Full-Chip, Three-Dimensional, Shapes-Based RLC Extraction

Dipak Sitaram, Yu Zheng, *Member, IEEE*, and Kenneth L. Shepard, *Senior Member, IEEE*

***Abstract*—In this paper, we report the development of a full-chip, three-dimensional, shapes-based, resistence–inductance–capacitance extraction tool, which was developed as part of a university–industry collaboration. The technique of return-limited inductances is used to provide a sparse, frequency-independent inductance and resistance network with self-inductances that represent sensible "nominal" values in the absence of mutual coupling. Mutual inductances are extracted for accurate crosstalk analysis. The tool exploits high-capacity scan-band techniques and disk caching. Accuracy is validated by comparison with full-wave finite-element field solvers.**

***Index Terms*—Inductance, interconnect modelling, parasitic extraction.**

I. Introduction

With technology scaling, on-chip frequencies are increasing as device f_T values exceed 50 GHz. In digital integrated circuits, slew times are being driven below 50 ps, corresponding to frequency content approaching 10 GHz. For many nets, the clock being the most notable [1], inductance must be included to accurately predict rise and fall times and delays in timing analysis. If an inductive net is overdriven, an underdamped ringing response can be observed, which can result in functional failure in receiving circuits or produce reliability problems through gate oxide stress. Moreover, inductive coupling, along with capacitive coupling, can be a significant source of noise on quiet nets due to the switching of nearby perpetrators.

In analog integrated circuits (ICs), on-chip frequencies for wireline and wireless applications are also pushing beyond 10 GHz, in optical communications circuits [2] and in RF circuits [3]. On-chip inductance extraction techniques have already been applied to spiral inductors [4]–[6], but on-chip transmission lines are finding places in both distributed feedback amplifiers [7] and oscillators [8]. Increasingly, the parasitic inductance of on-chip interconnect is also becoming a concern.

Manuscript received May 8, 2003; revised August 25, 2003. This work was supported in part by the National Science Foundation under Grant CCR-97-34216 and by a grant from Cadence Design Systems. This paper is an expanded version of a work presented at the International Conference on Computer-Aided Design, San Jose, CA, in 2000 and 2002. This paper was recommended by Associate Editor S. Sapatnekar.

D. Sitaram is with Cadence Design Systems, New Providence, NJ 07974 USA.

Y. Zheng was with the Columbia Integrated Systems Lab, Department of Electrical Engineering, Columbia University, New York, NY 10027 USA. He is now with Cadence Design Systems, San Jose, CA 95134 USA.

K. L. Shepard is with the Columbia Integrated Systems Lab, Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: shepard@ee.columbia.edu).

Digital Object Identifier 10.1109/TCAD.2004.826545

High-frequency interconnect analysis has traditionally been relegated to full-wave Maxwell's equations solvers. Field solvers, such as Ansoft's HFSS, which find broad use in the microwave community, solve the differential form of Maxwell's equations with volume discretization and finite-element techniques. By contrast, boundary-element techniques are integral-equation-based solutions, which rely on discretizing the sources. Method-of-moment-based solvers [9], such as Agilent's Momentum, are widely employed in the microwave community. In the IC and package communities, integral-equation-based solution are popularly represented in terms of partial-element equivalent circuits (PEEC) which can be solved with circuit simulation engines [10], [11]. If the coupling distances are short, relative to the wavelength (that is, the distances between conductor segments that are magnetically or electrically coupled), then the quasi-static approximation applies and no retardation is necessary in the analysis [12]. Fast integral equation solvers have been developed to provide accurate electrostatic, magnetostatic, and full-wave integral equation solutions [13]–[15]. Such solvers, while achieving high accuracy, are still very slow, except for the smallest problem sizes, and the resulting formulations are still intractably dense for large designs. As a result, parasitic extraction engines, which compromise some accuracy to achieve full-chip capacity, are being pushed to provide quasi-static high-frequency extraction capabilities. That is, the extractors are being extended to add inductance to the more familiar resistance and capacitance effects. This is the focus of the work described in this paper.

Full-chip extraction engines generally use pattern matching and interpolation from look-up tables to calculate capacitance, look-up tables that are generally calculated with use of fast integral equation solvers. Capacitances have very strong geometry dependence. Therefore, considerable care is necessary to achieve accurate values. The extracted capacitances are intrinsically sparse because capacitances are very "short-ranged." Electric field lines effectively terminate on the closest conductors and more distant couplings are negligibly small. Physically, these small capacitances can be discarded without any effect on the passivity of the resulting network.

Inductances, by contrast, have relatively weak geometry dependence, allowing for accurate calculation with relatively simple analytic formulas. Magnetic coupling, however, is dense, leading immediately to circuit-level intractability for large problem sizes. The inductance matrix is dense physically because magnetic fields induced by a current can spread much further and must be "terminated" by eddy currents induced in nearby conductors. Furthermore, the partial inductance formulation is mathematically dense. Partial inductances are defined

by flux areas that extend to infinity. Physically meaningful loop inductances are only obtained when the more distant flux areas are "cancelled" out by distant partial mutual inductances. Small mutual inductances cannot be discarded without disrupting the passivity of the network [16].

There have been two recent approaches proposed to "sparsify" the inductance problem, each with its own advantages and disadvantages [17]. The first approach recognizes that the inverse of the inductance matrix (variously called the K element [18][1] and the susceptance matrix [19]) is mathematically "local;" that is, one does not have to combine many distant couplings to obtain a physical answer. As such, as for the case of the capacitance matrix, elements can be discarded without disturbing the passivity of the network. While the formulation is now mathematically sparse, the problem is still physically dense and the approach provides no mechanism for determining the interaction window size. The most significant disadvantage of the technique, however, is that existing simulation tools (and designer intuition) are based on the concept of inductance rather than "inverse inductance."

An alternative to the inverse-inductance approach is the approach of return-limited inductances [20], [21], which is employed in the Assura RCX-PL extraction engine [22] described here. In this technique, the power-ground network is used to divide the chip into interaction regions. This approach recognizes the fact that power-ground nets are always available as fail-safe high-frequency current returns so that eddy currents in most cases will not be induced significantly beyond the nearest power-ground nets. The power-ground network is modeled *implicitly* in this approach with the power-ground nets acting as virtual ground planes. By "implicitly," we mean that an equivalent RLCK network is generated for the signal lines, which includes the effects of the power-ground lines without requiring an explicit network representation of these wires, a situation which could easily result in an intractable analysis.

Because whether the inductance of a given line is important depends, in general, on the entire line and its environment including the driving and receiving circuits, it is often difficult to know *a priori* all of the nets that need this consideration. This is much the same problem associated with a "selective-net" approach to detailed resistance and capacitance extraction, and most modern design methodologies have embraced "full-chip" analysis and extraction for interconnect [23], [24]. For inductance extraction, this puts the burden on the extractor to include inductance only where it is electrically significant. More details of this are presented in Section IV.

In this paper, we describe the practical details of implementing return-limited inductance modeling [21] into a full-chip extraction engine, validated with careful comparisons with full-wave finite-element field solvers. In addition to the basic implementation, we describe a number of features that can be selectively enabled by the user to achieve higher accuracy at the expense of longer extraction and analysis runtimes:

- *Consideration of eddy current losses in the power-ground network*. Earlier return-limited inductance extraction approaches [21] ignored eddy current losses in the power-ground distribution. This assumption is frequently not justified because high-frequency current returns may choose more resistive return paths through power-ground lines to minimize inductance, resulting in a higher effective resistance for the line. We show how these effects can be considered while preserving the *implicit* nature of the extraction as described above.
- *Consideration of nonuniform current distributions in wires due to skin and proximity effects*. At 20 GHz, the skin depth in aluminum wires is approximately 0.7 μm. For coplanar structures, proximity-effect current crowding along the width of the wire at high frequencies is also a concern. There is a common misconception that these current crowding effects for coplanar structures are significant when the wire width is greater than the skin depth; when the wire width is much greater than the wire thickness, these crowding effects actually become important at higher frequencies than predicted by the skin depth. Skin and proximity effects can be modeled within the context of the *implicit* network representation if the user seeks this level of accuracy.
- *Consideration of eddy current losses in the substrate*. In the metal-rich environment of digital ICs, eddy current loss in the substrate is not a concern because the on-chip power supply is always available as a lower impedance current return than the silicon substrate. In analog designs, however, this may not always be the case since routed power-ground distributions are more common and a sparse metal environment is often present in which the silicon substrate (particularly if an epitaxial substrate is used) may be the lowest impedance current return, often at the cost of considerable resistive losses. In general, this will be a concern only in the case of heavily-doped epitaxial substrates, which are commonly employed (particularly for digital design) for latch-up immunity. The user may choose to consider such substrate eddy current losses while preserving *implicit* extraction.

In the case of high-substrate resistivities, capacitances terminating on the substrate result in displacement currents in the substrate which can pick up additional resistive losses on their way to be collected by substrate plugs. Furthermore, capacitances to the substrate can become frequency-dependent since at higher frequencies, electric field lines will generally extend further into the substrate before terminating on charge, resulting in smaller capacitances. Modeling this would require a detailed resistance and capacitance extraction of the substrate along with the interconnect extraction. In this case, *explicit* modeling of the substrate would be required and only small problem sizes in which high accuracy is required could be accommodated. As a result, we do not model these effects in Assura RCX-PL.

Section II reviews the idea of return-limited inductance extraction, including modeling eddy current losses in the power-ground network and modeling skin- and proximity-effect current crowding in the wires. A multilayer Greens' function approach to model eddy currents in the substrate is

[1]The use of the term K element was perhaps a poor one since SPICE already defines a K element as simply a (normalized) mutual inductances. We use this definition of K in this paper.

described in Section III, while Section IV considers how the concomitant frequency dependence of the inductance and resistance resulting from implicit eddy-current loss mechanisms and current crowding is handled with frequency-independent elements with an equivalent ladder network fit. Implementation details in Assura RCX-PL are considered in Section V, along with filtering approached employed to limit extraction complexity. Section VI presents comparison of Assura RCX-PL extraction results with full-wave solution. A larger example is also presented to demonstrate the performance and capacity of the tool. Section VII concludes the paper.

II. Return-Limited Inductance Extraction and Power-Ground Eddy Current Loss Modelling

Practical inductance extraction (through a set of meaningful approximations based only on geometry) must provide a sparse inductance and resistance network with self-inductances that represent sensible "nominal" values in the absence of mutual coupling. Eddy current losses in the substrate and power-ground should be handled implicitly for tractability and the concomitant frequency dependence that comes with this implicit treatment modeled with frequency-independent elements to maintain compatibility with SPICE simulation. Return-limited inductance extraction achieves all of these goals. We refer the reader to [20] and [25] for details on the technique, but we review the aspects of the approach most important to its implementation here.

A. Return-Limited Inductance Extraction

In return-limited inductance extraction, the power-ground network of the chip is used to divide the interconnect into a set of disjoint interaction regions. Self-inductances are defined by loops formed with the nearest parallel power-ground lines. Signal lines within an interaction region are magnetically coupled as loop inductances, and signal lines contained in two different interaction regions do not couple. In this approach, the power-ground lines are implicit in the extraction and when losses in these lines are considered, the resulting signal-line models acquire a frequency-dependence because the exact distribution of current in the power-ground lines (and the associated eddy current losses) become frequency-dependent. At low frequency, low loss (but high inductance) current distributions are assumed, while at high frequency, low inductance (but high loss) current distributions are assumed.

We note than an important part of full-wave analysis is understanding the interplay of displacement, conduction, and eddy currents, and this analysis involves the simultaneous quasi-static PEEC extraction of the signal lines and the power-ground distribution. In this context, to get a truly accurate analysis of the high-frequency behavior of the power-ground nets one must correctly model the wires of the distribution as well as all the sources of on-chip decoupling capacitance. Depending on the amount of on-chip decoupling, package modeling might also be necessary. The resulting extraction and analysis easily become intractable. A better approach to handle large problem sizes (and that taken in return-limited inductance extraction) is to assume that the power-ground distribution has been well-designed and has a very low impedance compared to the signal lines being analyzed. For power-supply integrity considerations, this is achieved in practice with power-ground grid structures and adequate decoupling capacitance. With the assumption of a low-impedance power-ground distribution, we can ease the full-wave requirements on the power-ground distribution by treating capacitances to power and ground lines as capacitances to ideal ground. Furthermore, eddy currents induced in the nearest neighbor power-ground lines can be "sourced" from the power-ground distribution with zero impedance.

Within Assura RCX-PL, then, a given interaction region consists of a set of "clusters," each containing a signal line and its associated parallel power-ground lines (referred to as the power-ground aggregate), which define the return-limited loop self-inductance for the signal line. These clusters are then magnetically coupled to each other. Two such clusters are shown in Fig. 1(a). In this example, each cluster consists of a signal line and power-ground return and is characterized by an appropriate set of segment resistances and partial inductances. The current flowing through the signal line of the ith cluster is given by I_i, giving a voltage drop of V_i. The current flowing through the J ground returns of the cluster is given by I_j. Each cluster is also augmented (not shown in Fig. 1) by a "pseudoreturn" to crudely model all of the (low-resistance, but high-inductance) current returns outside the interaction region; this enables a frequency-dependent transition from a dc resistance defined only by the resistance of the signal line (power grid is lossless at dc) to a high-frequency resistance defined by returns confined to the interaction region. Low-frequency inductances are not correctly modeled, but they have no significance in determining interconnect response. The segment currents and voltages, independent of the cluster partitioning, are related by an impedance matrix

$$\boldsymbol{V} = \boldsymbol{Z}\boldsymbol{I} \tag{1}$$

where $\boldsymbol{Z}$, given by $\boldsymbol{R} + j\omega\boldsymbol{L}$. $\boldsymbol{R}$ as the resistance matrix, is diagonal. $\boldsymbol{L}$ is the dense matrix of partial inductances.

Let $\sum_j I_j = I_g$ be the (net) current in the power-ground lines of a cluster. The vector of the I_g currents, one for each cluster, is given by $\boldsymbol{I}_g$. $\boldsymbol{V}_i$ and $\boldsymbol{I}_i$ are the signal-line voltages and currents, respectively. From the assumption of a low-impedance power-ground distribution [shown by the shunts across the power-ground nets in Fig. 1(b)], we see that

$$\begin{pmatrix} \boldsymbol{I}_i \\ \boldsymbol{I}_g \end{pmatrix} = \boldsymbol{B}\boldsymbol{Z}^{-1}\boldsymbol{B}^T \begin{pmatrix} \boldsymbol{V}_i \\ 0 \end{pmatrix} \tag{2}$$

where $\boldsymbol{B}$ is an incidence matrix. Each column of $\boldsymbol{B}$ corresponds to a signal line or a power-ground line; each row of $\boldsymbol{B}$ corresponds to a signal line or power-ground line aggregate in one cluster. The ith column of B is all zero except for the ones in rows corresponding to the same signal line or the aggregate containing the power-ground line.

Inverting the current-voltage relation yields

$$\begin{pmatrix} \mathbf{V}_i \\ \mathbf{0} \end{pmatrix} = (\boldsymbol{B}\boldsymbol{Z}^{-1}\boldsymbol{B}^T)^{-1} \begin{pmatrix} \mathbf{I}_i \\ \mathbf{I}_g \end{pmatrix}. \tag{3}$$

Solving for $\boldsymbol{V}_i$ yields

$$\mathbf{V}_i = \boldsymbol{S}(\boldsymbol{B}\boldsymbol{Z}^{-1}\boldsymbol{B}^T)^{-1}\boldsymbol{S}^T\mathbf{I}_i \tag{4}$$

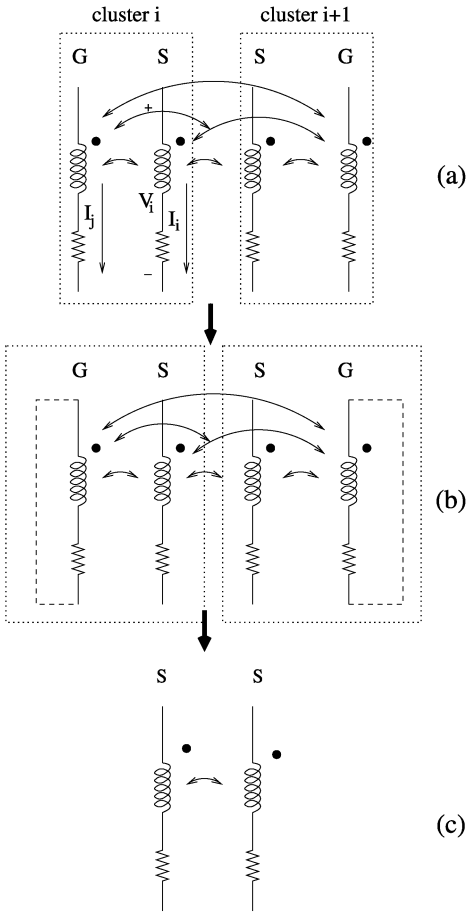

Fig. 1. Interaction-region cluster defined by a set of signal lines and the associated power-ground lines.

yielding the equivalent impedance matrix of the signal lines in the interaction region [as shown in Fig. 1(c)]

$$\boldsymbol{Z}_{eq} = \boldsymbol{S}(\boldsymbol{B}\boldsymbol{Z}^{-1}\boldsymbol{B}^{T})^{-1}\boldsymbol{S}^{T} \tag{5}$$

where $\boldsymbol{S}$ is an incidence matrix. Each column of $\boldsymbol{S}$ corresponds to a signal line or aggregate of power-ground lines; each row of $\boldsymbol{S}$ corresponds to a signal line. The element of the ith column and jth row of $\boldsymbol{S}$ is 1 if the column and row correspond to the same signal line; it is -1 if the column corresponds to the power-ground aggregate and the row corresponds to the associated signal line of the same cluster. Equation (5) is effectively a "double-inverse" operation, dominated by the time taken to perform the LU factorization of $\boldsymbol{Z}$. Restricting the interaction region size is essential for extraction performance. In Section IV, we show how this (frequency-dependent) impedance is represented in the SPICE extracted netlist.

B. Geometry-Based Inductance Matrix Decomposition

In return-limited inductance extraction, interaction regions are created with simple, geometry-based matrix decomposition rules, which we refer to as halo rules [20], [21]. A horizontal wire segment is one in which the current flow is known to be horizontal, while a vertical segment is one in which the current is known to be vertical. The halo of a segment consists of the six semi-infinite subregions shown in Fig. 2. The horizontal halo (for current flow in the y direction in this case) consists only of

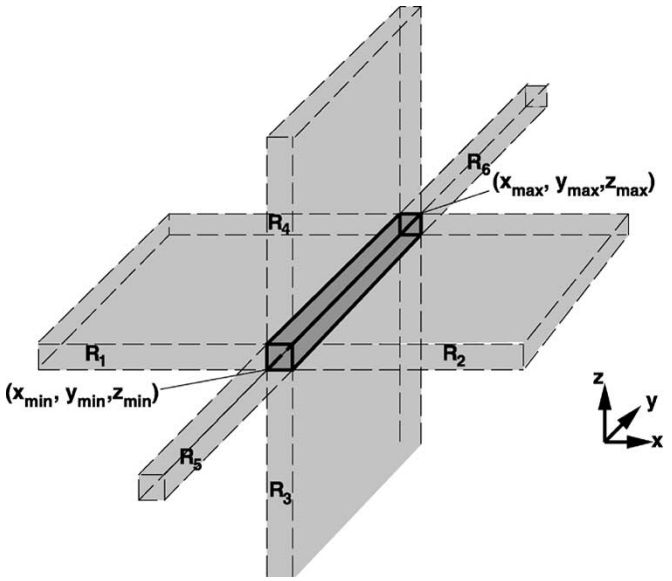

Fig. 2. Halo of a given segment consists of six semi-infinite regions.

regions R_3, R_4, R_5, and R_6, while the vertical halo (for current flow in the x direction) consists only of regions R_1–R_4. Clearly, the terms "horizontal" and "vertical" are interchanged when the design is rotated 90 degrees. The halo rules must, therefore, be invariant under such an interchange.

The halo rules are given as follows.

- Horizontal and vertical signal line segments are treated independently since they do not inductively couple to each other. Segments with horizontal currents can only couple inductively with other segments with horizontal currents. Similarly, vertical segments can only couple inductively with other vertical segments.[2]
- Horizontal halos of power and ground are "blocked" by horizontal signal segments while vertical halos of power and ground are "blocked" by vertical signal segments. If the halos are viewed as columnated beams emanating orthogonally from each face of a segment, then blocking occurs whenever these beams are interrupted by another segment.
- Inductive coupling between two horizontal segments is nonzero if and only if it is possible to connect two segments by a path which does not cross the horizontal halo of any ground or supply line. Similarly, inductive coupling between two vertical segments is nonzero if and only if it is possible to connect two segments by a path which does not cross the vertical halo of any supply or ground line.[3]

These halo rules divide the chip interconnect into a collection of disjoint *horizontal interaction regions* defined by the nonblocked horizontal halos of the power and ground distribution. Horizontal segments must be contained within the same *horizontal interaction region* to inductively couple. Independently,

[2]This means that the extraction can be done in two separate "passes," one to extract the horizontal segments and one to extract the vertical segments.

[3]Since when doing vertical signal line extraction, for example, we only need to consider vertical halos of ground or supply lines and these halos can only be blocked by vertical signal segments, horizontal signal segments do not have to be considered at all.

the chip is also divided into a collection of disjoint vertical interaction regions defined by the nonblocked vertical halos of the power and ground distribution. Similarly, vertical segments must be contained within the same *vertical interaction region* to inductively couple. The halo rules define the "baseline" for inductance sparsification. Assura RCX-PL allows interaction regions to be merged under user control when greater accuracy is required or structures problematic for the halo rules are encountered.

C. Modeling Skin and Proximity Effects

Nonuniform current distributions across wire cross sections are handled with a well-known volume filament decomposition [26]. The conductor cross section is divided into several filaments, each treated as a separate conductor. Equation (5) continues to define the equivalent impedance matrix. In this case, while each row of $\boldsymbol{B}$ still corresponds to a signal line or a power-ground lines aggregate in one cluster, a column of $\boldsymbol{B}$ may corresponds to a filament of a signal line or a power-ground line. Volume-filament decomposition results in larger $\boldsymbol{Z}$ and $\boldsymbol{B}$ matrices, producing longer calculation times. The dimensions of $\boldsymbol{S}$ and $\boldsymbol{Z}_{eq}$ matrices remain the same. The additional frequency dependence in $\boldsymbol{Z}_{eq}$ resulting from this skin and proximity-effect modeling is also represented in the ladder networks introduced in Section IV.

III. Magnetostatic Substrate Loss Modelling

In metal-sparse environments with heavily-doped (epitaxial) substrates, it is occasionally necessary to model eddy current losses in the substrate. Such a metal environment is not expected to be common for digital designs but may characterize analog designs with routed power-ground distributions. There are two basic modeling approaches one could take. One could mesh the substrate based on a volume filament approach [27], modeling the substrate as (in general) a coupled RLCK mesh (i.e., a full PEEC model, tantamount to a full-wave solution within the quasi-static approximation [10]) The advantage of this technique is that it is very accurate since all 3D effects can be handled (e.g., differences in substrate effects due to well diffusions and the influence of well or substrate plugs). Substrate losses can be accurately modeled in both the electrostatic and magnetostatic problems, and their interactions are also modeled. The main disadvantage to this approach, despite its accuracy, is that extractions will be "clogged" with complex substrate networks. As in the case of explicit treatment of power-ground, it will quickly result in a computational intractable problem for all but the most trivial problem sizes. An alternative is to treat the substrate *implicitly* by means of a Green's function treatment and combine this with the implicit power-ground treatment of return-limited inductance extraction.

To consider substrate effects, we return to Maxwell's equations. Ignoring the displacement term (quasi-static approximation), the magnetic field is determined by

$$\nabla \times \boldsymbol{B} = \mu \boldsymbol{J} \tag{6}$$

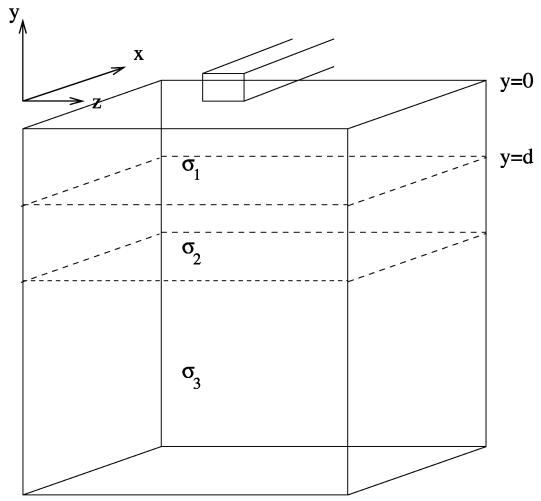

Fig. 3. Multilayer substrate. Metal layers are routed in layer 0. $y = 0$ defines the substrate-oxide interface.

where the current density $\boldsymbol{J} = \sigma \boldsymbol{E} + \boldsymbol{J}_{\rm src}$, where $\boldsymbol{J}_{\rm src}$ is the applied current density. The (time-harmonic) fields are related to the scalar and vector potentials by

$$\boldsymbol{E} = -j\omega \boldsymbol{A} - \nabla \phi \tag{7}$$

$$\boldsymbol{B} = \nabla \times \boldsymbol{A} \tag{8}$$

Using (7) and (8) (and Coulomb gauge) in (6) yields the following relation for the magnetic vector potential:

$$\nabla^2 \boldsymbol{A} = j\mu\omega\sigma \boldsymbol{A} - \mu \boldsymbol{J}_{\rm src} - \mu\sigma\nabla\phi \tag{9}$$

If we can assume the substrate is well plugged to a low-impedance power-ground distribution, it can be regarded as an equipotential and the last term on the right-hand-side of (9) can be ignored. All of the voltage induced (magnetically) in the substrate is dropped across the resistance of the substrate (or equivalently, the eddy currents of the substrate are sourced losslessly). With this assumption, capacitances to the substrate can be regarded as capacitances to the ideal ground reference, as is done for the case of capacitors coupled to the power-ground distribution. This leads to an implicit treatment of the substrate that is equivalent to that applied to the power-ground nets. This treatment ignores losses as displacement currents in the substrate are collected by plugs. The solution of (9) may then be written in integral form as

$$\boldsymbol{A}(\boldsymbol{r}) = \mu \int_V \boldsymbol{J}_{\rm src}(\boldsymbol{r}') G(\boldsymbol{r}, \boldsymbol{r}') d^3\boldsymbol{r}' \tag{10}$$

where $G(\boldsymbol{r}, \boldsymbol{r}')$ is the Green's function, found by solving the following equation:

$$\nabla^2 G(\boldsymbol{r}, \boldsymbol{r}') = -\delta(\boldsymbol{r} - \boldsymbol{r}') + j\mu\omega\sigma G(\boldsymbol{r}, \boldsymbol{r}'). \tag{11}$$

From (10), the Green's function allows one to calculate the magnetic field at the field point $\boldsymbol{r}$ as a result of a current at the source point $\boldsymbol{r}'$.

Multilayer Green's functions techniques are well known [28]–[30]. If one assumes a multilayer substrate stretching horizontally to infinity as shown in Fig. 3, then axial symmetry

applies and (11) can be reduced to a two-variable problem in cylindrical coordinates. Because this treatment assumes a uniform two-dimensional (2-D) substrate profile, three-dimensional (3-D) features such as substrate laminations sometimes used to reduce eddy current losses cannot be modeled. The spectral-domain transform can then be reduced to a one-dimensional Hankel transform [31]. We instead choose to keep the Green's function in Cartesian coordinates because we compute the filamentary inductance in the spectral domain, where axial symmetry no longer applies, before computing the transformation to spatial coordinates with a 2-D fast Fourier transform (FFT). This provides us with a mechanism for a quick precharacterized look-up table for determining inductances directly in the presence of a uniform multilayer substrate.

If the source point (x', y', z') and field point (x, y, z) are in the top oxide layer, denoted as $k = 0$, it is straightforward to show that the Green's function between points in layer 0 can be expressed as the double integral

$$G(\boldsymbol{r},\boldsymbol{r}') = \int_0^\infty dm \int_0^\infty \mathrm{dncos}(mx)\cos(mx')\cos(nz)\cos(nz') \times \frac{\left(\beta_0^u e^{\gamma_{\mathrm{mn}}^0 \max(y,y')} + \Gamma_0^u e^{-\gamma_{\mathrm{mn}}^0 \max(y,y')}\right)}{2\pi^2\gamma_{\mathrm{mn}}^0\left(\beta_0^l\Gamma_0^u - \beta_0^u\Gamma_0^l\right)} \times \left(\beta_0^l e^{\gamma_{\mathrm{mn}}^0 \min(y,y')} + \Gamma_0^l e^{-\gamma_{\mathrm{mn}}^0 \min(y,y')}\right) \quad (12)$$

where ${\gamma_{\mathrm{mn}}^k}^2 = m^2 + n^2 + j\omega\sigma_k\mu$. σ_k is the conductivity of layer k and μ is the permeability of free-space (all the materials are assumed to be nonmagnetic). $\beta_0^{u,l}$ and $\Gamma_0^{u,l}$ are coefficients determined by satisfying the appropriate boundary conditions at each material interface. This follows from similar derivations for the electrostatic problem [28], [29].

As shown in Fig. 3, $y = 0$ defines the oxide-silicon interface. To satisfy the requirement that the Green's function remain bounded as $y \to \infty$, $\beta_0^u = 0, \Gamma_0^u = 1$. Furthermore, if the bottom layer of the substrate (layer M) extends to $y \to -\infty$, $\beta_M^l = 1$, and $\Gamma_M^l = 0$.[4] With d_k being the y-distance to the interface between substrate layers k and $k+1$, the values of β_0^l and Γ_0^l can be found from the recursive formula

$$\begin{pmatrix} \beta_k^l \\ \Gamma_k^l e^{-2\gamma_{\mathrm{mn}}^k d_k} \end{pmatrix} = \boldsymbol{A}\begin{pmatrix} \beta_{k+1}^l \\ \Gamma_{k+1}^l e^{-2\gamma_{\mathrm{mn}}^k d_{k+1}} \end{pmatrix} \quad (13)$$

where the elements of $\boldsymbol{A}$ are given by

$$\boldsymbol{A} = \begin{pmatrix} a & b \\ b & a \end{pmatrix} \quad (14)$$

and a, b are expressed as

$$a = \frac{1}{2}\left(1 + \frac{\gamma_{\mathrm{mn}}^{k+1}}{\gamma_{\mathrm{mn}}^k}\right) e^{\left(\gamma_{\mathrm{mn}}^{k+1} - \gamma_{\mathrm{mn}}^k\right)d_k} \quad (15)$$

$$b = \frac{1}{2}\left(1 - \frac{\gamma_{\mathrm{mn}}^{k+1}}{\gamma_{\mathrm{mn}}^k}\right) e^{\left(\gamma_{\mathrm{mn}}^{k+1} - \gamma_{\mathrm{mn}}^k\right)d_k}. \quad (16)$$

[4]For very high-resistivity substrates, one could argue that the backside of the wafer must be modeled, for example, as an ideal ground place. If the backside were at $y = -D$, this would require that the Green's function (and the magnetic field) vanish for $y < -D$. We choose to model the substrate as infinitely thick in our context because of the presumption that other interconnect layers would always be present as a favored current return over a backside groundplane.

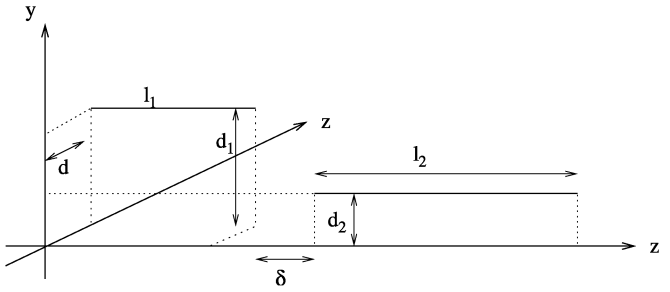

Fig. 4. Geometry and spacings of the coupled filaments.

We have found this formulation to be numerically robust.

We now use this multilayer Green's function to derive the partial inductance for filamentary segments ("dressed" by the presence of the substrate). The current density of a filamentary current I at $y = d_1, x = 0, -l_1/2 \le z \le l_1/2$ in the z direction is given by

$$\boldsymbol{J}(x',y',z') = I\delta(x')\delta(y'-d_1)\left[u\left(z'+\frac{l_1}{2}\right) - u\left(z'-\frac{l_1}{2}\right)\right]\hat{\boldsymbol{z}} \quad (17)$$

Fig. 4 shows the geometry of the source filament and coupled filament for this calculation. Combining (10), (12), and (17) yields the following expression for the magnetic vector potential

$$\boldsymbol{A}(x,y,z) = \hat{\boldsymbol{z}}\left(\frac{2\mu}{\pi^2}\right)\int_0^\infty \mathrm{dm}\int_0^\infty \mathrm{dncos}(mx)\cos(nz)\frac{\sin\left(\frac{nl_1}{2}\right)}{n} \times \frac{\left(\beta_0^u e^{\gamma_{\mathrm{mn}}^0 \max(y,d_1)} + \Gamma_0^u e^{-\gamma_{\mathrm{mn}}^0 \max(y,d_1)}\right)}{\gamma_{\mathrm{mn}}^0\left(\beta_0^l\Gamma_0^u - \beta_0^u\Gamma_0^l\right)} \times \left(\beta_0^l e^{\gamma_{\mathrm{mn}}^0 \min(y,d_1)} + \Gamma_0^l e^{-\gamma_{\mathrm{mn}}^0 \min(y,d_1)}\right) \quad (18)$$

Computing the mutual inductance to another filament at $y = d_2$, $x = d, l_1/2 + \delta \le z \le l_1/2 + \delta + l_2$ (see Fig. 4)

$$M = \int_{\frac{l_1}{2}+\delta}^{\frac{l_1}{2}+\delta+l_2} A(d,d_2,z')dz' \quad (19)$$

yields

$$M = \left(\frac{2\mu}{\pi^2}\right)\int_0^\infty \mathrm{dm}\int_0^\infty \mathrm{dncos}(md)\frac{\sin\left(\frac{nl_1}{2}\right)}{\gamma_{\mathrm{mn}}^0 n} \times \frac{\sin\left[n\left(\frac{l_1}{2}+\delta+l_2\right)\right] - \sin\left[n\left(\frac{l_1}{2}+\delta\right)\right]}{n} \times \frac{\left(\beta_0^u e^{\gamma_{\mathrm{mn}}^0 \max(d_1,d_2)} + \Gamma_0^u e^{-\gamma_{\mathrm{mn}}^0 \max(d_1,d_2)}\right)}{\left(\beta_0^l\Gamma_0^u - \beta_0^u\Gamma_0^l\right)} \times \left(\beta_0^l e^{\gamma_{\mathrm{mn}}^0 \min(d_1,d_2)} + \Gamma_0^l e^{-\gamma_{mn}^0 \min(d_1,d_2)}\right) \quad (20)$$

Equation (20) explicitly applies to filamentary currents. To extend this treatment to wires of finite cross section, we exploit the geometrical mean distance (GMD) approximation [32]. A volume filament decomposition as discussed in Section II is used to handle nonuniform current distributions. In the free-space case, $\beta_0^l = \Gamma_0^u = 1$ and $\beta_0^u = \Gamma_0^l = 0$ and this integral can be done exactly, to yield the well-known inductance formula of Grover [32].

We can approximate the integral of (20) by enclosing the x and z dimensions in a box of dimensions L_1 and L_2, respectively. Furthermore, we grid the substrate in x and z, so the L_1, L_2, d, l_1, l_2, and δ are integral multiplies of the grid spacing a. After some trigonometric manipulations, the summation becomes

$$\begin{aligned} M = \frac{\mu}{2L_1L_2a^2} \sum_{m=1}^{L_1-1} \sum_{n=1}^{L_2-1} f_{\rm mn} \\ \times \left\{ \cos\left(\frac{m\pi d}{L_1}\right) \cos\left(\frac{n\pi(l_1+\delta)}{L_2}\right) \right. \\ + \cos\left(\frac{m\pi d}{L_1}\right) \cos\left(\frac{n\pi(l_2+\delta)}{L_2}\right) \\ - \cos\left(\frac{m\pi d}{L_1}\right) \cos\left(\frac{n\pi(l_1+l_2+\delta)}{L_2}\right) \\ \left. -\cos\left(\frac{m\pi d}{L_1}\right) \cos\left(\frac{n\pi\delta}{L_2}\right) \right\} \end{aligned} \tag{21}$$

where f_{mn} is given by:

$$\begin{aligned} f_{\rm mn} = \frac{\left(\beta_0^u e^{\gamma_{\rm mn}^0 \max(d_1,d_2)} + \Gamma_0^u e^{-\gamma_{\rm mn}^0 \max(d_1,d_2)}\right)}{\left(\beta_0^l \Gamma_0^u - \beta_0^u \Gamma_0^l\right)} \\ \times \left(\beta_0^l e^{\gamma_{\rm mn}^0 \min(d_1,d_2)} + \Gamma_0^l e^{-\gamma_{\rm mn}^0 \min(d_1,d_2)}\right). \end{aligned} \tag{22}$$

Defining $\tilde{M}_{\rm uv}$ as

$$\tilde{M}_{\rm uv} = \frac{\mu}{2L_1L_2a^2} \sum_{m=1}^{L_1-1} \sum_{n=1}^{L_2-1} \cos\left(\frac{m\pi u}{L_1}\right) \cos\left(\frac{n\pi v}{L_2}\right) f_{\rm mn} \tag{23}$$

(21) becomes

$$M = \tilde{M}_{d,l_1+\delta} + \tilde{M}_{d,l_2+\delta} - \tilde{M}_{d,l_1+l_2+\delta} - \tilde{M}_{d,\delta}. \tag{24}$$

$\tilde{M}_{\rm uv}$ can be efficiently calculated by means of a FFT by creating an extended sequence $y_{\rm mn}$ from $f_{\rm mn}$ as follows:

$$y_{\rm mn} = \begin{cases} f_{\rm mn} & m=0,\dots L_1-1; n=0,\dots L_2-1 \\ f_{2M-m,n} & m=L_1,\dots 2L_1-1; n=0,\dots L_2-1 \\ f_{m,2N-n} & m=0,\dots L_1-1; n=L_2,\dots 2L_2-1 \\ f_{2M-m,2N-n} & m=L_1,\dots 2L_1-1; n=L_2,\dots 2L_2-1 \end{cases}. \tag{25}$$

With this, $\tilde{M}_{\rm uv}$ is given by

$$\begin{aligned} \tilde{M}_{\rm uv} = \frac{\mu}{8L_1L_2a^2} \sum_{m=1}^{2L_1-1} \sum_{n=1}^{2L_2-1} f_{mn} \\ \times \exp\left(\frac{jm2\pi u}{2L_1}\right) \exp\left(\frac{jn2\pi v}{2L_2}\right). \end{aligned} \tag{26}$$

For N metal layers, $N(N+1)/2$ FFTs are calculated as part of technology characterization (i.e., they only need to be calculated once and do not degrade the computational efficiency of the extraction engine) to provide look-up tables for $\tilde{M}_{\rm uv}$ to allow calculation of M according to (24). It is well known that a fine enough gridding must also be chosen (smaller a and more points in FFT) to cover adequate spectral content to achieve reasonable accuracy. This is particularly true for small values of d for which, $M_{\rm uv}$ has a (logarithmic) singularity. For the results presented here, we have chosen a as 0.25 μm and a maximum interaction region size of 256 $\times$256 μm, requiring a 2048 $\times$ 2048 FFT. We postprocess the FFT with a cubic-order interpolation formula [33], which allows us to achieve several percent accuracy down to the smallest filament spacing required for self-inductance calculation in the GMD approximation. This allows us to avoid the complexity of alternate solution techniques in the near field, the approach that is generally taken in more accurate field solvers [30].

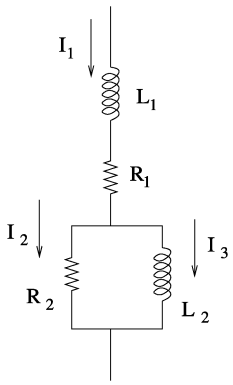

Fig. 5. Ladder network for a signal segment.

IV. Ladder Network

To model the frequency dependence of (5), including the additional frequency dependence due to substrate effects as represented by the "dressed" inductances of (21), a ladder network as shown in Fig. 5 is introduced for each signal line branch. The elements, in this case, are frequency-independent, making the extracted netlist compatible with SPICE simulation.

Representing all n signal branches within an interaction region, the current components become vectors, where each element of the vector corresponds to a different signal line. Consequently, $\boldsymbol{L_1}$, $\boldsymbol{L_2}$, $\boldsymbol{R_1}$, $\boldsymbol{R_2}$ are $n \times n$ matrices. $\boldsymbol{R_2}$ has off-diagonal elements, resulting in transresistances modeled in SPICE as current-dependent voltage sources. Physically, this models the common-return crosstalk since the signal lines share the same "lossy" implicit current returns. We note that RL ladder networks have been employed in the past to model proximity [34] and skin effects [35]. Note that the network of Fig. 5 generalizes prior approaches by offering a *coupled* equivalent circuit representation including both mutual inductances and transresistances.

The $n \times n$ impedance matrix of this reduced network is given by $\boldsymbol{Z}_{\rm fit}$ where $\boldsymbol{V} = \boldsymbol{Z}_{\rm fit}\boldsymbol{I}$ and $\boldsymbol{Z}_{\rm fit} = \boldsymbol{R}_1 + s\boldsymbol{L}_1 + s\boldsymbol{L}_2(\boldsymbol{R}_2 + s\boldsymbol{L}_2)^{-1}\boldsymbol{R}_2$. In the high-frequency limit, this ladder network yields inductance and resistance matrices given by

$$\boldsymbol{L}_{hf} = \boldsymbol{L}_1 \tag{27}$$

$$\boldsymbol{R}_{hf} = \boldsymbol{R}_1 + \boldsymbol{R}_2. \tag{28}$$

Similarly, in the low-frequency limit

$$\boldsymbol{L}_{\rm dc} = \boldsymbol{L}_1 + \boldsymbol{L}_2 \tag{29}$$

$$\boldsymbol{R}_{\rm dc} = \boldsymbol{R}_1 \tag{30}$$

We seek a wide-band fit of (5) to the ladder network over the frequency range from zero to a user-specified $f_{\rm max}$. Because the current distributions are uniform across the wire cross-sections at dc and there are no substrate eddy current losses at dc,

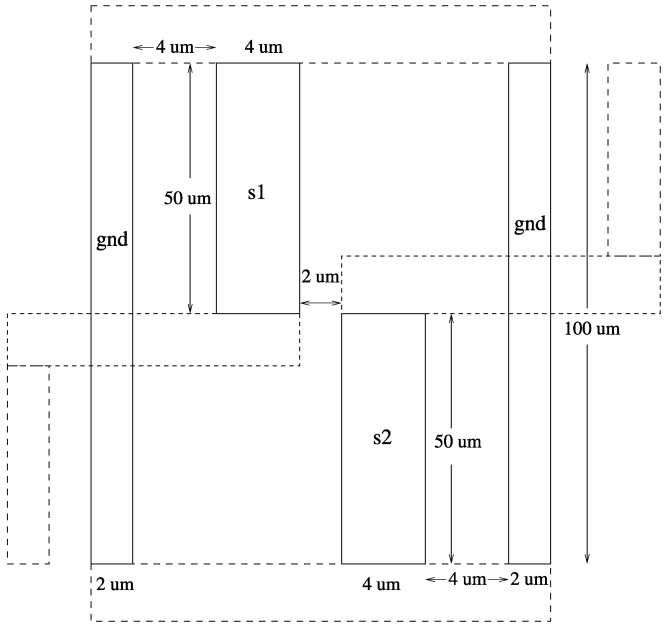

Fig. 6. Example to test the ladder fitting.

(5) yields the following expressions for the dc resistance and inductance

$$\tilde{\boldsymbol{L}}_{\rm dc} = \lim_{\omega\to 0}\frac{\Im\{\boldsymbol{Z}_{\rm eq}(\omega)\}}{\omega} = \boldsymbol{S}(\boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{B}^T)^{-1}\boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{L}\boldsymbol{R}^{-1}\boldsymbol{B}^T \times (\boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{B}^T)^{-1}\boldsymbol{S}^T \quad (31)$$

$$\tilde{\boldsymbol{R}}_{\rm dc} = \Re\{\boldsymbol{Z}_{\rm eq}(0)\} = \boldsymbol{S}(\boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{B}^T)^{-1}\boldsymbol{S}^T. \quad (32)$$

$\boldsymbol{R}$ is the diagonal resistance matrix and $\boldsymbol{L}$ is the dense matrix of partial inductances for the segments of the interaction region. Expressions for the high-frequency resistance and inductance come from evaluating (5) at $\omega_{\max} = 2\pi f_{\max}$

$$\tilde{\boldsymbol{L}}_{\rm hf} = \frac{\Im\{\boldsymbol{Z}_{\rm eq}(\omega_{\max})\}}{\omega_{\max}} \quad (33)$$

$$\tilde{\boldsymbol{R}}_{\rm hf} = \Re\{\boldsymbol{Z}_{\rm eq}(\omega_{\max})\}. \quad (34)$$

The fitting can now be performed trivially. $\tilde{\boldsymbol{L}}_{\rm hf}$ determines $\boldsymbol{L}_1, \tilde{\boldsymbol{R}}_{\rm dc}$ determines $\boldsymbol{R}_1, \tilde{\boldsymbol{L}}_{\rm dc}$ determines $\boldsymbol{L}_2$, and $\boldsymbol{R}_{\rm hf}$ determines $\boldsymbol{R}_2$. Specifically, $\boldsymbol{R}_2 = \tilde{\boldsymbol{R}}_{\rm hf} - \boldsymbol{R}_1$ and $\boldsymbol{L}_2 = \tilde{\boldsymbol{L}}_{\rm dc} - \boldsymbol{L}_1$

The ladder model parameters are also used to formalize a simple set of filter criterion used to generate ladder networks only where necessary and to remove inductors from nets, on which they are electrically insignificant. If $|\boldsymbol{R}_1|_i/|\boldsymbol{L}_1|_i \gg 2\pi f_{\max}$, then inductors are removed for segment i of the interaction region. The norm $|A|_i$ is defined as the 2-norm [36] of the vector corresponding to the ith row (or column) of the matrix $\boldsymbol{A}$. If $|\boldsymbol{R}_2|_i/|\boldsymbol{R}_1|_i \ll 1$ or $|\boldsymbol{L}_2|_i/|\boldsymbol{R}_2|_i \ll 2\pi f_{\max}$, then the ladder is filtered for segment i and a simple R–L model is used. It should be noted that more aggressive filtering would be achievable if it were possible to analyze the entire net, including its capacitance and the driving and receiving circuits. This is not easily done within the context of an extraction tool.

We use a simple example as shown in Fig. 6, consisting of an interaction region with two signal lines and two ground lines, to demonstrate the accuracy of the fitting procedure. $f_{\max}$ is chosen as 20 GHz in this example. The dashed rectangles are not included in the vertical interaction region because they are horizontal or they are blocked out by vertical halos. We use technology parameters that correspond to the top-level metal in a TSMC 0.25-μm CMOS process; skin and proximity effect modeling is included but this example does not include substrate loss. Fig. 7 compares the resistance and inductance values that come from direct evaluation of (5) with those that result from the ladder network fit. $R11$ and $R12$ represent the resistance of signal line $s1$ and the transresistance between $s1$ and $s2$, respectively. $L11$ and $L12$ represent the inductance of $s1$ and the mutual inductance between $s1$ and $s2$, respectively. The fit is reasonable and representative of a larger set of testcases examined; better fits can only be achieved with more complex network representations.

V. Assura RCX-PL Implementation

Assura RCX-PL is derived from the GOALIE2 shapes-processing engine described elsewhere [37], [38]. In this section, we give a brief overview of the software architecture as it has evolved in the base Assura RCX system and discuss the extensions to support inductance extraction in RCX-PL.

The Assura RCX system consists of a number of separate program modules which operate on geometric information (rectangles) stored in *edge files*. Other binary files are used to store elements such as transistors (*device files*), capacitors (*capacitance files*), and resistors (*resistance files*). A simplified data flow diagram of the system is shown in Fig. 8. The blocks to the left show the basic system, which makes extensive use of disk caching and intermediate files to handle large data volumes. An input GDSII file is processed through the layout versus schematic (LVS) engine to generate a set of edge files and device files, which contain the extracted layout of the input GDSII. The resistance extraction program (`rex`) modifies the edge files with cuts for capacitance distribution and generates the resistance files. Current directions identified by `rex` in these edge files are used to identify horizontal and vertical segments in inductance extraction. `rex` also associates *terminals* with certain rectangles or collections of rectangles in order to define the connectivity of the resulting netlist. These terminals could exist anywhere within a rectangle, but there is at most one per rectangle. The modified edge files are then passed to the capacitance extraction module, which generates a capacitance file. The device files, resistance files and the capacitance files are then input to the SPICE netlist generator which generates a SPICE netlist for the design.

lextract module. The blocks in Fig. 8 enclosed in dotted lines are the modules we have added in this work to support inductance extraction. The edge files after resistance extraction are passed as input to the interaction region generation program (`lextract`) which is responsible for applying the halo rules described in Section II, dividing the chip into disjoint horizontal and vertical interaction regions. A "scanband" approach, commonly employed in capacitance extraction and layout verification, is used in `lextract`. The width of the schanband is chosen to accommodate the largest interaction region size,

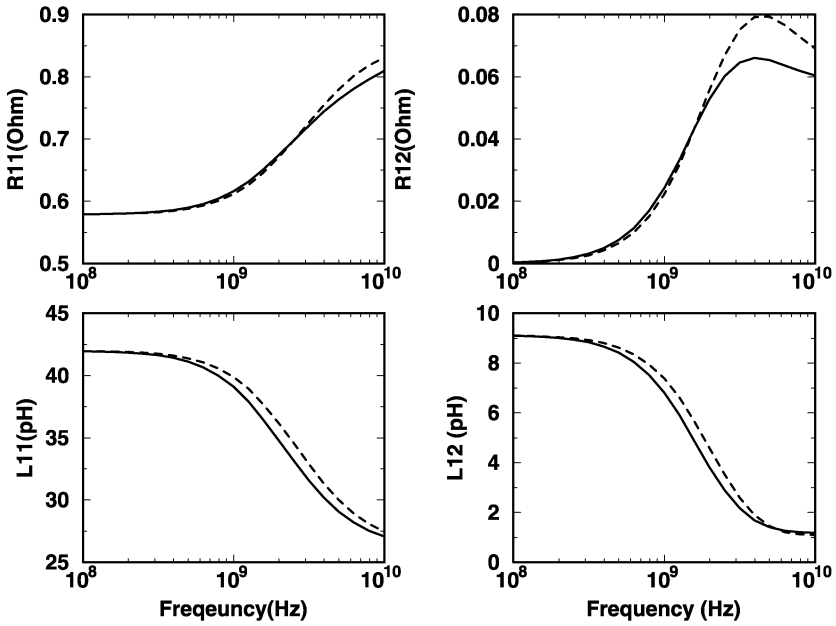

Fig. 7. Comparison of the impedance of the fitted ladder network with the actual impedance of (5). The solid curve is the exact result; the dashed curve is the fit.

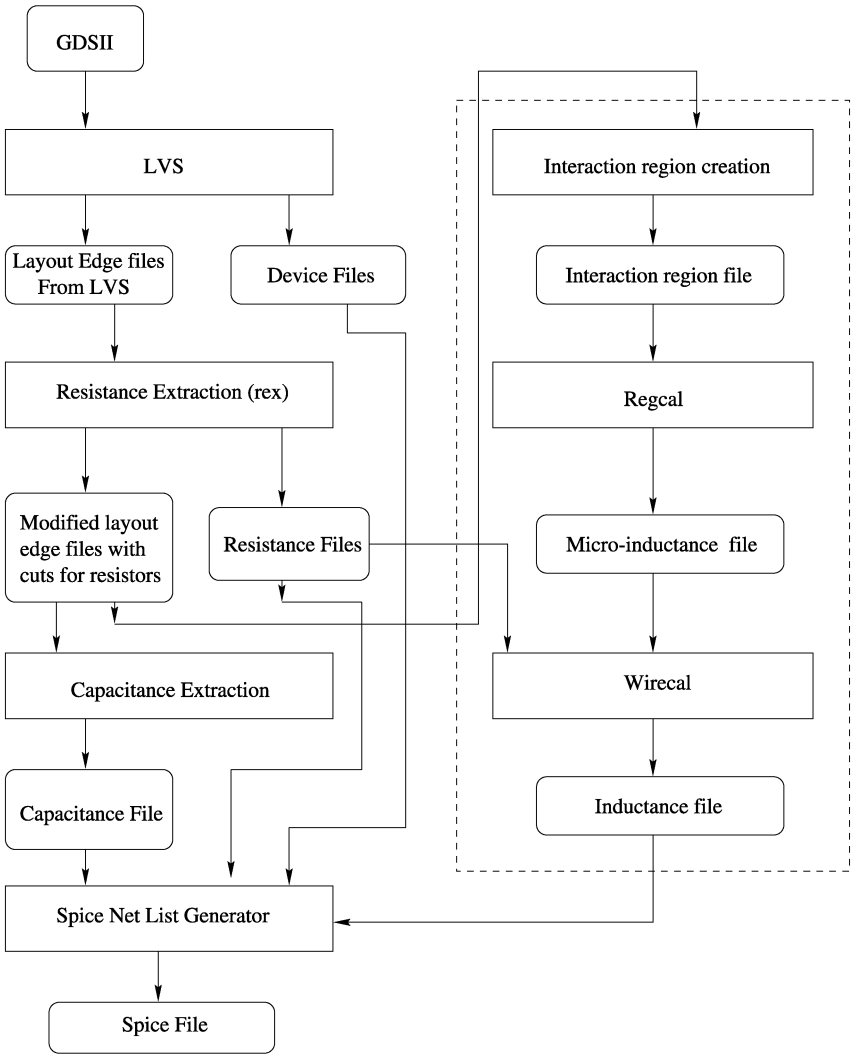

Fig. 8. Assura RCX-PL architecture. The modules in the dotted box are added to support inductance extraction.

as generally determined by power-grid pitch. `lextract` produces two *interaction region files*. The horizontal (vertical) interaction region file contains a geometric specification of the interaction region and all the horizontal (vertical) signal and power-ground lines contained in the interaction region. A left-to-right scan produces the vertical interaction region file. To produce the horizontal file, `lextract` is run as a second pass with the input rectangles resorted and transposed to effectively perform a bottom-to-top scan. The interaction region geometry is specified on each metal layer with a set of rectangular *subregions*. To ensure that this specification is unique, the subregions are defined in a *vertically dominant* way. By this, we mean that every horizontal side of a subregion will constitute part of the interaction region boundary.

Following `lextract`, the inductance calculation is implemented with two engines `regcal` and `wirecal` as shown in Fig. 8. There are also various other secondary modules in the flow to sort and merge intermediate files.

regcal module. Acting one interaction region at a time, `regcal` divides nets into smaller rectangles based on a set of fracturing rules and calculates $\tilde{\boldsymbol{L}}_{\rm hf}$, $\tilde{\boldsymbol{R}}_{\rm hf}$, $\tilde{\boldsymbol{L}}_{\rm dc}$, and $\tilde{\boldsymbol{R}}_{\rm dc}$ as defined in (31)–(34) for each interaction region. This information is represented in the *microinductance file* noted in Fig. 8.

Signal and power-ground conductors are fractured to define the clusters described in Section II. There are three successive fracturing steps for the signal lines. (We reference all of our discussion to the vertical interaction region file. The horizontal interaction region file is independently processed in a similar manner.)

- If a rectangle overlaps the horizontal boundary of a subregion, then it is divided into two rectangles by the boundary. Rectangles outside the current interaction region are discarded, since they will be analyzed as part of the adjacent interaction region.
- If a signal rectangle overlaps one of the ends of a power rectangle in the vertical (y) direction (in other words, if the y-value defining the end of the power rectangle lies between the y values of the two ends of the signal rectangle), then the signal rectangle is divided into two rectangles by the y-value of the end of the power rectangle. This effectively breaks signal lines when their associated power-ground "environment" changes.
- Every terminal divides the associated signal rectangle into two rectangles.

Fig. 9 is an example illustrating the fracturing process. One signal rectangle in the original interaction region is divided into six segments, two of them discarded. Breaks a and d are created by the boundary of subregion 2. b is created by the end of a power-ground rectangle. c is formed by the terminal of the net. All of the signal line fractures are then projected onto the power-ground lines. Based on these fractures, each signal rectangle has associated with it a set of parallel power-ground rectangles in the interaction region overlapping it in the y-direction. Note that the same power-ground segment may be part of multiple clusters even though it appears once in $\boldsymbol{Z}$.

In addition to $\tilde{\boldsymbol{L}}_{\rm hf}$, $\tilde{\boldsymbol{R}}_{\rm hf}$, $\tilde{\boldsymbol{L}}_{\rm dc}$, and $\tilde{\boldsymbol{R}}_{\rm dc}$, the microinductance file contains the midpoint of each associated signal rectangle and its net label, information required by `wirecal`, which "reduces" the inductance matrix to match the terminals in the resistance file created by `rex` and then performs the appropriate ladder network synthesis. This reduction preserves the netlist "bookkeeping" infrastructure of the *RC* extractor. Resistance extraction is controlled with a requirement on the maximum number of squares between breaks to ensure an adequate lumped-element approximation for inductance modeling.

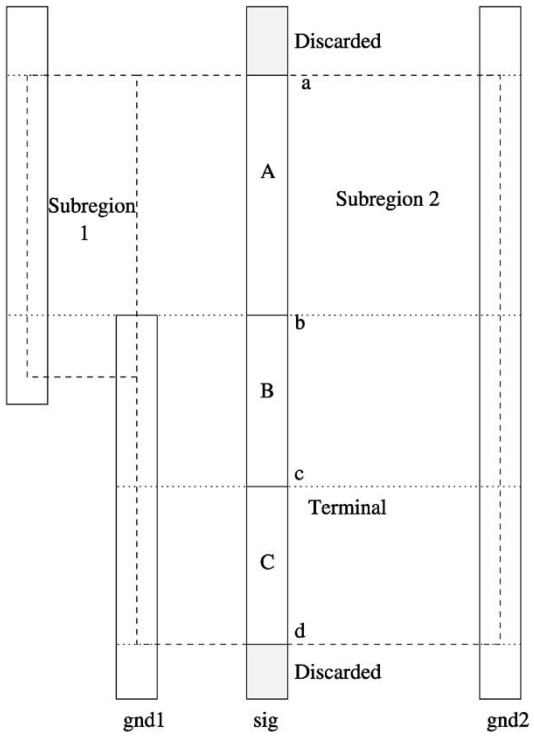

Fig. 9. Example illustrating the fracturing rules.

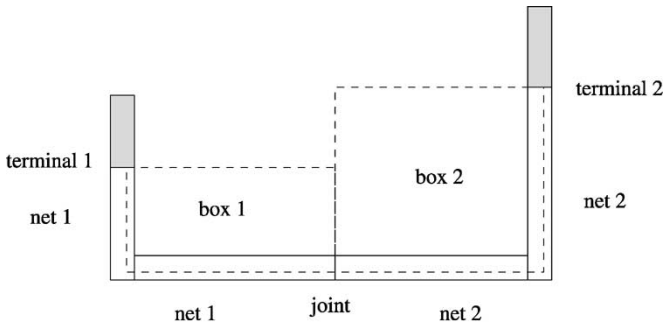

Fig. 10. Example showing how "micro-inductances" are combined.

wirecal module. A simple congruence transformation is used to convert $\boldsymbol{Z}_{\rm eq}$ into $\boldsymbol{Z}'_{\rm eq}$, the impedance matrix that corresponds to the branches defined by `rex`

$$\boldsymbol{Z}'_{\rm eq} = \boldsymbol{V}\boldsymbol{S}(\boldsymbol{B}\boldsymbol{Z}^{-1}\boldsymbol{B}^T)^{-1}\boldsymbol{S}^T\boldsymbol{V}^T. \tag{35}$$

Each column of $\boldsymbol{V}$ corresponds to a segment in the microinductance file, while each row of $\boldsymbol{V}$ corresponds to a segment as defined by `rex`. The ith row of $\boldsymbol{V}$ is all zero except for ones in columns corresponding to segments which combine to form the associated `rex` segment. To determine $\boldsymbol{V}$, simple geometry operations are used to determine which microinductance segments must be combined to constitute the `rex`-determined

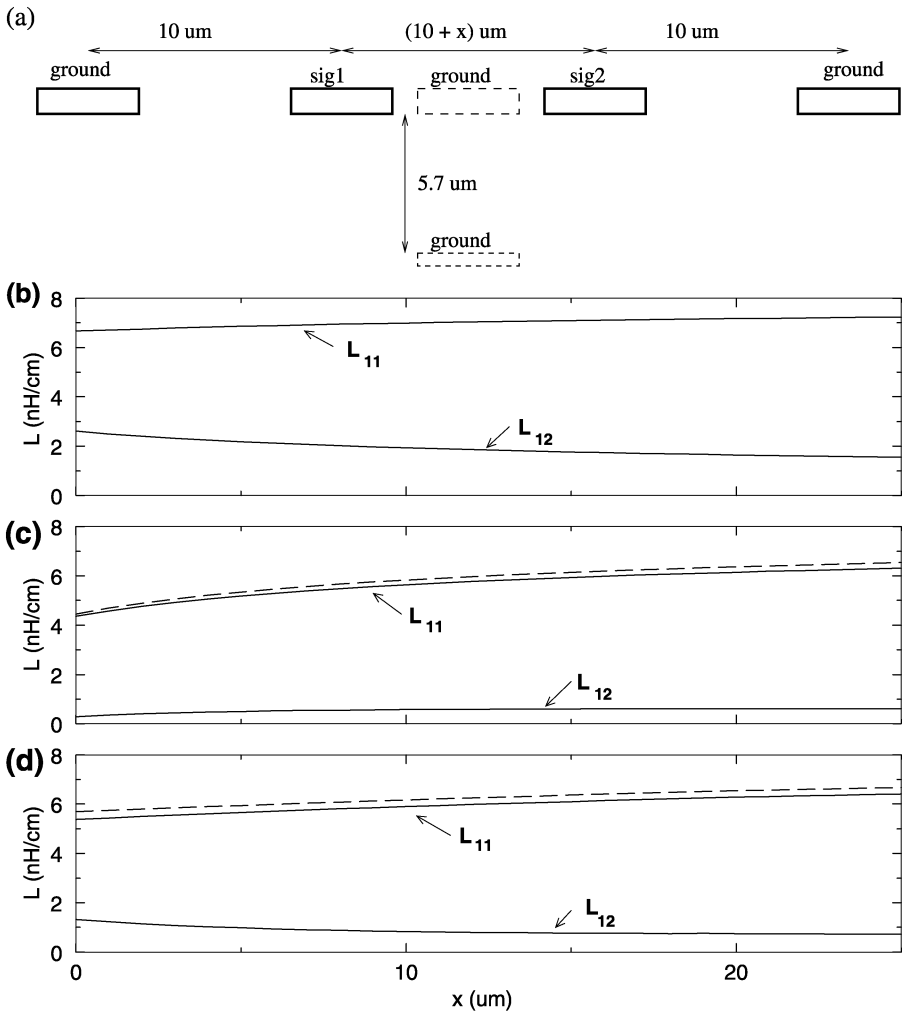

Fig. 11. Simple testcase to consider the efficacy of ground returns in shielding magnetic coupling. (a) Base testcase consists of two signal lines and two ground lines (drawn with solid boundaries). (b) Self and mutual inductance for the base testcase. (c) Self and mutual inductances in the presence of a third ground lines on the same metal layer between the two signal lines (solid lines are exact results, while the dashed lines are the Assura RCX-PL results). (d) Self and mutual inductances in the presence of a third ground line that is on first-level metal between the two signal lines.

branches. In general, a single terminal as created by `rex` may be associated with several connected rectangles. The result may be an "L" shape (as in Fig. 10) but will not be a "T", "+", or "U" shape. Each net has one and only one terminal, and every `rex`-defined branch is between the terminals of two nets. Two bounding boxes are used to determine which segments must be combined for a branch as shown in the example of Fig. 10. Each bounding box is defined by the location of the terminal of the associated net and the *joint* location, the point at which the shapes associated with each of the two nets adjoin. Every rectangle labeled with the same net number as the terminal whose midpoint is contained within the associated bounding box is combined. In Fig. 10, four such rectangles are associated with the connection between `net1` and `net2`. The gray rectangles (outside the bounding boxes) are not included.

With the $\boldsymbol{V}$ matrix so determined, the microinductance file matrices are then transformed accordingly

$$\tilde{\boldsymbol{L}}'_{\rm hf} = \boldsymbol{V}^T \tilde{\boldsymbol{L}}_{\rm hf} \boldsymbol{V} \tag{36}$$

$$\tilde{\boldsymbol{R}}'_{\rm hf} = \boldsymbol{V}^T \tilde{\boldsymbol{R}}_{\rm hf} \boldsymbol{V} \tag{37}$$

$$\tilde{\boldsymbol{L}}'_{\rm dc} = \boldsymbol{V}^T \tilde{\boldsymbol{L}}_{\rm dc} \boldsymbol{V} \tag{38}$$

$$\tilde{\boldsymbol{R}}'_{\rm dc} = \boldsymbol{V}^T \tilde{\boldsymbol{R}}_{\rm dc} \boldsymbol{V}. \tag{39}$$

These transformations reduce to simple addition of appropriate matrix elements of the original matrices to form the elements of the new matrices. These values are then used to synthesize ladder networks as described in Section IV. Filtering techniques, also described in Section IV, are used to simplify the resulting network. The result is resistance and inductance files, which are combined with capacitance and device files for netlist generation as shown in Fig. 8.

VI. Results

The block sparsification of the inductance matrix in return-limited extraction is based on the efficacy of ground lines in "screening" magnetic coupling. To address the accuracy of this approximation, consider the testcase of Fig. 11(a). Two ground lines and two signal lines—2-μm-thick, 4-μm-wide, and 10 μm apart, representing the top metal layer of a five-layer metal process—have the self-inductance $(L_{11} = L_{22})$ and mutual inductance $(L_{12} = L_{21})$ shown in Fig. 11(b) as a function of the distance between the two signal lines. Fasthenry

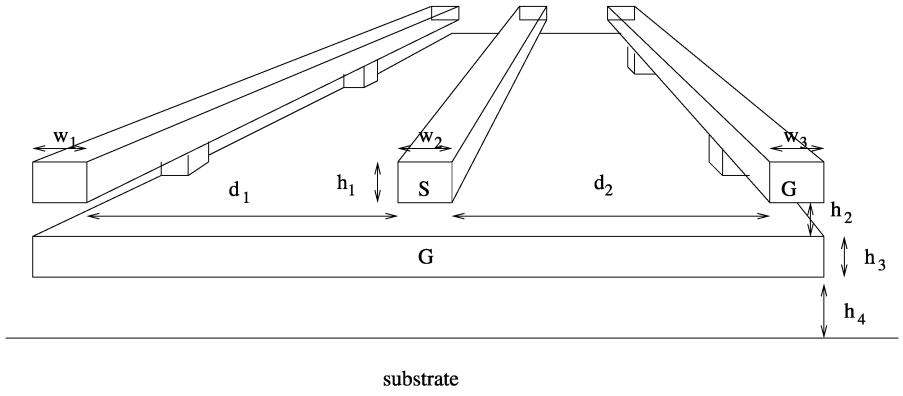

Fig. 12. Simple GSG coplanar waveguide structure with a metal ground plane.

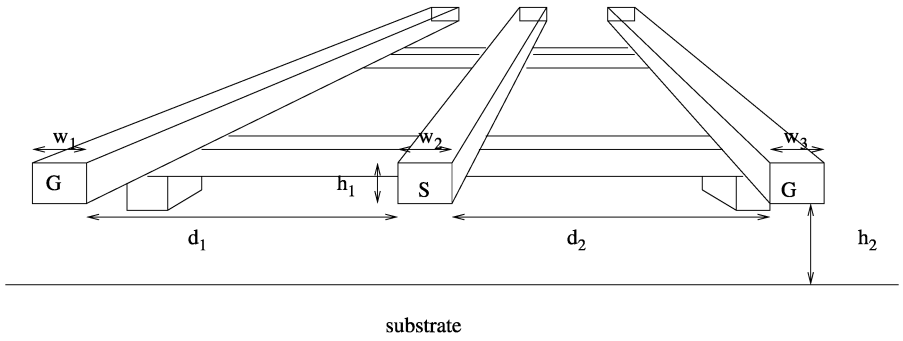

Fig. 13 Simple GSG coplanar waveguide structure running over the silicon substrate.

[14] and Assura RCX-PL offer indistinguishable results in this case. In Fig. 11(c), we add an additional ground lines, also 4 μm wide, between the two signal lines on the same metal layer. In this case, the solid lines indicate the Fasthenry results[5] and the dashed lines indicate the Assura RCX-PL results. Assura RCX-PL in this case is using the new ground line to break the magnetic coupling between the two signal lines ($L_{12} = L_{21} = 0$). The actual mutual inductance of almost 0.6 nH/cm is therefore ignored in this approximation, but this is only 10% of the self-inductance value in this case. Assura RCX-PL slightly overestimates L_{11} because of the consideration of only two of the three ground lines for each signal lines in determining the self-inductance. The interaction region approximation in this testcase gets slightly worse when the ground line between the two signal lines is moved down to the first metal layer (4-μm-wide and 1-μm-thick) in Fig. 11(d), but the ignored mutual inductance still remains approximately 10% of the self-inductance.

Correct validation of the extraction approach presented here requires comparison with full-wave analysis because of the complex interactions of eddy currents, conduction currents, and displacements currents. We have developed a set of two-port ground-signal-ground (GSG) coplanar waveguide structures of the form shown in Figs. 12 and 13. The structure of Fig. 12 contains a metal ground plane not present in the structure of Fig. 13. These structures are simple enough to enable full-wave solution with Ansoft's HFSS engine, a finite-element solver,

[5]These results are indistinguishable from Assura RCX-PL results in which a single interaction region is user-specified. That is, the "center" ground line is not allowed to divide the problem into two interaction regions.

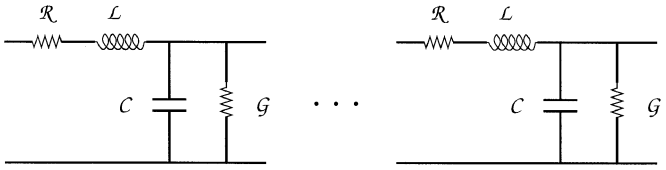

Fig. 14. Equivalent model for transmission line.

yielding S parameters. Ports are defined at the near-end and far-end. We similarly calculate the S-parameters from Spectre simulation of the Assura RCX-PL SPICE netlist.

Since these structures are long enough to be considered homogenous, we extract four distributed transmission line parameters as shown in Fig. 14 from the S-parameters. $\mathcal{R}$, $\mathcal{L}$, $\mathcal{C}$, and $\mathcal{G}$ correspond to the series resistance, series inductance, shunt capacitance, and shunt conductance per unit length, respectively.

These circuit parameters can be extracted from the propagation constant γ and characteristic impedance Z of the line, which can be extracted from the S parameters according to the relationships [39]

$$e^{-\gamma l} = \left\{ \frac{1 - S_{11}^2 + S_{21}^2}{2S_{21}} \pm K \right\}^{-1} \tag{40}$$

and

$$Z^2 = Z_0^2 \frac{(1 + S_{11})^2 - S_{21}^2}{(1 - S_{11})^2 - S_{21}^2} \tag{41}$$

where

$$K = \sqrt{\frac{(S_{11}^2 - S_{21}^2 + 1)^2 - (2S_{11})^2}{(2S_{21})^2}}. \tag{42}$$

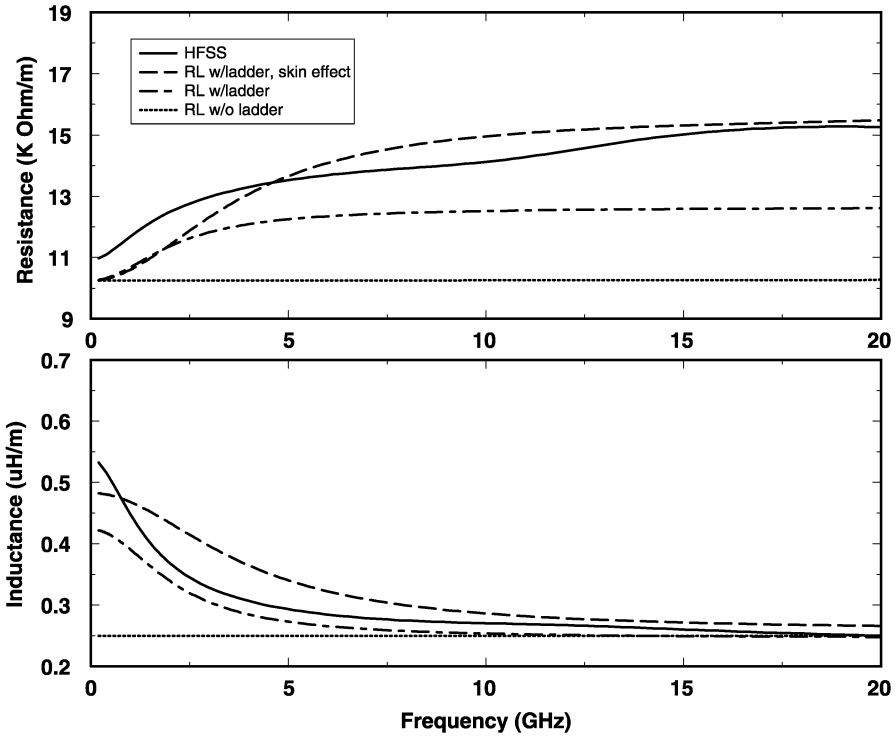

Fig. 15. Extracted resistance and inductance per unit length for Example 1.

In (40) and (41), l is the length of the structure, and Z_0 is the reference impedance for the S parameters (usually 50 Ω). $\mathcal{R}$ and $\mathcal{L}$ are determined as follows:

$$\mathcal{R} = Re(\gamma Z) \tag{43}$$

$$\mathcal{L} = \frac{Im(\gamma Z)}{\omega}. \tag{44}$$

We use these $\mathcal{R}$ and $\mathcal{L}$ values to compare the accuracy of Assura RCX-PL with full-wave analysis in HFSS.

A. Example 1. Coplanar Structure With Metal Ground Plane

The first example is a grounded coplanar waveguide structure as shown in Fig. 12, where $h_1 = 0.99$ μm, $h_2 = 1$ μm, $h_3 = 0.57$ μm, $h_4 = 6.855$ μm, $d_1 = 4$ μm, $d_2 = 20$ μm, $w_1 = 2$ μm, $w_2 = 4$ μm, and $w_3 = 10$ μm. The metallization is aluminum, and the dimensions are typical of the top two metal layers in a TSMC 0.25-μm process. The structure is 5-mm long; the two ground lines are connected with the ground plane every 500 μm with vias. Due to the shielding action of the ground plane, the substrate has negligible effect on the signal line characteristics. $f_{\max}$ is chosen as 20 GHz for the Assura RCX-PL extraction.

The $\mathcal{R}$ and $\mathcal{L}$ values for Example 1 are shown in Fig. 15. HFSS indicates the HFSS result. "RL w/o ladder" is the return-limited Assura RCX-PL result without consideration of power-ground losses and skin and proximity effects. "RL w/ladder" is the return-limited Assura RCX-PL result with consideration of power-ground losses. "RL w/ladder, skin effect" is the return-limited Assura RCX-PL result with consideration of power-ground losses and skin and proximity effects. Both resistance and inductance show better agreement with HFSS as loss mechanisms are added into the Assura extraction. Including the ladder network and the skin effect are essential to good agreement between HFSS and Assura in the high-frequency resistance. HFSS and Assura disagree in the dc resistance because Assura includes a pseudoreturn to drive the dc resistance to that of the signal line alone (modeling the more distant returns that inevitably exist in any real power distribution). HFSS sees only the coplanar ground returns given, even at dc. One could argue that Assura is "more accurate" than HFSS because it has a more "natural" view of the power-ground distribution. We comment that the HFSS results require in excess of 18 hours of compute time. The extraction in Assura and simulation in Spectre requires a couple of minutes.

B. Example 2. Coplanar Waveguide Over Nonepitaxial Silicon Substrate

The second structure we consider is a coplanar waveguide structure running over a silicon substrate as shown in Fig. 13, where $h_1 = 0.99$ μm, $h_2 = 7.425$ μm, $d_1 = 4$ μm, $d_2 = 20$ μm, $w_1 = 2$ μm, $w_2 = 4$ μm, and $w_3 = 10$ μm. In this case, the ground plane of Example 1 is replaced by straps connecting the two ground lines together every 500 μm. The substrate is a "nonepitaxial" substrate with a 0.1-$\Omega-\mathrm{cm}$ field implant 0.7-μm -thick on a bulk 20-$\Omega-\mathrm{cm}$ wafer.[6] The substrate is plugged (or tapped) every 500 μm from both ground lines. $f_{\max}$ is chosen as 20 GHz for the Assura RCX-PL extraction.

[6]This is characteristic of the nonepitaxial substrates used in TSMC's 0.25-μm process.

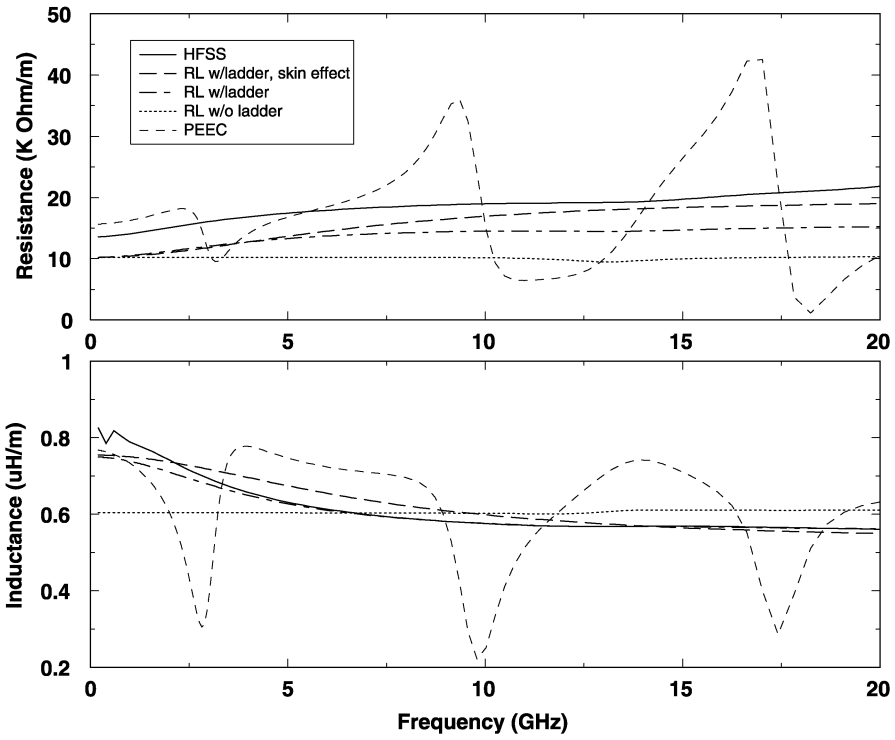

Fig. 16. Extracted resistance and inductance per unit length for Example 2.

The $\mathcal{R}$ and $\mathcal{L}$ values for Example 2 are shown in Fig. 16. HFSS indicates the HFSS result. "RL w/o ladder" is the return-limited Assura RCX-PL result without consideration of power-ground losses.[7] "RL w/ ladder" is the return-limited Assura RCX-PL result with consideration of power-ground losses. "RL w/ ladder, skin effect" is the return-limited Assura RCX-PL result with consideration of power-ground losses and the proximity and skin effects. In this case, because the substrate is very resistive and distant, the substrate corrections are negligible. As in Example 1, the results show steadily better agreement with HFSS as loss mechanisms are added to the extraction.

There is an essential consistency between modeling the capacitances to the substrate and power-ground grid as capacitances to ground and the magnetostatic assumptions manifest in the shunt connections in Fig. 1(b) (for power-ground line modeling) and ignoring the last term in (9) (for magnetostatic substrate modeling). To illustrate the importance of this consistency to the accuracy of the extraction, we have added an additional analysis in Fig. 16 labeled "PEEC." In this case, we model the ground lines explicitly; that is, we have an explicit RLC model of the ground line included in the extraction. However, we continue to model the substrate as an "ideal" ground; that is, capacitances to the substrate are capacitances to ground. The resulting network is wildly inaccurate because the displacement currents flowing into the substrate are actually being collected by plugs and contribute to the return currents flowing in the ground lines. By not allowing these currents to flow into the power-ground lines, we have "starved" these lines of return current, resulting in a very inaccurate extraction result.

[7]Both the resistance and inductance should be frequency-independent in this case. The slight frequency dependence observed for this curve is an artifact of the fitting procedure to the simulated S-parameter data.

C. Example 3. Coplanar Waveguide Over Epitaxial Silicon Substrate

The third example is also a coplanar waveguide structure running over a silicon substrate as shown in Fig. 13, where $h_1 = 0.57\ \mu$m, $h_2 = 2.715\ \mu$m, $w_1 = 2\ \mu$m, $w_2 = 4\ \mu$m, $w_3 = 10\ \mu$m, $d_1 = 30\ \mu$m, and $d_2 = 50\ \mu$m. This is typical of the bottom two metal layers in a TSMC 0.25 μm process. In addition to placing the signal line closer to the substrate, the ground lines have also been pulled further away to accentuate substrate effects. The ground lines are neither equal width nor equidistant from the signal lines, given an additional proximity-effect frequency dependence to the resistance and inductance. The substrate is an epitaxial substrate modeled as three layers. The top layer is 1 μm-thick with 2 $\Omega-$ cm resistivity; the second layer is 10 μm-thick with 15-$\Omega-\mathrm{cm}$ resistivity; the third layer is a bulk with 1-m$\Omega-$ cm resistivity. The ground lines are strapped every 500 μm and the substrate is plugged from both ground lines every 500 μm. $f_{\max}$ is chosen as 20 GHz for the Assura RCX-PL extraction.

Once again, Fig. 17 compares $\mathcal{R}$ and $\mathcal{L}$ from 50 MHz to 20 GHz. HFSS gives the HFSS result. "RL w/o ladder" gives the Assura RCX-PL result without any consideration of substrate or power-ground losses. "RL w/ ladder, skin effect" includes the power-ground losses and the skin and proximity effects but no effects from the substrate. "RL w/ ladder, skin, substrate" includes substrate loss as well as power-ground loss and the proximity and skin effects. In this case, the substrate correction pro-

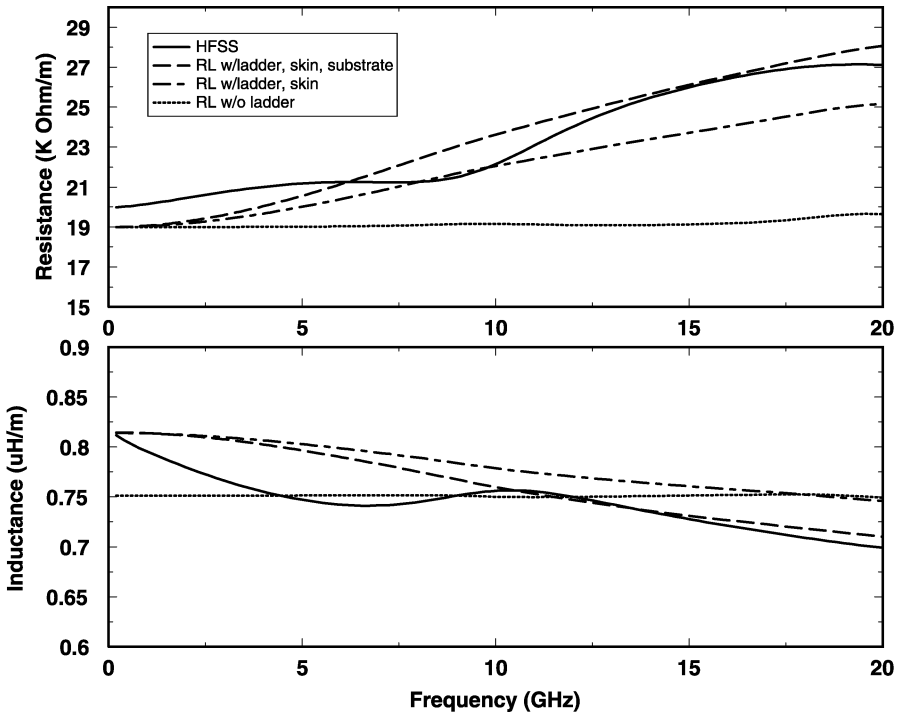

Fig. 17. Extracted resistance and inductance per unit length for Example 3.

vides an important component to the loss and an important correction to the high-frequency inductance.

D. Example 4. Large Testcase to Consider Runtime Performance and Number of Extracted Elements

We also present results on a large testcase to consider the extraction runtime performance and complexity of the extracted netlists. The design is a high-frequency processor core with an area of approximately 15 mm^2 in a 0.18-μm, five-level-metal process. $f_{\max}$ is chosen as 20 GHz for this run, which is also used for inductance and ladder network filtering as described in Section IV. Runtime statistics were gathered for a run which considers losses in the power-ground network but ignores volume filament modeling in the conductors or substrate losses. These higher accuracy options still result in prohibitively high runtimes for full-chip analysis and can currently only be employed for small problem sizes. Runtime statistics are for a Sunfire 880 (750 MHz). Peak memory usage is approximately 500 MB in `regcal`. As shown in Table I, full-chip inductance extraction does add considerably to the runtime of the full-chip extraction but is still reasonable for the design size. With the extensive "blocking" mechanisms in the halo rules, there was a concern that on "real" designs, the interaction regions could grow quite large, significantly slowing down the *regcal* module (which contains a matrix factorization whose size is based on the number of segments in the interaction region). However, this has not proved to be the case. In this example, the largest number of signal rectangles in an interaction region (based on *rex* fracturing) is 1146.

TABLE I
STATISTICS FOR "LARGE" TESTCASE

Number of nets	22354
Number of transistors	68062
Technology	0.18μm, 5-layer
Inductance extraction run time	286 minutes
RC extraction runtime	26 minutes
Average interaction region x extent	50.53 μm
Average interaction region y extent	56.47 μm
Average number of power rectangles per interaction region	11
Average number of signal rectangles per interaction region	20
Maximum number of power rectangles per interaction region	467
Maximum number of signal rectangles per interaction region	1146
Number of resistors	158340
Number of capacitors	1590059
Number of inductors	27646
Number of mutual inductors	444033
Number of transresistance	70802
Number of ladders	11230

VII. CONCLUSION

In this paper, we report the development of a full-chip, 3-D, shapes-based, RLC extraction tool. The technique of return-limited inductances is used to provide a sparse, frequency-independent inductance and resistance network with self-inductances that represent sensible "nominal" values in the absence of mutual coupling. Mutual inductances are extracted for accurate crosstalk analysis. The tool exploits high-capacity scan-band techniques and disk caching. Accuracy is validated by comparison with full-wave finite-element field solvers.

REFERENCES

[1] P. J. Restle *et al.*, "A clock distribution network for microprocessors," *IEEE J. Solid-State Circuits*, vol. 36, pp. 792–799, May 2001.

[2] S. Voinigescu *et al.*, "Circuits and technologies for highly integrated optical networking IC's at 10 Gb/s to 40 Gb/s," in *Proc. Int. Custom Integrated Circuits Conf.*, 2001, pp. 331–338.

[3] K. W. J. Chew, S.-F. S. Chu, and C. C. C. Leung, "Driving CMOS into the wireless communications arena with technology scaling," in *Proc. Int. Custom Integrated Circuits Conf.*, 2001, pp. 571–574.

[4] J. N. Burghartz, M. Soyuer, and K. A. Jenkins, "Integrated RF and microwave components in BiCMOS technology," *IEEE Trans. Electron Devices*, vol. 43, pp. 1559–1570, Sept. 1996.

[5] C. P. Yue and S. S. Wong, "On-chip spiral inductors with patterned ground shields for Si-based RF IC's," *IEEE J. Solid-State Circuits*, vol. 33, pp. 743–752, May 1998.

[6] A. M. Niknejad and R. G. Meyer, "Analysis, design, and optimization of spiral inductors and transformers for Si RF IC's," *IEEE J. Solid-State Circuits*, vol. 33, pp. 1470–1481, Oct. 1998.

[7] B. Kleveland, C. H. Diaz, D. Vook, L. Madden, T. H. Lee, and S. S. Wong, "Exploiting CMOS reverse interconnect scaling in multigigahertz amplifier and oscillator design," *IEEE J. Solid-State Circuits*, vol. 36, pp. 1480–1488, Oct. 2001.

[8] H. Wu and A. Hajimiri, "Silicon-based distributed voltage-controlled oscillators," *IEEE J. Solid-State Circuits*, vol. 36, pp. 493–502, Mar. 2001.

[9] B. F. Harrington, *Field Computation by Moment Methods*. New York: MacMillan, 1968.

[10] A. E. Ruehli, "Equivalent circuit models for three-dimensional multiconductor systems," *IEEE Trans. Microwave Theory Tech.*, vol. MTT-22, pp. 216–221, Mar. 1974.

[11] P. J. Restle, A. H. Ruehli, S. G. Walker, and C. Papadopoulos, "Full-wave PEEC time-domain methods for the modeling of on-chip interconnects," *IEEE Trans. Computer-Aided Design*, vol. 20, pp. 877–887, July 2001.

[12] H. Heeb and A. Ruehli, "Retarded models for PC board interconnects—or how the speed of light affects your SPICE simulation," in *Proc. IEEE/ACM Int. Conf. Computer Aided-Desiqn*, 1991, pp. 70–73.

[13] K. Nabors and J. White, "Fastcap: a multipole accelerated 3-D capacitance extraction program," *IEEE Trans. Computer-Aided Design*, vol. 10, pp. 1447–1459, Nov. 1991.

[14] M. Kamon, M. J. Tsuk, and J. K. White, "FAS-THENRY: A multipole-accelerated 3-D inductance extraction program," *IEEE Trans. Microwave Theory Tech.*, vol. 42, pp. 1750–1758, Sept. 1994.

[15] S. Kapur and D. Long, "IES^3: A fast integral equation solver for efficient 3-dimensional extractions," in *Proc. IEEE/ACM Int. Conf. Computer-Aided Design*, 1997, pp. 448–455.

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

[17] K. Gala, D. Blaauw, J. Wang, V. Zolotov, and M. Zhao, "Inductance 101: Analysis and design issues," in *Proc. ACM/IEEE Design Automation Conf.*, 2001, pp. 329–334.

[18] A. Devgan *et al.*, "How to efficiently capture on-chip inductance effects: Introducing a new circuit element K," in *Proc. IEEE/ACM Int. Conf. Computer-Aided Design*, 2000, pp. 150–155.

[19] M. W. Beattie and L. T. Pileggi, "Modeling magnetic coupling for on-chip interconnect," in *Proc. ACM/IEEE Design Automation Conf.*, 2001, pp. 335–340.

[20] K. L. Shepard and Z. Tian, "Return-limited inductances: A practical approach to on-chip inductance extraction," *IEEE Trans. Computer-Aided Design*, vol. 19, pp. 425–436, Apr. 2000.

[21] K. L. Shepard, D. Sitaram, and Y. Zheng, "Full-chip, three-dimensional, shapes-based RLC extraction," in *Proc. Int. Conf. Computer-Aided Design*, 2000, pp. 142–149.

[22] Assura Parasitic Extractor (RCX) Datasheet. Cadence Design Systems Tech. Rep. [Online] Available: http://www.cadence.com/datashee.ts/daLpdf/asmirascx.pdf

[23] K. L. Shepard, S. Carey, E. Cho, B. Curran, R. Hatch, D. Hoffman, S. McCabe, G. Northrop, and R. Seigler, "Design methodology for the G4 S/390 microprocessors," *IBM J. Res. Develop.*, vol. 21, no. 4/5, pp. 515–548, 1997.

[24] R. M. Averill *et al.*, "Chip integration methodology for the IBM S/390 G5 and G6 custom microprocessors," *IBM J. Res. Develop.*, vol. 45, no. 5/6, pp. 681–706, 1999.

[25] K. L. Shepard and Z. Tian, "Return-limited inductances: A practical approach to on-chip inductance extraction," in *Proc. Custom Integrated Circuits Conf.*, 1999, pp. 453–456.

[26] A. Ruehli, C. Paul, and J. Garrett, "Inductance calculations using partial inductances and macromodels," in *Proc. Int. Symp. Electromagn. Compat.*, 1995, pp. 23–28.

[27] Y. Massoud and J. White, "Simulation and modeling of the effect of substrate conductivity on coupling inductance," in *Proc. IEDM*, 1995, pp. 491–494.

[28] R. Gharpurey and R. G. Meyer, "Modeling and analysis of substrate coupling in integrated circuits," *IEEE J. Solid-State Circuits*, vol. 31, pp. 344–353, Mar. 1996.

[29] A. M. Niknejad, R. Gharpurey, and R. G. Meyer, "Numerically stable Green function for modeling and analysis of substrate coupling in integrated circuits," *IEEE Trans. Computer-Aided Design*, vol. 17, pp. 305–315, Apr. 1998.

[30] J. Zhao, W. Dai, S. Kapur, and D. E. Long, "Efficient three-dimensional extraction based on static and full-wave layered Green's functions," in *Proc. ACM/IEEE Design Automation Conf.*, 1998, pp. 224–229.

[31] J. Zhao, W. Dai, R. C. Frye, and K. L. Tai, "Green function via moment matching for rapid and accurate substrate parasitics evaluation," in *Proc. IEEE Custom Integrated Circuits Conf.*, 1997, pp. 371–374.

[32] F. Grover, *Inductance Calculations: Working Formulas and Tables*. New York: Dover, 1962.

[33] J. N. Lyness, "Some quadrature rules for finite trigonometric and related integrals," in *Numerical Integration*, P. Keast and G. Fairweather, Eds. Dordrecht, Germany: D. Reidel, 1987, pp. 17–33.

[34] B. Krauter and S. Mehrotra, "Layout based frequency dependent inductance and resistance extraction for on-chip interconnect timing analysis," in *Proc. 35th ACM/IEEE Design Automation Conf.*, 1998, pp. 303–308.

[35] S. Kim and D. P. Neikirk, "Compact equivalent circuit model for the skin effect," presented at the Int. Microwave Symp., June 1996, pp. 17–21.

[36] G. H. Golub and C. F. Van Loan, *Matrix Computation*, 3rd ed. Baltimore, MD: Johns Hopkins Univ. Press, 1996.

[37] K.-W. Chiang, S. Nahar, and C.-Y. Lo, "Time-efficient VLSI art-work analysis algorithms in GOALIE2," in *Proc. 25th ACM/IEEE Design Automation Conf.*, 1988, pp. 471–475.

[38] K.-W. Chiang, "Resistance extraction and resistance calculation in GOALIE2," in *Proc. 26th ACM/IEEE Design Automation Conf.*, 1989, pp. 682–685.

[39] W. R. Eisenstadt and Y. Eo, "S-parameter based IC interconnect transmission line characterization," *IEEE Trans. Comp., Hybrids, Manufact. Technol.*, vol. 15, pp. 483–490, Aug. 1992.

Dipak Sitaram received the B.S. degree in electrical engineering from the Birla Institute of Technology, Mesra, India, and the M.S. degree in electrical engineering from the University of Massachusetts, Amherst, 1989.

He has worked in the area of computer-aided design for IBM, Mentor Graphics, and Lucent Technologies. He is currently working on the development of extraction tools at Cadence Design Systems, New Providence, NJ. His research interests include high-frequency interconnect analysis and inductance modeling.

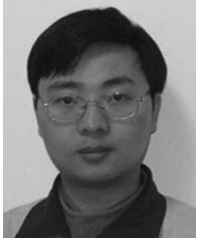

Yu Zheng (S'01–M'04) received the B.S. and M.S. degrees in microelectronics from Peking University, Beijing, China, in 1996 and 1999, respectively, and the Ph.D. degree in electrical engineering from Columbia University, New York, NY, in 2003.

He is currently a Senior Member of the technical staff at Cadence Design Systems, San Jose, CA.

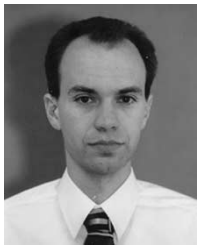

Kenneth L. Shepard (S'85–M'91–SM'03) received the B.S.E. degree from Princeton University, Princeton, NJ, in 1987 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1988 and 1992, respectively.

From 1992 to 1997, he was a Research Staff Member and Manager in the VLSI Design Department at the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he was responsible for the design methodology for IBM's G4 S/390 microprocessors. Since 1997, he has been at Columbia University, where he is now an Associate Professor. He also served as Chief Technology Officer of CadMOS Design Technology, San Jose, CA, until its acquisition by Cadence Design Systems in 2001. His current research interests include design tools for advanced CMOS technology, on-chip test and measurement circuitry, low-power design techniques for digital signal processing, low-power intrachip communications, and CMOS imaging applied to biological applications.

Dr. Shepard received the Fannie and John Hertz Foundation Doctoral Thesis Prize in 1992. At IBM, he received Research Division Awards in 1995 and 1997. He was also the recipient of an NSF CAREER Award in 1998 and IBM University Partnership Awards from 1998 through 2002. He was also awarded the 1999 Distinguished Faculty Teaching Award from the Columbia Engineering School Alumni Association. He has been an Associate Editor of IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION SYSTEMS and was the Technical Program Chair and General Chair for the 2002 and 2003 International Conference on Computer Design, respectively. He has served on the program committees for ICCAD, ISCAS, ISQED, GLS-VLSI, TAU, and ICCD.