

On-chip Inductance Issues in Multiconductor Systems

Shannon V. Morton
Alpha Development Group
Compaq Computer Corporation
Shrewsbury, MA 01545
+1 508 841 2281
Shannon.Morton@compaq.com

ABSTRACT

As the family of Alpha microprocessors continues to scale into more advanced technologies with very high frequency edge rates and multiple layers of interconnect, the issue of characterizing inductive effects and providing a chip-wide design methodology becomes an increasingly complex problem. To address this issue, a test chip has been fabricated to evaluate various conductor configurations and verify the correctness of the simulation approach. The implementation of and results from this test chip are presented in this paper. Furthermore the analysis has been extended to the upcoming EV7 microprocessor, and important aspects of the derivation of its design methodology, as pertains to these inductive effects, are discussed.

Keywords

Alpha microprocessor, semiconductor, interconnect, buses, inductance, resistance, capacitance, RLC, noise, cross-talk, transmission line.

1. INTRODUCTION

The Alpha family of microprocessors has consistently maintained leadership performance since the inception of the 21064 microprocessor in 1992. This performance advantage has been achieved through a variety of advanced architectural and circuit design techniques. Each progressive implementation of the Alpha microprocessor has required innovative solutions to a number of design challenges, such as clocking, latches, power management, race analysis, and capacitive coupling [1,2]. With the ongoing trend towards ever-increasing clock frequencies, the management of inductive coupling becomes another important concern.

The use of a true 64-bit super-scalar multiple instruction issue architecture with multiple parallel functional units results in a vast number of buses to be routed throughout the chip, potentially across long electrical distances. This therefore necessitates a sufficient number of metal layers for routing signal wires as well

as the power and ground network. Typically a coarse pitch metal layer would be used for global routing (perhaps between functional units) and a fine pitch metal layer with greater wiring density would be used for local routing (perhaps within functional units). The extensive use of advanced dynamic circuit techniques also facilitates the use of a very high frequency clock.

In such an environment it is absolutely critical to manage any noise that can be induced on these buses, since both short and long term functional errors can result. Noise that couples a node above VDD (or below VSS) can result in a number of reliability concerns including hot carrier effects and time dependent dielectric breakdown. Noise that couples a node below VDD (or above VSS) may cause functional errors if the receiving latch stores the incorrect noisy state, either directly from this node or through the propagation of a glitch. Dynamic nodes are especially prone to this form of error since there is no static pull-up device during activation that can recover from the noise. Noise can also manifest itself as a delay variation rather than a logic variation if the victim node happens to be switching during the noise injection.

In a high performance microprocessor there may be many sources of noise, such as DC and AC variations in the supply rails, Miller coupling (from C_{gs} or C_{gd} of subsequent devices), charge sharing, and hot carrier collection. One of the most predominant sources of noise however is capacitive coupling from adjacent nodes as illustrated in Figure 1a. The victim node V should ideally remain at VDD however the aggressor nodes A on either side are all switching low. Due to the lateral capacitance between these nodes, V will experience a voltage dip (Figure 1b) whose magnitude is dependent upon the relative values of its lateral and total capacitances (including any gate capacitance). A dynamic node may not recover from this induced noise (termed AV noise), whereas a static node is likely to recover in time. Nonetheless, if the magnitude of the noise is beyond the unity gain point of the receiving logic (often defined as or related to the noise margin) then this glitch will be amplified rather than attenuated. It is therefore of importance to ensure that any induced noise on a victim node is well below this noise margin (allowing some additional margin for other noise sources).

The phenomenon of capacitive coupling is well understood and can be managed relatively easily since the electrostatic fields predominantly terminate on adjacent nodes. This means that one can simply model the capacitance to a victim's nearest neighbors (as illustrated in Figure 1) and effectively ignore any capacitance to nodes beyond this. Obviously, if the layers above and below do not run orthogonally with sufficient metal density then this

Permission to make digital/hardcopy 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, the copyright notice, the title of the publication and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 99, New Orleans, Louisiana
(c) 1999 ACM 1-58113-109-7/99/06..\$5.00

assumption may not be valid, however this is often the exception rather than the rule.




Figure 1: (a) 5 conductors in a stripline (b) Example of AV capacitive coupling.

At short lengths the phenomenon of inductive coupling is negligible since the edge rate of the signal is long compared to the signal's time of flight [3]. Beyond this length, and for the very fast edge rates found in modern technologies, these effects may be significant. This phenomenon is considerably more difficult to model than capacitive coupling since the magnetic fields may extend with sufficient amplitude well beyond the victim's nearest neighbors. In essence, the capacitance matrix for a system of N parallel conductors has off tri-diagonal terms ≈ 0 whereas the inductance matrix does not. Furthermore, the resistance matrix is tightly coupled to the inductance matrix due to the frequency dependent dispersion of the return current flow and its depth of penetration into the conductors, and both therefore vary with frequency. The resistance matrix includes the resistance of the return current flow and contains both self and mutual terms, of which the effects of the latter may not be negligible (as is often assumed in standard RC interconnect models).

In essence then, it is of importance to investigate and quantify the effects of inductance (on delay and AV noise) in the presence of multiple conductors rather than for the simple case of a single conductor. Furthermore, with the advent of multiple metal layers in modern microprocessors, the effects of conductors on other parallel layers also need to be analyzed. A test chip has been fabricated to address some of these issues as well as to ascertain the relative merits of various conductor configurations in terms of their ability to minimize signal noise.

2. GENERAL PRINCIPLES

Consider now the example of Figure 2 with just one conductor shielded on both sides by wide (5x the conductor width) supply rails and flanked above and below by orthogonal wires which are assumed unable to conduct currents in the return path. Assuming a uniform resistivity of $2.0 \mu\Omega/\text{cm}$, a dielectric constant of 4.0, and all spatial dimensions of $1\mu\text{m}$, then at a frequency of 1GHz the RLC matrices for this conductor can be calculated as: $R_{\text{eff}}=220\Omega/\text{cm}$; $L=4.5\text{nH}/\text{cm}$; $C=2.4\text{pF}/\text{cm}$ (this frequency corresponds to a modest edge rate of $\approx 300\text{ps}$ via the well-known approximation: $f=1/\pi\tau_r$). At this frequency $R_{\text{eff}} \gg \omega L$ and $\sqrt{LC} = 105\text{ps}/\text{cm}$ implying that for most wire lengths the effects of inductance can be ignored and the wire can be modeled using RC elements. Note that the effective resistance of $220\Omega/\text{cm}$ is 1.1x the resistance of the wire because of the additional resistance in the return path. This ratio is termed the effective resistance factor, or ERF.




Figure 2: Multiconductor system with wide returns

Now consider the multiconductor system of Figure 2 with ten conductors shielded on both sides by supply rails. A 1GHz extraction of this system results in RLC matrices each containing 10×10 elements. These can be simplified into unity terms by considering the effective RLC matrices for the two conditions of all wires switching in unison ($V_i = V_{i+1}$), and alternate wires switching in opposition ($V_i = -V_{i+1}$). These effective RLC matrix elements are shown in Table 1.

Switching Condition	R_{eff} (Ω/cm)	L_{eff} (nH/cm)	C_{eff} (pF/cm)
1. In Unison	410	31.0	1.55
2. In Opposition	201	4.7	3.25

Table 1: Equivalent RLC parameters for 10 wires

For condition 2 above we see that once again $R_{\text{eff}} \gg \omega L_{\text{eff}}$ and $\sqrt{L_{\text{eff}}C_{\text{eff}}} = 125 \text{ps}/\text{cm}$. Hence, the system can be modeled using RC elements with $R_{\text{eff}}=200 \Omega/\text{cm}$ and an effective capacitance which adds in the Miller component. Note that in this instance the ERF is 1.0 (excluding skin effects) since opposing current flows cause all of the mutual resistance terms to cancel out. Note also that the effective inductance remains low since neighboring signals provide for tight return paths.

For condition 1 above we now see that $R_{\text{eff}} \approx \omega L_{\text{eff}}$ and $\sqrt{L_{\text{eff}}C_{\text{eff}}} = 220 \text{ps}/\text{cm}$. As such, the effects of inductance are significantly exaggerated for this condition despite the additive effect of the mutual resistances (resulting in an ERF of 2.0 excluding skin effects). The effective inductance is large in this case due to the wide expanse of current injection and the distant return currents in the supply rails, and the effective capacitance is low since the Miller component is subtracted from the nominal. In this case the system *must* be modeled with the effective RLC elements. It is clear then that to exacerbate the effects of inductance on delay and AV noise, one should consider the case in which all wires switch in unison.

3. TEST CHIP CONFIGURATION

As shown in Figure 3, the design used to evaluate the performance of a given system of conductors was essentially a ring oscillator in which two legs of the oscillator traverse a reasonably long interconnect path (whose chosen length will be discussed in Section 6). These two interconnect paths were widely separated and additionally isolated by wide, low resistance supply rails to minimize unwanted coupling. Additional logic was inserted into one end of the ring oscillator to enable external clocking (needed primarily for e-beam probing as discussed in Section 5) and to ensure a known state upon start up (to prevent multiple wavefronts from propagating). This logic also enabled only one of sixteen different conductor configurations to be active at any time.




Figure 3: Generalized structure of the test bed

In all cases, the conductor configuration consisted of 10 aggressor wires on M4 (and/or M2) and a victim wire on M1 which was tied to VDD at a given position along its length. A 4 layer metal process was used with M1/2 at fine pitch and M3/4 at coarse pitch. This circuit enabled the effects of inductance on delay and overshoot to be measured on the aggressor wires and for AV noise to be measured on the victim wire. Table 2 shows the various conductor configurations that were incorporated onto the test chip. The heading “#M4” indicates the number of M4 aggressors, “M4 S:R” indicates the corresponding signal to return (SR) ratio, and “M4R” indicates the width of the return wire relative to the signal wires. Similar headings apply to M2 (for which the return width was always 1x), and the heading “M3P?” indicates whether or not a reference plane was present on M3.

Structure	#M4	M4 S:R	M4R	M3P?	#M2	M2 S:R
A	10	None		No		
B	10	None		Yes		
C	10	10:1	5x	No		
D	10	10:1	1x	No		
E	10	1:1	1x	No		
F	10	1:1	1x	Yes		
G	10	2:1	1x	No		
H	10	2:1	3x	No		
I	10	5:1	1x	No		
J	4	None		No	6	None
K	4	None		Yes	6	None
L	4	None		No	6	6:1
M	4	4:1	1x	No	6	None
N	4	1:1	1x	No	6	None
O				No	10	None
P				Yes	10	None

Table 2: Structures incorporated onto the test chip

Structure A was designed to produce the worst case behavior for a system of conductors since no nearby return paths exist. Structure B enabled the effects of a reference plane to be evaluated whilst structures C through I evaluated the effects of various SR ratios with various return widths. Structures J through N consider similar scenarios for conductors on both M2 and M4, whilst structures O and P consider conductors only on M2. This vast array of configurations was designed to enable the relative performance of each method of noise minimization to be measured, and to verify the correctness of the simulation procedure across a range of different conductor configurations.

For all layers, lateral fingers were incorporated into the layout of the interconnect structures as shown in Figure 4a. These were employed for three reasons. Firstly, this ensured that no auto-generated metal filler would be placed between the interconnect and the distant supply wires which might inadvertently introduce return paths not considered in the simulation. Secondly, it

ensured that conductors at the edge of the array had approximately the same capacitance as if a signal wire were adjacent to it, which would better approximate the case for a typical bus as well as more accurately mirror the capacitance values used in simulation. Lastly, it improved the uniformity of the dielectric interface between layers, which again provided a better correlation between simulation and measurement.




Figure 4a: Lateral fingers adjacent to the conductors




Figure 4b: Overlapping fingers on M2 and M3

In the absence of a M3 plane it is necessary to ensure that M1 and M4 are still completely capacitively isolated so that any noise measured on the victim wire is due solely to inductive coupling. A pure RC model would therefore predict that this node remains stable when the M4 wires switch. To effect this, the lateral fingers on M2 and M3 were overlapped as shown in Figure 4b to efficiently isolate M1 and M4. Note that for the structures involving M2 aggressors, capacitive isolation to M1 could not be completely guaranteed. Finally, a significant amount of decoupling capacitance was placed beneath the distant supply rails to ensure good supply integrity. This also ensured an effective AC short for return currents.

4. EXTRACTION AND SIMULATION

An in-house 2D extraction engine was used to generate the $N \times N$ RL matrices at various frequencies of interest for each system of N conductors. The extraction engine partitioned each conductor into $J \times K$ sub-elements in which uniform current density (uniform voltage distribution and resistivity) was assumed. The area of these elements nearer to the surface, and especially nearer to the corners of the conductor, is smaller than in the center of the conductor. This enabled greater accuracy in the extraction process at high frequencies, when the skin effect becomes relevant and there exists a significant gradient in current density from the edge of the conductor to its center. For most conductors, $J=K=7$ was found to give sufficient accuracy. Simulations also showed that modeling the RL conductor matrix at 5 logarithmic frequencies between 0.1 GHz and 10 GHz (inclusive) gave sufficient accuracy.

For simulation purposes the $N \times N$ RL matrices (with $P=N(N+1)/2$ independent elements due to symmetry about the leading

diagonal) is modeled as a set of P voltage controlled current sources (conductances) whose relationship is defined by the time domain response to a ramp input over the SPICE simulation timestep. This response requires various parameters to be provided for each conductance, and another in-house tool is used to convert the 2D RL matrices (via partial fraction expansion) into these modeling parameters and subsequently generates a SPICE subcircuit with P frequency dependent conductances.

A user defined SPICE model was created which enabled the conductance parameters to be used to generate the output current for each conductor based upon the voltages across all other conductors, according to the ramp response referred to earlier. This is only made possible through the use of an in-house version of SPICE and highlights one of the many advantages of having an internal CAD team. It remains then to incorporate this SPICE subcircuit, together with the SPICE user model, into the conventional simulation procedure.

The 2D extraction engine generates all of its parameters on a per cm basis, however the conductor segments to be modeled for a given length of interconnect are usually significantly shorter than this. For example, a $700\mu\text{m}$ wire might be modeled as four lumped segments of $175\mu\text{m}$. To effect this length transformation one could write a script to downsize all parameters as necessary, however a better option is to use current multipliers. Given that a 1cm wire presents an impedance (matrix) of Z , then a current of $I=V/Z$ is produced by the voltage controlled current source. At a length of $175\mu\text{m}$ ($1/57$ cm) the impedance of the wire is $Z/57$ and hence a current of $I=57(V/Z)$ should be produced. We can therefore maintain the same SPICE subcircuit generated for a length of 1cm and implement shorter (or longer) lengths merely by introducing a current multiplier of magnitude $1\text{cm}/\text{length}$. An example model with 3 conductors is shown in Figure 5.




Figure 5: Example macrobox for a 3-conductor system of $175\mu\text{m}$ length

The discussion so far has addressed the incorporation of the LR frequency dependent matrix but has made no mention of capacitance. As can be observed in Figure 5, self and lateral capacitances are explicitly incorporated into the model in a π formation. Capacitance extraction was performed using another in-house tool and is an electrostatic (frequency independent) phenomenon. Furthermore the capacitance matrix is essentially tri-diagonal, with only self and nearest neighbor capacitances being of relevance. As such it can be explicitly incorporated into the conductor model. The RL matrix cannot be so incorporated because of its frequency dependent nature (and consequent conductance modeling), mutual resistance terms (which preclude the use of a lumped element ladder structure) and complexity of symbol connections even for modest values of N .

Occasionally, a 3D extraction engine was used to verify the validity of certain phenomenon approximated (or ignored) in the 2D model. Examples include the effect of fringing fields in the direction of wire length or the presence of orthogonal buses over a given layer (instead of a plane for capacitance extraction or zero conduction for an inductance/resistance extraction).

For each conductor in the system the source inverter was sized to drive the *maximum* capacitance according to a nominal sizing rule, as would be required in a practical system in which adjacent conductors might switch in opposing directions. For the test chip, which exaggerates inductive effects by switching all wires in the same direction, the driver sees the resulting *minimum* capacitance. Hence, the wires will be overdriven with a low impedance driver and a fast edge rate, which further exacerbates the inductive effects that could occur on chip.

5. ON CHIP MEASUREMENTS

Of particular interest in regards to the measurement of signal noise is the aggressor overshoot above VDD for a positive transition, the overshoot of the victim M1 wire, and the ring oscillator frequency. Measuring this latter component, which gives a measure of the time of flight effects, is trivial since it can merely be buffered from the intermediate inverters at either side of the interconnect ring and driven to the output pads (after appropriate frequency division). From there it can be wire-bonded to the package and driven to the output pins of the chip.

Gauging the overshoots on the metal wires however is not so trivial. One cannot simply connect these wires to an output pad since the additional loading will result in significantly different noise levels at the points of interest, and clearly buffering these signals will attenuate the noise levels.

A twofold solution to this problem was employed. Firstly, small probe points (square patches of M4 surrounded by a supply rail) were attached to each of the conductors of interest at the source and receiver to enable e-beam probing of the signal noise. The introduction of these probe points added negligible error to the simulation model.




Figure 6: Rectifier used to gauge peak noise levels

A second, more cost effective solution was also employed. The circuit of Figure 6 was connected to a few select conductors (adding negligible loading) and is used to rectify the peak noise level into a DC voltage that can be sent directly to the output pads. This voltage provides an indication of the peak noise level on each conductor, but is not absolute due to the voltage loss across the diode and its AC characteristics for the short duty cycle of the noise peak. It provides a *measure* of the relative performance between different conductor configurations as well as an inexpensive means by which to compare the simulation results with the actual results.

6. RESULTS

As mentioned in Section 1, short lengths of interconnect exhibit no noticeable inductive phenomenon and similarly, long lengths of interconnect become resistively damped (although they still exhibit transmission delays). It is important to choose a wire length that is both representative of typical interconnects on chip as well as exhibiting significant inductive noise levels. Figure 7 plots the simulated waveforms at the receiver for ten M4 lines with distant, low resistance return paths. From this plot, a length of 3.5mm was chosen for the interconnect loop.



Figure 7: Plot of signal voltage vs. length

Figures 8 and 9 show the simulation results for a variety of conductor configurations, detailing in particular the overshoot on the ten M4 aggressors and the inductive noise coupled to the M1 victim wire (which was tied to VDD at the halfway point along its length and measured at the source).




Figure 8: Waveforms with and without planes for (a) the M4 aggressors (b) the M1 victim wire

Clearly, if no nearby return paths are provided (Fig. 8a) then a significant overshoot on the M4 aggressor results. In this example, the overshoot is as much as 150% of VDD, and would likely violate the reliability rules for almost any technology. Furthermore, the noise coupled to the M1 victim wire (Fig. 8b) is also significant. The peak deviation of 30% of VDD is already likely to violate allowable noise margins and result in functional errors. Note that in the worst case this noise could add to any other capacitively or inductively coupled noise from aggressors on layers other than M4. Note also that the waveform would be flipped (interchanging overshoots and undershoots) for a negative transition on the input.

When a single plane is inserted below the M4 aggressors, the overshoot on this wire is significantly reduced (to 120% of VDD). The transmission delay is also seen to improve by more than 50%. This is because the plane provides a low resistance, localized (low inductance) return path for the injected currents, which also helps minimize the effects of mutual resistance. Furthermore the magnetic fields barely penetrate through the plane and as such, the noise coupled to the M1 victim wire is effectively zero.

Inserting planes above *and* below every two orthogonal conductor layers (forming a stripline) is beneficial for a number of reasons. Firstly, both planes conduct return currents for both orthogonal

conductor layers, which results in low self and mutual inductances and resistances. Secondly, as evidenced in Fig. 8b, conductors within one stripline can be analyzed independently of conductors in another since negligible field interactions occur between them. This simplifies the analyses necessary for developing a design methodology. Furthermore, a plane provides a very low impedance supply path to the core of the chip from a peripheral pad ring, as was utilized in the Alpha 21264 microprocessor [4]. Although the test chip was designed with just one plane below the M4 conductors, simulations showed that with another plane 2 layers above (hence forming the stripline), the peak M4 overshoot was reduced to 115% of VDD (the M1 noise remained negligible).

One might argue that with the advent of flip-chip technology, which alleviates the need for a wide supply network (since smaller lengths between bumps significantly reduces the IR drop to the logic), then perhaps a better solution to minimizing on-chip noise is to provide explicit return wires interspersed amongst the conductors. The ratio of signal to return wires is termed the SR ratio, and to enable the same number of routing channels as the stripline approach, a minimum SR ratio of 2:1 is required (assuming identical signal and return widths). Figure 9 shows the simulation results for a range of SR ratios from 1:1 to 10:1, as were incorporated onto the test chip.




Figure 9: Waveforms at 1:1, 2:1, 5:1, and 10:1 SR ratios for (a) the M4 aggressors (b) the M1 victim wire

As expected, the aggressor and victim noise levels decrease as the SR ratio decreases (due to the tighter return loops, reduced return resistance, and reduced mutual coupling). At a SR ratio of 10:1, the noise levels (135% of VDD) are only slightly better than for when no returns are provided at all. This would still be an unacceptable solution for most technologies. The results for a SR ratio of 5:1 are comparable to the single plane of Fig.8a for the M4 overshoots, but still worse than for a stripline. Furthermore, the M1 noise in this instance is still significant (15% of VDD).

At a SR ratio of 2:1 the M4 and M1 noise levels are now quite good, with the former improving upon the stripline approach but the latter being worse. Another advantage of this scheme is that as well as providing the same number of routing channels, each aggressor is capacitively coupled on just one side. This halves the variation in data dependent min/max capacitance and simplifies critical path and race analyses. It may appear then that this approach is preferable to a stripline, however one also needs to be aware of the following issues. Firstly, in a practical system, the next lowest layer of parallel routing would most likely be on M2 rather than on M1. This would increase the AV noise on this layer and might even prohibit signal routes on these layers from being analyzed independently, thereby complicating the simulation and verification procedures or exaggerating allowable noise margins. Secondly, the maximum ERF is 3.0 for a 2:1 SR ratio but is just 1.25 for a stripline (assuming planes of equivalent thickness and resistivity to the conductor metals). In both cases the minimum

ERF is 1.0, when adjacent conductors switch in opposition. As such, the variation in effective resistance for the 2:1 SR ratio is significantly higher than for the stripline approach, which further complicates the task of electrical verification. Clearly, a 1:1 SR ratio results in even less M4 and M1 noise, eliminates capacitive cross-talk, and reduces the maximum ERF to 2.0, although the number of available routing channels is now reduced to 75% of the stripline approach.

Table 3 below compares the pin-out measurements from the loop frequency (in MHz) and peak detector (DC voltage from the rectifier) for a variety of conductor configurations. These results are quoted both from simulation and from actual (averaged) measurements of over 20 packaged parts.

Structure	Frequency		M4 noise		M1 noise	
	Sim	Meas	Sim	Meas	Sim	Meas
A	113	112	2.17	2.05	1.79	1.82
B	127	123	1.62	1.47	1.47	1.51
C	121	119	1.91	1.72	1.61	1.62
D	118	118	1.78	1.71	1.63	1.65
E	106	105	1.40	1.47	1.52	1.54
F	109	109	1.39	1.44	1.47	1.50
G	113	112	1.47	1.49	1.56	1.57
H	117	115	1.42	1.46	1.46	1.53
I	117	118	1.59	1.58	1.54	1.60
J	107	108	1.79	1.70	1.75	1.85
K	116	113	1.56	1.46	1.54	1.60
L	108	109	1.68	1.61	1.68	1.76
M	110	110	1.63	1.56	1.66	1.77
N	102	103	1.45	1.49	1.62	1.70
O	98	100	1.46	1.48	1.76	1.83
P	104	104	1.46	1.43	1.50	1.63

Table 3: Comparison of simulated and measured results

The excellent correlation between simulation and measurement implies that the simulation procedure is sufficiently accurate and that the results extracted from the simulations are reliable. E-beam probing of the on-chip signal waveforms is still to be scheduled and is expected to further validate the simulation results.

7. METHODOLOGY DEVELOPMENT

The test chip as presented was designed primarily to verify the accuracy of the simulation procedure, however it also served as a vehicle in which to investigate the effects of various conductor configurations on aggressor overshoots and AV coupling to a capacitively isolated victim wire. Although these analyses are important and useful, they are only a subset of the full range of analyses required to derive a comprehensive methodology for managing on-chip signal integrity issues.

A more comprehensive analysis requires an investigation into AV coupling to victim wires on the same layer as well as on other layers. This effect is significant in the presence of both capacitive and inductive cross-talk, and depending on the return rules assumed (planes and/or SR ratios) one may have to incorporate the switching effects of signals on other layers as well. Another important issue is the effect of self and mutual inductance and resistance, together with capacitance, on signal delay (time of flight and RC delay) as compared to a raw RC

model, for a range of data dependent switching conditions. These effects demand investigation since ideally, circuit designers ought not to be encumbered with the need to incorporate inductance into all of their simulations. The return rules should be sufficient enough that the inductive effects are minimal and can simply be incorporated as an additional margin to both noise and delay. Furthermore, one needs to extract from simulation (and from the resistance matrices) the maximum ERF that designers need to consider for each layer of interconnect.

All of these analyses are still insufficient for developing an exhaustive methodology. One also needs to investigate how these noise and delay values vary with length, conductor width and spacing, and especially lumped and distributed gate loading. Each of these parameters adds another dimension of complexity to the task of managing on chip signal integrity. Without such a comprehensive analysis one risks settling on return rules which may prove to be inadequate, perhaps even catastrophic, for a particular set of unforeseen conditions.

8. CONCLUSIONS

The test chip has enabled the inductive effects of various conductor configurations to be analyzed as well as validating the simulation approach. This involves modeling systems of interconnect as a set of frequency dependent voltage controlled current sources, with each section appropriately and easily scaled to any desired length. The results of the test chip correlate well with the simulation results and highlight the performance tradeoffs between stripline and interstitial approaches to providing dedicated on-chip signal returns. Of particular importance is the need to analyze AV coupling in addition to overshoot and delay analyses, and thus arrive at a comprehensive solution to managing on-chip signal integrity.

9. ACKNOWLEDGEMENTS

The support of Don Priore, Michael Tsuk, and Jim St. Laurent in developing the test bed, extraction, and simulation procedures is greatly appreciated. Discussions with and suggestions from Bill Bowhill, Anil Jain, Stephen Thierauf, Michael Gowan, and Johann Nittman are also acknowledged.

10. REFERENCES

- [1] H. Fair, D. Bailey, "Clocking Design and Analysis for a 600 MHz Alpha Microprocessor", *ISSCC Digest of Technical Papers*, Feb 1998, pp. 398-399.
- [2] P. Gronowski, W. Bowhill, R. Preston, M. Gowan, R. Allmon, "High-Performance Microprocessor Design", *IEEE Journal of Solid-State Circuits*, May 1998, pp. 676-686.
- [3] Y. I. Ismail, E. G. Friedman, J. L. Neves, "Figures of Merit to Characterize the Importance of On-Chip Inductance", *DAC'98*, June 1998, pp. 560-565.
- [4] B. A. Gieseke et al., "A 600MHz Superscalar RISC Microprocessor with Out-Of-Order Execution", *ISSCC Digest of Technical Papers*, Feb 1997, pp. 176-177.