Bounding Bus Delay and Noise Effects of On-Chip Inductance

Michael Linderman, David Harris, David Diaz
Harvey Mudd College
301 E 12th St.
Claremont, CA 91711
mlinderman@stanford.edu, david_harris@hmc.edu, ddiaz@onehotlogic.com

Abstract

On-chip inductance depends on current return paths and is unreasonably computationally expensive to extract and model in the general case. A practical solution is to provide a well-defined power supply network so the current return paths are more predictable. This paper develops a model of bus delay and noise as a function of the physical dimensions of busses and the switching parameters. It applies this model to develop bounds on the inductive effects on delay and noise for on-chip busses in 180, 130 and 100 nm processes. If one power or ground line is interdigitated with every four bus lines, the RLC noise and delay are no more than 7% greater than RC models would predict. Designers may treat this delay and noise as small penalties for all busses rather than having to individually extract and model inductance on each bus.

Introduction

As process technology advances and signal rise times decrease, wire inductance has become an increasingly important design consideration. Once considered negligible, inductance can now be a significant factor in interconnect delay and coupling noise. The growing reach of inductance and increasing datapath widths require an understanding of inductive coupling on large multi-conductor arrays.

Noise and timing analysis has traditionally been performed with RC models [1]. These models are relatively easy to generate because capacitance is a localized phenomenon. Most of the electrostatic field terminates in neighboring conductors, allowing all other conductors to be ignored. Generating RLC models is considerably more difficult, requiring determining current return paths and including a larger set of conductors. Because extraction and simulation is so computationally expensive, it is advantageous to have a screening method to determine when inductance must be considered.

The very complexity of inductance extraction that makes a screening function necessary also makes it difficult to generate such a function. The large design space and strong geometric dependence of inductive coupling makes it impossible to extend screening functions developed for single conductors to on-chip busses [2]. Instead numerous authors have presented inductance modeling and screening techniques based on tables of precompiled results for specific bus geometries [3][4][5]. These methodologies increase extraction and modeling efficiency considerably, but require significant simulation prior to use. Thus these methods are not suitable as a first order tool for the designer attempting to select a bus topology with an acceptable level of inductive effects.

We propose limiting the allowed interconnect topologies by inserting a regular power supply network interleaved with the signal lines. The power and ground traces limit the range of inductive effects because return paths are close and well defined [6]. Using this topology we develop a model of bus delay and noise effects as a function of the physical dimensions of the interconnect for 180, 130 and 100 nm processes. We apply this model to develop bounds on the worst-case inductive effects on delay and noise. Designers can use these bounds to: 1) create accurate screening methods as a function of wire geometry 2) choose appropriate bus geometries for a given performance target 3) assess small delay and noise penalties for all busses without having to extract and model inductance on each bus.

Supply and Interconnect Topology

Fig. 1 shows the interconnect model used in this analysis. Each layer consists of wires with width *W*, spacing *S* and thickness *H*. The layers have conductivity σ. Between each layer is a dielectric of thickness *T* and dielectric constant ε. The layers alternate usage in the x and y directions, with orthogonal layers assumed to have approximately equal number of lines switching in each direction.

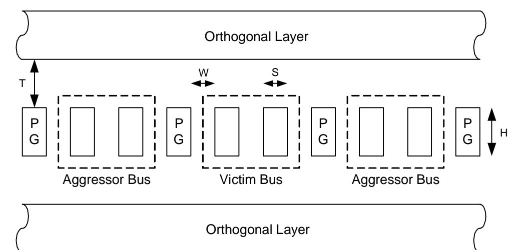

Fig. 1: Interconnect Model

The key to making this analysis tractable is to require a regular power supply network in which coplanar power/ground return lines are interleaved with the signal lines. The signal:reference (SR) ratio [6] indicates how many signals are routed between returns; for example Fig. 1 shows a stack with a SR = 2:1. In our frequency range of interest the return current is spread among multiple power/ground wires. As a result coupling from adjacent bus sets cannot be ignored, and so two neighboring aggressor busses are included. Ideal, zero resistance, returns are used to further simplify the analysis. The results remain physically meaningful, corresponding to interconnect strategies, like those in [6] and others, which use wide power/ground traces.

The interconnect topology was examined over a range of physical dimensions. The combinations of physical parameter are defined in Table 1.

Parameter	Range
Signal Width (W)	1x, 1.5x, 2x, 4x, 8x
Signal Spacing (S)	1x, 1.5x, 2x, 4x, 8x
Line Length (mm)	.5, 1, 2
Conductors	1, 2, 3, 4, 6, 8, 10, 20

Table 1: Physical Parameter Combinations

Circuit Model

The circuits were simulated in HSPICE at 70^o C using the TSMC 180 nm models and the 130 and 100 nm Berkeley Predictive transistor models [7]. Interconnect parameters are based on the TSMC process [8] and the 2002 International Technology Roadmap for Semiconductors (ITRS) projections [9]. The specifications used are defined in Table 2.

Parameter	180	130	100
V_{DD} (V)	1.8	1.5	1.2
FO4 Inverter Delay (ps)	75.6	53.3	42.7
Min. Wire Width W (µm)	.28	.19	.16
Min. Wire Spacing S (µm)	.28	.19	.16
Wire Height H (µm)	~2W	.38	.32
Dielectric Thickness T (µm)	~H	.38	.32
Dielectric Constant ε	~3.8	3.2	2.8
Wire Conductivity $(\Omega\text{-}\mu m)^{-1}$	24.8	45.4	45.4

Table 2: Model parameters

The inductance, capacitance and resistance matrices were extracted with the 2-D solver in HSPICE. The use of a 2-D solver requires approximating the effect of the orthogonal layers. The orthogonal conductors are parallel to the magnetic field generated by the aggressors so these layers can be ignored for inductance extraction. The orthogonal layers cannot be ignored for capacitance extraction, as a significant portion of the generated electric field terminates in adjacent layers. In a densely routed chip, defined as 50% metal on a given layer, the orthogonal layers can be modeled as solid planes. Because it is assumed that an equal number of traces on the orthogonal layer are switching in both directions and these wires can be considered static, the entire layer can be treated as a solid reference plane for the purpose of capacitance extraction. Hence two extractions are performed: one is done without reference planes to obtain inductance, while a second is done with planes to obtain capacitance.

Frequency dependent resistive effects and mutual resistance are ignored for all models. DC wire resistance is unaffected by the orthogonal layers, and is extracted along with the capacitance. RC models are generated from the RLC models by dividing all entries in the inductance matrix by 20.

The simulation setup is shown in Fig. 2. Wires are modeled using the HSPICE W-element, a multiconductor lossy frequency dependent transmission line [10]. Unit inverter size is selected for a nominal fanout-of-4 (FO4) to the total wire capacitance for a 1mm long, minimum width and spacing wire. For the 130 nm process, a 1 mm, minimum width and spacing wire has a total capacitance of 200 fF. The unit size inverter gate capacitance is thus 50 fF, which corresponds to 28.6 µm of gate width. Driver and load size is varied as defined in Table 3. Thus the driver ranges from 14 to 57 µm and the load ranges from 14 to 228 µm of gate width.

Parameter	Range
Driver (X unit size)	.5x, 1x, 2x
Load (X driver size)	1x, 2x, 4x

Table 3: Diver and Load Combinations

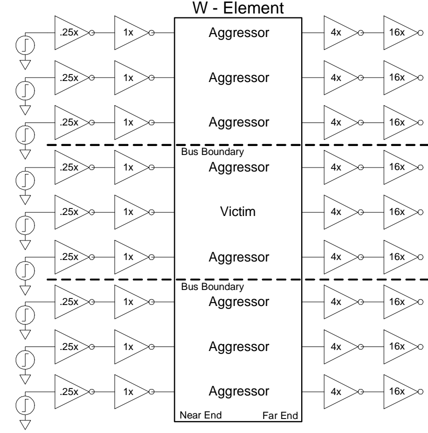

Fig. 2: Simulation setup for S:R = 3:1

The center trace in a bus is the victim, and is switching for delay simulations and quiet for noise. All other traces are considered aggressors and are driven according to the switching scenario being investigated. Table 4 lists the six scenarios considered. All aggressors switch simultaneously; staggered input scenarios were not considered to simplify the analysis.

Scenario	Description
1	All aggressors ↑
2	Adj. aggressors ↑, others ↓
3	All but adj. aggressors ↑
4	All but adj. aggressors ↓
5	Adj. aggressors ↓, others ↑
6	All aggressors ↓

Table 4: Switching scenarios

Propagation delay is measured as the time between 50% of V_{DD} on the input to the driving inverter and output of the W-element for the victim trace. Noise is measured as the maximum voltage of the victim trace at the far end of the W-element.

Results

Simulations were performed to consider all combinations of switching scenarios, signal width, spacing, and length, S:R ratio, and driver and receiver sizes in each process.

The maximum percentage difference between RLC and RC propagation delay for 180, 130 and 100 nm processes, all lengths, switching scenarios and driver and load combinations is plotted for each S:R ratio and width/spacing combination in Fig. 3 (a-c).

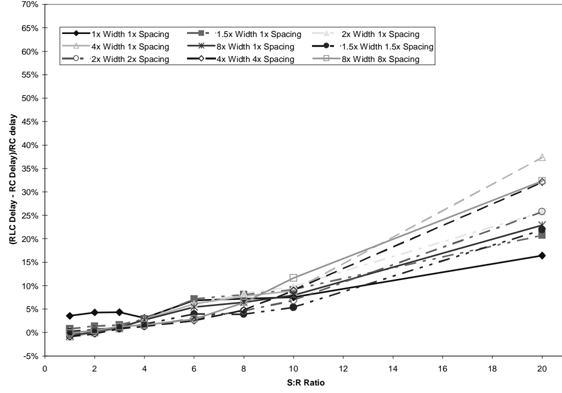

(a) TSMC 180 nm

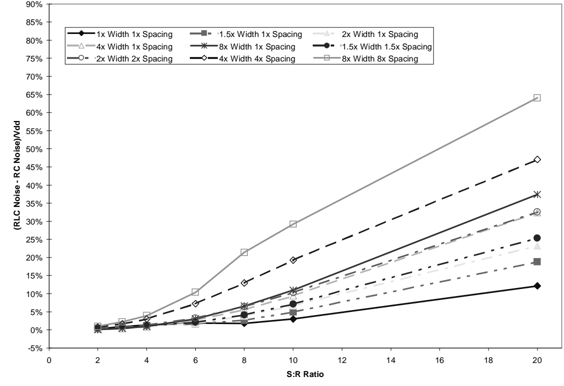

(a) TSMC 180 nm

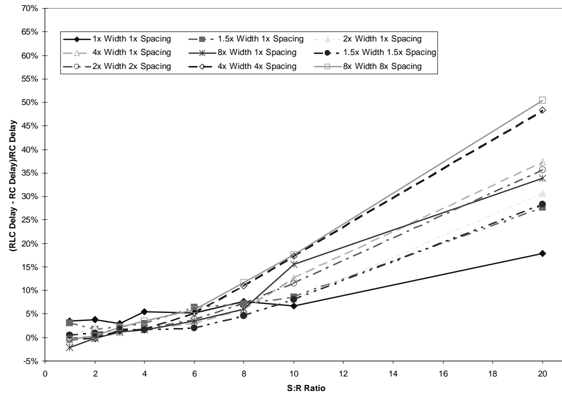

(b) Berkeley 130 nm

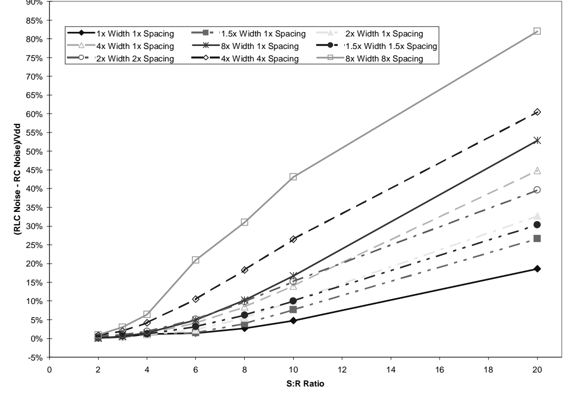

(b) Berkeley 130 nm

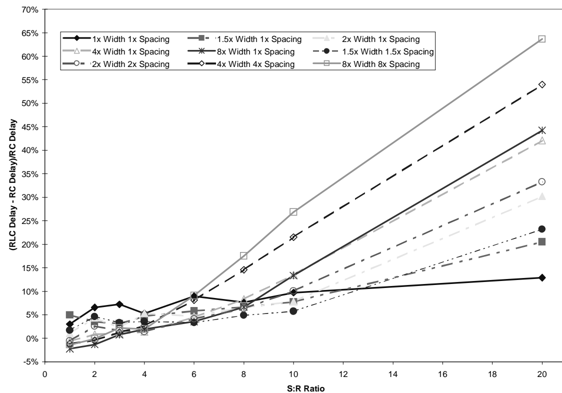

(c) Berkeley 100 nm

Fig. 3: Comparison of RLC and RC propagation delay

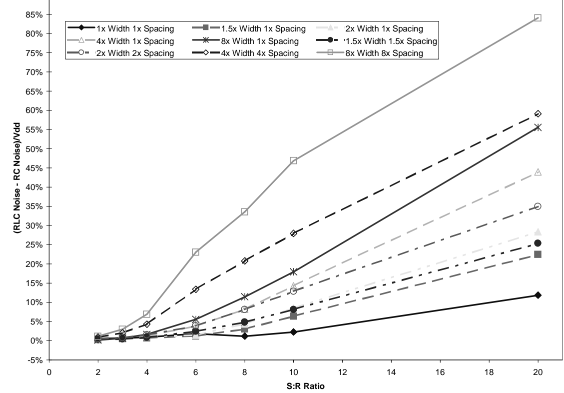

(c) Berkeley 100 nm

Fig. 4: Comparison of RLC and RC far-end noise

The maximum difference between RLC and RC far end victim noise as a percentage of V_{DD} is similarly plotted in Fig. 4 (a-c).

Discussion

The S:R ratio is a critical variable to determine inductive effects. As the plots show, the impact of inductance on propagation delay and victim noise depends strongly on this ratio. Whereas capacitive coupling is confined to a victim's nearest neighbors, inductive coupling falls off with distance at a much lower rate.

Additional signal traces in a bus, although negligible for capacitive coupling, will contribute significantly to inductive effects. The increase from a S:R = 10:1 to 20:1 can result in worst-case inductive effects growing by a factor of four or more.

Analysis that focuses on single conductors or small busses can significantly underestimate the inductive effects present in larger busses. At small S:R ratios the shielding effect is prominent [11]. The minimum total propagation delay occurs when the output resistance of the driver, R_{dr}, equals the lossy characteristic impedance of the wire, Z_{line}, defined in Eq. 1.

$$|Z_{line}| = \sqrt{\frac{\sqrt{R_{line}^{\ 2} + (\omega L_{line})^2}}{\omega C_{line}}} \quad (1)$$

At small to intermediate wire widths Z_{line} is comparable to R_{dr} and the system lies near the delay minimum. As the width increases, the capacitance increases, and R_{line} and L_{line} decrease such that Z_{line} grows smaller. Since $Z_{line} < R_{dr}$, the wire inductance tends to push the system towards to the delay minimum. Thus although the inductance is increasing relative to the wire resistance, the worst case negative effects of that inductance (increased overall delay) are reduced for wider wires.

However as the S:R ratio grows the system is no longer dominated by the shielding effect, and the above stated conditions no longer hold. For S:R > 10, mutual inductive coupling is a significant source of added delay. In a 100 nm process for a S:R = 20:1 bus inductive coupling results in relative error of 13% to 64%. Furthermore, the inductive effects grow with increasing width and spacing. Because the wire is now a significant contributor to the propagation delay, increasing width, and decreasing wire resistance makes inductive effects more prominent, not less.

Inductively coupled noise shows a similar dependence on width, spacing and S:R ratio. Noise is particularly sensitive to increasing wire width and spacing. As Fig. 4 (a-c) shows, the difference between RC and RLC noise for 8x and 4x width and spacing becomes significant at much smaller S:R ratios, and grows much faster than other geometries. In the 100 nm process, 8x and 4x width and spacing with a S:R = 6:1 have a noise difference of 23% and 14% respectively of V_{DD}, while all other geometries do not exceed 6%. The other geometries do not reach that level until S:R > 10.

Inductance-aware bus design is a compromise between all of these variables, attempting to maximize performance while minimizing area devoted to wiring and power/ground distribution. The primary application of the bounding plots is to assist in a selecting the appropriate bus geometry prior to layout and without extensive simulation. For example in the 100 nm process, a S:R = 10:1 bus has limited inductive effects for all but 8x and 4x width and spacing wires, less than 13% in noise and 10% in delay. Similar busses in the 130 nm and 180 nm processes show even smaller inductive effects.

Although careful selection of width and spacing reduces inductive effects, in general a S:R = 4:1 bus, independent of wire geometry and process will keep worst case relative error and the noise difference under 7% because each victim has few aggressors.

Limitations

This analysis is limited by the consideration of a just a single layer. Although there is no inductive coupling to an adjacent orthogonal layer, there is coupling between the parallel layers. Traces on these other layers would behave similarly to long distance aggressors, and would result in increased noise and coupling induced relative delay error, particularly for large width and spacing. A similar approach could be used to bound these effects, and would be even more effective since multi-layer extraction and simulation is even more computationally expensive.

Other limitations include considering only simultaneously switching aggressors and assuming ideal return paths with no mutual resistance.

Conclusion

This paper has presented an alternate method to determine worst-case inductive effects on delay and noise for multiconductor coplanar busses. Although the general case of inductance extraction is very difficult, limiting possible interconnect topologies by inserting a regular array of power/ground traces into busses makes current return paths more predictable and limits inductive effects. We develop a comprehensive model of bus delay and noise as a function of interconnect geometry and switching conditions and use this model to develop bounds on worst case inductive effects for 180, 130 and 100 nm processes. These results can allow designers to choose appropriate bus geometries for a given performance target prior to routing and without additional simulation.

The analysis showed that an S:R = 4:1 bus shows inductive effects of less than 7% for both delay and noise, independent of wire size, driver and load configuration and process.

References

[1] R. Ho, K. Mai, and M. Horowitz, "The future of wires," in *Proc. IEEE*, vol. 89 no. 4 April 2001, pp 490 – 50

[2] Y. I. Ismail, E. G. Friedman, and J. L. Neves, "Figures of merit to characterize the importance of on-chip inductance," in *Proc. DAC*, June 1998, pp. 560 – 565.

[3] L. He, N. Chang, S. Lin, O. S. Nakagawa, "An efficient inductance modeling for on-chip interconnects,' in *Proc. CICC*, 1999, pp. 457 – 460.

[4] Y, Cao et al., "Effective on chip inductance modeling for multiple signal lines and applications for repeater insertion," in *Trans. on VLSI*, vol. 10 no. 6 Dec. 2002, pp 799 – 805.

[5] T. Sato, T. Kanamoto, A. Kurokawa, Y. Kawakami, H. Oka, T. Kitaura, H. Kobayashi and M. Hashimoto, "Accurate prediction of the impact of on-chip inductance on interconnect delay using electrical and parameter-based RSF," in *Proc. ASP-DAC*, Jan. 2003, pp. 149 – 155.

[6] S. Morton, "On-chip inductance issues in multiconductor systems," in *Proc. DAC*, June 1999, pp. 921 – 926.

[7] Y. Cao, T. Sato, D. Sylvester, M. Orchansky, and C. Hu, "New paradigm of predictive MOSFET and interconnect modeling for early circuit designs," *Proc. CICC*, May 2000, p 9.3.

[8] C. Diaz et al., "A 0.18 μm CMOS logic technology with dual gate oxide and low-k interconnect for high-performance and low-power applications," *VLSI Technology Symposium*, 1999, pp. 11 – 12.

[9] SIA, "International technology roadmap for semiconductors," 2002 Edition.

[10] Avant!, "HSPICE User's Guide", 2001.4.2

[11] M . El-Moursy and E. Friedman, "Shielding effect of on-chip interconnect inductance," *Proc. Great Lakes Symp. VLSI*, Apr. 2003, pp. 165 – 170.