Return-Limited Inductances: A Practical Approach to On-Chip Inductance Extraction

Kenneth L. Shepard and Zhong Tian
Integrated Systems Laboratory, Department of Electrical Engineering
Columbia University, New York, NY 10027

Abstract

This paper proposes a practical approach for extracting approximate inductances of on-chip interconnect. This approach, which we call the method of return-limited inductances, is based on performing the inductance modeling of signal lines and power-ground lines independently and on taking advantage of the power-ground distribution of the chip to localize inductive coupling. Results from representative on-chip interconnect examples demonstrate the efficiency and accuracy of this approach over other sparsification techniques.

1 Introduction

To model inductances in complex integrated circuit environments in which current return paths are not known, *partial-element equivalent circuit* (PEEC) models are traditionally used[1]. Because *partial inductances* obey the same branch constitutive relations as closed-loop inductances, they can be conveniently applied in the context of modified nodal analysis and used internally in circuit simulators such as SPICE.

Nonetheless, the partial inductance approach, which assigns portions of the loop inductances to segments along the loop, results in a large, densely-coupled network representation, which makes subsequent circuit simulation practical only for small examples. One approach to make the inductance matrix sparse and, therefore, more tractable is simply to discard those mutual coupling terms of the partial inductance matrix which are below a certain threshold. This approach, however, does not guarantee the positive semi-definiteness of the resulting inductance matrix[2]. As an alternative to simple truncation, another proposal[3, 4] is to associate with every segment a distributed current return out to a shell r_0 away. Segments spaced more than r_0 apart have no inductive coupling. The value of r_0 that must be used to achieve a desired accuracy depends on the interconnect topology, the connectivity of the nets in question, and the particular net or coupling being considered. This leads to complicated schemes in which the value of r_0 is dynamically chosen based on accuracy in a particular transfer function calculation[2]. Because r_0 is a global parameter, as r_0 is adjusted, the entire inductance extraction must be redone. Moreover, this approach does not work well for long wires broken into many segments.

In this paper, we present *return-limited inductance extraction* as an approximate on-chip inductance extraction technique. This approach can be applied to complex geometries with the same generality as a partial inductance formulation, while resulting in a positive-semi-definite and structurally sparse inductance matrix.

2 Return-Limited Inductance Extraction

Return-limited inductance extraction achieves sparsity in two important ways:

- The inductances of the power and ground lines are extracted independently from the signal lines; therefore, there are no mutual inductances between power-ground lines and signal lines. This is consistent with the traditional approach in digital IC design of performing independent signal and power-grid analysis and takes advantage of the known availability of power-ground lines as high-frequency current returns.
- Mutual inductances between signal lines and between power-ground lines are further restricted by a set of simple geometry-based decomposition rules, which we call *halo rules.*

The return-limited inductance extraction approach consists of three steps:

- Using halo rules to determine the wire segments which inductively couple.
- Applying fracturing techniques to partition the wires of the chip for resistance and inductance extraction.
- Calculating the return-limited inductances and resistances.

2.1 Halo rules for signal-line extraction

We begin with a few definitions. A wire *segment* is a rectangular parallelepiped defined by coordinates $(x_{min}, y_{min}, z_{min})$ and $(x_{max}, y_{max}, z_{max})$. A *horizontal segment* is one in which the current flow is known to be horizontal, while a *vertical segment* is one in which the current is known to be vertical. The *halo* of a segment consists of the six semi-infinite subregions as shown in Figure 1.

The *halo rules* as applied to signal-line extraction are given as follows:

- Horizontal and vertical signal line segments are treated independently since they do not inductively couple to each other.
- Horizontal (vertical) halos of power-ground lines are "blocked" by horizontal (vertical) signal segments. If the halos are viewed as columned beams emanating orthogonally from each face of a segment, then blocking occurs whenever these beams are interrupted by another segment.

0-7803-5443-5/99/$10.00©1999 IEEE

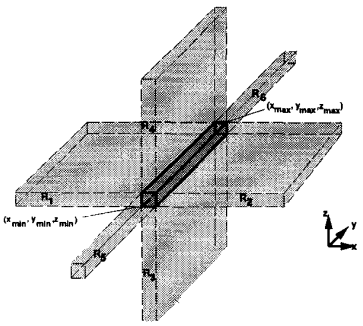

Figure 1: The *halo* of a given segment consists of six semi-infinite regions. The *horizontal halo* consists only of regions R_3, R_4, R_5, and R_6, while the *vertical halo* consists only of regions R_1, R_2, R_3, and R_4.

- Inductive coupling between two horizontal (vertical) segments is nonzero if and only if it is possible to connect two segments by a path which does not cross the horizontal (vertical) halo of any ground or supply line.

Because of space limitations, detailed consideration of the halo rules will be presented elsewhere[5]. In some cases, refinement of the halo rules is necessary to preserve couplings that would normally be discarded.

These halo rules divide the chip interconnect into a collection of disjoint *horizontal (vertical) interaction regions* defined by the non-blocked horizontal (vertical) halos of the power and ground distribution. Horizontal (vertical) segments must be contained within the same *horizontal (vertical) interaction region* to inductively couple.

In Figure 2, we are "looking" down on a wire topology in layers M3, M4, and M5 in a hypothetical five-layer-metal process. The M4 and M5 shapes are all associated with power or ground. Two signal lines (*net*1 and *net*2) are routed on M3. The interaction region shown in cross-hatch acts as a vertical interaction region for signal segments 1 through 4 and as a horizontal interaction region for signal segments 5 and 6. Note that segments 1 and 4 couple (and, therefore, belong to the same vertical interaction region), for example, because the vertical halo of the intervening ground line is blocked.

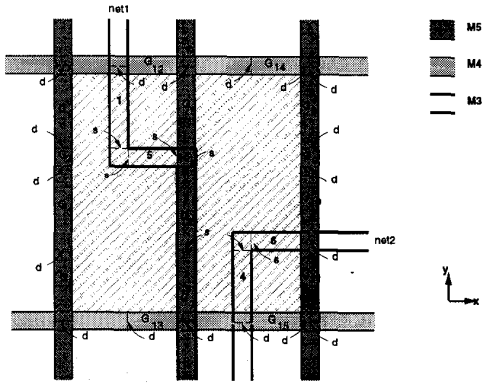

Figure 2: Application of halo rules in three dimensions.

2.2 Fracturing techniques

After the interaction regions are defined, extraction begins by *fracturing* each wire into a set of segments. Fracturing is itself a two-stage process – static fracturing to establish rectangular segments and current direction, followed by dynamic fracturing to define the geometries for inductance extraction. The static fracturing approach employed is very similar to that of the INDEX[6] extractor applied to superconducting circuits. If a given rectangle has two edges that are contacted by a neighbor, we determine if the current flow through the segment is horizontal, vertical, or mixed. Naturally, static fracturing seeks to minimize the wire area represented by mixed segments, while inductances are extracted only for horizontal and vertical segments.

Once static fracturing is completed for all of the interconnect, dynamic fracturing is used to create additional fractures in the signal, power, and ground wires required for the inductance calculation. First, static fractures in the power and ground lines must be projected onto the signal lines. Sometimes, additional signal line fractures may be necessary for long uniform wire runs to ensure a reasonable approximation of the distributed nature of the resistance and inductance. Also, new fractures must be created to split signal segments between interaction regions. At last, all of these signal line fractures, both static and dynamic, must be, in turn, projected onto the parallel adjacent power and lines. In the example of Figure 2, static fractures are labelled with "s", which the dynamic fractures necessary are labelled with "d."

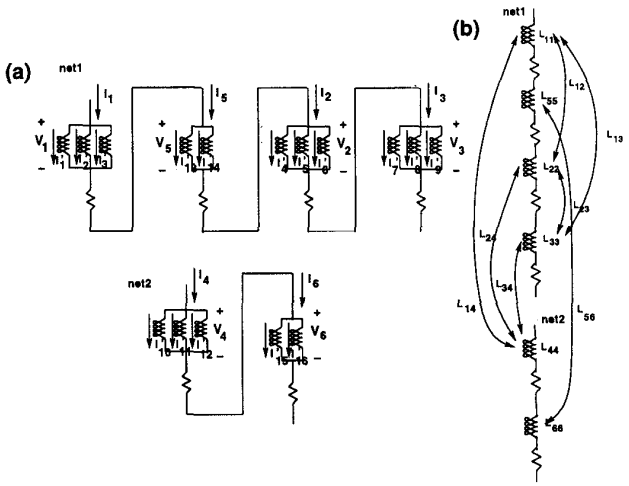

Figure 3: Signal-line resistance and inductance extractions for *net*1 of Figure 2.

2.3 Return-limited inductance formulation

The model for each signal segment which results from dynamic fracturing is a resistance in series with a parallel combination of *return-limited loop inductances*, where each loop inductance is defined by a signal line returning through one of the return segments defined on parallel, adjacent power or ground lines. For the example of Figure 2, the corresponding circuit representation for *net*1 and *net*2 is shown in Figure 3(a).

For the vertical segments of Figure 2, the return-limited inductance matrix, $s\mathbf{I}' = \mathbf{L}'\mathbf{V}$, is given by:

$$\mathbf{L}' = \begin{pmatrix} \mathbf{L}'_{11} & \mathbf{L}'_{12} & \mathbf{L}'_{13} & \mathbf{L}'_{14} \\ \mathbf{L}'_{21} & \mathbf{L}'_{22} & \mathbf{L}'_{23} & \mathbf{L}'_{24} \\ \mathbf{L}'_{31} & \mathbf{L}'_{32} & \mathbf{L}'_{33} & \mathbf{L}'_{34} \\ \mathbf{L}'_{41} & \mathbf{L}'_{42} & \mathbf{L}'_{43} & \mathbf{L}'_{44} \end{pmatrix}$$

where $\mathbf{L}'_{ij}$ are the blocks associated with the inductances of each of the four vertical segments of *net1* and *net2*. The currents $\mathbf{I}'$ and voltages $\mathbf{V}$ are as shown in Figure 3(a). In terms of the partial inductances, for example, $\mathbf{L}'_{11}$ is given by:

$$\mathbf{L}'_{11} = \begin{pmatrix} \mathcal{L}_{11} + \mathcal{L}_{\tilde{1}\tilde{1}} - \mathcal{L}_{\tilde{1}1} - \mathcal{L}_{1\tilde{1}} & \mathcal{L}_{11} + \mathcal{L}_{\tilde{1}\tilde{5}} - \mathcal{L}_{1\tilde{5}} - \mathcal{L}_{\tilde{1}1} & \mathcal{L}_{11} + \mathcal{L}_{\tilde{1}\tilde{8}} - \mathcal{L}_{1\tilde{8}} - \mathcal{L}_{\tilde{1}1} \\ \mathcal{L}_{11} + \mathcal{L}_{\tilde{5}\tilde{1}} - \mathcal{L}_{1\tilde{1}} - \mathcal{L}_{\tilde{5}1} & \mathcal{L}_{11} + \mathcal{L}_{\tilde{5}\tilde{5}} - \mathcal{L}_{1\tilde{5}} - \mathcal{L}_{\tilde{5}1} & \mathcal{L}_{11} + \mathcal{L}_{\tilde{5}\tilde{8}} - \mathcal{L}_{1\tilde{8}} - \mathcal{L}_{\tilde{5}1} \\ \mathcal{L}_{11} + \mathcal{L}_{\tilde{8}\tilde{1}} - \mathcal{L}_{1\tilde{1}} - \mathcal{L}_{\tilde{8}1} & \mathcal{L}_{11} + \mathcal{L}_{\tilde{8}\tilde{5}} - \mathcal{L}_{1\tilde{5}} - \mathcal{L}_{\tilde{8}1} & \mathcal{L}_{11} + \mathcal{L}_{\tilde{8}\tilde{8}} - \mathcal{L}_{1\tilde{8}} - \mathcal{L}_{\tilde{8}1} \end{pmatrix}$$

where the "tilde" indices refer to the labelled ground returns in Figure 2 (i. e., $\mathcal{L}_{\tilde{5}\tilde{5}}$ denotes the partial self inductance of G_5).

Both the partial inductance and return-limited loop inductance matrices are symmetric and positive semi-definite. Since the interaction regions do not couple in any way, when ordered appropriately, the return-limited loop inductance matrix of the entire chip is block diagonal, with each block representing a single interaction region. One can further collapse the circuit of Figure 3(a) into the circuit of Figure 3(b) represented by the *return-limited inductance matrix* $\mathbf{L}$. Let n be the number of return-limited loop inductances in the interaction region, and let m be the number of return-limited inductances in the interaction region, as determined by the number of signal segments. We then form the $n \times m$ matrix $\mathbf{B}$, where the ith column of $\mathbf{B}$ is all zero except for ones in the rows corresponding to the return-limited loop inductances associated with the given return-limited inductance. One finds

$$s\mathbf{I} = \mathbf{B}^T \mathbf{L}'^{-1} \mathbf{B} \mathbf{V}$$

where $\mathbf{I}$ are the currents shown in Figure 3(a). The return-limited inductance matrix of the interaction region is then given by:

$$\mathbf{L} = (\mathbf{B}^T \mathbf{L}'^{-1} \mathbf{B})^{-1}$$

The main computational cost here is two LU factorizations, one of the $\mathbf{L}'$ matrix associated with each horizontal and vertical interaction region and the second to find the final return-limited inductance matrix for the given interaction region (the $\mathbf{L}$ matrix).

2.4 Power-grid extraction

The return-limited inductance extraction of the power grid follows similarly to the signal-line extraction. In this case, we ignore the signal lines entirely. In creating a worst-case estimate, signal lines cannot be assumed to be decoupled to the power grid and, therefore, available as high-frequency current returns. The halo rules for power and ground lines can be stated as follows:

- Horizontal and vertical power and ground segments are treated independently since they cannot couple to each other.
- Inductive coupling between two horizontal (vertical) power and ground segments is nonzero if and only if it is possible to connect the two segments by a path which does not cross the horizontal (vertical) halo of any *other* power or ground segment.

There are no blocking rules since signal segments are not involved.

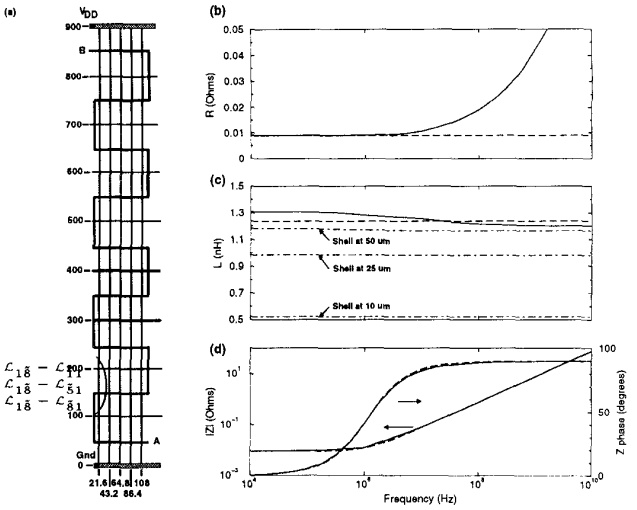

Figure 4: (a) Example geometry. (b) Resistance at the near end as a function of frequency. (c) Inductance at the near end as a function of frequency. (d) Magnitude and phase of the impedance at the near end as a function of frequency.

3 Results

To further study the efficiency and accuracy of return-limited inductance extraction, we have implemented a prototype resistance and inductance extractor which can handle complex three-dimensional geometries. The extractor takes a shapes text file input with established connectivity, similar in format to the Fasthenry input file[7]. For the examples presented in this section, we work with a hypothetical five-level-metal copper process. M5 is $2\mu m$ thick, while the other metal layers are $0.9\mu m$ thick. With this technology, we apply our extraction tool to signal lines routed within the gridded power and ground interconnect network shown in Figure 4(a). An image is chosen in which M5 and M3 run horizontally, while M4 and M2 run vertically. As is consistent with a typical distribution[8], power and ground are distributed in a grid spaced $100\mu m$ apart vertically and $21.6\mu m$ apart horizontally. V_{DD} is carried on the M3 lines running horizontally at $y = 900\mu m$, $700\mu m$, $500\mu m$, $300\mu m$, and $100\mu m$. Ground is routed on the remaining M3 horizontal lines. Similarly, V_{DD} is carried on the M2 and M4 lines running vertically at $x = 21.6\mu m$, $64.8\mu m$, and $108\mu m$. Ground is routed on the remaining alternate M2 and M4 vertical lines. For the results presented here, we assume ideal decoupling (i. e., a short) at each point the M2-M4 grid crosses the M3 grid. We choose shorts rather than actual decoupling capacitances to facilitate straightforward comparison with Fasthenry.

In Figure 4 we show a single signal line, beginning at A and ending at B, routed within the power and ground distribution network. Node B is shorted to the power grid at the far end. We compare the results of Fasthenry (solid curves) with the return-limited inductance extraction (dashed curves) for the impedance seen between A and the adjacent near-end power grid. Figure 5 shows a more complex example in which two signal lines, one beginning at A and ending at D, and the other beginning at C and ending at B, are routed within the same power and ground distribution network. Nodes B and D are shorted at the far end to the power grid.

Two competing trends in the approximation determine

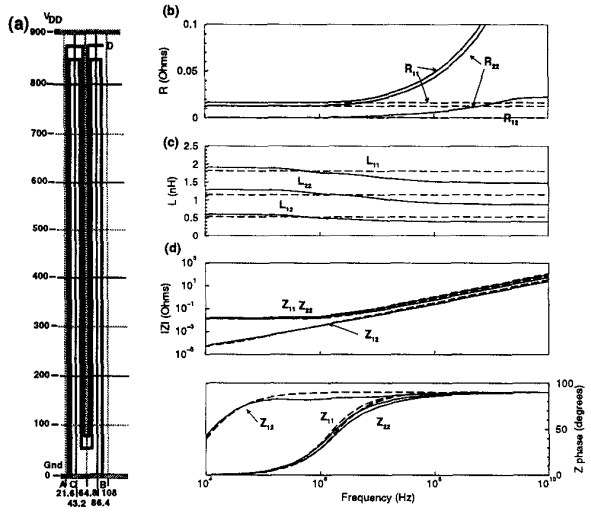

Figure 5: (a) Example geometry. (b) Resistance at the near end as a function of frequency. R_{11} denotes the resistance at port A, R_{22} denotes the resistance at port B, while R_{12} denotes the coupling resistance between the ports. (c) Inductance at the near end as a function of frequency. L_{11} denotes the inductance at port A, L_{22} denotes the inductance at port B, while L_{12} denotes the coupling inductance between the ports. (d) Magnitude and phase of the impedance at the near end as a function of frequency. Z_{11} denotes the impedance at port A, Z_{12} denotes the impedance at port B, while Z_{12} denotes the coupling impedance between the ports.

the accuracy of the return-limited inductance value. The failure to consider more distant parallel return paths (which results in a slight overestimation of the inductance) is counterbalanced by the discarding of mutual inductances between segments that fall into different interaction regions (which results in a slight underestimation of the inductance). Because the frequency-dependence of the resistance does not become important until frequencies at which the impedance is dominated by the inductance, the skin effect can be safely ignored in nearly all cases. As a result, we get good agreement in the impedance (magnitude and phase) between the return-limited values and the Fasthenry result, as shown in Figure 4(d) and 5(d).

In Table 1, we compare the number of nonzero elements in the inductance matrix for the return-limited inductance results against the full partial inductance extraction and various applications of the shell sparsification algorithm[3, 4]. The number of nonzero elements directly translates into the efficiency of simulation and analysis of the resulting network. The full partial inductance extraction is shown under two conditions, one in which single filament is used to represent each segment and a second in which each segment cross section is broken into 20 filaments. The Fasthenry results of Figures 4 and 5 were calculated using the 20 filament decomposition to properly model skin-effect to $10GHz$. In Figure 4, we also shown the inductance as a function of frequency (dash-dotted curves) for each of the shell approximations. The result at $r_o = 10\mu m$ is far worse than the return-limited inductance result even through it requires 859 nonzero elements as compared to only 178 for the return-limited case. Good accuracy is not achieved until $r_o = 50\mu m$, requiring over 3000 elements. Furthermore, we have no way of knowing *a priori* what value of r_o is required to achieve tolerable accuracy. The full partial inductance matrix, even with only single filament decomposition, has more than 80,000 nonzero elements.

	1 filament per segment	20 filaments per segment	Shell at $50\mu m$	Shell at $25\mu m$	Shell at $10\mu m$	Return-limited
Fig. 4	82137	very big	3075	1427	859	178
Fig. 5	117069	very big	6103	2513	681	597

Table 1: Number of nonzero elements in the inductance matrix. "very big" indicates that the number of nonzero elements was too large to even practically count (> 50 million).

4 Conclusions

In this paper, we have presented an approximate inductance extraction approach that can be practically applied for full-chip extraction of complex integrated circuits. The approach results in a positive-semi-definite inductance matrix and achieves sparsity by breaking inductive couplings between the power-ground lines and signal lines and by using a set of geometry-based decomposition rules to discard insignificant coupling interactions between signal lines and between power and ground lines. By taking advantage of the fail-safe availability of power-ground lines as high-frequency current returns while preserving inductive coupling between signal lines, we have demonstrated remarkable efficiency and accuracy over other sparsification techniques.

References

[1] A. E. Ruehli. Inductance calculations in a complex integrated circuit environment. *IBM Journal of Research and Development*, 16(5):470 – 481, 1972.

[2] Zhijiang He, Mustafa Celik, and Lawrence Pileggi. SPIE: Sparse Partial Inductance Extraction. In *Proceedings of the 34th Design Automation Conference*, pages 137–140, Anaheim, California, June 1997.

[3] B. Krauter and L. T. Pileggi. Generating sparse partial inductance matrices with guaranteed stability. In *Proceedings of the International Conference on Computer-Aided Design*, pages 45 – 52, 1995.

[4] Mattan Kaman, Byron Krauter, Joel Phillips, Lawrence T. Pileggi, and Jacob White. Two optimizations to accelerated method-of-moments algorithms for signal integrity analysis of complicated 3-d packages. In *Proceedings of the IEEE 4th Topical Meeting on Electrical Performance of Electronic Packaging*, pages 213 – 216, Portland, OR, October 1995.

[5] K. L. Shepard and Zhong Tian. Return-limited inductances: a practical approach to on-chp inductance extraction. *submitted to IEEE Transactions on CAD.*

[6] P. H. Xiao, E. Charbon, A. Sangiovanni-Vincentelli, T. van Duzer, and S. R. Whitely. INDEX: An Inductane Extractor for Superconducting Circuits. *IEEE Transactions on Applied Superconductivity*, 3(1):2629–2632, March 1993.

[7] M. Kamon, M. J. Tsuk, and J. White. FastHenry, a Multipole-Accelerated 3-D Inductance Extraction Program. In 30^{th} *ACM/IEEE Design Automation Conference*, pages 678–683, Dallas, June 1993.

[8] K. L. Shepard, S. Carey, E. Cho, B. Curran, R. Hatch, D. Hoffman, S. McCabe, G. Northrop, and R. Seigler. Design Methodology for the G4 S/390 Microprocessors. *IBM Journal of Research and Development*, 21(4/5):515 – 548, 1997.