Spatially Distributed 3D Circuit Models†

Michael Beattie[a], Hui Zheng[b], Anirudh Devgan[b] and Byron Krauter[a]
Electronic Design Automation[a]
Austin Research Lab[b]
IBM Corporation
11400 Burnet Road, Austin, Texas 78758

Abstract

Spatially distributed 3D circuit models are extracted with a segment-to-segment BEM (Boundary Element Method) algorithm for both capacitance and inverse inductance couplings rather than using the traditional net-to-net approach. Critical issues regarding the extraction efficiency and accuracy of segment-to-segment BEM capacitance models are explored. An adaptive discretization scheme is developed for segment-to-segment capacitance extraction and also applied to segment-to-segment high-frequency inverse inductance extraction. We demonstrate the limitations of the duality between capacitance and inverse inductance. Examples demonstrating the accuracy of these models are presented for real packaging cases.

Categories and Subject Descriptors

J.6 [Computer-Aided Engineering]: Computer-aided design (CAD).

General Terms

Design, Theory, Verification.

Keywords

Distributed Circuit Models, Boundary Element Method (BEM), Capacitance, Inverse Inductance.

1. Introduction

Traditional BEM extractors reduce segment-to-segment potential matrices, which contain detailed spatial coupling information, to net behaviors and in the process discard this spatial information. This reduces segment-to-segment electrostatic potential coupling matrices to smaller net-to-net capacitive coupling matrices. Similarly, a traditional BEM inductance extractor reduces partial inductance matrices (i.e. magnetic vector potential matrices) to smaller loop inductance matrices or some other port model. But knowing how capacitances and inductances spatially distribute on a fine scale is critically important if the dynamic electromagnetic behavior in either the time or frequency domain needs to be modeled.

Generating spatial coupling from a more traditional BEM extractor requires inherently approximate heuristics. For example, Yu and Shi [1] introduced a method that considers small windows and even smaller line segments. They place short end caps on the segment for which couplings are to be determined (active segment) and solve for the panel charges when this segment and its end caps are held at unit potential and the surrounding conductors are grounded.

Their method uses heuristics to deal with the extra end cap charges and what portion of the surrounding panels couple to the active segment versus its end caps. Moreover, because they alter the shapes close to the active segment, the end caps do not match the real conductor shapes and so their method introduces errors in the capacitance calculation, because electrostatic and magnetostatic potentials can be superimposed linearly, but their inverses cannot. In other words, the rows and columns of the electrostatic coupling matrix P are independent of the number of shapes being considered, but the dominant terms in the inverse, P^{-1}, strongly depend on whether or how much of the adjacent shapes are included.

Inverse inductance, however, is spatially distributed naturally. Inverse inductance couplings, when obtained by inverting a partial inductance matrix, are by construction segment-to-segment couplings. The BEM capacitance algorithm is modified from a net-to-net method to a segment-to-segment method in similar fashion and consequently avoids approximate distribution heuristics.

2. Background

Inverse inductance extraction is related to but still differs from traditional BEM techniques. Traditional 3D BEM extractors start by computing potential coupling matrices - electrostatic for capacitance extraction and partial inductances for inductance extraction. However, they discard spatial information when they reduce the large potential coupling matrices to smaller net behaviors that strictly adhere to either the electrostatic behavior of capacitance or the magnetostatic behavior of inductance. During BEM capacitance extraction, every net surface is held at a single potential and during BEM inductance extraction, Kirchoff's current law is strictly enforced for DC connected circuits.

As first introduced in [3], inverse inductance models are created by simply inverting the low-frequency partial inductance matrix. Since the matrix size is not reduced, e.g. to a smaller loop inductance matrix, the fine-scale spatial information found in the partial inductance matrix is retained. Because shielding effects make off-diagonals decay rapidly with distance, the inverse inductance models can be sparsified and the larger sparse model used directly in simulation after some modifications are made to the simulator [4].

While the truncation process actually alters the magnetic model [5], the changes are similar to equipotential shell approximations. In the latter, close and equal length conductors undergo similar shifts in their partial self and partial mutual inductances. Loop inductances, which are proportional to the difference $L–M$, do not change much and so the subsequent circuit solution is accurate. These similarities

† This work was supported by IBM Corporation.

DAC 2005, June 13-17, 2005, Anaheim, California.

aside, inverse inductance proved superior to sparse partial inductance methods for two reasons.

First, sparse inverse inductance models are more accurate than equally sparse partial inductance models because of an indirect coupling mechanism. Second, inverse inductance models, when combined with Norton sources, make nodal analysis possible for which the circuit equations are *symmetric* and *positive definite*. Solving these matrices can now be done using *Cholesky factorization*, exploiting symmetry, rather than using more general matrix decomposition algorithms. A more detailed discussion of both benefits as well as additional examples can be found in [6].

The idea of inverting the potential matrix without imposing the electrostatic constraints found in traditional BEM extraction, can also be applied to capacitance extraction. The result is a spatially distributed capacitance model which can be combined with a spatially distributed inverse inductance model to produce a spatially distributed $RL^{-1}C$ circuit model. Windowed extraction can be performed and the capacitance made more sparse via truncation, since shielding still occurs.

In the following, Section 3 describes the basic segment-to-segment BEM capacitance extraction algorithm and discusses the relevance of a new self-coupling term. Beattie and Pileggi [7][8] introduced but only obliquely mentioned this algorithm. Section 4 discusses the discretization and gridding requirements of segment-to-segment BEM capacitance extraction. Two separate requirements are recognized: *segment lengths* which are determined by the circuit speeds and *panel sizes* within a segment which are determined by the extraction accuracy target.

An adaptive gridding scheme is introduced to improve extraction efficiency in windowed distributed BEM extractions. Active and touching segments are finely discretized while far and already shielded segments are minimally discretized. Approximate coupling coefficients — derived assuming all segments are spheres — are used to determine which are near and which are far or already shielded segments. These adaptive gridding improvements are then applied to inverse inductance extraction in Section 5 to efficiently create a more accurate high frequency inverse inductance model. In addition, we demonstrate the limitations of the duality between capacitance and inverse inductance.

In Section 6 we present two examples: A simple test case for which we could compare our results to two-dimensional surface extraction results. This example illustrates the accuracy of circuit models combining segment-to-segment capacitance models with high-frequency inverse inductance models. In a second example, hardware TDR measurements are compared to the simulation model of a ceramic BGA package. This further demonstrates the accuracy of our approach and the viability of distributed 3D BEM circuits to model tens of thousands of shapes and wires at arbitrary angles.

3. Distributed BEM Capacitance Extraction

In a traditional BEM capacitance algorithm, conductor surfaces are divided into panels and Poisson's integral equations are approximated by the linear equations

$$\boldsymbol{P}q = v \tag{1}$$

where $\boldsymbol{P}$ is a matrix of panel-to-panel coupling coefficients, v is a vector of panel potentials, and q a vector of panel charges. The total capacitance for net i is obtained by applying unit potential to *every* panel on net i and zero potential to all other panels on all other nets in its vicinity or potentially the entire the system, solving for the panel charges, and finally summing up these panel charges. The total capacitance between net i and j, C_{ij}, is given by the sum of the panel charges on net j, when net i is held at unit potential and all other panels on all other are grounded.[1] In summary,

$$\boldsymbol{C} = A^T \boldsymbol{P}^{-1} A \tag{2}$$

with $\boldsymbol{A}$ being the *incidence matrix*, having as many rows as panels and columns as conductors. A_{ij} is one if panel i is part of conductor j, zero if not. While assigning constant surface potentials to every net physically matches electrostatics, it also means the traditional BEM capacitance extractor only computes total net capacitance and cannot determine how the net-to-net couplings are spatially distributed.

Although their papers focus on the efficacy and accuracy of windowed BEM extraction and sparse potential methods, the capacitance extraction algorithm proposed by Beattie and Pileggi [7][8] actually apply discontinuous surface potentials and computed segment-to-segment capacitances. Net-to-net capacitances are subsequently obtained by summing the segment-to-segment capacitances. The linear equations $\mathbf{P}q = v$, sparse and localized, are solved for every segment of every net by applying unit potential to just one segment while all other segments, *including those on the same net as the active segment*, are grounded.

For example, when a net i is divided into N continuous segments their algorithm solves the localized linear equations N times, each time assigning unit potential to the panels of the active segment only, while grounding the panels on all other conductor segments — including the other N-1 segments on net i. In each of these N solutions, the panel charges on other net segments (including the N-1 grounded segments on net i) will equal the distributed coupling capacitance to the active segment on net i that is held at unit potential.

Charge vectors associated with net-to-net capacitances, which are normally found by solving (1) when an entire net was assigned to unit potential and all other nets are grounded, are found by summing the N segment-to-segment solutions:

$$q_{\text{net-to-net}} = \sum_{i=1}^{N} q_{i_{\text{segment-to-segment}}} \tag{3}$$

Since (1) is a set of linear equations, superposition applies and net-to-net capacitances obtained via (3) are as accurate as the traditionally obtained net-to-net capacitances — for all but large and ill-conditioned matrices.

It must be noted, however, that segment-to-segment BEM capacitance extraction considers physically unrealizable surface potential discontinuities along the net. Furthermore, these discontinuities show up as charge density singularities which when viewed as capacitances represent forward and backward couplings between touching segments on the same net. This coupling has been historically ignored in capacitance modeling for three reasons. First, it is absent in the 2D limit. Second, it is small for most practical 3D cases unless a net wraps back on itself. And finally, its importance is negligible in circuit analysis since the potential differences between adjacent segments is typically negligible for even the fastest signals.

1. Multi-dielectrics are handled by either discretizing the dielectric interfaces into panels and modeling these panels as fixed charge surfaces or if the dielectric interfaces are planar, employing more complicated Greens functions when computing the coupling coefficient matrix.

The four shaded rectangles in Fig. 1 represent two wires discretized into two segments each. The RC circuit superimposed on these rectangles is what a segment-to-segment BEM capacitance extraction will create. Touching segments on the same net couple to one another. However, because this self coupling capacitance is in series with the wire resistance of the segment, R_{seg}, its impact will be negligible in simulation.

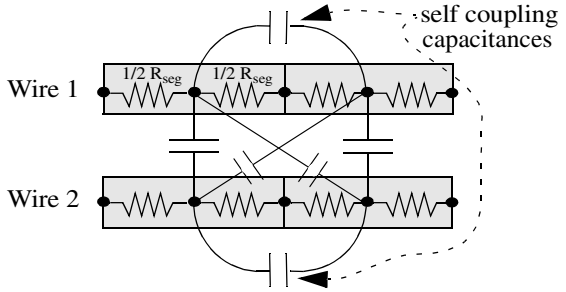

Fig. 1: Distributed BEM capacitive couplings for two wires discretized into two segments each.

4. Discretizing Distributed BEM Models

In the traditional BEM extraction/simulation flow, the net-to-net capacitances and loop inductances are extracted first. Then, via heuristics, a lumped circuit model distributing these couplings is constructed. Hence, during the traditional BEM extraction the only concern are the maximum panel or filament sizes that accurately extract net capacitance or loop inductance. The lumped circuit time constants found in simulation are subsequently determined by distribution heuristics. To avoid the uncertainties of these heuristics, which are only well established for 2D structures, we use the segment-to-segment method to create the distributed lumped model accurately during the extraction itself. Therefore, distributed BEM extraction has two discretization requirements.

First, the *maximum segment length* must be electrically short relative to the fastest circuit edge rate. The faster the circuits operate, the shorter the segment lengths must be. The following rule of thumb can be used for determining maximum segment length

$$\tau_f < \tau_r / 10 \tag{4}$$

where τ_r is the rise time of the signal and τ_f is the minimum time-of-flight delay along the segment. Because signal speeds in areas where TEM propagation exists, are equal to the speed of light in that dielectric medium, $v = c / \sqrt{\varepsilon_r}$, the maximum segment length is given by

$$l_{max} = (\tau_r c)/(10\sqrt{\varepsilon_r}). \tag{5}$$

Second, panel sizes must be made small enough to accurately approximate the boundary element integral equations through matrix equations such as (1). Because these equations are linear, superposition applies and the panel sizes required for accurate segment-to-segment extraction are no smaller than those required for accurate net-to-net extraction. But simply using the same panel sizes for all shapes is not the most efficient approach for segment-to-segment extraction.

The purpose of segment-to-segment BEM extraction is *not* primarily to improve the extraction times, but to create a more accurate, spatially distributed lumped $RL^{-1}C$ circuit model *without heuristics*. If N nets are each divided into M segments, segment-to-segment BEM extraction must formulate, factor, and solve NM potential coupling matrices. On the other hand, net-to-net extraction may require as few as N if the system is small. Windowing techniques can be applied to both methods and net-to-net BEM extraction can be sped up in ways that are not practical for segment-to-segment extraction [9][10].

However, similar to the hierarchical methods [10], windowed segment-to-segment BEM extractions can be made more efficient by realizing that the panel discretization can be lowered the further a neighbor segment is away from the active segment. The smallest panels are used on the active segment. For touching segments slightly larger panels are used, and for nearby — but not touching — segments an even smaller resolution will suffice. Finally, more remote conductor segments are subdivided with even larger panels or are modeled by volume elements in the case of the most remote conductors. Fig. 2 qualitatively depicts the two separate discretization steps we propose in distributed BEM capacitance extraction.

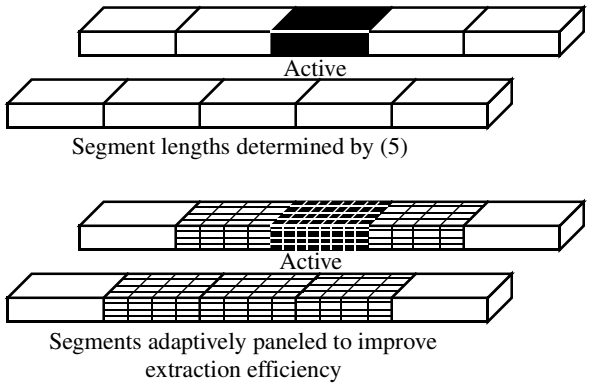

Fig. 2: Adaptive two step discretization for distributed BEM.

Wires and ground planes are initially discretized according to (5). These segments are further adaptively paneled so as to improve extraction accuracy without unduly degrading extraction efficiency. To decide which segments in a window, apart from the central segment or segments touching it, are near to or far from the central shape, different strategies may be applied. A novel strategy we use involves performing a very fast, approximate pre-extraction of the set of shapes in the window with respect to the central conductor. For this, each shape is represented by a simpler form, for instance a sphere centered at the center of each conductor, sized as large as possible so that no two spheres overlap. Setting up the potential matrix for this scenario is very simple: the self terms are proportional to $1/\textit{radius}$, the mutual terms are proportional to $1/\textit{center distance}$ of the spheres. This approach, as opposed to the more standard distance-based method, has the advantage of capturing shielding. If nearness is based only on the distance being smaller than some threshold, shapes which have a insignificant coupling to the central shape due to metal in between may be discretized finely. With the pseudo-coupling method this inefficiency is removed without sacrificing accuracy.

TABLE 1. Line-to-line capacitances for wires in Fig. 3

	Traditional BEM with 0.33 mm panels	Adaptive Distributed BEM with 0.33 & 1.0 mm panels
$C_{9\text{-}10}$	0.604 pF	0.604 pF
$C_{9\text{-}15}$	0.122 pF	0.120 pF

The results in Table 1, for a IBM Power3 CPU running AIX 5.1, demonstrate the efficacy of the adaptive method. For the traditional

BEM extraction, the wires in Fig. 3 are discretized into 0.33 mm panels. For the adaptive distributed BEM extraction, wire segment lengths are 1.0 mm and panel sizes are 0.33 mm for active, touching, and near segments and 1.0 mm for remote segments. Accuracy is affected insignificantly. For a representative measure of the speed-up for the adaptive method, see the example in Fig. 9.

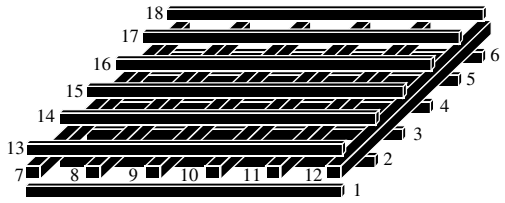

Fig. 3: Eighteen wire example with wire widths, heights, and spaces all 1.0 mm in a uniform dielectric, $\varepsilon_r = 4.0$.

5. High Frequency Inverse Inductance

Uniform interconnect current density implies signal frequency close to zero and therefore negligible skin and proximity effects [12]. In this domain the magnetic model is not significant, since the dominant impedance contribution is resistive. So, while inverting a partial inductance matrix implies *high frequency*, assuming uniform current distribution within each wire implies *low frequency.* This is a *self-contradicting* proposition.

We improve on such previous approaches as [3] by splitting each wire segment into many filaments with small cross-sections, similar to [13], and invert this larger filament-to-filament partial inductance matrix to find the correct high frequency inverse inductance. Since we use the same set of segments for both C and L^{-1} extraction, the segment-to-segment C and L^{-1} couplings can be easily combined in the final netlist. To get the *conductor-to-conductor* inverse inductance, we accumulate all filament currents for each conductor

$$L_{hf}^{-1} = A_p^T (L^f)^{-1} A_p \tag{6}$$

with A_p being the *incidence matrix* for parallel wires. A_p has as many rows as filaments and columns as conductors. Its elements are one if a filament i is part of conductor j, zero if not. Shielding by outer layer filaments de-couples the inner filaments from the rest of the system, thus only the surface contributes to the high frequency inverse inductance model.

The summation of filament-to-filament inverse inductance is possible in this case, since it satisfies $L^{-1}\boldsymbol{v} = d\boldsymbol{i}/dt$. All filaments have the same induced voltage drop since they are shorted at the ends. This high frequency inverse inductance is *not* the same as the inverse of the conductor-to-conductor partial inductance *with uniform current density.* Rather, it implies that *current is only present at the surface*.

Splitting wire segments into multiple filaments in a high frequency L^{-1} extraction is similar to the additional panelization of conductor segments in capacitance extraction. Therefore, similar adaptive discretization methods can be applied: We split the active and touching segment finely and the more remote segments more coarsely. An approximate pre-extraction to determine which are near and far shapes, can still assume spherical shapes, but now current directions in the wire segments should be considered.

A second, more general example demonstrates that inverse inductance is not the dual of capacitance, despite earlier claims [3,5,11]. The differences stem from how the filament/panel coefficients are processed, not how these coefficients are calculated. Unlike capacitance, inductance extraction must satisfy both KCL and KVL and therefore treat *cross-section (parallel) discretizations* differently from *length (serial) discretizations*.

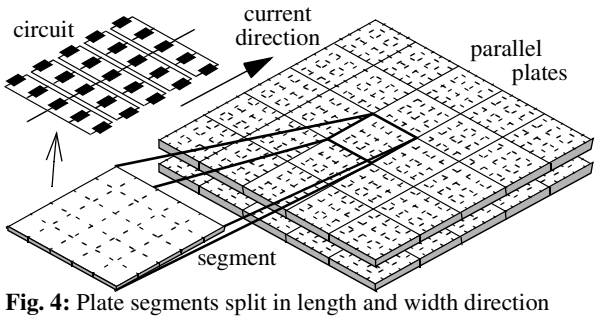

Fig. 4: Plate segments split in length and width direction

Two parallel plates shown in Fig. 4 are split into segments. Each segment is split into panels/filaments for electric/magnetic extraction. For capacitance, we create the potential matrix P^p for the panels, invert it to get the panel capacitance matrix C^p and then sum up all couplings between panels on the same pair of segments to get the segment capacitance C^s

$$C^s = A^T (P^p)^{-1} A \tag{7}$$

where the incidence matrix element A_{ij} is one if panel i is on segment j, zero otherwise.

For inverse inductance *we need to distinguish parallel and serial filaments*. Inductances create voltage drops, so they can be added up if they are in series, inverse inductances create inductor currents, so they can be added up if they are in parallel. Therefore, the segment inverse inductance $L^{s,-1}$ is

$$L^{s,-1} = (A_s (A_p^T (L^f)^{-1} A_p)^{-1} A_s)^{-1} \tag{8}$$

where in addition to the parallel incidence A_p from (6) there is the serial incidence A_s mapping parallel groups of filaments to segments. For parallel filaments inverse inductances are added, while for serial groups the inductances are added. This matches the equivalent circuit in Fig. 4 for segments. Eq. (8) reverts to Eq. (6) if there are no length splits for $(L^f)^{-1}$.

For the example in Fig. 4, C^s is diagonally dominant and has only negative off-diagonals if the segments are sufficiently discretized[2]. This would apply to the magnetic case as well, if it also adhered to (7), but it does not. Eq. (8) applies, because the current directions matter for magnetic couplings. Therefore, $L^{s,-1}$ might not be diagonally dominant and may have positive off-diagonal elements, although the number and magnitude of these are much smaller than for the filament inverse inductance matrix $(L^f)^{-1}$. For the example in Fig. 4, $L^{s,-1}$ is not diagonally dominant and has positive off-diagonal elements.

The proof that the capacitance matrix is diagonal dominant and all its off-diagonals are negative assumes the surface panel potentials for each conductor to be constant and either 0 or 1 Volts. Eq. (7) embodies this assumption, Eq. (8) does not. Inverse inductance extraction equates the integral of the vector potentials along the length to either 1 or 0, therefore does not set the vector potential for the entire surface to 1 or 0.

2. Two 1x1 mm^2 parallel plates separated by 40 µm are divided into 5x5 segments, each further divided into 5x5 panels.

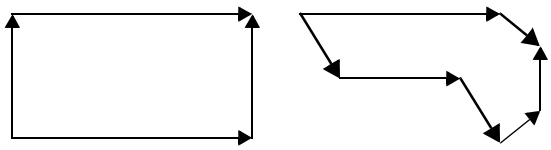

Fig. 5: Rectilinear and arbitrary angle wires

It must be pointed out that there is for inverse inductance yet another source of positive off-diagonal terms, *wires at arbitrary angles*. Rectilinear wires can always be oriented such that all mutual inductances are positive. But for wires oriented at arbitrary angles, Fig. 5, right side, no matter how the current directions are chosen, there are always some segments for which either the assumed *x*- or the *y*-current directions are opposed to those of other segments. *Between these segments, the mutual inductances will be negative, and the mutual inverse inductances positive.* Since our tool, *PATS (Package Analysis Tool Suite),* analyzes wires at arbitrary angles, we do not handle positive off-diagonals in any special way. Truncation is determined by distance because even for inverse inductance modeling the current loops must be coupled.

6. Application

To demonstrate the advantages of distributed BEM circuit models and the necessity of modeling the high frequency inverse inductance properly, we present two examples: A simple test case which we compare to two-dimensional surface extraction methods illustrates the benefits of high frequency inverse inductance extraction. A second more complex extraction and simulation of a TDR measurement on a real ceramic BGA package demonstrates the accuracy and the viability of distributed BEM methods to model tens of thousands of shapes and wires at arbitrary angles. Both examples were run on IBM Power4 CPUs running AIX 5.1.

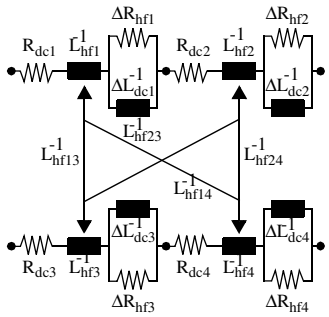

Fig. 6: Filter circuits for four conductors.

Our modeling approach begins by combining low frequency resistance and high frequency inverse inductance. It is, however, useful for more than just low and high frequency analysis. The most significant changes with frequency in impedance are caused by return current crowding. This effect is captured naturally by any three-dimensional partial inductance or inverse inductance model that discretizes and carries a wide return path into simulation. One could also capture signal-to-signal proximity effects by keeping multiple signal wire filaments in simulation, but currently we use uncoupled $\boldsymbol{L^{-1}R}$ *Foster filters* (see Fig. 6) to include skin effect behavior in the individual conductor cross-sections.

Our first example is a three-dimensional test case easily compared to two-dimensional extraction and simulation results. Shown in Fig. 7, it consists of two 8 mm parallel signal lines between two wide ground planes within a homogenous dielectric material, $\varepsilon_r = 3.8$. A 2D BEM surface extraction predicted a 46 Ω characteristic impedance and a 51 ps time of flight for the active line, and when ideally terminated, near and far-end coupling noises of 0.084 and 0.0 times the incident wave voltage, $(K_C \pm K_L)/4$. The active line is driven with a 46 Ω, 21.74 mA Norton source at one end and terminated in 46 Ω at the other end. The quiet line is terminated in 46 Ω at both ends. Fig. 7 also shows the layout after ground plane discretization in preparation for extraction.

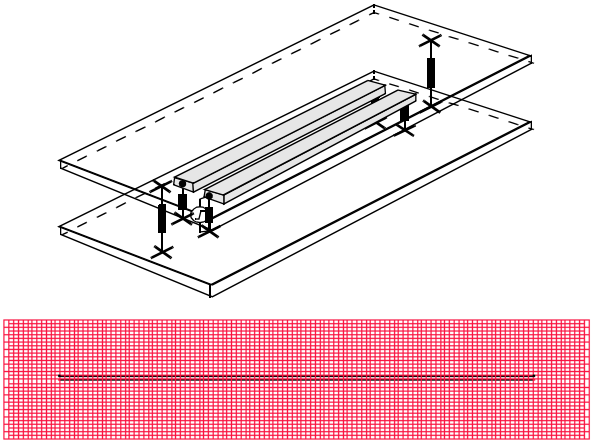

Fig. 7: Layout for test case for high frequency inductance

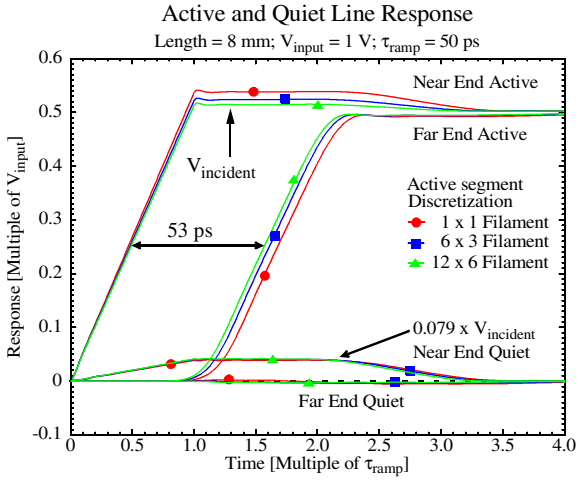

Fig. 8: Voltage responses for case in Fig. 7.

The voltage responses in Fig. 8 show the importance of using high frequency inverse inductance, obtained using (6), when operating at high frequencies. Active line impedance and delay predictions were inaccurate using low frequency inverse inductance (1x1 filament case). For the 12x6 active segment case however, the results were very close to the 2D predictions: 50 Ω and 53 ps respectively. For noise prediction, however, the amount of cross-section discretization was apparently less important. The absolute magnitude of the near and far end noise changed little when the inverse inductance model was further discretized, but it should also be noted that the relative error of the far end noise changed considerably. Far end noise is proportional to the difference in capacitive and inductive coupling coefficients. For two symmetric strip lines in a homogenous material these components are equal, $K_C = K_L$. In the example in Fig. 7 the far end noise should be zero and indeed the simulation results in Fig. 8 with the high frequency inverse inductance model are much closer to zero.

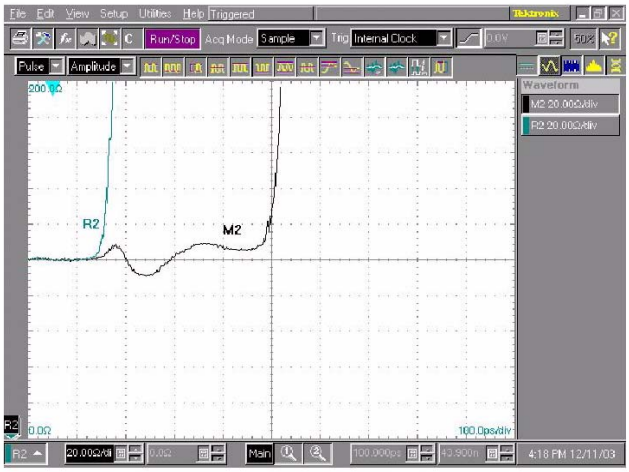

Open circuited differential pair TDR measurements on a ceramic BGA package.

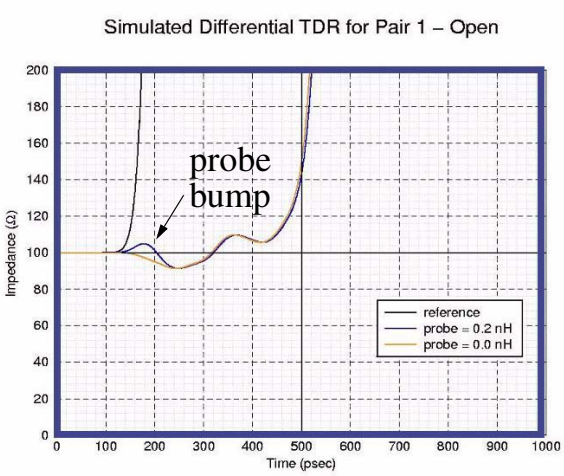

Distributed BEM simulation results.

Fig. 9: Measured and simulated TDR measurements on a differential net pair on a ceramic BGA package.

The second example compares a TDR measurement taken on a differential net pair on an actual ceramic BGA package to distributed BEM simulation results. The *halo* for this net, the set of all shapes in the vicinity of the differential net, contained about 8,000 individual shapes. Detailed modeling of electric and magnetic interactions with sufficient accuracy is made possible only by applying adaptively gridded, distributed BEM methods. These methods do not rely on any type of regularity in the layout, as many other heuristic approaches using net capacitances and loop inductances do, and can therefore be applied to complex circuits containing wires at any angle and very irregular ground planes. The average number of shapes in each extraction window was 80. Using the adaptive gridding method outlined earlier, extracting the capacitances and inverse inductances required a runtime of about one hour, while a conventional, non-adaptive extraction of the same system was far from completion even after three days on the same IBM Power3 CPU running AIX 5.1. The accuracy of the adaptive method is demonstrated in Table 1 and in Fig. 9.

The two waveforms in top half of Fig. 9 are an open-circuited reference *measurement* and the reflections when the differential pair was driven from the top surface of the package and the bottom surface pins were open-circuited. Except for a missing initial bump, which indicates the measurement saw an inductive discontinuity near the top surface, our simulations matched the measurements both qualitatively and quantitatively. And the missing bump could be produced in simulation if we assumed the measurements did not zero-out 0.1 nH parasitic inductance on each probe.

7. Summary

Distributed BEM circuit models that require no distribution heuristics are presented as a novel means of modeling the increasingly critical electric and magnetic interactions in very complex interconnect structures. Key issues relating to discretization and the proper application of these methods to interconnect modeling are discussed, as well as a novel means to improve extraction efficiency of these models. The limitations of the duality between capacitance and inverse inductance and their consequences for extraction are demonstrated. Finally, the methods developed here were applied to examples which demonstrate the practical impact of this modeling approach.

8. References

[1] F. Yu, W. Shi, *A Divide-and-Conquer Algorithm for 3D Capacitance Extraction*, 5th ISQED (Mar. 2004)

[2] B. Krauter, L. Pileggi, *Generating Sparse Partial Inductance Matrices with Guaranteed Stability*, ICCAD 1996 (Nov. 1996)

[3] A. Devgan, H. Ji, W. Dai, *How to Efficiently Capture On–Chip Inductance Effect: Introducing a New Circuit Element K*, ICCAD 2000 (Nov. 2000)

[4] H. Ji, A. Devgan, W. Dai, *KSim: A Stable and Efficient RKC Simulator for Capturing On-Chip Inductance Effect*, ASP-DAC 2001 (Jan. 2001)

[5] M. Beattie, L. Pileggi, *Efficient Inductance Extraction via Windowing*, DATE 2001 (Mar. 2001)

[6] H. Zheng, L. Pileggi, M. Beattie, B. Krauter, *Window-Based Susceptance Models for Large-Scale RLC Circuit Analyses*, DATE 2002 (Mar. 2002)

[7] M. Beattie, L. Pileggi, *Bounds for BEM Capacitance Extraction*, 34th DAC (Jun. 1997)

[8] M. Beattie, L. Pileggi, *Error Bounds for Capacitance Extraction via Window Techniques*, IEEE TCAD, vol. 18, No. 3 (Mar. 1999)

[9] K. Nabors, J. White, *FASTCAP: A Multipole Accelerated 3-D Capacitance Extraction Program*, IEEE TCAD, vol. 10, No. 11 (Nov. 1991)

[10] W. Shi, J. Liu, N. Kakani, T. Yu, *A Fast Hierarchical Algorithm for 3-D Capacitance Extraction*, 35th DAC (Jun. 1998)

[11] T. Chen, C. Luk, C. Chen, *INDUCTWISE: Inductance-Wise Interconnect Simulator and Extractor*, IEEE TCAD, vol. 22, No. 7 (Jul. 2003)

[12] J. Jackson, *Classical Electrodynamics*, 2nd Edition, John Wiley & Sons, New York (1975)

[13] W. Weeks, L. Wu, M. McAllister, A. Singh, *Resistive and inductive skin effect in rectangular conductors*, IBM J. Res. Dev., vol. 30, No. 6 (Nov. 1979)