

Clock Net Optimization using Active Shielding

Himanshu Kaul, Dennis Sylvester, and David Blaauw
EECS Department, University of Michigan, Ann Arbor, MI 48109
{hkaul, dmcs, blaauw}@umich.edu

Abstract - We propose the use of active shields, which are shield wires that are switched concurrently with the signal wire, to improve performance and reduce inductive ringing for RLC wires. This technique significantly reduces ringing behaviour (up to 4.5X) and offers better slopes (up to 40% reduction) and signal propagation delays than the traditional (passive) shielding approach, all of which are shown in the context of a clock net optimization.

1. INTRODUCTION

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

Shields (additional GND or VDD wires) are frequently inserted to reduce loop inductance and the associated inductive effects [1,2]. Global clock nets are shielded on both sides to keep the current return paths as close as possible. For very wide clock wires, simply placing shields on both sides is not sufficient to reduce inductive ringing to acceptable levels. In this case the shields are inter-digitated, where the clock net is split into two or more fingers with shields inserted between the fingers [3]. As the number of fingers increases, the current loops become smaller and the net becomes less inductive. Inter-digitating the clock net with conventional (passive) shields comes at the expense of delay and slope degradation due to increased capacitance and the RC nature of the line.

In this paper, we address this problem using *active* shields, where the shield wires are actively switched. Capacitive coupling can be used to aid transitions on RC dominated wires by ensuring neighbouring wires switch in-phase with the signal wire [4]. For inductive wires, where inductive coupling is stronger than capacitive coupling, active shields can be switched in the opposite direction to aid the transitions on the signal wire (through inductive coupling). At the same time they also serve as better current return paths than passive shields, reducing ringing. This approach obtains the benefits of differential signaling while maintaining the simplicity of single-ended signaling. Note that since the receiver senses voltages only on the signal wire, there is no requirement of balanced loads or wire widths. To demonstrate the concept, a simple setup is used in Section 3 to show the region of feasibility (in terms of wire widths) where opposite-phase active shields out perform passive shields. Optimizations on various clock nets with the passive and active shielding schemes are carried out and compared in Section 4. The impact of process variation is analyzed in Section 5 to demonstrate the feasibility of the approach at worst-case process corners. Section 6 compares




Fig. 1. Power grid and interconnect structure.




Fig. 2. Circuit used for modeling the signal wires, shields and return path. A 7mm line was represented by 14 lumped segments.

active shielding to differential signaling and discusses the inherent trade-offs between the two approaches.

2. EXTRACTION AND SIMULATION SETUP

Fig. 1 shows the power grid and interconnect structure used for resistance, capacitance, 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 and spacing of 0.5µm. FastHenry [5] 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 (Fig. 2) for the wires. The 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. Industrial 0.18µm MOS models are used with a VDD of 1.8V. HSPICE is used for all simulations.

3. CONCEPT FEASIBILITY REGION

In order to determine the region of feasibility of active shields, we initially use a simplified setup without drivers to gauge the possible performance improvement and suppression of the ringing behaviour with active shields. Fig. 3 shows the test setups using only voltage ramps (100ps rise/fall times) as inputs. For a fair comparison of the passive and active shielding schemes, we initially constrained the total shield width to equal the total signal

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

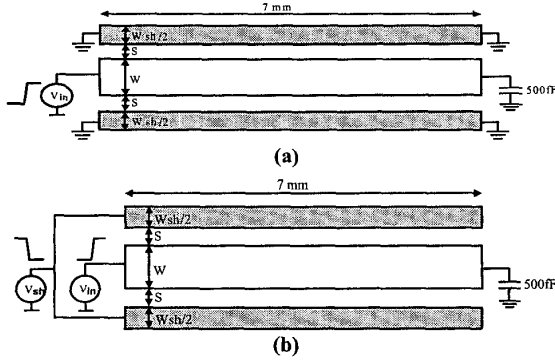


Fig. 3. Test case setup for analyzing (a) passive shielding and (b) active shielding.

width ($W=W_{sh}$), while interconnect spacing was kept minimum ($0.5\mu\text{m}$). The middle wire carries the signal while the side wires act as shields. When used as passive shields, the side wires are connected to the GND grid. V_{sh} , which is opposite in phase to V_{in} , is the signal applied for active shielding. The width (W) is swept to observe the line behaviour in different impedance domains under the two shielding schemes. Delay is measured between the 50% points of V_{in} and the voltage at the end of the line and slope is measured as the 10-90% transition times at the end of the line (for a falling transition). Inductive effects are measured using the peak undershoot at the end of the line. The results from this analysis are shown in Figs. 4 and 5. The ringing with active shielding is consistently lower, while the transition times are also lower for signal widths greater than $4\mu\text{m}$, along with a modest reduction in delay.

Relaxing the constraints on shield width and inter-metal spacing, the total shield width (W_{sh}) was varied from 0.5 to 3 times the signal width (W) and spacing (S) was varied from $0.5\mu\text{m}$ to $2\mu\text{m}$. The goal is to find an optimal spacing and shield size for passive and active shielding within this solution space to minimize peak undershoot. We then compare the performance of the two shielding schemes at their respective optimal points. Simulations showed that peak undershoot is smallest with minimum spacing of $0.5\mu\text{m}$ for all combinations (for active and passive shielding). For passive shielding, shield width has a monotonic effect, with peak undershoot decreasing as shield width increases for all signal widths (Fig. 6a). With active shielding there is a distinct optimal point, occurring when the total shield width is equal to the signal width ($W_{sh}=1.0\bullet W$) for all signal widths

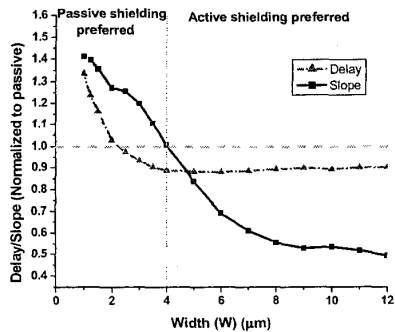


Fig. 4. Active shielding delay and slope for various line widths for setup in Figure 3 ($W=W_{sh}$, $S=0.5\mu\text{m}$). Delays and slopes are normalized with respect to passive shielding delays and slopes.




Fig. 5. Peak undershoot with passive and active shielding for various line widths for setup in Fig. 3 ($W=W_{sh}$, $S=0.5\mu\text{m}$). Peak undershoots are normalized to VDD.

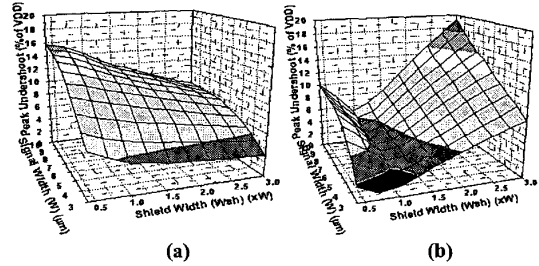


Fig. 6. Effect of varying W and W_{sh} (with $S=0.5\mu\text{m}$) for the setups in Fig. 3 for (a) passive shielding and (b) active shielding.

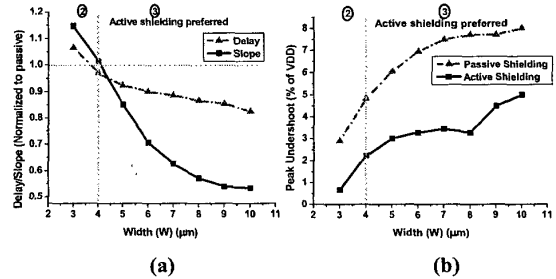


Fig. 7. (a) Delays and slopes (normalized to passive shielding) at optimal undershoot points in Fig. 6. (b) Peak undershoot at optimal undershoot points in Fig. 6.

greater than $4\mu\text{m}$ (Fig. 6b). Fig. 7 compares the active and passive shielding at their respective optimized points in this solution space. The active shielding design point has smaller delays, slopes, and ringing than the optimized passive shielding setup (for $W>4\mu\text{m}$). Furthermore, these benefits are obtained using much less shielding resources – about 1/3 the shielding requirements for the optimal passive shielding configuration in this solution space.

4. CLOCK NET OPTIMIZATION

Next, we apply the active shielding approach to three 7mm clock net structures with signal widths (W) of 5, 10, and $20\mu\text{m}$. Fig. 8 shows the setup for the cases where the clock net has two fingers. Clock nets with more fingers have a similar topology to the one shown in Fig. 8. Clock nets with three different signal widths (W) are optimized and for each value of W the total width of the clock net and dedicated shield wires are each kept equal to W under different topologies (i.e., number of fingers). The metal spacing is always $0.5\mu\text{m}$. The driver size (D) of the clock net is scaled up linearly with W . For active shielding, the active shield driver size (D_{sh}) is swept to meet the constraints on the clock signal as described below. The complementary clock signal




Fig. 8. (a) Passively shielded and (b) actively shielded clock net with 2 fingers.

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. This constraint is important since delay or clock latency translates directly to clock skew. The slope (measured as 10-90% delay at the end of the net) should be no greater than 25% of the cycle time. Peak ringing, measured as the maximum deviation from 0V after a single falling transition, should be less than 5% of VDD (90mV). For the passively shielded case, the only optimization variable is the number of fingers. Since the primary goal is to reduce ringing, the clock net is split into an increasing number of fingers until the ringing constraint is met. The reduction in ringing saturates as the number of fingers increase and the splitting procedure is stopped at a point when adding an extra finger results in less than a 15% reduction in ringing for the passively shielded clock. The actively shielded clock net has an extra variable for optimization – namely the active shield driver size (D_{sh}). As the active shield driver size increases, delay and peak ringing reduce monotonically but the transition time shows an optimal point. The minimum value of D_{sh} that meets the above constraints is used. Table 1 lists the results for W values of 5, 10, and 20 μ m. The constraints could not be met with passive shielding for any of the clock nets, while active

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

Signal Width (W) Clock Driver Size (D) Target Clock frequency (f) Delay Constraint (T_D) Slope Constraint (T_S)	Number of fingers	Active shield Driver Size (D_{sh})	Delay (ps)	Slope (ps)	Peak Under- shoot (mV)	Energy per clock period (pJ)
W=5 μ m D=500X f=1GHz T_D =200ps T_S =250ps	1	300X	141	166	56	38.2
	2	-	144	216	118	18.0
W=10 μ m D=1000X f=1.25 GHz T_D =160ps T_S =200ps	1	1000X	121	133	106	62.7
	2*	723X	128	179	38	69.9
	1	-	129	237	223	31.1
	2	-	123	261	143	31.5
	3	-	134	280	111	33.4
W=20 μ m D=2000X f=1.25 GHz T_D =160ps T_S =200ps	1	1900X	135	97	349	116.6
	2	2000X	117	133	99	121.8
	3*	1650X	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	

* \Rightarrow Clock constraints met

$$\text{Driver Size } 1X \Rightarrow W_{NMOS} = 2 \cdot L_{min} \quad (W_{PMOS} = 2 \cdot W_{NMOS})$$




Fig. 9. Delay and slope of optimized actively shielded clock net normalized to delay and slope of optimized passively shielded clock nets.

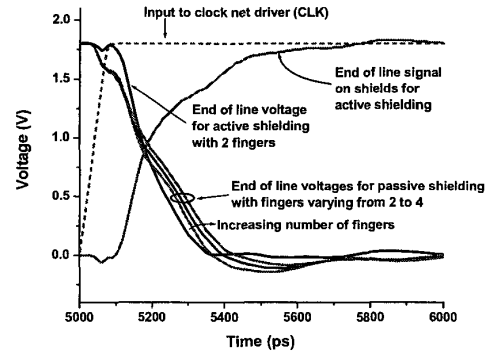


Fig. 10. Voltage waveforms at the end of a clock net with a signal width (W) of 10 μ m.

shielding allowed the constraints to be met for each of the three different clock nets.

Fig. 9 represents the relative performance (delay and slope) and peak undershoot of the best actively shielded setup compared to the best passively shielded configuration (the one with lowest ringing) for different values of W. Active shielding resulted in consistent gains of almost 40% in transition times, while ringing was kept below 5% of VDD. The delays for the clock nets with signal widths (W) of 5 and 10 μ m are ~14% lower with active shielding. Optimal active shielding resulted in up to 4.5X decrease in ringing for W=20 μ m. It is clear from the waveforms in Fig. 10 that the signal obtained with active shields has faster transitions with highly suppressed ringing.

The actively shielded setup can consume as much as twice 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 distribution at the global level. If a comparison is made across different topologies in Table 1, the actively shielded configuration outperforms (in terms of transition times) the passively shielded configuration for approximately the same energy consumption. Relevant examples for such a comparison are between W=5 μ m/Active Shields/1 finger and W=10 μ m/Passive Shields/4 fingers and between W=10 μ m/Active Shields/1 finger and W=20 μ m/Passive Shields/2 fingers.

Ringing for the passively shielded clock can also be decreased by decreasing the size of the clock driver (D). This




Fig. 11. Drive chain structure used to analyze the impact of process variation on actively shielded clock nets. Sizes D and D_{sh} are the same as listed in Table 1.

results in even slower transition times. Only one configuration ($W=5\mu\text{m}$ /Passive Shields/1 finger) allowed us to reduce the main clock driver size, since the only violation of the constraints was the peak undershoot. A range of slower drivers allowed this configuration to meet the clock constraint. However, the optimal active shield setup still resulted in a 29% faster transition time than the fastest transition time obtainable with the range of smaller clock driver sizes that allowed the passively shielded clock to meet the constraints.

5. PROCESS VARIATION IMPACT

In the clock net optimization process of the previous section, the clock signal and its complementary signal input to the active shield driver, have the same slew rate and zero skew between them. With no process variation, the complementary signals can be generated from a common input signal with minimal skew between them and similar slew rates. Since the buffers in the drive chain are spatially close together, process variation is not expected to create considerable mismatch between the two signals. However, to analyze the effects of worst-case channel-length variation on the performance of active shields, we simulated a drive chain (Fig. 11) that generates the signals to drive the clock net and the active shields for the optimal configurations obtained in the previous section. The channel-length variation of each inverter is assumed to be independent of the others and the maximum deviation from nominal is assumed to be $\pm 5\%$, which is intended to be representative of spatially proximate intra-die fluctuations [7]. We simulated all combinations of extreme variations in the channel lengths of individual inverters to obtain worst-case positive deviation in slopes, delays, and ringing. Simulations showed that ringing and delays were relatively insensitive to these variations – worst-case ringing was under 5% of VDD and the maximum deviation in delay was 5%. Transition times were more sensitive to process variation. Worst case variability reduced the gains (with respect to passive shielding at the worst process corner) from $\sim 40\%$ for $W=5, 10,$ and $20\mu\text{m}$ to 35%, 34%, and 34%, respectively.

6. DIFFERENTIAL SIGNALING COMPARISON

Since the active shielding approach has similarities to differential signaling we compare the active shielding and differential signaling schemes in this section. We used voltage ramps to simulate the drivers so as not to bias the results due to a particular driver/receiver architecture. The setup for differential signaling (Fig. 12) used the same total metal width as the active shielding setup with equal shield and signal widths. To simulate a low-swing differential driver, the complementary signal lines were driven by opposite-phase voltage ramps, each with a voltage swing of $0.5 \cdot V_{DD}$. The receiver is an ideal (zero-delay) low-swing differential to full-swing single-ended converter. The delays and slopes with the differential scheme (Fig. 13) are compared using the full-swing output of the receiver. The delays with active shielding are smaller when $W \geq 6\mu\text{m}$ while the transition times are clearly faster with differential




Fig. 12. Setup for differential signaling comparison. Total metal width is the same as the setup in Fig. 3 (when $W=W_{sh}$). The ideal receiver at the end of the line converts low-swing differential signals to full-swing single-ended signals.

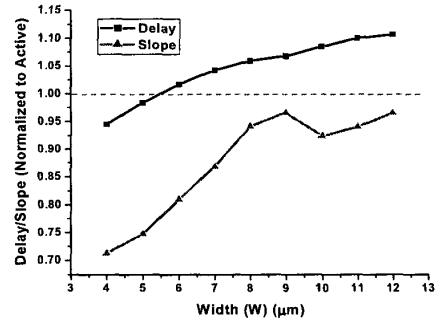


Fig. 13. Delays and transition times with differential signaling setup of Fig. 12 normalized to active shielding setup of Fig. 3b.

signaling for $W < 8\mu\text{m}$. At larger signal widths, the transition times with the two schemes are comparable. Also, the differential signaling scheme resulted in half the power consumption of the actively shielded scheme since the voltage swing on the lines is only $0.5 \cdot V_{DD}$, as opposed to full swing for active shielding. Active shielding results in performance that is better than passive shielding but not clearly better than differential signaling. Active shielding results in much simpler driver and receiver circuits (i.e., standard CMOS inverters) than differential signaling and provides a reasonable trade-off between design complexity and performance.

7. CONCLUSIONS

A new shielding method for clock nets that results in better signal integrity and performance than conventional routing approaches has been proposed. This improvement is achieved at the expense of global clock distribution power, which is typically much smaller than local clock-related power consumption. Various clock nets were optimized with passive and active shielding to achieve performance and signal integrity constraints. Active shielding always resulted in significantly faster transition times (up to 40% reduction) and lower ringing (typically 2-4X).

8. REFERENCES

- [1] S. Morton, "On-chip Signaling", Special Topic Evening Session, *Intl. Solid State Circuits Conference (ISSCC)*, 2002.
- [2] L. He and K. M. Lepak, "Simultaneous shield insertion and net ordering for capacitive and inductive coupling minimization", *Proc. Intl. Symposium on Physical Design*, pp. 55-60, 2000.
- [3] Y. Massoud, et al., "Layout techniques for minimizing on-chip interconnect self inductance," *Proc. DAC*, pp 566-571, 1998.
- [4] S. van Dijk and D. Hély, "Reduction of Interconnect Delay by Exploiting Cross-talk", *Proceedings of European Solid State Circuits Conference*, 2001.
- [5] M. Kamon, M. Tsuk, and J. White, "FastHenry: A multipole-accelerated 3-D inductance extraction program," *IEEE Trans. on MTT*, pp. 1750-1758, Sep. 1999.
- [6] P. Restle, et al., "A clock distribution network for microprocessors," *IEEE J. of Solid State Circuits*, pp 792-799, May 2001.
- [7] <http://public.itrs.net>, ITRS, 2001 update