# Circuit-Aware On-chip Inductance Extraction

*Haitian Hu and Sachin S. Sapatnekar*
Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, USA

**Abstract**

In this paper we propose a practical approach for on-chip inductance extraction. This approach differs from previous methods in that it uses circuit characteristics to obtain a sparse, stable and symmetric inductance matrix, using the concept of resistance dominant and inductance dominant lines. Experimental results show that only the important inductance terms related to strong inductance couplings are included in the sparsified inductance matrix to ensure a specified predefined accuracy. For a good design, the sparsification can reach 95% by setting an acceptable delay error of 10% and oscillation magnitude error of 2%.

## 1. Introduction

As technology shrinks, inductive effects become more prominent, particularly in the uppermost metal layers, as lines become longer and more closely packed. With next generation technologies projected to use low-K dielectrics, capacitive effects will be diminished and on-chip inductances play an even more significant role. Inductive effects have become important in determining power supply integrity, timing and noise analysis, especially for global clock networks, signal buses and supply grids in upper several layers for high-performance microprocessors. There are two types of lines that are impacted by inductive effects:

- *switching lines*, i.e., clock nets and signal nets
- *supply lines*, i.e., $V_{dd}$ and ground lines

It is important to integrate the analysis of switching and supply lines since (a) the supply lines act as return paths for switching lines, and their distribution affects the signals on supply lines, and (b) the magnitude of the return currents impacts the supply lines' integrity.

The concept of inductance is defined over a current loop. However, return paths for the loop are difficult to predict as they are impacted by factors such as RC parasitics, pad locations, the operating frequency and the switching patterns on neighboring lines. The traditional method for representing a complex multiconductor topology without predetermined current return paths is to use the PEEC model [1]. This uses the concept of partial inductance associated with line segments in the wire. The formulae for partial self and mutual inductance are available in [2]. The partial inductance can be defined as the inductance of a line segment which is in the magnetic field of its own current and/or other line segments' currents, which forms a loop with infinity. For two line segments *k* and *m*, the partial mutual inductance[1] is given by:

$$L_{km} = \frac{1}{I_m a_k}\left(\int_{a_k}\int_{l_k} \vec{A}_{km} \bullet d\vec{l}_k da_k\right) \quad (1)$$

where $\underline{a_k}$ is the cross section of segment *k*, $\vec{l}_k$ is the length vector along segment *k* and $\vec{A}_{km}$ is the magnetic vector potential along segment *k* due to the current $I_m$ in segment *m*, given by:

$$\vec{A}_{km} = \frac{m_0}{4pa_m}\left(\int_{a_m}\int_{l_m} \frac{I_m}{r_{km}} d\vec{l}_m da_m\right) \quad (2)$$

where $r_{km}$ is the distance between two points on segment *k* and *m*.

However, the blind use of this method can result in a dense inductance matrix that causes a high computational overhead for a simulator. Although many entries in this matrix are small and have negligible effects, discarding them may cause the resulting inductance matrix to no longer remain positive semidefinite. The shift and truncate method [3] finds an approximate sparse positive semidefinite inductance matrix, and operates by assuming that the current return of each line segment is distributed on a shell of radius $R_0$. The drawback of this approach is that the value of $R_0$ must be constant for the whole chip, and the work in [4] dynamically determines this global value. An alternative approach that uses return-limited inductances [5] is a shape-based method to sparsify the inductance matrix, using "halo rules." While this method is a good first order approach, it considers supply lines to be perfect conductors: this is not always a valid approximation since these return paths can influence the response of the switching lines. Moreover, it is possible to observe cases where mutual inductance with the non-nearest supply line can affect the waveform on a switching line. Another recent approach [6] introduces a heuristic sparsification technique based on a simple partition of the circuit topology, and neglects mutual inductances between partitions.

In this paper, we propose an alternative method to sparsify the inductance matrix resulting from the PEEC model of the whole circuit. In order to accurately estimate the current return paths and inductance effect, the comprehensive PEEC model includes: interconnect partial resistance, capacitance and inductance, signal line drivers and receivers, supply pad resistance, capacitance, inductance and location, via resistance. Unlike most previous techniques that largely neglect the circuit characteristics, our procedure is "circuit-aware" and explicitly examines the RLC's and classifies the switching lines into two categories:

- ♦ *inductance dominant lines (ID lines)*: a self/mutual inductance for the line strongly affects a waveform in the circuit.
- ♦ *resistance dominant lines (RD lines)*: inductive effects are partially or completely damped out by the driver resistance, so that both the self and mutual inductances associated with this line have a weak impact on all waveform in the circuit.

Subsequently, only inductance dominant lines and lines that are strongly influenced by the inductance dominant lines (including the nearby supply lines and some of the resistance dominant lines) are included in the sparsified inductance matrix. A worst case switching pattern is used in determining the sparsified inductance matrix so that a worst case inductance matrix[2] can be found that can safely be used under other input switching patterns. Finally, the inductance matrix is further sparsifed by the shift and truncate method [3].

## 2. Circuit Model

The circuit model used in this work includes supply grids, signal buses and clock nets on all the metal layers. Pads are located on the top layer to connect supply grid to the external supply supply. Each signal bus and clock net is connected to drivers and receivers. The specifics of the models are detailed below.

*Line models*: Each line is divided into line segments using an RLC model for each segment. The model also includes mutual inductance

This work was supported in part by the SRC under contract 99-TJ-714 and by the NSF under award CCR-9800992.

[1] Setting *k* = *m* yields the partial self-inductance for the line segment.

[2] The term "worst case" her only refers to the fact that under specific switching patterns, further sparsification of the inductance matrix is possible.

between any two non-perpendicular line segments, capacitance between any two adjacent line segments, via resistance, pad resistance, pad self-inductance and pad-to-ground capacitance.

The resistance of any line segment is calculated as $R = R_s L/W$; $R_s$, $L$ and $W$ are, respectively, the sheet resistance, length and width. The inductance of any line segment is calculated by Geometrical Mean Distance (GMD) formulae in [2]. The line-to-ground and line-to-line capacitances are calculated by Chern's model [8].

*Driver and receiver models*: The drivers are modeled by a voltage source, an effective driver resistance and an output capacitance. The receivers are modeled as a load capacitance connected to ground. The effective resistance of the driver is inversely proportional to the size, and the output capacitance of the driver and the load capacitance of the receiver are each proportional to the size of the corresponding entity, with differing constants of proportionality.

*Pad model*: Pads are located on top metal layer and are modeled by a pad resistance, pad self-inductance and a pad capacitance.

## 3. Proposed sparsification method

Until recently, when on-chip inductances were insignificant, RC modeling was adequate for all on-chip lines since the RC elements overwhelmed any inductive coupling. A very simple but important observation for sparsifying the inductance matrix is based on this idea: the value of the resistances and capacitances in the circuit determine whether a line has a significant inductance or not. To ascertain if the inductance is significant in determining the response of a switching line, the traditional method is to compare the characteristic impedance of the line with the line resistance and the effective strength of the driver. For a lossy transmission line with quasi-TEM propagation in it, inductance is important if

$$R_d, R_l < Z_0 = \sqrt{\frac{L}{C}} \qquad (3)$$

where $R_d$ is the effective resistance of the driver, $R_l$ is the total resistance of the line, $Z_0$ is the characteristic impedance of the line, and $L$ and $C$ are, respectively, the self inductance and the line-to-ground capacitance per unit length of the line.

Unfortunately, equation (3) cannot be used directly for on-chip interconnect because it does not include mutual inductances and it is not applied to the non-uniform transmission lines which characterize on-chip inductance. However, if we assume uniform interconnect structure in all cases, (3) does reflect whether a line is ID or not in the absence of mutual inductances; if mutual inductances are added, inductive effects can only increase. If a line does not satisfy the condition, it is still possible that for it to be ID because the mutual inductance of two line segments along the line may be significant.

To demarcate ID lines from RD lines, we use a relative criterion to define ID lines, called the *ID criterion*, as follows:

1. A line that satisfies the criterion in (3) is ID.
2. A line that does not satisfy (3) is ID if the behavior of the output waveform in the presence of inductances is significantly different from the waveform when a pure RC model is used and inductances are ignored. Specifically, if the percentage variation in the oscillation magnitude and delay of the output response are larger than $\varepsilon$ and $\delta$, respectively, the line is ID.

RD lines include all those lines that are not inductance dominant. In this way, we separate all the on-chip lines into three categories: ID switching lines, RD switching lines and supply lines.

We use these ideas of RD and ID lines to identify *clusters*. Formally, we define a cluster as a group of on-chip interconnects for which mutual inductances must be calculated between any pair of line segments in this group. A cluster corresponds to a full inductance sub-matrix and there is no mutual inductance between line segments within and without a cluster. Finally, we use the shift-and-truncate method to further sparsify this full sub-matrix. *Therefore, by construction, the resulting sparse inductance matrix is positive semidefinite.*

### 3.1 Foundations for the algorithm

We have performed a series of experiments to create a set of foundations on which our extraction procedure is based. In each of these experiments, we have two parallel lines of equal length that drive signals in the same direction. To see the maximum inductive effects, both lines are simultaneously excited by rising inputs.

The criterion of equation (3) can be used to determine whether a line is ID on the basis of its self-inductance. Our objective here is to develop criteria to draw similar conclusions based on mutual inductances. Therefore, we have performed experiments that compare the effects of including mutual inductances between the two lines, as compared to excluding them[3]. In all cases, the line segment resistances, the driver resistance, and the line-to-line and line-to-ground capacitances are considered in the simulations.

The RLC parameters that are used correspond to a 0.1μm technology and are extrapolated from [9]. These parameters are summarized as follows:

Minimum line width= 0.1 μm
Minimum line spacing=0.14 μm
Driver resistance for minimum buffer size=23.9 KΩ
Driver input capacitance for minimum buffer size=0.07 fF
Supply voltage $V_{dd}$ = 1.05 V

#### 3.1.1 Coupling inductance between switching lines

The width and spacing of all lines are set to twice or ten times of the minimum value of each technology. The driver resistance is set to 5x, 50x, 100x and 200x of the minimum driver size. The change in the delay and the oscillation magnitude before and after considering the mutual inductance between two lines are compared.

It is observed that, as expected, for small drivers, the large driver resistance causes the behavior of the line to be RD. The smaller the resistance of the driver, the more likely it is that the line is ID. Moreover, the effect of its mutual inductance with the other line becomes more significant when that line is not terminated by a medium or large driver; for very small drivers, all inductive effects are damped out. For example, the response of switching line driven by a 5x driver does not change even if the mutual inductance effect of the other line, driven by a 200x driver is considered. However, if the driver size is changed from 5x to 50x or larger, it is seen that the mutual inductances can perceptibly affect the waveform, both in terms of the delay and the oscillation magnitude of the overshoot.

From our simulation results (not shown here due to space limitations), we can infer the first set of foundations:

***Foundation 1***: ID lines have strong mutual inductance effects on other ID lines. The more likely it is that a switching line is an ID line, the more significant the effect is.

***Foundation 2***: RD lines, especially highly RD lines, have very little mutual inductance effects on other lines.

***Foundation 3***: ID lines may have mutual inductance effects on RD lines, depending on whether the RD lines are highly RD[4] or not.

---

[3] Since we work with partial inductances, we divide each line into multiple segments. The self-inductance of a line consists of the self-inductance of each segment and mutual inductances between segments of the same line [7].

[4] Recall that the differentiation between RD and ID lines is a gray area.

### 3.1.2. Coupling between switching lines and supply lines

In these experiments, the experimental setup is the same as before, except that there is only one switching line and the other line is a supply line with two ends connected to perfect ground through pads. The driver sizes are changed from 5x to 250x, and we perform three sets of experiments:

1. using an RC model for both lines.
2. using a PEEC model for the switching line, but no mutual inductance between the switching line and the supply line.
3. using a PEEC model for the entire circuit.

It is observed that the RC model fails as the driver size is increased since inductive effects become prominent. The inclusion of mutual inductances between the switching line and the supply line permits a nearby current return path for the switching lines and greatly reduces the inductance effect of the ID line, both in terms of delay and oscillations associated with the overshoot. However, if the switching lines are highly RD lines (for example, a driver size of 5x), supply lines have little effect on these parameters.

From our simulation results (again, not shown here due to space limitations), we can infer the second set of foundations:

***Foundation 4***: Supply lines have significant mutual inductance effect on ID lines nearby, which greatly reduces the inductance effect on ID lines.

***Foundation 5***: Supply lines do not have significant mutual inductance effects on RD lines.

### 3.2 Formation of clusters

Based on the above foundations, we separate the mutual inductance interaction into two types:

1. Strong mutual inductance interactions between
   - ID lines and nearby ID lines
   - ID lines and nearby supply lines
2. Weak mutual inductance interactions between
   - ID lines and nearby RD lines
   - RD lines and nearby supply lines
   - RD lines and nearby RD lines
   - supply lines and nearby supply lines

Since strong mutual inductance interactions are the most important, our algorithm first identifies strong mutual inductance terms and forms clusters, and then adds weak mutual inductance terms into those clusters if necessary. In order to reduce as many the number of mutual inductance terms as possible, we always find the supply return paths for a cluster before we determine which other clusters it will affect. The reason is that if there is a supply return path nearby, the magnetic field vector $\vec{B}$ and magnetic vector potential $\vec{A}$ of the aggressor cluster are strongly weakened by the magnetic field induced by supply return path. The magnetic vector potential drop along the victim line, as well as the inductance effect of aggressor cluster on the victim line, is also greatly reduced. However, if we consider the mutual inductance effect between the aggressor cluster and victim line before we reduce the magnetic field of aggressor cluster by its supply return path, it is very possible that we may overestimate the inductance effect of the aggressor cluster.

### 3.3 Our procedure for on-chip inductance extraction

We proceed by selectively including a new set of inductive effects in each iterative step. Simulations are carried out before and after new mutual inductance is added into the circuit to check if a switching line is an ID line or not, to check if a supply line is the return path and to check if two clusters should be grouped into one new cluster. To excite the worst case, all lines are made to switch simultaneously in such a way that the currents are carried in the same direction in order to enable the largest $L_{jk}\ dI_k/dt$ drop on the lines. In our experiments, we use PRIMA to perform the simulations efficiently.

The algorithm is summarized as follows:

1. Use an RC model for all lines and simulate the circuit.
2. Set all the switching lines that satisfy condition (3) as ID lines.
3. For each other switching line, use a model that includes self-inductances and mutual inductances between segments on the same line. Re-simulate the circuit and compare its behavior with an all-RC model to test if the line is ID, using the ID criterion listed at the beginning of Section 3.
4. For each ID line, find its supply return paths starting from the nearest supply lines and moving out towards lines farther away until at least one ID line is encountered in each direction. Return paths are identified by including mutual inductances with the supply lines and performing a simulation to test the ID criterion. Group each ID line with its return paths to form a cluster. If two clusters have the same supply return paths, group them into one cluster.
5. Check if any two of these ID clusters should be grouped into one larger cluster. To perform this test, group adjacent clusters together by including mutual inductances between them and test the ID criterion. At the end of this process, if any two of the newly formed clusters have common lines, group them into one cluster.
6. Test, as in step 4, for additional return paths for each cluster, starting form the nearest supply lines. As stated at the end of Section 3.1.2, creating larger clusters has the effect of allowing a larger set of return paths.
7. Each RD line can be thought of as a cluster with only one line in it. Test the ID criterion to see if any singleton cluster (RD line) should be grouped with an ID cluster. As before, this test is performed by grouping the clusters if they show strong mutual inductance effects. Repeat steps 6 and 7 until no new cluster is formed.

## 4. Experimental results

We have carried out a set of experiments on a 0.1μm technology to study the effect of dedicated power/ground lines on reducing inductive effects and to compare with our algorithm with the shift and truncate method. The circuit topologies used in this section correspond to the top three metal layers with wide and long switching lines routing on the uppermost metal layer, M5.

The first experiment uses a configuration that has 10 signal buses on M5 with 8 power/ground grid lines on each side. There are 7 power/ground lines in the orthogonal direction on M4 and 16 power/ground lines on M3 in the same direction. The supply grid is connected by vias with a specified via resistance. The power/ground grid lines have width 12.0μm and spacing 108.0μm, while the width and spacing of signal buses are both 0.9μm. The thickness of metal layers and oxide layers are 0.5μm and 0.6μm respectively. There are 16 pads located on the M5 with spacings of 240μm. The resistance, capacitance and inductance of the pad are, respectively, 0.0003Ω, 390fF and 0.15nH respectively. We assume the input rise times to be 10ps. The signal buses are driven by different size of drivers. The driver size are, respectively, 5×, 60×, 100×, 200×, 40×, 220×, 200×, 70×, 10×, 5× times of minimum buffer size for the ten signal lines in sequence. We use a $V_{dd}$ of 1.05V.

Since there are no dedicated power/ground lines near the signal lines (some of which are ID) and the spacing of grid power/ground

lines is rather large, the magnetic field induced by the ID lines is not reduced effectively by supply return paths, so that even the highly RD lines are influenced by mutual inductances and must be included into the same cluster with ID lines. The application of our algorithm shows that there is only one cluster formed in this circuit, including all the signal lines as well as the nearest four supply lines on the uppermost layer, four middle lines on M3 and also part of lines on M4. Although the lines on M4 are perpendicular to those on M5 and M3 and do not directly contribute to the response of signal lines, the mutual inductance of line segments on M4 enlarges the oscillation on the supply lines which worsens the integrity on the supply lines. Fig. 1 shows a comparison of HSPICE waveforms from the fully coupled (exact) inductive system and our sparsified one. The errors in oscillation and delay are less than $\varepsilon = 20$mv and $\delta = 10\%$.

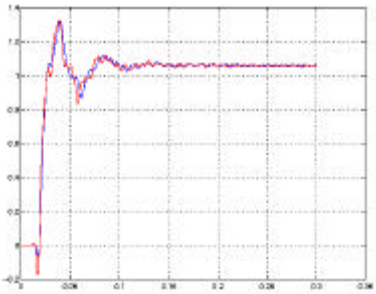

Fig. 1. Response of the signal line using the exact and sparsified L matrix.

In order to further reduce the inductance effect of ID lines, one dedicated power/ground line is added next to each ID line, with the same width and spacing as those of signal buses. These dedicated power/ground buses are connected to the power/ground grid. The oscillation magnitude of ID line is reduced significantly, by 250mV, and delay is reduced by 15%. In this circuit, two smaller clusters are formed, with the lines driven by 60×, 100×, 200× drivers in one cluster and the 40×, 220×, 200×, 70× driven lines in another. The highly RD lines, such as lines with 5× or 10× drivers can now be modeled with an RC model.

Both the shift and truncate method and our algorithm are applied to two circuits with 16 pads on each circuit. Circuit1 has 4 signal lines on M5 and 8 power/ground grid lines on each side, 7 grid lines on M4 and 16 grid lines on M3. There is no dedicated power/ground line near the signal buses. The width and spacing of grid lines are 6.0μm and 54.0 μm, respectively, and the width and spacing of signal buses are each 0.9μm. The thickness of metal layers and oxide layers are the same as before. The driver sizes for the signal buses are 220×, 150×, 200×, 5× times the minimum buffer size for all the signal buses in sequence. Circuit2 is an optimized version of circuit 1, with one dedicated power/ground lines on each side of the four signal buses. The width and spacing of these dedicated power/ground lines are, again, the same as in the previous experiment. Table 1 compares the results of our sparsification method as compared to only using the shift-and-truncate method for the two circuits. The value of $R_0$ for the shift-and-truncate method is determined using the technique in [4].

Circuit1 can be represented by 9 clusters, with the largest cluster including all the signal lines, 4 grid power/ground lines on each side

| | Our method | Shift and truncate method |
|---|---|---|
| Circuit 1 | **79%** | **< 68%** |
| Circuit 2 | **95.25%** | **< 90%** |

Table 1. Comparison of sparsification of two methods on test circuits.

on the M5 layer and the 8 middle grid power/ground lines on M3. Since the mutual inductance between two lines on M4 is smaller than the mutual inductance between two nearest line segments on these lines, only the largest inductance terms are kept for lines on M4. Here again we see that without dedicated power/ground lines, the inductance effects of ID lines are quite strong and the inductance matrix cannot be significantly sparsified. The number of clusters for Circut2 is more than that in Circuit 1 and each cluster is significantly smaller. The RD line (5x driver) can be analyzed under an RC model. The largest cluster includes the other three ID lines and their dedicated power/ground lines. The mutual inductances of two nearest line segments of the nearest supply grid lines on M5 are also included into the final inductance matrix representation.

## 5. Conclusion

A circuit-aware sparsification methodology for fully coupled PEEC inductance representation is proposed in this paper by analyzing the circuit characteristics and clustering the inductances according to their relative importance to the whole circuit. In this algorithm, all the signal lines are roughly defined as inductance dominant (ID) lines or resistance dominant (RD) lines. Strong couplings are taken care of first and weak couplings are then added to the clusters. Our experimental results show the effectiveness of this sparsification method, which helps to determine current return paths for a design and identify the most critical inductance terms for optimization.

## References


[1] A. E. Ruehli, *et al.*, "Inductance Calculations in a Coupled Integrated Circuit Environment," *IBM Journal of Research and Development*, pp. 470-481, Vol. 16, No. 5, September 1972.

[2] F. W. Grover, *Inductance calculations; Working Formulas and Tables*, Dover Publications, New York, NY, 1946.

[3] B. Krauter and L. T. Pileggi, "Generating Sparse Inductance Matrices with Guaranteed Stability." *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pp. 45-52, Nov. 1995.

[4] Z. He, *et al.*, "SPIE: Sparse Partial Inductance Extraction," *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 137-140, 1997.

[5] K. L. Shepard and Z. Tan, "Return-limited Inductances: A Practical Approach to On-Chip Inductance Extraction," *Proceedings of the IEEE Custom Integrated Circuits Conference*, pp. 453-456, 1999.

[6] K. Gala *et al.*, "On Chip Inductance Modeling and Analysis," *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 63-68, 2000.

[7] A. E. Ruehli, *Circuit Analysis, Simulation and Design*, Vol. 3, Part 2, North-Holland, Amsterdam, The Netherlands, 1987.

[8] J. H. Chern, *et al.*, "Multilevel Metal Capacitance Models for CAD Design Synthesis Systems," *IEEE Electron Device Letters*, Vol. 13, no. 1, pp. 32-34, Jan. 1992.

[9] Cong J., "Challenges and Opportunities for design Innovations in Nanometer Technologies," SRC Design Science Concept Paper, 1997.