

Algorithms for Coupled Transient Simulation of Circuits and Complicated 3-D Packaging

Luís Miguel Silveira, *Student Member, IEEE*, Mattan Kamon, and Jacob White, *Member IEEE*

Abstract—Techniques are described for coupled simulation of complicated 3-D interconnect and nonlinear transistor drivers and receivers. The approach is based on combining: multipole-accelerated method-of-moments techniques for extracting frequency-dependent inductances and resistances for the interconnect; a sectioning method for fitting the frequency-domain data with a rational function; a balanced-realization approach to reducing the order of the rational function in a guaranteed stable manner; and an implementation of fast recursive convolution to incorporate the rational function in SPICE3. Results are presented to demonstrate some of the frequency-dependent effects in a packaging analysis problem.

Index Terms—Interconnect, packaging, simulation, multipole, balanced realization, recursive convolution.

I. INTRODUCTION

THE DENSE 3-D PACKAGING now commonly used in compact electronic systems may produce magnetic interactions which interfere with system performance. Such effects are difficult to simulate because they occur only as a result of an interaction between the field distribution in a complicated geometry of conductors, and the circuitry connected to those conductors. Effective simulation techniques which combine interconnect and circuitry have been developed based on simplified physical models and Padé style methods [1], but in this paper we examine an approach more tuned to packaging problems. In particular, multipole accelerated algorithms are used to efficiently compute frequency-dependent coupling resistances and inductances of the complicated 3-D packaging. Then, a section-by-section plus balancing approach is used to approximate the frequency-dependent elements with a rational function. Finally, the rational function is converted to a SPICE-compatible circuit model and combined with the connecting circuitry to perform the coupled simulation.

II. INDUCTANCE EXTRACTION

The frequency dependent resistance and inductance matrices describing the terminal behavior of a set of conductors can be rapidly computed with the multipole-accelerated mesh-

Manuscript received February, 1994; revised July 14, 1994. This work was supported by the Defense Advanced Research Projects Agency contract N00 014-91-J-1 698, the National Science Foundation contract (MIP-8 858 764 A02), an NSF Graduate Research Fellowship, the Portuguese "Junta Nacional de Investigação Científica e Tecnológica" under project "Cisência" and grants from Digital Equipment Corporation and IBM.

The authors are with the Research Laboratory of Electronics, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

IEEE Log Number 9407816.

formulation approach as implemented in FastHenry [2], [3]. To describe the approach, consider that each conductor is approximated as piecewise-straight sections. The volume of each straight section is then discretized into a collection of parallel thin filaments through which current is assumed to flow uniformly. The interconnection of these current filaments can be represented with a planar graph, where the n nodes in the graph are associated with connection points between conductor segments, and the b branches in the graph represent the current filaments into which each conductor segment is discretized (Fig. 1).

A. Mesh Formulation

To derive a system of equations for the filament currents we start by assuming that the applied currents and voltages are sinusoidal, and that the system is in sinusoidal steady-state. Following the partial inductance approach in [4], [5], the branch current phasors can be related to branch voltage phasors (hereafter, phasors will be assumed and not restated) by

$$\mathbf{Z}\mathbf{I}_b = \mathbf{V}_b, \quad (1)$$

where $\mathbf{V}_b, \mathbf{I}_b \in C^b$, b is the number of branches (number of current filaments), and $\mathbf{Z} \in C^{b \times b}$ is the complex impedance matrix given by

$$\mathbf{Z} = \mathbf{R} + j\omega\mathbf{L}, \quad (2)$$

where ω is excitation frequency. The entries of the diagonal matrix $\mathbf{R} \in R^{b \times b}$ represent the dc resistance of each current filament, and $\mathbf{L} \in R^{b \times b}$ is the dense matrix of partial inductances [6].

Kirchhoff's voltage law, which implies that the sum of branch voltages around each mesh (a mesh is any loop of branches in the graph which does not enclose any other branches) in the network is represented by

$$\mathbf{M}\mathbf{V}_b = \mathbf{V}_s, \quad (3)$$

where \mathbf{V}_b is the vector of voltages across each branch, $\mathbf{V}_s \in R^m$ is the mostly zero vector of source branch voltages, and $\mathbf{M} \in R^{m \times b}$ is the mesh matrix. The mesh currents, that is the currents around each mesh loop, satisfy

$$\mathbf{M}^t \mathbf{I}_m = \mathbf{I}_b, \quad (4)$$

where the superscript t denotes matrix transpose, and $\mathbf{I}_m \in R^m$ is the vector of mesh currents. Note that one of the entries in the mesh current vector will be identically equal




Fig. 1. One conductor. (a) Piecewise-straight sections. (b) Discretized into filaments. (c) Modeled as a circuit.

to the source branch current, shown as I_{m_source} in Fig. 1. Combining (4) with (3) and (1) yields

$$\mathbf{M}\mathbf{Z}\mathbf{M}^t\mathbf{I}_m = \mathbf{V}_s. \quad (5)$$

The complex admittance matrix which describes the terminal behavior of the conductor system, denoted $\mathbf{Y}_r = \mathbf{Z}_r^{-1}$, can be derived from (5) by noting that

$$\tilde{\mathbf{I}}_s = \mathbf{Y}_r \tilde{\mathbf{V}}_s, \quad (6)$$

where $\tilde{\mathbf{I}}_s$ and $\tilde{\mathbf{V}}_s$ are the vectors of terminal source currents and voltages. Therefore, to compute the i^{th} column of \mathbf{Y}_r , solve (5) with a \mathbf{V}_s whose only nonzero entry corresponds to $\tilde{\mathbf{I}}_{s,i}$, and then extract the entries of \mathbf{I}_m associated with the source branches.

B. Multipole Acceleration

To solve (5) by Gaussian Elimination would require order m^3 operations. To improve the situation, FastHenry uses a preconditioned GMRES iterative algorithm [7]. In general, the cost for each iteration of an iterative algorithm applied to solving (5) is order m^2 . This follows from the fact that each iteration requires computing the dense matrix-vector product, $(\mathbf{M}\mathbf{Z}\mathbf{M}^t)\mathbf{I}_m$. However, it is possible to approximately compute $\mathbf{M}\mathbf{Z}\mathbf{M}^t\mathbf{I}_m^k$ in order b operations using a hierarchical multipole algorithm [8]. Such algorithms also avoid explicitly forming $\mathbf{M}\mathbf{Z}\mathbf{M}^t$, and so reduce the memory required to order b .

To show how a multipole algorithm can be applied to computing $\mathbf{M}\mathbf{Z}\mathbf{M}^t\mathbf{I}_m^k$, consider expanding the matrix-vector product using (2),

$$\mathbf{M}\mathbf{Z}\mathbf{M}^t\mathbf{I}_m^k = \mathbf{M}\mathbf{R}\mathbf{M}^t\mathbf{I}_m^k + j\omega\mathbf{M}\mathbf{L}\mathbf{M}^t\mathbf{I}_m^k. \quad (7)$$

Except for $\mathbf{L}\mathbf{M}^t\mathbf{I}_m^k$, all of the products in (7) involve sparse matrices and can be computed in order m operations. To

compute $\mathbf{L}\mathbf{M}^t\mathbf{I}_m^k$, or equivalently $\mathbf{L}\mathbf{I}_b$, it is shown in [3] that

$$(\mathbf{L}\mathbf{I}_b)_i = \frac{1}{a_i} \int_{\text{fil}_i} \mathbf{A}(\mathbf{X}_i) \cdot \mathbf{l}_i(\mathbf{X}_i) d^3x_i, \quad (8)$$

where $\mathbf{X}_i \in R^3$ is the position in filament i , $\mathbf{A}(\mathbf{X}_i)$ is the vector potential at \mathbf{X}_i , $\mathbf{l}_i \in R^3$ is the unit vector in the direction of current flow in the filament, a_i is the cross sectional area, and $\text{fil } i$ represents the volume of filament i . Furthermore,

$$A_k(\mathbf{X}) = \frac{\mu_0}{4\pi} \sum_j \frac{I_{b_j}}{a_j} \int_{\text{filament}_j} \frac{(I_j(\mathbf{X}_j))_k}{|\mathbf{X} - \mathbf{X}_j|} d^3x_j, \quad (9)$$

where A_k is the k th component of the vector potential. Thus by viewing $(I_{b_j}/a_j)I_j(\mathbf{X}_j)_k$ as a ‘charge’ then computing each component of \mathbf{A} to eventually give $\mathbf{L}\mathbf{I}_b$ involves the evaluation of electrostatic potential along b filaments due to b filament charges for three separate sets of filament charges. It is the evaluation of these electrostatic potentials which can be accelerated to order b operations with the hierarchical multipole algorithm [8].

To see roughly what the multipole algorithm exploits to achieve its efficiency, consider the two configurations given in Figs. 2 and 3, depicted in 2-D for simplicity. 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. 2 can be computed in many fewer operations using multipole expansions, which exploit the fact that $r \gg R$ (defined in Fig. 2). That is, the details of the distribution of the charges in the inner circle of radius R in Fig. 2 do not strongly affect 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




Fig. 2. Evaluation point potentials are approximated with a multiple expansion.




Fig. 3. Evaluation point potentials are approximated with a local expansion.

points in the inner circle of Fig. 3 in many fewer than $n_1 * n_2$ operations using *local expansions*, which again exploit the fact that $r \gg R$ (as in Fig. 3). In this second case, what can be ignored are the details of the evaluation point distribution. This brief description of the hierarchical multipole algorithm is only intended to make clear that the algorithm's efficiency stems from coalescing charges and evaluation points using multipole and local expansions.

III. RATIONAL FUNCTION APPROXIMATION

The most commonly used approaches to fitting rational functions to frequency domain data are the Padé or moment-matching methods. These methods compute the coefficients of a rational function by matching that approximation to the value of the system function and its derivatives around $s = 0$ or $s = \infty$. Here, we describe a different approach to the approximation problem which is more suitable when matching tabulated data like that produced by FastHenry.

B. Sectioned ℓ_2 Minimization

One approach to generating a rational function which best matches a frequency response $F(s)$ specified at a set of frequencies $\{s_1, s_2, \dots, s_m\}$, is to set up and solve, as accurately as possible, the following set of equations:

$$\mathbf{H}(s_j) = \mathbf{F}(s_j) \quad j = 1, 2, \dots, m \quad (10)$$

where

$$\mathbf{H}(s) = \frac{\mathbf{U}(s)}{\mathbf{V}(s)} = \frac{u_q s^q + \dots + u_1 s + u_0}{v_p s^p + \dots + v_1 s + 1} \quad (11)$$

is the low-order approximation.

Typically, the system in (10) will be over-determined as the number of frequency points will exceed the number of unknown coefficients in the approximation (11), that is if $m > p + q + 1$. In that case there will be, in general, no exact solution, and the best that we can expect is that the approximation error be minimal in some sense. For instance we can force the 2-norm of the error to be minimized, that is, make sure that the coefficients of the polynomials $\mathbf{U}(s)$ and $\mathbf{V}(s)$ are chosen such that

$$\|\mathbf{H}(s) - \mathbf{F}(s)\|_2 = \left\| \frac{\mathbf{U}(s)}{\mathbf{V}(s)} - \mathbf{F}(s) \right\|_2 \quad (12)$$

is minimized for all $s \in \{s_1, s_2, \dots, s_m\}$. However, this is a nonlinear optimization problem whose solution is difficult to compute. Instead, the problem can be made linear by weighting the 2-norm by $\mathbf{V}(s)$. Then, the minimization problem becomes

$$\min_{\mathbf{U}, \mathbf{V}} \|\mathbf{U}(s_j) - \mathbf{V}(s_j)\mathbf{F}(s_j)\|_2 \quad j = 1, \dots, m. \quad (13)$$

Minimizing this weighted 2-norm does not guarantee that the resulting rational function will be accurate at any particular frequency, and this is unacceptable for use in circuit simulation. In order to insure that the steady-state will be computed exactly, the ℓ_2 minimization must be constrained as follows:

$$\begin{cases} \frac{\mathbf{U}(0)}{\mathbf{V}(0)} = \mathbf{F}(0) \\ \min_{\mathbf{U}, \mathbf{V}} \|\mathbf{U}(s_j) - \mathbf{V}(s_j)\mathbf{F}(s_j)\|_2 \quad j = 1, \dots, m \\ \lim_{s \rightarrow \infty} \frac{\mathbf{U}(s)}{\mathbf{V}(s)} = \lim_{s \rightarrow \infty} \mathbf{F}(s) \end{cases} \quad (14)$$

This constrained ℓ_2 -minimization however presents some difficulties, especially in the case when the natural frequencies of the problem are spread out over a wide range. In that situation, the minimization can become ill-conditioned, and also the weighted ℓ_2 minimization improperly focuses too much attention on the high frequencies. To avoid both these problems, it is possible to perform local approximations in a repeated fashion. Initially, the frequency range of interest, $\Omega = [w_{\min}, w_{\max}]$, is partitioned into small sections, $\Omega_1, \Omega_2, \dots, \Omega_M$, such that $\Omega = \bigcup_{i=1}^M \Omega_i$, where each Ω_i is a decade or two long. Then, starting with the lowest frequency range Ω_1 , with frequency values $\mathbf{F}(w_{11}), \mathbf{F}(w_{12}), \dots, \mathbf{F}(w_{1m_1})$, a constrained ℓ_2 minimization is performed and a local approximant is computed. Once the first local approximation, $\tilde{\mathbf{H}}_1(s)$, is obtained in the form of a collection of poles and their corresponding residues, it is examined and the stable poles are retained while the unstable ones are discarded, leaving us with a *forced stable* approximation, $\mathbf{H}_1(s)$. Next, the second section Ω_2 , the values $\mathbf{F}(w_{21}), \mathbf{F}(w_{22}), \dots, \mathbf{F}(w_{2m_2})$ are computed and fit again using the constrained weighted ℓ_2 minimization. Note that since the previous fit at the lower frequencies has captured the low frequency dynamics, $\mathbf{F}(s) - \mathbf{H}_1(s)$ will contain primarily the higher-frequency error information. This results in a new approximation for $\mathbf{F}(s) - \mathbf{H}_1(s)$, $\mathbf{H}_2(s)$ and therefore $\mathbf{F}(s) \approx \mathbf{H}_1(s) + \mathbf{H}_2(s)$. The procedure is repeated until data in the last frequency section, Ω_M , is approximated.

B. Model Order Reduction

The section-by-section approach just described is reliable in that it obtains a stable collection of pole-residue pairs which form an accurate approximation to $F(s)$. Unfortunately since $H(s)$ is represented as a sum of local approximations the approach introduces redundancies resulting in many more poles than necessary. However it is possible to further reduce the order of the approximation using robust model order reduction techniques.

In particular, the high order model can be converted to a state-space form as in

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{Ax} + \mathbf{Bu}, & \mathbf{x} &\in \mathbf{R}^n, \mathbf{u} \in \mathbf{R}^m \\ \mathbf{y} &= \mathbf{Cx}, & \mathbf{y} &\in \mathbf{R}^p \end{aligned} \quad (15)$$

whose transfer function is $G(s) = C(s\mathbf{I} - \mathbf{A})^{-1}\mathbf{B}$. Since we assume that (15) corresponds to passive interconnect, the transfer function is stable, i.e., all the poles (or equivalently all the eigenvalues of the matrix \mathbf{A}) are in the left half plane. It is well known [9] that for a given transfer function, the choice of the triplet $[\mathbf{A}, \mathbf{B}, \mathbf{C}]$ is not unique. Indeed, a linear coordinate transformation \mathbf{T} in the state space modifies the triplet $[\mathbf{A}, \mathbf{B}, \mathbf{C}]$ to $[\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]$ without modifying the transfer function. For the specific purpose of extracting *stable* reduced-order models from the state-space representation, it is desirable that the new triplet $[\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]$ be in a form that allows such an extraction using simple state truncation. In a seminal paper in the field of system theory, Moore has shown [10] that such a triplet exists and called it a “balanced realization” of the transfer function $G(s)$. The balancing transformation \mathbf{T} can be computed explicitly for any triplet $[\mathbf{A}, \mathbf{B}, \mathbf{C}]$. The numerical cost of such a computation is that of solving two matrix Lyapunov equations, one Cholesky factorization, and one symmetric eigenvalue problem. The resulting triplet $[\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}]$ has the remarkable property that simple reordering and truncation of the state vector $\tilde{\mathbf{x}}$ with the corresponding reordering of the system matrices *necessarily* produces *stable* reduced-order models at any desirable order. Let k be this order, and let $[\tilde{\mathbf{A}}_k, \tilde{\mathbf{B}}_k, \tilde{\mathbf{C}}_k]$ be the reduced-order model with a transfer function $G_k(s)$. It can then be shown [10], [11] that the impulse response of the error transfer function $E_k(s) = G(s) - G_k(s)$ has a norm that monotonically decreases to zero as k is increased to n , the order of the original model.

Judging the validity of the reduced-order model depends not only on meeting the error criterion mentioned above but also on meeting the goals of the circuit simulation task for which this reduced model is used. Typically, in circuit simulations, it is essential that the reduced model match the original transfer function at $s = 0$ so that the steady-state behavior of both the reduced and full models is identical. To achieve this, we apply a least-squares/collocation technique to match the reduced-order model with the full model at zero frequency.

IV. RESULTS

In this section we describe an example that demonstrates the value of using the reduced order models with the frequency dependent data acquired with FastHenry. The example is an investigation of crosstalk between a small set of package pins




Fig. 4. Seven pins of a cerquad pin package.




Fig. 5. General configuration for the connection between received and driver chips. All the circuit elements inside the same chip share that chip's power and ground.

connecting on-chip drivers to off-chip receivers. The frequency dependent resistance and inductance data for the package pins computed with FastHenry is approximated with low order rational functions, and then the time domain responses are computed with recursive convolution [12] implemented in SPICE3 [13].

Consider the crosstalk between seven adjacent pins of a 68-pin cerquad package as shown in Fig. 4. Assume the five middle lines carry output signals from the chip and the two outer pins carry power and ground. The signals are driven and received with CMOS inverters. The drivers are capable of driving a large current to compensate for the impedance of the package pins. The inductance of the pins is computed with FastHenry and the capacitance is assumed to be 8 pF. The interconnect from the end of pin to the receiver is modeled with a capacitance of 5 pF. The overall configuration is illustrated in Fig. 5 and a more detailed view for a single pin is given in Fig. 6. A 0.1 μF decoupling capacitor is connected between the driver's power and ground to minimize supply fluctuations.




Fig. 6. Detailed view of the connection between driving and receiving chips, showing the power and ground connections. Decoupling capacities between the power and ground lines are also shown. Pin capacitance and receiver interconnect capacitance are also modeled as small capacitive loads.

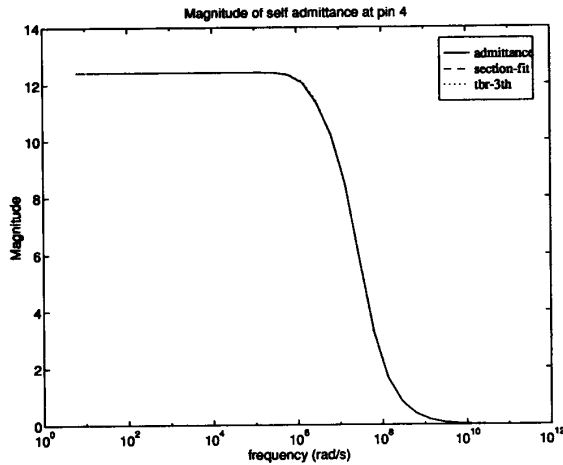


Fig. 7. Magnitude of the self-admittance term at pin 4. Shown in the plot are data points computed with FastHenry, the 12th order section-by-section approximant and the 3rd order reduced model computed by truncating the balanced realization. The error in both approximations is less than 0.5%.

To compute the resistance and inductance matrix at each frequency with FastHenry, the pins were discretized into five filaments along their height and nine along their length. This allows accurate modeling of changes in resistance and inductance due to skin and proximity effects. Matrices were generated for the frequency range 1 MHz to 10 MHz, with three points per decade.

The frequency dependence of each element in the admittance matrix is fit with a rational approximation using the algorithms described in Section III. First, the section-by-section approach is used to obtain approximations which have orders in the range of 12 to 24. Following the section-by-section algorithm a realization is determined and balanced. We have found that truncated models of 3rd order are sufficiently accurate to provide approximation with less than 5% error. The following two figures demonstrate this fact. Fig. 7 shows the magnitude of the self-admittance term at pin 4. Shown in the plot are data points computed with FastHenry, the 12th order section-by-section approximant and the 3rd order reduced model computed by truncating the balanced realization. As can be seen on the plot, the three curves match each other very closely.

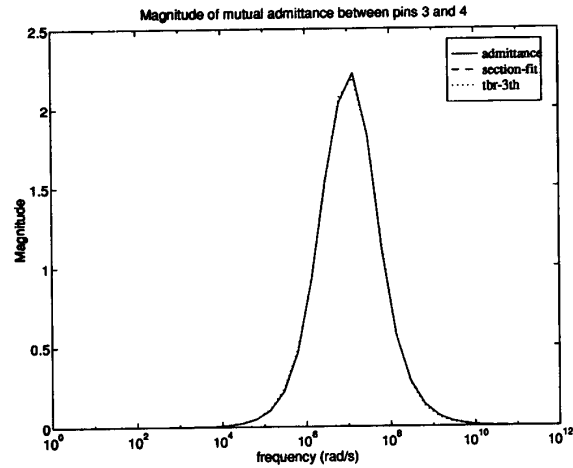


Fig. 8. Magnitude of the mutual admittance term between pins 3 and 4. Shown in the plot are data points computed with FastHenry, the 20th order section-by-section approximant and the 3rd order reduced model. The error in both approximations is less than 1%.

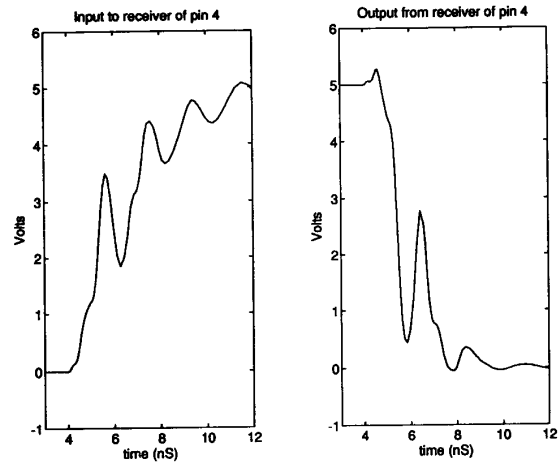


Fig. 9. Pin 4's receiver when four adjacent pins switch 1 ns after pin 4.

Fig. 8 shows the magnitude of the mutual admittance term between pins 3 and 4. Again, shown in the plot are data points computed with FastHenry, the 20th order section-by-section approximant and the 3rd order reduced model. As on the previous plot, the three curves match each other almost perfectly.

As a sample time domain simulation, imagine that at time $t_0 = 4$ ns the signal on pin 4 of Fig. 5 is to switch from high to low and pins 2, 3, 5, and 6 are to switch from low to high but that due to delay on chip, pins 2, 3, 5, and 6 switch at $t_1 = 5$ ns. In this case, significant current will suddenly pass through the delayed pins, 2, 3, 5 and 6 while pin 4 is in transition. Due to crosstalk, this large transient of current has significant effects on the input of the receiver on pin 4, as shown in Fig. 9. Note that the input does not rise monotonically. Fig. 9 also shows that the bump in the waveform is carried through to the output of receiver, as a large glitch.




Fig. 10. Pin 4's receiver using frequency dependent data (full), constant value from low frequency (low), and constant value from high frequency (high).




Fig. 12. Pin 4's receiver with ground pin 7 and signal line 5 swapped but no current through ground pin.




Fig. 11. Pin 4's receiver with ground pin 7 and signal line 5 swapped.

As a demonstration of the importance of frequency dependent analysis, consider the same scenario but with constant inductance and resistance matrices. That is, each pin will be modeled as a single resistor in series with an inductor whose values do not change with frequency. Fig. 10 is a comparison of the waveforms from Fig. 9 to waveforms resulting from simulations using constant resistance and inductance values corresponding to the high or low frequency limit. Note that for the receiver input waveforms, the large voltage bump swings by approximately 0.5 V more for the full frequency-dependent case. While this is small on the input, this is a very sensitive region for the receiver and doubles the size of the output glitch.

Now consider changing the design by swapping the ground pin, pin 7, with signal pin 5. Now the ground pin sits between signal lines and adds greater separation between pin 4 and the signals which are now on lines 6 and 7. As might be expected, the crosstalk is significantly reduced and the voltage bump does not exceed 1.5 V as shown in Fig. 11.




Fig. 13. Total current traveling toward drivers on the four late signal lines (signal), and the current traveling back along ground pin (ground).

Next, consider adding many extra ground and power pins outside of seven pins of this example. This would effectively eliminate the inductance and resistance effects of the power and ground pins. For the above example, this could be modeled by letting V_{dd} and ground for both the drivers and the receivers be 5 V and 0 V, respectively, and then no current will flow along power pin 1 or ground pin 5. In this case, the crosstalk appears to have worsened as shown in Fig. 12. This effect can be explained by realizing that current through the ground pin flows in the opposite direction to the current in the delayed pins, as shown in Fig. 13. Then since the ground path is now adjacent to pin 4, its current counteracts the inductive effects on pin 4 from the other pins' current.

V. CONCLUSION AND ACKNOWLEDGMENTS


In this paper techniques are described for coupled simulation of complicated 3-D interconnect and nonlinear transistor drivers and receivers. The approach is based on combining

multipole-accelerated method-of-moments techniques for extracting frequency-dependent inductances and resistances for the interconnect, and a sectioning plus balancing method for fitting the frequency-domain data with a rational function. Results are presented to demonstrate both the effectiveness of the algorithms and some of the frequency-dependent effects in a particular packaging problem.

The authors gratefully acknowledge the help of Ignacio McQuirk, Michael Tsuk, and Robert Armstrong in finding relevant examples.

REFERENCES

- [1] V. Raghavan, J. E. Bracken, and R. A. Rohrer, "AWESpice: a general tool for the accurate and efficient simulation of interconnect problems," in *29th ACM/IEEE Design Automation Conf.*, 1992, pp. 87-92.
- [2] M. Kamon, M. Tsuk, C. Smithisler, and J. White, "Efficient techniques for inductance extraction of complex 3-D geometries," in *Proc. Int. Conf. on Comp. Aided Design*, Nov. 1992.
- [3] M. Kamon, M. J. Tsuk, and J. White, "FastHenry, a multipole-accelerated 3-D inductance extraction program," in *Proc. ACM/IEEE Design Automation Conf.*, June 1993.
- [4] W. T. Weeks, L. L. Wu, M. F. McAllister, and A. Singh, "Resistive and inductive skin effect in rectangular conductors," *IBM J. Res. Develop.*, vol. 23, pp. 652-660, Nov. 1979.
- [5] A. E. Ruehli, "Inductance calculations in a complex integrated circuit environment," *IBM J. Res. Develop.*, vol. 16, pp. 470-481, Sept. 1972.
- [6] P. A. Brennan, N. Raver, and A. Ruehli, "Three dimensional inductance computations with partial element equivalent circuits," *IBM J. Res. Develop.*, vol. 23, pp. 661-668, Nov. 1979.
- [7] Y. Saad and M. H. Schultz, "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM J. Scientific and Statistical Computing*, vol. 7, pp. 856-869, July 1986.
- [8] L. Greengard, *The Rapid Evaluation of Potential Fields in Particle Systems*. Cambridge, Massachusetts: M.I.T. Press, 1988.
- [9] T. Kailath, *Linear Systems*. Information and System Science Series, 1st ed. Englewood Cliffs, New Jersey: Prentice-Hall, 1980.
- [10] B. Moore, "Principal component analysis in linear systems: controllability, observability, and model reduction," *IEEE Trans. Automat. Contr.*, vol. AC-26, pp. 17-32, Feb. 1981.
- [11] K. Glover, "All optimal Hankel-norm approximations of linear multi-variable systems and their l^∞ error bounds," *Int. J. Contr.*, vol. 39, pp. 1115-1193, June 1984.
- [12] S. Lin and E. S. Kuh, "Transient simulation of lossy interconnects based on the recursive convolution formulation," *IEEE Trans. Circuits Syst.*, vol. 39, pp. 879-892, Nov. 1992.
- [13] T. L. Quarles, "The SPICE3 implementation guide," *Tech. Rep. ERL M89/44*, Electronics Research Laboratory Report, University of California at Berkeley, Berkeley, California, Apr. 1989.



Luis Miguel Silveira (S'85) was born in Lisbon, Portugal. He received the Engineer's (summa cum laude) and Master's degrees in electrical and computer engineering in 1986 and 1989, from Instituto Superior Técnico at the Technical University of Lisbon, and the M.S., E.E. and Ph.D. degrees in 1990, 1991 and 1994 from the Massachusetts Institute of Technology, Cambridge, MA. His doctoral thesis was entitled "Model Order Reduction Techniques for Circuit Simulation".

From 1992 to 1994, he was supported by an IBM doctoral fellowship. Since 1986, Mr. Silveira has been on leave from Instituto Superior Técnico at the Technical University of Lisbon, whose faculty he will join in the spring of 1995. His research interests are in various aspects of computer-aided design of integrated circuits with emphasis on parallel computer algorithms and the theoretical and practical issues concerning numerical simulation methods for circuit design problems.

Dr. Silveira is a member of Sigma Xi.



Mattan Kamon received the B.S. degree in engineering science from the Pennsylvania State University in 1991, the M.A. degree in mathematics also from the Pennsylvania State University in 1991, and the S.M. degree in electrical engineering and computer science from the Massachusetts Institute of Technology in 1994. He is presently working toward the Ph.D. degree at the Massachusetts Institute of Technology.



Jacob White (S'80-M'83) received the B.S. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge. He received the S.M. and Ph.D. degrees from the University of California, Berkeley, both in electrical engineering and computer science.

He worked at the IBM T. J. Watson research center from 1985 to 1987, was the Analog Devices Career Development Assistant Professor at the Massachusetts Institute of Technology from 1987 to 1989, and was a 1988 Presidential Young Investigator. Jacob White is currently an associate Professor at M.I.T., and his research interests are in serial and parallel numerical algorithms for problems in circuit, interconnect, device, and microelectromechanical system design.