

Efficient Formulation and Model-Order Reduction for the Transient Simulation of Three-Dimensional VLSI Interconnect

Mike Chou and Jacob K. White, *Associate Member, IEEE*

Abstract—Accurately accounting for three-dimensional (3-D) geometry and distributed RC effects in on-chip interconnect is important for predicting crosstalk in memory cells, analog circuits, and regions of congested routing in digital circuits. In this paper we describe a multipole-accelerated, mixed surface-volume formulation, and a preconditioned model-order reduction algorithm for distributed RC, or electroquasistatic, simulation of 3-D integrated circuit interconnect. The difficulties arising from the ill conditioning inherent in the dynamic problem is effectively resolved by a combined surface-volume approach. Results are presented to demonstrate that the computational cost for extracting a complete reduced-order model is order N , where N is the number of surface unknowns. Finally, the multipole-accelerated code is used to investigate the accuracy of the one-dimensional diffusion equation for long RC lines.

Index Terms—Diffusion equation, integral equations, Krylov-subspace methods, model-order reduction, transient analysis.

I. INTRODUCTION

WHEN analyzing high-performance integrated circuit designs, it is well known that the single lumped resistor-capacitor model of interconnect is insufficiently accurate. It has been shown [1] that reasonably accurate electro-quasistatic, or transient interconnect, simulations could be performed by computing the time evolution of the electric field both inside and outside the conductors via a finite-difference discretization of Laplace's equation. More recently, a boundary-element approach [2] based on Green's theorem was proposed which performs the calculation using the same surface discretization used for ordinary capacitance extraction, thereby avoiding the large exterior domain mesh and computation. However, the latter approach generates dense matrix problems, which require $O(N^3)$ operations to solve directly, and at least $O(N^2)$ to solve iteratively, where N is the number of surface unknowns. Therefore it is necessary to accelerate such methods when solving large problems. The direct application of the $O(N)$ fast-multipole algorithm on the boundary-element formulation produces unacceptable results because the multipole errors are magnified by the ill conditioning in the linear system which results from the wide range of time constants in the dynamics.

Manuscript received August 12, 1996. This work was supported by the Defense Advanced Research Projects Agency under Contract B146103 and the Semiconductor Research Corporation under Contract SJ-558. This paper was recommended by Associate Editor D. Ling.

The authors are with the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: mtchou@rle-vlsi.mit.edu; white@rle-vlsi.mit.edu).

Publisher Item Identifier S 0278-0070(97)09361-5.

To overcome this difficulty, we derive a mixed surface-volume formulation and show how it prevents the magnification of the multipole error. In this formulation, the interior finite-difference method is used to solve Laplace's equation *inside* the conductors, and the boundary-element method is used to solve the *exterior* Laplace problem.

For three-dimensional (3-D) interconnect structures to be included along with the actual transistors in a coupled SPICE-level circuit simulation, it is necessary to construct low-order macromodels whose terminal behaviors essentially capture the complicated 3-D field interactions among the interconnect. Most model order reduction techniques, such as asymptotic waveform evaluation (AWE) [3] and the more recent Padé-via-Lanczos (PVL) [4] and Arnoldi [5] algorithms, have been successful because it is feasible to carry out an LU decomposition of the associated *sparse* system matrix, after which each solve can be performed cheaply. For problems involving large *dense* matrices, direct factorization is computationally intractable. Iterative methods can also be expensive if many solution iterations, or matrix-vector product computations are required for convergence, as is the case for ill-conditioned linear systems. We show how the surface-volume formulation can be modified slightly to allow effective *preconditioning* which produces rapid convergence in the iterative solution.

The outline of this paper is as follows. The surface-integral formulation of the transient-interconnect problem is briefly outlined in Section II. The phenomenon of ill conditioning is described in Section III. The surface-volume formulation is derived in Section IV, and the resulting error control is demonstrated in Section V. The guaranteed stable Arnoldi algorithm for model-order reduction is reviewed in Section VI. The modified surface-volume formulation and preconditioning techniques are presented in Section VII. Section VIII describes the method of images for including groundplanes. Computational results are presented in Section IX where we show that the cost associated with generating a q th-order model is order N and is less than that of performing q capacitance extractions. The conclusions and acknowledgments are given in Section X. Finally, in the Appendix, we compare the popular diffusion model of the distributed RC line to 3-D calculations.

II. THE SURFACE INTEGRAL FORMULATION

For the transient interconnect problem, the system is assumed to be in the electro-quasistatic (EQS) regime. The scalar potential $\psi(x, t)$ satisfies Laplace's equation in all of space

except on conductor surfaces where charge can accumulate [2]

$$\nabla^2 \psi(x, t) = 0, \quad x \notin S \quad (1)$$

where S is the union of all conductor surfaces. Since Laplace's equation (1) holds both inside and outside of the conductors, all charges in the system reside on the conductor surfaces S . Therefore, the potential ψ is related to the conductor surface charge density ρ_s through the superposition integral

$$\frac{\partial \psi(x, t)}{\partial t} = \int_S \frac{1}{4\pi\epsilon \|x - x'\|} \frac{\partial \rho_s(x', t)}{\partial t} da' \quad (2)$$

where the regions inside and outside the conductors are assumed to have uniform permittivity ϵ . Charge conservation [6] at the surface yields the continuity condition

$$\frac{\partial \rho_s(x)}{\partial t} = J^{\text{internal}}(x) - J^{\text{external}}(x) \quad (3)$$

where J^{internal} and J^{external} are the normal current densities taken just inside and just outside the conductor surface. Inside a conductor, the current obeys the constitutive relation

$$J^{\text{internal}}(x) = -\sigma \frac{\partial \psi}{\partial n}(x) \quad (4)$$

where σ is the conductivity and n is the outward normal to the surface S .

Combining (2)–(4) results in an integral formulation

$$-4\pi\tau \frac{\partial \psi(x, t)}{\partial t} = \int_S \frac{1}{\|x - x'\|} \frac{\partial \psi}{\partial n'}(x', t) da' + \frac{1}{\sigma} \int_S \frac{1}{\|x - x'\|} J^{\text{external}}(x', t) da' \quad (5)$$

where $\tau = \epsilon/\sigma$ is the dielectric relaxation time of the conductors, x is a point on a conductor surface, n' is the outward normal to the conductor surface, and $\|x - x'\|$ is the Euclidean distance between x and x' . Careful application of Green's theorem [2], [7] to the first integral on the right-hand side of (5) yields

$$-4\pi\tau \frac{\partial \psi(x, t)}{\partial t} = 2\pi\psi(x, t) + \int_S \psi(x', t) \frac{\partial}{\partial n'} \frac{1}{\|x - x'\|} da' + \frac{1}{\sigma} \int_S \frac{J^{\text{external}}(x', t)}{\|x - x'\|} da'. \quad (6)$$

Let S_{contact} be the subset of S which is in contact with external voltage sources, and let $S_{\text{free}} = S \setminus S_{\text{contact}}$ be the noncontact or free surfaces. Then $\psi(x, t)$ for $x \in S_{\text{contact}}$ is known *a priori*. Since there is no external current flow at noncontact surfaces, we also have *a priori* that $J^{\text{external}}(x, t) = 0$ for $x \in S_{\text{free}}$.

To numerically solve (6) for ψ at noncontact surfaces and for J^{external} at contact surfaces, the conductor surfaces are broken into N small tiles, or panels. It is then assumed that on each panel l , there is a constant potential ψ_l and a constant external supply current density J_l . A collocation scheme [8], in which (6) is enforced at the centroid of each panel, is used to generate a system of N equations. The result is a $N \times N$ dense linear system

$$-4\pi\tau \frac{d}{dt} \Psi(t) = (2\pi\mathbf{I} + \mathbf{D})\Psi(t) + \mathbf{P} \frac{1}{\sigma} \mathbf{J}^{\text{ext}}(t) \quad (7)$$

where $\Psi \in \mathbb{R}^N$, $\mathbf{J}^{\text{ext}} \in \mathbb{R}^N$ represent the discretized panel potentials and external supply current densities. The elements of the dense matrices $\mathbf{D} \in \mathbb{R}^{N \times N}$ and $\mathbf{P} \in \mathbb{R}^{N \times N}$ are

$$\mathbf{P}_{kl} = \frac{1}{a_l} \int_{\text{panel}_l} \frac{1}{\|x' - x_k\|} da' \quad (8)$$

$$\mathbf{D}_{kl} = \int_{\text{panel}_l} \frac{\partial}{\partial n'} \frac{1}{\|x' - x_k\|} da' \quad (9)$$

where x_k is the center of the k th panel, and a_l is the area of the l th panel. Mathematically, \mathbf{P}_{kl} is the potential at x_k due to a unit *charge* distributed uniformly over panel l . Similarly, \mathbf{D}_{kl} is the potential at x_k due to a unit *dipole* oriented along the normal to and distributed uniformly over panel l . The integrals in (8) and (9) are often referred to as *single-layer* and *double-layer* integrals [9], respectively.

Suppose N_c of the N surface panels are connected to voltage contacts whose potentials $\Psi_c \in \mathbb{R}^{N_c}$ are known but whose supply currents $\mathbf{J}_c^{\text{ext}} \in \mathbb{R}^{N_c}$ are unknown. It is then clear that (7) is an index-one differential-algebraic equation (DAE), solvable with backward-differencing formulas (BDF). In addition to the N_c elements of $\mathbf{J}_c^{\text{ext}}$, the unknowns also include the $N_f = (N - N_c)$ elements of $\Psi_f \in \mathbb{R}^{N_f}$, which correspond to the noncontact panel potentials. Discretization of (7) in time with the backward-Euler method yields the linear system

$$\mathbf{H} \begin{pmatrix} \Psi_f \\ \mathbf{J}_c^{\text{ext}} \end{pmatrix}_{t=(m+1)h} = \frac{4\pi\tau}{h} \Psi_{t=mh} - \left(\frac{4\pi\tau}{h} \mathbf{I} + 2\pi\mathbf{I} + \mathbf{D} \right) \cdot \begin{pmatrix} 0 \\ \Psi_c \end{pmatrix}_{t=(m+1)h} \quad (10)$$

where h is the timestep. The matrix, or linear operator, $\mathbf{H} \in \mathbb{R}^{N \times N}$, is defined by the transformation rule

$$\mathbf{H} \begin{pmatrix} \mathbf{v} \\ \mathbf{w} \end{pmatrix} = \left(\frac{4\pi\tau}{h} \mathbf{I} + 2\pi\mathbf{I} + \mathbf{D} \right) \begin{pmatrix} \mathbf{v} \\ 0 \end{pmatrix} + \mathbf{P} \begin{pmatrix} 0 \\ \mathbf{w} \end{pmatrix} \quad (11)$$

where $\mathbf{v} \in \mathbb{R}^{N_f}$ and $\mathbf{w} \in \mathbb{R}^{N_c}$.

Since \mathbf{H} is defined in terms of \mathbf{P} and \mathbf{D} , the unknowns can be interpreted as a distribution of monopoles and dipoles, with the panels associated with the elements of $\mathbf{J}_c^{\text{ext}}$ acting as uniform monopoles (single layers), and the panels associated with Ψ_f acting as uniform dipoles (double layers).

III. DIFFICULTIES WITH MULTIPOLE ACCELERATION

Consider using a Krylov-subspace based iterative algorithm, such as GMRES [10] to solve (10) at each timestep. The k th iteration of the GMRES algorithm requires computing the matrix-vector product $\mathbf{H}\mathbf{u}^k$, where \mathbf{u}^k is the k th GMRES search direction. Since \mathbf{H} is dense, computing $\mathbf{H}\mathbf{u}^k$ directly requires N^2 operations. However, forming $\mathbf{H}\mathbf{u}^k$ is equivalent to computing potentials at N points due to a distribution of N monopoles and dipoles. *Fast-multipole* algorithms [11]–[13] can be used to compute approximate values of the N potentials in βN operations where β is independent of N but dependent on the required accuracy.




Fig. 1. Single wire ($L = 4$) connected to voltage source at one end.

If (10) is solved by using a fast-multipole algorithm to approximate \mathbf{H} in (11), then

$$\tilde{\mathbf{H}} \begin{pmatrix} \tilde{\Psi}_f \\ \tilde{\mathbf{J}}_c^{\text{ext}} \end{pmatrix} = \mathbf{b} \quad (12)$$

where $\tilde{\mathbf{H}}$ is the multipole approximation to \mathbf{H} , \mathbf{b} is the right-hand-side of (10), and $\tilde{\Psi}_f, \tilde{\mathbf{J}}_c^{\text{ext}}$ are approximations to the true solution $\Psi_f, \mathbf{J}_c^{\text{ext}}$ in (10). The relative error in the computed potentials and currents is given by

$$\frac{\left\| \begin{pmatrix} \Psi_f \\ \mathbf{J}_c^{\text{ext}} \end{pmatrix} - \begin{pmatrix} \tilde{\Psi}_f \\ \tilde{\mathbf{J}}_c^{\text{ext}} \end{pmatrix} \right\|}{\left\| \begin{pmatrix} \Psi_f \\ \mathbf{J}_c^{\text{ext}} \end{pmatrix} \right\|} \leq \mathcal{K}(\mathbf{H}) \frac{\|\mathbf{H} - \tilde{\mathbf{H}}\|}{\|\mathbf{H}\|} \quad (13)$$

where $\|\mathbf{H} - \tilde{\mathbf{H}}\|$ is the multipole error and $\mathcal{K}(\mathbf{H})$ is the condition number [14] of \mathbf{H} . As is clear from (13), the error from the multipole algorithm is magnified by the condition number of \mathbf{H} . To see the impact of even mild ill conditioning in \mathbf{H} on multipole algorithm errors, consider the first model problem, a rectangular wire with dimensions $L : 1 : 1$, which is connected to a step voltage source at one end, shown in Fig. 1. The steady-state voltage at any point on the conductor surface is 1 V. Fig. 2 is a plot of the steady-state voltage at the opposite end of the wire (labeled $v1$) versus wire length computed using a multipole-accelerated algorithm.

For the algorithm used (second-order multipole expansions), the multipole approximation errors in the potential calculation is between 0.1–1%, but the steady-state error is much larger because of the magnification due to the ill conditioning in \mathbf{H} . As further evidence of this explanation, the condition number of \mathbf{H} is plotted as a function of wire length in Fig. 3.

For multiconductor systems, the condition number of \mathbf{H} grows as the spacing between conductors is reduced. Fig. 4 shows a simple two-conductor problem. Each conductor has voltage boundary conditions at one end. Fig. 5 shows that the condition number for the system increases as the spacing between the conductors is reduced. At very large separations, the two conductors are decoupled, and the condition number approaches that of the single-wire example in Fig. 1.

We comment here that while higher-order multipole expansions can be used (at a much greater computational expense) to improve the accuracy, it only serves to *delay* the onset of error magnification, and since the condition number is observed to grow *quadratically* with the length of the conductors, we shall pursue other means of resolving this difficulty.

IV. THE MIXED SURFACE-VOLUME FORMULATION

We derive here a mixed surface-volume formulation which can be multipole-accelerated without loss in solution accuracy, although it does not change the condition of the system matrix. Consider the interior Dirichlet-to-Neumann operator \mathcal{X} , defined by the linear map between the surface potential ψ and its normal derivative $(\partial/\partial n)\psi$, where the limit for $(\partial/\partial n)\psi$ is approached from the interior of the conductor surfaces

$$\mathcal{X}\psi(x) \equiv \frac{\partial\psi}{\partial n}(x), \quad x \in S. \quad (14)$$

This relation allows the surface-integral formulation (5) to be written as

$$-4\pi\tau \frac{\partial\psi(x, t)}{\partial t} = \int_S \frac{1}{\|x - x'\|} \mathcal{X}\psi(x', t) da' + \frac{1}{\sigma} \int_S \frac{1}{\|x - x'\|} J^{\text{external}}(x', t) da'. \quad (15)$$

We now discretize the conductor surfaces into N panels and assume uniform potentials and currents on each panel as described in Section II. The resulting matrix equation is

$$-4\pi\tau \frac{d}{dt} \Psi(t) = \mathbf{P} \left[\mathbf{X}\Psi(t) + \frac{1}{\sigma} \mathbf{J}^{\text{ext}}(t) \right] \quad (16)$$

where \mathbf{P} is as defined in (8). The matrix $\mathbf{X} \in \mathfrak{R}^{N \times N}$ approximates the continuous operator \mathcal{X} and is defined by

$$\mathbf{X}\Psi \equiv \Psi_n \quad (17)$$

where $\Psi_n \in \mathfrak{R}^N$ corresponds to $(\partial/\partial n)\psi$ at the N panels. Given Ψ at the surface nodes of a conductor, Laplace's equation can be solved in the *interior* domain with an interior finite-difference method to yield Ψ_n at each surface node. Hence, applying \mathbf{X} implies solving the interior problems.

As before, a fixed-timestep backward-Euler method is used to solve the DAE derived from (16). The resulting linear system is

$$\mathbf{A} \begin{pmatrix} \Psi_f \\ \mathbf{J}_c^{\text{ext}} \end{pmatrix}_{t=(m+1)h} = \frac{4\pi\tau}{h} \Psi_{t=mh} - \left(\frac{4\pi\tau}{h} \mathbf{I} + \mathbf{P}\mathbf{X} \right) \cdot \begin{pmatrix} 0 \\ \Psi_c \end{pmatrix}_{t=(m+1)h}. \quad (18)$$

The new operator \mathbf{A} is defined by the transformation rule

$$\mathbf{A} \begin{pmatrix} \mathbf{v} \\ \mathbf{w} \end{pmatrix} = \frac{4\pi\tau}{h} \begin{pmatrix} \mathbf{v} \\ 0 \end{pmatrix} + \mathbf{P} \left\{ \mathbf{X} \begin{pmatrix} \mathbf{v} \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ \mathbf{w} \end{pmatrix} \right\}. \quad (19)$$




Fig. 2. Steady-state voltage at v_1 versus wire length.




Fig. 3. Condition number grows with wire length.

The computation associated with applying \mathbf{X} can be performed efficiently. Since the interior Laplace problem is solved *independently* for each conductor, the action of the \mathbf{X} operator corresponds to solving a *block-diagonal* and *sparse* linear system. Thus the dominant cost of applying \mathbf{A} in (19) comes

from applying \mathbf{P} , which is a *dense* matrix operation since it couples every panel to all panels on all conductors. But as described in Section III, the application of \mathbf{P} to a vector can be multipole-accelerated. Therefore the combined surface-volume approach can be made very efficient.




Fig. 4. Parallel wires ($L = 8$), each with voltage contacts at one end.

The mixed surface-volume method provides an important guarantee on the solution accuracy. This is stated in the theorem below.

Theorem 1: If the steady-state solution of (15) is such that the surface potential on each conductor is a constant, and none of the conductors is floating, then the steady-state solution computed by the mixed surface-volume method is exact, regardless of multipole approximation error and discretization error.

Proof: Consider first the single conductor problem. From (16), the steady state solution satisfies

$$\frac{d}{dt} \Psi = 0 = \mathbf{P} \left(\mathbf{X} \Psi + \frac{1}{\sigma} \mathbf{J}^{\text{ext}} \right). \quad (20)$$

From the theory of fractional Sobolev spaces, it can be shown that the potential coefficient matrix \mathbf{P} is nonsingular given a sufficiently fine discretization [15]. It then follows that $\mathbf{X} \Psi + (1/\sigma) \mathbf{J}^{\text{ext}} = 0$ in the steady state. In the finite-difference implementation of \mathbf{X} , this is equivalent to a resistor network connected to external voltage sources [1]. Assuming that all voltage sources are at 1 V, the solution satisfies

$$\mathbf{X} \begin{pmatrix} \Psi_f \\ 1 \end{pmatrix} + \begin{pmatrix} 0 \\ \frac{1}{\sigma} \mathbf{J}_c^{\text{ext}} \end{pmatrix} = 0. \quad (21)$$

In the equivalent resistor network picture, N_c of the surface nodes are connected to unit-voltage sources while the remaining N_f surface nodes are left open-circuited. Network analysis immediately yields $\Psi_f = 1$ and $\mathbf{J}_c^{\text{ext}} = 0$, the exact steady-state solution. For many-conductor problems, the same result holds since each conductor is treated independently by the \mathbf{X} operator. ■

Since (6) and (15) are both derived from (5), the Green's theorem based and the surface-volume based formulations are equivalent in their integral equation form. If we define the integral operators $\mathcal{P}\mathcal{X}$ and \mathcal{D} as

$$\mathcal{P}\mathcal{X}\psi(x) \equiv \int_S \frac{1}{\|x - x'\|} \frac{\partial}{\partial n'} \psi(x') da' \quad (22)$$

$$\mathcal{D}\psi(x) \equiv \int_S \left(\frac{\partial}{\partial n'} \frac{1}{\|x - x'\|} \right) \psi(x') da' \quad (23)$$

then formally $\mathcal{P}\mathcal{X} = (2\pi I + \mathcal{D})$ by Green's theorem [7], where I is the identity operator. Thus it follows that in

the limit as the mesh becomes very fine, (i.e., $N \rightarrow \infty$), the discretized versions of these operators approach each other, $\mathbf{P}\mathbf{X} \approx (2\pi\mathbf{I} + \mathbf{D})$. Since $\psi(x) \equiv \text{constant}$ implies $\mathcal{X}\psi(x) = (\partial/\partial n)\psi(x) = 0$, both $\mathbf{P}\mathbf{X}$ and $(2\pi\mathbf{I} + \mathbf{D})$ are singular matrices with the vector $\{1, 1, \dots, 1\}$ in the null space. The surface-volume formulation essentially factors the matrix $(2\pi\mathbf{I} + \mathbf{D})$ into the product of a singular \mathbf{X} and a well-conditioned, nonsingular \mathbf{P} . When the action of \mathbf{P} is multipole accelerated in the mixed formulation, errors are introduced only in the capacitance matrix of the surface panels which does not alter the physical character of the system. This error appears only during the transient and will be shown experimentally to be small and independent of condition number. This is expected since approximations are made only on \mathbf{P} , the well-conditioned part. The null space of $\mathbf{P}\mathbf{X}$ is preserved. The same is not true for the Green's theorem based pure-boundary formulation, since multipole approximations are made on \mathbf{D} which alters the null space of $(2\pi\mathbf{I} + \mathbf{D})$.

V. COMPUTATION RESULTS

To show that both the pure boundary-element (BE) formulation and the mixed finite-difference/boundary-element (FD/BE) formulation produce similar results *without* multipole acceleration, we performed simulations on the single-wire conductor in Fig. 1 using the dimensions $L = 4$, $W = 1$, and $H = 1$. Here we introduce a discretization parameter m , which represents the number of sections into which each unit-length is divided. For example, $m = 3$ in Fig. 1. One end of the conductor is connected to a step voltage source, and the voltage waveform at the other end is shown in Fig. 6. For the coarse mesh $m = 3$, both methods produce small discretization errors. For the fine mesh $m = 7$, the two methods converge to the same waveform. This confirms the validity of the new mixed formulation.

Multipole acceleration is performed on both techniques for the double-wire example in Fig. 4, with actual discretization ($m = 3$) shown. At their near ends, one wire is connected to a step-voltage source, while the other is grounded. Simulated voltage waveforms at their far ends are shown in Fig. 7. For the mixed FD/BE formulation, the multipole-accelerated result produces the correct steady state, and is practically indistinguishable from the nonaccelerated explicit calculations. The multipole-accelerated pure BE technique is seen to produce obviously erroneous results, as reported in Section III. Experimentally, for the mixed surface-volume formulation, we find that second-order multipole acceleration always produces results matching those of the explicit calculations, independent of the condition number.

A fairly complex 3-D interconnect example is presented here to demonstrate that the multipole-accelerated surface-volume method is necessary for large problems. The GMRES [10] iterative method without preconditioning is used to solve the linear systems (10) and (18). Polysilicon resistivity of $\rho = 0.02 \Omega \cdot \text{cm}$ is assumed for all conductors, and oxide permittivity of $\epsilon_r = 3.2$ is assumed throughout space. All computations are performed on a 266 MHz DEC AXP3000/900 workstation with one gigabyte of physical memory.




Fig. 5. Condition number increases as wire separation decreases.

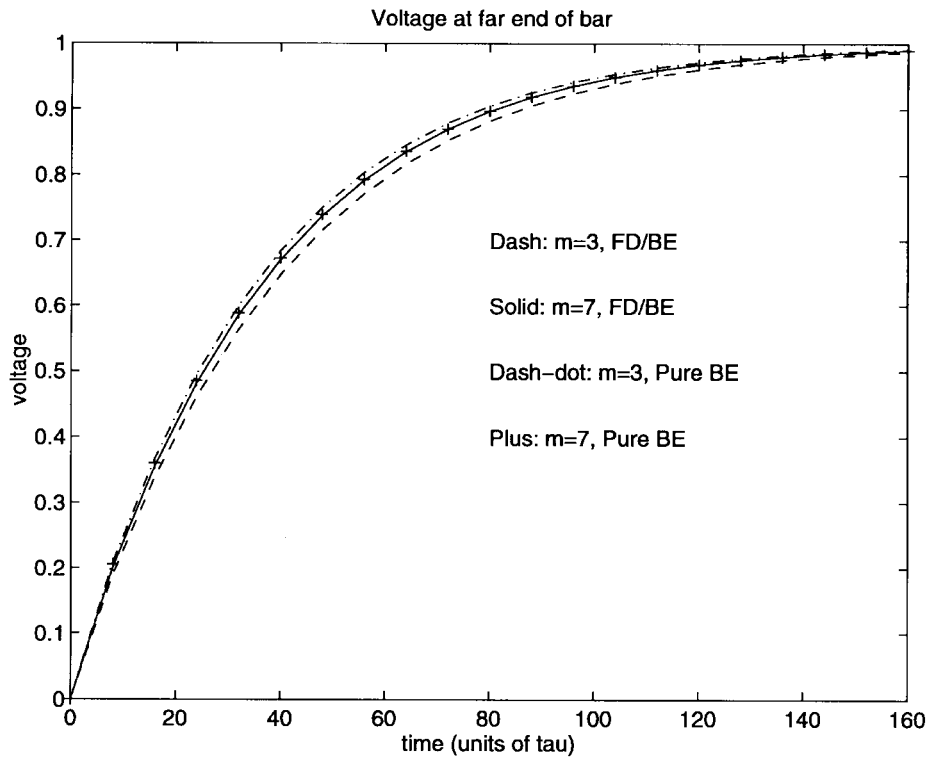


Fig. 6. Without acceleration, both techniques produce correct results.

Fig. 8 displays a model of a six-conductor SRAM cell. The groundplane is shown but not used in this example. Each conductor is connected at a port labeled 1-6. An additional port, labeled 7, is connected to conductor 3. Table I lists the number of surface unknowns for four successive refinements,

with Fig. 8 corresponding to $(m = 3)$. The pair of *L*-shaped conductors (1 and 2) are the clock lines, while the pair of Π -shaped conductors (5 and 6) are the data lines. A third pair of intertwined, interior conductors (3 and 4) make interconnections between transistors in the cell. Assume




Fig. 7. Multipole errors get magnified only in the pure BE method.




Fig. 8. SRAM cell ($m = 3$ mesh).

that ports 3 and 4 are grounded and that ports 5, 6, and 7 are floating during a particular fetch cycle. We simulate the cross-talk noise induced on these lines by a unit-step voltage source on both ports 1 and 2. Such spurious signals must be minimized at the design phase to ensure error-free operation.

Results of the time-domain backward-Euler simulation are shown in Fig. 9. The multipole-FDBE method is applied to four successively finer meshes ($m = 1, 2, 3, 4$), while the explicit-BE method, due to CPU time and memory limitations,

TABLE I
PROBLEM SIZE AND FD TIME FOR SRAM

m	1	2	3	4
panels	986	3,944	8,874	15,776
FD time	0.3%	1.0%	2.5%	6.7%

is used only for the coarsest mesh ($m = 1$). From the voltage waveforms $v5$ and $v7$, it is seen that the ($m = 1$) mesh results in significant discretization error. The finer meshes generate large numbers of unknowns which necessitates using the multipole-accelerated FD/BE method. CPU times for the multipole-FD/BE method are plotted in Fig. 10, which can be seen to exhibit linear, or $O(N)$, growth. CPU times for the explicit-BEM approach are extrapolated from the ($m = 1$) mesh computation, and grows as $O(N^2)$ since it is a dense-matrix method. For the ($m = 4$) mesh, with 15 776 panels, the multipole-accelerated method is seventeen times faster than the dense-matrix approach.

We make the additional note here that the CPU time consumed by the interior finite-difference computation as a percentage of the total CPU time grows with increasing mesh refinement, but remains small as shown in Table I. This confirms the earlier assertion that the cost of the BE calculation is dominant.

VI. THE ARNOLDI ALGORITHM FOR MODEL-ORDER REDUCTION

In order to fully evaluate the effects of interconnect on overall circuit performance, it is necessary to perform a *coupled* circuit-interconnect simulation at the SPICE level. It is impractical to incorporate the large, dense matrices associated




Fig. 9. Waveforms computed using various mesh refinements.




Fig. 10. Comparing CPU times.

with the 3-D interconnect directly into the circuit simulator. Instead, reduced-order models, which use small matrices to capture the current-voltage relations at the terminal ports of the interconnect, can be extracted from the full model and then used in the coupled simulation. Techniques such as asymptotic

AWE [3] and the PVL algorithm [4] have been used successfully for this purpose. In this section, we summarize previous work on the similar Arnoldi [16] algorithm, a numerically robust orthogonal-projection based scheme which generates guaranteed stable reduced-order models [17].

Consider the single-input-single-output (SISO), linear time-invariant system described by a system of first-order ordinary differential equations of the form

$$\begin{aligned}\dot{\mathbf{x}}(t) &= \mathbf{A}\mathbf{x}(t) + \mathbf{b}u(t) \\ y(t) &= \mathbf{c}^T \mathbf{x}(t)\end{aligned}\quad (24)$$

where the N -vector \mathbf{x} represents the circuit variables or the detailed internal voltages of the interconnect, and the $N \times N$ matrix \mathbf{A} represents the detailed interactions among internal elements, $\mathbf{b} \in \mathbb{R}^N$ is the excitation vector corresponding to the input terminal, and $\mathbf{c} \in \mathbb{R}^N$ is the observation vector corresponding to the output terminal. The scalar quantities $u(t)$ and $y(t)$ are the input and output terminal-port variables, through which the linear system “interfaces” with external circuitry. The state-space representation of (24) is

$$\begin{aligned}s\mathbf{X}(s) &= \mathbf{A}\mathbf{X}(s) + \mathbf{b}U(s) \\ Y(s) &= \mathbf{c}^T \mathbf{X}(s)\end{aligned}\quad (25)$$

where \mathbf{X} , U , and Y denote the Laplace transforms of \mathbf{x} , u , and y , respectively. The transfer function $F(s) \equiv Y(s)/U(s)$ can be written

$$F(s) = \mathbf{c}^T \cdot (\mathbf{I} - s\mathbf{A}^{-1}) \cdot \mathbf{p} = \sum_{k=1}^N \frac{\nu_k}{s - \lambda_k} \quad (26)$$

where $\mathbf{p} = -(\mathbf{A}^{-1}) \cdot \mathbf{b}$.

Since N can be of the order of tens of thousands, it is desirable to *reduce* the large and dense matrix \mathbf{A} or \mathbf{A}^{-1} in a manner that captures the low-frequency behavior of the transfer function. This is done by matching Taylor series terms at $s = 0$. It has been shown in [16] that an Arnoldi-based orthogonalization process can be used to construct an orthonormal basis for the Krylov subspace

$$\mathcal{K}_q(\mathbf{A}^{-1}, \mathbf{p}) = \text{span}\{\mathbf{p}, \mathbf{A}^{-1}\mathbf{p}, \mathbf{A}^{-2}\mathbf{p}, \dots, \mathbf{A}^{-(q-1)}\mathbf{p}\}. \quad (27)$$

After q steps, the Arnoldi algorithm returns a set of q orthonormal vectors, as the columns of the matrix $\mathbf{V}_q \in \mathbb{R}^{N \times q}$, where N is the size of \mathbf{A} , and typically $q \ll N$. The reduced-order transfer function can then be constructed as

$$\begin{aligned}\tilde{F}(s) &= \tilde{\mathbf{c}}^T \cdot (\mathbf{I} - s\mathbf{H}_q)^{-1} \cdot \tilde{\mathbf{p}} \\ \mathbf{H}_q &= \mathbf{V}_q^T (\mathbf{A}^{-1}) \mathbf{V}_q \\ \tilde{\mathbf{p}} &= \mathbf{V}_q^T \mathbf{p} = \|\mathbf{p}\| \mathbf{e}_1 \\ \tilde{\mathbf{c}}^T &= \mathbf{c}^T \mathbf{V}_q\end{aligned}\quad (28)$$

where \mathbf{H}_q is a $q \times q$ upper Hessenberg matrix. The transfer function $\tilde{F}(s)$ of the reduced q th-order system (28) has been shown in [16] to match $(q - 2)$ derivatives, or moments, of the exact transfer function in (26) at $s = 0$, the low-frequency limit. The triplet $[\mathbf{H}_q, \tilde{\mathbf{p}}, \tilde{\mathbf{c}}]$ is said to be the *reduced-order model* of the triplet $[\mathbf{A}^{-1}, \mathbf{p}, \mathbf{c}]$.

It is possible to extend the present work to the multiple-input-multiple-output (MIMO) case using block algorithms similar to those described in [5] and [18].

VII. PRECONDITIONED MODEL-ORDER REDUCTION

Both the AWE and PVL algorithms have been successfully applied to reduce circuit networks for the lumped-element model of the interconnect since the associated large, sparse matrices can be factored to solve for the low-frequency moments of the transfer function [19]. The difficulty with applying the AWE, PVL, or Arnoldi algorithm to reduce 3-D interconnect models is that the associated large, dense matrices are too expensive to store and factor. Matrix-implicit iterative solution can also be expensive since many matrix-vector product computations are required for ill-conditioned problems. We recall here that the matrix ill conditioning results from the wide range of time constants associated with typical interconnect, and is thus independent of the problem formulation. In Section VII-A, we show that the straightforward iterative solution converges slowly, and in Section VII-B, we reformulate the mixed surface-volume approach slightly and derive an effective preconditioner which allows for rapid convergence of the iterative solution. Section VIII describes how to include ideal groundplanes in the problem, and Section IX presents the computational results.

A. Application of Arnoldi

To simplify the notation in (16), let $\hat{\mathbf{P}} = (-1/4\pi\epsilon)\mathbf{P}$ and $\hat{\mathbf{D}} = (-1/4\pi\tau)\mathbf{P}\mathbf{X}$, which results in

$$\frac{d}{dt} \Psi(t) = \hat{\mathbf{D}}\Psi(t) + \hat{\mathbf{P}}\mathbf{J}^{\text{ext}}(t). \quad (29)$$

Since $\hat{\mathbf{D}}$ is singular, the steady-state voltage $\Psi(t)$ is not uniquely determined by the external current $\mathbf{J}^{\text{ext}}(t)$. Thus we recast (29) as a DAE system. This is done by using voltage sources instead of current sources and then computing the resulting n -port frequency-dependent admittance matrix, which is then well behaved near zero frequency. Recall that the N unknowns are the first N_f entries in the potential vector Ψ corresponding to the free potentials Ψ_f plus the N_c nonzero externally supplied currents $\mathbf{J}_c^{\text{ext}}$. The last N_c entries Ψ_c in Ψ are given *a priori* and correspond to external voltage sources. In frequency domain, the result is a system of equations

$$\begin{bmatrix} s\mathbf{I} - \hat{\mathbf{D}}_{ff} & -\hat{\mathbf{P}}_{fc} \\ -\hat{\mathbf{D}}_{cf} & -\hat{\mathbf{P}}_{cc} \end{bmatrix} \begin{bmatrix} \Psi_f(s) \\ \mathbf{J}_c^{\text{ext}}(s) \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{D}}_{fc}\Psi_c(s) \\ (\hat{\mathbf{D}}_{cc} - s\mathbf{I})\Psi_c(s) \end{bmatrix} \quad (30)$$

where $\hat{\mathbf{D}}_{ff} \in \mathbb{R}^{N_f \times N_f}$, $\hat{\mathbf{D}}_{fc} \in \mathbb{R}^{N_f \times N_c}$, $\hat{\mathbf{D}}_{cf} \in \mathbb{R}^{N_c \times N_f}$, and $\hat{\mathbf{D}}_{cc} \in \mathbb{R}^{N_c \times N_c}$ are partitions of the $\hat{\mathbf{D}}$ matrix

$$\hat{\mathbf{D}} = \begin{bmatrix} \hat{\mathbf{D}}_{ff} & \hat{\mathbf{D}}_{fc} \\ \hat{\mathbf{D}}_{cf} & \hat{\mathbf{D}}_{cc} \end{bmatrix}. \quad (31)$$

Similarly, $\hat{\mathbf{P}}_{ff} \in \mathbb{R}^{N_f \times N_f}$, $\hat{\mathbf{P}}_{fc} \in \mathbb{R}^{N_f \times N_c}$, $\hat{\mathbf{P}}_{cf} \in \mathbb{R}^{N_c \times N_f}$, and $\hat{\mathbf{P}}_{cc} \in \mathbb{R}^{N_c \times N_c}$ are partitions of the $\hat{\mathbf{P}}$ matrix. The subscript f denotes the free-floating panels and the subscript c denotes panels in contact with voltage sources. Since the contacts are typically at the ends of long conductors, the number of contact panels is typically much smaller than that of floating panels. Therefore, $\hat{\mathbf{P}}_{cc}$ is a small matrix and can be inexpensively inverted. Using $\hat{\mathbf{P}}_{cc}^{-1}$ allows (30) to be recast in the standard form for reduced-order modeling. Let $\mathbf{v} \in \mathbb{R}^{N_c}$, $\mathbf{w} \in \mathbb{R}^{N_c}$ be vectors of ones and zeros which selects the

TABLE II
MATRIX-VECTOR MULTIPLIES REQUIRED PER ORDER VERSUS LENGTH OF WIRE

Wire Aspect Ratio	16	32	64	128
Mat-Vecs (Direct apply)	24	37	60	102
Mat-Vecs (Preconditioned)	4	5	5	6

input voltage and output current panels, respectively. Then $\hat{\Psi}_c(s) = \mathbf{v}u(s)$ and $y(s) = \mathbf{w}^T \cdot \hat{\mathbf{J}}_c^{\text{ext}}(s)$, where $u(s)$ is the scalar voltage input and $y(s)$ is the scalar current output. After some amount of algebra, the admittance transfer function $g(s) = y(s)/u(s)$ is

$$g(s) = (k_0 + k_1s) + \mathbf{c}^T \cdot (s\mathbf{I} - \mathbf{A})^{-1} \cdot \mathbf{b} \quad (32)$$

where

$$\mathbf{A} = \hat{\mathbf{D}}_{ff} - \hat{\mathbf{P}}_{fc} \hat{\mathbf{P}}_{cc}^{-1} \hat{\mathbf{D}}_{cf} \quad (33)$$

$$\mathbf{b} = (\hat{\mathbf{D}}_{fc} - \hat{\mathbf{P}}_{fc} \hat{\mathbf{P}}_{cc}^{-1} \hat{\mathbf{D}}_{cc} + \mathbf{A} \hat{\mathbf{P}}_{fc} \hat{\mathbf{P}}_{cc}^{-1}) \cdot \mathbf{v}, \quad \mathbf{c}^T = -\mathbf{w}^T \cdot \hat{\mathbf{P}}_{cc}^{-1} \hat{\mathbf{D}}_{cf} \text{ and } k_0, k_1 \text{ are scalar constants.}$$

Rewriting the second term in (32) as $f(s) = \mathbf{c}^T \cdot (\mathbf{I} - s\mathbf{A}^{-1}) \cdot \mathbf{p}$, with $\mathbf{p} = -(\mathbf{A}^{-1}) \cdot \mathbf{b}$, we apply the Arnoldi method to reduce the triplet $[\mathbf{A}^{-1}, \mathbf{b}, \mathbf{p}]$ by matching low-frequency moments. For each additional order in the model, a new vector in the Krylov subspace $\mathcal{K}_q(\mathbf{A}^{-1}, \mathbf{p})$ in (27) is generated by applying GMRES to perform an iterative solution of the system

$$\mathbf{A} \cdot \mathbf{x} = \text{RHS} \quad (34)$$

using only multipole-accelerated matrix-vector multiplication as described in Section II.

As a numerical experiment, we directly apply the above Arnoldi algorithm to the simple interconnect in Fig. 1, a single wire, with the supply voltage as input variable $u(s)$ and the supply current as the output variable $y(s)$. The same calculation is performed for wires of varying lengths, keeping the other two dimensions fixed. Our numerical results, summarized in Table II, show that the number of iterations, or matrix-vector product calculations, required for GMRES convergence in solving (34) grows quickly as the wire length, or aspect ratio, is increased. This is caused by the system matrix \mathbf{A} becoming more ill-conditioned as the range of time constants, or eigenvalues, grows with the wire length. It is well-known that the rate of convergence for Krylov-subspace style algorithms deteriorates with growing matrix condition number [20].

B. Preconditioned Formulation

We derive here a slightly modified version of (16), which can be easily preconditioned to accelerate convergence of the iterative method used to compute the Krylov subspace vectors. In this formulation, we assume that the panels in contact with voltage sources store no charge, or equivalently, that the contact capacitances have been removed. This model may also be supported based on physical arguments: terminal ports of interconnects are not exposed surfaced when the connections to transistors or other circuitry have been made. The contact panels in practice exist *inside* conducting material where Laplace's equation holds, and hence cannot store charge.

We start from the original integral formulation of (5). By writing the surface integrals over S as a direct sum of integrals over the contact and noncontact surfaces S_{contact} and S_{free} , (5) becomes

$$\begin{aligned} -4\pi\tau \frac{\partial\psi(x, t)}{\partial t} &= \int_{S_{\text{free}}} \frac{1}{\|x - x'\|} \\ &\cdot \left[\frac{\partial\psi}{\partial n'}(x', t) + \frac{1}{\sigma} J^{\text{external}}(x', t) \right] da' \\ &+ \int_{S_{\text{contact}}} \frac{1}{\|x - x'\|} \\ &\cdot \left[\frac{\partial\psi}{\partial n'}(x', t) + \frac{1}{\sigma} J^{\text{external}}(x', t) \right] da'. \end{aligned} \quad (35)$$

The assumption that the charge density ρ_s is zero at contact surfaces S_{contact} , combined with the continuity condition (3) and the constitutive relation (4), implies

$$\frac{\partial\psi}{\partial n'}(x', t) + \frac{1}{\sigma} J^{\text{external}}(x', t) = 0, \quad x' \in S_{\text{contact}}. \quad (36)$$

Thus, the second surface integral in (35) vanishes. In addition, since there are no external supply currents at noncontact surfaces, $J^{\text{external}}(x', t) = 0 \forall x' \in S_{\text{free}}$. The unknown potentials on S_{free} then satisfy

$$\begin{aligned} -4\pi\tau \frac{\partial\psi(x, t)}{\partial t} &= \int_{S_{\text{free}}} \frac{1}{\|x - x'\|} \frac{\partial\psi}{\partial n'}(x', t) da' \\ &= - \int_{S_{\text{free}}} \frac{1}{\|x - x'\|} \frac{1}{\sigma} J^{\text{internal}}(x', t) da' \\ &\quad x \in S_{\text{free}} \end{aligned} \quad (37)$$

where (4) has been used in the second equality.

As before, we discretize S_f into N_f elements and S_c into N_c elements using the collocation scheme. The interior Dirichlet-to-Neumann operator defined in (17) can be rewritten as

$$\mathbf{X} \begin{bmatrix} \Psi_f \\ \Psi_c \end{bmatrix} \equiv \begin{bmatrix} \mathbf{X}_{ff} & \mathbf{X}_{fc} \\ \mathbf{X}_{cf} & \mathbf{X}_{cc} \end{bmatrix} \begin{bmatrix} \Psi_f \\ \Psi_c \end{bmatrix} = -\frac{1}{\sigma} \begin{bmatrix} \mathbf{J}_f^{\text{int}} \\ \mathbf{J}_c^{\text{int}} \end{bmatrix} \quad (38)$$

where $\mathbf{J}_f^{\text{int}} \in \mathfrak{R}^{N_f}$ and $\mathbf{J}_c^{\text{int}} \in \mathfrak{R}^{N_c}$ correspond to normal current densities *just inside* the N_f noncontact panels and the N_c contact panels, respectively. Discretization of (37) yields the $N_f \times N_f$ system

$$\frac{d}{dt} \Psi_f(t) = -\hat{\mathbf{P}}_{ff} \mathbf{J}_f^{\text{int}}(t) \quad (39)$$

where $\hat{\mathbf{P}}_{ff} \in \mathfrak{R}^{N_f \times N_f}$ has been defined previously. Combining (38) and (39) and again letting $\Psi_c(s) = \mathbf{v}u(s)$ and $y(s) = \mathbf{w}^T \cdot \mathbf{J}_c^{\text{int}}(s)$, we have the state-space form

$$s\Psi_f(s) = \mathbf{A}\Psi_f(s) + \mathbf{b}u(s) \quad (40)$$

$$y(s) = \mathbf{c}^T \cdot \Psi_f(s) + d \cdot u(s) \quad (41)$$

where

$$\mathbf{A} = \sigma \hat{\mathbf{P}}_{ff} \mathbf{X}_{ff} \equiv \left(\frac{1}{4\pi\tau} \right) \mathbf{P}_{ff} \mathbf{X}_{ff} \quad (42)$$

$\mathbf{b} = (1/4\pi\tau) \mathbf{P}_{ff} \mathbf{X}_{fc} \mathbf{v}$, $\mathbf{c}^T = \mathbf{w}^T \mathbf{X}_{cf}$, and $d = \mathbf{w}^T \mathbf{X}_{cc} \mathbf{v}$. The new expression for \mathbf{A} in (42) is to be compared




Fig. 11. Compare formulations with and without contact-port capacitances.




Fig. 12. Solution converges with discretization.

with that in (33). Notice that $\hat{\mathbf{D}}_{ff} = [1/(4\pi\tau)\mathbf{P}\mathbf{X}]_{ff} \neq (1/4\pi\tau)\mathbf{P}_{ff}\mathbf{X}_{ff}$. Time-domain solutions show that for reasonably long wires, the two formulations yield the same results since the capacitances associated with the contact ports are comparatively small. See Fig. 11 for the far-end voltage

waveforms computed for a length = 64 wire, in the absence of a groundplane, excited by a unit-step voltage source at the near-end.

Proceeding with the Arnoldi algorithm as in Section VII-A, the central task is to solve linear systems of the form




Fig. 13. Reduced-order models converge in frequency domain.

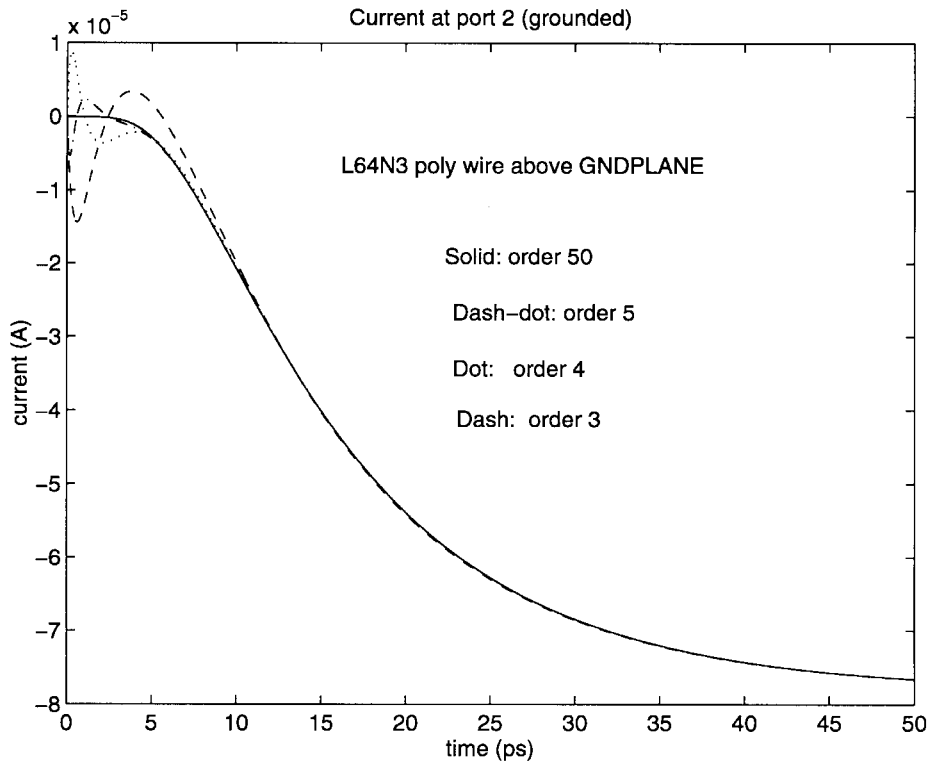


Fig. 14. Port 2 current versus time computed with reduced models.

$\mathbf{A} \cdot \mathbf{x} = \mathbf{RHS}$ for arbitrary right-hand sides. Since the operator \mathbf{A} has now a product form, it is easy to reduce its condition number by making the substitution

$$\mathbf{x} = \mathbf{X}_{ff}^{-1} \mathbf{\Pi}_{ff}^{-1} \cdot \mathbf{y} \tag{43}$$

where $\mathbf{\Pi}_{ff}^{-1}$ is a *sparse* matrix approximation to the inverse of the *dense* matrix \mathbf{P}_{ff} and is constructed by explicitly inverting local, overlapping blocks of \mathbf{P}_{ff} . For details on this computation, which fits naturally in the fast-multipole algorithm as demonstrated in the capacitance extraction program




Fig. 15. Port 2 voltage versus time computed with reduced models.




Fig. 16. Poly-to-poly coupling voltage noise in frequency domain.

FASTCAP, see [12]. The operator \mathbf{X}_{ff}^{-1} is the exact inverse of \mathbf{X}_{ff} , and its action is effected by solving the interior Laplace problem with mixed boundary conditions: Dirichlet on the contact panels ($\Psi_c = 0$) and Neumann on the free panels (\mathbf{J}_f^{int} arbitrary). As in the pure Dirichlet-to-Neumann problem \mathbf{X} , the mixed problem is solved by LU-factoring the associated sparse matrix generated from finite differences. The

free-panel potentials Ψ_f are computed along with \mathbf{J}_f^{int} as a by-product.

Using the preconditioner $\mathbf{X}_{ff}^{-1}\mathbf{\Pi}_{ff}^{-1}$, we apply GMRES to solve for \mathbf{y} in

$$\frac{1}{4\pi\tau} (\mathbf{P}_{ff}\mathbf{\Pi}_{ff}^{-1}) \cdot \mathbf{y} = \text{RHS} \tag{44}$$




Fig. 17. Poly-to-poly coupling voltage noise in time domain.




Fig. 18. Metal-to-poly coupling voltage noise in frequency domain.

and then compute the final solution \mathbf{x} by applying (43). This preconditioned Arnoldi algorithm is applied to the single-wire interconnect test case in Section VII-A. Table II displays the number of iterations, or matrix-vector product calculations, required for the iterative solution of (44). The rapid convergence shows that the condition number of the operator $\mathbf{P}_{ff}\mathbf{\Pi}_{ff}^{-1}$

is much smaller than that of \mathbf{A} and nearly independent of conductor length. The ill conditioning caused by the wide range of time constants has been removed by explicit solution of the interior problem \mathbf{X}_{ff}^{-1} , and the ill conditioning caused by the proximity of conductors is removed by the overlapping preconditioner $\mathbf{\Pi}_{ff}^{-1}$.

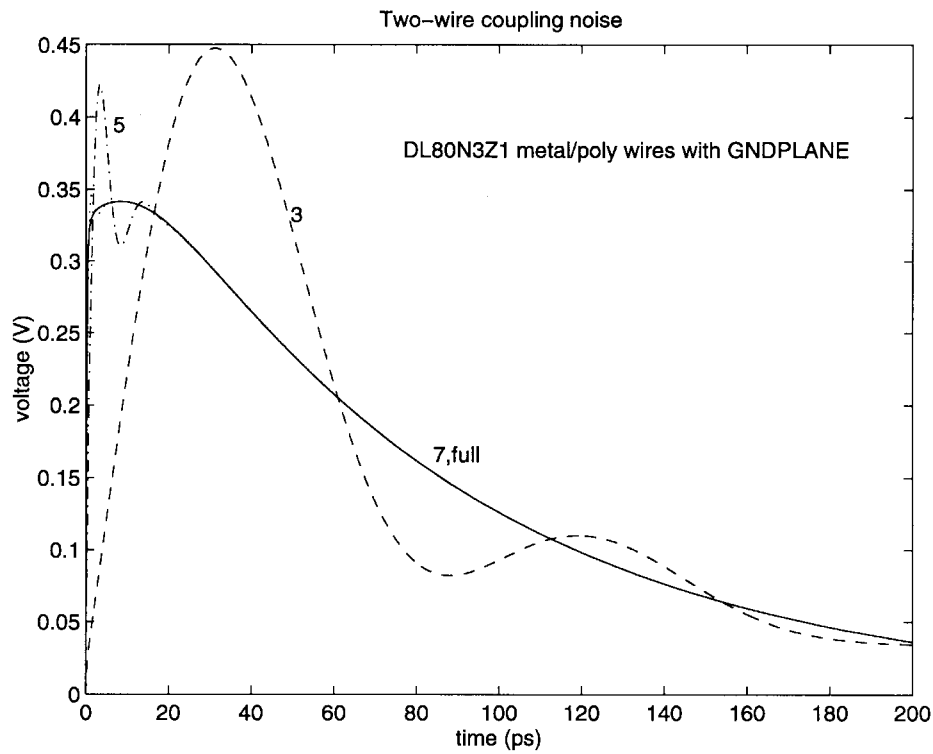


Fig. 19. Metal-to-poly coupling voltage noise in time domain.




Fig. 20. Convergence with discretization for SRAM.

We make a note here that the starting Arnoldi vector \mathbf{p} in (27) is computed by $\mathbf{p} = -\mathbf{X}_{ff}^{-1}\mathbf{X}_{fc} \cdot \mathbf{v}$ rather than an iterative solve involving \mathbf{A} . Hence, a q -order reduced model requires q GMRES iterative solutions rather than $(q + 1)$ solutions.

VIII. GROUND-PLANE IMPLEMENTATION

The potential variation of the grounded silicon substrate is typically of the order of tens of millivolts due to the many local, grounded body-plugs. Since this is small compared




Fig. 21. Reduced-order models for SRAM.

with the 3- or 5-V power supply, we will assume an *ideal groundplane* in this paper. To include the groundplane in the preconditioned formulation of the previous section, the only modification to make is the charge-to-potential operator $\hat{\mathbf{P}}_{ff}$. The interior Dirichlet-to-Neumann operator \mathbf{X} remains unaffected. Let the groundplane be approximated by a finite sheet, and assume it is explicitly discretized into N_g panels. Let $\Psi_g \in \mathbb{R}^{N_g}$ be the vector of ground panel potentials, and let $\mathbf{J}_g \in \mathbb{R}^{N_g}$ be the vector of corresponding panel currents. To include the groundplane, additional terms are introduced into (39)


$$\begin{bmatrix} \hat{\mathbf{P}}_{ff} & \hat{\mathbf{P}}_{fg} \\ \hat{\mathbf{P}}_{gf} & \hat{\mathbf{P}}_{gg} \end{bmatrix} \begin{bmatrix} \mathbf{J}_f^{\text{int}} \\ \mathbf{J}_g \end{bmatrix} = -\frac{d}{dt} \begin{bmatrix} \Psi_f \\ \Psi_g \end{bmatrix} \quad (45)$$

where $\hat{\mathbf{P}}_{fg} \in \mathbb{R}^{N_f \times N_g}$, $\hat{\mathbf{P}}_{gf} \in \mathbb{R}^{N_g \times N_f}$, $\hat{\mathbf{P}}_{gg} \in \mathbb{R}^{N_g \times N_g}$ describe capacitive interactions among conductor surfaces and the groundplane, and are similarly defined as $\hat{\mathbf{P}}_{ff}$. The condition $(d/dt)\Psi_g(t) = 0$ in the dynamic (45) implies that

$$\frac{d}{dt} \Psi_f = -\hat{\mathbf{P}}_{ff} \mathbf{J}_f \quad (46)$$

$$\hat{\mathbf{P}}_{ff} = \hat{\mathbf{P}}_{ff} - \hat{\mathbf{P}}_{fg} \hat{\mathbf{P}}_{gg}^{-1} \hat{\mathbf{P}}_{gf} \quad (47)$$

where $\hat{\mathbf{P}}_{ff} \in \mathbb{R}^{N_f \times N_f}$ is the new charge-to-potential operator in the presence of a groundplane. Since N_g may be large, it is impractical to factor $\hat{\mathbf{P}}_{gg}$, and since $\hat{\mathbf{P}}_{ff}$ is applied multiple times in an iterative solve, it is impractical to apply $\hat{\mathbf{P}}_{gg}^{-1}$ via an inner-loop iterative solve. Hence we will use the *method-of-images* [7] to apply the operator $\hat{\mathbf{P}}_{ff}$. Fictitious image charge panels are created by reflecting real charge panels across the groundplane and are always assigned the opposite charge. A similar procedure applies to the overlapping preconditioning


 Fig. 22. Three-level interconnect ($m = 1$ mesh).

operation. Since the $O(N)$ fast-multipole algorithm is used, the net cost is twice that of the problem without the groundplane.

It would also be possible to use precorrected-FFT methods with a modified Green's function to include the groundplane [21].

IX. MODEL-ORDER REDUCTION RESULTS

In this section, we present numerical results of our multipole-accelerated, preconditioned model-order reduction algorithm and demonstrate its accuracy and efficiency. Throughout, polysilicon conductivity and oxide permittivity will be assumed unless otherwise noted. The groundplane is also included in all following examples. Fig. 12 shows the frequency response for a 2 port, computed from the full model, for a rectangular conductor with aspect ratios $1 \mu\text{m} \times 1 \mu\text{m} \times 64 \mu\text{m}$, sitting one μm above the groundplane.

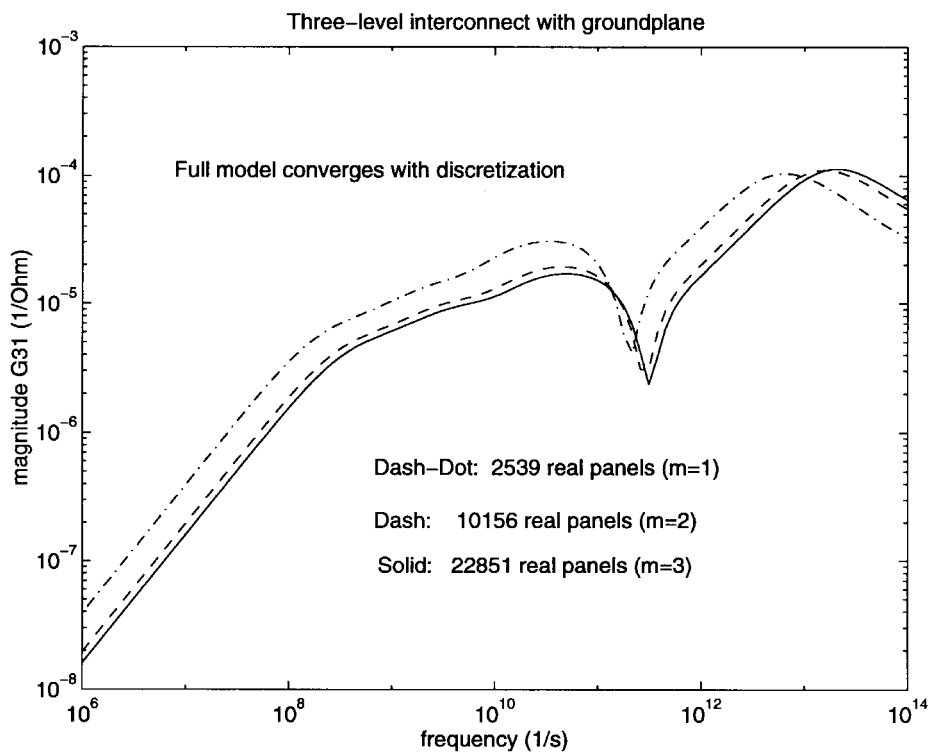


Fig. 23. Convergence with discretization for three-level interconnect.

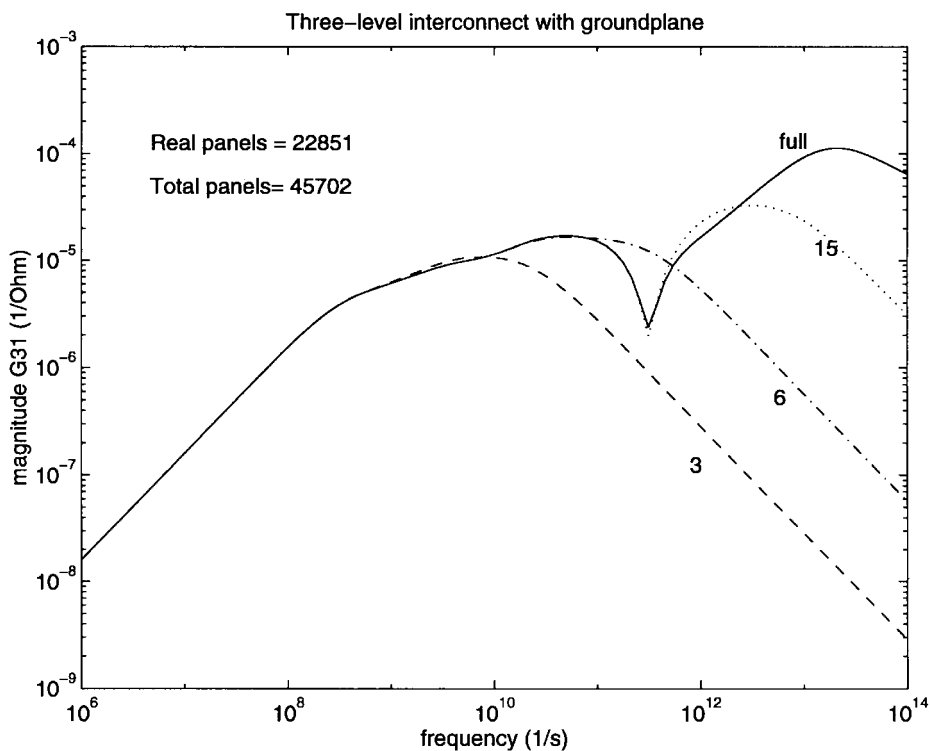


Fig. 24. Reduced-order models for three-level interconnect.

Two discretizations are used: the coarser one divides each unit square into nine equal panels, and the finer one divides each unit square into 16 panels. It is seen that up to a frequency of 10 THz (10^{13}), the results for the two discretizations are nearly identical. Henceforth, we shall use the coarser discretization.

Fig. 13 is a plot of the frequency response of the reduced-order models for the same conductor and shows that a twentieth-order model produces virtually identical results as the full model of order 2304. Time-domain data generated by the full-order models and the reduced-order models are




Fig. 25. CPU time grows linearly with problem size.

also given for comparison. Fig. 14 displays the short-circuit current in port 2 (held at ground), and Fig. 15 displays the open-circuit voltage at port 2 for various reduced models; the excitation in both is a unit-step voltage source at port 1. We see from Figs. 13 and 14 that third-order models are accurate enough if there are no signals in the system faster than 10–30 ps. Fig. 15 shows that the third-order model captures most of the essential features of the true response, while the first-order model, which is equivalent to a single-lumped RC model, fails miserably.

Next, we perform two-conductor coupling experiments using the same configuration as in Fig. 4, with the driven conductor connected to a voltage source and the “victim” conductor grounded at the near ends. We are interested in the voltage noise v_2 . Both wires have dimensions $1 \mu\text{m} \times 1 \mu\text{m} \times 80 \mu\text{m}$ and sit $1 \mu\text{m}$ above the groundplane. Fig. 16 shows the magnitude of v_2 in the frequency domain generated from several reduced models. It is seen that a fifth-order model is necessary to capture the full model up to 10 GHz. Fig. 17 shows the time-domain response v_2 to a unit-step voltage source. The fifth-order model is nearly indistinguishable from the full model. Similar experiments were performed with the driven polysilicon line replaced by aluminum and the victim line material unchanged. Results are shown in Figs. 18 and 19. The metal line introduces a much smaller timescale due to its high conductivity and, as a result, the coupling noise remains significant up to a much higher frequency 10 THz. The time response also shows a much faster risetime. For purposes of SPICE-level simulation in which the excitation is bandwidth-limited to, for example, below 10 GHz, a fifth-order model is sufficient.

TABLE III
TOTAL (REAL+IMAGE) PANEL COUNT

m	1	2	3	4
SRAM	1,952	7,808	17,568	31,232
3-level	5,078	20,312	45,702	NA

Next, we apply our algorithm to two large interconnect examples. The first is the six-conductor SRAM structure shown in Fig. 8. The structure is treated as a six-port problem, with the excitation ports labeled 1–6 in the figure. The simulation is now run with the groundplane with its approximate position shown in the figure. (Refer to Section IV for the discretization scheme and labeling.) Total panel counts, including real and image panels, are shown in Table III. Fig. 20 shows the frequency response of the conductance G_{61} computed using various discretizations. It is seen that up to 10^{13} Hz, the results are nearly identical for the mesh refinements $m = 3$ and $m = 4$; refer to Section V for the definition of the mesh parameter m . The coarser meshes $m = 1$ and $m = 2$ may be used for quick estimates. The model-order reduction results for $m = 4$ is plotted in Fig. 21, which shows that a sixth-order model is necessary to capture the first *knee* in the frequency response at ~ 100 GHz. We make a note here that the straight-line section of Fig. 21 corresponds to the low-frequency approximation

$$\mathbf{Y} = j\omega\mathbf{C} \tag{48}$$

where \mathbf{Y} , \mathbf{C} are the admittance and capacitance matrices, respectively. Resistance plays no part in the interconnect conductance until the higher-frequency components are excited.




Fig. 26. Memory use grows linearly with problem size.




Fig. 27. RC ladder circuit.

A three-level interconnect structure with a coarse discretization ($m = 1$) is shown in Fig. 22, in which each unit square is $1 \mu\text{m} \times 1 \mu\text{m}$ and the groundplane is $1 \mu\text{m}$ below the bottom-level interconnect. Polysilicon is used for the bottom level and aluminum for the top two. Each interconnect layer is excited at a single port, as shown in the figure. In this example, we compute the first column of the admittance matrix by connecting port 1 to a voltage source and grounding ports 2 and 3. Three discretizations were used ($m = 1, 2, 3$), and the total panel count is shown in Table III.

Fig. 23 shows the convergence with discretization of the frequency response for G_{31} , and Fig. 24 shows results of reduced-order modeling. Although the frequency dependence is complicated, a third-order model is accurate up to 1 GHz, and a sixth-order model accurate up to 10 GHz.

To demonstrate that the entire multipole-accelerated preconditioned model-order reduction algorithm has order N complexity, we plot CPU time and memory used versus the total number of panels (real and image) in computing a single-input-multiple-output (SIMO), sixth-order model for the SRAM and the three-level interconnect examples. See Figs. 25 and 26.

Since a SIMO reduced-order model corresponds to one column of the frequency-dependent admittance matrix, we compare this cost to that of computing one column of the capacitance matrix using FASTCAP [12], a multipole-accelerated




Fig. 28. Global capacitive coupling model.

TABLE IV
RATIO OF REDUCED-MODEL TO CAPACITANCE-EXTRACTION CPU TIMES

	m=1	m=2	m=3	m=4
3-level	1.6	1.6	2.4	NA
SRAM	1.8	2.2	3.0	3.1

capacitance extraction program. Table IV displays the ratio of the CPU times. While a sixth-order model essentially solves the capacitance problem six times, the actual CPU time overhead is seen to be only a factor of two to three. This is because the significant set-up time associated with the multipole algorithm, common to both procedures, is better amortized in the reduced-order model computation. Memory requirements are nearly identical in both cases.

X. CONCLUSIONS

We have shown that while the boundary-element formulation of the transient interconnect problem eliminates the need for exterior volume meshing, it produces large errors when multipole-accelerated due to the somewhat poor conditioning of the problem. The ill conditioning is inherent to problems with a large range of time constants, or natural frequencies. This magnified error is eliminated in




Fig. 29. Voltage attenuation: diffusion model versus 3-D.




Fig. 30. Diffusion model less accurate for large Z.

the alternative mixed surface-volume formulation, in which the ill-conditioned interior Laplace problem is separated from the well-conditioned capacitance problem and solved explicitly at a small additional cost.

To construct reduced-order models by matching Taylor series terms of the transfer function at $s = 0$, iterative solutions of a linear system must be performed repeatedly. A large num-




Fig. 31. Coupled RC ladders.




Fig. 32. Coupling noise ($Z = 1$): diffusion model versus 3-D.




Fig. 33. Diffusion model less accurate for large Z .

ber of iterations are required for ill-conditioned problems, such as those involving long wires. By reformulating the surface-volume approach slightly, we found natural preconditioners which produce rapid convergence in the iterative solve. We presented results which demonstrate that the cost of computing

a fixed-order reduced model is order N , independent of condition number, and is only several times that of a multipole-accelerated capacitance extraction.

Finally, we used our multipole-accelerated code to investigate the accuracy of the one-dimensional diffusion equation

for long RC lines. Our simulations show that the diffusion equation is accurate up to relatively high frequencies, unless the line is some distance from the groundplane.

APPENDIX COMPARING DIFFUSION AND 3-D MODELS

For analyzing 2-D interconnect problems, such as a single long wire or a collection of parallel wires, the *diffusion equation* is often used in the electro-quasistatic approximation, or RC regime. In the context of interconnect analysis, the basic single-conductor diffusion equation can be written as

$$\mathcal{R}\mathcal{C} \frac{\partial}{\partial t} \Phi(x, t) = \frac{\partial^2}{\partial x^2} \Phi(x, t) \quad (49)$$

where \mathcal{R} , \mathcal{C} are the resistance and capacitance, respectively, per unit length, and $\Phi(x, t)$ is the electric potential along the wire as a function of position and time. Equation (49) can be easily derived by taking the continuous limit of the discrete RC ladder circuit shown in Fig. 27. Similarly, numerical solutions which are accurate up to a given excitation frequency can be obtained from a circuit solution of the discrete RC ladder network if a large enough number of sections or lumps are used.

The diffusion model differs from the full 3-D model in several ways. First, the diffusion equation assumes capacitive coupling only between each node and the groundplane and not among the nodes themselves, whereas the 3-D picture models capacitive coupling among all panels as well as the groundplane. Fig. 28 displays the capacitive interaction between one particular node and all other nodes. Second, the diffusion picture models only current flow parallel to the wire, whereas the 3-D picture models current flow in all three directions in the conductor produced by possible potential differences in the transverse directions. Results from the two models approach each other as the conductor approaches the groundplane and as the excitation frequency is lowered, since the former effectively reduces the relative strength of panel-to-panel interactions and the latter makes the conductor potential more uniform in both transverse and longitudinal directions.

We present results from numerical experiments which compare the diffusion and the 3-D models. A 3-D capacitance extraction is performed on a single, rectangular wire over a groundplane using the mesh in Fig. 1. A long enough wire is used to ensure that the capacitance per unit length is within one percent of the long-wire limit. This capacitance value, along with the wire resistance computed from its cross-sectional area, is used in the diffusion model. For all experiments in this section, the conductor dimensions are $80 \mu\text{m} \times 1 \mu\text{m} \times 1 \mu\text{m}$, with a distance above the groundplane at $1 \mu\text{m}$ ($Z = 1$) or $10 \mu\text{m}$ ($Z = 10$). First, we perform the single-wire experiment in which the near-end of the conductor is excited by a voltage source and the resulting far-end voltage is measured as a function of frequency. The setup is similar to that in Fig. 1. The results from the diffusion and 3-D model are shown in Fig. 29. A close-up view is shown in Fig. 30. It is seen that up to a frequency of 100 GHz (10^{11} Hz), the diffusion results give a fair approximation to the 3-D results. Also, the

approximation becomes worse as the distance between the wire and groundplane is increased.

Next, we perform the two-conductor coupling experiment in a set up similar to that of Fig. 4. The parameters used for the diffusion model are extracted from the two-conductor capacitance matrix, computed with the groundplane included, and the coupled diffusion equation is solved numerically with the coupled RC ladders shown in Fig. 31. The noise voltage v_2 is plotted as a function of excitation frequency for the case $Z = 1$ in Fig. 32, and a magnified view is shown in Fig. 33 for both the $Z = 1$ and $Z = 10$ cases. The same observations can be made here as in the single-wire experiments. The low-frequency, straight-line section in the figures correspond to the capacitive limit described by (48) where the wire resistance plays no role.

We conclude from the above experiments that the diffusion model and the 3-D model yield similar results when the conductors are in close proximity to the groundplane, in which case, the relative importance of global capacitive coupling is minimized.

ACKNOWLEDGMENT

The authors would like to thank T. Korsmeyer for making his FASTLAP program available. The multipole-accelerated algorithms in the code were mostly based on FASTLAP. They would also like to acknowledge D. Ling of IBM-Yorktown for many insightful discussions and for suggesting the comparison to the diffusion equation.

REFERENCES

- [1] S. Kumashiro, R. Rohrer, and A. Strojwas, "A new efficient method for the transient simulation of 3-D interconnect structures," in *Proc. Int. Electron Devices Meeting*, Dec. 1990, pp. 193-196.
- [2] D. Ling, S. Kim, and J. White, "A boundary-element approach to transient simulation of 3-D integrated circuit interconnect," in *Proc. 29th Design Autom. Conf.*, June 1992, pp. 93-98.
- [3] L. T. Pillage and R. A. Rohrer, "Asymptotic waveform evaluation for timing analysis," *IEEE Trans. Computer-Aided Design*, vol. 9, pp. 352-366, Apr. 1990.
- [4] P. Feldmann and R. W. Freund, "Efficient linear circuit analysis by Padé approximation via the Lanczos process," in *Proc. Euro-DAC*, Sept. 1994, pp. 170-175.
- [5] L. M. Silveira, M. Kamon, and J. K. White, "Efficient reduced-order modeling of frequency-dependent coupling inductances associated with 3-D interconnect structures," in *Proc. 32nd ACM/IEEE Design Autom. Conf.*, San Francisco, CA, June 1995, pp. 376-380.
- [6] H. Haus and J. Melcher, *Electromagnetic Fields and Energy*. Englewood Cliffs, NJ: Prentice-Hall, 1989.
- [7] J. D. Jackson, *Classical Electrodynamics*, 2nd ed. New York: Wiley, 1975.
- [8] R. F. Harrington, *Field Computation by Moment Methods*. New York: Macmillan, 1968.
- [9] R. Kress, *Linear Integral Equations*. New York: Springer-Verlag, 1989.
- [10] Y. Saad and M. H. Schultz, "GMRES: A generalized minimum residual algorithm for solving nonsymmetric linear systems," *SIAM J. Sci. Stat. Comput.*, vol. 7, pp. 856-869, 1986.
- [11] L. Greengard, *The Rapid Evaluation of Potential Fields in Particle Systems*. Cambridge, MA: MIT Press, 1988.
- [12] K. Nabors and J. White, "Fastcap: A multipole accelerated 3-D capacitance extraction program," *IEEE Trans. Computer-Aided Design*, pp. 1447-1459, Nov. 1991.
- [13] *FastLap: Version 1.0*, T. Korsmeyer, K. Nabors, and J. White, Computational Hydrodynamics Facility and Research Laboratory of Electronics, MIT, Cambridge, MA, Apr. 1993.
- [14] G. Dahlquist and A. Björck, *Numerical Methods*. Englewood Cliffs, NJ: Prentice-Hall, 1974.

- [15] A. H. Schatz, V. Thomée, and W. L. Wendland, *Mathematical Theory of Finite and Boundary Element Methods*. Berlin: Birkhäuser-Verlag, 1990.
- [16] L. M. Silveira, M. Kamon, and J. K. White, "Efficient reduced-order modeling of frequency-dependent coupling inductances associated with 3-D interconnect structures," in *Proc. Eur. Design Test Conf.*, Paris, France, Mar. 1995, pp. 534–538.
- [17] K. J. Kerns, I. L. Wemple, and A. T. Yang, "Stable and efficient reduction of substrate model networks using congruence transforms," in *Int. Conf. Computer Aided-Design*, San Jose, CA, Nov. 1995, pp. 207–214.
- [18] P. Feldmann and R. W. Freund, "Reduced-order modeling of large linear subcircuits via a block Lanczos algorithm," in *32nd ACM/IEEE Design Autom. Conf.*, San Francisco, CA, June 1995, pp. 474–479.
- [19] C. L. Ratzlaff, N. Gopal, and L. T. Pillage, "RICE: Rapid interconnect circuit evaluator," in *28th ACM/IEEE Design Autom. Conf.*, June 1991, pp. 555–560.
- [20] J. Stoer and R. Bulirsch, *Introduction to Numerical Analysis*, 2nd ed. New York: Springer-Verlag, 1993.
- [21] J. R. Phillips and J. K. White, "Efficient capacitance extraction of 3-D structures using generalized precorrected FFT methods," in *Proc. IEEE 3rd Topical Meeting on Electrical Performance Electron. Packaging*, Nov. 1994, pp. 253–256.



Mike Chou received the B.S. degree in physics from the California Institute of Technology, Pasadena, CA, and the M.S. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, in 1991 and 1994, respectively. Currently, he is working towards the Ph.D. degree in electrical engineering at the Massachusetts Institute of Technology.

His research interests include computational electromagnetics, multigrid and wavelet methods, integral equation methods, stochastic systems, and computational finance.



Jacob K. White (A'88) received the B.S. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, and the S.M. and Ph.D. degrees in electrical engineering and computer science from the University of California, Berkeley.

From 1985 to 1987, he was with the IBM T. J. Watson Research Center, and from 1987 to 1989, he was the Analog Devices Career Development Assistant Professor at the Massachusetts Institute of Technology. He is currently a Professor at the Massachusetts Institute of Technology where his research interests are in serial and parallel numerical algorithms for problems in circuit, interconnect, device, and microelectromechanical system design.

Dr. White was a 1988 Presidential Young Investigator, and an Associate Editor of the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS from 1992 to 1996.