Active Shields: A New Approach to Shielding Global Wires

Himanshu Kaul
University of Michigan, Ann Arbor
hkaul@engin.umich.edu

Dennis Sylvester
University of Michigan, Ann Arbor
dennis@eecs.umich.edu

David Blaauw
University of Michigan, Ann Arbor
blaauw@umich.edu

ABSTRACT
A new shielding scheme, active shielding, is proposed for reducing delays on interconnects. As opposed to conventional (passive) shielding, the active shielding approach helps to speed up signal propagation on a wire by ensuring in-phase switching of adjacent nets. Results show that the active shielding scheme improves performance by up to 16% compared to passive shields and up to 29% compared to unshielded wires. When signal slopes at the end of the line are compared, savings of up to 38% and 27% can be achieved when compared to passive shields and unshielded wires, respectively.

Categories and Subject Descriptors
B.7.2 [**Integrated Circuits**]: Design Aids – *Layout, Placement and Routing.*

General Terms
Design

1. INTRODUCTION
With increasing die sizes and shrinking wire dimensions, wires are becoming longer and more resistive and at the same time clock frequencies are rising. Robust on-chip global signaling in the face of heightened coupling capacitance is beginning to place fundamental limits on global clock frequencies. The growing coupling capacitance increases the amount of functional noise introduced on a wire as well as enhances the sensitivity of wire delays to aggressor switching. In the worst-case scenario, when the aggressor is switching at the same time as the victim but in the opposite direction, the victim must charge an effective coupling capacitance twice that of the nominal one (according to [1], a factor of 3 results from the absolute worst case scenario). In a recent 1 GHz commercial microprocessor, designers found a 50 MHz reduction in achievable clock frequency due to coupling capacitance effects [2]. To alleviate the problem of noise injected by aggressors, shields can be placed on either or both sides [3]. These are ground or power (V_{dd}) lines placed between two wires to prevent direct coupling between them. Other methods include increasing wire spacing to reduce the coupling capacitance or increasing wire widths to reduce the ratio of coupling to ground capacitance. Buffer insertion algorithms already exist to reduce RC delays and noise on long lines [4,5].

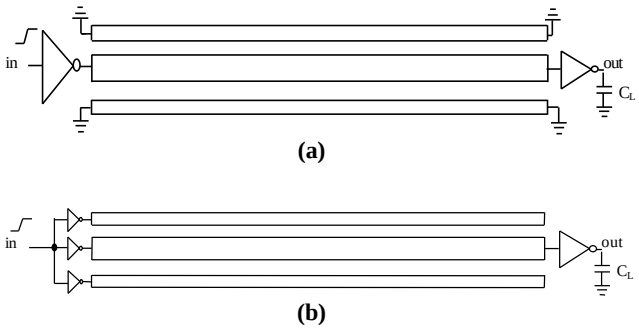

Figure 1. a) Passive shields b) Active shields

The shielding methodology used today is passive in that shield wires are tied statically to V_{dd} or ground. A more useful approach would guarantee a best-case switching scenario for a wire. The concept of active shielding uses shields on either side of the wire that help to speed up signal propagation through the Miller effect. The Miller effect states that the effective coupling capacitance between two nodes is zero if the transitions at the two nodes occur at the same time and in the same direction (a best case scenario can result in a factor of –1). Figures 1a and 1b illustrate the concept of active shielding. Allowing the shields on either side of a wire to switch in the same direction helps to reduce the total line capacitance. This approach is scalable since the increasing line resistance can be offset by the increasing coupling capacitance. We demonstrate that using active shields as proposed in this work results in better performance when compared to other methods like passive shielding and wire spacing/sizing under the same area constraint and capacitive load on the previous stage. Comparisons with buffer insertion are not made since this approach is not meant to replace the buffer insertion methodology but to complement it. For very long lines buffers would still be required to meet performance requirements.

The rest of this paper is structured as follows. Section 2 describes the active shielding approach and develops a simplified analytical model of the theory behind it. Section 3 describes the simulation setup used to compare active shielding to other approaches, and provides results and optimization approaches to obtain the maximum gains from active shielding. Section 4 details the limitations of this work and future work needed to overcome these limitations. Finally, Section 5 provides conclusions.

GLSVLSI'02, April 18-19, 2002, New York, New York, USA.

This work was supported by the MARCO/DARPA Gigascale Silicon Research Center and by Semiconductor Research Corporation (SRC) contract 959.001.

2. ACTIVE SHIELDING

The active shielding approach uses the effective Miller capacitance to reduce the total capacitance on a line. If the transitions on wires adjacent to the wire of interest can be ensured to switch simultaneously and in-phase, the effective coupling capacitance on the middle wire is reduced. This will result in smaller delays for the wire of interest. At the same time the neighboring wires act as shields when the line is quiet. To ensure in-phase switching of the side wires, they must be driven by the same input signal as the middle wire. This approach can be used to speed up signal propagation in two cases – a wire with shields on its sides and a wire wide enough to be split up into three wires while maintaining the same footprint (to avoid area penalties). In the case of the wide (unshielded) wire, the total ground capacitance of the middle wire is reduced (and the effective coupling capacitance is reduced through the Miller effect) while the resistance of the wire increases. The decrease in effective capacitance on the line must overcompensate for the increase in resistance due to the splitting of the wire for this approach to be practical. The unshielded wire will be referred to as the fat wire henceforth. When a fat wire is converted to an actively shielded one, the noise immunity will improve since any aggressors that were previously coupled to the fat wire will no longer be coupled directly to the middle wire.

Since the input signal now must drive three inverters instead of one (as in the two reference cases) we must consider the capacitive load presented to the prior stage. The approach we follow when comparing active shielding to passive shielding and unshielded fat wires is to set the total device width of all three new drivers to that of the original single driver.

In order to obtain analytical insight into the tradeoffs involved in keeping the capacitive load on the previous stage the same, we used a simple model (based on work in [1]) to compare the delays of the passively shielded and actively shielded configurations. The model assumes there are no coupled aggressors. Figures 2a and 2b show the equivalent circuits used for modeling the delays using passive and active shields (the load capacitance C_L is not included in the model). The labeled wire parasitics are denoted as per unit length. The equivalent resistance of an inverter with NMOS width of 1μm is R_0. Thus, if the driver size in the passively shielded configuration is W, then the driver resistance is

$$R_{drv}=R_0/W. \quad (1)$$

The resistances of the drivers for the side and middle wires are

$$R_{drv\_side}=R_0/W_{side} \quad (2)$$

$$R_{drv\_mid}=R_0/W_{mid} \quad (3)$$

The constraint that is used when converting from passive to active shields is $W=W_{mid}+2*W_{side}$.

Delay due to passive shields (for a wire of unit length) is given by

$$T_{passive}= 0.693R_{drv}(C_{gmid}+2*C_c) + 0.378R_{wmid}*(C_{gmid} + 2*C_c) \quad (4)$$

Delay due to active shields is given by

$$T_{active}=0.693R_{drv\_mid}(C_{gmid} + 2*k*C_c) + 0.378R_{wmid}*(C_{gmid} + 2*k*C_c), \quad (5)$$

$$k = 1 - t_{rmid}/t_{rside} \quad (6)$$

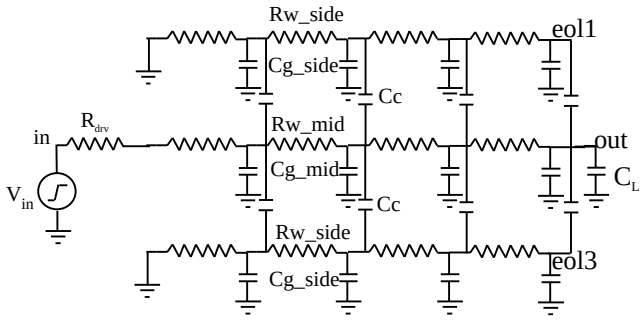

Figure 2a. Passive shields

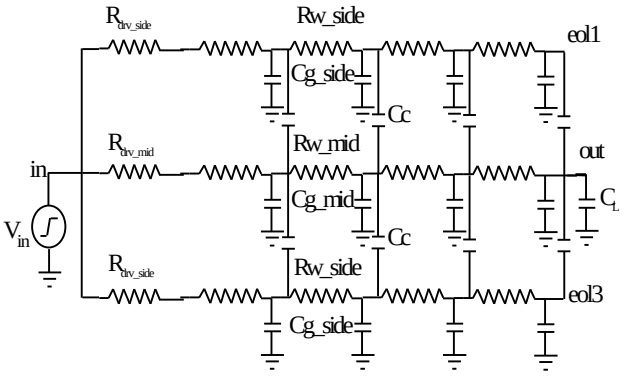

Figure 2b. Active shields

and

$$t_{rmid} = 2.2R_{drv\_mid}(C_{gmid} + 2*k_{mid}*C_c) + 0.9R_{wmid}* (C_{gmid} + 2*k_{mid}*C_c) \quad (7)$$

$$t_{rside} = 2.2R_{drv\_side}(C_{gside}+k_{side}*C_c)+0.9R_{wside}*(C_{gside}+k_{side}*C_c) \quad (8)$$

Obtaining the rise/fall times on the middle and side wires (t_{rmid}, t_{rside}) involves iterations since these times depend on the effective coupling capacitance during switching and the effective coupling capacitance depends on the rise/fall times.

Since the k term is less than one, the effective coupling capacitance on the middle line becomes smaller. To minimize this capacitance, the drivers of the side lines should be strengthened to reduce their slew rates. In this case, the effective resistance of the middle driver increases (since the total W is fixed), creating an inherent tradeoff in designing for active shields. Thus, R_{drv_mid} in (7) increases while the C_c term decreases as driving power is shifted to the side wires. We expect that there is a particular distribution of driver sizes for which the delay of the middle wire can be optimized.

This optimal point depends heavily on the values of the ground and coupling capacitances. As the C_c term increases, it is more favorable to shift more of the driving capability onto the side wires. Figure 3 shows the delays using active shields normalized to that with passive shields. The total driver size (W/L of NMOS) was fixed at 200 while the size of the driver for the middle line (in the actively shielded case) was swept. The model does not yield realistic trends for driver sizes less than half the total driver size but it does point to the existence of an optimal delay point. The figure shows that as the coupling capacitance rises (by decreasing

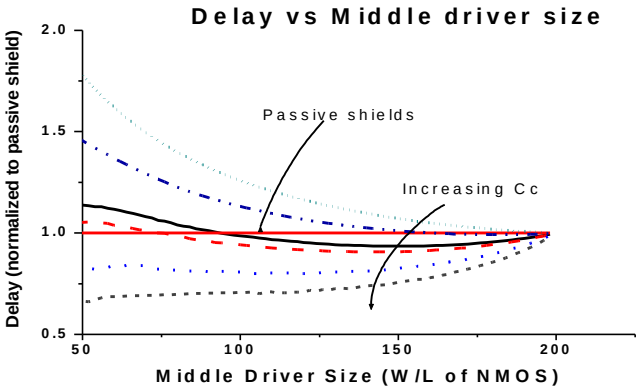

Figure 3. Delay with active shields as driver distribution is changed for different coupling capacitances

the spacing or increasing the thickness) the delay gains compared to the passively shielded case increase. There are cases where the effective decrease in capacitance does not offset the decrease in driving power of the middle wire. In such cases the actively shielded configuration is always slower than the passively shielded one. To get the maximum gains out of active shielding, the coupling capacitance must be maximized (by minimizing the spacing between the side and middle wires).

3. SIMULATION RESULTS

3.1 Simulation Setup

Typically fat wires are used to reduce signal delays by limiting the line resistance. For our analysis, the proposed and alternative approaches are compared to the fat wire in terms of the achievable delay. In all approaches, the footprint of the signal route is kept the same so as not to incur any area penalty. The total capacitive load on the previous stage (which drives *in2*) is kept the same by maintaining the total device width of the driver(s) used for transmitting the signal. The only area overhead incurred with active shielding is due to the splitting of one large driver into three smaller drivers.

Figures 4-6 show the simulation setups for three cases – fat wire, active shielding, and passive shielding. The signal is applied at *in2*, while *in1* and *in3* provide switching activity for the aggressors (*in1* and *in3* switch simultaneously in the opposite direction of *in2*).

Another configuration we considered was the case in which the wire of interest ('2' in Figure 4a) is reduced in width while the spacing to the aggressors is increased by the same amount. This represents the wire sizing/spacing methodology to reduce delays and injected noise on wires, under an area constraint. This leads to less coupling capacitance but at the same time increases the ground capacitance. The resistance on the wire increases rapidly but, unlike the active shielding case, the entire current drive can be used for one wire. In contrast the active shielding approach aims to reduce the ground capacitance by increasing the coupling capacitance (which can be effectively reduced during switching).

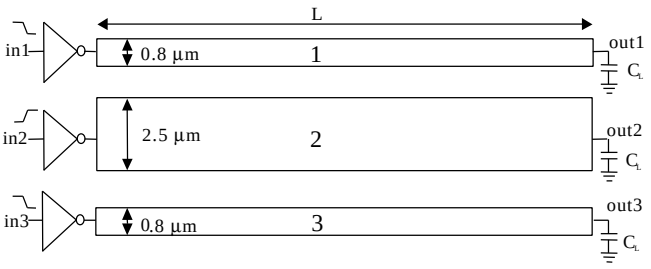

Figure 4a. Setup for fat wire scheme.

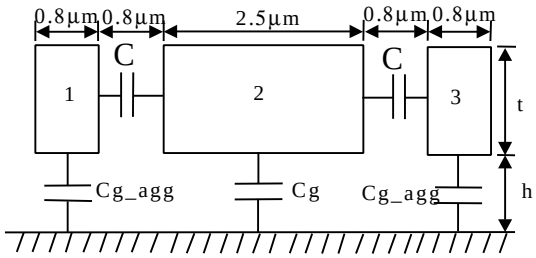

Figure 4b. Cross-section of interconnect structure for fat wire scheme.

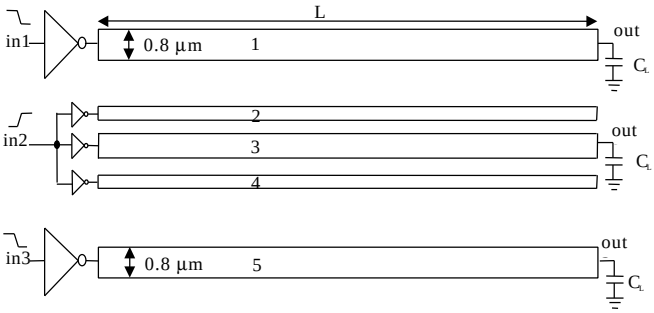

Figure 5a. Setup for actively shielded wire.

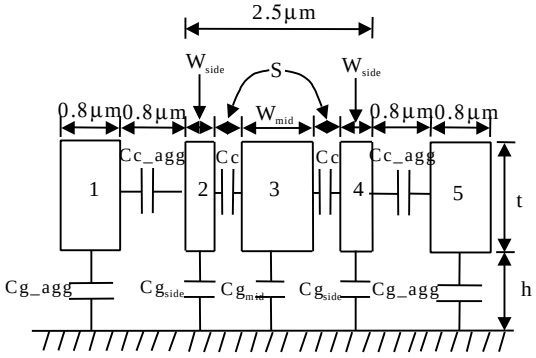

Figure 5b. Cross section of interconnect structure for actively and passively shielded wire.

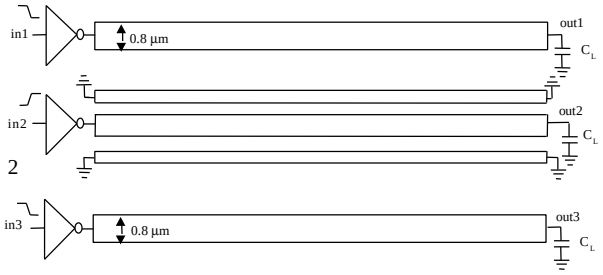

Figure 6. Setup for passively shielded wire.

For the active shielding approach, the delay depends on C_c, C_g, and the resistance of each of the three wires (all of which depend on the wire widths) along with the optimal driver size distribution among the three wires used for signal transmission. Optimizing the C_c, C_g, and line resistances yields optimal widths of the three wires under a constant area constraint. This is further constrained by the minimum width and spacing rules for a particular technology. The simulations considered different technologies (denoted by different wire thicknesses t) with the inter-level dielectric (ILD) thickness (h) fixed at 0.55 times the wire thickness for each technology. The width of the footprint of the wire was kept constant at 2.5um for each technology and the aspect ratio was fixed at 2.4 (which dictated the minimum width and spacing). For each technology five different configurations in which the wire could be split were used to compare the passively and actively shielded wires. The spacing between the wires was always kept minimum to maximize C_c and minimize ground capacitance which is not affected through active shielding. For each wire sizing configuration, the optimal delay for the actively shielded case was obtained by sweeping the driver size distributions. The line lengths simulated were 7.5mm and 3.75mm. For the 7.5mm (3.75mm) line length, the driver size being driven by node *in2* was 200X (125X) and the aggressor driver sizes were 200X (125X) each. Delays were measured from *in2* to *out2* and slopes were measured as 10-90% delays at node *out2*. The load capacitance (C_L) corresponds to the input capacitance of a 25X inverter. All capacitance values in this paper were extracted using a 2-D field solver. To calculate power, a switching activity of 0.5 for an 800MHz operating frequency was applied at node *in2* and the current drawn from the supplies for driving the signal i*n2* was measured.

3.2 Analysis of Results

For the active and passive shielding approaches the delay and slopes were normalized to that of the fat wire for each technology (represented by a wire thickness). The results presented show the delays and slopes obtained with optimized wire and driver sizing configurations (in the case of active shielding). Typical waveforms for both active and passive shielding are shown in Figure 7. Active shielding clearly demonstrates superior delay and slew rate characteristics, which is the strength of the active shielding approach. Further, Figures 8-11 show that the optimal delay using active shields is always better than that of the fat wire and the passive shields. The absolute numbers for delays of passive and active shields with or without switching aggressors are very similar. However, for the fat wire they can vary as much as 22% (delay) and 13% (slope) from the case when the

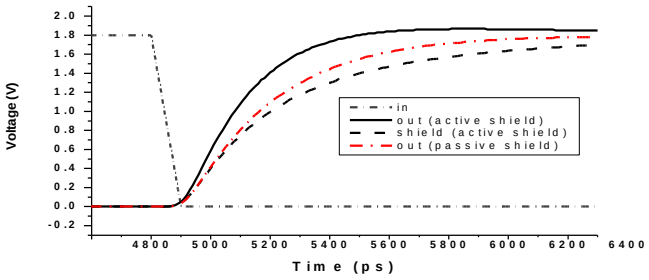

Figure 7. Voltage waveforms at the end of the wire for active and passive shields. The signal on the active shields is also shown.

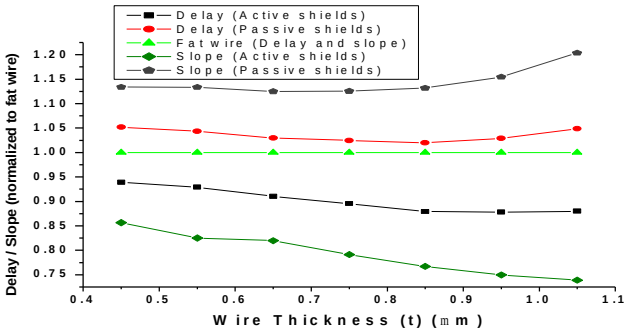

Figure 8. Optimal delay/slope vs. wire thickness for copper wire of length 7.5mm and aggressors switching.

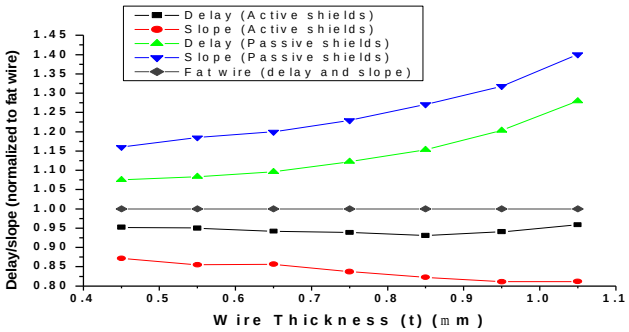

Figure 9. Optimal delay/slope vs. wire thickness for copper wire of length 7.5mm and aggressors not switching.

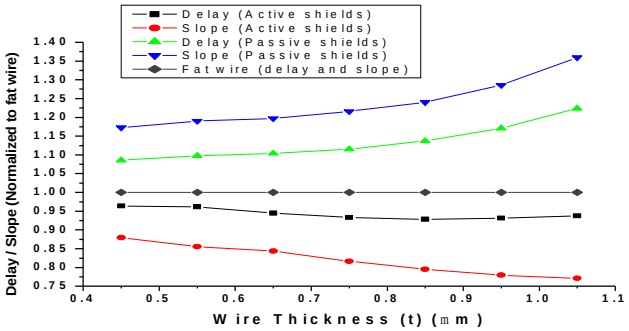

Figure 10. Optimal delay/slope vs. wire thickness for aluminum wire of length 7.5mm and aggressors switching.

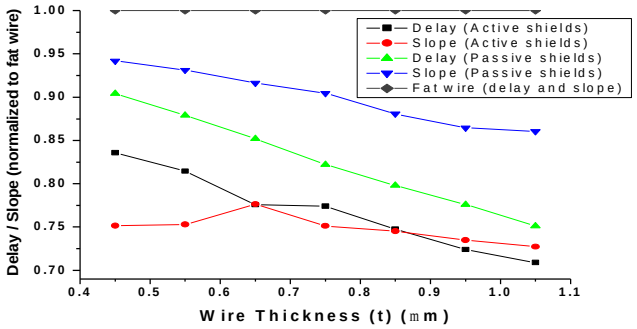

Figure 11. Optimal delay/slope vs. wire thickness for copper wire of length 3.75mm and aggressors switching.

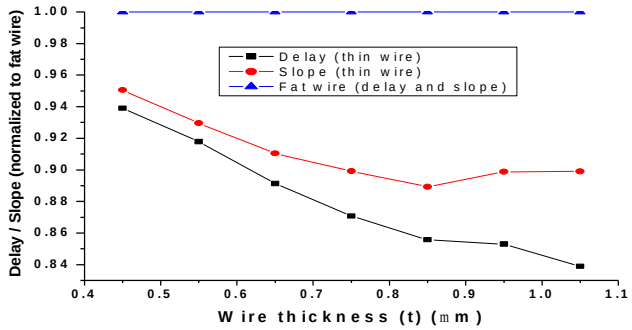

Figure 12. Delay/slope vs. wire thickness for narrowed middle wire (copper) of length 7.5mm and aggressors switching.

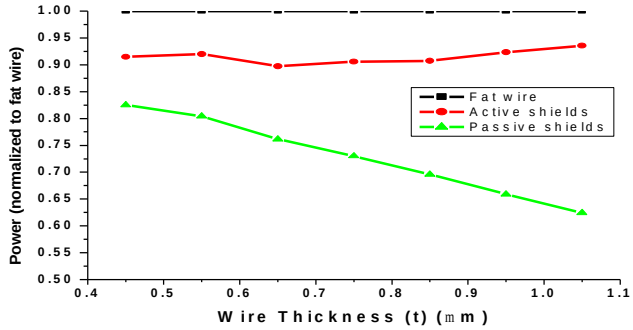

Figure 13. Power consumption for a 7.5mm copper wire with aggressors switching.

aggressors are not switching to the case of aggressors simultaneously switching in the opposite direction. Therefore, gains (when compared to fat wires) are reduced when there is no worst-case switching on the aggressors especially as wires become less resistive (greater thickness). For the same reason passive shields show worsened results as wires become less resistive when there is no switching on the aggressors. All the results indicate that as the total resistance of the wire increases, the gains through active shielding reduce. This is demonstrated by the reduction in gains for each setup as the wire thickness is reduced. For the same reason, the 7.5mm line with aluminum shows less gains than the copper 7.5mm line and the 3.75mm long copper line shows more gains than the 7.5mm copper line. In current 0.18 and 0.13µm processes, repeaters are typically placed every 3-5mm to reduce delay – these distances maximize the usefulness of active shielding.

For the case in which the fat wire is reduced in width to increase spacing from the aggressors, simulations (Figure 12) show that the delay gain is the same as active shielding but the improvements in slopes are much less than with active shielding. The noise analysis shows that the noise immunity in this case is comparable to that of active shields but not as good as passive shields.

Figure 13 shows the power consumption for the configurations in which the delay was optimized. The power consumption was always smallest for the passively shielded case. This is due to the fact that in the actively shielded case extra driving power is consumed by the outer wires in fighting the aggressors. The fat wire always consumes the most power since it has the largest capacitance associated with it and must also fight the aggressors.

A noise analysis confirmed that the active shielding results in enhanced noise immunity compared to the fat wire. However, active shields are not as good at screening functional noise as passive shields. The drivers on the side wires present a highly resistive path to the ground and result in degraded shielding properties.

3.3 Driver and Wire Sizing Issues

Converting a fat wire or a wire with passive shields on its sides to one with active shields while maintaining the same footprint and capacitive load on the previous stage involves two optimizations – wire widths and spacings and distribution of the total driver size over the three wires. As was shown in Figure 3, maximizing the coupling capacitance results in smaller delays (due to the increase in the useful Miller effect). So, the spacing between the side and middle wires should be kept minimum (as dictated by the technology constraints.) The simulation data indicates that for long lines, for which the interconnect delay becomes comparable to the gate delay, the optimal solution is obtained with fatter middle wires since the useful coupling effect is not able to overcome the delay due to line resistance. The middle wire is as wide as possible in nearly all the cases simulated (this leads to minimum width side lines). In a few cases when it is not as wide as possible, the difference in delays between when it is widest and the optimal width is less than 6%. Thus, optimal (or near optimal) delays can be obtained by keeping the middle wire as wide as possible. A figure of merit would be useful in determining the potential gains from different wire sizings.

The optimal driver size distribution can be obtained with the delay model described in Section 2. Though the delay model is not accurate in terms of the delay gains, the trends in terms of optimal driver sizing are similar to the simulations. Figure14 shows the delay gains compared to passive shields for different technologies using both the analytical model and SPICE simulations. There is a close match between the predicted optimal driver sizes from the model and SPICE simulations. Table 1 shows the optimal size of the middle driver obtained through both methods for the setup used in Figure 5a. Though the model does not consider aggressors, the optimal driver size distribution should be the same independent of the presence of aggressors. This trend is shown in the simulations as well. The right-most column indicates the maximum difference in delays as a percentage of optimal delay (as

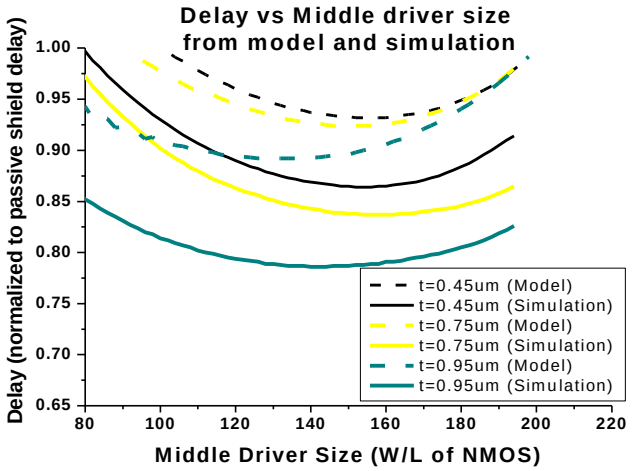

Figure 14. Comparison of delay obtained from model and simulation with active shields for various driver sizes.

obtained through simulations) if the optimal driver size predicted by the model is used. The differences are less than 1%. Once the optimal wire sizing is decided upon, obtaining the optimal driver size distribution is a relatively simple task with this fairly accurate model.

A major discrepancy between the model and the simulation is that the delays predicted by the model converge to the delays using passive shields (when the drivers of the side wires become very weak) whereas in the simulations they do not. The addition of a weak driver to a shield wire results in a highly resistive path to the ground through the coupling capacitance. As a result, less current is injected by the middle wire into the coupling capacitances and the delay improves even when the useful Miller effect is negligible. An improvement in delay (though not by the Miller effect) can be achieved just by adding very weak inverters onto the shield wires.

4. LIMITATIONS AND FUTURE WORK

The biggest limitation of this work is that inductive effects are not considered. The simulated cases involved long lines in which the resistance may dampen out much of the inductive effects, maintaining the results validity. However, inductance may have a large impact on the feasibility of active shields for shorter and much fatter wires, which should be considered [6]. The wire sizing problem needs to be addressed to find optimal wire sizes without resorting to simulations. Future work will consist of including inductive effects into the analysis and developing a figure of merit for analytically determining the optimal wire sizing.

Table 1. Comparison of optimal driver sizes for active shields obtained by model and SPICE.

Wire Thickness (mm)	Optimal middle driver size(s) (W/L of NMOS)		Maximum difference in delays (%)
	Model	Simulation	
0.45	152 – 160	152 – 156	0.19
0.55	136 – 142	142 – 146	0.25
0.65	156 – 160	160 – 168	0.30
0.75	148 – 150	154 – 162	0.24
0.85	148 – 150	148 – 156	0.00
0.95	126 – 138	140 – 144	0.45
1.05	108	130 – 134	0.95

5. CONCLUSIONS

The concept of active shielding was proposed. Simulation results show that converting a fat wire to an actively shielded wire results in lower delays, faster signal slopes, lower power consumption, and better noise immunity. Actively shielded wires were shown to have better performance in terms of delay and signal slopes than passively shielded wires at the expense of higher power consumption and slightly degraded noise properties. A simple yet fairly accurate model was developed to obtain optimal driver size distribution for actively shielded wires but the optimal wire sizing issue needs to be resolved. Future work should consider on-chip inductance.

REFERENCES

[1] A.B. Kahng, S. Muddu, and E. Sarto, "On switch factor based analysis of coupled RC interconnects," *Proc. ACM/IEEE Design Automation Conference*, pp. 79-84, 2000.

[2] P.K. Green, "A GHz IA-32 architecture microprocessor implemented on 0.18? m technology with aluminum interconnect," *Proc. International Solid-State Circuits Conference*, pp. 98-99, 2000.

[3] A. Vittal and M. Marek-Sadowska, "Crosstalk reduction for VLSI," *IEEE Trans. on Computer-Aided Design*, vol. 16, pp. 290-298, 1997.

[4] L.P. P.P. van Ginneken, "Buffer placement in distributed RC-tree networks for minimal Elmore delay", *IEEE International Symposium on Circuits and Systems*, vol.2, pp. 865 –868, 1990.

[5] C.J. Alpert, A. Devgan and S.T. Quay, "Buffer Insertion for Noise and Delay Optimization", *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 18, pp. 1633-1645, Nov. 1999.

[6] Y. Massoud, S. Majors, T. Bustami, and J. White, "Layout techniques for minimizing on-chip interconnect self-inductance," *Proc. ACM/IEEE Design Automation Conference*, pp. 566-571, 1998.