IC Analyses Including Extracted Inductance Models†

Michael W. Beattie
Carnegie Mellon University
Dept. of ECE
5000 Forbes Ave.
Pittsburgh, PA 15213
+ 1 412 268 5238
beattie@ece.cmu.edu

Lawrence T. Pileggi
Carnegie Mellon University
Dept. of ECE
5000 Forbes Ave.
Pittsburgh, PA 15213
+ 1 412 268 6774
pileggi@ece.cmu.edu

Abstract
IC inductance extraction generally produces either port inductances based on simplified current path assumptions or a complete partial inductance matrix. Combining either of these results with the IC interconnect resistance and capacitance models significantly complicates most IC design and verification methodologies. In this tutorial paper we will review some of the analysis and verification problems associated with on–chip inductance, and present a subset of recent results for partially addressing the challenges which lie ahead.

Keywords
Interconnect; Inductance; Model Order Reduction.

1. Introduction
Due to the global nature of inductive coupling, extracted inductance models come in various forms and are derived using several simplifying approximations. For IC inductance extraction the models are generally either in the form of port inductances based on simplified current path assumptions, or described by a complete partial inductance matrix. For certain regular, overdesigned structures, there is also the possibility of a 2D (two–dimensional) infinite line approximation model.

In all cases, combining these inductance models with the resistance and capacitance models for IC design and/or verification purposes is an extremely difficult task. When the models are based on 2D field solutions, the inaccuracies associated with the non-infinite line effects can cause substantial errors. While complete 3D (three-dimensional) inductance matrices are easily generated using partial inductance equations [10][11] (even more easily than the corresponding 3D capacitance matrices), the resulting matrix is generally of unmanageable size and density to be useful for analysis purposes.

For example, unlike capacitance matrices which can be truncated to represent only localized couplings, simply discarding faraway mutual inductances can result in an unstable equivalent circuit model (positive poles). For this reason, inductance extraction often produces port inductances [7] instead of using the complete partial inductance matrix. Partial inductance extraction assumes that displacement currents are negligible, yet current paths are required to calculate port inductances. For package models this assumption is valid, but for CMOS ICs *all* currents are actually displacement (capacitor) currents. Therefore, in either the case of the complete partial inductance matrix or the port inductance models, assumptions are made which impact the way in which the models can be used for subsequent design and verification problems.

Model order reduction can help control the complexity of a provably stable circuit model which includes inductive couplings, but by adding such couplings the problem of N-port interconnect circuit model passivity [1][2] becomes much more difficult. In addition, low frequency poles from the partial inductance models can significantly complicate the model order reduction process [14].

In this tutorial paper we will explain some of the aforementioned problems in greater detail, then present a subset of partial solutions to cope with the IC inductance analysis and modeling challenges which lie ahead.

2. Design for Simplified Inductance Models
To date, the inductance extraction and modeling problem has been suppressed by design strategies which minimize the formation of significant long–range induction on chip. These approaches to avoid inductance as much as possible, however, can be costly in terms of IC area and performance.

One straightforward attempt to limit on–chip induction is found in the recent Alpha chip designs [15]. The wiring layers containing lines with high current density (hence high inductive capability) are sandwiched in between isolating metal planes above and below, as shown in Fig. 1. This design methodology has the following benefits:

- The magnetic field is blocked by the metal planes since currents induced in the isolation layer compensate the field outside completely. The coupling to adjacent layers and the substrate is suppressed.
- Since the layer distance is smaller than the wire width for the levels being isolated, the inductive coupling distance

† This work was supported in part by a Semiconductor Research Corporation Fellowship for Michael Beattie, by a grant from Intel Corporation and by funding from the National Science Foundation

DAC 99, New Orleans, Louisiana

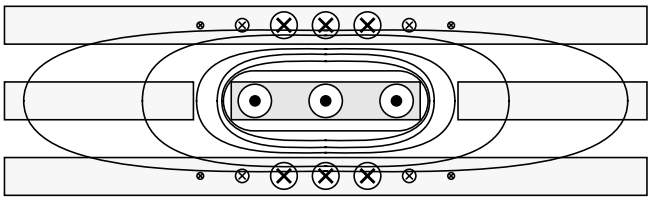

Fig. 1: Isolating layer approach (cross-section view). The magnetic field is blocked in the vertical direction and weakens quickly in the horizontal direction.

within the isolated layer is also reduced significantly. Hence, the inductance coupling is localized.

- The current return paths are now known by design. This makes calculating the loop inductance straightforward and a 2D model is reasonably accurate.

These advantages are at the manufacturing cost of two new metal layers for each original layer susceptible to major inductive interaction. These isolating layers also add significantly to the capacitive coupling of the surrounding circuitry which adversely impacts speed and power dissipation.

A more selective method of reducing the inductance between high–current wires is to introduce additional inverters (not buffers) in the lines at regular intervals. This causes the current to change direction in each adjacent interval. The self inductance of the entire line is reduced significantly, since the mutual coupling between neighboring segments of the same line subtract from the total inductance of the line rather than adding to it. The same effect can be found for neighboring, parallel lines if the buffer inverters are staggered with respect to each other, as depicted in Fig. 2.

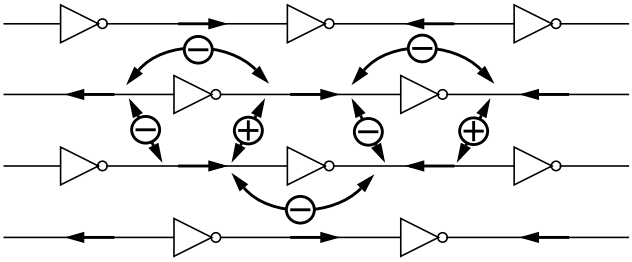

Fig. 2: Staggered inverter pattern with assumed current directions and mutual inductance polarities.

With enough added inverters, the reduction of the self inductance by inverter insertion can cause the inductive impedance to become insignificant compared with the ohmic impedance. At such a point, the inductance models are no longer required, and RC propagation dominates.

The negative side effects of buffer insertion are slightly higher power consumption and increased IC area, but primarily a reduction in circuit performance due to the additional delay introduced by the inverters.

3. Partial Inductance IC Models

The design strategies to avoid inductance outlined in the previous section are only temporary fixes to a phenomenon which will only become more pronounced as on–chip interconnect feature sizes decrease and operating frequencies increase. To model general 3D interconnect for which the return paths are not known *a priori* requires partial inductance elements. The partial inductance concept which was developed by Rosa [10] was introduced to the circuit design field by Ruehli [11]. Since the actual current and flux linkage loops are unknown, partial inductance is defined as the flux created by an aggressor segment through the virtual loop which a victim segment forms with infinity (see Fig. 3). The loop inductances are then represented by sums of the partial self and mutual inductances of the segments which form loops. Referring to Fig. 4:

$$L_{loop}^{exact} = \sum_i \sum_j S_{ij} L_{ij} \quad \text{with} \quad S_{ij} = \pm 1 \tag{1}$$

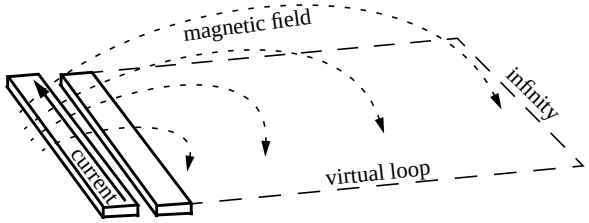

Fig. 3: Loop definition of partial inductance.

By *defining* each segment as forming its own return loop with infinity, partial inductances are used to represent the eventual loop interactions *without* prior knowledge of the actual current loops. The partial inductance values tend to be much larger than the corresponding loop inductances, which does reduce the numerical robustness of the inductance representations (ill–conditioned matrices). It is, however, the size and density of the partial inductance matrix that is most problematic for any circuit analysis.

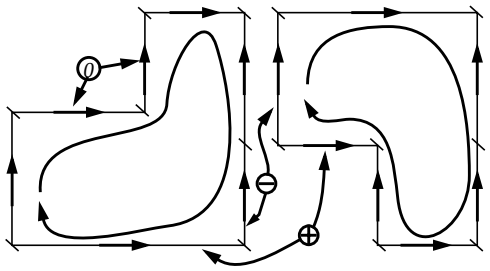

Fig. 4: Representation of current loops through partial inductance elements. The short arrows show the current directions, the long arrows identify the loop interactions. The signs of the S_{ij} terms are shown for selected segment pairs.

Loop inductance models are more compact, and intuitive, but more difficult to generalize for IC interconnect structures since the actual current return paths are not known *a priori*. Accurate return paths could be found only through circuit analysis, which requires the inductance models. Trying to avoid this *chicken-egg* problem by analyzing simpler RC or RC(self)L interconnect models can help to identify the return paths under certain circumstances. But once

inductance becomes dominant it can alter the current return paths significantly with respect to those found via the RC circuit.

4. IC Current Return Paths

A widely used approximation of the IC current return paths is to assume that the signal nets are shorted at the far end to a ground or that a return line is placed nearby so that the entire return current path is known without capacitance in the model. Tools such as *FastHenry* [4] can then be used to efficiently find loop inductances for long wires based on this simplified interconnect model. To analyze inductive coupling among IC bus lines, for example, these signal wires are shorted to a ground plane at the far end to form an RL loop. This assumes that the displacement currents through the capacitive couplings to neighboring nets are negligible.

Of course for on–chip CMOS circuits (contrary to packaging level systems), displacement currents are not negligible (see Fig. 5). To model the electromagnetic interaction accurately, long segments must be subdivided sufficiently in the current direction to take account of the capacitive coupling at intermediate points of the line.

As a consequence, assuming that all the current returns at the far end of the line overestimates the wire inductance significantly in most cases. In reality, coupling capacitors at intermediate points in the line can drain ac current from the wire to one of its neighbors before it reaches the end of the original conductor. This reduces the size of the current loops in the interconnect and reduces the loop inductance.

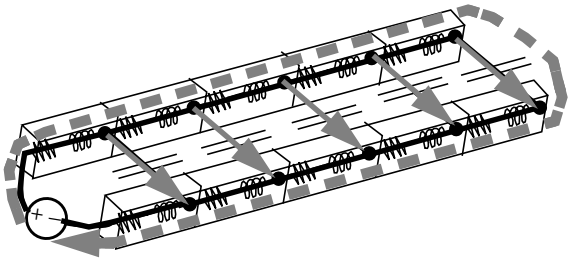

Fig. 5: Current return paths for loop inductance calculation. With capacitance (solid thick lines) and ignoring capacitance (dashed thick line).

This suggests that for CMOS IC extraction the capacitive couplings should be known prior to inductance extraction. However, with this capacitance coupling information, the increased number of current loops makes calculation of the port inductances difficult, if not impossible.

5. Sparsity and Stability

To avoid the need to combine the port inductances with the extracted capacitance, an alternative strategy is to work directly with the dense partial inductance matrix, and sparsify it as necessary. For example, there are numerous mutual inductance terms that are of negligible magnitudes. In practical terms this usually means that only couplings between segments which are closer to each other than a given window size are considered.

Such sparsification approaches can lead to indefinite partial inductance matrices, hence unstable equivalent circuit models, even for the simplest on–chip structures [9]. As empirical examples, Fig. 6 shows an unstable bus current pattern and Fig. 7 shows the minimal window radius required for this bus (normalized by the bus length) to maintain the partial L-matrix stability. The dependence on the normalized bus–width does not vary substantially with changing bus lengths (various curves in Fig. 7).

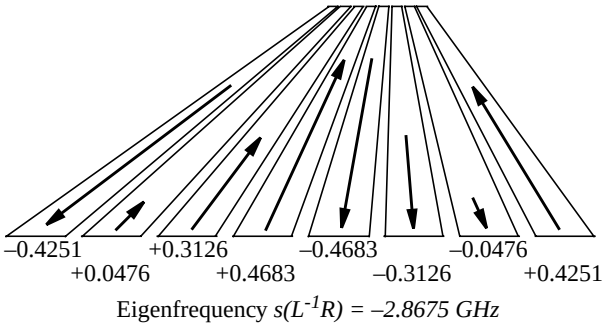

Fig. 6: Example bus with current pattern associated with a negative eigenvalue of $L^{-1}R$. The eight bus lines are 32 μm wide, spaced 12 μm, 1 μm thick and 1 cm long. The normalized eigenvector components are shown below the bus.

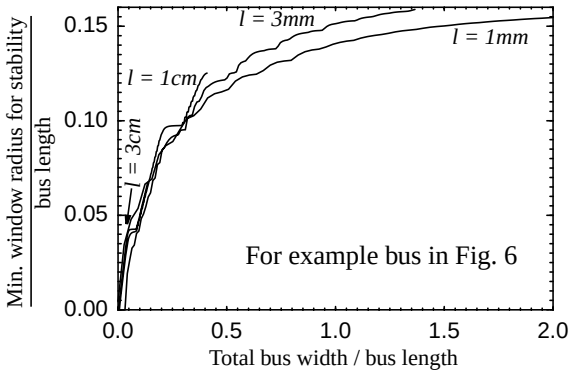

Fig. 7: Minimal radius for simple trunctation in order to preserve partial inductance matrix stability. Results for bus widths up to twice the bus length for various lengths.

For bus widths less than 5% of the bus length, the dependency in Fig. 7 is roughly linear and the window needs to include nearly the entire bus to ensure stability. This case is encountered in on–chip bus structures, and truncation will generally lead to unstable inductance approximations.

6. Sparsification via Equipotential Shells

A more general approach to providing L-matrix sparsity without compromising stability is via equipotential shells [8]. Artificially imposed equipotential shells shift the magnetic vector potential due to a given current distribution by a constant inside the shell and compensate it outside the shell, thus providing a stable sparse approximation. These shells are extremely easy to implement when the equipotential surface model is known [9].

The equipotential shell concept adds an artificial current distribution to the extraction problem which compensates the filament current's vector potential outside of this artificial current shell. Importantly, to simplify rather than com-

plicate the extraction process, only shell functions that are constant throughout the interior of the shell are applicable [9]. Simply stated, the shell potential compensates the filament potential outside the shell and shifts it by a constant value inside. This constant depends on a potential $A_{0,i}$ (which specifies the equipotential shell surface), the length of the victim line within the shell and the angle between aggressor and victim filament.

6.1 Potential Shell Model Stability

It has been shown that using such equipotential shells for generating sparse partial inductance matrices preserves stability [8]. For example, the magnetostatic energy, $I^T LI/2$, of a system of currents I is positive, since the exact inductance matrix L is positive definite for every physically realizable system. When shell currents are placed on equipotential surfaces, the total energy is

$$\frac{1}{2}\begin{bmatrix} I^T & -I^T \end{bmatrix}\begin{bmatrix} L_c & L_{c,sh} \\ L_{c,sh}^T & L_{sh} \end{bmatrix}\begin{bmatrix} I \\ -I \end{bmatrix} > 0 \qquad (2)$$

where L_c is the original inductance matrix, L_{sh} is the inductance among the shell currents and $L_{c,sh}$ is the coupling between shells and original currents. The inductance matrix using shells is $L^{eq} := L_c - L_{c,sh}$ and here the system energy is:

$$\frac{1}{2}I^T L^{eq} I = \frac{1}{2}I^T L_c I - \frac{1}{2}I^T L_{c,sh} I \qquad (3)$$

Since the current distribution of the shells is less dense than the original distribution, it follows that L^{eq} is positive definite [8].

6.2 Conservation of Loop Inductances

Since the shell reduces the magnetic vector potential everywhere, all partial inductances are underestimated. However, partial inductance does not contain information about the actual current flow in the system. This information is captured only by the loop inductance. If the shells of the conductors belonging to a specific loop are big enough that within the shell of *every* conductor *every* other filament of the loop is contained, then it was shown in [9] that:

$$L_{loop}^{shell} \equiv L_{loop}^{exact} \qquad (4)$$

If the lengths of the filaments become infinitesimally small, the loops become smooth, and this property also holds for arbitrary closed loops.

6.3 Guidelines for Shell Sizing

For the reasons stated above, we must make the shells large enough to encompass the significant loops. Given the design rules for on–chip power line spacings, one should make the shells large enough so that the each wire-segment's shell includes the nearest power and ground lines. A simple metric for determining the need to model the self inductance is the following. If the total line resistance (including the driver resistance) is smaller than the maximal self inductance impedance, the wire segment may comprise portions of non-negligible loops. For mutual couplings a similar metric can be applied, but the test is more difficult. The current in the aggressor can differ substantially from the current in the victim that it is mutually coupled to. Therefore, estimates for these currents must be obtained (actually estimates for the ratio of the time derivatives thereof) in order to estimate the voltage induced by the aggressor filament across the victim as compared with the ohmic voltage drop along the victim. One possibility is to use technology information that is available in form of *design rules* to find shell sizes which are appropriate for the given structure (such as a few periods of a return grid structure or a periodic signal bus).

6.4 Example

In [9] equipotential shells were used to calculate the loop inductance for the system shown in Fig. 8. In Fig. 9, L = 1.0 represents the exact inductance value, and ρ_0 is the shell size. Note that with increasing window size ρ_0 the inductance increases nearly logarithmically (relatively linear graphs in log–linear plots) and reaches the final value well before the system size of 1.4 mm is reached. Even at 1 MHz a shell size of 200 µm produces a very accurate approximation, but with 90 % sparsity. These equipotential shell approximations enable a sparse yet stable and accurate representation of the inductive interaction within the interconnect.

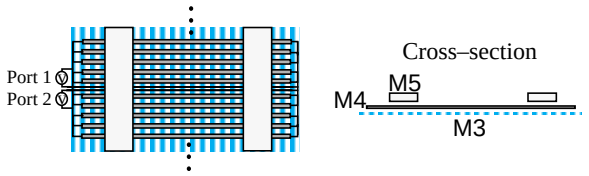

Fig. 8: Periodic return grid on M4 with two signal lines. Seventy lines in parallel in M4.

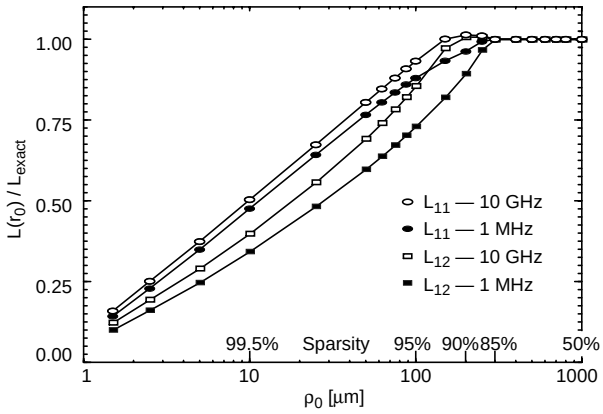

Fig. 9: Loop self and mutual inductance convergence for 1 MHz and 10 GHz vs. ellipsoid radius and corresponding sparsity. The curves for the higher frequency reaching 1.0 for smaller ellipsoids indicates a smaller current return

7. Model Order Reduction

Once the inductance models are combined with the resistance and capacitance models for the interconnect structures, the circuit complexity can be overwhelming [8]. For

this reason reduced order models of the resulting interconnect RLC circuit are necessary to analyze the circuit responses in a reasonable amount of time [1][2][3][4].

The addition of inductance to the system creates new challenges for the reduction process. With coupling, the full RLC representation of the interconnect has multiple drivers (ports) which make it much more difficult to ensure the passivity [1][2] of reduced order equivalent circuits. PRIMA [2] can theoretically guarantee passivity in terms of multi–input–multi–output (MIMO) models, but nothing developed to date has been able to do so in terms of the simpler single–input–multi–output (SIMO) N-port representations.

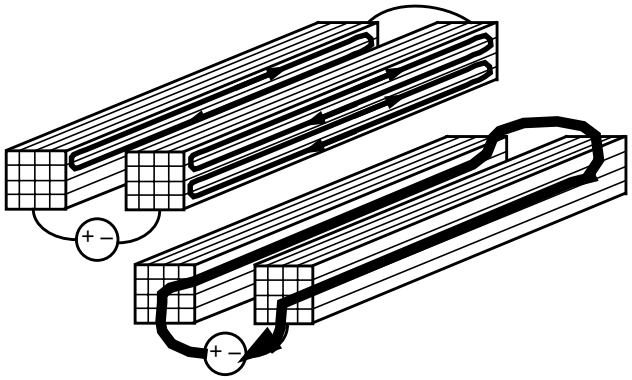

Fig. 10: Induced eddy currents (top) and the main segment loop current (bottom) for a two–wire example with four subdivisions on each edge orthogonal to the current direction.

An additional difficulty with handling mutual inductance models via model order reduction is the additional poles created due to modeling frequency dependence. IC interconnect inductance has a frequency dependence due to proximity effect and skin effect [12]. This frequency dependence is modeled by dividing wire segments into multiple parallel filaments [12], as shown in Fig. 10. These additional poles, which have been dubbed *bogus poles* [14], are inconsequential in terms of their residues, but can significantly complicate the model order reduction process.

As seen in Fig. 10 (top), the subdivision of wire segments has the side effect of allowing new current loops *within* the segments to form. This is possible since all filaments of each segment are shorted together at each end of the segment. These wire–internal current loops model the eddy currents within the conductors up to the resolution of the given subdivision. Eddy currents can only be generated through induction, therefore the "bogus" pole effect only arises when inductance is included in the interconnect model. The "real" poles are caused by the regular segment loop current, as shown in Fig. 10 (bottom).

Evidence that these added poles correspond to eddy currents can be seen in Fig. 11. For a pure RL circuit, the smallest pole will belong to the current loop with the largest loop inductance and the smallest resistance. This is the loop formed by the segments themselves, as shown in Fig. 10 (bottom). Since this loop does not change with

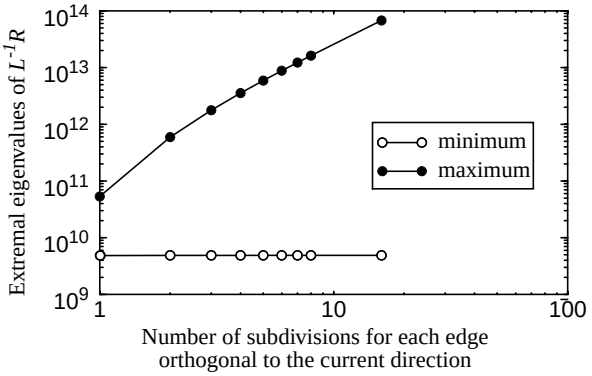

Fig. 11: Variation of smallest and largest pole of a system of two parallel conductors subdivided into parallel filaments to model frequency dependence.

increasing subdivision of the segments we expect the smallest eigenvalue of $L^{-1}R$ to be independent of the number of subdivisions, which is the case.

In contrast, the largest pole will be associated with the narrowest loop (smallest loop inductance) having the largest resistance. These are the loops formed by any two directly adjacent smallest subdivisions of the given segments as shown in Fig. 10 (top). Since the loop inductance (for small subdivision cross–sections) is roughly independent of (and the resistance increases quadratically with) the number of subdivisions, we expect a quadratic increase of the largest eigenvalue of $L^{-1}R$. This is exactly what we see in Fig. 11.

8. Driver Models for Timing Verification

Any of the sparse inductance models described above, or the corresponding reduced-order N-port models generated from them, can be combined with nonlinear driver/load models within a circuit simulator. As inductance becomes more prevalent on-chip, however, greater attempts will be made to analyze these models at higher levels of abstraction.

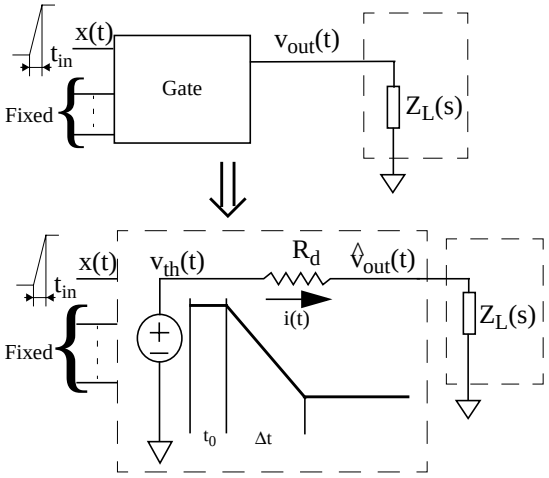

Fig. 12: The C_{eff} gate delay model from [13].

Although most gate models are pre–characterized as a function of capacitance loading only, the C_{eff} concept can be used to analyze RLC interconnects with these empirical gate

models [13]. As shown in Fig. 12, the most recent C_{eff} methodology is based on modeling the switching gate using a linear time-varying voltage source in series with a resistance. The linear resistance selected is independent of the load and input waveform. The parameters for the Thevenin voltage, $v_{th}(t)$, namely t_0 and Δt, can be pre–characterized as a function of C_{load} and t_{in} via SPICE. The values may be stored in a table or as regression fitted equations in terms of C_{load} and t_{in}. Refer to [13] for more details.

Once the model is pre–characterized, given a reduced order model for the RLC interconnect, an 'effective capacitance' is found which would draw the same average current as the reduced order model during the period of time when the Thevenin voltage is in transition. The initial guess for this C_{eff} is taken as the total capacitance of the RLC line, and the Thevenin voltage parameters for that capacitance are extracted from the table or equation model. Then, equating the average currents obtained by using these Thevenin voltage parameters, a new C_{eff} is computed. This iterative procedure is repeated until convergence is obtained.

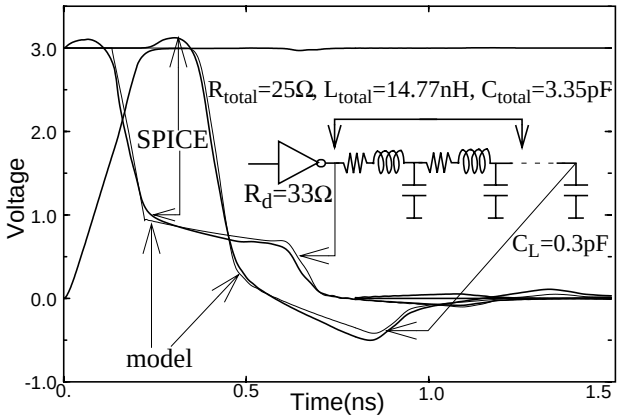

Fig. 13: Output waveforms for an inverter (L_{eff} = 0.5 μm, WN = 96 μm, WP = 240 μm) driving an RLC electronic package interconnect.

Fig. 13 shows an example of an inverter with an RLC load. The RLC circuit was approximated by a 8th order driving point admittance model [3]. The Thevenin voltage in the model was characterized by a single ramp for this example. As shown in the plot, the model captures the driving point and load responses very well.

9. Conclusions

The nature of inductive coupling and its relationship to capacitance makes extraction, analysis, design and verification of high-speed interconnects a challenging problem. This tutorial paper highlighted some of the key issues surrounding these challenges and hopefully proposed some promising partial solutions.

10. References

[1] K. Kerns, I. Wemple, A. Yang, *Stable and Efficient Reduction of Substrate Model Network using Congruence Transformation*, Proc. ICCAD 1995 (Nov. 1995).

[2] A. Odabasioglu, M. Celik, L. Pileggi, *PRIMA: Passive Reduced-order Interconnect Macromodeling Algorithm*, Proc. ICCAD 1997 (Nov. 1997).

[3] L. Pillage, R. Rohrer, *Asymptotic Waveform Evaluation for Timing Analysis*, IEEE. Trans. Computer-Aided Design, **9**, No. 4 (Apr. 1990).

[4] L. Silveria, M. Kamon, J. White, *Efficient Reduced-Order Modeling of Frequency-Dependent Coupling Inductance Associated with 3-D Interconnect Structure*, Proc. 32nd DAC (June 1995).

[5] A. Deutsch, et al., *Modeling and characterization of long on–chip interconnections for high–performance microprocessors*, IBM J. Res. Dev., **39**, No. 5, pg. 547–567 (Sept. 1995).

[6] F. Grover, *Inductance Calculations*, Dover Publications, New York (1946).

[7] M. Kamon, M. Tsuk, J. White, *FASTHENRY: A Multipole Accelerated 3–D Inductance Extraction Program*, IEEE Trans. Microwave Theory and Techniques, **42**, No. 9, pp. 1750–1758 (Sept. 1994).

[8] B. Krauter, L. Pileggi, *Generating Sparse Partial Inductance Matrices with Guaranteed Stability*, Proc. ICCAD 1996 (Nov. 1996).

[9] M. Beattie, L. Alatan, L. Pileggi, *Equipotential Shells for Efficient Partial Inductance Extraction*, Proc. 1998 IEDM (Dec. 1998).

[10] E. Rosa, *The Self and Mutual Inductance of Linear Conductors*, Bulletin of the National Bureau of Standards, **4**, pp. 301-344 (1908).

[11] A. Ruehli, *Inductance Calculations in a Complex Integrated Circuit Environment*, IBM J. Res. Dev., **16**, No. 5, pg. 470-481 (Sept. 1972).

[12] W. Weeks, L. Wu, M. McAllister, A. Singh, *Resistive and Inductive Skin Effect in Rectangular Conductors*, IBM J. Res. Dev., **23**, No. 6, pg. 652-660 (Nov. 1979).

[13] R. Arunachalam, F. Dartu and L. Pileggi, *CMOS Gate Delay Models for General RLC Loading*, Proc. ICCD 1997 (Oct. 1997).

[14] M. Kamon, N. Marques, L. Silveira, J. White, *Generating Reduced Order Models via PEEC for Capturing Skin and Proximity Effects*, Proc. 6th Meeting on Electr. Perform. of Electr. Packaging, San Jose (Nov. 1997).

[15] D. Bailey, B. Benschneider, *Clocking Design and Analysis for a 600–MHz Alpha Microprocessor*, IEEE J. Solid–State Circuits, **33**, No. 11 (Nov. 1998).