

Inductance 101: Analysis and Design Issues

Kaushik Gala, David Blaauw, Junfeng Wang, Vladimir Zolotov, Min Zhao

Motorola Inc., Austin TX 78729

kaushik.gala@motorola.com

Abstract

With operating frequencies approaching the gigahertz range, inductance is becoming an increasingly important consideration in the design and analysis of on-chip interconnect. In this paper, we give a tutorial overview of the analysis and design issues related to on-chip inductance effects. We explain the complexity of the current flow in VLSI circuits. We discuss the applicability of the PEEC approach in a detailed circuit model of the signal and power grid interconnect, switching devices, power pads and the package. Further, we explain techniques that can be used to speed-up simulation of the large PEEC model. We then discuss a simplified model that uses the so-called loop inductance approach, and compare it with the detailed model. We present experimental results, obtained from simulations of industrial circuits, for both the PEEC and loop models. We also cover design techniques that can help tackle the on-chip inductance issues.

1 Introduction

Inductance effects in on-chip interconnect structures have become increasingly significant[1] due to longer metal interconnects, reductions in wire resistance (as a result of copper interconnects and wider upper-layer metal lines) and higher clock frequencies. These effects are particularly significant for global interconnect lines such as those in clock distribution networks, signal buses, and power grids for high-performance microprocessors. On-chip inductance impacts these in terms of delay variations, degradation of signal integrity due to overshoots, undershoots and oscillations, aggravation of signal crosstalk, and increased power grid noise.

The main difficulty in the extraction and simulation of on-chip inductance is the fact that inductance is a function of a closed current loop. Therefore, it is required that both the current through a signal net and the return currents be considered simultaneously instead of being analyzed in isolation. Most of the return currents flow through the power grid, but some can also flow in other signal nets or through the substrate. The current distribution in the entire circuit, including the grid, must be known in order to obtain a correct estimate of loop inductance.

However, actual chip topologies consist of complex power grid and signal line structures, and current distribution depends on many elements, including device and interconnect decoupling capacitance, power grid resistance and inductance, pad locations, and operating frequency. Thus, the determination of current paths and, hence, the inductance is quite difficult, since it requires the accurate modeling and simulation of the complete signal net and power grid topology. In Section 2 of this paper, we examine the common current paths that arise when a signal is switched. These paths cannot be determined a priori in VLSI circuits, except in special cases like transmission line structures. Therefore, a circuit model that does not require the predetermination of current loops must be used.

The use of Partial Equivalent Elements Circuit (PEEC)[2] method, based on partial inductances, has been proposed to handle this issue[3]. We discuss in Section 3, how a detailed PEEC-based model [4] can be constructed using interconnect resistance, inductance and capacitance, device decoupling capacitances, quiescent activity in the grid and pad/package inductance.

Such a detailed PEEC model leads to a dense RLC circuit matrix requiring large SPICE simulation times. Hence, it must be combined with an acceleration technique like matrix sparsification or reduced order modeling. Section 4 discusses such techniques, ranging from simple truncation of the inductance matrix to complex shell, halo methods. We also present a methodology that combines block-diagonal sparsification with model order reduction.

In Section 5, we discuss a simplified approach which uses the loop inductance model[5] [6], as an alternative to the detailed PEEC model. The loop inductance and resistance are extracted by defining ports at the driving gate, and then solving the current distribution for an RL model of the circuit using tools such as FastHenry[7]. The extracted inductance and resistance are then combined with lumped capacitance to construct a netlist. The inductance extraction neglects the contribution of capacitance to current distribution. This can lead to inaccuracies, since the interconnect and device decoupling capacitances strongly affect current return paths. However, the simpler loop inductance model is faster to simulate, and can be used in a pre-layout estimation methodology.

We provide experimental results in Section 6, based on simulations of the global clock net of a high-performance microprocessor, using both the loop inductance and the detailed PEEC model.

Finally, Section 7 discusses some design guidelines that can be used to minimize inductive effects. In Section 8 we draw our conclusions.

2 Current flow in the circuit

Consider a typical on-chip topology including a signal net along with the power and ground grid in its vicinity. When the signal switches due to change in state of its driving gate, there are various currents that flow in the signal line and the neighboring grid. As shown in Figure 1, the currents I_1 and I_2 form loops throughout the package and grid decoupling capacitances, while I_3 forms a current loop from the driver output, through the grid and back.




FIGURE 1. Currents in Driver-Receiver-Grid topology

1. I_1 - Short circuit current flowing from power grid to ground while the gate is switching.
2. I_2 - Charging current, flowing from Vdd to ground, for the interconnect capacitance and gate capacitance between signal line and ground grid.
3. I_3 - Discharging current for the interconnect capacitance and gate capacitance between signal line and power grid.

The loops for the currents flowing from one grid to another are completed via the package and external supply, or through the decoupling capacitance between the power and ground grids.

3 PEEC circuit model

A typical circuit topology consists of two supply grids (power, ground) and signal lines laid out over multiple metal layers. The gates draw power from the lowest metal layer, while external power and ground are supplied via pads to the uppermost layer.

Figure 2 shows the proposed partial inductance based circuit model for the study of on-chip inductance effects. The circuit model shown above consists of

- Resistance, partial self-inductance and grounded capacitance (RLC- π) model for each metal segment.
- Mutual inductances between all pairs of parallel segments.
- Coupling capacitance between all pairs of adjacent lines.
- Via resistances between adjacent metal layers.
- Resistance and decoupling capacitance (to model non-switching gates)
- Time-varying current sources (to model switching gates)
- Pad resistances and inductances.

This model can also easily be extended to include substrate models, N-well capacitance and explicit decoupling capacitance.




FIGURE 2. Typical power grid topology and corresponding partial-inductance circuit model

Interconnect RLC

Each segment of the topology is modeled as an RLC- π circuit. The resistance is frequency independent and is computed as a function of geometry and sheet resistance. The ground and coupling capacitances for the interconnect are computed using Chern[8] models or commercial extraction tools. The partial self and mutual inductances are computed using analytical formulae [9], [10], [11]. These do not consider skin effect, hence very wide conductors must be split into narrower lines before computing inductance.

Device decoupling capacitance

During normal chip operation, approximately 10-20% of the gates switch while the remaining 80-90% remain static. The parasitic device capacitance of these non-switching gates results in a significant decoupling capacitance effect, which reduces IR-drop and changes current distribution by allowing current to jump from one grid to the other. This capacitance is estimated by using a statistical model that can be applied at both pre- and post-layout stages[12]. The estimation is done via small-signal analysis of representative circuit blocks. The capacitance values of one block can be easily translated to other circuit blocks based on the relative circuit sizes (total transistor widths) of the blocks.

Current sources

In addition to the signal of interest, other signals switch simultaneously. Those gates draw current from the power grid and inject it into the ground grid, causing voltage fluctuations and affecting current distribution. This effect is modeled by using time-varying current sources connected at random locations on the lowest metal layer. The current value changes with time during the simulation, to account for different parts of the chip switching at different times.

Pad/Package models

External power and ground are routed to a chip via package leads and pads. The parasitic inductances associated with the package must be modeled, since they affect on-chip behavior significantly. In the PEEC model, it is assumed that the package planes are ideal, since the voltage difference across these planes is typically of order of few mV. The package is modeled as a bar, including the pad and a via between the pad and package. More detailed package models can also be accommodated in the PEEC model.

4 Acceleration/Sparsification

Since the PEEC model includes mutual inductances between every pair of conductors, the resulting circuit matrix is very dense. As an example, large clock net topologies along with power grid can lead to the number of self inductances of the order of 100k and mutual inductance of the order of 10G. Hence the SPICE simulation is infeasible due to impractical time and memory requirements. This has been the main bottleneck in the use of PEEC models. The following techniques can be used to sparsify the partial inductance matrix or generate reduced order models.

Truncation

The simplest approach to sparsifying the inductance matrix is to discard all mutual coupling terms falling below a certain threshold. This translates to removing entries from the partial inductance matrix, thus making it sparse and faster to process. However, the resulting matrix can become non-positive definite, and the sparsified system becomes active and can generate energy. Since there is no guarantee on either the degree of sparsity or stability, truncation is not a feasible solution. This is in contrast to the capacitance matrix, which can be truncated without passivity problems.

Block-diagonal sparsification

Block-diagonal sparsification is a simple partitioning technique based on circuit topology, which guarantees the sparsified matrix to be positive definite. The topology is split into multiple sections as follows:

- Each section is stamped using self inductances and all the mutual inductances between elements of the same section.
- There exists no mutual coupling between elements from different sections.
- The signal bus of interest lies in the middle of the corresponding section, to capture the most significant inductive coupling between signal lines and power grid.
- Sections away from the signal of interest can be modeled as RC instead of RLC.

The section size depends on a trade-off required between run-time and accuracy. There is no direct measure of the accuracy lost as a result of removing mutual inductances. However, this approach is easy to use and guarantees passivity.

Shell method

As an alternative to simple truncation, one approach associates each segment with a distributed current return path out to a shell of some radius[13]. Segments with spacing more than this radius are assumed to have no inductive coupling. The inductance values of the segments within the radius are shifted to account for those entries that were dropped as a result of truncation. This shift-truncate method can guarantee to generate positive definite sparse approximations of the original matrix. However, this approach leads to complications in determining the value of the shell radius. An extension of this work[14] uses a moment-based algorithm to compute the shell radius.

Halo method

Another approach for limiting the inductive interaction is proposed by [15], which introduces return-limited inductances for sparsification and the use of “halos” to limit the number of mutual inductances. It is based on the assumption that the currents of signal lines return within the region enclosed by the nearest same-direction power-ground lines.

Hierarchical models

Hierarchical interconnect models[16] have been proposed to utilize the existing hierarchical nature of parasitic extractors. The concept of “global circuit node” is introduced to separate the electrical interaction into local and global interaction. These two are linked using controlled voltage/current sources.

‘K’ Matrix

A recent approach[17] defines a circuit matrix K , as the inverse of the partial inductance matrix L . K has a higher degree of locality and sparsity, similar to the capacitance matrix, and hence is amenable to sparsification and simulation. However, it requires inversion of the partial inductance matrix, and a special circuit simulator that can handle the K matrix.

Reduced-order modeling

Reduced-order models for the linear portion of the circuit can be combined with the gate models and simulated in SPICE [18] [19]. Model order reduction algorithms such as PRIMA[20] require matrix-vector multiplications, which are expensive for the fully-dense matrix of the PEEC model. However, reduced order models are very efficient in terms of simulation time and can match the original large model quite accurately. They are well suited to handle large topologies or longer simulation times and also provide a control over the accuracy via the order of the reduced system.

Combined technique (Block diagonal sparsification+PRIMA)

In [4], the acceleration in simulation of the detailed PEEC model is achieved by combining block-diagonal sparsification with reduced order modeling (PRIMA). The efficiency of PRIMA depends on the number of ports and the order of the reduced model. This technique is further improved by using the following.

- Active ports are differentiated from passive sinks in the model. A variant of the PRIMA algorithm is used to reduce the computation time by applying excitation sources only to the active ports, and not to the sinks.
- The MNA circuit matrix for the linear part of the model can be manipulated such that the matrix to be inverted is made positive-definite. This matrix can then be solved very fast using a direct solver based on the Cholesky method.
- Current models for switching devices, other than the signal net under consideration, are not included, to reduce the number of active ports. However, experiments show that the effect of these devices on signal integrity is not significant, and they contribute mainly in terms of IR-drop.

5 Loop inductance approach

Figure 3(a) shows a typical signal net and its neighboring ground grid. The loop inductance model defines a port at the driver side of the signal line and shorts the receiver side (which actually sees a capacitive load) to the local ground, since inductance extraction is performed independent of capacitance. Typically, an extraction tool such as FastHenry[7] is used to obtain the impedance over a frequency range, as shown in Figure 3(b). A netlist is then constructed with the resistance and loop inductance of the signal and ground grid, at one frequency, as shown in Figure 3(c).





FIGURE 3. Typical grid topology, R & L vs. frequency

Note that all the interconnect and load capacitance is modeled as a lumped capacitance at the receiver end of the signal interconnect. [5] proposes the construction of a ladder circuit, Figure 3(d), to model the frequency dependence of resistance and inductance. The loop impedance is extracted at two frequencies, and the parameters R_0 , L_0 , R_1 and L_1 used in the ladder circuit in Figure 3 are computed. The lumped RLC circuit representation can be improved by increasing the number of RLC- π segments. After the interconnect model is constructed, driver and receiver gates are connected and the complete circuit is simulated in SPICE.

6 Experimental results

To compare the detailed PEEC model with the loop inductance model, we construct and simulate both models for the same circuit topology. The topologies of interest to us are those having long and wide signal lines, since inductive effects dominate for such interconnect lines. Hence, we consider a global clock net in the presence of a multi-layer power grid. Figure 4 shows the simulation results for the two approaches.

The simulations were also performed on a RC PEEC model and as seen below, there are significant changes in delay and skew due to inductance. In the PEEC model, the delay increased by 10ps, compared with the RC model, while in the loop model, the delay increased by 3ps.



A - Driver input, B - Driver output, C - Receiver input




FIGURE 4. Top-level clock net: Loop vs. PEEC

The run-time shown in Table 1 includes the time for inductance extraction, matrix sparsification and model order reduction. The SPICE simulation for the reduced-order models took about 30sec in each case. Table 1 shows that the loop approach trades off lower accuracy for faster run-time, compared with the PEEC mode.

Table 1: Simulation of global clock net

	PEEC (RC)	PEEC (RLC)	LOOP (RLC)
Num. of R	220k	220k	3k
Num. of C	400k	400k	6k
Num. of L	-	190k	2k
# mutuals	-	4000k	-
Worst delay	86ps	131ps	116ps
Worst skew	9ps	19ps	12ps
Run-time	20 min.	45 min.	5 min.

7 Design issues and guidelines

Since inductance is directly related to interconnect length, short/medium length wires show resistive behavior, while long and wide wires exhibit inductive behavior. Inductance increases with the area of the current loop, hence inductive effects are reduced by the use of closer power/ground return paths. Inductive coupling noise and crosstalk can be minimized by controlling the overlap of the magnetic fields for parallel wires. Below, we have listed a number of common design techniques for reducing inductance.

Shielding

Capacitive coupling can be reduced by increasing the spacing between two metal lines. On the other hand, loop inductance can be reduced by sandwiching a signal line between ground return lines or guard traces. This forces the high-frequency current return paths to be close to the signal line, thus minimizing inductance.




FIGURE 5. Shielding

Dedicated ground planes

Dedicated ground planes or meshes in the layers above and below the signal line can be used to reduce inductance. Although they do not significantly lower the inductive effect at low frequencies, since resistance dominates and currents take wide return paths, at high frequencies, the ground planes provide excellent return paths for the signal current, thus reducing inductive behavior.




FIGURE 6. Ground Planes

Shield insertion and net ordering

Coupling noise can be reduced by simultaneously inserting shields and ordering nets, subject to constraints on area, and bounds on inductive and capacitive noise[21]. This optimization problem was found to be NP-hard and hence was solved by algorithms based on greedy approaches or simulated annealing.

Inter-digitated wires

Wider wires can be split into multiple thinner wires with shields in between. Such inter-digitizing reduces self-inductance, increases resistance and capacitance[22]. However, it increases the amount of metallization used for the interconnect.




FIGURE 7. Inter-digitated wires

Staggered inverter patterns

By using patterns of staggered inverters, the coupling capacitance and inductance effects can be reduced. The length of the overlapping portion between adjacent wires is reduced, thus reducing the amount of capacitive and inductive coupling. Also, the signal polarities alternate with each inverter, and hence the impact of the coupling tend to cancel out.




FIGURE 8. Staggered Inverters

Twisted-bundle layout structures

A recent approach[23] proposes a twisted-bundle layout structure for minimizing inductive coupling noise. The chip is divided into several routing regions and the routing of nets is reordered in each of these regions. This is done to create complementary and opposite current loops in the twisted bundle layout structure, such that the magnetic fluxes arising from any signal net within a twisted group cancel each other in the current loop of a net of interest.




FIGURE 9. Twisted-Bundle Layout

8 Conclusions

In this paper, we covered the analysis and design issues related to on-chip inductance effects. We explained the current flow in the power grid, which determines the loops and hence the effective inductance. A detailed circuit model consisting of resistance, partial inductance and distributed capacitance, device capacitance, currents, and pad/package model was examined. We presented acceleration approaches that can be combined with the large PEEC models to obtain a practical inductance analysis solution. The loop inductance approach was presented and compared with the PEEC approach. Finally, we discussed design guidelines that can minimize inductive effects.

References

- [1] Deutsch, A., et al, "When are Transmission-Line Effects Important for On-Chip Interconnections?," IEEE Transactions on Microwave Theory and Techniques, Oct. 1997, pp 1836-1847
- [2] Ruehli, A. E., "Inductance Calculations in a Complex Integrated Circuit Environment," IBM Journal of Research and Development, Sept. 1972, pp 470-481
- [3] He, L., et al, "An Efficient Inductance Modeling for On-chip Interconnects," CICC, May 1999, pp 457-460
- [4] Gala, K., et al, "On Chip Inductance Modeling and Analysis," DAC, June 2000, pp 63-68
- [5] Krauter, B., et al., "Layout Based Frequency Depended Inductance and Resistance Extraction for On-Chip Interconnect Timing Analysis," DAC, June 1998, pp 303-308
- [6] Sinha, A., et al, "Mesh-Structured On-Chip Power/Ground: Design for Minimum Inductance and Characterization for Fast R, L Extraction," CICC, May 1999, pp 461-464
- [7] Kamon, M., et al, "FASTHENRY: A Multipole-Accelerated 3-D Inductance Extraction Program," IEEE Transactions on MTT, Sept. 1994, pp 1750-1758
- [8] Chern, J. H., et al, "Multilevel metal capacitance models for CAD design synthesis systems," IEEE Electron Device Letters vol.13, no.1 pp 32-40
- [9] Sinclair, A. J., et al, "Analysis and Design of Transmission-Line Structures by means of the Geometric Mean Distance," IEEE Africon, Sept. 1996, pp 1062-1065
- [10] Grover, F. W., Inductance Calculations: Working Formulas and Tables, Dover Publications, New York, 1946
- [11] Hoer C., et al, "Exact Inductance Equations for Rectangular Conductors with Applications to More Complicated Geometries," Journal of Research of the National Bureau of Standards, April-June 1965, pp 127-137
- [12] Panda, R., et al., "Model and analysis for combined package and on-chip power grid simulation," ISLPED, July 2000, pp179-184
- [13] Krauter, B., et al., "Generating Sparse Partial Inductance Matrices with Guaranteed Stability," ICCAD, Nov. 1995, pp45-52
- [14] He Z., et al, "SPIE: Sparse Partial Inductance Extraction," DAC, June 1997, pp 137-140
- [15] Shepard, K. L., et al, "Return-Limited Inductances: A Practical Approach to On-Chip Inductance Extraction," IEEE Trans. CAD, Vol 19, April 2000, pp 425-436
- [16] Beattie, M., et al, "Hierarchical Interconnect Circuit Models," ICCAD, November 2000, pp 215-221
- [17] Devgan, A., et al, "How to Efficiently Capture On-Chip Inductance Effects: Introducing a New Circuit Element K," ICCAD, November 2000, pp 150-155
- [18] Beattie, M. W., et al, "IC Analyses Including Extracted Inductance Models," DAC, June 1999, pp 915-920
- [19] Krauter, B., et al., "Including Inductive Effects in Interconnect Timing Analysis," CICC, May 1999, pp 445-452
- [20] Odabasioglu A., et al, "PRIMA: Passive Reduced-order Interconnect Macromodeling Algorithm," ICCAD 1997, pp 58-65
- [21] He, L., et al, "Simultaneous shield insertion and net ordering for capacitive and inductive coupling minimization", ISPD, April 2000, pp 55-60
- [22] Massoud, Y., et al, "Layout Techniques for Minimizing On-Chip Interconnect Self-Inductance," DAC, June 1998, pp 566-571
- [23] Zhong, G., et al, "A Twisted-Bundle Layout Structure for Minimizing Inductive Coupling Noise," ICCAD, November 2000, pp 406-411