Inductance Model and Analysis Methodology for High-Speed On-Chip Interconnect

Kaushik Gala, David Blaauw, Vladimir Zolotov, Pravin M. Vaidya, and Anil Joshi

***Abstract*—With operating frequencies entering the multi-gigahertz range, inductance has become an important consideration in the design and analysis of on-chip interconnects. In this paper, we present an accurate and efficient inductance modeling and analysis methodology for high-performance interconnect. We determine the critical elements for a PEEC based model by analyzing the current flow in the power grid and signal interconnect. The proposed model includes distributed interconnect resistance, inductance and capacitance, device decoupling capacitances, quiescent switching currents in the grid, pad connections, and pad/package inductance. We propose an efficient methodology for extracting these elements, using statistical models for on-chip decoupling capacitance and switching currents. Simulation results show the importance of various elements for accurate inductance analysis. We also demonstrate the accuracy of the proposed model compared to the traditional loop-based inductance approach. Since the proposed model can consist of hundreds of thousands of *RLC* elements, and a fully dense mutual inductance matrix, we propose a number of acceleration techniques that enable efficient analysis of large interconnect structures. We use block-diagonal matrix sparsification that guarantees the passivity of the sparsified circuit while maintaining good accuracy. We also employ reduced-order modeling using the PRIMA algorithm. To accelerate the reduced order modeling, we introduce a new formulation of the moment calculation that employs a matrix reduction technique and also ensures that the resulting matrix is positive-definite. This allows the factorization of this matrix to be performed using efficient Cholesky factorization instead of the more time consuming LU decomposition, traditionally used. The proposed methods were implemented and used on the clock network of a gigahertz microprocessor. The combined sparsification and reduced order modeling approaches allow the analysis of a circuit model consisting of over 720 thousand *RLC* elements and 8 million mutual inductances in less than 1 hour. The presented analysis results emphasize the importance of inductance on the signal behavior in high performance processor designs and demonstrate the accuracy and efficiency of the proposed inductance model and analysis methodology.**

I. Introduction

With the advance of process technology, inductance effects in on-chip interconnect structures have become increasingly significant. Interconnect wire lengths are increasing, while the introduction of copper and wider upper-layer metal lines are reducing wire resistance. At the same time, higher frequency operation is leading to faster transition times. Combined, these trends are leading to a significant increase of inductance effects in on-chip interconnect [1], [2]. This is particularly the case for global interconnect lines, such as clock distribution networks, signal buses, and power grids for high-performance processors. On-chip inductance impacts these interconnects in a number of ways [3]. Signal wire delays are increased, causing possible performance problems. Signal transition times are decreased, which, together with inductive crosstalk can result in signal integrity related problems. Inductive signal overshoots result in high gate input voltages raising thin-oxide reliability concerns. Finally, inductance in the power grid can increase the power grid noise. In high-performance design, it is therefore critical in that inductance is accurately extracted and modeled for on-chip interconnects.

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 through the power grid are considered simultaneously instead of being analyzed in isolation. The current distribution in the entire circuit, including the grid, must be known in order to obtain a correct estimate of 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 accurate modeling and simulation of the signal net and power grid topology.

Traditional approaches to inductance analysis have been based on simple loop inductance models [4]–[6]. This approach is illustrated above in Fig. 1(a), where a typical signal net and its neighboring power and ground grid are shown. For illustration purposes, the power and ground grids are shown separate from each other, although in the design they are interleaved. An equivalent RL model is constructed by omitting all capacitors from the circuit, defining a port at the driver output, and shorting the signal net to the power and ground grids at the receiver input, as illustrated in Fig. 1(b). From this simplified RL model, the loop inductance and resistance are extracted by solving the current distribution for the circuit at a particular frequency 'ω' using tools such as FastHenry [7] and computing the port impedance $Z = R + j\omega L$. The equivalent inductance and resistance are then combined with lumped interconnect capacitance to construct an *RLC* netlist, as shown in Fig. 1(c). A more recent approach suggests the construction of a ladder circuit to model the frequency dependence of resistance and inductance [4] over a range of frequencies. The lumped *RLC*

Manuscript received October 31, 2001; revised March 31, 2002.

K. Gala and V. Zolotov are with Motorola, Inc., Austin, TX 78731 USA (e-mail: kaushik.gala@motorola.com)

D. Blaauw is with the University of Michigan, Ann Arbor, MI USA 48109 USA.

P. M. Vaidya is with the CA&SI, Urbana, IL USA and also with University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA.

A. Joshi is with the CA&SI, Urbana, IL USA.

Digital Object Identifier 10.1109/TVLSI.2002.801619

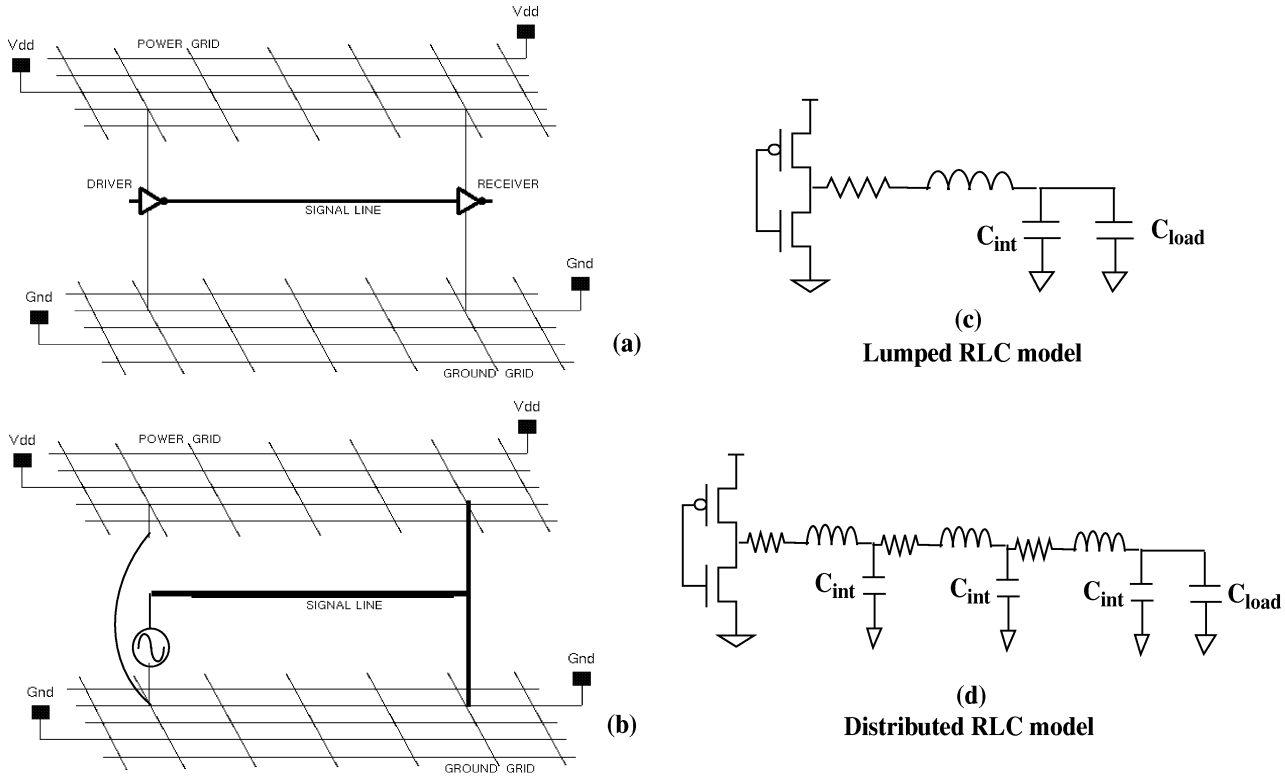

Fig. 1. Loop inductance approach for a simple signal net.

circuit representation can also be distributed using many *RLC*-π segments, as shown in Fig. 1(d).

The primary issue with the loop inductance approach is that while extracting the inductance, the current distribution is determined solely by the resistance and inductance in the circuit, while all capacitance is omitted. This leads to significant inaccuracies, since the coupling capacitance of the signal net and the device decoupling capacitances in the power grid strongly affect the current return paths. Also, defining a port at the driving gate and shorting the input of receiver gate alters the current distribution in the circuit and ignores certain current paths, such as the short-circuit gate current, current through the signal net coupling capacitances, and power grid current due to the switching activity of other gates in the circuit. The loop inductance approach therefore has limited accuracy. However, its simplicity leads to fast simulation times, making it applicable in pre-layout inductance estimation for highly controlled signal net topologies.

We propose a new inductance modeling and analysis methodology that uses the Partial Electrical Equivalent Circuit (PEEC) [9] method. In this approach, partial self and mutual inductances are defined for individual wire segments. The PEEC method can be used to construct a circuit model that does not require the predetermination of current loops, and can include capacitance in the circuit model. PEEC models have been used to obtain a more accurate current distribution [11], [12]. However, such techniques have been applied to highly simplified structures like coplanar waveguides. In addition, they ignore important components that determine current paths in on-chip interconnect structure, such as the short-circuit current, currents to the power pads, and through decoupling capacitances, and hence, lack accurate estimation capability.

In this paper, we determine the critical components that must be included in a PEEC-based model, by analyzing the current paths that arise in the signal net and power grid, when an on-chip signal switches. We propose a comprehensive model that includes the following elements:

- interconnect resistance, capacitance, and partial self and mutual inductances;
- explicit decoupling capacitance structures and implicit device decoupling capacitance between the power and ground grids;
- power and ground pad locations and package inductance models;
- quiescent power and ground currents due to the switching of other signal net drivers in the design;
- the signal net driver and receiver gates and their connections to the power grid.

These elements have a strong impact on current distribution in a power grid and their inclusion in the model leads to a significantly more accurate analysis of signal nets. We propose a effective methodology for extracting the listed elements, including statistical models for modeling the device decoupling capacitance and quiescent currents from other signal drivers. We also present a simple package model that can be included in the analysis and propose a numerically stable approach for calculating the partial self and mutual inductances.

The proposed approach was used on industrial circuits to study the effect of on-chip inductance on clock nets and signal busses. We compared the PEEC-based model with the simplified loop inductance model and found that the latter dramatically overestimates the impact of signal inductance. It is important for the circuit designer to know the accurate impact of inductance to avoid over-design of signal nets or shields. We also studied the impact of the different model elements on the signal behavior. We found that the number of pads and their locations, the package inductance, the decoupling capacitance, and the quiescent switching activity in the grid, significantly affect the signal behavior. This underscores the importance of including these elements in an accurate model and also allows the designer to optimize the signal performance by correctly controlling these elements in the design.

The proposed PEEC model leads to a dense *RLC* circuit matrix requiring large SPICE simulation times. Hence, we further propose a new analysis methodology for efficiently simulating large PEEC-based models. We employ a number of acceleration techniques. In order to overcome the computational complexity imposed by the dense mutual inductance matrix, we propose a simple block sparsification method which dramatically increases the run time efficiency while guaranteeing the passivity of the model. We then employ reduced order modeling, using the PRIMA [22] algorithm, to further increase the analysis efficiency of the sparsified model. The moment computation in the PRIMA algorithm requires factorization of the conductance matrix, which can be nonpositive definite when partial inductances are present in the model. This requires the use of LU decomposition, instead of the more efficient Cholesky factorization [29] that can be used for models without inductance. However, we propose a new formulation of the moment computation that results in both the reduction of the matrix size, and also in it being positive-definite. Therefore, our formulation can use efficient Cholesky factorization of a reduced matrix for models that include inductance, yielding a significant speed advantage. To further improve the efficiency of the reduced order model generation, we reduce the number of ports in several ways including the use of active and passive ports, collapsing close-by power/ground ports and replacing receiver gates with equivalent capacitive loads.

The proposed inductance extraction and analysis methodology was implemented in an industrial analysis tool called Lyric and was used on a number of high-performance microprocessor and DSP designs. We demonstrate the accuracy of the proposed PEEC-based model and the proposed acceleration techniques. Results are presented for a global clock net of a giga-hertz processor, consisting of 720 thousand *RLC* elements and over 8 million mutual inductances. This interconnect structure could be extracted and analyzed in less than 1 hour, clearly demonstrating the efficiency of the proposed method. We also show that the on-chip inductance for this structure significantly impact its signal behavior, underscoring the importance of accurate and efficient on-chip inductance modeling.

The remainder of this paper is organized as follows: In Section II, we present the PEEC-based inductance model and discuss the extraction and modeling methodology of its various elements. In Section III, we present the proposed analysis methodology. We present our technique for increasing the efficiency of the analysis using mutual inductance sparsification and accelerated reduced order modeling and compare these techniques to other methods. Finally, in Section IV we present our simulation results and in Section V we draw our conclusions.

II. PEEC-Based Inductance Model

A PEEC-based model allows a wide variety of circuit elements to be represented in the model. In order to determine whether the inclusion of an element is critical for the accuracy of the model, we first analyze the current paths that arise when a signal net switches. Fig. 1(a) shows an abstract representation of a signal net topology. A typical circuit topology consists of two supply grids (power and ground) and signal lines laid out over multiple metal layers. The driving gates draw power from the lowest metal layer, while external power and ground are supplied via pads to the uppermost metal layer. Although in Fig. 1(a) the two power grids are shown separate from each other, the actual power and ground grids are interleaved and the signal net is embedded among them. Power and ground shields, which are connected to the power and ground network may also be present on either side of the signal net.

In Fig. 2 three primary current loops are shown that arise in the power and ground grid when a gate drives a loaded signal line. In the example, the signal net is switching high. Fig. 2(a) shows the short circuit current loop I_1, consisting of current flowing from the power pads, through the power grid and the P and N devices of the driver gate and then through the ground grid to the ground pads. Fig. 2(b) shows the charging current loop I_2, consisting of current flowing from the power pads through the power grid and P devices of the driver gate to the coupling capacitance between the signal net and the ground grid, and then through the ground grid to the ground pads. This current loop is responsible for charging the capacitances between the ground grid and the signal net. Note that both currents I_1 and I_2 complete their current loop through the package and outside power supply. The inductance model therefore requires accurate modeling of the pad inductances and the pad locations in the grid. These current loops are also completed through on-chip decoupling capacitance between the power and ground grid which acts as a local supply voltage. Such decoupling capacitances can take the form of either explicitly designed decoupling capacitor structures, or are contributed by the parasitic capacitances of the nonswitching devices in design. Since such decoupling capacitances impact the current distribution in the power grid, it is necessary to include them in the inductance model. Also, the power and ground grid structures that connect the driver gate and signal coupling capacitances to the nearest power and ground pads need to be represented in the model.

The third current loop I_3 is shown in Fig. 2(c) and consists of current flowing from the coupling capacitances between the signal net and the power grid, through the power grid and P devices of the driver gate and back to the decoupling capacitances. This current forms a self-contained current loop that is responsible for discharging the coupling capacitance between the power grid and signal net. In addition to the current generated by the transition of the signal net of interest, other gates

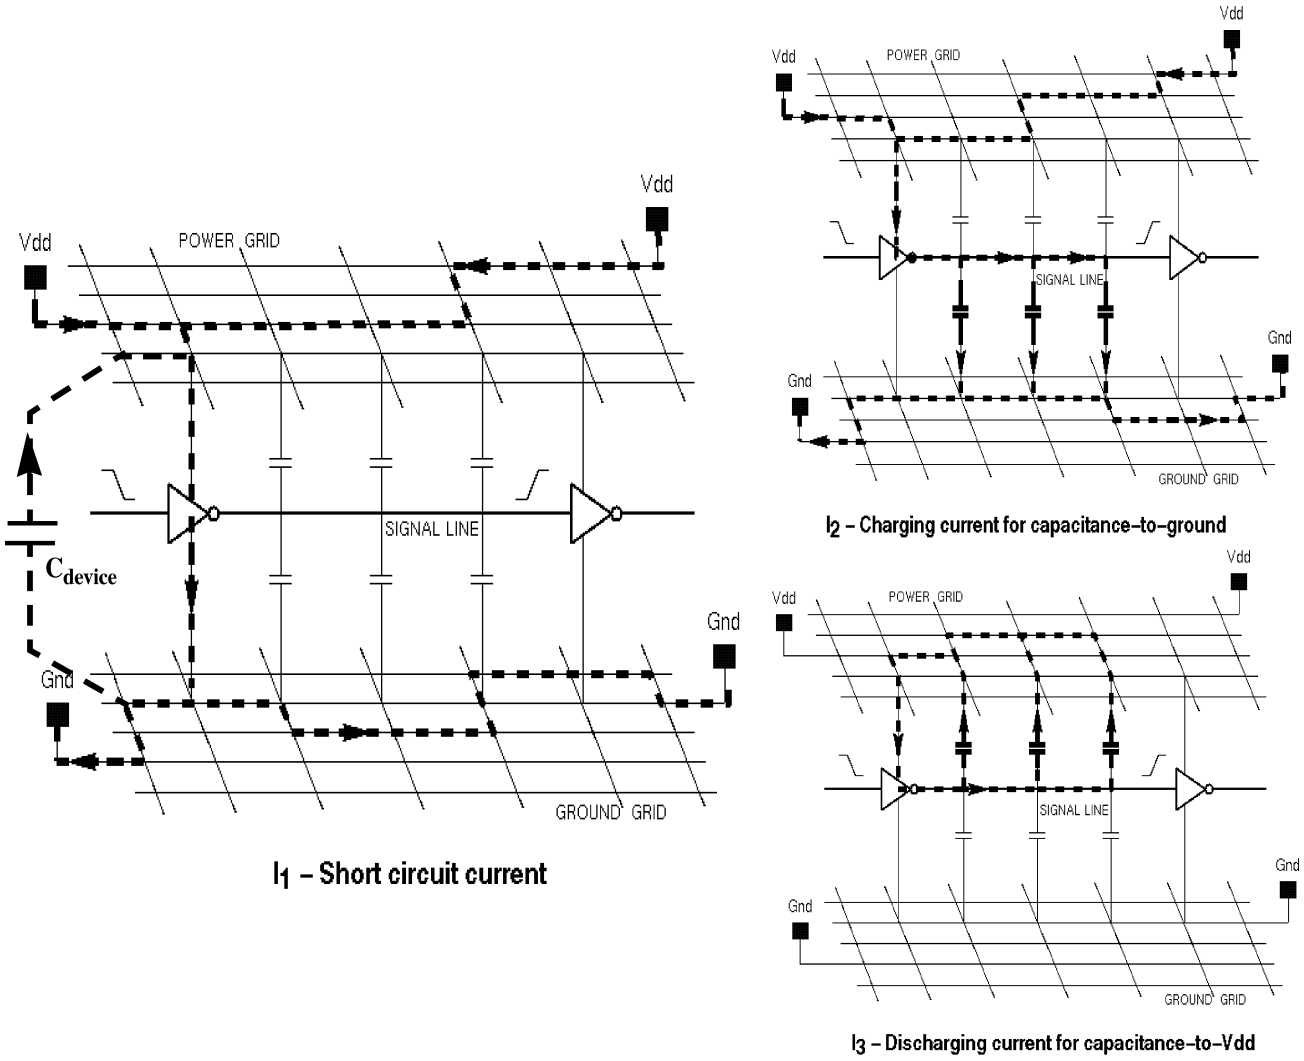

Fig. 2. Currents in driver-receiver-grid topology.

in the design are also switching and producing currents in the power and ground grids. Since these so-called *quiescent* power grid currents affect the voltages in the grid and its current distribution, they also need to be represented in the inductance model.

Our proposed PEEC-based inductance model therefore consists of the following elements:

- resistance, partial self-inductance and capacitance (*RLC*-π) for each metal segment in the signal net and in those portions of the power and ground grid that connect the signal driver and receiver to their nearby pad locations or that lie in the vicinity of the signal net;
- mutual inductances between all pairs of parallel segments;
- coupling capacitances between all pairs of adjacent metal lines;
- via resistances between adjacent metal layers;
- implicit and explicit decoupling capacitance between the power and ground grids;
- quiescent currents due to other gates switching in the design;
- pad resistances and package inductance models;
- the driver and receiver gates.

In addition to these, our model can also be extended to include a substrate model for enhanced accuracy, but at the cost of additional model complexity. The structure of the inductance model is illustrated in Fig. 3. For clarity, the power grid is removed and the mutual inductance between each pair of wire segments is omitted. The extraction of the resistance, coupling capacitance, and partial self and mutual inductances is discussed below in Section II-A. The power grid decoupling capacitances are represented with a series *RC* circuit and the quiescent switching current by current sources. The extraction and modeling of these two elements is discussed in Sections II-B and C, respectively. The package model for the power and ground pads is discussed in Section II-D. Finally, in Section IV, we present empirical results showing the impact of each model element on the behavior of the signal net transition. These results underscore the importance of the proposed detailed PEEC-based model.

A. Interconnect RLC Extraction

Each grid segment is modeled as an *RLC*-π circuit. The resistance is computed as a function of length, width and sheet resistance. The segment capacitance to ground and the coupling

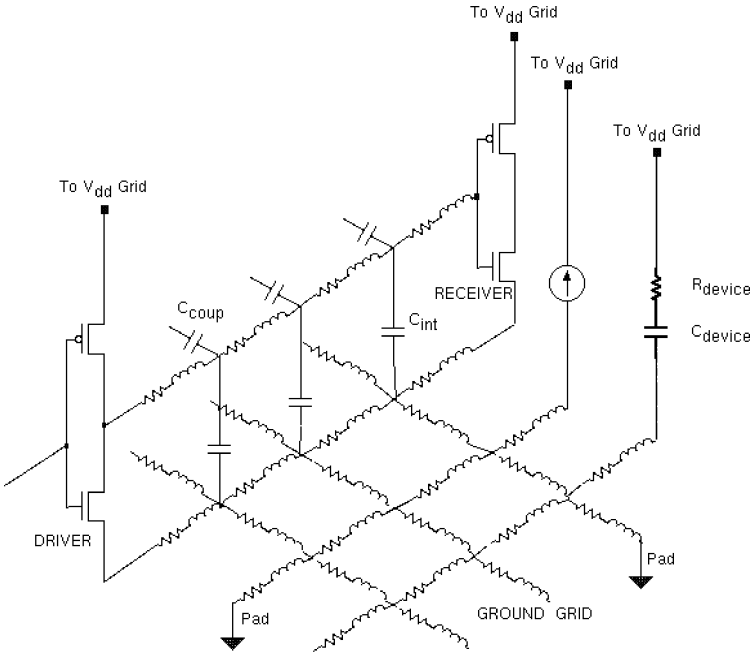

Fig. 3. PEEC-based inductance model and its elements.

capacitances between each pair of adjacent metal lines are computed using either Chern models [27], or using commercial *RC* extraction tools [31].

The partial self and mutual inductances are a function solely of the geometry of the wire segments and can be computed by taking a double volume integral over the two conductors. For on-chip VLSI structures, we take advantage of the fact that conductors are rectangular and parallel with arbitrary relative positioning, as shown in Fig. 4. In [15], a closed form solution for the partial self and mutual inductances of two such wires was presented. By choosing the appropriate limits of integration for function "f" later, we obtain either the partial self inductance or the partial mutual inductance shown in the equation at bottom of page. This analytic formula is exact, under the assumption that current is uniformly distributed across the cross-section of the wire. This formulation therefore ignores the skin effect and proximity effect within the conductor. For a rise-time of 100 ps (which is typical for a 2.5 GHz processor), the maximum frequency of interest is 3.2 GHz and the skin depth is 1.53 μm, allowing for a maximum wire thickness of 3.06 μm. Since all on-chip interconnect structures are currently less than this thickness, skin-effect does not need to be considered for on-chip interconnects. However, for very wide wires, proximity effect can pose a problem. In this case, a wire must be split in its width to form multiple narrower wire segments which are connected in parallel. Although this operation is straightforward, it results in a significant increase in the model size, which can cause run time issues. Fortunately, current interconnect technologies use chemical mechanical polishing (CMP) methods which strictly limit the maximum allowed width of interconnects. Thus additional discretization of wires in their widths is not needed in practice.

For wire segment pairs lying far from each other, the earlier formula becomes numerically unstable. This results from operations such as subtraction of z^4 and y^4 from y^2z^2 in the first term. For two wires separated by 1000 times their individual

$$
\begin{aligned}
f = \frac{0.001}{abcd}\Bigg[\Bigg[\Bigg[& -\frac{xyz^3}{6}a\tan\frac{xy}{x\sqrt{x^2+y^2+z^2}} - \frac{xy^3z}{6}a\tan\frac{xz}{y\sqrt{x^2+y^2+z^2}} - \frac{x^3yz}{6}a\tan\frac{yz}{x\sqrt{x^2+y^2+z^2}} \\
& + \left(\frac{y^2z^2}{4} - \frac{y^4}{24} - \frac{z^4}{24}\right)x\log\left(\frac{x+\sqrt{x^2+y^2+z^2}}{\sqrt{y^2+z^2}}\right) + \left(\frac{x^2z^2}{4} - \frac{x^4}{24} - \frac{z^4}{24}\right)y\log\left(\frac{y+\sqrt{x^2+y^2+z^2}}{\sqrt{x^2+z^2}}\right) \\
& + \left(\frac{x^2y^2}{4} - \frac{x^4}{24} - \frac{y^4}{24}\right)z\log\left(\frac{z+\sqrt{x^2+y^2+z^2}}{\sqrt{x^2+y^2}}\right) \\
& + \frac{1}{60}(x^4+y^4+z^4-3x^2y^2-3y^2z^2-3x^2z^2)\sqrt{x^2+y^2+z^2}\Bigg](x)\Big|_{E+d-a,E}^{E-a,E+d}\Bigg](y)\Big|_{P+c-b,P}^{P-b,P+c}\Bigg](z)\Big|_{l_3+l_2-l_1,l_3}^{l_3-l_1,l_3+l_2}
\end{aligned}
$$

where, $\Big[\Big[[f(x,y,z)](x)\big|_{q_2,q_4}^{q_1,q_3}\Big](y)\big|_{r_2,r_4}^{r_1,r_3}\Big](z)\big|_{s_2,s_4}^{s_1,s_3} \equiv \sum_{i=1}^{4}\sum_{j=1}^{4}\sum_{k=1}^{4}(-1)^{i+j+k+1}f(q_i,r_j,s_k).$

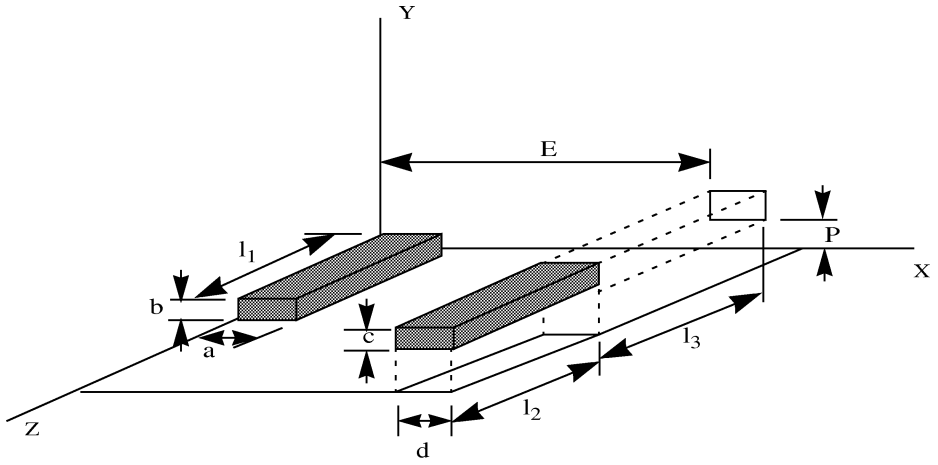

Fig. 4. Two parallel rectangular conductors, placed in any relative position.

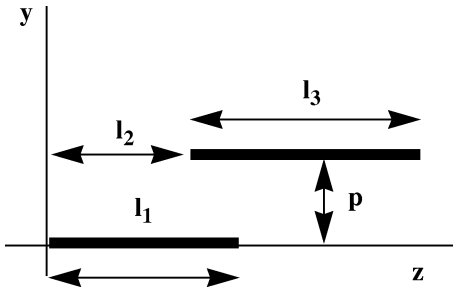

$$f = 0.001[z\log(z+\sqrt{z^2+p^2}) - \sqrt{z^2+p^2}](z)\Big|_{l_3+l_2-l_1,\,l_3}^{l_3-l_1,\,l_3+l_2}$$

$$where, [f(z)](z)\Big|_{s_2,s_4}^{s_1,s_3} \equiv \sum_{k=1}^{4}(-1)^{k+1} f(s_k)$$

Fig. 5. Far-away parallel rectangular conductors treated as FILAMENTS.

wire widths, the ratio between the magnitude of the terms and their difference reaches 10^{20} orders and the computation becomes numerically unstable, unless expensive high precision arithmetic is used. Therefore, for far-away conductors, we approximate the conductors as filaments as shown in Fig. 5 and use a simpler formula, which allows for fast and stable computations.

B. Power Grid Decoupling Capacitance

The primary decoupling capacitance between the power and ground grid is contributed by three sources: explicit decoupling structures, n-well capacitance, and device decoupling capacitance. Explicit decoupling structures can be extracted using commercial extraction tools, while n-well capacitance can be characterized using a process simulator [28] and represented using a simple lumped model for each N-well. Device decoupling capacitance is contributed by the parasitic device and interconnect capacitances of the gates and signal routes that are static (i.e., nonswitching) in a design. Since at any point in time as much as 80–90% of all gates are static, these nonswitching gates result 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. The difficulty in modeling this decoupling capacitance is that the exact set of gates that are static during a particular transition of a signal net can vary greatly depending on the overall operation of the design. Therefore, we use a statistical model for the device decoupling capacitance, which reflects the average decoupling capacitance expected during the operation of the design.

We estimate the decoupling capacitance directly using SPICE simulation of several representative circuit blocks. Fig. 6(a) depicts the simulation setup and Fig. 6(b) shows the equivalent *RC* circuit for a circuit block. The input terminals of the circuit are set arbitrarily at logic 0 or 1 and the power terminals are set at a dc bias equal to the operating voltage. A small sinusoidal voltage is then superimposed on the supply rails to cause fluctuation in the supply voltage. The decoupling action of the circuit is studied by monitoring the input current. Since no devices are switched, the input current is solely in response to the fluctuation in the grid voltage, thus representing the current in the underlying *RC* decoupling circuit. The C_{eff} and R_{eff} values determined for a block represent the combined decoupling action of the device capacitances and the extracted parasitic capacitances of its interconnects. The obtained equivalent decoupling capacitance is then scaled to account for the switching activity in the circuit, since in practice not all gates are static. We repeat the above measurements for a set of random states and take the average values for resistance and capacitance. The values of one block can be easily translated to other circuit blocks based on the size of the blocks (total transistor widths). For each circuit block the total decoupling capacitance is computed and then distributed across the area of the block using a number of small equivalent decoupling capacitor structures.

C. Quiescent Power Grid Current

In addition to driver gate of the signal line, other gates switch simultaneously in the design, drawing current from the V_{dd} grid and injecting it into the ground grid. These quiescent power grid

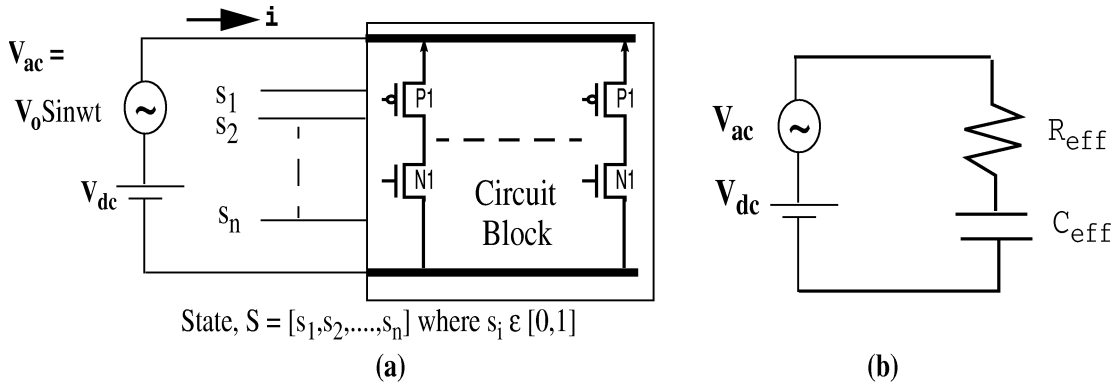

Fig. 6. Determination of device decoupling cap.

currents cause voltage fluctuations in the power and ground grid and affect the current return paths of the signal of interest and hence, its inductance. Different gates draw current at different times and in varying amounts, causing a continuously changing current profile in the grid. Explicit modeling of all switching gates during the signal transition would lead to an intractably large models. We therefore again use a statistical model, consisting of time-varying current sources distributed through the power grid and connected at the lowest metal layer. Each current source represents the combined current of a small number of switching gates. Each switching gate is represented by a sharp triangular current pulse, obtained through characterization of typical gates using SPICE simulation. In a typical 0.13 μm design, a current pulse with a peak current of 1 μA and 200 ps pulse width was used. Since the transition time of the signal of interest is relatively fast, the combined current consumption of the entire model is assumed to be constant through the duration of the transition. Therefore, alignment times of the pulse currents are assigned randomly using a uniform probability distribution across the simulation time. For longer simulation times, a more complex current profile can be used [28].

D. Pad/Package Inductance Model

External signals are routed to a chip via package leads and pads. The parasitic inductances associated with the package geometries must be modeled, since they affect on-chip behavior significantly. In our circuit model, it is assumed that the planes in the package are ideal, since the voltage difference across these planes is typically of order of few mV. Thus, for a flip-chip-bump package technology, the vertical via which connects a pad to a power supply plane in the package is modeled as a rectangular bar. For wire-bonded packages or to obtain a higher level of accuracy, commercial package modeling tools [32] can be employed which produce a reduced multiport inductance model of the package, which can be incorporated in the PEEC-based inductance model.

III. Analysis Methodology and Acceleration Techniques

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 their surrounding power grid can result in a models with 100 000 self inductances, which would require 10 G mutual inductances. Computing such a large number of mutual inductances is impossible. Even for smaller systems, this computation might be done given sufficient time but, SPICE simulation would be infeasible due to extremely large time and memory requirements. Consequently, the large number of model elements and the dense mutual inductance matrix have been the main bottleneck in the use of PEEC models.

A number of techniques have been proposed to address these problems. To accelerate the mutual inductance computation, both explicit and implicit sparsification methods have been proposed. Explicit sparsification methods aim to reduce the number of nonzero entries in the partial inductance matrix by computing only a sub-set of all possible mutual inductances. The simplest approach is to discard all mutual coupling terms falling below a certain threshold. However, the resulting matrix can become nonpositive definite, and the sparsified system can become active and can generate energy. As an alternative to simple truncation, one approach associates each segment with a distributed current return path out to a shell of some radius [16]. Segments with spacing more than this radius are assumed to have no inductive coupling while the mutual inductance to segments within the radius are modified to guarantee the positive definite nature of the sparsified matrix. However, the computation of the radius is difficult. An extension of this work [17] uses a moment-based algorithm to compute the shell radius. Another approach for limiting the inductive interaction is proposed by [18], which introduces return-limited inductances for sparsification and the use of "halos" to limit the number of mutual inductances. This approach is based on the assumption that the currents of signal lines return within the region enclosed by the nearest same-direction power-ground lines. A recent approach [23] defines a circuit matrix K, as the inverse of the partial inductance matrix L which is more amenable to sparsification and simulation. However, it requires inversions of partial inductance submatrices, which can be computationally expensive, and requires a special circuit simulator that can model K elements.

Implicit techniques do not discard mutual inductances, but rather approximate the far-off inductive couplings without explicitly computing them. The Fast Multipole algorithm, which has been used in FastHenry [7], replaces a cluster of

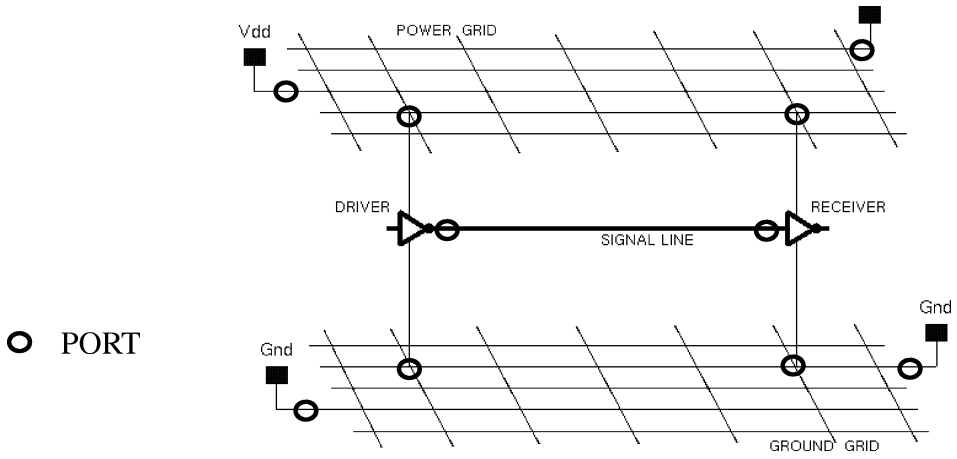

Fig. 7. Port locations for reduced order model generation.

elements with a single representative element to compute its far-off inductive couplings, thereby reducing the number of coupling terms. The Pre-corrected FFT method uses a regular grid upon which the vector potential of the elements is projected to compute approximate interactions between far-away conductors. This method has only been used for capacitance extraction [8], but can be applied to inductance extraction as well. Similar to the explicit methods, the implicit methods compute exact mutual inductances for nearby couplings in addition to approximate far-off coupling. Their run-times are therefore typically slower than the explicit methods, although they have a higher accuracy.

A. Proposed Analysis Methodology

In this paper, we propose the use of a block-diagonal sparsification method to reduce the density of the mutual inductance matrix, combined with a reduced-order modeling technique for an efficient and accurate inductance analysis. The circuit is partitioned (geometrically) using a rectangular grid, and the grid dimensions are controlled to achieve a tradeoff between model accuracy and run-time. The partitioned circuit thus results in block-diagonal sparsification, and the mutual inductance matrix is constructed as follows:

- each section is stamped using self inductances and all mutual inductances between elements of 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 the surrounding power grid;
- sections far away from the signal of interest can be modeled as *RC* instead of *RLC*.

Each section can be thought of as an independent circuit with all mutual coupling preserved and is therefore passive. Since different sections are only connected through resistive and capacitive couplings, the complete circuit is also passive and the block-diagonal sparsification preserves the positive definite nature of the inductance matrix. Also, note that the entries that are not discarded do not need to be modified in order to preserve the passivity of the model and that the approach is simple and efficient to implement.

Even after the mutual inductances have been sparsified, the remaining circuit model can consist of hundreds of thousands of *RLC* elements and several million mutual inductances. Therefore, direct simulation of such a model with SPICE is still prohibitively expensive. We therefore combine the sparsification with reduced order modeling using the PRIMA algorithm [22] to further increase the efficiency of the analysis. The reduced order model is constructed by defining ports at the power grid supply locations and the driver and receiver gate connections, as shown in Fig. 7.

Since the driver and receiver gates connect to local power and ground connections, three ports are required for each gate (local V_{dd}, ground and signal net), and a total of six ports is required for each signal line. Since the efficiency of the reduced order model generation depends linearly on the number of ports, it is necessary to reduce the number of ports for circuits with a large number of signal nets, such as large bus structures, or for signal nets with a large number of sinks, such as clock nets. We reduce the number of ports in several ways:

- The receiver gates are replaced with equivalent linear capacitors, which are then included in the interconnect model. This substitution introduces a small error, since the nonlinear parasitic device capacitances are modeled with a linear capacitor. However, for large signal nets, the interconnect capacitance is significantly greater than the device loading capacitance and, hence, the error is typically insignificant.
- We differentiate between the passive ports representing the receiver gate inputs and the active ports at the driver gate outputs. A variant of the PRIMA algorithm is then used to reduce the computation time by applying excitation sources only to the active ports, and not to passive sink ports [30]. In this case, the reduced order model generation complexity is linear in the number of active ports, and voltage controlled voltage sources can be constructed for the passive ports with only a marginal computational overhead.
- For driver and receiver gates that are placed closely together, power and ground ports can be shared. The power

and ground grid is typically very well connected such that adjacent gates observe virtually the same supply voltage and only a small error is incurred. This technique is particularly useful for large busses where drivers are closely spaced together.

The number of current sources needed to model the quiescent power grid currents can easily exceed tens of thousands. Modeling this elements would therefore result in an excessively large number of ports in the reduced order model. However, as shown in Section IV, the presence of these quiescent currents reduces the inductance effects on the signal behavior. Their omission therefore introduces a slight overestimation of the inductance effect on the signal net behavior.

Even after the number of ports has been sufficiently diminished, reduced order model generation remains computationally expensive due to the large problem size. In Section IV, we therefore show how to increase the efficiency of the reduced order model generation itself.

B. Acceleration of Reduced Order Model Generation

The main computation time in the reduced order model generation using the PRIMA algorithm is devoted to factorization of the conductance matrix. Due to the presence of inductances in the proposed model, this matrix is no longer positive definite. This requires the use of standard LU decomposition, instead of the significantly more efficient Cholesky factorization. In this section, we therefore transform the problem such that it involves the factorization of a positive definite matrix. In the process, the matrix size is also significantly reduced, further accelerating the reduced order model generation time.

PRIMA is based on the Modified Nodal Analysis formulation for linear circuits using circuit equations of the following form:

$$\begin{aligned} Cx_n + Gx_n &= Bu_n \\ L^T x_n &= i_n \end{aligned}$$

where G is the conductance matrix, C is the susceptance matrix, x_n are the state variables, i_n are the port currents and u_n are the port voltages. Further,

$$C = \begin{bmatrix} Q & 0 \\ 0 & H \end{bmatrix} \quad G = \begin{bmatrix} N & E^T \\ -E^T & 0 \end{bmatrix} \quad x_n = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

here, x_1 and x_2 are the MNA variables corresponding to the node voltages and branch currents respectively. The matrices N, Q and H contain the stamps for resistors, capacitors and inductors respectively. E consists of one and minus one entries, which represent the current variables in the KCL equations.

PRIMA requires factorization of the G matrix during moment computation, which can be formulated as follows. Solve the linear system

$$Gx = b \tag{1}$$

where

$$G = \begin{bmatrix} N & E \\ -E^T & 0 \end{bmatrix}, \quad x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}, \quad b = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}.$$

The conductance part of the G matrix (i.e., N) is symmetric positive definite. However, for general *RLC* circuits, the presence of E makes the full admittance matrix G nonpositive definite, preventing the use of the Cholesky method. Our approach for reducing the system size and for allowing the use of Cholesky factorization is outlined as follows:

1) demonstrate that solving this system is equivalent to minimizing a convex function, with equality constraints;
2) demonstrate that the equality constraints allows the linear system to be reduced to an equivalent, unconstrained system with lower order;
3) demonstrate that the admittance matrix G for the reduced system is positive-definite and hence can be factorized using Cholesky decomposition.

Equivalence to Minimization of Convex Function: We first rewrite the linear system (1) as follows:

$$Nx_1 + Ex_2 = b_1 \tag{2}$$

$$E^T x1 = -b_2. \tag{3}$$

We show that solving the original problem (1) is equivalent to solving the constrained minimization problem (4), defined as

$$\text{minimize } \frac{1}{2} x_1^T N x_1 - b_1^T x_1, \text{ subject to } E^T x_1 = -b_2. \tag{4}$$

To solve problem (4), let us consider the equivalent Lagrangian augmented function:

$$F(x_1, x_2) = \frac{1}{2} x_1^T N x_1 - b_1^T x_1 + x_2^T \left(E^T x_1 + b_2 \right) \tag{5}$$

where, the Lagrange multipliers x_2 have been introduced to handle the equality constraints.

Setting the partial derivatives of (5) with respect to x_1 to zero gives: $Nx_1 + Ex_2 = b_1$.

Setting the partial derivatives of (5) with respect to x_2 to zero gives: $E^T x_1 = -b_2$.

The constraints $E^T x_1 = -b_2$ indicate that any change or variation Δx_1 in x_1 must satisfy $E^T \Delta x_1 = 0$. Thus, these constraints force the solution (i.e., the minimum) to satisfy the imposed constraints. The equations $Nx_1 + Ex_2 = b_1$ say that at the minimum, the gradient of the quadratic function in problem (4) must be orthogonal to the subspace defined by the constraints $E^T x_1 = 0$. In other words, the gradient has a zero component in this subspace.

Note that the function in (4) is convex, since it is quadratic and N is positive semidefinite. Setting the partial derivative of a convex function to zero provides us with a minimum of the function. Thus, setting the partial derivatives of the Lagrangian augmented function to zero, will provide us with a solution to problem (4). Also, since the original system has a unique solution, the solution to (4), obtained by setting the partial derivative of (5) to zero, is also unique. Therefore solving (4) is equivalent to solving the original system of (1).

Using the Equality Constraints in Problem (4) to Reduce the Linear System: We consider the equality constraint $E^T x_1 = -b_2$ of (4). Each row of E^T contains at most 2 nonzero entries which are ± 1. A constraint is therefore of the form $x_{1i} = -b_{2k}$ or $x_{1i} - x_{1j} = -b_{2k}$. We can assume that the rows of E^T are linearly independent; since otherwise there is a redundant constraint and the matrix G is singular. A constraint of the form $x_{1i} = -b_{2k}$ fixes the voltage at node i and we assume that such constraints have been eliminated and the dimensionality of the problem has been appropriately reduced. Such reduction is implemented by using a queue type data structure, which keeps track of constraints with a single variable of the current

$$x_{1a} - x_{1b} = \alpha_1$$

$$x_{1b} - x_{1c} = \alpha_2$$

$$x_{1c} - x_{1d} = \alpha_3$$

$$x_{1d} - x_{1e} = \alpha_4$$

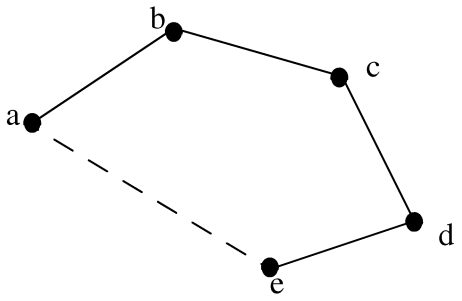

Fig. 8. Graph representation of equality constraints.

step during the elimination process. Once constraints with single variables have been eliminated, we are left with constraints with exactly 2 variables, of the form $x_{1i} - x_{1j} = -b_{2k}$.

Consider a graph where x_{1i}'s are the nodes and each constraint corresponds to an edge. An edge corresponding to $x_{1i} - x_{1j} = -b_{2k}$ is between nodes i and j, as illustrated in the example in Fig. 8. If there is a path between a & e in this example graph, then it automatically imposes a constraint on the difference $x_{1a} - x_{1e}$. This additional constraint $x_{1a} - x_{1e} = \alpha_5$ is then either linearly dependent and redundant or inconsistent and the graph is therefore acyclic.

Each constraint of this type fixes the voltage across an inductor to a specified value and in each connected component, there is exactly one independent nodal voltage. A connected component with m nodes has $m-1$ constraints associated with it. Suppose for a connected component, we choose nodal voltage x_{1i} corresponding to node i in the component as the independent variable. Then, all constraints corresponding to the component may be put in the form $x_{1i} - x_{1j} = -\beta_{ij}$ for each node j in the component other than i. We refer to the node i whose voltage is the independent variable as a *master node* and the remaining nodes in the component as the *slave nodes*. The voltage variables corresponding to all the slave nodes can be eliminated from the problem and we can solve the a reduced problem which is unconstrained, as follows: Let $x_1 = \begin{bmatrix} x_{1s} \\ x_{1m} \end{bmatrix}$ where x_{1s} are the slaves and x_{1m} are the masters. We can rewrite the constraints as $x_{1s} = Hx_{1m} + \beta$, where H has the following special structure:

- each row of H has exactly one nonzero entry with value 1;
- all rows for slave nodes in the same component have their single nonzero entry in the same column.

$$\text{Let } N = \begin{bmatrix} N_s & N_{sm} \\ N_{sm}^T & N_m \end{bmatrix}.$$

Substituting $x_{1s} = Hx_{1m} + \beta$ in problem (4), the reduced problem becomes

$$\text{Minimize } \frac{1}{2} \begin{bmatrix} (Hx_{1m}+\beta)^T & x_{1m}^T \end{bmatrix} \begin{bmatrix} N_s & N_{sm} \\ N_{sm}^T & N_m \end{bmatrix} \times \begin{bmatrix} Hx_{1m}+\beta \\ x_{1m} \end{bmatrix} - \begin{bmatrix} b_{1s}^T & b_{1m}^T \end{bmatrix} \begin{bmatrix} Hx_{1m}+\beta \\ x_{1m} \end{bmatrix}.$$

We can further rewrite the reduced problem as

$$\text{Minimize } \frac{1}{2} x_{1m}^T N_R x_{1m} - b_R^T x_{1m}$$

where N_R is the reduced conductance matrix and b_R is the reduced right-hand side

$$N_R = N_m + H^T N_s H + N_{sm}^T H + H^T N_{sm}$$
$$b_R = -\left(H^T N_s \beta + N_{sm}^T \beta\right) + \left(H^T b_{1s} + b_{1m}\right).$$

Because of the special structure of H, N_R has the same sparsity as N and both N_R and b_R can be efficiently computed. The complete solution therefore consists of the following steps:

1) Solve the reduced problem (using the fast Choleksy method) $N_R x_{1m} = b_R$
2) Using the constraints compute the voltages x_{1s} at the slave nodes component by component.
3) Once the nodal voltages x_1 have been obtained, compute the inductor current by solving for x_2 in $Ex_2 = b_1 - Nx_1$.

Each column of E corresponds to an inductor current. The graph induced by the inductive links alone is acyclic and corresponds to a tree when restricted to a connected component. The row corresponding to a leaf has a single nonzero entry and we can directly solve for the corresponding "inductor current" in x_2. Substituting for this inductor current and eliminating that variable corresponds to deleting that node, so we always have a row (i.e., a node) with a single nonzero entry during the solution. This directly follows from the fact that a tree has at least two leaf nodes. Thus solving for x_2 can be efficiently executed.

Positive Definiteness of N_R: Based on the definition of N_R and the fact that N is symmetric positive semidefinite, i.e., $x_1^T N x_1 \geq 0$, it can be proven that $x_{1m}^T N_R x_{1m} \geq 0$. Thus N_R will also be symmetric positive semidefinite. Further, the original problem has a unique solution, and the equivalent problem (4) has a unique minimum.This implies that N_R is not singular and hence is *symmetric positive definite*. This property allows the application of Cholesky decomposition for factoring the G matrix. Thus, the run-time for the PRIMA algorithm can be significantly reduced.

IV. Simulation Results

The proposed inductance and analysis methodology was implemented in an industrial analysis tool, called Lyric. It was used on a number of high performance microprocessor and DSP designs. Below, we present the following simulation results:

- comparison of the proposed PEEC-based model with the loop inductance model for a signal bus;
- investigation into the impact of model components on the signal behavior;

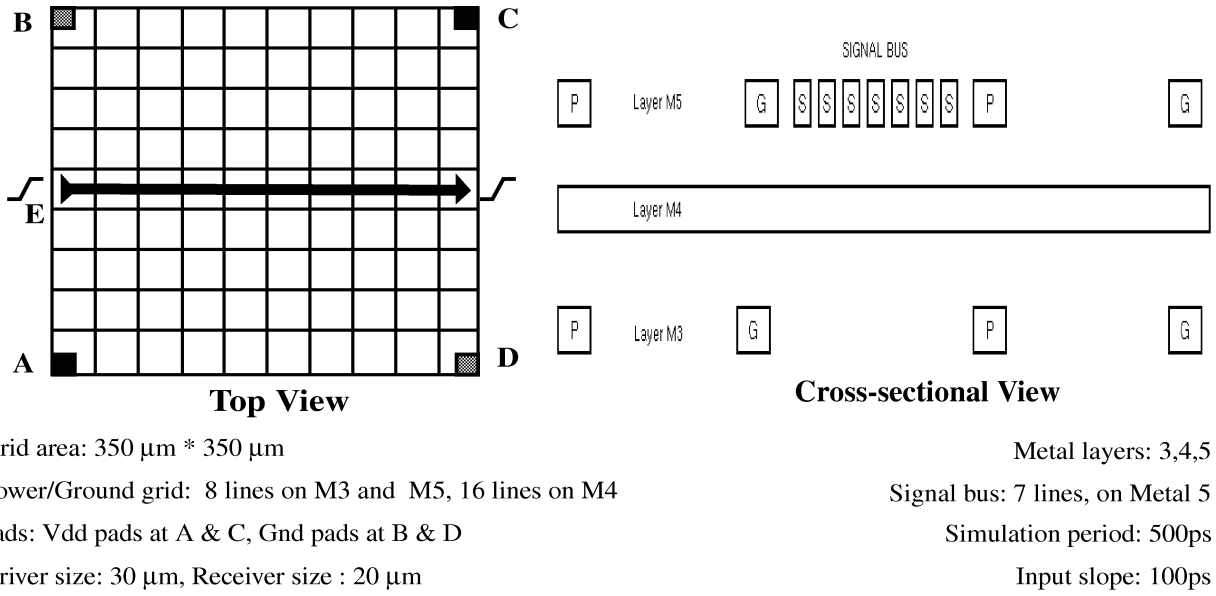

Fig. 9. Signal bus and power grid topology.

- comparison of the analysis with acceleration techniques (block diagonal sparsification and reduced order modeling) compared with direct SPICE simulation;
- simulation results for a large microprocessor clock net.

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 signal lines routed on the uppermost layer, which typically carry global signals such as clocks and buses in the presence of a multilayer power grid. For the first three experiments, we use a 7 bit bus in a three-layer power and ground grid, as shown in Fig. 9. This topology is based on a recent high-performance microprocessor design. The power and ground grid occupy all three metal layers, while the signal bus lies on the uppermost layer M5. Each signal line is connected to driver and receiver inverters. Power and ground pads are connected as shown. Metal M1 and M2 are not included in the model since they consist of short wires in a dense grid and do not significantly affect the accuracy of the analysis. Their omission dramatically reduces the model size and allows us to perform the analysis without the proposed acceleration techniques, for comparison.

A. Comparison of the PEEC Model With the Loop Inductance Model

We compared the proposed PEEC model with the traditional loop inductance model for the bus topology shown in Fig. 9. For the loop inductance model, the complete topology was analyzed with FastHenry and the loop inductance and resistance were extracted by defining ports for each signal line. These were then combined with interconnect capacitances using a 10-section distributed π-model, which was simulated in SPICE. The delay for the signal interconnect was measured from 50% V_{dd} at the driver output to 50% V_{dd} at the receiver input. Undershoot is measured at the receiver input and signal slope is measured from 10% to 90% of V_{dd}. Fig. 10 shows the simulation results for the two approaches. The loop inductance approach significantly overestimates delay and undershoot, which is undesirable since it might prompt the designer to overcompensate in the power grid or shielding structures, thereby yielding an inefficient interconnect topology. This overestimation is in part due to the lack of capacitances in the loop inductance model when inductance is extracted. This forces current to travel to the end of the signal lines, increasing the loop current area and resulting in a higher inductance estimate.

Fig. 11 shows the same comparison for a larger circuit (700 μm$*$350 μm). This larger topology demonstrates a worse overestimation of inductive effects in the loop approach.

B. Impact of Model Components on Signal Behavior

In Table I, we present simulation results demonstrating the impact of the different model components on the signal behavior for the bus topology shown in Fig. 9. In row 1, the delay, slope, and undershoot of the middle signal line in the bus are recorded when all discussed model components, except the package inductance, are included in the analysis. In rows 2 through 6, we present the same signal measures when a model component is omitted, added, or altered. We discuss each case in more detail below:

1) Device Decoupling Capacitance/Explicit Decoupling Capacitance: In row 2, the analysis was performed without the device decoupling capacitance. The total device decoupling capacitance in the model was 75 pF for the 350 μm by 350 μm die area, which corresponds to a device decoupling capacitance of 6 nF for a full size design of 1 cm^2. The removal of the device decoupling capacitance from the model introduces high-frequency oscillations into the transient voltages, resulting in a 12% increase in the signal undershoot and a 26% decrease in the delay. The device decoupling capacitance therefore has a significant impact on the delay and undershoot of the signal and must be accounted for in the inductance model. Conversely, in row 3, the

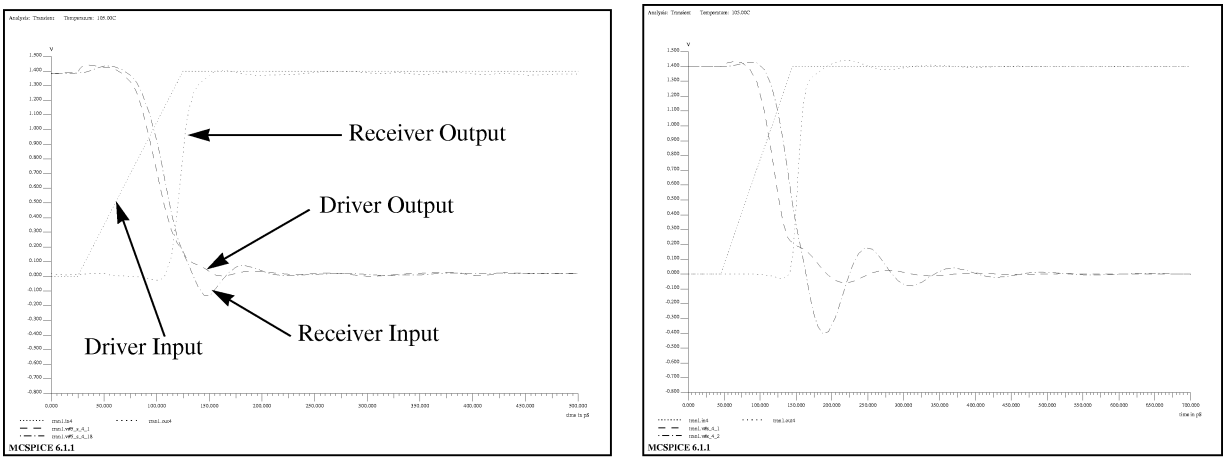

Fig. 10. PEEC-based (left) versus loop inductance model (right) for 350 μm $*$ 350 μm circuit area.

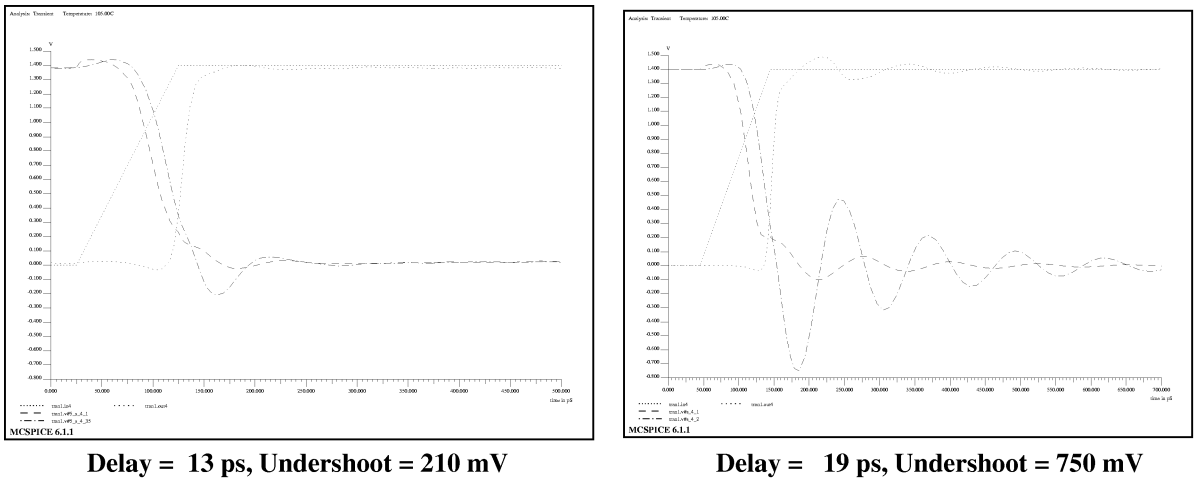

Fig. 11. PEEC-based (left) versus loop inductance model (right) for (700 μm $*$ 350 μm circuit area.

TABLE I
EFFECT OF MODEL COMPONENTS ON SIGNAL BEHAVIOR

Row	Experiment setup	Delay (ps)	Slope (ps)	Undershoot (mV)
1	Complete model	6.1	42	140
2	Remove device decoupling cap	4.5	45	157
3	Add extra decoupling cap	6.2	40	138
4	Include pad inductances	6.0	39	240
5	1 Vdd pad at A, 1 Gnd pad at B	6.3	42	152
6	No current sources	6.5	43	165

analysis is performed when 75 pF of explicit decoupling capacitance is added to the model. In this case, the inductive ringing is slightly decreased, while the interconnect delay is increased.

2) Package Inductance: In row 4 of Table I, the analysis results are shown when a package model is included in the analysis. The C4 pads were modeled as 100 μm by 100 μm vertical bars of 400 μm length. The inclusion of package inductance in the model introduces significant lower-frequency oscillations into the transient waveforms, as shown in Fig. 12(a), where the reported voltages are measured with respect to system ground. However, for the operation of the receiver gate, the gate input voltage with respect the local power supply grid connections is of more importance. When measuring with respect to the local ground node at the receiver gate, the voltage fluctuations are

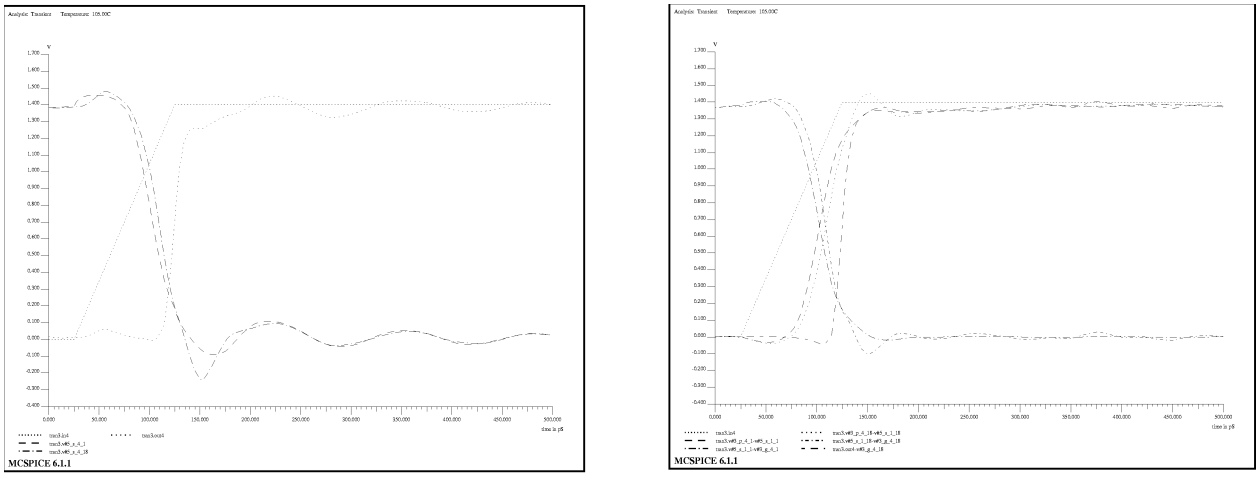

(a) (b)

Fig. 12. Signal response with and without package inductance.

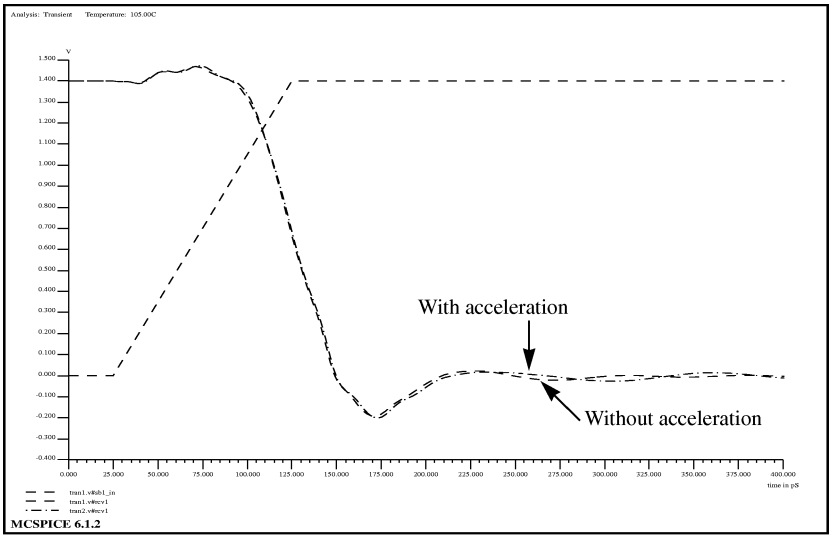

Fig. 13. Signal response with and without acceleration.

significantly reduced and the undershoot becomes 140 mV. The corresponding voltage waveforms are shown in Fig. 12(b). Note that the signal delay is relatively unchanged by the package inductance, and therefore the package model is primarily important when analyzing signal integrity issues.

3) Pad Number/Location: In the standard analysis, two V_{dd} and G_{nd} pads are used to connect the power grids. In row 5, the number of supply connection was reduced to one V_{dd} and G_{nd} pad. Reducing the number of pads worsens the IR drop. It also influences current return paths and, hence, the effective inductance of the signal net. Row 5 shows that a reduced number of pads increases the inductive undershoot significantly as well as slightly increasing the interconnect delay.

4) Current Sources: Finally, the analysis was performed without the current sources representing the quiescent currents in the power grid. Row 6 in Table I shows that the quiescent currents have a significant dampening effect on the inductive ringing of the signal net. The signal undershoot increased by 18% and the signal delay increased by 7% when the quiescent currents are removed.

C. Analysis of Acceleration Methods

In order to test the effectiveness and accuracy of the proposed acceleration techniques, three bus structures were analyzed both with and without acceleration. The tested bus topology is shown in Fig. 9 and the circuit model sizes ranged between 350×350 μm and 1050×350 μm. Table II shows the run time results when SPICE simulation was performed directly on the circuit model with a fully dense mutual inductance matrix, and when the block diagonal sparsification and accelerated reduced order modeling were first performed. The reported run time with acceleration, includes both the sparsification and PRIMA run time and the subsequent SPICE run time of the reduced model. The largest bus topology, shown in the third row, could not be analyzed without the discussed acceleration techniques. Since this structure is relatively small compared to large on-chip structures, it clearly demonstrates the need for the proposed acceleration. Fig. 13 shows that the waveforms for the 700×350 μm bus, with and without acceleration, almost overlap each other. Thus the accuracy of the proposed acceleration methods is very high.

TABLE II
RUN TIME RESULTS WITH AND WITHOUT ACCELERATION TECHNIQUES

Circuit area	Number of Partial Mutual Inductances		RunTime	
	w/o sparsification	w/ sparsification	w/o acceleration	w/ acceleration
350μm * 350μm	250k	55k	12 hr.	4 min.
700μm * 350μm	1,000k	110k	80 hr.	7 min.
1050μm * 350μm	2,400k	165k	*	8 min.

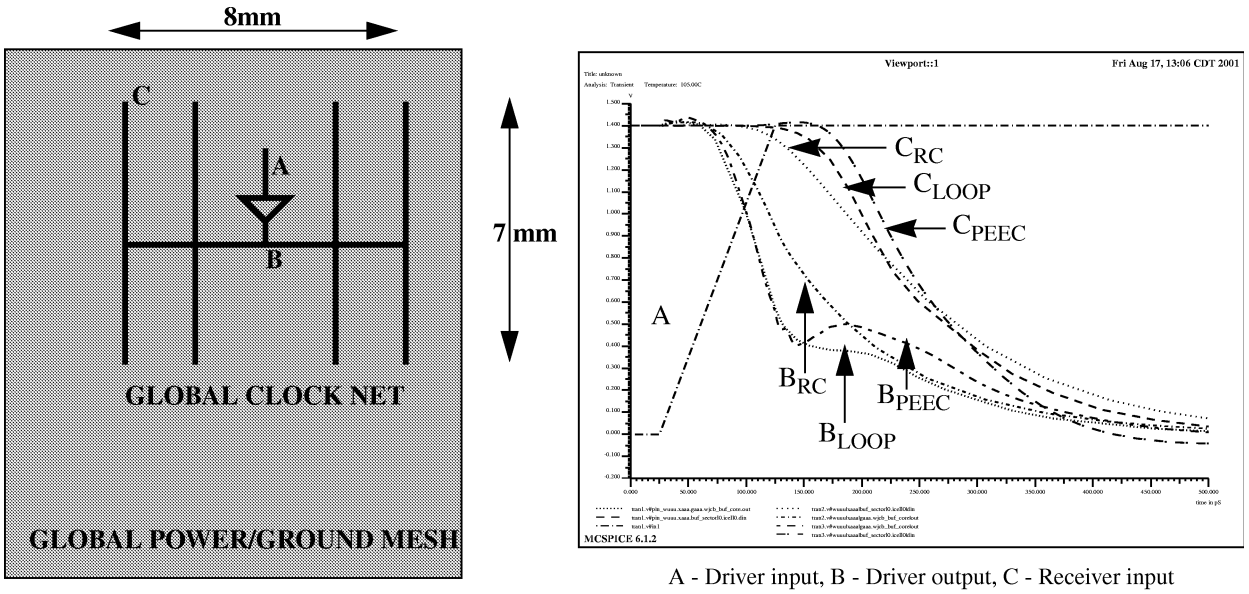

Fig. 14. Clock net topology and PEEC model versus loop model simulation results.

TABLE III
COMPARISON OF ANALYSES FOR A HIGH-PERFORMANCE CLOCK NET

	#R	#C	#L	#M	Delay (Diff w.r.t. RC)	Skew (Diff w.r.t. RC)	Run-time
RC	3k	6k	-	-	86ps	9ps	2min
PEEC	160k	400k	160k	8000k	131ps (45ps)	19ps (10ps)	60min
Loop	3k	6k	2k	-	116ps (30ps)	12ps (3ps)	4min

D. Comparison for the Top-Level Clock-Net of a High-Performance Microprocessor

Finally, we present analysis results for a large global clock net of a giga-hertz microprocessor. The topology of the clock net and the power grid is shown in Fig. 14(a). The clock net was simulated in three ways: with the proposed PEEC-based model and acceleration techniques, with the traditional loop-based analysis approach, and without inductance, using a standard *RC* model. The simulation results are shown in Fig. 14(b). The clock delay and skew for each of the three analyses are shown in Table III, as well as the different model sizes. Each entry shows the total measure of delay and skew, with the contribution due to inductance shown in parentheses. The results show that on-chip inductance has a significant impact on the delay and skew of the clock, and also shows that the loop based approach underestimates the delay and skew contributed by inductance by 33% and 77% respectively, for this particular example. The reported run times in Table III include the time for extracting the model, sparsification, reduced order model generation, and SPICE simula-

tion. The results demonstrate the effectiveness of the proposed acceleration techniques, in that it allows the analysis of a circuit model consisting of approximately 150 thousand circuit nodes, 720 thousand *RLC* elements and 8 million mutual inductances in only 60 minutes on a 400 MHz Sun UltraSparc-II compute server.

V. Conclusion

We have presented a new methodology for accurate modeling and efficient analysis of on-chip inductance in high-performance VLSI circuits. We proposed the inclusion of various parasitics including distributed interconnect resistance, inductance and capacitance, device decoupling capacitances, quiescent switching currents in the grid, pad connections, and pad/package inductance, along with techniques for extracting these elements. Simulation results show that our PEEC-based model more accurately determines the current distribution and hence inductive effects, while the traditional simplified loop inductance model significantly overestimates the inductive effects. Further, we have used the PEEC model to demonstrate the importance of the various model components on signal behavior. Since the PEEC model leads to a fully dense inductance matrix, we used block-diagonal matrix sparsification that guarantees the passivity of the sparsified circuit while maintaining good accuracy. We also employed reduced-order modeling using the PRIMA algorithm to allow the simulation of large circuit structures. To accelerate the reduced order modeling, we introduced a new formulation of the moment calculation that allows the factorization of the conductance matrix to be performed using efficient Cholesky factorization instead of the more time consuming LU decomposition. The combined sparsification and reduced order modeling approaches allow the analysis of global clock nets and busses for large IC designs with high accuracy and efficiency.

References

[1] A. Deutsch, P. W. Coteus, G. V. Kopcsay, H. H. Smith, C. W. Surovic, B. L. Krauter, and D. C. Edelstein, "On-chip wiring design challenges for gigahertz operation," *Proc. IEEE*, vol. 89, pp. 529–555, Apr. 2001.

[2] P. J. Restle and A. Deutsch, "Designing the best clock distribution network," in *Proc. Symp. VLSI Circuits*, Honolulu, HI, June 1998, pp. 2–5.

[3] S. Morton, "On-chip inductance issues in multiconductor systems," in *Proc. Design Automation Conference*, June 1999, pp. 921–926.

[4] B. Krauter and S. Mehrotra, "Layout based frequency depended inductance and resistance extraction for on-chip interconnect timing analysis," in *Proc. Design Automation Conference*, June 1998, pp. 303–308.

[5] A. Sinha and S. Chowdhury, "Mesh-structured on-chip power/ground: Design for minimum inductance and characterization for fast R, L extraction," in *Proc. Custom Integrated Circuits Conference*, May 1999, pp. 461–464.

[6] Y. Massoud, S. Majors, T. Bustami, and J. White, "Layout techniques for minimizing on-chip interconnect self-inductance," in *Proc. Design Automation Conference*, June 1998, pp. 566–571.

[7] M. Kamon, M. J. Tsuk, and J. K. White, "FASTHENRY: A multipole-accelerated 3-D inductance extraction program," *IEEE Trans. Microwave Theory Tech.*, pp. 1750–1758, September 1994.

[8] J. R. Phillips and J. K. White, "Efficient capacitance extraction of 3D structures using generalized pre-corrected FFT methods," in *Proc. IEEE 3rd Topical Meeting on Electrical Performance of Electronic packaging*, 1994, pp. 253–256.

[9] A. E. Ruehli, "Inductance calculations in a complex integrated circuit environment," *IBM J. Res. Dev.*, pp. 470–481, September 1972.

[10] H. Heeb and A. E. Ruehli, "Retarded models for PC board interconnects—Or how the speed of light affects your SPICE circuit simulation," in *Proc. International Conference on Computer-Aided Design*, Nov. 1991, pp. 70–73.

[11] L. He, N. Chang, S. Lin, and O. S. Nakagawa, "An efficient inductance modeling for on-chip interconnects," in *Proc. Custom Integrated Circuits Conference*, May 1999, pp. 457–460.

[12] B. Young, "Return path inductance in measurements of package inductance," *IEEE Trans. Compon., Packag., and Manufact. Technol.*, vol. 20, pp. 50–55, Feb. 1997.

[13] A. J. Sinclair and J. A. Ferreira, "Analysis and design of transmission-line structures by means of the geometric mean distance," *IEEE Africon*, pp. 1062–1065, Sept. 1996.

[14] F. W. Grover, *Inductance Calculations: Working Formulas and Tables*. New York: Dover, 1946.

[15] C. Hoer and C. Love, "Exact inductance equations for rectangular conductors with applications to more complicated geometries," *J. Res. Nat. Bur. Stand.*, pp. 127–137, April 1965.

[16] B. Krauter and L. T. Pileggi, "Generating sparse partial inductance matrices with guaranteed stability," in *Proc. International Conference on Computer-Aided Design*, November 1995, pp. 45–52.

[17] Z. He, M. Celik, and L. T. Pileggi, "SPIE: Sparse partial inductance extraction," in *Proc. Design Automation Conference*, June 1997, pp. 137–140.

[18] K. L. Shepard and Z. Tian, "Return-limited inductances: A practical approach to on-chip inductance extraction," *IEEE Trans. Comput.-Aided Design*, pp. 425–436, Apr. 2000.

[19] M. W. Beattie and L. T. Pileggi, "IC analyses including extracted inductance models," in *Proc. Design Automation Conference*, June 1999, pp. 915–920.

[20] B. Krauter, S. Mehrotra, and V. Chandramouli, "Including inductive effects in interconnect timing analysis," in *Proc. Custom Integrated Circuits Conference*, May 1999, pp. 445–452.

[21] K. Gala, D. Blaauw, J. Wang, M. Zhao, and V. Zolotov, "Inductance 101: Analysis and design," in *Proc. Design Automation Conference*, June 2001, pp. 329–334.

[22] A. Odabasioglu, M. Celik, and L. T. Pileggi, "PRIMA: Passive reduced-order interconnect macromodeling algorithm," in *Proc. International Conference on Computer-Aided Design*, 1997, pp. 58–65.

[23] A. Devgan, H. Ji, and W. Dai, "How to efficiently capture on-chip inductance effects: Introducing a new circuit element K," in *Proc. International Conference on Computer-Aided Design*, November 2000, pp. 150–155.

[24] L. He and K. M. Lepak, "Simultaneous shield insertion and net ordering for capacitive and inductive coupling minimization," in *Proc. International Symposium on Physical Design*, April 2000, pp. 55–60.

[25] G. Zhong, C.-K. Koh, and K. Roy, "A twisted-bundle layout structure for minimizing inductive coupling noise," in *Proc. International Conference on Computer-Aided Design*, November 2000, pp. 406–411.

[26] M. W. Beattie and L. T. Pileggi, "Modeling magnetic coupling for on-chip interconnect," in *Proc. Design Automation Conference*, June 2001, pp. 335–340.

[27] J. H. Chern, "Multilevel metal capacitance models for CAD design synthesis systems," *IEEE Electron Device Lett.*, pp. 32–40.

[28] R. Panda, D. Blaauw, R. Chaudhry, B. Young, and R. Ramaraju, "Model and analysis for combined package and on-chip power grid simulation," in *Proc. International Symposium on Low Power Electronics and Design*, July 2000, pp. 179–184.

[29] G. H. Golub and C. F. Van Loan, *Matrix Computations*, Second ed. Baltimore, MD: The Johns Hopkins University Press, 1989.

[30] B. Tutuianu, D. Lehther, M. Pandey, and R. Baldick, "Efficient RLC macromodels for digitial IC interconnect," in *Proc. VLSI Conference*, 1999.

[31] *Star-RCXT User Manual*, Avant! Corporation,, 2001.

[32] *Maxwell User Manual*, Ansoft Corporation,, 2000.

Kaushik Gala received the B.E. degree in Instrumentation from College of Engineering, Pune, India, and the M.S. degree in electrical engineering from the University of Minnesota, Minneapolis. He is currently pursuing the M.B.A. degree at the University of Texas, Austin.

He is with the Advanced Tools Group at Motorola-SPS, Austin, TX, since January 1999. He currently works on research and development of design automation tools and methodologies, primarily in the area of signal integrity for high-performance VLSI circuits. His research interests include on-chip inductance and noise analysis and avoidance, clock skew analyis, power grid analysis, and studying the effects of process variations on circuit performance and reliability.

David Blaauw received the B.S. degree in physics and computer science from Duke University in 1986, the M.S. degree in computer science from the University of Illinois, Urbana, in 1988, and the Ph.D. degree in computer science from the University of Illinois, in 1991.

He was with the Engineering Accelerator Technology Division, IBM Corporation, Endicott, as a Development Staff Member, until August 1993. From 1993 to August 2001, he was with Motorola, Inc., Austin, TX, were he was the manager of the High Performance Design Technology group. Since August 2001, he has been with the faculty of the University of Michigan as an Associate Professor. His work has focused on VLSI design and CAD with particular emphasis on circuit analysis and optimization problems for high performance and low power designs.

Dr. Blaauw was the Technical Program Chair and General Chair for the International Symposium on Low Power Electronic and Design in 1999 and 2000, respectively, and was the Technical Program Co-Chair and member of the Executive Committee the ACM/IEEE Design Automation Conference in 2000 and 2001.

Vladimir Zolotov (M'97) received Master degree on electrical engineering from Moscow Institute of Electronics, Moscow, Russia, and the Ph.D. degree on electrical engineering from Scientific Research Institute of Micro Devices, Moscow.

He has been with Advanced Tools group, Motorola, Inc., Austin, TX, since 1998. He is involved in the development of EDA tools and methodology for high performance VLSI designs. Previously, he was with Moscow Research Laboratory, Motorola, Inc., Moscow. His research interests include signal integrity, reliability, on chip inductance, timing analysis, and optimization of VLSI.

Pravin M. Vaidya, photograph and biography not available at the time of publication.

Anil Joshi received the B.S. degree in electronics and communications engineering, Osmania University, India, in 1980, the M.S. degree in industrial engineering, NITIE, Bombay, 1982, the M.S. degree in computer science, Polytechnic University of New York, in 1989, and the Ph.D. degree in computer science, University of Illinois at Urbana-Champaign, 1997.

He has over 15 years of experience in Computer Programming, Systems Analysis and Design, Software Project Management, Database Design, Online Transaction Processing, in both business applications and systems. His current research interests include iterative and direct methods for solving large sparse systems of linear equations, which arise in engineering, scientific, and financial applications.