Submitted to the 44st ECTC, Electronic Components and Technology Conference, Washington D.C, USA, May 1–4, 1994

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

L. Miguel Silveira Mattan Kamon Jacob White
Research Laboratory of Electronics
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139

Abstract

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; 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.

1 Introduction

The dense three-dimensional 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 three-dimensional 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.

2 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-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 (See Fig. 1).

2.1 Mesh Formulation

To derive a system of equations for the filament currents we start by assuming 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

$$\boldsymbol{Z}\boldsymbol{I}_b = \boldsymbol{V}_b, \tag{1}$$

where $\boldsymbol{V}_b$, $\boldsymbol{I}_b \in \mathbb{C}^b$, b is the number of branches (number of current filaments), and $\boldsymbol{Z} \in \mathbb{C}^{b \times b}$ is the

complex impedance matrix given by

$$\boldsymbol{Z} = \boldsymbol{R} + j\omega\boldsymbol{L}, \tag{2}$$

where ω is excitation frequency. The entries of the diagonal matrix $\boldsymbol{R} \in \mathbb{R}^{b\times b}$ represent the dc resistance of each current filament, and $\boldsymbol{L} \in \mathbb{R}^{b\times b}$ 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

$$\boldsymbol{M}\boldsymbol{V}_b = \boldsymbol{V}_s \tag{3}$$

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

$$\boldsymbol{M}^t\boldsymbol{I}_m = \boldsymbol{I}_b, \tag{4}$$

where the superscript t denotes matrix transpose, and $\boldsymbol{I}_m \in \mathbb{R}^m$ is the vector of mesh currents. Note that one of the entries in the mesh current vector will be identically equal to the source branch current, shown as $\boldsymbol{I}_{m\_source}$ in Fig. 1. Combining (4) with (3) and (1) yields

$$\boldsymbol{M}\boldsymbol{Z}\boldsymbol{M}^t\boldsymbol{I}_m = \boldsymbol{V}_s. \tag{5}$$

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

$$\tilde{\boldsymbol{I}}_s = \boldsymbol{Y}_r\tilde{\boldsymbol{V}}_s, \tag{6}$$

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

2.2 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, $(\boldsymbol{M}\boldsymbol{Z}\boldsymbol{M}^t)\boldsymbol{I}_m$. However, it is possible to approximately compute $\boldsymbol{M}\boldsymbol{Z}\boldsymbol{M}^t\boldsymbol{I}_m^k$ in order b operations using a hierarchical multipole algorithm [8]. Such algorithms also avoid explicitly forming $\boldsymbol{M}\boldsymbol{Z}\boldsymbol{M}^t$, and so reduce the memory required to order b.

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

$$\boldsymbol{M}\boldsymbol{Z}\boldsymbol{M}^t\boldsymbol{I}_m^k = \boldsymbol{M}\boldsymbol{R}\boldsymbol{M}^t\boldsymbol{I}_m^k + j\omega\boldsymbol{M}\boldsymbol{L}\boldsymbol{M}^t\boldsymbol{I}_m^k. \tag{7}$$

Except for $\boldsymbol{L}\boldsymbol{M}^t\boldsymbol{I}_m^k$, all of the products in (7) involve sparse matrices and can be computed in order m operations. To compute $\boldsymbol{L}\boldsymbol{M}^t\boldsymbol{I}_m^k$, or equivalently $\boldsymbol{L}\boldsymbol{I}_b$, it is shown in [3] that

$$(\boldsymbol{L}\boldsymbol{I}_b)_i = \frac{1}{a_i}\int_{fil_i} \boldsymbol{A}(\boldsymbol{X}_i)\cdot\boldsymbol{l}_i(\boldsymbol{X}_i)d^3x_i. \tag{8}$$

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

$$A_k(\boldsymbol{X}) = \frac{\mu_0}{4\pi}\sum_j \frac{I_{b_j}}{a_j}\int_{filament_j} \frac{(\boldsymbol{l}_j(\boldsymbol{X}_j))_k}{|\boldsymbol{X}-\boldsymbol{X}_j|}d^3x_j, \tag{9}$$

where A_k is the k^{th} component of the vector potential. Thus by viewing $(\boldsymbol{I}_{b_j}/a_j)\boldsymbol{l}_j(\boldsymbol{X}_j)_k$ as a 'charge' then computing each component of $\boldsymbol{A}$ to eventually give $\boldsymbol{L}\boldsymbol{I}_b$ 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 $r >> R$ (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 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 $r >> R$ (as in Fig. 3). In this second case, what can

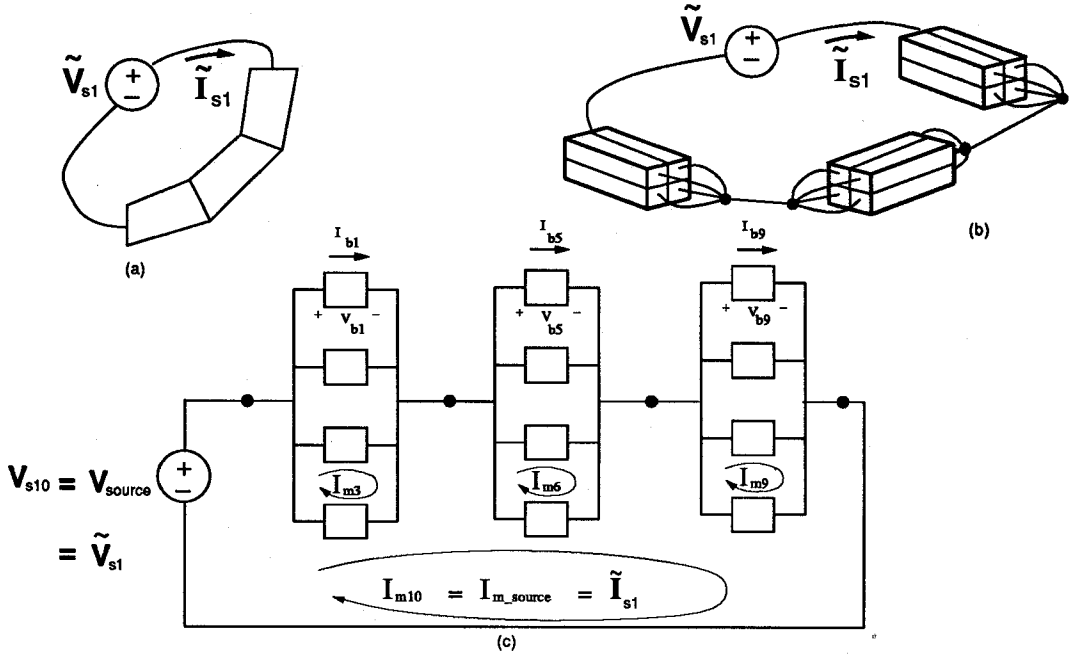

Figure 1: One conductor, (a) as piecewise-straight sections, (b) discretized into filaments, (c) modelled as a circuit.

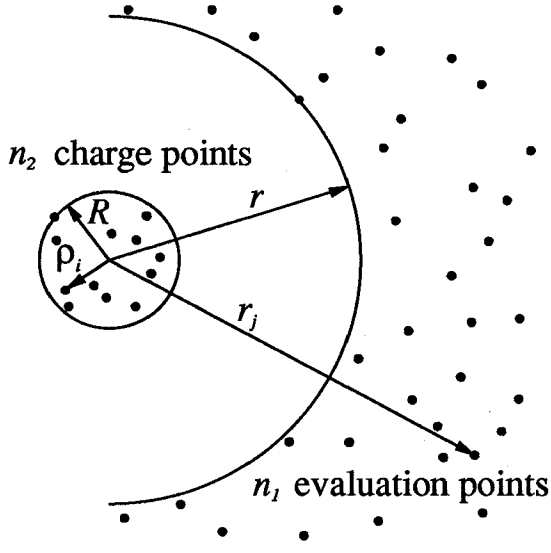

Figure 2: The evaluation point potentials are approximated with a multipole expansion.

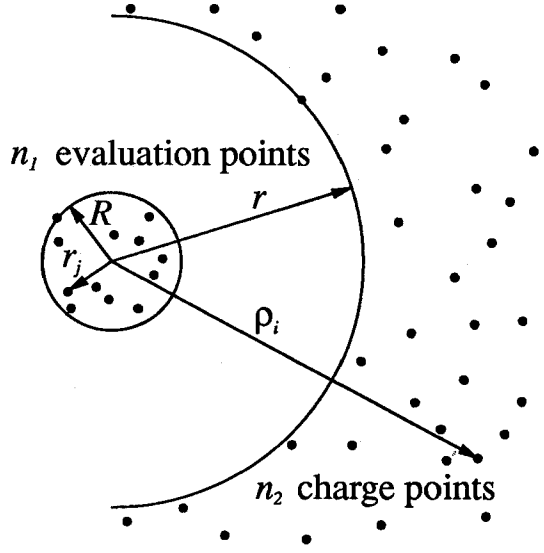

Figure 3: The evaluation point potentials are approximated with a local expansion.

be ignored is 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.

3 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.

3.1 Sectioned ℓ_2 Minimization

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

$$\boldsymbol{H}(s_j) = \boldsymbol{F}(s_j) \qquad j = 1, 2, \cdots, m \tag{10}$$

where

$$\boldsymbol{H}(s) = \frac{\boldsymbol{U}(s)}{\boldsymbol{V}(s)} = \frac{u_q s^q + \cdots + u_1 s + u_0}{v_p s^p + \cdots + v_1 s + 1} \tag{11}$$

is the low-order approximation.

Typically, the system in (10) will be overdetermined 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 $\boldsymbol{U}(s)$ and $\boldsymbol{V}(s)$ are chosen such that

$$\|\boldsymbol{H}(s) - \boldsymbol{F}(s)\|_2 = \|\frac{\boldsymbol{U}(s)}{\boldsymbol{V}(s)} - \boldsymbol{F}(s)\|_2 \tag{12}$$

is minimized for all $s \in \{s_1, s_2, \cdots, s_m\}$. 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 $\boldsymbol{V}(s)$. Then, the minimization problem becomes

$$\min_{\boldsymbol{U},\boldsymbol{V}} \|\boldsymbol{U}(s_j) - \boldsymbol{V}(s_j)\boldsymbol{F}(s_j)\|_2 \qquad j = 1, \cdots, m. \tag{13}$$

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:

$$\begin{cases} \frac{\boldsymbol{U}(0)}{\boldsymbol{V}(0)} = \boldsymbol{F}(0) \\ min_{\boldsymbol{U},\boldsymbol{V}} \|\boldsymbol{U}(s_j) - \boldsymbol{V}(s_j)\boldsymbol{F}(s_j)\|_2 \qquad j = 1, \cdots, m \\ \lim_{s\to\infty} \frac{\boldsymbol{U}(s)}{\boldsymbol{V}(s)} = \lim_{s\to\infty} \boldsymbol{F}(s) \end{cases} \tag{14}$$

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, $\Omega_1, \Omega_2, \cdots, \Omega_M$, 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 $\boldsymbol{F}(w_{11}), \boldsymbol{F}(w_{12}), \cdots, \boldsymbol{F}(w_{1m})$, a constrained ℓ_2 minimization is performed and a local approximant is computed. Once the first local approximation, $\hat{\boldsymbol{H}}_1(s)$, 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, $\boldsymbol{H}_1(s)$. Next, the second section Ω_2, the values $\boldsymbol{F}(w_{21}), \boldsymbol{F}(w_{22}), \cdots, \boldsymbol{F}(w_{2m})$ 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, $\boldsymbol{F}(s) - \boldsymbol{H}_1(s)$ will contain primarily the higher-frequency error information. This results in a new approximation for $\boldsymbol{F}(s) - \boldsymbol{H}_1(s)$, $\boldsymbol{H}_2(s)$ and therefore $\boldsymbol{F}(s) \approx \boldsymbol{H}_1(s) + \boldsymbol{H}_2(s)$. The procedure is repeated until data in the last frequency section, Ω_M, is approximated.

3.2 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 $\boldsymbol{F}(s)$. Unfortunately since $\boldsymbol{H}(s)$ 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

$$\begin{aligned} \dot{\boldsymbol{x}} &= \boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{u}, & \boldsymbol{x} \in \mathbb{R}^n,\ \boldsymbol{u} \in \mathbb{R}^m \\ \boldsymbol{y} &= \boldsymbol{C}\boldsymbol{x}, & \boldsymbol{y} \in \mathbb{R}^p \end{aligned} \tag{15}$$

whose transfer function is $\boldsymbol{G}(s) = \boldsymbol{C}(s\boldsymbol{I} - \boldsymbol{A})^{-1}\boldsymbol{B}$. 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 $\boldsymbol{A}$) are in the left half plane. It is well known [9] that for a given transfer function, the choice of the triplet $[\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C}]$ is not unique. Indeed, a linear coordinate transformation $\boldsymbol{T}$ in the state space modifies the triplet $[\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C}]$ to $[\tilde{\boldsymbol{A}}, \tilde{\boldsymbol{B}}, \tilde{\boldsymbol{C}}]$ *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 $[\tilde{\boldsymbol{A}}, \tilde{\boldsymbol{B}}, \tilde{\boldsymbol{C}}]$ 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 $\boldsymbol{G}(s)$. The balancing transformation $\boldsymbol{T}$ can be computed explicitly for any triplet $[\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C}]$. 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 $[\tilde{\boldsymbol{A}}, \tilde{\boldsymbol{B}}, \tilde{\boldsymbol{C}}]$ has the remarkable property that simple reordering and truncation of the state vector $\tilde{\boldsymbol{x}}$ with the corresponding reordering of the system matrices *necessarily* produce *stable* reduced-order models at any desirable order. Let k be this order, and let $[\tilde{\boldsymbol{A}}_k, \tilde{\boldsymbol{B}}_k, \tilde{\boldsymbol{C}}_k]$ be the reduced-order model with a transfer function $\boldsymbol{G}_k(s)$. It can then be shown [10, 11] that the impulse response of the error transfer function $\boldsymbol{E}_k(s) = \boldsymbol{G}(s) - \boldsymbol{G}_k(s)$ has an energy norm that consistently 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 L_∞ 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.

4 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 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 8pF. The interconnect from the end of pin to the receiver is modelled with a capacitance of 5pF. The overall configuration is illustrated in Fig. 5 and a more detailed view for a single pin is given in Fig. 6. A $0.1\mu F$ decoupling capacitor is connected between the driver's power and ground to minimize supply fluctuations.

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 modelling of changes in resistance and inductance due to skin and proximity effects. Matrices were generated for the frequency range 1MHz to 10MHz, 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 3. 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

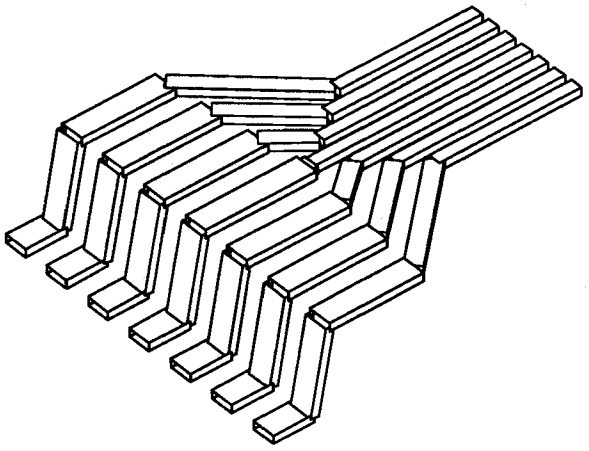

Figure 4: Seven pins of a cerquad pin package.

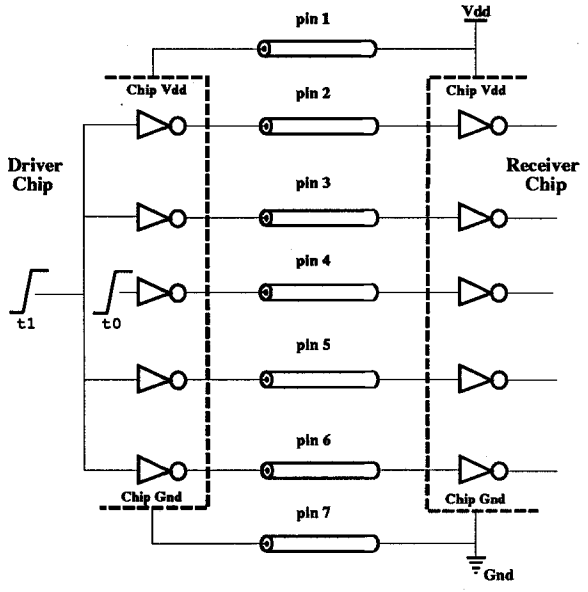

Figure 5: General configuration for the connection between received and driver chips. All the circuit elements inside the same chip share that chip's power and ground.

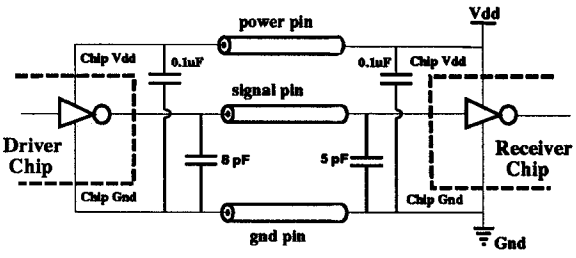

Figure 6: Detailed view of the connection between driving and receiving chips, showing the power and ground connections. Decoupling capacitance between the power and ground lines are also shown. Pin capacitance and receiver interconnect capacitance are also modeled as small capacitive loads.

truncated models of 3^{rd} order are sufficiently accurate to provide approximation with less than 5% error. The following two figures demonstrate this fact. Figure 7 shows the magnitude of the self-admittance term at pin 4. Shown in the plot are data points computed with FastHenry, the 12^{th} order section-by-section approximant and the 3^{rd} order reduced model computed by truncating the balanced realization. As can be seen on the plot, the three curves match each other almost perfectly.

Figure 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, a 20^{th} order section-by-section approximant and the 3^{rd} 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 late pins 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.

As a demonstration of the importance of frequency dependent analysis, consider the same scenario but with constant inductance and resistance matrices.

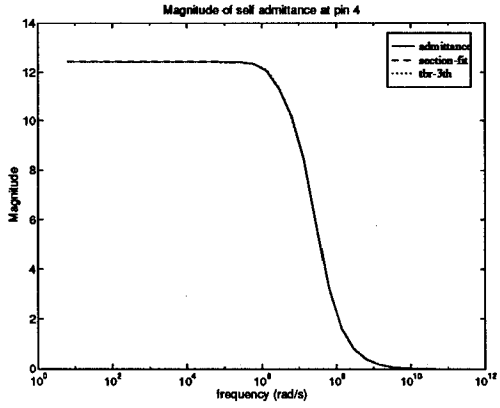

Figure 7: Magnitude of the self-admittance term at pin 4. Shown in the plot are data points computed with FastHenry, the 12^{th} order section-by-section approximant and the 3^{rd} order reduced model computed by truncating the balanced realization. The error in both approximations is less than 0.5%.

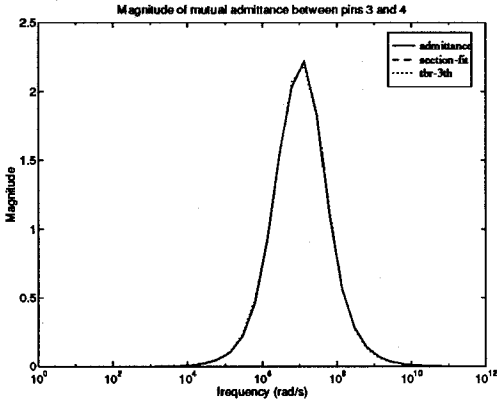

Figure 8: Magnitude of the mutual admittance term between pins 3 and 4. Shown in the plot are data points computed with FastHenry, a 20^{th} order section-by-section approximant and the 3^{rd} order reduced model. The error in both approximations is less than 1%.

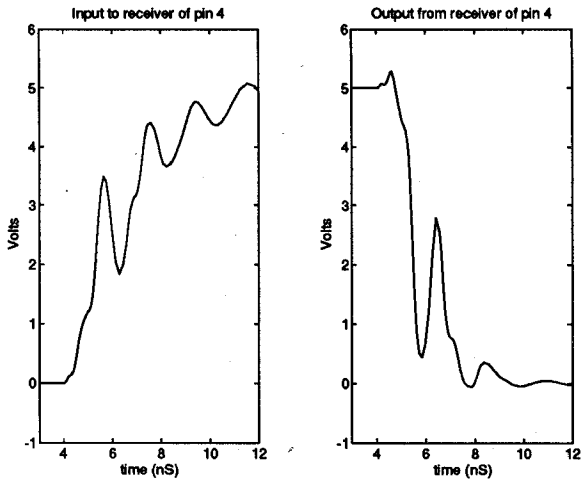

Figure 9: Pin 4's receiver when four adjacent pins switch 1 ns after pin 4.

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 limit. Note that for the receiver input waveforms, the large voltage bump swings by approximately $0.5V$ 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 $1.5V$ as shown in Fig.11.

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 modelled by letting V_{dd} and ground for both the drivers and the receivers be $5V$ and $0V$, 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 late 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.

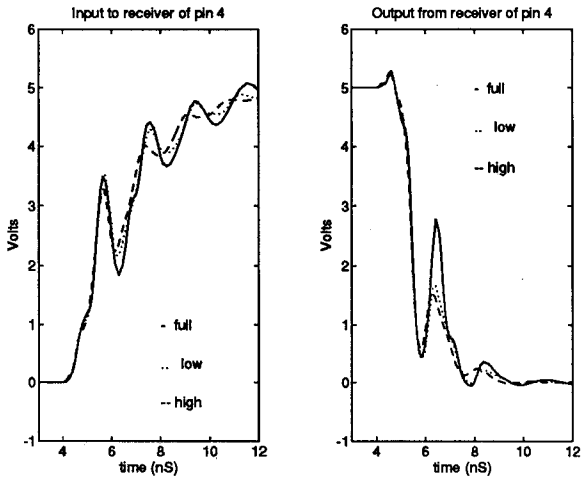

Figure 10: Pin 4's receiver using frequency dependent data (full), constant value from low frequency (low), and constant value from high frequency (high).

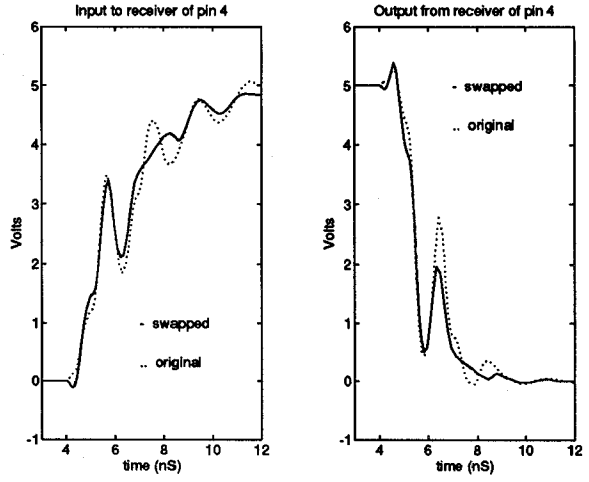

Figure 12: Pin 4's receiver with ground pin 7 and signal line 5 swapped but no current through ground pin.

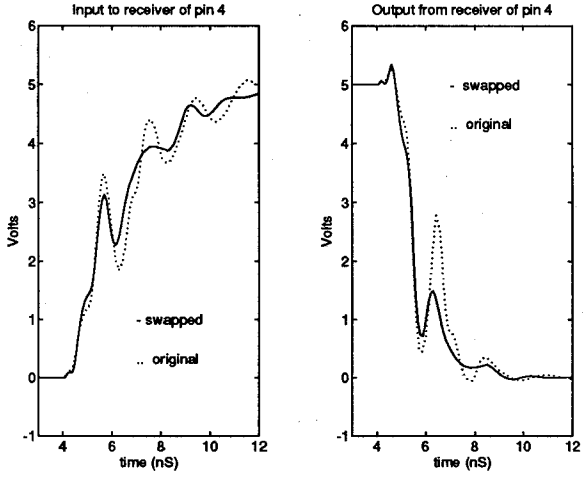

Figure 11: Pin 4's receiver with ground pin 7 and signal line 5 swapped.

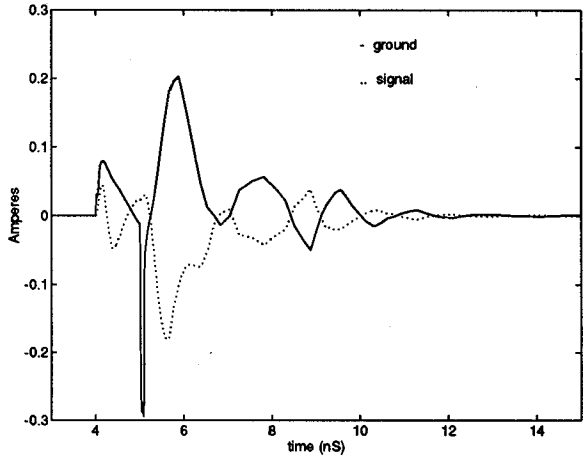

Figure 13: Total current travelling toward drivers on the four late signal lines (signal), and the current travelling back along ground pin (ground).

5 Conclusions and Acknowledgements

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. This work was supported by the Defense Advanced Research Projects Agency contract N00014-91-J-1698, the National Science Foundation contract (MIP-8858764 A02), an NSF Graduate Research Fellowship, the Portuguese "Junta Nacional de Investigação Científica e Tecnológica" under project "Ciência" and grants from Digital Equipment Corporation and IBM.

References

[1] Vivek Raghavan, J. Eric Bracken, and Ronald A. Rohrer. AWESpice: A General Tool for the Accurate and Efficient Simulation of Interconnect Problems. In 29th *ACM/IEEE Design Automation Conference*, pages 87–92, Anaheim, California, June 1992.

[2] M. Kamon, M. Tsuk, C. Smithhisler, and J. White. Efficient techniques for inductance extraction of complex 3-d geometries. In *Proceedings of the Int. Conf. on Comp. Aided Design*, November 1992.

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

[4] W. T. Weeks, L. L. Wu, M. F. McAllister, and A. Singh. Resistive and inductive skin effect in rectangular conductors. *IBM Journal of Res. and Develop.*, 23(6):652–660, November 1979.

[5] A. E. Ruehli. Inductance calculations in a complex integrated circuit environment. *IBM J. Res. Develop.*, 16:470–481, September 1972.

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

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

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

[9] Thomas Kailath. *Linear Systems.* Information and System Science Series. Prentice-Hall, Englewood Cliffs, New Jersey, First edition, 1980.

[10] Bruce Moore. Principal Component Analysis in Linear Systems: Controllability, Observability, and Model Reduction. *IEEE Transactions on Automatic Control*, AC-26(1):17–32, February 1981.

[11] Keith Glover. All optimal Hankel-norm approximations of linear multivariable systems and their L^∞-error bounds. *International Journal on Control*, 39(6):1115–1193, June 1984.

[12] Shen Lin and Ernest S. Kuh. Transient Simulation of Lossy Interconnects Based on the Recursive Convolution Formulation. *IEEE Trans. Circuits Syst.*, 39(11):879–892, November 1992.

[13] Thomas L. Quarles. The SPICE3 Iimplementation Guide. Technical Report ERL M89/44, Electronics Research Laboratory Report, University of California, Berkeley, Berkeley, California, April 1989.