Active Shielding of RLC Global Interconnects

Himanshu Kaul
University of Michigan
hkaul@umich.edu

Dennis Sylvester
University of Michigan
dennis@eecs.umich.edu

David Blaauw
University of Michigan
blaauw@umich.edu

ABSTRACT

In this paper we apply the concept of active shielding to inductive global interconnections. Active shields, which are shield wires that are switched at the same time as the signal wire, were initially developed to speed global signal propagation in RC dominated lines by ensuring in-phase switching of adjacent nets. This work further investigates this concept by examining RLC wiring. We find a distinct line width crossover point at which in-phase switching of neighbors no longer offers benefits and where the increased inductive behavior introduces substantial ringing. We propose the use of out-of-phase active shielding for such wide inductive lines. This technique is shown to significantly reduce ringing behavior (up to 4.5X) and offer better slopes (up to ~40% reduction) and signal propagation delays, all of which are shown in the context of a clock net optimization.

Categories and Subject Descriptors

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

General Terms: Design.

1. INTRODUCTION

Cross-coupling capacitance and inductance are becoming more prominent as technologies scale. Both can cause noise problems as well as degrade performance. To address these problems, shields (additional GND or VDD wires) are typically inserted between wires. When inserted between two wires, almost all the capacitive coupling noise injected between wires is eliminated by the shield wire. It also reduces worst-case delay since the cross-coupling capacitance that the victim must switch is always the nominal Cc (instead of 2*Cc when the aggressor switches simultaneously in the opposite direction). Shields also help to reduce the self-inductance of a wire by providing a current return path that is very close to the signal wire. This also reduces the inductive coupling between wires on either side of a shield. The number of shield wires has increased considerably to cope with signal integrity problems in the face of increasing clock frequencies and signal slew rates. According to [1], the recommended ratio of signal to shield wires has diminished to 2:1 for high-performance processor designs.

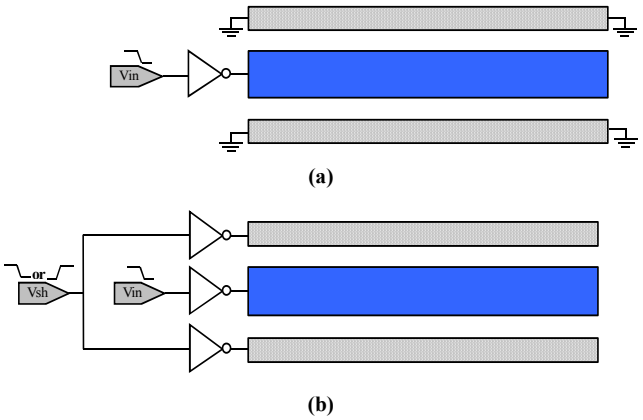

Figure 1. (a) Signal wire with passive shields (b) Signal wire with active shields. Active shields switch either in the same phase or in the opposite phase of the signal.

Since the effective capacitance and inductance of a wire is determined not just by the geometry but also by the switching of neighboring wires, the shields can be switched to further reduce delay and/or loop inductance of critical wires. These *active shields* can be used to change the effective coupling capacitance and loop inductance of a wire through switching rather than layout optimizations. For thin wires, when the resistance dampens out almost all the inductive effects, the delay and slope of the wire can be improved by switching the shields in the same direction as the signal wire. This is achieved by reducing the effective coupling capacitance compared to a passive, non-switching, shield design. We previously explored this approach in [2]. A similar approach was described in [3], where the active shields are called "booster wires". However, [3] does not consider the delay penalty due to added input capacitance of the drivers for the "boosted wires" on the previous stage. In [2] we showed that even under strict input capacitance and area constraints, actively shielded wires enhance performance (delay and slope) for various technologies.

In [2] and [3] the performance gains from switching neighboring wires in the same direction as the signal wire are examined, assuming a distributed RC model of the network. However, when inductance is included in the analysis, switching the shields in the same direction worsens the current return paths and results in a more inductive wire behavior. The increased inductive nature of the signal line can be tolerated for lines that are RC dominated but for wider wires, this increase in inductance leads to an

This work was supported in part by the MARCO/DARPA Gigascale Silicon Research Center (http://www.gigascale.org). Their support is gratefully acknowledged.

TAU'02, December 2–3, 2002, Monterey, California, USA.

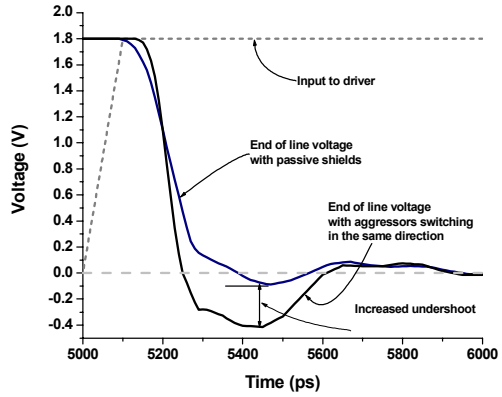

Figure 2. Voltage waveforms at the end of a RLC line when aggressors are switching in the same direction and when passive shields are used.

increase in over/under shoots (shown in figure 2) which directly affects device oxide reliability as well as increased inductive cross-talk. Also, as wires get wider the gains from in-phase switching of neighbors decreases as the ratio of coupling to ground capacitance (Cc/Cg) decreases. Though this limitation can be overcome by placing wires at the upper or lower layer in parallel (and in-phase) with the signal wire, this is not a practical solution with standard design methodologies dictating that wires in adjacent layers run in orthogonal directions. For wires where inductance cannot be neglected, it would be more favorable to switch the shields in the opposite direction of the signal line. The active shields then serve as better current return paths than passive shields and result in lower ringing.

In this paper we investigate the performance gains through active shielding for RLC wires while taking into account the ringing at the end of the line. The next section analyzes a simple setup to obtain the trade-offs of delay/slope vs ringing for in-phase and opposite-phase active shielding in various impedance domains. Using the insight from section 2, we employ in-phase and opposite-phase active shielding for the wire widths for which they are expected to show gains in performance and/or signal integrity over passive shields in sections 3 and 4. Section 3 contains results from simulations of RC-dominated wires with in-phase active shielding. Since our work in [2] presents a detailed analysis of in-phase active shielding, section 3 is restricted to showing that the method works even when inductance is taken into account for RC-dominated wires. We also present a buffering topology for in-phase active shielding and show that gains over passive shielding increase with increasing number of buffered stages. Section 4 describes detailed results for opposite-phase active shielding used for tuning clock nets to reduce ringing, skew and transition time. Section 5 concludes the paper with a summary of the results of our work.

2. CONCEPT

The direction in which the active shields should switch for reducing delays and transition times on the signal wires depends on the nature of capacitive or inductive coupling between the wires. Figure 3 shows that for a RC-dominated line (where the capacitive coupling is stronger) the noise spike generated on the line is in the same direction that the aggressor is switching. In this

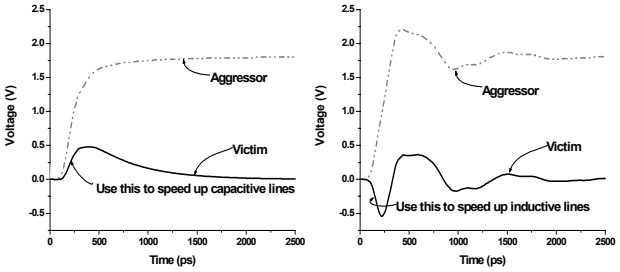

Figure 3. (a) Noise injected on a quiet victim when aggressors switch for a RC-dominated line. (b) Noise injected on a quiet victim when aggressors switch for an inductive line.

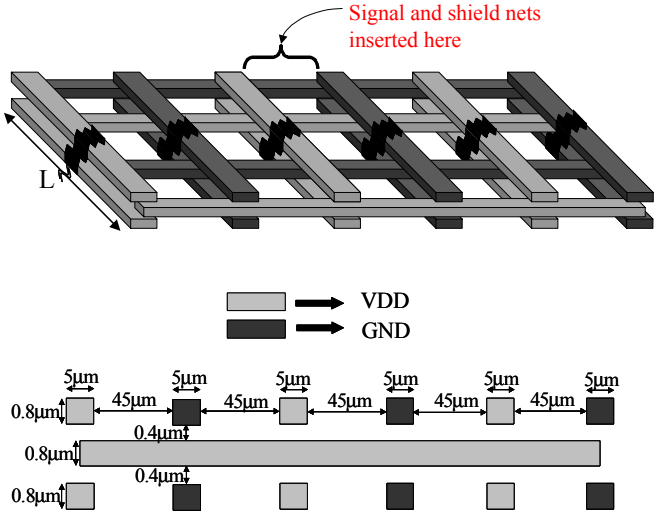

Figure 4. Power grid and interconnect structure.

case, the active shields should be switched in the same direction as the signal to reduce delays. For a more inductive line, the initial and sharper noise spike generated on the victim is in the opposite direction of the aggressor. This can speed up signal propagation along the signal wire if it is switching in the opposite direction to the shields. For opposite-phase active shielding, the decrease in loop inductance also results in reduced wire delays (due to a decrease in inductive impedance) and reduced ringing. In effect, the active shielding approach tries to make RC dominated lines more inductive and inductive lines more RC in nature. Our conjecture is that the in-phase and opposite-phase switched shields will be beneficial in different line impedance domains as mentioned in the previous section. In this section we set up a simple test case to determine where each of the active shielding approaches gives gains over passive shields. Figure 4 shows the power grid and interconnect structure used for resistance and inductance extraction for all simulation setups described in this paper. The top level metal (which was used for all test cases) has a minimum width of 0.5µm and a minimum spacing of 0.4µm.

FastHenry[4] was used for the resistance and inductance extractions. Capacitance extraction was performed using Raphael. All simulations use an inductively and capacitively coupled distributed RLC model (Figure 5) for the wires. The return path

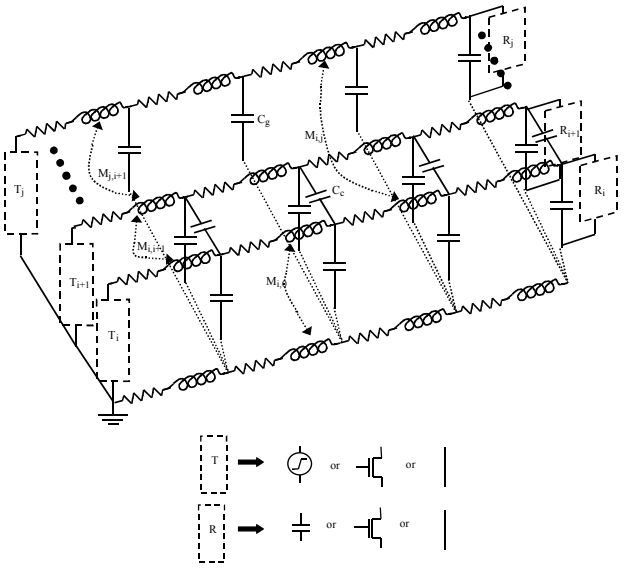

Figure 5. Circuit used for modeling the signal wires, shields and return path.

provided by the power grid is assumed to be the same for all wires. The return path is assumed to originate from the mid-point of the bundle of signal and shield wires used in any setup. This introduces negligible error in the analysis while maintaining simplicity for the purpose of simulation. For any setup, the transmission side of the wire is connected to a voltage source , an active device or is shorted to GND (for passive shields). The receiver side of the wire is connected to the return path through an active device, a capacitor or is shorted to the return path (for passive shields).

Figure 6 shows the simple test setup with only voltage ramps as inputs. The middle wire carries the signal while the side wires act as shields. When used as passive shields, the side wires are connected to GND at both ends. Vsh is the signal applied for active shielding. It is either in-phase or opposite-phase to Vin for the two active shielding approaches. The width (W) was swept so as to observe the line behavior in different impedance domains under the different switching schemes of active shields. In the results, performance will be measured through delays (measured between the 50% points of V_{in} and voltage at the end of the line with respect to a clean ground) and slopes (10-90% transition times at the end of the line) for a falling transition. Inductive effects will be measured as the peak undershoot at the end of the line.

Figures 7, 8 and 9 clearly show that the in-phase switching always results in faster slopes and skews with acceptable ringing when the line width is 1.5µm or less (region 1). As the lines become wider (and more inductive) the opposite-phase switched shields result in lower ringing (more than 50% reduction compared to passive shields), delay and slopes. The improvement in slopes through opposite-phase active shielding for wide wires (region 3) is considerable. There is a cross-over region (2) after which the in-phase switched shields result in slower signals and opposite phase shields result in faster signals than passive shields. This switch occurs due to the fact that the inductive impedance of the line increases for the in-phase switched shields and decreases for the opposite phase switched shields. The delay due to inductive impedance is greater than the contribution to delay by the cross-coupling capacitance and any reduction in inductive impedance results in greater savings in interconnect delay.

For the purpose of further analysis with real drivers at the line inputs, we divide the spectrum of wire widths into 3 domains. In the domain where the wire is RC dominated (region 1), in-phase switched shields are used since they are clearly superior in terms of performance with acceptable levels of ringing. These are typical wire widths used for buses. For wire widths in region 3, opposite-phase switched shields are superior since they reduce ringing, delays and transition times. These wire widths correspond to those used for global level clock nets. Wire widths

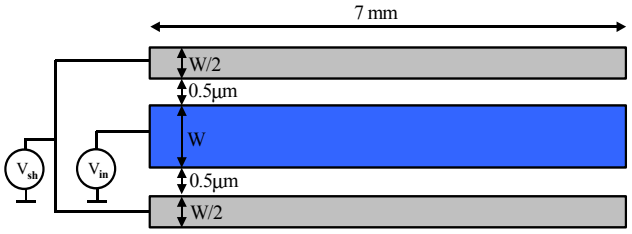

Figure 6. Test case setup

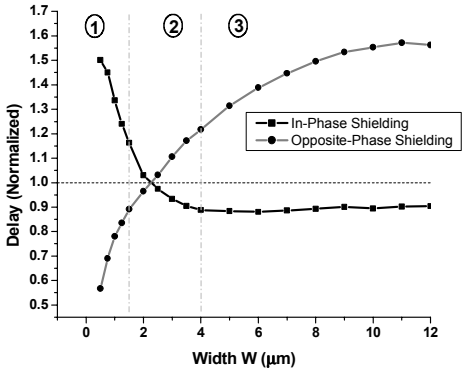

Figure 7. Wire delay for in-phase and opposite-phase active shielding. The delay is normalized to the delay obtained with passive shielding.

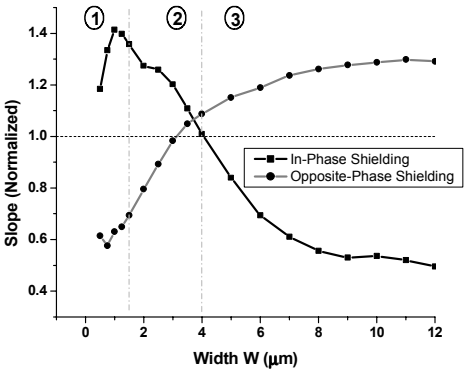

Figure 8. Signal slopes for in-phase and opposite-phase active shielding. Signal slopes are normalized to slopes obtained with passive shielding.

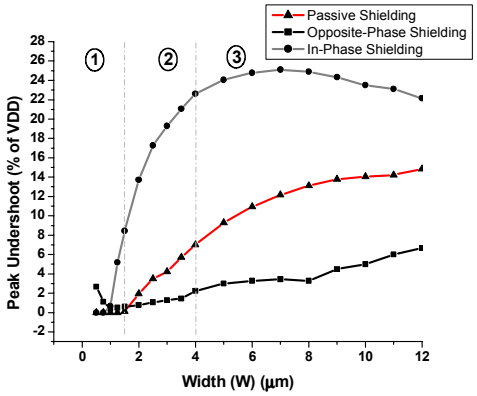

Figure 9. Peak undershoot obtained with passive and active shielding schemes. For inductive lines, opposite-phase active shielding results in lowest ringing.

in region 2 should use passive shields since the increased power consumption with active shields would not justify the marginal improvement in performance.

3. RC DOMINATED WIRES

Figures 10 and 11 show the setup used for comparing the active shields to passive shields for RC dominated wires. Figure 11 shows the setup for a longer wire with buffers inserted every 3.5mm, which is a typical buffering distance at the 0.18µm technology node. All wires are kept at minimum width with minimum spacing between them. Delay is measured at the 50% points between the input ramp Vin and the output of the receiver. Vagg1 and Vagg2 act as aggressors, switching simultaneously in the opposite direction. For a fair comparison, the input driver size

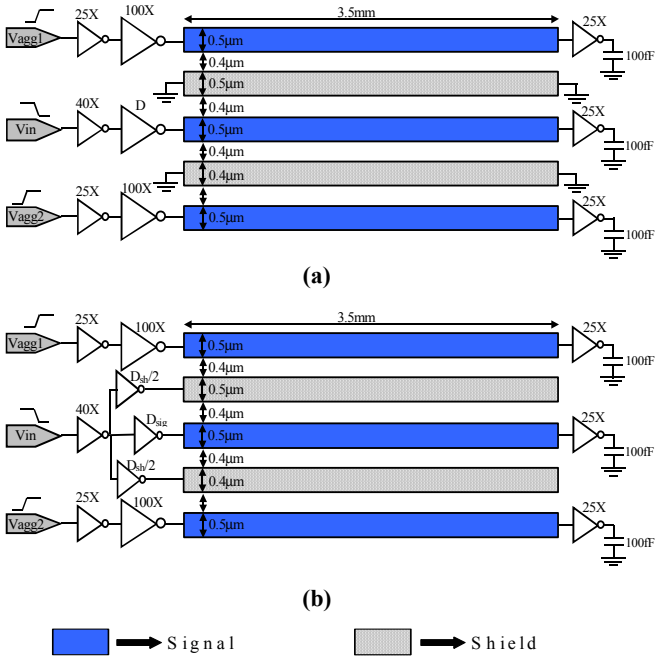

Figure 10.
(a) Passive shielding setup for a 3.5mm unbuffered line.
(b) Active shielding setup setup for a 3.5mm unbuffered line.

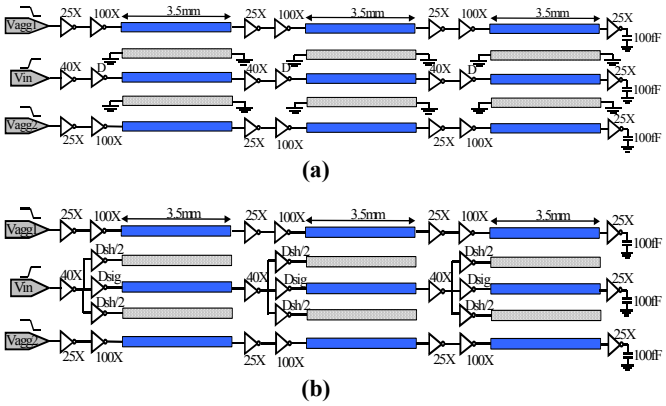

Figure 11. (a) Passive shielding for a 10.5mm line with buffers. (b) Active shielding for a 10.5mm line with buffers.

and the receiver driver size (along with its capacitive load) are kept constant at 40X for both cases. For the passively shielded wire, the size (D) of the second stage inverter is sized to obtain minimum delay. For the actively shielded configuration, the inverter sizes Dsh and Dsig are optimized to obtain minimum delay. There is an optimum driver size distribution among the shields and signal wire that results in minimum delay and slope. As the driving power is shifted from the signal to the shield wires, the effective cross-coupling capacitance seen by the middle wire decreases. There is an optimum point beyond which the reduction in effective cross-coupling capacitance cannot compensate for the reduction in drive strength on the signal wire. The optimization is performed for all line lengths in both setups and the results are shown in figure 12.

The gains from active shielding increase as the number of repeaters on the line increase, with up to 12% delay reduction for a 14mm line. This is due to the increased slew rates (from active shielding) at the end of each stage, which translates to delay gains in successive stages. The slew rates at the output of the receivers for the actively shielded wires are up to 33% faster than the passively shielded wires for the 100X aggressor case.

The ability of the active shields to enhance performance also depends on the strength of the aggressors on the active shields. Figures 10 and 11 show setups with fairly strong aggressors. The delay reductions are higher with weaker aggressors. Figure 12 also contains plots for cases where the aggressors are actively shielded (with the same signal and shield driver sizes as the victim) and when the aggressors are quiet. Using actively shielded aggressors results in an area penalty.

The major drawback of active shielding is the increased power consumption from switching the shields. For our setup, the power consumption for actively shielded wires was on an average almost 2X compared to passively shielded wires. The increase in power is heavily dependent on the ratio of shield width to signal wire width. If a wider signal wire (compared to the shields) is used, the increase in power consumption is less. Another tradeoff for increased performance is the degraded shielding property of the shields. The increased resistance (due to the driver) of the shield allows more noise to be injected onto the signal wire when the signal wire is quiet.

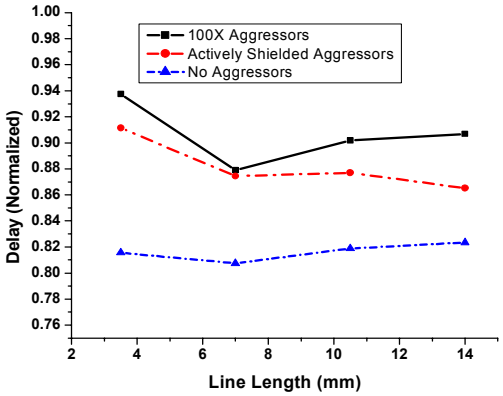

Figure 12. Active Shielding delay normalized to passive shielding delay. Buffers are inserted every 3.5mm on the line.

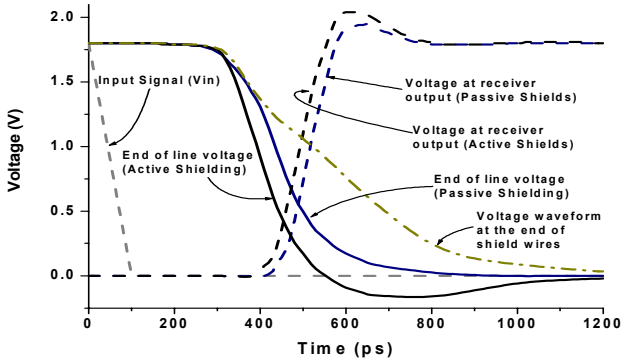

Figure 13. Typical waveforms at the end of the line and output of the receiver for an in-phase actively shielded 7mm long wire. The voltages are with respect to the local ground of the receiver.

4. INDUCTIVE WIRES

Clock nets are typically wide to reduce resistance, which helps to reduce skew and obtain faster transition times. With reduced resistance, inductive effects are very prominent in clock nets. The faster transition times in inductive wires are good for clock signals but the peak under/over shoots along with the associated ringing are undesirable from a signal integrity and device oxide reliability standpoint.

Global clock nets are shielded on both sides to keep the current return paths as close as possible. This reduces the loop inductance of the net. For very wide clock wires, just placing shields on either side is not sufficient to reduce the inductive ringing to an acceptable level. These clock nets are split into 2 or more fingers with shields inserted between the fingers [5]. As the number of fingers increases, the current loops become smaller and the net becomes less inductive. Splitting up the clock net into fingers comes at the expense of skew and slope degradation due to increased capacitance. Our goal with active shielding is to reduce the number of required fingers so as not to incur the associated performance penalty. From our analysis in section 2 we expect better skews and slopes when compared to the passively shielded clock net. Note that the actively shielded scheme differs from differential signaling. Only single-ended voltage sensing is required at the end of the line and the complementary shield

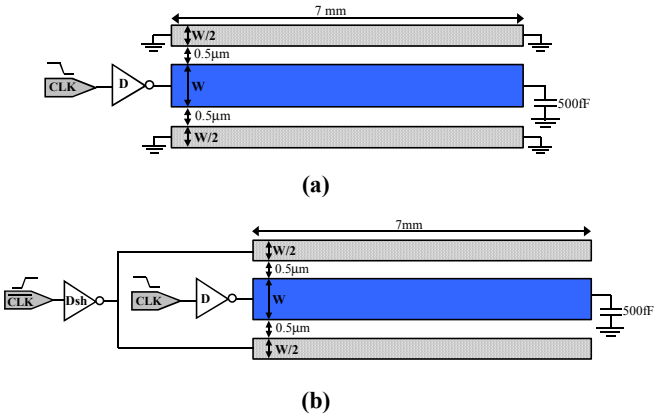

Figure 14. (a) Passively shielded clock net with one finger (no interdigitation). (b) Actively shielded clock net with one finger.

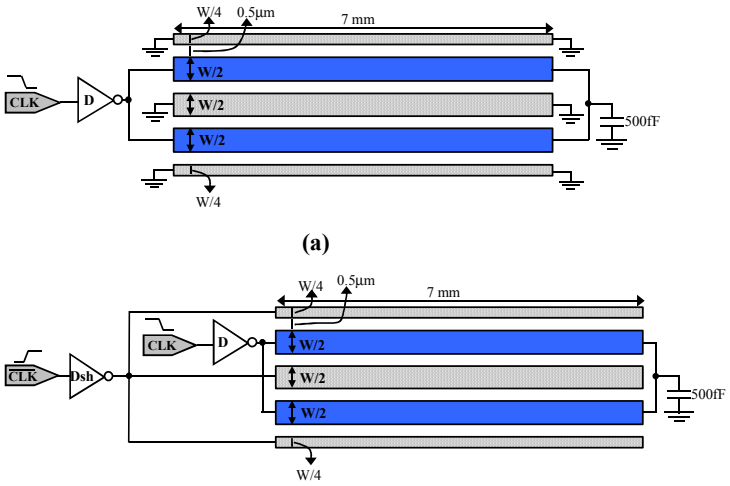

Figure 15. (a) Passively shielded clock net with two fingers. (b) Actively shielded clock net with two fingers.

signal does not need to be balanced in terms of wire width or driver size. The active shield signal is used to "condition" the actual clock signal.

Figures 14 and 15 show the setup we used for the cases where the clock net had 1 and 2 fingers, respectively. Certain details have not been included in the figures for clarity. For reliable fabrication purposes, wires in actual designs cannot be wider than a certain limit. We have taken this into account by splitting any single metal wire wider than 7µm, into 2 or more wires of equal width and a spacing of 0.5µm between these wires.[1] Clock nets with more fingers have a similar topology to the one shown in figure 15. Different values for W, which is the total signal width of the clock, were used. The line length and capacitive load at the end of the line were always kept constant.For a particular value of W and under different topologies (or number of fingers) the total

[1] The actual width after which a wire has to be split, varies from one fabrication company to the other. The limit can vary from 4µm to 12µm.

width of the clock net and dedicated shield wires are each W. The metal spacing is always 0.5µm. The driver size (D) of the clock net is scaled up linearly with W. For the active shielding, the active shield driver size (Dsh) is swept so as to meet the constraints on the clock signal. The complementary clock signal for the active shield driver is assumed to be generated by the drive chain leading up to the clock driver.

The constraints for a particular target clock frequency are that wire delay (measured between the 50% points from clock driver output to the end of the line) should be within 20% of the cycle time.[2] The slope (measured as 10-90% time at the end of the net) should be within 25% of the cycle time. The peak ringing, which is measured as the maximum deviation from 0V after a single falling transition, should be within 5% of VDD (90mV). For the passively shielded case, the only optimization process to meet these constraints is to change the number of fingers. The actively shielded clock net has an extra variable for optimization – the active shield driver size (Dsh). As the active shield driver size increases, delay and peak ringing decrease but the transition time has an optimum point. The value of Dsh that allows us to meet the constraints is used. Table 1 lists the results for W values of 3, 5, 10 and 20µm. Since our main goal is to reduce ringing, the clock net is split up into increasing number of fingers till the ringing constraint is met. The reduction in ringing decreases as the number of fingers increase and the splitting is stopped at a point when adding an extra finger results in less than 15% reduction in ringing.

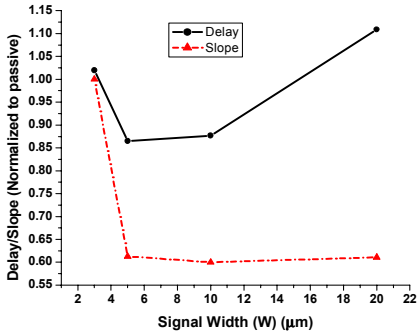

Figure 16. Skew and slope of actively shielded clock net normalized to skew and slope of passively shielded clock net.

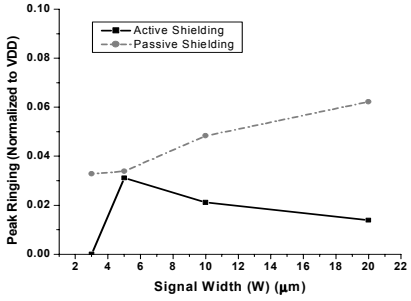

Figure 17. Peak ringing for actively and passively shielded clock nets.

[2] Delay and skew are used interchangeably for the clock net.

Figure 16 represents the relative performance (delay and slope) of the best actively shielded setup when compared to the best passively shielded setup (i.e., the one with lowest ringing) for different values of W. The peak ringing for these setups is shown in figure 17. Gains of up to almost 14% in delay and ~40% in transition times are obtained while ringing is always maintained within 5% of VDD. The constraints on the clock signal could be met for all values of W with active shielding while passive shielding could meet the constraint only for the case where W=3µm. The configuration with W=3µm also strengthens our analysis in section 2, where we conjectured that active shielding for wire widths between 1.5-4µm would not yield significant performance improvements to justify the extra power consumption. Figures 18 and 19 show the waveforms (for a single transition and periodic wave) obtained from HSPICE simulations for the case where W=10µm. It is clear that the signal obtained with active shields has better delay and slope with highly suppressed ringing.

The actively shielded setup can consume as much as 2X the power of the passively shielded setup for the same topology. However, the increase in power consumption at the global clock level does not have a major impact on the total power consumption for the clock distribution network. According to [6], the power consumption at the local clock level is at least an order of magnitude higher than clock at the global level. If a comparison is made across different topologies in table 1, the

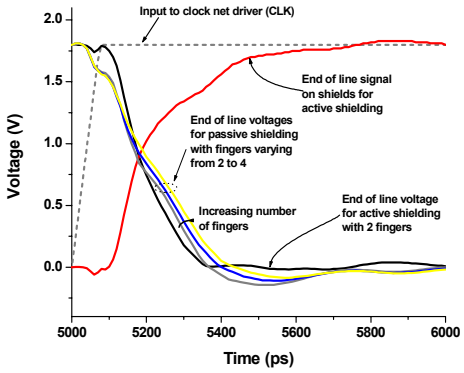

Figure 18. Voltage waveforms at the end of a clock net with a signal width (W) of 10µm.

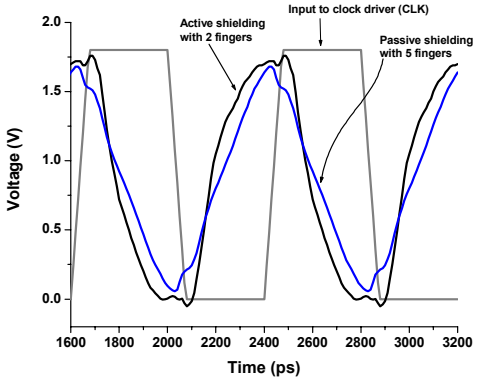

Figure 19. Clock signal at the end of the clock net with a signal width (W) of 10µm. The clock is operating at 1.25GHz.

Table 1. Results for actively and passively shielded clock nets. Rows in gray indicate active shielding.

Signal Width (W), Clock Driver Size (D) Target Clock frequency (f) Maximum Delay(T_D), Maximum Slope (T_{RF})	Number of fingers	Active shield Driver Size (Dsh)	Delay (ps)	Slope (ps)	Peak Under-shoot (mV)	Energy per clock period (pJ)
W = 3µm D = 600X f = 1GHz	1*	150	165	244	0	23.2
T_D = 200ps T_{RF} = 250ps	1*	-	161	245	59	12.5
W = 5µm D = 1000X f = 1GHz	1*	600	141	166	56	38.2
T_D = 200ps T_{RF} = 250ps	1	-	144	216	118	18.0
	2	-	163	271	61	20.7
W = 10µm D = 2000X f = 1.25 GHz T_D = 160ps T_{RF} = 200ps	1	2000	121	133	106	62.7
	2*	1450	128	179	38	69.9
	1	-	129	237	223	31.1
	2	-	123	261	143	31.5
	3	-	134	280	111	33.4
	4	-	146	299	87	35.1
W = 20µm D = 4000X f = 1.25 GHz T_D = 160ps T_{RF} = 200ps	1	3800	135	97	349	116.6
	2	4000	117	133	99	121.8
	3*	3300	122	182	25	126.4
	1	-	141	264	332	60.2
	2	-	115	273	170	60.1
	3	-	112	291	127	58.4
	4	-	110	298	112	60.8

* ➡ Clock constraints met

actively shielded configuration out-performs (in terms of transition times) the passively shielded configuration for approximately the same energy consumption. Examples for such a comparison are between W=5µm/Active Shields/1 finger and W=10µm/Passive Shields/4 fingers and between W=10µm /Active Shields/1 finger and W=20µm/Passive Shields/4 fingers.

5. CONCLUSION

The concept of active shielding was examined in the context of RLC global interconnects. A simple analysis showed that for RC-dominated wires, in-phase switching of shields still helps to speed up signal propagation with an acceptable increase in ringing. For wide inductive wires, signal propagation is enhanced and ringing is reduced by switching the shields in the opposite direction. Shields switching in the same phase as the signal have shown up to ~12% and ~33% reduction (compared to passive shields) in delay and transition times for realistic test cases. The opposite-phase active shielding concept was used to optimize clock nets and resulted in up to ~40% reduction in transition times. It also resulted in much lower ringing (up to 4.5X reduction). Active shielding consumes extra power and should be used when the performance gains justify the increase in power.

6. REFERENCES

[1] S. Morton, "On-chip Signaling", Special Topic Evening Session, *International Solid State Circuits Conference (ISSCC*), 2002.

[2] H. Kaul, D.Sylvester and D. Blaauw, "Active Shields: A New Approach to Shielding Global Wires", Proc. Great Lakes Symposium on VLSI (GLSVLSI), pp 112-117, April 2002.

[3] S. van Dijk and D. Hély, "Reduction of Interconnect Delay by Exploiting Cross-talk", *Proceedings of European Solid State Circuits Conference*, 2001.

[4] M. Kamon. M. Tsuk and J. White, "FastHenry: A Multipole-Accelerated 3-D Inductance Extraction Program", IEEE Transaction on MTT, Vol. 42, No. 9, pp. 1750-1758, September 1999.

[5] Y. Massoud, et. al., "Layout Techniques for Minimizing On-Chip Interconnect Self Inductance", *Proceedings of Design Automation Conference (DAC)*, pp 566-571, 1998.

[6] P. Restle, et. al., "A Clock Distribution Network for Microprocessors", *IEEE Journal of Solid State Circuits*, Vol.36, pp 792-799, May 2001.