re^n$ is a vector of mesh currents, $M \in \Re^{n\times b}$ is the mesh matrix, V_s is the mostly zero vector of source voltages, and n is the number of fundamental loops or meshes.

*This work was supported by the Advanced Research Projects Agency contract N00014-91-J-1698, a National Science Foundation Fellowship and contract (MIP-8858764 A02), and grants from I.B.M. and Digital Equipment Corporation.

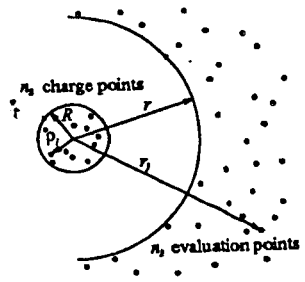

Figure 1: Exploiting charge clusters using multipole expansions.

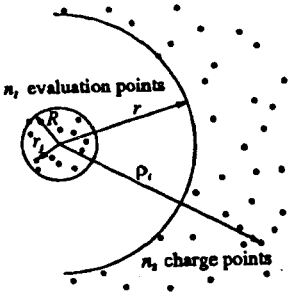

Figure 2: Exploiting evaluation point clusters using local expansions.

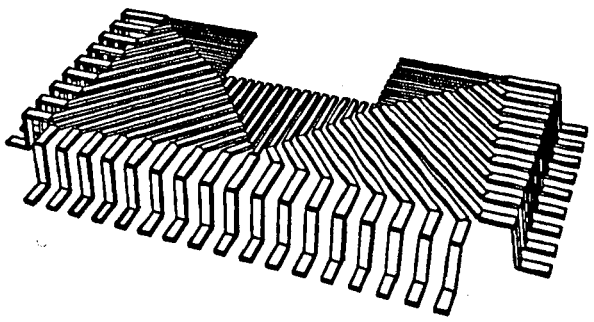

Figure 3: Half of a pin-connect structure. Thirty-five pins shown.

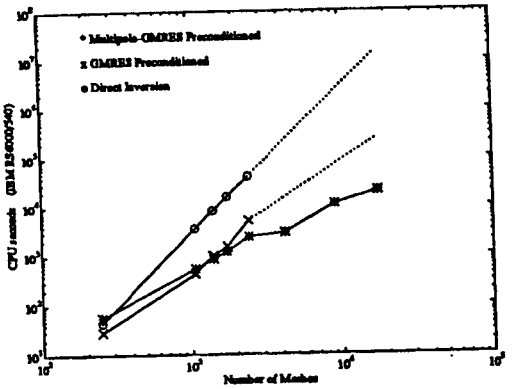

Figure 4: CPU times (IBM RS6000/540).

If an iterative algorithm, typically GMRES [2], is used to solve (4), then each iteration of the algorithm will cost n^2 operations. This follows from the fact that in each GMRES iteration, MZM^t is multiplied by a vector, and although MRM^t is very sparse, MLM^t is dense. However, it is possible to reduce the matrix-vector product cost to order n operations by using a fast multipole algorithm [3, 4, 5]. The fast multipole algorithm will compute n potentials due to n charges in order n operations, and the algorithm can be applied to reducing the cost of forming $MLM^t I_m$ by decomposing the product into three component potential evaluations.

As a brief explanation of how the fast multipole algorithm can be used to evaluate n potentials due to n charges in order n operations, consider the two configurations, depicted in 2-D for simplicity, given in Figs. 1 and 2. In either figure, the obvious approach to determining the electrostatic potential at the n_1 evaluation points from the n_2 point-charges involves $n_1 * n_2$ operations; at each of the n_1 evaluation points one simply sums the contribution to the potential from n_2 charges. An accurate approximation for the potentials for the case of Fig. 1 can be computed in many fewer operations using *multipole expansions*, which exploit the fact that $r >> R$ (defined in Fig. 1). That is, the details of the distribution of the charges in the inner circle of radius R in Fig. 1 do not strongly effect the potentials at the evaluation points outside the outer circle of radius r. It is also possible to compute an accurate approximation for the potentials at the evaluation points in the inner circle of Fig. 2 in many fewer than $n_1 * n_2$ operations using *local expansions*, which again exploit the fact that $r >> R$ (as in Fig. 2). In this second case, what can be ignored is the details of the evaluation point distribution.

To see the effectiveness of multipole acceleration, consider a typical industrial example, a 35-pin package shown in Fig. 3. The cost of computing the associated inductance matrix using direct factorization, preconditioned GMRES, and multipole-accelerated preconditioned GMRES, as a function of discretization refinement is show in Fig. 4. The preconditioner used is described in the next section. Note that for a 10,000 filament discretization, which is barely sufficient to properly model the finite penetration depth and proximity effects, the multipole accelerated algorithm is two orders of magnitude faster than direct factorization, and an order of magnitude faster than explicit GMRES.

3 Approaches to Preconditioning

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

$$(MZM^t)Px = V_s \tag{5}$$

for the unknown vector x. The mesh currents are then computed with $I_m = Px$. Clearly, if P is precisely $(MZM^t)^{-1}$, then (5) is trivial to solve, but then P will be very expensive to compute.

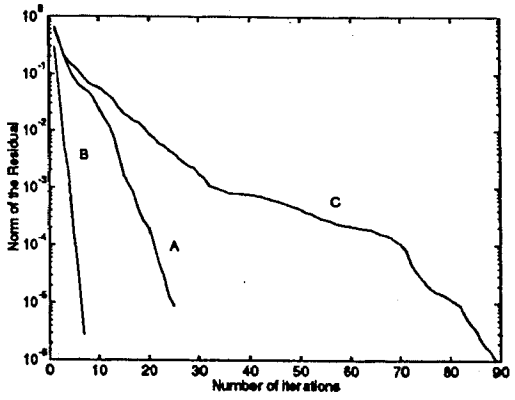

Figure 5: Convergence of GMRES applied to the packaging example with inverse-R preconditioning (A), local inversion preconditioning (B), and no preconditioner (C).

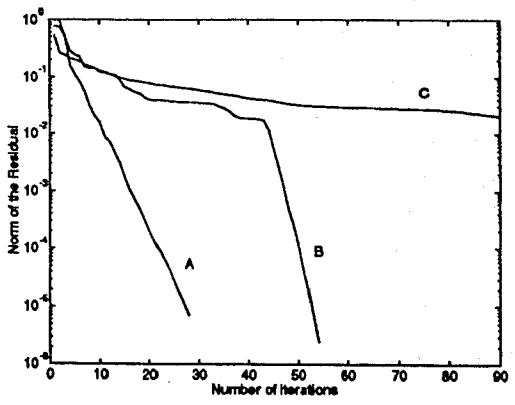

Figure 6: Convergence of GMRES applied to the ground plane example with inverse-R preconditioning (A), local inversion preconditioning (B), and no preconditioner (C).

An easily computed good approximation to $(MZM^t)^{-1}$ can be constructed by noting that the most tightly coupled meshes are ones which are physically close. To exploit this observation, for each mesh i, the submatrix of MZM^t corresponding to all meshes near mesh i is inverted directly. Then, the row of the inverted submatrix associated with mesh i becomes the i^{th} row of P. We refer to this preconditioner as a "local-inversion" preconditioner, because it is formed by inverting physically localized problems.

Another approach to preconditioning exploits the fact that MZM^t can be decomposed into $MRM^t + j\omega MLM^t$, where MRM^t is very sparse, and is therefore reasonably easily factored. This suggests using an "inverse-R" preconditioner, in which case the preconditioned matrix is $I + j\omega(MRM^t)^{-1}MLM^t$. That this matrix is easier to solve with GMRES can be seen by noting the following theorem.

Theorem 1 *The matrix $I + j\omega(MRM^t)^{-1}MLM^t$ is normal, and all of its eigenvalues have a real part $= 1$.*

The above result is easily proved by noting that $(MRM^t)^{-1}MLM^t$ is within a similarity transform of a symmetric matrix. The theorem also implies that the eigenvalues of the preconditioned matrix lie on a line in the complex plane, and it is well-known that GMRES converges more quickly when applied to normal matrices whose eigenvalues are restricted to a line [2].

To compare the relative merits of the above preconditioners, consider the packaging example in Fig. 3 and a resistive ground-plane problem with thirty external contacts (the external contacts fix potentials, and can impact convergence). For this experiment, the package was discretized into 1000 filaments, and the ground plane was discretized into a 33×33 grid of filaments. The solution error as a function of GMRES iteration is plotted in Fig. 5 for the packaging example, and in Fig. 6 for ground plane example. As the plots clearly show, preconditioning substantially accelerates convergence, but the packaging example converges more rapidly with the local-inversion preconditioner, and the ground-plane example converges more rapidly with the inverse-R preconditioner.

4 Conclusions and Acknowledgements

With a good preconditioner, the multipole-accelerated 3-D inductance extraction algorithm is extremely efficient, and can be used to accurately model skin and proximity effects. Which preconditioner is best is problem dependent however, and future work is on finding a preconditioner which is effective in all cases. The authors would like to thank Joel Phillips and Khalid Rahmat for their timely assistance.

References

[1] P. A. Brennan, N. Raver, and A. Ruehli. Three dimensional inductance computations with partial element equivalent circuits. *IBM Journal of Res. and Develop.*, 23(6):661–668, November 1979.

[2] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, July 1986.

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

[4] 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*, 10(11):1447–1459, November 1991.

[5] M. Kamon, M. J. Tsuk, and J. White. Fasthenry, a multipole-accelerated 3-d inductance extraction program. In *Proceedings of the ACM/IEEE Design Automation Conference*, Dallas, June 1993.