

A Fast Simulation Approach for Inductive Effects of VLSI Interconnects

Xiaoning Qi, Goetz Leonhardt, Daniel Flees, Xiao-Dong Yang, Sangwoo Kim,
Stephan Mueller, Hendrik Mau and Lawrence T. Pileggi*

Sun Microsystems, Inc.
410 North Mary Ave., SUN02-101
Sunnyvale, CA 94086
xiaoning.qi@sun.com

*Carnegie Mellon University
Department of Electrical & Computer Engineering
Pittsburgh, PA 15232
pileggi@ece.cmu.edu

ABSTRACT

Modeling on-chip inductive effects for interconnects of multi-gigahertz microprocessors remains challenging. SPICE simulation of these effects is very slow because of the large number of mutual inductances. Meanwhile, ignoring the non-linear behavior of drivers in a fast linear circuit simulator results in large errors for the inductive effect. In this paper, a fast and accurate time-domain transient analysis approach is presented, which captures the non-linearity of circuit drivers, the effect of non-ideal ground and de-coupling capacitors in a bus structure. The proposed method models the non-linearity of drivers in conjunction with specific bus geometries. Linearized waveforms at each driver output are incorporated into an interconnect reduced-order simulator for fast transient simulation. In addition, non-ideal ground and de-coupling capacitor models enable accurate signal and ground bounce simulations. Results show that this simulation approach is up to 68x faster than SPICE while maintaining 95% accuracy.

Categories and Subject Descriptors

J.6 [Computer-aided engineering]: Computer-aided design

General Terms

Algorithms

Keywords

VLSI interconnects, inductance, circuit simulation

1. INTRODUCTION

In multi-gigahertz microprocessor designs, ignoring the parasitic inductance of on-chip interconnects may cause inaccurate delay and noise estimations [1][2]. Inductive effects are quite evident in on-chip buses where signals may switch

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GLSVLSI'03, April 28-29, 2003, Washington, DC, USA.

Copyright 2003 ACM 1-58113-677-3/03/0004...\$5.00.




Figure 1. A stage-delay pushout is due to inductance. There is no significant inductive effect when only the victim switches.

simultaneously. Self- and mutual inductance between parallel bus wires can cause additional timing delay variations, overshoots, and significant inductive noise. This is due to the additive inductive coupling from all parallel wires. Figure 1 shows a victim's far-end waveform in a bus structure on three layers. When all the signals switch in the same direction, there are delay push-out and overshoots compared with the case when only the victim signal switches. Due to the long-range nature of the mutual inductive coupling, wider buses can cause more inductive problems for delays and noise. Therefore, the proposed return limited assumption [3] may not be valid for on-chip bus simulations. In addition, if all the wires in a bus are cut into n segments for the RLC distributed effect, the number of mutual inductances is $n(n-1)/2$. Because of the huge number of mutual inductances in a bus, SPICE runs very slowly and is not feasible for large industry examples. Meanwhile, the simple truncation of mutual inductance couplings in an inductance matrix can lead to large simulation errors and even result in an unstable circuit model [4]. Because of the in-efficiency of SPICE simulation, a fast and yet accurate time-domain transient simulation including inductance is desirable. This paper focuses on fast and accurate inductance modeling and transient simulation for on-chip buses.

2. SYSTEM OVERVIEW

There has been some previously reported work on improving the efficiency of simulating interconnects including inductance. The effective capacitance method for RC interconnect simulation [5] has been extended to RLC

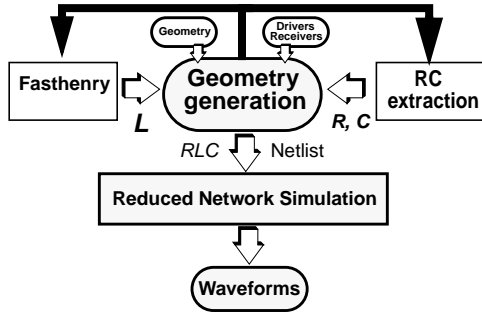


Figure 2. Flow chart of the *RLC* interconnect simulation tool.

interconnect simulation [6]. The efficient *Ceff* model works in terms of pre-characterizing the parameters of a time varying Thevenin voltage source model (in series with a fixed resistor) over a wide range of effective capacitance load values. To avoid the expensive procedure in [6] due to a Padé approximation, Kashyap [7] proposed a synthesis procedure for *RLC* circuits that guarantees a realizable reduced order circuit using the first four moments of the input admittance. While these methods work well with self-inductance, the mutual inductance is not included. For on-chip buses, it is the mutual inductance that makes signal cross-talk worse and jeopardizes the signal integrity. Pre-characterization needs to include the mutual inductance for each specific bus structure. In addition, the *Ceff* method and its related techniques are based on the timing analysis theory and may not work properly with noise estimations. Meanwhile, the non-linearity of devices plays an important role in inductive noise. There is no reported work on fast simulation for the impact of device non-linearity on the mutual inductive effects.

In this paper, an efficient transient analysis tool to model the *RLC* effect on timing and noise is presented. It captures the self-/mutual inductance effect, the non-linearity of drivers, the effect of non-ideal ground and de-coupling capacitors (decaps), and the 3-D geometries of interconnects. To accurately model the inductive effect, a fast and practical pre-characterization technique for on-chip buses is proposed. The non-linearity of drivers is one of the keys to accurately model the inductive effect in a bus structure. In order to have a realistic bus simulation, packaging parasitics (i.e. non-ideal grounds) are included with on-chip de-coupling capacitors (decaps). Figure 2 illustrates the flow chart of the implemented tool. Based on the user-specified bus geometries (i.e., signal wires and shields on the top, victim and bottom layers plus power/grounds on a power layer), a complete multi-bus structure is constructed automatically. Fasthenry [8] is used to calculate self- and mutual inductance. Together with *RC* wire models, a circuit netlist for a reduce-order interconnect simulation can be generated. A fast reduce-order circuit simulator [9] is used to generate the waveforms for timing and noise analysis. Among the other parameters that users can input are drivers/receivers, taps along each power/ground wire, decaps and signal input waveforms.

In Section 3, a non-linear driver modeling approach for accurate inductive effect is proposed, which is one of the key contributions of this work. Incorporating the device non-linear behavior in a reduce-order linear circuit simulator significantly increases inductance simulation speed while maintaining a good accuracy. The non-ideal ground modeling with decaps is

also discussed. In Section 4, simulation results from the SUN UltraSparc™ microprocessor designs are presented to validate the *RLC* simulation capabilities. Conclusions are in Section 5.

3. MODELING FOR ACCURATE INDUCTIVE EFFECT

3.1 Modeling for the Non-Linearity of Drivers

Traditionally, a linear resistor model is used for a device in a linear circuit simulator, which is sufficient for an *RC* simulation. For inductance simulation, a linear driver model is unable to capture the inductance (especially the mutual inductance) interaction with the drivers. Simulations show that a linear driver model leads to the underestimating of the signal inductive noise by 70% (Figure 3). In terms of delay

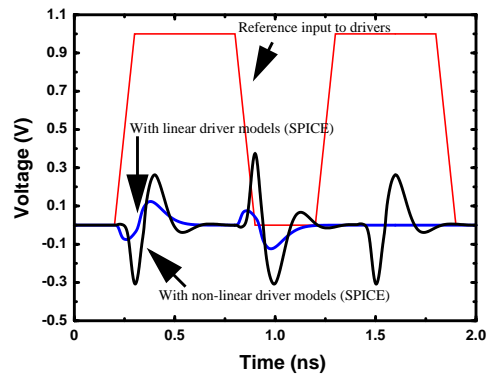


Figure 3. Cross-talk at a victim wire within a 16-bit bus on M6: linear driver models lead to the underestimating of the inductive noise by 70%.

simulation, although the linear model can predict delay with about 15% error, the waveform is not accurate - the overshoot/undershoot and fast ramp-up time are not properly captured. This is due to the fact that the linear driver model does not capture the fast di/dt behavior and can not produce the correct waveforms at the driver outputs. Figure 4 shows a

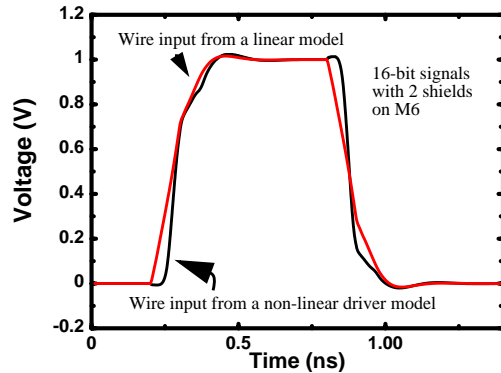


Figure 4. A non-linear driver model results in a much faster ramp-up time and a waveform with a small ledge.

driver output waveform of a 16-bit bus on M6. The difference between the two waveforms is due to the non-linearity of drivers/devices. The interaction between the self-/mutual inductance and the driver output voltage are not properly captured by a linear driver model.

To accurately model the waveforms at the driver output and capture the non-linearity of a device, a fast pre-characterization is done based on the actual interconnect geometries and the driver models. A non-linear circuit simulator, like SPICE, is used to extract the waveforms based on *simplified* circuit models for the interconnects to speed up the pre-characterization, but all self- and mutual inductance between the wires are preserved. For a transmission line system, the impedance which a driver sees is

$$Z_0 = \sqrt{\frac{R + j\omega L}{G + j\omega C}} = \sqrt{\frac{nR + j\omega nL}{nG + j\omega nC}}$$

where R , G , L , C are the resistance, conductance, inductance and capacitance per length of the transmission line. Parameter n is the number of segments used to model the distributed effects. The simplified circuit model uses the total resistance, conductance, inductance and capacitance values, respectively, instead of the values for per length. It presents the same impedance information to the driver as a complete distributed wire model, but with fewer segments to have a fast pre-characterization using SPICE. The simplified model is accurate to calculate the near-end waveforms, but it is not accurate enough to calculate the far-end signal waveforms. Results show that the pre-characterized waveforms at the interconnect near-end have a very good match against SPICE simulation with wire models using n segments. Each active driver can have a different output waveform depending on the wires which it connects to and their coupling. In other words, drivers on different layers and at a different bit position within a bus will have different waveforms. A piece-wise linear (PWL) approximation is then used to represent the characterized waveforms for each driver. A fast linear circuit simulator uses these waveforms as interconnect inputs. With a complete wire circuit model, the interconnect far-end waveforms can be accurately obtained (i.e., noise and delay). In Section 4, the pre-characterized piece-wise linear waveforms and the real driver waveforms will be plotted for accuracy comparison.

3.2 Modeling for the Non-Ideal Ground and Decaps

Unlike an RC extraction/simulation where power/ground wires may be modeled as ideal voltage connections, the ground/power wires should be modeled by distributed RLC elements in RLC simulation. The connection between any ground/power wires to the on-board VSS/VDD (e.g., C4 bumps) needs to be modeled for any inductance ground bounce. The non-ideal ground is modeled using RLC tap models.

Decaps provide charging and discharging currents at very high frequencies when the on-chip ground/power can not draw enough currents through the taps from the on-board power sources. In this work, decaps are modeled based on the area (that the bus structure covers) and the estimated decap values per chip area. The decap values are distributed along the VSS/VDD wires. With the help of decaps, the local VSS/VDD becomes more stable (i.e., less ground bounce).

4. SIMULATION EXAMPLES

Various examples from the SUN UltraSparc™ microprocessor designs are simulated to demonstrate the validity of the

proposed scheme. Test cases uses 0.18 μm technology. Results from the new tool are compared with the ones from the full SPICE simulation which includes all the non-linear drivers. An excellent waveform match is achieved, and the peak noise error is usually about 5% only.

In Figure 5, a 16-bit same layer bus with only two ground shields is constructed in order to exaggerate the inductive effects so that the accuracy can be better compared. The waveforms show the noise at the far end of the victim wire. The first peak noise is negative which is mainly due to the inductive effect while the first positive noise is largely due to the capacitive coupling.

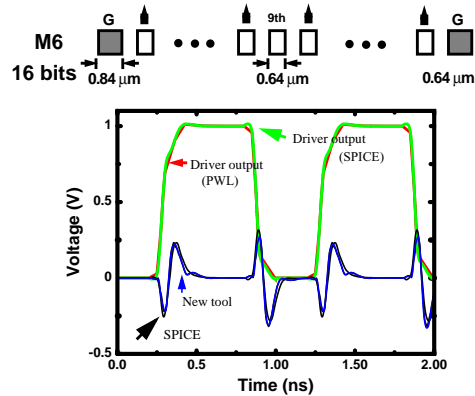


Figure 5. All the signals switch from low to high, and only the ninth signal keeps low. The pre-characterized driver PWL waveform is plotted with the one from SPICE simulation for accuracy comparison.

A more complicated example consists of three-layer buses with 12 bits on M6, 48 bits on M4 and M2. All signals are fully shielded, but the shield widths are comparable to the signal wires and they do not serve as completely effective shields for inductive effects. In the simulation, all the signals switch in the same direction at the same time. The victim wire is in the middle of the bus in M4. The signal delay (Figure 6(a)) and noise (Figure 6(b)) are plotted for the victim wire. In Figure 6(b), the noise waveform shows that the major noise is from the inductive coupling since there is no capacitive coupling between the signals in a fully shielded co-planar structure. Both signal delay and noise waveforms are almost indistinguishable from the ones with the full SPICE simulation. For this example, SPICE runs about 17 hours and the new tool runs about 15 minutes.

Figure. 7 shows the two noise waveforms of a signal wire within a bus where an ideal and a non-ideal ground tap model are used. Due to the parasitic inductance and resistance associated with the ground/power taps, a larger inductive noise is observed with a non-ideal ground model. The signal on a local ground wire on M6 is also plotted to illustrate the ground bounce. For decap effect, one simulation shows the delay of a bit in the center of a three layer bus without decaps is 15% larger than the one with decaps.

Table 1 shows the peak noise values of three different test circuits, and the comparison between our method and SPICE on the accuracy and run time. The run-time comparison indicates that smaller circuits gain less run-time speed-up. All examples are run on the SUN UltraSparc™ II machines.

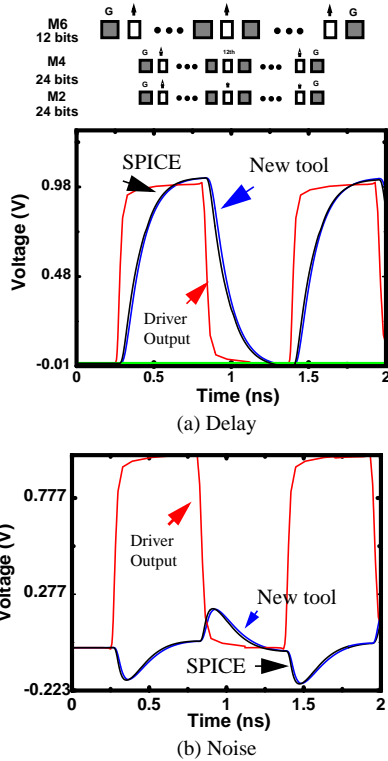


Figure 6. The delay and noise waveform of the victim wire on M4 is compared with the ones from the full SPICE simulation. The relative error of delay is only 3.9%.

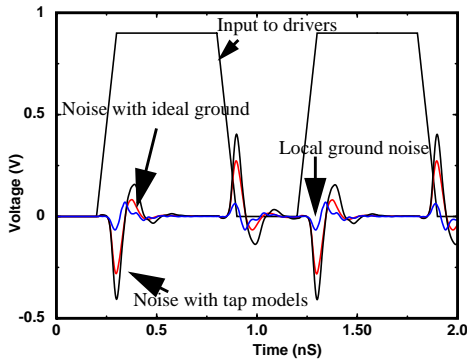


Figure 7. A three-layer bus example: the noise becomes larger due to the non-ideal ground. The inductance in the tap model is 500 pH and the resistance is 5 Ω .

5. CONCLUSIONS

Due to parallel routing and possible simultaneous switching of signals in a bus, mutual and self-inductance can cause additional delays and inductive noise for bus signals. A fast and accurate time-domain transient waveform simulation approach for *RLC* interconnect is presented. The capability to include the non-linearity of drivers in a reduced-order linear circuit simulator enables a much faster yet accurate simulation of on-chip interconnects with inductance. With tap models and decaps, bus simulation becomes more realistic, and ground bounce can also be observed. Results from the proposed method are in a

very good agreement with the ones from SPICE simulation while the simulation is 68X faster for large examples.

Table 1. COMPARISON OF THE PEAK NOISE FOR TEST CASES (V)

	Ctk1 (44 conductors)	Ctk2 (36 conductors)	Ctk3 (126 conductors)
SPICE	0.191	0.368	0.169
New tool	0.195	0.352	0.166
Error	-2.1%	4.3%	1.8%
Run-Time* SPICE:New Tool	25 min : 40 sec.	15 min : 30sec.	17 hr. : 15 min

* The run-time includes the pre-characterization and simulation time.

6. ACKNOWLEDGMENTS

The authors would like to thank Andrew Demas for his suggestions and testing of the program.

7. REFERENCES

- [1] A. Deutsch et al, "On-chip wiring design challenges for gigahertz operation", *IEEE Proceedings of the IEEE*, Vol. 89, No. 4, pp. 529-555, April 2001.
- [2] B. Kleveland, X. Qi, L. Maden, T. Furusawa, R. W. Dutton, M. A. Horowitz, and S. S. Wong, "High-frequency characterization of on-chip digital interconnects", *IEEE Journal of Solid-State Circuits*, Vol. 37, No. 6, pp. 716, 2002.
- [3] K. L. Shepard and Z. Tian, "Return-limited inductances: a practical approach to on-chip inductance extraction", *IEEE Trans. on CAD of Integrated Circuits and Systems*, Vol. 19, No. 4, pp. 425-436, April 2000.
- [4] M. Beattie and L. Pileggi, "Efficient inductance extraction via windowing", *Proc. DATE*, pp. 430-436, March, 2001.
- [5] J. Qian, S. Pullela, and L. T. Pileggi, "Modeling the 'effective capacitance' for the *RC* interconnect of CMOS gates", *IEEE Transactions on CAD*, Vol. 13, pp. 1526-1535, Dec., 1994.
- [6] R. Arunachalam, F. Dartu, and L. T. Pileggi, "CMOS gate delay models for general *RLC* loading", *Proc. ICCAD*, pp. 224-229, 1997.
- [7] C. V. Kashyap and B. L. Krauter, "A realizable driving point model for on-chip interconnect with inductance", *Proc. DAC*, 2000.
- [8] M. Kamon, M.J. Tsuk, and J. White, "FASTHENRY: a multipole accelerated 3D inductance extraction program", *IEEE Trans. Microwave Theory & Techniques*, pp. 1750, 1994.
- [9] A. Odabasioglu, M. Celik and L. T. Pileggi, "PRIMA: passive reduced-order interconnect macromodeling algorithm", *IEEE ICCAD '97*, pp. 58-65, Nov. 1997.