Driver Modeling and Alignment for Worst-Case Delay Noise

Supamas Sirichotiyakul, David Blaauw, Chanhee Oh, Rafi Levy*, Vladimir Zolotov, Jingyan Zuo

Motorola Inc. Austin, TX, *Motorola Semiconductor Israel Ltd. Tel Aviv, Israel

david_blaauw@email.mot.com

Abstract

In this paper, we present a new approach to model the impact of cross-coupling noise on interconnect delay. We introduce a new linear driver model that accurately models the noise pulse induced on a switching signal net due to cross coupling capacitance. The proposed model effectively captures the non-linear behavior of the victim driver gate during the transition and has an average error below 8% whereas the traditional approach using a Thevenin model incurs an average error of 48%. We also discuss the worst case alignment of the aggressor net transitions with respect to the victim net transition, emphasizing the need to maximize not merely the delay of the interconnect alone but the combined delay of the interconnect and receiver gate. We show that the worst case alignment of an aggressor net transition is a function of the receiver gate output loading, victim transition edge rate, and the noise pulse width and height and hence propose a pre-characterization approach to efficiently predict the worst-case alignment. The proposed methods were implemented in an industrial noise analysis tool called ClariNet. Results on industrial designs are presented to demonstrate the effectiveness of our approach.

1 Introduction and Previous Work

Due to process scaling, cross-coupling capacitance has become a dominant portion of the total parasitic interconnect capacitance. As previously observed [1], the interconnect delay of such nets are strongly dependent on whether their neighboring nets are simultaneously switching or not. The net under consideration is referred to as the *victim* net, and the neighboring nets that are capacitively coupled to it are referred to as *aggressor* nets. When the aggressor nets switch, they induce a noise pulse on the victim net. If the victim net is stable when the aggressors switch, the resulting noise pulse can cause a functional failure. This situation is referred to as *functional noise*. If the victim net itself is also switching when the aggressors switch, its delay can either increase or decrease depending on the aggressor and victim switching directions. This is referred to as *delay noise* and is the focus of this paper.

Since the interconnect is connected to non-linear driver and receiver gates, the coupled delay analysis is an inherently non-linear problem. However, non-linear simulation is not practical due to the large size of the interconnect model which can consist of many thousands of elements for a single victim net. To efficiently address the analysis of large designs, linear models of the driver and receiver gates are typically constructed to allow the use of efficient linear simulation and superposition. The driver gate is tradi-

tionally modeled with a Thevenin model consisting of a Thevenin resistance and a voltage source with a ramp transition. The receiver gate loading is modeled with a grounded capacitor. Figure 1(a) shows the linear model for a victim net with two aggressor nets. Using superposition, each of the driver gates is simulated in turn, while other Thevenin voltage sources are shorted. Figure 1(b) shows the simulation model when aggressor driver A is simulated, i.e. switching. A similar model is used to simulate aggressor driver B. Figure 1(c) shows the simulation model used to simulate the victim driver. The voltage waveforms observed at the receiver input from all simulations are added together using superposition to obtain the noisy waveform as shown in Figure 1(d). Typically, the amount of delay noise has been measured as the difference of the 50% Vdd crossing times of this noisy waveform and the noiseless waveform obtained from the simulation in Figure 1(c). Alternatively, the receiver gate can be simulated with the noisy waveform as shown in Figure 1(d) and the delay noise is then measured at the receiver output instead. Using linear driver and receiver models has the advantage that a reduced-order model of the network needs to be created only once with methods such as PRIMA [2], and is then reused in all different driver simulations.
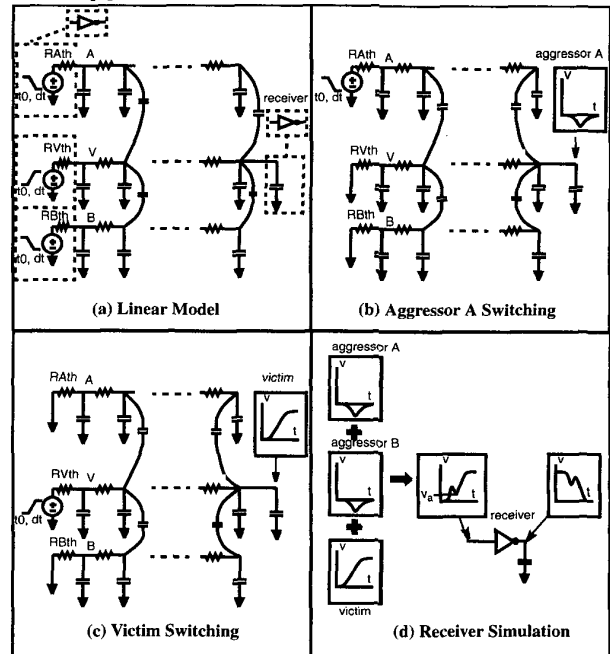

Figure 1. Interconnect analysis using linear simulation and superposition

This use of the Thevenin model is based on the standard approach in the analysis for non-coupled interconnects. Thevenin gate model parameters (t_0, dt, and R_{th}) are a function of the effective load that the driver gate sees. It reflects the fact that the driver is actually a non-linear device, which is approximated with different Thevenin models at different effective gate loading values. The effective

loading of a gate is calculated using so-called *C-effective* iterations [3][4] and captures the resistive shielding of the interconnect. For a particular effective load, the Thevenin model parameters are optimized to obtain a good correspondence with the non-linear driver simulation at the 10%, 50%, and 90% transition times.

The described approach for a coupled interconnect network simulates the victim driver and each aggressor driver in turn. When an aggressor driver is simulated (Figure 1(b)), the victim and other aggressor drivers are modeled with their Thevenin voltage source grounded, i.e. their Thevenin resistances are connected to ground. These grounded resistances, or *holding* resistances, represent the ability of the these drivers to hold their signal lines steady while the simulated aggressor gate injects noise through the coupling capacitances. However, the Thevenin resistance has been calculated to model the aggregate resistance of the driver over an entire transition of a gate whereas the noise from the simulated aggressor is injected for only a short period of time while the victim is switching. Since the small signal conductance of the driver gate varies dramatically during the transition, an accurate holding resistance is a function of the duration of the injected noise and its alignment relative to the gate transition. It is thus clear that the standard Thevenin resistance is not a good approximation to model the grounded drivers in the superposition flow for coupled interconnects. In Figure 2, we show an example of a coupled victim and aggressor network whose noise pulse computed using the Thevenin resistance for victim driver significantly underestimates the actual noise injected on the victim net.
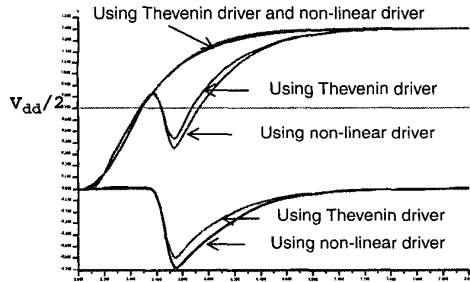

Figure 2. Simulation results using Thevenin model

To address this issue, a modified C-effective calculation approach has been proposed in [5], [6]. It accounts for the additional charge that a switching driver gate sees when other gate drivers are switching simultaneously. In this approach, the Thevenin model parameters are updated using a modified effective loading capacitance that accounts for the charge injected due to the switching aggressor nets. Since this charge is averaged over the entire transition, the exact small signal conductance of the non-linear victim driver gate during the short period that the aggressor switches cannot be accounted for.

We propose a new approach which models the victim driver gate with a different resistance when its voltage source is shorted in the superposition flow (Figure 1(b)). This new resistance model is referred to as the "*transient holding resistance*" - R_{tr}. The transient holding resistance is a function of the noise width, height, and alignment relative to the victim transition. It is computed with a single non-linear simulation of the driver gate using a single effective output load. For efficiency, it can be precharacterized and stored in a table similar to that for the Thevenin model.

Note that a similar problem occurs when the victim driver is switching and the aggressor drivers are shorted in the superposition flow (Figure 1(c)). In this case, the noise pulse injected on the aggressor net by the victim will be underestimated due to the Thevenin resistance used for the aggressor driver. However, the voltage on the aggressor net is not of direct interest to our analysis and has only an indirect effect on the victim net. This explains why the noiseless victim transition using a standard Thevenin model shown in Figure 2 is quite accurate. However, the proposed approach can also be extended to the shorted aggressor driver models to calculate their transient holding resistances if needed.

After all linear models are calculated for the driver gates, the remaining issue is how to align the transition of the aggressor nets relative to the transition of the victim net. The aggressor nets must be aligned within the constraints of the switching timing windows that are calculated during timing analysis [1]. One difficulty is that the timing windows are a function of the added delay due to cross coupling noise, and this added delay is in turn a function of the aggressor timing windows. In [8][9], it was shown that iteratively calculating the timing windows and the added noise delay will converge on the correct solution. In practice, very few iterations are needed for convergence.

The task that we examine in this paper is to determine the switching time within the constraints of the timing windows that produces the worst case victim delay in each iteration. We approach this problem in two steps: First, we determine the worst alignment of the aggressors relative to each other. This will produce a composite noise pulse which is the superposition of all aggressor induced noise pulses. Second, we determine the worst case alignment of the composite noise pulse with respect to the victim transition time. Typically, the objective has been to maximize interconnect delay, which is measured from 50% crossing time of the victim driver output to the 50% crossing time of the victim receiver gate *input*. In [6] it was shown that under reasonable assumptions, this delay is maximized by aligning all aggressor noise pulses such that their peaks occur at the same time. The peak of this composite noise pulse is then aligned at the point where the noiseless victim transition reaches $Vdd/2 + V_n$ for a rising transition, where V_n is the height of the composite noise pulse, as shown in Figure 3.

In timing analysis, however, the true objective is not to maximize only the interconnect delay, but the combined delay of the interconnect and the receiver gate, measured from 50% crossing time of the victim driver output to the 50% crossing time of the victim receiver gate *output*. Figure 3 also shows that aligning the composite noise pulse for the worst interconnect delay may result in a combined interconnect and receiver delay not being increased at
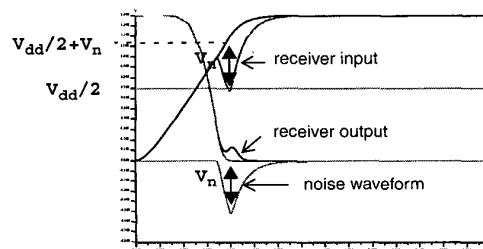

Figure 3. Worst case alignment at receiver input

all. This occurs when the alignment for maximizing the interconnect delay places the aggressor transition too late and the receiver gate has already completed its transition. In this situation the noise pulse at the receiver input is quite large and a correct alignment of the aggressor would have significantly increased the delay at the receiver output. Note that in this case, due to the effective filtering of the receiver gate, the noise pulse at the receiver output is less than 100mV and does not constitute a *functional* noise failure.

It is therefore clear that aligning the aggressor transition based solely on maximizing the interconnect delay is not valid and that the effect of the alignment on the receiver output transition must be considered. When the receiver delay is included in the aggressor alignment objective, the worst case aggressor alignment becomes a function of the receiver gate type, size, P/N ratio, and output load. Furthermore, the receiver gate is highly non-linear, making efficient closed form solutions difficult. In this paper, we first examine the worst case alignment of the aggressor transitions with respect to each other. We show that the worst case alignment does not always occur when all aggressor noise pulses have coincident peaks. In these cases, however, the receiver delay is relatively insensitive to the exact alignment of the aggressor peaks, and we show that using aligned noise peaks introduces only a very small amount of error. Second, we propose an effective pre-characterization approach to calculate the alignment of the composite noise pulse relative to the victim transition time. Although the worst case alignment is a function of the noiseless transition time, the noise pulse height and width, and the receiver gate loading, we can predict the worst case alignment of any possible condition using the precalculated alignment of a small set of conditions. This pre-characterization approach has the advantage that for a particular type of receiver gate, we precalculate the worst case alignment for a small set of conditions, after which the alignment for any instantiation of the gate is obtained easily through table lookup and interpolation.

In a coupled interconnect network, the linear driver models are a function of aggressor alignment and, conversely, the alignment is a function of the noise width and height, which are in turn a function of the linear driver models. Hence, it is impossible to determine one without first determining the other. In the overall approach, we iterate between the linear model calculation and the alignment calculation to reach convergence. The overhead in each iteration is relatively small because the linear model calculation involves only one non-linear simulation of the victim driver circuit and the alignment calculation involves only table lookup and interpolation operations. In practice we find that only one or two iterations are needed.

The remaining of this paper is organized as follows. Section 2 presents the method for calculating the transient holding resistance needed to model the grounded victim driver model in the superposition flow. Section 3 presents the methods for calculating aggressor alignment. Section 4 presents the results of the proposed approach, and Section 5 presents our conclusions.

2 Victim Driver Model

In the proposed superposition flow, the voltage source of the victim driver model is shorted when simulating the noise injected by an aggressor driver as shown in Figure 1(b). The victim driver is then represented only by the Thevenin resistance R_{th}. This model introduces a significant error as it does not represent the gate conductance during the time of the noise injection, as illustrated earlier in

Figure 2. We propose a more accurate model by replacing the standard Thevenin resistance R_{th} with a *transient holding resistance* R_{tr}. We determine this transient holding resistance such that it produces a matching noise waveform with noise injected on the non-linear victim driver. Our approach is outlined as follows. First, we obtain the aggressor noise on the victim net by performing a linear simulation using the standard Thevenin resistance for the shorted victim driver as in the original approach shown in Figure 1(b). From the simulation result, we calculate the associated noise current that is injected into the victim driver output. Then, we simulate the non-linear victim driver switching with an effective load, both with and without this injected noise current. Since the victim driver is switching when the aggressor driver injects noise on the victim, we cannot directly observe the noise injected on the victim line. We can only construct it from the difference of the driver responses with and without injected noise. Thus, we subtract the two driver output waveforms to obtain the noise waveform at the victim driver output. We then calculate a transient holding resistance that yields a matching noise pulse. Each of these steps is explained in more detail below:

1. Using Thevenin models for the victim and aggressor drivers, we simulate one aggressor driver at a time while grounding the victim and all other aggressor models. In each simulation we record the voltage waveform at the victim driver output and then calculate the total noise voltage V_n as the sum of all voltage waveforms.

2. Using the simplified model shown in Figure 4(a), we calculate the current waveform I_n injected into the driver gate as follows:

   $I_n = V_n/R_{th} + C_{load} \cdot (\partial V_n)/(\partial t)$, where R_{th} is the victim driver Thevenin resistance, and C_{load} is the effective load capacitance as calculated with C-effective iterations.


Figure 4. Transient holding resistance (R_{tr}) calculation

3. We perform a non-linear simulation of the victim driver gate with C_{load} at the output to obtain a noiseless transition V_1. We repeat this simulation with the added current source I_n obtained from Step 2 connected at the gate output, and measure the voltage waveform V_2 at the gate output, as shown in Figure 4(b).

4. We calculate the noise voltage response of the non-linear model, V'_n, by subtracting the two non-linear simulation results: $V'_n = V_1 - V_2$.

5. Finally, we construct the equivalent linear model with the transient holding resistance R_{tr} by replacing R_{th} in Figure 4(a) with R_{tr}. We determine the value of R_{tr} such that the area under the resulting noise voltage waveform V''_n matches the area under V'_n. The value of R_{tr} is calculated as follows:

$$I_n = V''_n / R_{tr} + C_{load} \cdot (\partial V''_n) / (\partial t)$$

$$\int I_n dt = \frac{1}{R_{tr}} \int V''_n dt + C_{load} \cdot V''_n \Big|_{t=0}^{t=t_{end}}$$

Since V''_n is a noise waveform which will return to its original value at t_{end}, $V''_n(0) = V''_n(t_{end})$, i.e. $V''_n \Big|_{t=0}^{t=t_{end}} = 0$.

Also, to match the area of V'_n and V''_n, we replace $\int V''_n dt$ with $\int V'_n dt$. Thus:

$$R_{tr} = \frac{\int V'_n dt}{\int I_n dt}$$

where $\int I_n dt$ and $\int V'_n dt$ are obtained from Step 2 and Step 4, respectively.

6. We calculate the noise waveform by performing a linear simulation using R_{tr} in place of the victim driver Thevenin resistance R_{th} in the circuit shown in Figure 1(b). Then we continue with the traditional linear simulation and superposition approach described earlier.

Note that after Step 6, the noise current I_n has changed requiring a recalculation of R_{tr}. However, in practice a single or at most two iterations are necessary. Also, when the alignment of the aggressor transition changes with respect to the victim transition, the non-linear noise waveform will be affected, and R_{tr} must be recalculated. In Figure 5, we show the simulation results when the proposed approach is applied on the circuit producing the waveforms shown in Figure 2. The result shows that the voltage waveforms match closely with the full non-linear simulation results. In this case, the calculated transient holding resistance, R_{tr} is 1463 Ohms, whereas the original thevenin resistance was 1203 Ohms.

Figure 5. Linear noise simulation using R_{tr}

3 Aggressor Alignment for Worst-Case Delay

The interconnect and receiver delay are strongly dependent on how the noise waveforms are aligned with respect to the victim transition. We approach this problem in two steps. First we determine the alignment of the aggressor with respect to each other, forming a composite noise waveform. Then, we align this composite noise waveform with respect to the victim transition. We discuss each of these two issues in more detail below.

3.1 Alignment Among Aggressors

Traditionally, the noise waveforms induced on the victim net are aligned such that their peaks coincide. Such an alignment will produce a composite noise pulse with a maximum pulse height and

minimum noise pulse width. Conversely, shifting the alignment of the individual noise peaks will result in a wider and less high composite noise pulse. When considering only the interconnect delay, a composite noise waveform with aligned aggressor noise pulses will typically result in the maximum delay. However, as we have discussed earlier that considering only the interconnect delay is not meaningful, and the receiver delay must be included. Since the receiver gate acts as a low pass filter, the maximum height composite noise pulse may not always result in the maximum response at the receiver output. Especially when the receiver gate has a large capacitive load, a composite noise pulse with a lower peak voltage and wider width can result in a stronger response at the output.

Figure 6 shows the combined interconnect and receiver delay of a circuit with two aggressor nets similar to Figure 1(a) under varying alignments. In one case, the receiver gate has a small output load, allowing it to pass a high frequency noise pulse relatively well. In this case, the worst aggressor alignment occurs when the noise peaks of both aggressor nets coincide. However, when the same receiver gate has a large output load, it acts more strongly as a low pass filter, and the worst aggressor alignment occurs when their peaks are not aligned and a wider and less high composite noise pulse is presented to the receiver gate.

Figure 6. Delay vs. alignment for 2 aggressors

Having to consider non-aligned aggressor peaks greatly expands the search space for the worst case aggressor alignment and makes the problem more complicated. However, in the cases where the worst case delay occurs with non-aligned aggressor noise peaks, the error introduced by using aligned noise peaks is very small. In Figure 6, for example, the delay difference at the receiver output is only 2.7 ps. In general, the cases where the worst case delay is produced by non-aligned aggressor noise peaks are when the victim transition is relatively fast, the aggressor transition is relatively slow, or the receiver output load is large. In all of these cases, the extra delay is relatively small and insensitive to the alignment. Therefore, we can align all aggressor peaks together without incurring a large error in the delay calculation. In all our simulations, the error introduced by this approximation is less than 5 percent.

3.2 Alignment with Respect to the Victim Transition

After the composite noise pulse has been constructed, its alignment relative to the victim transition must be determined. The worst case alignment is complicated by the non-linear behavior of

723

the receiver gate and by the fact that it is a function of the receiver gate type, gate size, output load, and the composite noise waveform and noiseless transition edge rate at the receiver input. Finding the actual worst case alignment involves performing an expensive search using a large number of non-linear simulations. This is clearly too expensive to perform during timing analysis.

We therefore propose a new pre-characterization method where the worst case noise alignment is stored in a lookup table. Since the number of parameters that influence the worst case alignment is very high, the number of data points needed to build a simple linearly interpolated lookup table would be unacceptable. For instance, if for a particular gate the four dimensions (output load, input noise pulse width / height, and input edge rate) were precharacterized at 10 points each, a total of 10,000 pre-characterization points would be required, which is unacceptable. In our approach, however, we were able to simplify the lookup table, such that only 8 pre-characterization points are required, while maintaining an accuracy within 10%. The dependence of the worst case alignment on all four parameters is discussed in more detail below.

Receiver output load capacitance. To understand the behavior of the receiver gate with respect to its output load, Figure 7(a) shows the total delay (the combined interconnect and receiver delay) as a function of the composite noise pulse alignment for different receiver output load capacitance values. The simulation shows that for small receiver loads, the alignment is very sensitive and even a small shift in alignment can produce a dramatic change in the delay. However, for large output loads, the delay is relatively insensitive to the alignment and a deviation in the alignment results in only a small error in the added delay. In our approach we therefore use the worst case alignment at minimum receiver output load for all loading conditions for the receiver gate. From Figure 7(a), it is clear that this will introduce only a small error for the case where the receiver gate has a large capacitive load.

Victim Edge Rate. The worst-case alignment exhibits a non-linear relationship as a function of the edge rate, if the alignment is measured from the start of the noiseless victim transition. However, when we measure the alignment with respect to the 50% crossing time of the victim transition, the relationship closely approximates a linear function. To illustrate this, Figure 7(b) shows the total delay as a function of the composite noise pulse alignment for different victim transition times with the alignment measured relative to the 50% Vdd crossing time of the victim transition. Since the worst case alignment is nearly linear with respect to the victim transition time, we need to precharacterize a gate for only minimum and maximum victim transition time and linearly interpolate for points in between.
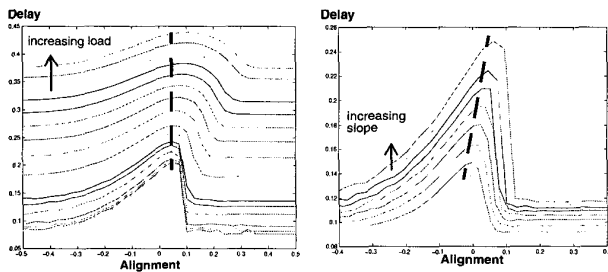
To determine the worst case alignment for different victim slopes and receiver output loads, we therefore need only two pre-characterization points, one at maximum victim slope and one at minimum victim slope. Both pre-characterizations are performed with minimum receiver output load. Figure 9(a) shows the accuracy of this alignment prediction approach for all possible victim slopes and receiver loads for a typical gate. In all cases the error is less than 7%.

Noise height and width. The worst case noise alignment time is not a linear function of the noise pulse height and width, complicating the generation of an efficient table. Instead, we use the receiver input voltage at the time point when the noise pulse reaches its peak, which we refer to as the alignment voltage (voltage V_a in Figure 1(d)). When considering the noise alignment in terms of its alignment voltage, the noise pulse width and height is linearly dependent on the alignment. Figure 8(a) and Figure 8(b) show the total delay as a function of the alignment voltage for varying noise pulse widths and heights, respectively. We precharacterize the alignment voltage at the four conditions corresponding to the minimum and maximum pulse width and pulse height. Note that we can always calculate the alignment time from the alignment voltage and the victim transition time. Figure 9(b) shows the error in the calculated delay using this alignment prediction approach for all possible noise pulse widths and heights. In all cases the error is less than 8%.
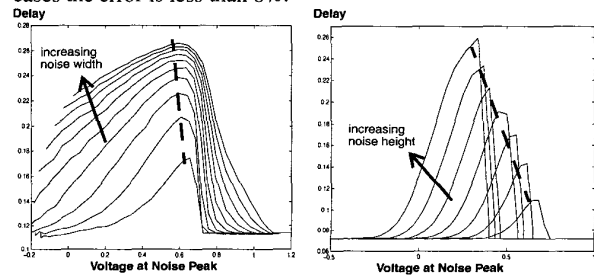

(a) For various pulse widths   (b) For various pulse heights

Figure 8. Delay as a function of alignment voltage

The overall pre-characterization approach uses 8 receiver gate conditions - 2 points in each of the pulse width, pulse height, and victim slope dimensions, all with the minimum receiver output load. For each case, a worst case alignment voltage is recorded in a table. During the actual noise analysis run, the actual alignment is calculated by performing linear interpolations of the alignment voltages in noise width and height dimensions, mapping the resulting alignment voltages to alignment times, and then linearly interpolating the alignment time in the receiver input dimension.


(a) For different receiver loads   (b) For different victim slopes

Figure 7. Delay as a function of noise alignment

Figure 9. Error plot for predicted alignment with
(a) varying receiver load and victim slope
(b) varying noise height and width

4 Results

The proposed algorithms were implemented in an industrial noise analysis tool called ClariNet, which has been used on a number of chip designs [7]. The results for 300 nets from a high performance microprocessor block are shown in Figure 13 and Figure 14. Figure 13 shows the accuracy obtained with the proposed transient holding resistance calculation. The calculated delay using linear simulation with either the original Thevenin resistance or our proposed transient hold resistance are plotted on the Y-axis, and are compared with the delay obtained using Spice simulation of the full non-linear circuit, plotted on the X-axis. The results show that the transient holding resistance has a significantly higher accuracy, with an average error of 7.41% compared to the Thevenin resistance, with a average error of 48.63%. Moreover, the Thevenin resistance incurs a higher error for nets with a larger delay and in all cases underestimates the delay, which is undesirable for noise analysis.
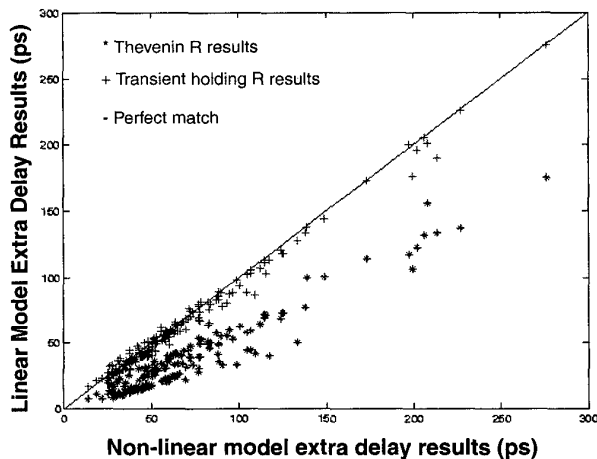

Non-linear model extra delay results (ps)

Figure 13. Linear model results vs. non-linear simulation

In Figure 14, the delay using predicted alignment (plotted on the Y-axis) is compared with the delay using an exhaustive search of the worst case alignment (plotted on the X-axis). The predicted alignment that maximizes the delay at the receiver input using the method presented in [5] is compared with our method which aims to maximize the delay at the receiver gate output. Our proposed method has a significantly higher accuracy, with a worst case error of 15ps, compared to the approach in [5] which incurs a worst case error of 31ps.

5 Conclusions

In this paper, we have presented a new approach to accurately calculate the extra delay due to cross coupled noise injection. We have proposed a new linear model that accurately captures the non-linear behavior of the victim driver gate when noise is injected from aggressor nets. Results show that this model significantly reduces the error in the calculated noise. The model is obtained through a simple simulation of the driver gate and can be precharacterized for gates prior to noise analysis. For determining the alignment of the aggressor noise pulses relative to the victim transition, we have demonstrated the need to include the victim receiver gate delay in the alignment objective function. We have shown that while in some cases non-aligned aggressor noise peaks will result in the worst case delay noise, aligned aggressor noise

peaks can be used with a small error. To determine the alignment of the composite noise pulse relative to the victim transition, we proposed an effective pre-characterization method that requires only 8 pre-characterization points and allows for accurate calculation of the extra delay. Finally, results were shown on industrial circuits demonstrating that the proposed methods significantly increase the accuracy of the analysis.

6 References

[1] K. L. Shepard, V. Narayanan, P. C. Elemendorf and G. Zheng, "Global Harmony: Coupled noise analysis for full-chip RC interconnect networks," Proc. Intl. Conf. Computer-Aided Design, pp. 139-146, 1997.

[2] A. Odabasioglu, M. Celik and L. T. Pileggi, "PRIMA: Passive reduced-order interconnect macromodeling algorithm," Proc. Intl. Conf. Computer-Aided Design, pp. 58-65, 1997.

[3] F. Dartu, N. Menezes, and L. T. Pileggi, "Performance Computation for Precharacterized CMOS Gates with RC Loads," IEEE Transaction on Computer-Aided Design of Integrated Circuits and Systems, Vol.15, No. 5, pp. 544-553, May 1996

[4] J. Qian, S. Pullela and L. T. Pillage, "Modeling the effective capacitance for the RC interconnect of CMOS gates," IEEE Trans. Computer-Aided Design, pp. 1526-1555, December 1994.

[5] F. Dartu, L. T. Pileggi, "Calculating Worst-Case Gate Delays Due to Dominant Capacitance Coupling," Proc. DAC, pp. 46-51, June 1997.

[6] P.D. Gross, R. Arunachalam, K. Rajagopal, L.T. Pileggi, "Determination of worst-case aggressor alignment for delay calculation," Proc. ICCAD, pp. 212-219, November 1998.

[7] R. Levy, D. Blaauw, G. Braca, A. Dasgupta, A. Grinshpon, C. Oh, B. Orshav, S. Sirichotiyakul, V. Zolotov, "Clarinet: A noise analysis tool for deep submicron design," Proc. DAC, pp. 233-238, June 2000.

[8] S. Sapatnekar, "Capturing the Effect of Crosstalk on Delay," Proc. VLSI Design 2000, pp. 364-369, January 2000.

[9] R. Arunachalam, K. Rajagopal, L. T. Pileggi, "TACO: timing analysis with coupling" Proc. Design Automation Conference, pp. 266-269, June 2000.
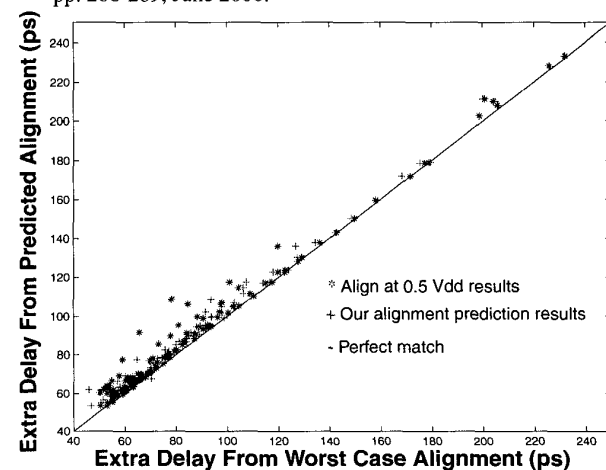
Extra Delay From Worst Case Alignment (ps)

Figure 14. Alignment prediction results