Performance Optimization of Critical Nets Through Active Shielding

Himanshu Kaul, *Student Member, IEEE*, Dennis Sylvester, *Member, IEEE*, and David Blaauw, *Member, IEEE*

***Abstract*—We propose the concept of active shields—shields that switch concurrently with a signal wire of interest. Active shields aid signal transitions through the coupling between the signal wire and shields. For *RC* dominated wires, the active shields switch in the same phase as the signal wire since capacitive coupling is the dominant coupling mechanism. For wires with dominant inductive coupling, active shields switch in the opposite phase of the signal wire. We show that under fixed area and input capacitance constraints, in-phase active shielding outperforms traditional (passive) shielding and wire sizing/spacing techniques for minimizing delays and transition times on *RC*-dominated wires. For *RLC* wires, we demonstrate a region of feasibility (in terms of signal wire widths) for which opposite-phase active shielding outperforms the passive shielding technique. Opposite-phase active shielding suppresses ringing behavior to a greater degree than passive shields, providing similar performance to differential signaling while maintaining the simplicity of single ended signaling. The benefits of opposite-phase active shielding as compared to passive shielding are shown in the context of various clock net optimizations where reductions in ringing behavior (up to 4.5X) and transition times (up to 40% reduction) are achieved.**

***Index Terms*—Capacitive coupling, inductive coupling, on-chip interconnects, shields.**

I. Introduction

CROSS-COUPLING capacitance and inductance have become very prominent in scaled CMOS technologies [1]–[3]. Both can cause noise problems as well as degrade performance. To address these problems, shields (additional GND or VDD wires) are typically inserted between critical global wires [4]–[6]. When inserted between two wires, almost all capacitively coupled noise injected between wires are 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 (or SF•Cc where SF is a switch factor that may vary between -1 and 3 [7]) 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 [8], [9]. This also limits 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 faster signal slew rates. According to [8], the recommended ratio of signal to shield wires has diminished to 2:1 for high-performance processor designs. Since the effective capacitance and inductance of a wire is determined not just by the geometry but also by the switching behavior of neighboring wires, the shields can be switched to further reduce delay and/or loop inductance of critical wires. Thus, *active shields* can be used to change the effective coupling capacitance and loop inductance of a wire through switching rather than layout optimizations (Fig. 1). In this paper, we will show that active shields can reduce the resources used in shielding signals while maintaining the same performance or improve performance under an area constraint. When the shields are assumed to be switching, a reoptimization of the shield wire widths may be required.

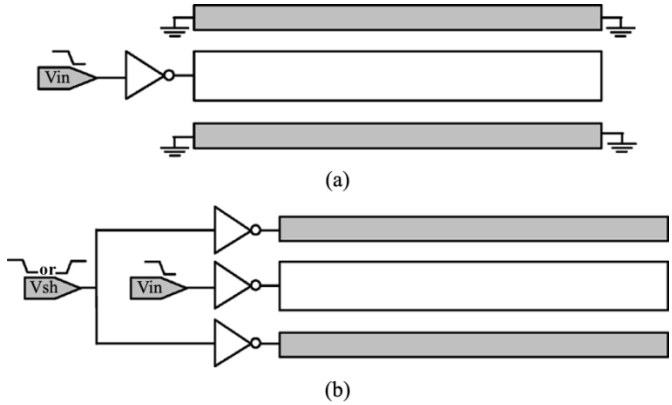

Fig. 1. (a) Signal wire with passive shields. (b) Signal wire with active shields. Active shields switch either in the same phase or in the opposite phase of the signal.

For thin and narrow wires the resistance dampens out most of the inductive effects. In these cases 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-witching) shield design. A similar approach was described in [10], where the active shields are called "booster wires." However, [10] does not consider the delay penalty due to added input capacitance of the drivers for the "boosted wires" on the previous stage. In [11], the authors describe the advantages of redundant switching. Their proposed method results in increased capacitive loading on the previous stage—while they mention that the increased delay due to this can be offset by improved delays on the wire, there are no results shown to justify this. Also the analysis in [11] was limited to a particular technology. We show that even under strict input capacitance and area constraints, in-phase actively shielded wires improve

Manuscript received March 8, 2003; revised January 24, 2004. This work was supported in part by the MARCO/DARPA Gigascale Systems Research Center and equipment donations from Intel Corporation and Sun Microsystems. This paper was recommended by Associate Editor I. Verbauwhede.

The authors are with the Electrical Engineering and Computer Science (EECS) Department, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: hkaul@umich.edu; dmcs@umich.edu; blaauw@umich.edu).

Digital Object Identifier 10.1109/TCSI.2004.838247

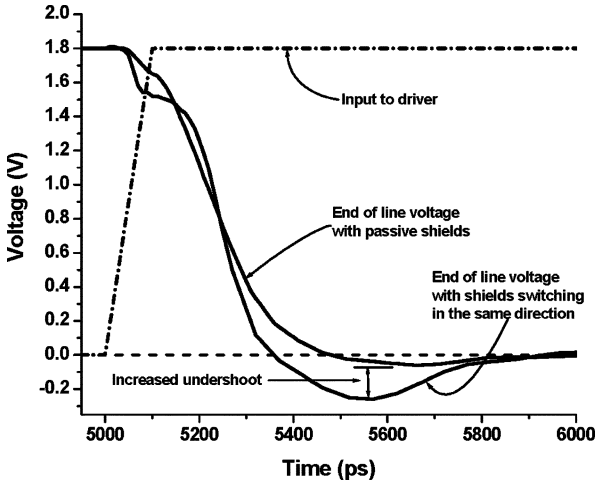

Fig. 2. Voltage waveforms at the end of a *RLC* line when the shields are switching in the same direction and when passive shields are used.

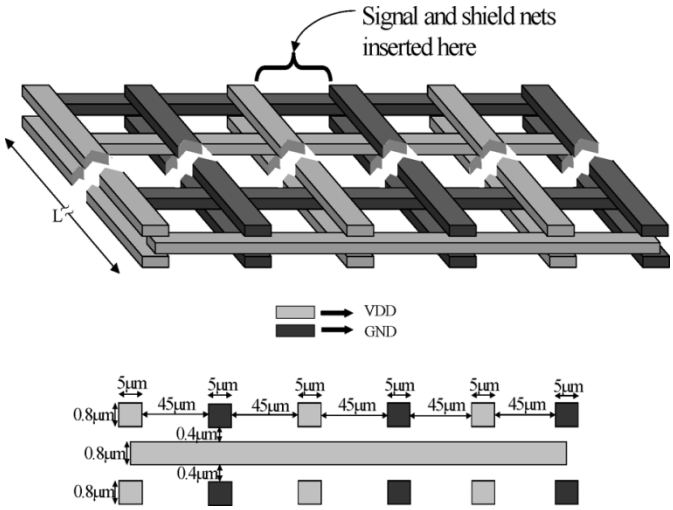

Fig. 3. Power grid and interconnect structure.

performance (delay and slope) for *RC*-dominated wires for various technologies.

In general, switching the shields in the same direction worsens the current return paths and results in more inductive wire behavior. Although the work in [11] claims that redundant switching in the same direction results in delay gains even for very inductive wires, there is no mention about the reliability concerns that stem from such inductive wire behavior. The increased inductive nature of the signal line can be tolerated for lines that are *RC* dominated, but for wider lines inductive effects are exacerbated, such as over/undershoots, which directly affects gate oxide reliability, and increased crosstalk. This phenomenon is illustrated in Fig. 2 which shows the increased undershoot at the end of a 7-mm signal line (with signal and shield widths of 3 μm each) when in-phase active shielding is used. Also, as wires become wider the gains from in-phase switching of neighbors deteriorate as the ratio of coupling-to-ground-capacitance (C_c/C_g) decreases. Though this limitation can be overcome by placing wires at the upper or lower layer in parallel (and switching them in-phase) with the signal wire [10], [11], this is not a practical solution with traditional preferred-direction routing methodologies dictating that wires in adjacent layers run in orthogonal directions. For wires where inductance cannot be neglected, it can be more favorable to switch the shields in the opposite direction of the signal line. For *RC*-dominated lines, the capacitive coupling is much stronger than mutual inductive effects, while in inductive lines the inductive coupling is dominant. The opposite-phase active shielding scheme can use this dominant inductive coupling to aid transitions on the signal wire. The opposite-phase active shields can also serve as better current return paths than passive shields and result in lower ringing. This approach obtains the benefits of differential signaling [12] 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.

The rest of the paper is organized as follows. Section II details the methodology used for wire parasitic extraction and technology information for the simulations in this work. In Section III we analyze the wire impedance region (in terms of wire widths for a typical 0.18 μm global interconnect) to gauge when (and by how much) the two active shielding approaches are more beneficial than passive shielding when performance and inductive behavior (ringing) is taken into account. We also show that the increase in ringing due to in-phase active shielding is tolerable for *RC*-dominated wires. We then explore the active shielding approach for *RC*-dominated wires in Section IV of the paper and present a detailed analysis of in-phase active shielding. The benefits of opposite-phase active shielding as compared to passive shielding are shown in the context of clock net optimizations in Section V. In Section V we also point out that opposite-phase active shielding is similar to differential signaling but it avoids the need for careful design of differential drivers/receivers, matching effects, etc. Section VI concludes with a summary of our results.

II. Parasitic Extraction Methodology

In Sections III and V inductance is included in the wire models and in Section IV a distributed *RC* model is used for the wires to reduce extraction and simulation time. All simulations for *RLC* wires in this work are performed for a single interconnect technology. Fig. 3 shows the power grid and interconnect structure used for resistance, capacitance, and inductance extraction for the simulation setups in Sections III and V. The top level metal (which was used for all test cases) has a minimum width and spacing of 0.5 μm [13]. FastHenry [14] was used for the resistance and inductance extractions. Raphael [15] was used for all capacitance extractions. All simulations in Sections III and V use an inductively and capacitively coupled distributed *RLC* model (Fig. 4) 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 midpoint 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

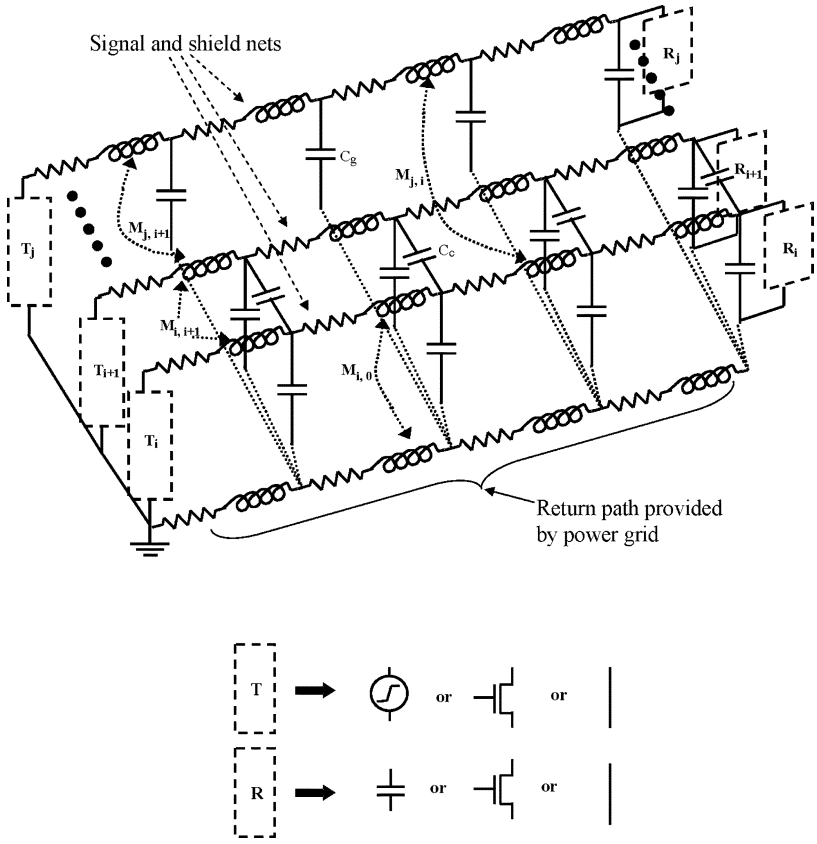

Fig. 4. Circuit used for modeling the signal wires, shields and return path. A 7-mm line was represented by 14 lumped segments.

to the return path through an active device, a capacitor, or is shorted to the return path (for passive shields). The signal and shield wires are coupled to the return path through their ground capacitance and the mutual inductance between the self-inductance of the wire and return path. Since the loop inductance is not predetermined, this model allows the actual current path (and loop inductance) to be determined only during simulation and resembles a PEEC-based model for *RLC* wires. Industrial 0.18-μm MOS models are used with a VDD of 1.8 V. HSPICE [16] is used for all simulations.

III. Concept Feasibility Region

In order to study the effect of the different switching behaviors of the shields on performance and signal integrity as the wire impedance is changed, we initially use a simplified setup without drivers. Fig. 5 shows the 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 the GND grid at both ends. $V_{\rm sh}$ is the signal applied for active shielding. It is either in-phase or opposite-phase to $V_{\rm in}$ for the two active shielding approaches. The width (W) is swept to observe the line behavior in different impedance domains under the three switching behaviors (passive, in-phase, or opposite-phase) of the shields. We constrained the total shield width to equal the total signal width ($W = W_{\rm sh}$), while interconnect spacing was kept minimum ($S = 0.5$ μm). A more elaborate analysis on wire sizing and spacing is undertaken in Sections IV and V. Delay is measured between the 50%

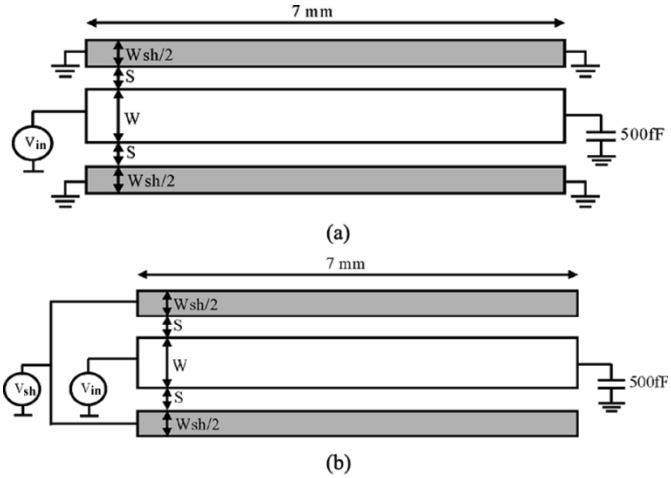

Fig. 5. Test case setup for analyzing: (a) passive shielding and (b) active shielding for a *RLC* wire.

points of $V_{\rm in}$ 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.

Fig. 6 clearly shows that in-phase switching always results in faster slopes and delays with acceptable ringing when the line width is 1.5 μm 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) and faster delays and slopes. The improvement in slopes through opposite-phase active shielding for wide wires

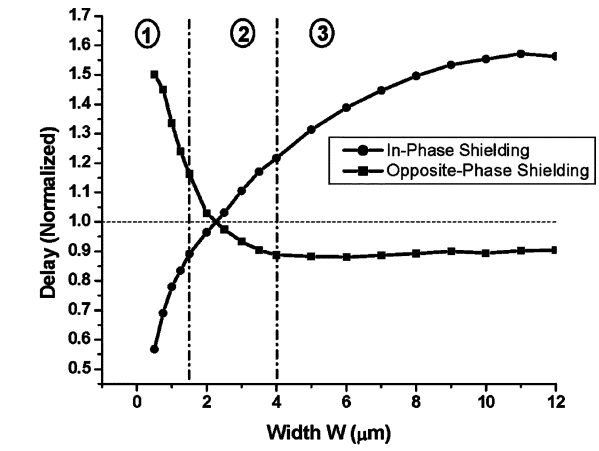

(a)

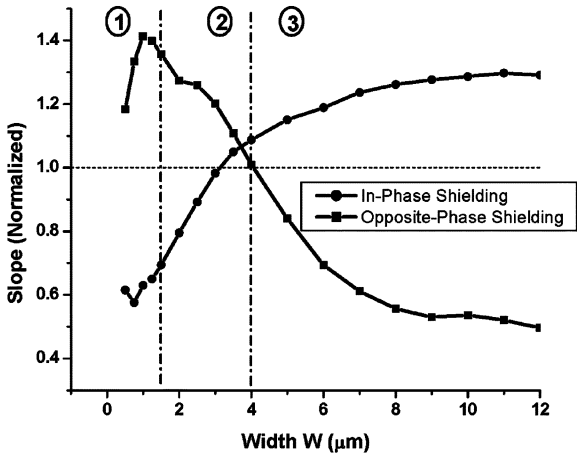

(b)

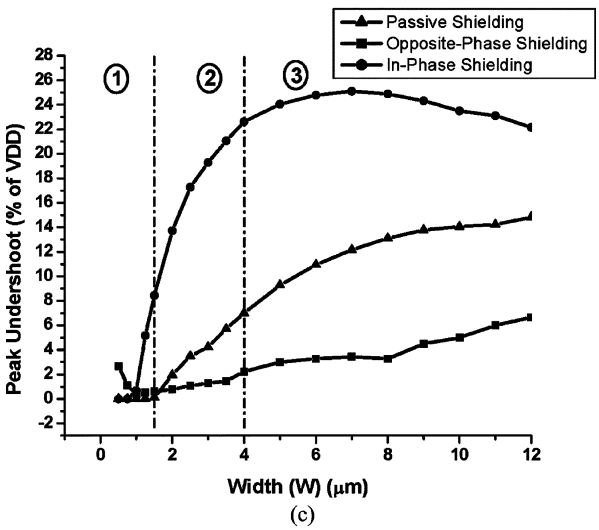

(c)

Fig. 6. (a) Wire delay for in-phase and opposite-phase active shielding. The delay is normalized to the delay obtained with passive shielding. The line width design space can be divided into three distinct regions, labeled 1, 2, and 3 above. (b) Signal slopes for in-phase and opposite-phase active shielding. Signal slopes are normalized to slopes obtained with passive shielding. (c) Peak undershoot obtained with passive and active shielding schemes. For inductive lines, opposite-phase active shielding results in lowest ringing.

(Region 3) is considerable. There is a crossover region (2) after which the in-phase switched shields result in slower transitions 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 a reduced inductive impedance results in an overall lowering of interconnect delay. Wire widths between these two domains (i.e., Region 2) should use passive shields since the increased power consumption with active shields would not justify the marginal improvement in performance. In the next two sections we explore in much greater depth the wire-impedance regions where the two active-shielding approaches show a potential for improvement in wire performance over passive shields.

IV. Active Shielding for *RC* Wires

The concept of in-phase 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 simultaneously in the same direction with equal ramp rates (a best case scenario can result in a factor of -1 according to [7]). 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 in-phase active shields as proposed in this work results in better performance when compared to other methods like passive shielding and wire spacing/sizing [17] under the same area constraint and capacitive load on the previous stage. Though this approach is meant to complement and not replace the buffer insertion methodology [18]–[21], we do show that even for buffered lines active shielding results in improved delays. For very long lines buffers would still be required to meet performance requirements.

Section IV-A describes the active shielding approach and develops a simplified analytical model of the theory behind it. Section IV-B describes the simulation setup used to compare active shielding to other approaches, and provides results. Optimization approaches to obtain the maximum gains from active shielding are discussed in Section IV-C.

A. Active Shielding Model

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 greatly reduced. This can 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 signal propagation in two cases: 1) a wire with shields on its sides and 2) a wire wide enough to be split up into three wires while maintaining the same footprint (to avoid routing area penalties). In the case of the wide (unshielded) wire, the total ground capacitance of the middle wire is reduced (and the

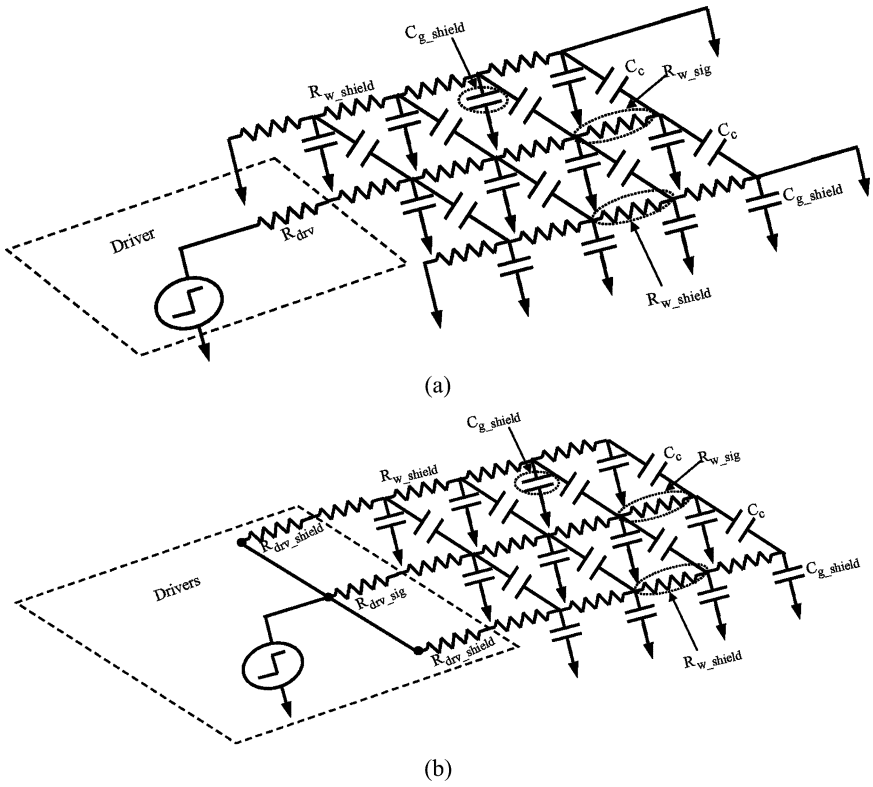

Fig. 7. *RC* network used for modeling: (a) passive shields and (b) active shields.

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. Henceforth, the wide unshielded wire will be referred to as the fat wire. 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 an analytical insight into the tradeoffs involved in maintaining a constant capacitive load on the previous stage, we use a simple model (based on work in [7]) to compare the delays of the passively shielded and actively shielded configurations. The model shown below assumes there are no coupled aggressors. Fig. 7(a) and (b) shows the equivalent circuits used for modeling the delays using passive and active shields. The labeled wire parasitics are denoted as per unit length. The equivalent resistance of an inverter with nMOS width of 1 μm is R_0. Thus, if the driver size (in terms of nMOS device width) in the passively shielded configuration is $W_{\rm drv}$, then the driver resistance is

$$R_{\rm drv} = \frac{R_0}{W_{\rm drv}}. \tag{1}$$

The resistances of the drivers for the side and middle wires for the actively shielded configuration are

$$R_{\rm drv\_shield} = \frac{R_0}{W_{\rm drv\_shield}} \tag{2}$$

$$R_{\rm drv\_sig} = \frac{R_0}{W_{\rm drv\_sig}} \tag{3}$$

where $W_{\rm drv\_shield}$ is the nMOS device width of each of the shield drivers and $W_{\rm drv\_sig}$ is the nMOS device width of the signal line driver.

To maintain the same input capacitance when converting from passive to active shields the constraint

$$W_{\rm drv} = W_{\rm drv\_sig} + 2 \bullet W_{\rm drv\_shield} \tag{4}$$

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

$$T_{\rm passive} = 0.693 R_{\rm drv}(C_{\rm g\_sig} + 2 \bullet C_{\rm c}) + 0.378 R_{W\_\rm sig} \bullet (C_{\rm g\_sig} + 2 \bullet C_{\rm c}) \tag{5}$$

where $C_{\rm g\_sig}$, $\rm R_{w\_sig}$ and $C_{\rm c}$ are the ground capacitance, wire resistance and coupling capacitance associated with the signal wire per unit length.

Delay due to active shields is given by

$$T_{\rm active} = 0.693 R_{\rm drv\_sig}(C_{\rm g\_sig} + 2 \bullet {\rm k_{sig}} \bullet C_{\rm c}) + 0.378 R_{\rm w\_sig} \bullet (C_{\rm g\_sig} + 2 \bullet {\rm k_{sig}} \bullet C_{\rm c}) \tag{6}$$

$$k_{\rm sig} = 1 - \frac{t_{\rm r\_sig}}{t_{\rm r\_shield}} \tag{7}$$

and $t_{\mathrm{r\_sig}}$ and $t_{\mathrm{r\_shield}}$ are the 10%–90% signal transition times on the signal and shield wires as computed below.

$$t_{\mathrm{r\_sig}} = 2.2R_{\mathrm{drv\_sig}}(C_{\mathrm{g\_sig}} + 2 \bullet k_{\mathrm{sig}} \bullet C_{\mathrm{c}}) + 0.9R_{\mathrm{w\_sig}} \bullet (C_{\mathrm{g\_sig}} + 2 \bullet k_{\mathrm{sig}} \bullet C_{\mathrm{c}}) \quad (8)$$

$$\mathrm{t_{r\_shield}} = 2.2R_{\mathrm{drv\_shield}}(C_{\mathrm{g\_shield}} + \mathrm{k_{shield}} \bullet C_{\mathrm{c}}) + 0.9R_{\mathrm{w\_shield}} \bullet (C_{\mathrm{g\_shield}} + \mathrm{k_{shield}} \bullet C_{\mathrm{c}}) \quad (9)$$

$$k_{\mathrm{shield}} = 1 - \frac{t_{\mathrm{r\_shield}}}{t_{\mathrm{r\_sig}}}. \quad (10)$$

Equations (5), (6), (8), and (9) are derived from the delay and far-end transition time equations for distributed-*RC* wires driven by a resistive driver in [22]. Equations (7) and (10) are derived from the $C_$eqexpression in [7]. The delay and far-end transition time equations assume a step input to the driver and result in an optimistic model. However, some of the error from this optimistic assumption is negated by the fact that the far-end slew rates (which are much slower than slew rates anywhere else on the wire) are used for computing the effective coupling capacitance along the entire line. Approaches that can be used to improve the accuracy of the model are discussed in Section IV-C. Obtaining the rise/fall times on the signal and shield wires (t_{r_sig}, t_{r_shield}) involves iterations since these times depend on the effective coupling capacitance during switching and the effective coupling capacitance in turn depends on the rise/fall times. The effect of aggressors on the active shields can be taken into account by adding $2 \bullet C_{\mathrm{c\_agg}}$ (twice the coupling capacitance between the active shield and aggressor) to the $C_{\mathrm{g\_shield}}$ term. This is a simplified assumption that does not take into account the slew rates on the shield and aggressors. Calculating a more accurate effective coupling capacitance to the shields with respect to the aggressors would involve more iterations and we show in the section on driver sizing that this simplified assumption is still useful enough to derive optimal driver size distributions.

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 increase their slew rates. In this case, the effective resistance of the middle driver increases (since the total device width is fixed at W_{drv}), creating an inherent trade-off in designing with active shields. Thus, $R_{\mathrm{drv\_sig}}$ in (6) increases while the C_{c} term decreases as drive strength is shifted to the side wires. We expect that there is a particular distribution of driver sizes for which the delay of the signal wire is minimized.

This optimal point depends heavily on the values of the ground and coupling capacitances. As the C_{c} term increases, it is desirable to shift more of the driving capability to the active shield wires. Fig. 8 shows the delay using active shields normalized to that with passive shields for a typical configuration of wire widths and spacing. The only difference between the curves is the value of the coupling capacitance C_{c}. The total driver size was fixed at 100X[1] while the size of the driver for the signal line (in the actively shielded case) was swept. For driver sizes, the model breaks down less than half the total driver size, but it does point to the existence of an optimal

[1]Driver size 1X $\Rightarrow W_{\mathrm{nMOS}} = 2 \bullet L_{\mathrm{min}}$ ($W_{\mathrm{pMOS}} = 2 \bullet W_{\mathrm{nMOS}}$).

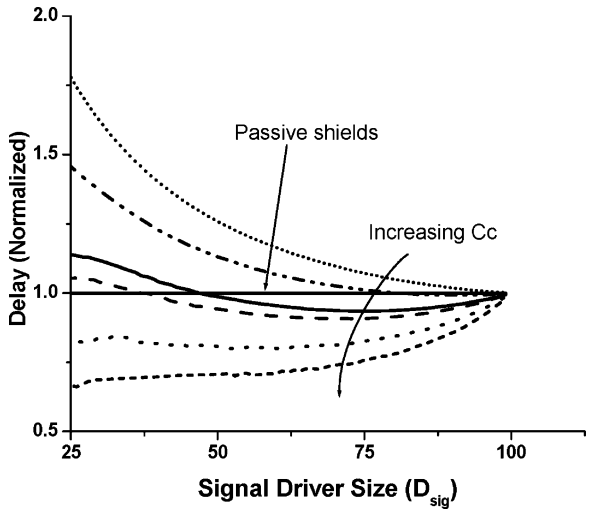

Fig. 8. Delay with active shields as driver size distribution (between signal and shield wires) is changed for different coupling capacitances.

delay point. The figure shows that as the coupling capacitance is increased (by reducing the spacing) the delay gains compared to the passively shielded case rise. There are scenarios 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 active shields and signal wire). An immediate implication of this is that design rules could be relaxed to allow sub-minimum spacing at the global layers so as to achieve higher gains in performance through active shielding. The minimum spacing at these layers is not always dictated by technology constraints, but by limits on cross-coupling capacitance. Routes that use in-phase active shielding could benefit from tighter spacing design rules (e.g., equivalent to those at lower level metals which can be two to three times smaller [23]) to create faster and denser interconnections. Note that minimum line width and aspect ratio would remain unchanged with this design rule alteration, maintaining a highly manufacturable design (as long as spacing constraints are not set by the manufacturing process).

B. Simulation Results

The setups for comparing the various performance optimization techniques for *RC* wires are first described, followed by a detailed analysis of the results.

1) Simulation Setup: Typically fat wires are used to reduce signal delays by suppressing the line resistance. For our analysis, the passive and active shielding 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* in Fig. 9) is held constant 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, which should have a minimal impact.

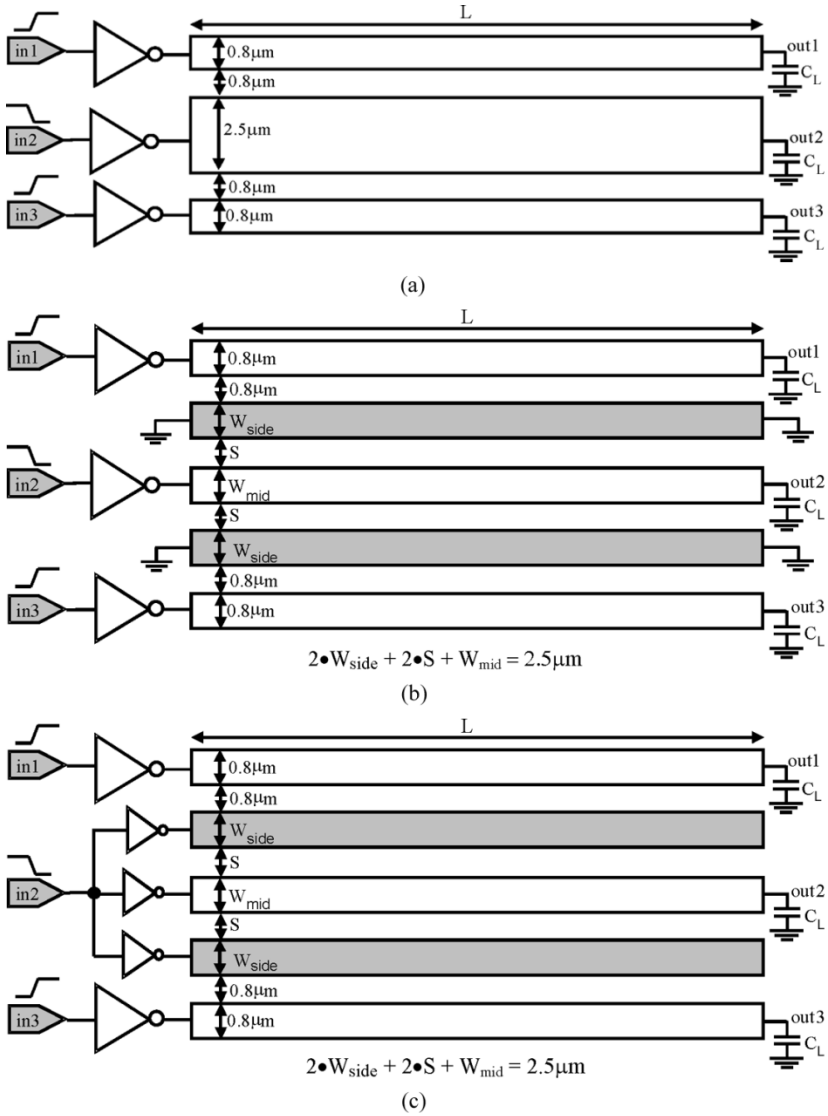

Fig. 9. Setup for (a) fat wire, (b) passively shielded wire, and (c) actively shielded wire. Wires in grey are shield wires.

Fig. 9 shows the simulation setups for three cases—fat wire, passive shielding, and active shielding, respectively. 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 [driven by *in2* in Fig. 9(a)] is reduced in width while the spacing to the aggressors is increased by the same amount. This represents a wire sizing/spacing methodology to reduce delays and injected noise on wires under an area constraint. 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 then be effectively reduced during switching).

For the active shielding approach, the delay depends on the ground and coupling capacitances and the resistance of each of the three wires (all of which depend on the wire width distribution among the signal and shield wires) along with the optimal driver size distribution among the three wires used for signal transmission. The optimum delay with the passive shielding approach depends on the ground capacitance and resistance of only the signal wire (along with the coupling capacitance to the passive shields). Optimizing the ground and coupling capacitances and line resistances yields optimal widths of the three wires under a constant area constraint for the two shielding schemes. 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. Fig. 10 shows the 2-D interconnect topology used for extraction purpose. A top level metal layer is assumed with one ground plane.

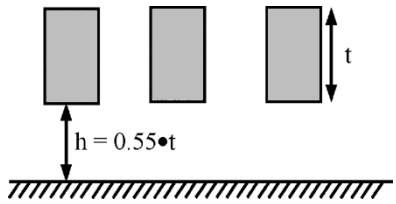

Fig. 10. Cross section of interconnect structure.

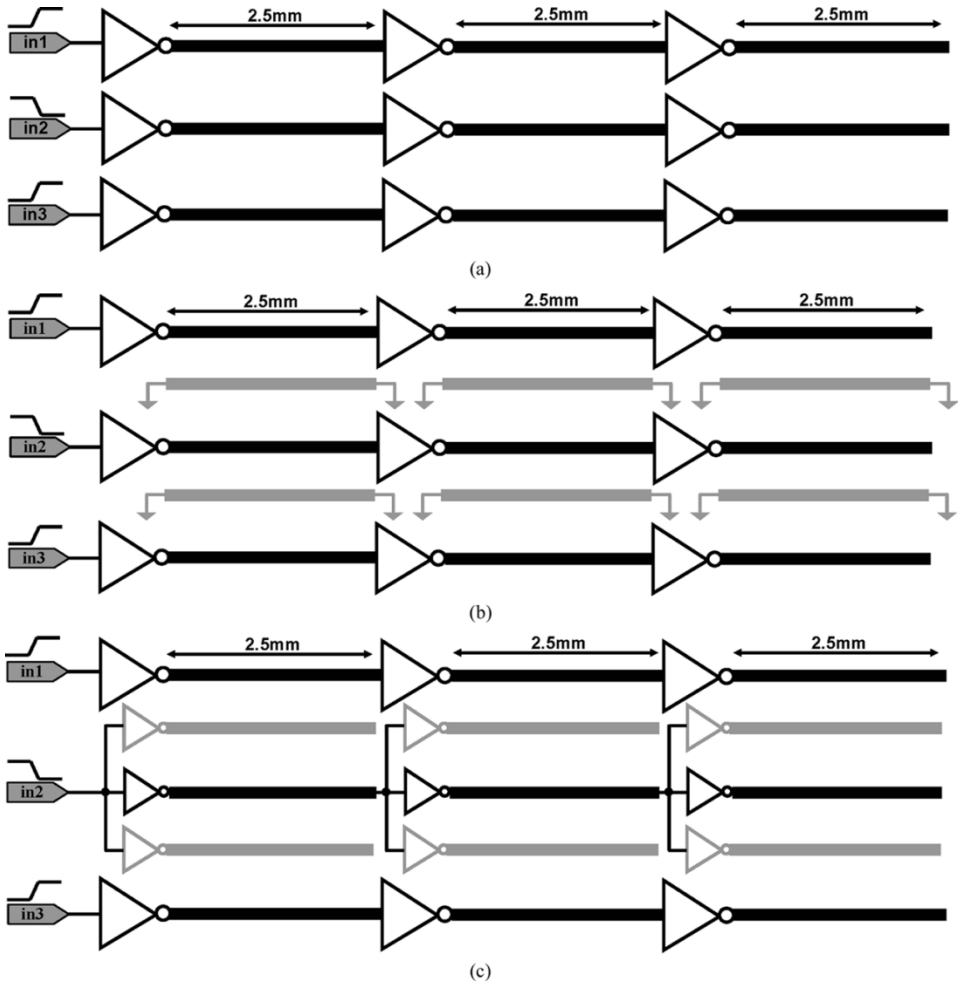

Fig. 11. Schematic for setup with repeaters every 2.5 mm for a 7.5-mm line for: (a) fat wire; (b) passively shielded wire; and (c) actively shielded wire.

The footprint (including any shields) for transmitting the signal was kept constant at 2.5 μm for each technology and the aspect ratio was fixed at 2.4 (which dictated the minimum width and spacing). This footprint constraint allows the shielded wires to lie in the *RC*-dominant region where the increase in ringing can be considered tolerable. In Section III, the region demarcated as *RC* dominant had a footprint of less than 4 μm when including the area used by the shields. Since the inductive effect can be assumed to be negligible for these wires, only distributed *RC* models are used for the wires in this section. For each technology, five different configurations in which the wire could be split were used to obtain the optimum shield and signal widths for each of the passively and actively shielded wires. The spacing between the wires was always set to the minimum allowable in order to achieve the smallest resistance for the wires (by allowing them to be as wide as possible) within the area constraint. 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 7.5 and 3.75 mm, respectively. For the 7.5-mm (3.75 mm) line length, the driver size being driven by node *in2* was 100X (62.5X) and the aggressor driver sizes were 100X (62.5X) each. A 7.5 mm line length with an inter-repeater distance of 2.5 mm was also simulated with driver and repeater sizes of the signal and aggressor wires fixed at 50X each. The schematic of this setup for the different configurations is shown in Fig. 11. Delays were measured between the 50% points of *in2* and *out2* and slopes were measured as 10%–90% delays at node *out2*. The load capacitance (C_{L}) corresponds to the input capacitance of a 12.5X inverter. To calculate power, a switching activity of 0.5 for an 800-MHz operating frequency was applied at node *in2* and the current drawn from the supplies by the driver(s) driven by the signal *in2* was measured. We also ran test cases with both aluminum and copper wiring to investigate the impact of the line resistivity.

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) for each of the shielding schemes. Typical waveforms for both active and passive shielding are shown in Fig. 12. Active shielding clearly demonstrates superior delay and slew rate characteristics, which is the strength of the active shielding approach.

Simulation results showed that the optimal delay using active shields is always better than that of the fat wire and passive shields. Fig. 13 shows the performance of actively and pas-

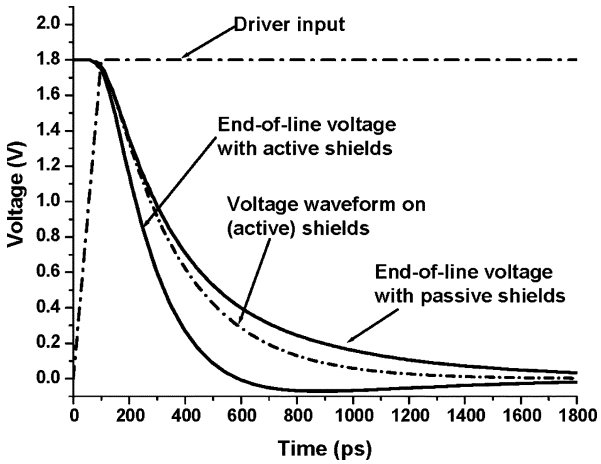

Fig. 12. Voltage waveforms at the end of the signal wire for active and passive shields. The signal on the active shields is also shown.

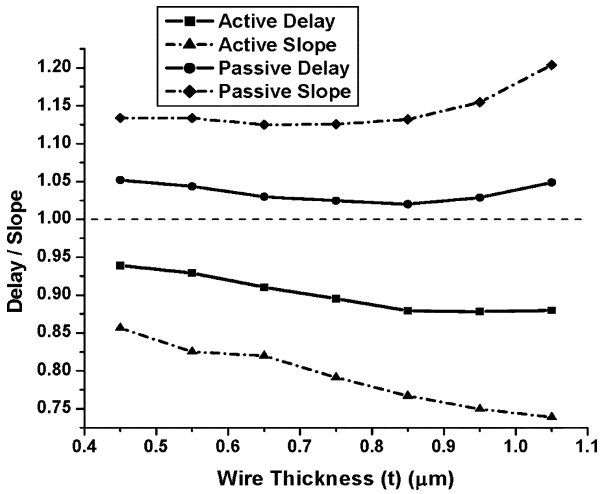

Fig. 13. Optimal delay/slope versus wire thickness for copper wire of length 7.5 mm and aggressors switching. Delays and slopes are normalized to the fat wire delays and slopes.

sively shielded wires (normalized to the fat wire) for one particular configuration as the wire thickness is varied. The results for the other configurations for two distinct wire thicknesses are given in Table I. The delays with passive shielding are insensitive to switching activity of the aggressors. For active shielding some of this insensitivity is lost since the ability of the active shields to aid the signal wire during transitions depends on the switching activity of the aggressors. The fat unshielded wire is the most sensitive to aggressors as expected, with delay and slope varying by as much as 22% and 13% (respectively) between the quiet aggressor case and the case where the aggressors switch simultaneously 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 the gains through active shielding reduce as the total resistance of the wire increases. As the wire resistance increases, the slopes at the end of the wire degrade and the gains from active shielding reduce. This is demonstrated by smaller gains for each setup as the wire thickness is reduced. For the same reason, the aluminum wiring cases show less improvement than with copper and the shorter wire lengths show more gains than the 7.5 mm cases (without repeaters). In the current 0.18- and 0.13-μm processes, repeaters are typically placed every 1–3 mm to reduce delay—these distances maximize the usefulness of active shielding as is evident in the 7.5-mm case with repeater insertion.

The results in Table I are obtained using a fixed load of a 12.5$\times$ inverter at the end of the wire. To gauge the sensitivity of the delay gains of active and passively shielded wires, this load was changed to a 100$\times$ inverter for the 7.5-mm unrepeated copper wire and the 7.5-mm copper wire with repeaters every 2.5 mm. This was simulated for the configuration with a thickness (t) of 0.45 μm since that configuration has a higher wire resistance and would result in the delays and slopes being more sensitive to the load at the end of the wire. The actively shielded configuration still yielded the least delay and transition times with the increased load. The reduction in delay gains was not more than 2% for either case (active or passive shielding), while the degradation in gains for transition times was limited to 6%–7% for either case.

For the case in which the fat wire is reduced in width to increase spacing from the aggressors, simulations (see Fig. 14) show that the delay gain [obtained by the optimal width of signal wire in Fig. 9(a)] is the same as active shielding but the improvements in slopes are much less than with active shielding. Noise analysis showed that noise immunity in this case is comparable to that of active shields but not as good as passive shields.

Fig. 15 shows the power consumption for a configuration for which the delay was optimized. The power consumption (normalized to the fat wire) for all the delay-optimized configurations are also shown in Table I. The fat wire always consumes the most power since it has the largest capacitance associated with it and must also fight the aggressors. The actively shielded wires consume more power than the passively shielded configurations since the shields also switch. In Table I this increase in power with respect to passive shielding ranges from 18% to 40%. A noise analysis confirmed that active shielding results in enhanced noise immunity compared to the fat wire. The far-end noise produced by switching aggressors when the victim wire is quiet is shown for two of the configurations from Table I in Table II. Active shields are not as good at screening functional noise as passive shields. The drivers on the (active) shield wires present a highly resistive path to ground and result in degraded shielding properties.

C. 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: 1) wire widths and spacings and 2) distribution of the total driver size over the three wires. Fig. 8 showed that maximizing the coupling capacitance results in smaller delays (due to the increased Miller effect). As a result we concluded that the spacing between the side and middle wires should be kept minimum (as dictated by minimum spacing constraints.) As mentioned in Section IV-A, design rules can be relaxed to allow de-

TABLE I
IN-PHASE ACTIVE SHIELDING AND PASSIVE SHIELDING DELAYS AND SLOPES FOR VARIOUS LINE CONFIGURATIONS. DELAYS, SLOPES, AND POWER ARE NORMALIZED TO DELAYS, SLOPES, AND POWER OF THE FAT WIRE. ROWS IN GREY INDICATE ACTIVE SHIELDING AND UNSHADED ROWS INDICATE PASSIVE SHIELDING

Material	Length (mm)	Aggressor switching activity	Wire Thickness (t)					
			t = 0.65μm			t = 0.95μm		
			Delay	Slope	Power	Delay	Slope	Power
Copper	7.5	Switching	0.91	0.82	0.90	0.88	0.75	0.92
			1.02	1.12	0.76	1.03	1.15	0.66
Copper	7.5 (repeated every 2.5mm)	Switching	0.68	0.73	0.75	0.67	0.74	0.80
			0.74	0.84	0.61	0.72	0.89	0.64
Copper	7.5	Quiet	0.94	0.86	0.91	0.94	0.81	0.92
			1.10	1.20	0.80	1.20	1.32	0.76
Copper	3.75	Switching	0.78	0.74	0.85	0.74	0.71	0.88
			0.85	0.92	0.65	0.78	0.86	0.67
Aluminum	7.5	Switching	0.94	0.84	0.89	0.93	0.78	0.91
			1.10	1.20	0.74	1.17	1.29	0.65

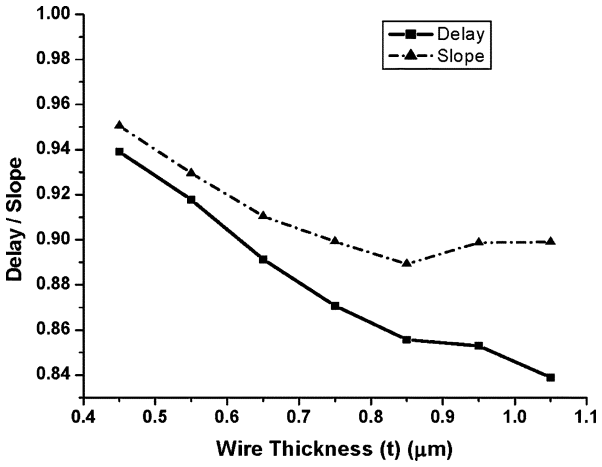

Fig. 14. Delay and slope versus wire thickness for narrowed middle wire (copper) of length 7.5 mm and aggressors switching. Delays and slopes are normalized to fat wire delays and slopes.

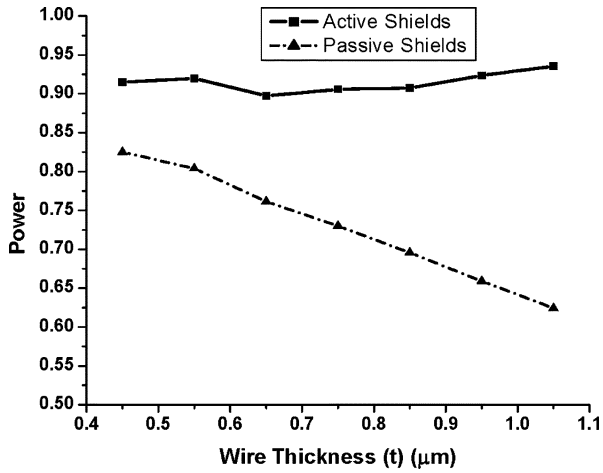

Fig. 15. Power consumption with active and passive shielding for a 7.5-mm copper wire with aggressors switching. The power consumption is normalized to the power consumption of the fat wire.

creased inter-metal spacing at the global layers so that active shielding can provide additional gains beyond those shown in Section IV-B. Simulation data indicates that for long lines (for which the interconnect delay becomes comparable to the gate delay) the optimal solution is obtained with wider signal 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 we investigated (this leads to minimum width shields). 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 setting the middle wire as wide as possible given the area limitations.

The optimal driver size distribution can be obtained with the delay model described in Section IV-A. Though the delay model is not accurate in terms of the delay improvements, its trends for optimal driver sizing are similar to simulation results. Fig. 16 shows the delay gains compared to passive shields for different technologies using both the analytical model and HSPICE simulations for the configuration with a 7.5-mm long copper wire with aggressors switching in the opposite direction. Table III shows the optimal size of the middle driver obtained through both methods. Though the model assumes a simplified effective coupling capacitance of twice the nominal between the active shield and aggressors, there is a close match between the predicted optimal driver sizes from the model and SPICE simulations. The rightmost column in Table III indicates the maximum difference in delays as a percentage of optimal delay (as obtained through simulations) if the optimal driver size predicted by the model is used. The differences are less than 1%. The predicted delays from the model can be further improved (at the cost of increased complexity and run-time) by using analytical models that calculate slew rates at the output of the driving gate rather than at the end of the line, for obtaining the effective

TABLE II
MAXIMUM NOISE (IN mV) PRODUCED BY AGGRESSOR SWITCHING AT THE FAR-END OF THE VICTIM WIRE

Configuration	Wire Thickness (t)					
	t = 0.65μm			t = 0.95μm		
	Passive	Active	Fat	Passive	Active	Fat
7.5mm copper wire	12	28	54	27	80	128
7.5mm copper wire with inter-repeater distance of 2.5mm	2	27	51	3	68	114

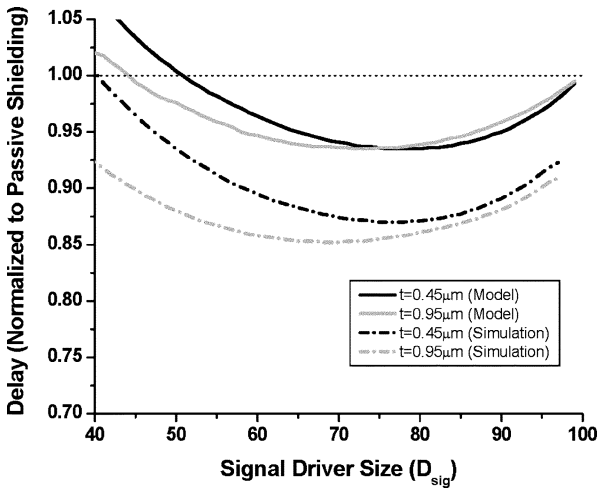

Fig. 16. Comparison of delay obtained from model and simulation with active shielding as signal driver size is varied for different wire thicknesses for the configuration with 7.5-mm copper wire and aggressors switching in the opposite direction. Delays are normalized to passive shielding delays.

TABLE III
COMPARISON OF OPTIMAL DRIVER SIZES FOR ACTIVE SHIELDS OBTAINED BY MODEL AND SPICE SIMULATIONS

Wire Thickness (μm)	Optimal signal driver size (X min. size)		Maximum difference in delays (%)
	Model	Simulation	
0.45	78X	75X - 79X	0.00
0.55	70X – 74X	68X – 73X	0.25
0.65	79X – 85X	78X – 84X	0.30
0.75	78X – 81X	73X – 82X	0.00
0.85	72X – 81X	70X – 77X	0.38
0.95	67X – 79X	67X – 71X	0.82
1.05	63X – 72X	58X – 68X	0.43

coupling capacitance between the signal wire and active shield. These models compute the effective load capacitance [24], [25] "seen" by the driving gate to compute the slew rate at the output of the driving gate. Using slew rates at the end of the line to compute effective capacitance is a pessimistic approach since the slew rates at the start of the wire are faster. On the other hand, computing effective capacitance with slew rates calculated at the driving end of the wire can be optimistic since the slew rates degrade as the signal travels down the wire and the skew between the active shield and signal wire increases toward the far end of the wire. Once the optimal wire sizing is determined, obtaining the optimal driver size distribution becomes possible using the analytical 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 shield 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 signal wire into the coupling capacitances and the delay improves (∼8%–10%) even when the useful Miller effect is negligible. Thus, an improvement in delay (though not by the Miller effect) can be achieved just by adding very weak inverters onto the shield wires. The fact that the effective cross-coupling capacitance depends on the driver resistance has been used in [26] for crosstalk noise estimation.

V. ACTIVE SHIELDING FOR INDUCTIVE WIRES

In Section III it was shown that opposite-phase active shielding provides higher performance and lower ringing than passive shielding for wire widths greater than 4 μm for the chosen global interconnect structure (which is typical for the 0.18 μm technology node [13]). Nets with such wire widths are usually global level clock nets and would, thus, be a primary application space for opposite-phase active shielding to be beneficial over the traditional shielding technique. Section V-A describes and compares detailed results for passive shielding and opposite-phase active shielding used for tuning clock nets to reduce ringing and achieve the desired skew and transition times. In Section V-B we discuss some fine tuning strategies that can be used to further optimize the passively shielded clock nets of Section V-A and show that even when these techniques are employed, the actively shielded clock nets provide a better tradeoff between performance and signal-integrity. Signal and wire sizing strategies for actively shielded clock nets are discussed in Section V-C. The optimized clock nets of Section V-A do not include the impact of nearby aggressors since global level global clock nets are well spaced from any aggressor nets. However, we analyze the impact on and of nearby signal nets on the optimized clock nets in Section V-D. The impact of process variation on the optimized clock nets is analyzed in Section V-E to demonstrate the feasibility of the approach at worst-case process corners. In Section V-F we show that active shielding provides a reasonable performance alternative to differential signaling while maintaining the simplicity of full-swing single-ended signaling.

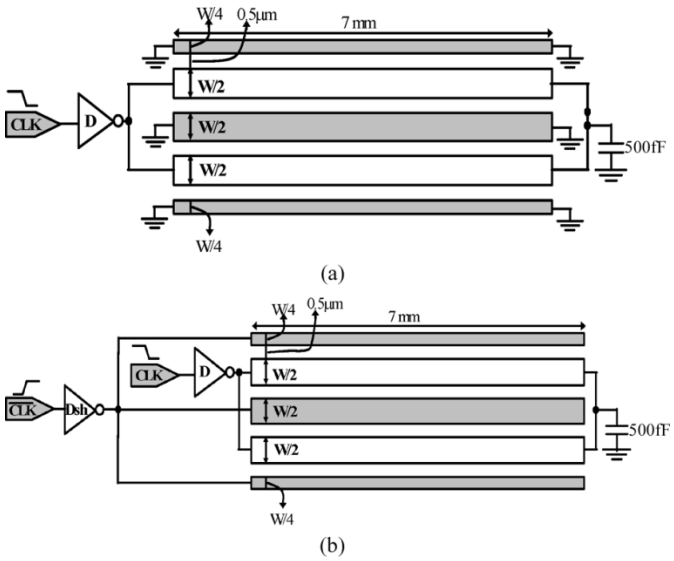

Fig. 17. (a) Passively shielded clock net with two fingers. (b) Actively shielded clock net with two fingers.

A. Clock Net Optimization

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/overshoot along with the associated ringing are undesirable from a signal integrity and 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, simply placing shields on both sides is not sufficient to reduce the inductive ringing to an acceptable level. In this case clock nets are split into two or more fingers with shields inserted between the fingers [9]. As the number of fingers increases, the current loops become smaller and the net becomes less inductive. Splitting the clock net into fingers comes at the expense of delay (leading to a higher likelihood of clock 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 III we expect better delay 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 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.

We apply the active shielding approach to four 7-mm-long clock net structures with signal widths (W) of 3, 5, 10, and 20 μm. Fig. 17 shows the setups for cases where the clock net has two fingers. Clock nets with more fingers have a similar topology to the one shown in Fig. 17. Certain details are not included in the figures for clarity. For reliable fabrication purposes, wires in actual designs cannot be wider than a certain limit. We have considered this by splitting any single metal wire wider than 7 μm into two or more wires of equal width with a spacing of 0.5 μm between these wires.[2] Clock nets with the four different signal widths (W) are optimized. 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 μm. The driver size (D) of the clock net is scaled up linearly with W. For active shielding, the active shield driver size ($D_{\rm sh}$) is swept to meet the constraints on the clock signal as described below. The complementary clock signal for the active shield driver is assumed to be generated by the drive chain leading up to the clock driver.[3]

The constraints for a particular target clock frequency are as follows.

1) Wire delay (measured between the 50% points from clock driver output to the end of the line) should be $\leq$20% of the cycle time. This constraint is important since delay or clock latency translates directly to clock skew when process variability is considered.
2) The slope (measured as 10%–90% delay at the end of the net) should be no greater than 25% of the cycle time.
3) Peak ringing, measured as the maximum deviation from 0 V after a single falling transition, should be less than 5% of VDD (90 mV).

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 increases and the splitting procedure is stopped when adding an extra finger yields 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 ($D_{\rm sh}$). 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 $D_{\rm sh}$ that meets the listed constraints is used. Table IV lists the results for the W values investigated. The constraints on the clock signal can be met for all values of W with active shielding while passive shielding can meet the constraint only for $W = 3$ μm. The results for $W = 3$ μm also support the analysis in Section III, where we found that active shielding for wire widths in the range of 1.5–4 μm in this process technology would not yield significant performance improvements to justify the added power consumption.

Fig. 18 represents the relative performance (delay and slope) of the best actively shielded clock setup compared to the best passively shielded clock net 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 μm are $\sim$14% lower with active shielding. The delay with active shielding is higher for $W = 20$ μm since the footprint of the optimal clock net with passive shields now occupies the entire area between two consecutive power supply lines with minimum spacing from them.

[2]The actual width beyond which a wire must be split, varies from one manufacturer to another. The limit can range from 4 to 12 μm.

[3]There is no constraint on the total input capacitance, since the effect of increased input capacitance to the clock net driver can be easily adjusted for by the drive chain that drives it.

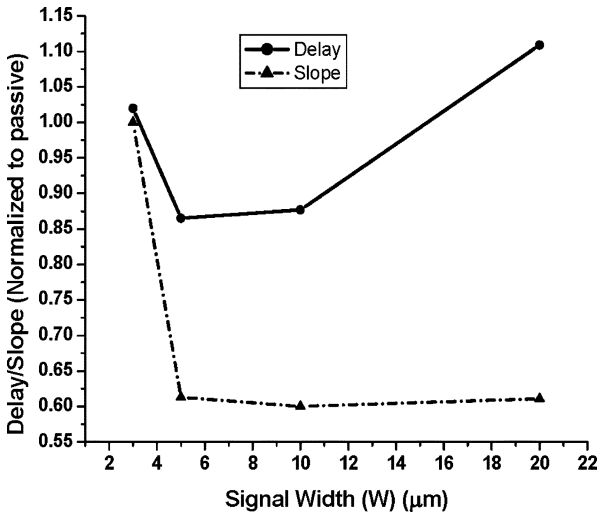

Fig. 18. Skew and slope of actively shielded clock net normalized to skew and slope of passively shielded clock net.

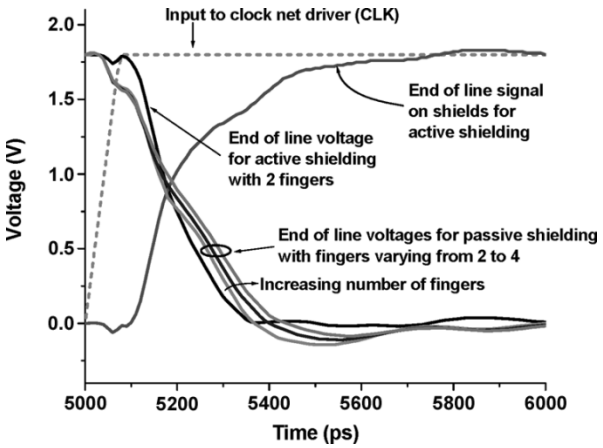

Fig. 19. Voltage waveforms at the end of a clock net with a signal width (W) of 10 μm.

The power supply lines act as additional shields for the passively shielded clock net, reducing the inductive impedance beyond what is achievable by the explicit shields. This effect reduces the delay on the passively shielded clock net but the use of inductive coupling to speed up transition times on the actively shielded clock is still clearly evident. Optimal active shielding resulted in up to 4.5X decrease in ringing for $W = 20$ μm. It is clear from the waveforms (for a single transition and periodic signal) in Figs. 19 and 20 that the signal obtained with active shields has faster transitions with highly suppressed ringing.

The opposite-phase 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 [27], the power consumption at the local clock level is at least an order of magnitude higher than clock distribution at the global level. Furthermore, if a comparison is made across different topologies in Table IV, the actively shielded configuration outperforms (primarily in terms of transition times) the passively shielded configuration for approximately

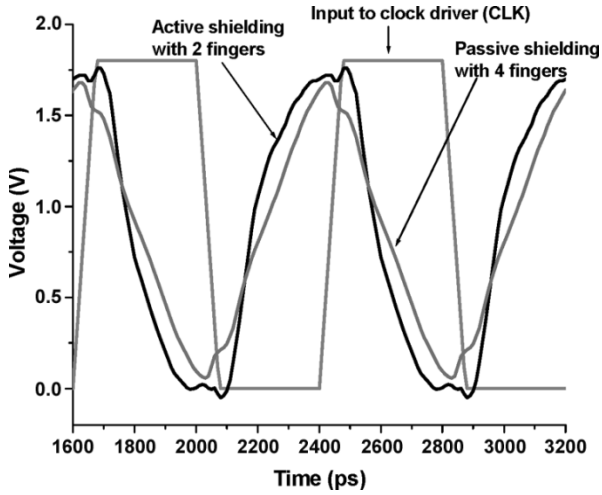

Fig. 20. Clock signal at the end of the clock net with a signal width (W) of 10 μm. The clock is operating at 1.25 GHz.

the same energy consumption. Relevant examples for such a comparison are between $W = 5$ μm/ActiveShields/1 finger and $W = 10$ μm/PassiveShields/4 fingers and between $W = 10$ μm/ActiveShields/1 finger and $W = 20$ μm/PassiveShields/2 fingers.

B. A Discussion on the Optimization of Passively Shielded Clock Nets

Ringing for the passively shielded clock nets could also be decreased by reducing the size of the clock driver (D). This results in even slower transition times. Only one passively shielded configuration ($W = 5$ μm/PassiveShields/1 finger) allowed for the reduction of the main clock driver size, since the only violation of the constraints in this case was the peak undershoot. A range of slower drivers allowed this configuration to meet the ringing constraint. However, the optimal actively shielded topology for $W = 5$ μm still resulted in 29% better slopes than the fastest downsized passively shielded case.

For the passively shielded clock nets in the comparisons in Section V-A, the shields were connected to the local grounds only at two points (either end). Increasing the number of such taps results in a reduction of the current loop and, hence, the inductive nature of the line. This reduction in the loop inductance results in lower ringing for the passively shielded nets with diminishing returns as the number of connections to local grounds are increased. Table V shows the impact of connecting the passive shields to the local ground every 500 μm for the clock net with $W = 10$ μm. The results of the actively shielded clock net are duplicated in this table for comparison purposes. Though the passively shielded clock net can now meet the ringing constraint with 3 fingers, it still cannot meet the slope constraint. Actively shielded clock nets result in 34% faster transition times when compared to the passively shielded configuration that meets the ringing constraint.

The addition of a repeater on the passively shielded clock nets can improve the transition time and undershoot trade-off, at the expense of higher delay. The higher the global clock delay, the greater is the impact from variations on the absolute clock skew. It is desirable to reduce global clock delays so as to minimize

TABLE IV
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_{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 = 300X f = 1GHz	1*	75X	165	244	0	25.7
T_D = 200ps T_{RF} = 250ps	1*	-	161	245	59	14.5
W = 5μm D = 500X f = 1GHz	1*	300X	141	166	56	43.6
T_D = 200ps T_{RF} = 250ps	1	-	144	216	118	21.4
	2	-	163	271	61	24.1
W = 10μm D = 1000X f = 1.25 GHz	1	1000X	121	133	106	76.1
T_D = 160ps T_{RF} = 200ps	2*	725X	128	179	38	81.5
	1	-	129	237	223	37.8
	2	-	123	261	143	38.2
	3	-	134	280	111	40.1
	4	-	146	299	87	41.8
W = 20μm D = 2000X f = 1.25 GHz	1	1900X	135	97	349	142.8
T_D = 160ps T_{RF} = 200ps	2	2000X	117	133	99	148.7
	3*	1650X	122	182	25	150.9
	1	-	141	264	332	73.6
	2	-	115	273	170	73.5
	3	-	112	291	127	71.8
	4	-	110	298	112	74.2

* ⇒ Clock constraints met

Driver Size 1X ⇒ $W_{NMOS} = 2 \bullet L_{min}$ $(W_{PMOS} = 2 \bullet W_{NMOS})$

TABLE V
RESULTS FOR ACTIVELY AND PASSIVELY SHIELDED CLOCK NETS, WITH PASSIVE SHIELDS CONNECTED TO LOCAL GROUND EVERY 500 μm. ROWS IN GREY INDICATE ACTIVE SHIELDING

Signal Width (W), Clock Driver Size (D) Target Clock frequency (f) Delay Constraint (T_D) Slope Constraint (T_{RF})	Number of fingers	Active shield Driver Size (Dsh)	Delay (ps)	Slope (ps)	Peak Under-shoot (mV)
W = 10μm D = 1000X f = 1.25 GHz	1	1000X	121	133	106
T_D = 160ps T_{RF} = 200ps	2*	725X	128	179	38
	1	-	140	188	195
	2	-	131	250	106
	3	-	134	272	86
	4	-	140	295	71

* ⇒ Clock constraints met

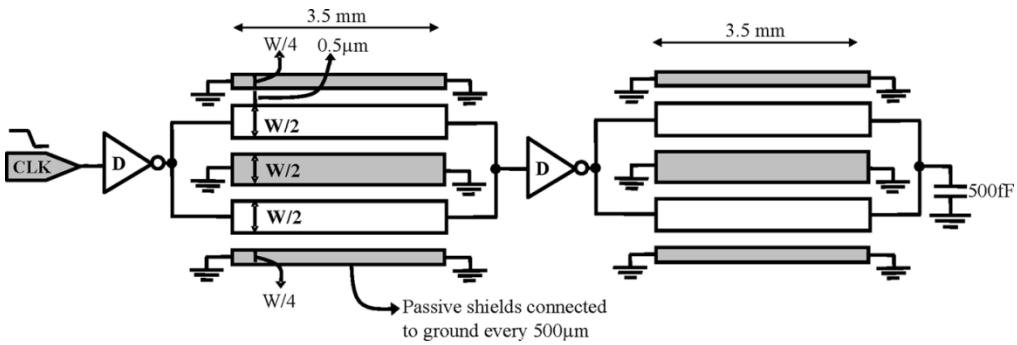

Fig. 21. Passively shielded clock net with a repeater and 2 fingers.

skew that results from variations in the clock loading, process, power supply or temperature. To assess the impact of repeater insertion on the passively shielded clock nets, a setup (shown in Fig. 21) for a repeated clock net with $W = 10$ μm was analyzed. The driver and repeater size (D) were varied to study the impact of repeater insertion on the passively shielded clock net performance and signal integrity. Table VI shows the results for the driver and repeater sizes that met the ringing constraint.

TABLE VI
RESULTS FOR PASSIVELY SHIELDED CLOCK NET ($W = 10\ \mu$m) WITH ONE REPEATER

Number of fingers	Driver/ Repeater Size (D)	Delay (ps)	Slope (ps)	Peak Undershoot (mV)	Energy per clock period (pJ)
1	220X	199	173	87	29.1
2	350X	179	141	82	35.4
3	525X	169	132	89	42.2
4	700X	170	133	90	51.8

Although, the passively shielded clock net now meets the slope and ringing constraints, the addition of a repeater increases the delay. With the addition of a repeater, the passively shielded clock nets could not meet the delay constraints that were used to optimize the clock nets of the previous sub-section.

The optimization of the clock nets has shown that even though the use of active shields for global clock nets results in larger driver area and power consumption, it can allow the designer to meet constraints for high-performance and reliability that cannot be met simultaneously with passive shielding.

C. Wire Sizing Issues

The optimized clock nets used shield widths that were equal to the signal width. In order to gain insight into the effect of varying shield widths and shield spacing on actively and passively shielded clock nets, we used the setup of Fig. 5 with relaxed constraints on shield width and inter-metal spacing. The total shield width ($W_{\rm sh}$) was varied from 0.5 to 3 times the signal width (W) and spacing was varied from 0.5 to 2 μm. The goal was to find an optimal spacing and shield size for passive and opposite-phase active shielding within this solution space to minimize peak undershoot. We then compared the performance of the two shielding schemes at their respective optimal points. The results indicate that peak undershoot is smallest with minimum spacing of 0.5 μm 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. 22(a)]. With active shielding there is a distinct optimal point, occurring when the total shield width is equal to the signal width ($W_{\rm sh} = 1.0 \bullet W$) for all signal widths greater than 4 μm [Fig. 22(b)]. Fig. 23 compares active and passive shielding at their respective optimized points in this solution space. The active shielding design has smaller delays, slopes, and ringing than the optimized passive shielding setup (within Region 3). 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. An optimal distribution of wire widths between the signal and shield nets (under an area constraint) for active shielding requires that the widths be equal with minimum spacing between the nets.

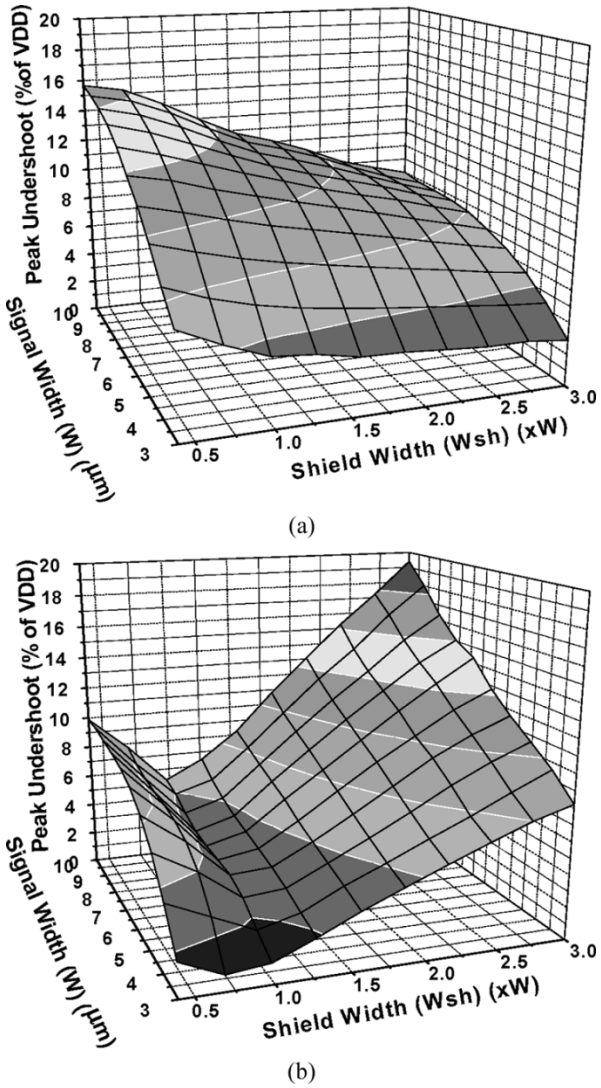

Fig. 22. Effect of varying W and $W_{\rm sh}$ (with $S = 0.5\ \mu$m) for the setups in Fig. 5 for (a) passive shielding and (b) active shielding.

D. Impact of Aggressors

Clock nets are typically well spaced from any neighboring aggressors not just to protect the clock nets from any noise but also to protect neighboring data or control signals from noise injected by the clock net which usually has sharper slew rates, from a voltage and current perspective. To study the impact of closely spaced aggressors on the optimized clock nets we placed aggressors as shown in Fig. 24 for clock nets with $W = 5\ \mu$m and 10 μm. Since the aggressors are not clock nets they have a smaller wire width and driver size than the clock net. The configurations that met the ringing constraint for the passively (with

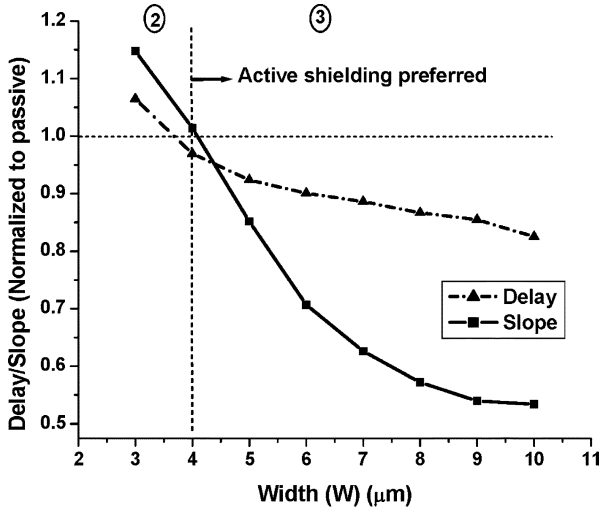

Fig. 23. (a) Delays and slopes at optimal undershoot points in Fig. 22. Delays and slopes for active shielding are normalized to passive shielding delays and slopes. (b) Peak undershoot at optimal undershoot points in Fig. 22.

ground taps every 500 μm) and actively shielded clock nets were compared for the impact of these aggressors. When the aggressors switch in the same direction as the clock signal, the passive and active shields are less effective as return paths and result in increased undershoots. They also increase transition times on the clock nets due to mutual inductive coupling. The results are tabulated in Table VII. Though the performance and undershoots of all the configurations degrades, the actively shielded nets still meet the constraints on the clock nets. Table VII also contains the peak noise injected on the clock net when only the aggressors are switching and vice-versa. The active shields provide a more resistive path to ground for any injected noise and hence, result in higher injected noise when the aggressors switch. When the clock net is switching, the active shields can reduce the amount of noise injected due to mutual inductive coupling onto closely spaced signal nets. However, the active shields also add capacitive noise onto the neighboring signal net. For the case where $W = 5$ μm the capacitive noise injected by the active shields is greater than the reduction in inductive noise and the actively shielded clock net introduced a higher noise peak on the signal net than the passively shielded clock net. For the more inductive clock net with $W = 10$ μm the two effects balance each other and the resultant noise is slightly less than with the passively shielded clock net. As the spacing of the signal nets from the clock net increases we can expect greater reduction in the noise injected by the (actively shielded) clock net onto the neighboring signal net since most of that noise will be through inductive coupling.

E. Process Variation Impact

In the clock net optimization process of Section V-A the clock signal and its complementary signal applied to the active shield driver have the same slew rate and zero skew/offset 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. 25) that generates the signals to drive the clock net and the active shields for the optimal configurations obtained above. 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 $\pm 5\%$. This value is small compared to worst-case wafer-to-wafer and lot-to-lot variation but is intended to be representative of spatially proximate intra-die fluctuations [28]. All combinations of extreme variations in the channel lengths of individual inverters were simulated in order to obtain the worst-case positive deviations in slopes, delays, and ringing. The channel length of each of the drive chain inverters, (whose drive sizes are labeled D, D2, Dsh, Dsh2, and Dsh3 in Fig. 25) was allowed to vary independently to obtain the worst-case corner in this analysis. Simulations showed that ringing and delays were relatively insensitive to these variations—worst case ringing was kept under 5% of VDD and the maximum deviation in delay was 5%. Transition times were more sensitive to process variation. Assuming worst case intra-die variability reduced the gains (with respect to worst-case corner of passive shielding) from $\sim$40% for $W = 5$, 10, and 20 μm to 35%, 34%, and 34%, respectively.

To assess the impact of inter-die variation, the channel lengths of the drive chain inverters tracked each other and the maximum allowed deviation was $\pm 10\%$ from the nominal. The transition times were insensitive to inter-die variations ($<3\%$ variation), while the delays had a maximum variation of 6.5% among all the three (actively shielded) test cases. The transition times are more sensitive to the offset between the complementary clock signals and, therefore, showed greater variations for intra-die process variations.

F. Differential Signaling Comparison

Since the opposite-phase active shielding approach is similar to differential signaling we compare the active shielding and

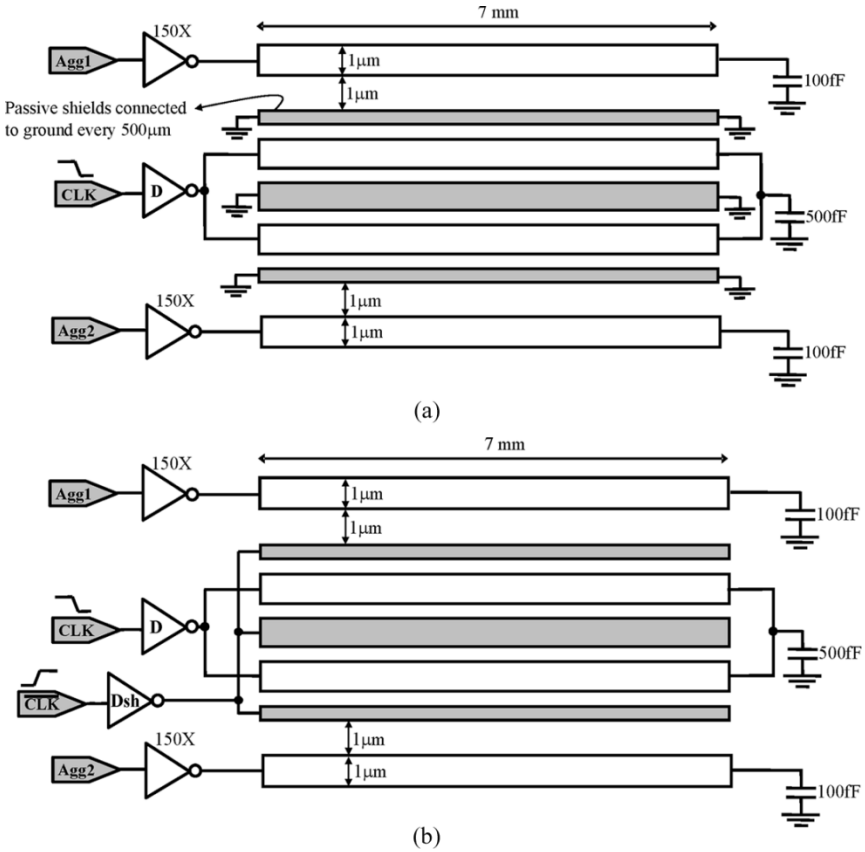

Fig. 24. Schematic of the setup used for studying the impact of aggressors on (a) passively shielded clock net with two fingers and (b) actively shielded clock net with two fingers.

TABLE VII
RESULTS FOR ACTIVELY AND PASSIVELY SHIELDED CLOCK NETS IN THE PRESENCE OF AGGRESSORS. ROWS IN GREY INDICATE ACTIVE SHIELDING. NUMBERS IN PARANTHESES INDICATE CHANGE WITH RESPECT TO THE CONFIGURATION WITHOUT AGGRESSORS

Signal Width (W), Clock Driver Size (D) Target Clock frequency (f) Delay Constraint (T_D) Slope Constraint (T_{RF})	Number of fingers	Active shield Driver Size (Dsh)	Delay (ps)	Slope (ps)	Peak Under-shoot (mV)	Peak noise on CLK (mV)	Peak noise on aggressor (mV)
W = 5μm D = 500X f = 1GHz T_D = 200ps T_{RF} = 250ps	1*	300X	163 (+15.6%)	188 (+13.3%)	78	191	219
	2	-	157 (+5.4%)	317 (+15.3%)	89	106	141
W = 10μm D = 1000X f=1.25 GHz T_D = 160ps T_{RF} = 200ps	2*	725X	141 (+10.1%)	200 (+11.7%)	51	109	207
	3	-	134 (0%)	298 (+9.6%)	115	96	219

* ⇒ Clock constraints met

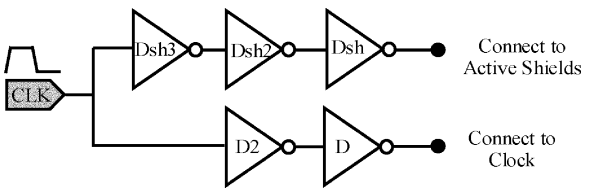

Fig. 25. Drive chain structure used to analyze the impact of process variation on actively shielded clock nets. Sizes D and D_{sh} are the same as listed in Table IV.

differential signaling schemes in this section. We used voltage ramps to simulate the drivers to avoid biasing the results toward a particular driver/receiver architecture. The setup for differential signaling (Fig. 26) 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 • VDD, while the current was drawn from the VDD power supply (which is similar to the operation of an intermediate swing driver). 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. 27) are compared using the full-swing output of the receiver. The ideal receiver does not drive any load and is primarily used for com-

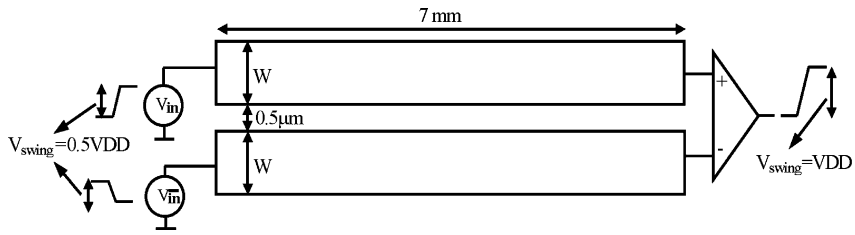

Fig. 26. Setup for differential signaling comparison. Total metal width is the same as the setup in Fig. 5 (when $W = \mathrm{Wsh}$). The ideal receiver at the end of the line converts low-swing differential signals to full-swing single-ended signals.

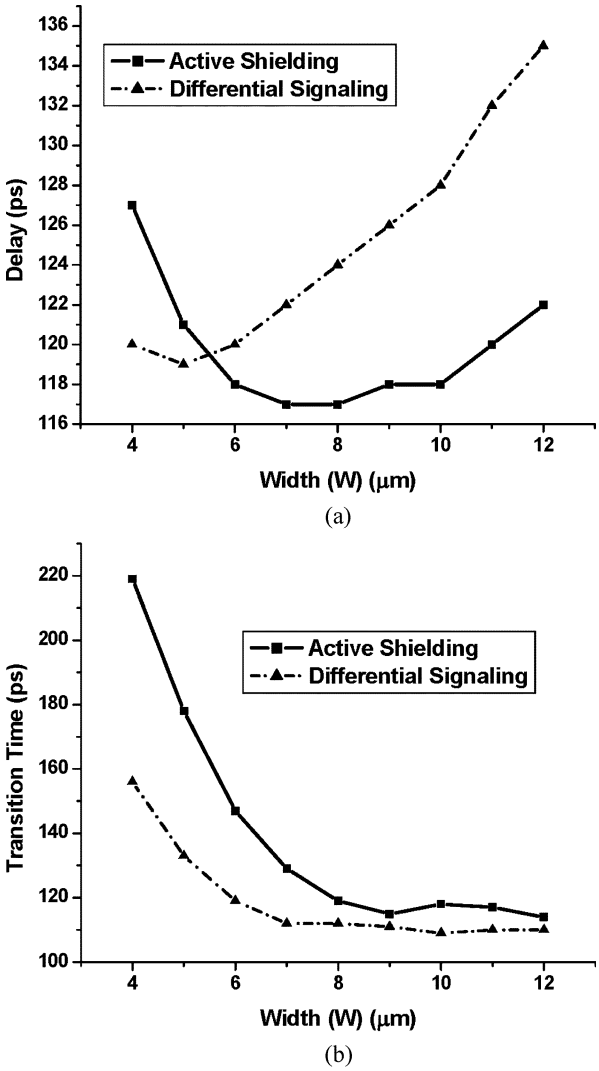

Fig. 27. (a) Delays and (b) transition times with active shielding setup of Fig. 5(b) and differential signaling setup of Fig. 26.

paring the effective slew rate and delay that would be seen by an actual differential receiver to one that would be seen by a single ended receiver (for the active shielding setup). The delays with active shielding are smaller when $W \geq 6\ \mu$m while the transition times are faster with differential signaling for $W < 8\ \mu$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 • VDD (with the power supply as VDD), as opposed to full swing for active shielding. In general, active shielding results in performance that is much better than passive shielding but not clearly better than differential signaling. The advantage of active shielding lies in much simpler driver and receiver circuits (i.e., standard CMOS inverters) than differential signaling, providing a reasonable trade-off between design complexity and performance.

VI. Conclusion

The concepts of in-phase and opposite-phase active shielding were proposed for *RC* and inductive wires, respectively. 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. Simulation results show that for *RC* wires, converting a wide wire to an actively shielded (in-phase) wire results in lower delays, faster signal slopes, lower power consumption, and better noise immunity. In-phase actively shielded wires were also 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 analytical model was developed to obtain the optimal driver size distribution for in-phase actively shielded wires. The opposite-phase active shielding concept was used to optimize clock nets and resulted in up to 40% reductions in transition times. It also resulted in much lower ringing (up to a 4.5× reduction). Opposite-phase active shielding was shown to achieve considerable gains even when worst-case process corners were assumed. Optimal wire sizing rules for opposite-phase active shielding were discussed. We also point out that active shielding consumes extra power (as compared to passive shielding) and should be used when the performance gains justify the increase in power such as global clock distribution.

References

[1] D. Sylvester *et al.*, "Interconnect scaling: signal integrity and performance in future high-speed CMOS designs," in *Symp. on VLSI TechnologyDig. Tech. Papers*, 1998, pp. 42–43.

[2] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," *Proc. IEEE*, vol. 89, pp. 490–504, 2001.

[3] M. H. Chowdhury *et al.*, "Performance analysis of deep sub micron VLSI circuits in the presence of self and mutual inductance," in *Proc. Int. Symp. Circuits and Systems (ISCAS)*, 2002, pp. 197–200.

[4] L. He and K. M. Lepak, "Simultaneous shield insertion and net ordering for capacitive and inductive coupling minimization," in *Proc. Int. Symp. Physical Design (ISPD)*, 2000, pp. 55–60.

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

[6] S. Khatri *et al.*, "A novel VLSI layout fabric for deep sub-micron applications," in *Proc. Design Automation Conf. (DAC)*, 1999, pp. 491–496.
[7] A. B. Kahng, S. Muddu, and E. Sarto, "On switch factor based analysis of coupled RC interconnects," in *Proc. Design Automation Conf. (DAC)*, 2000, pp. 79–84.
[8] S. Morton, "On-chip signaling," presented at the Int. Solid-State Circuits Conf. (ISSCC), San Francisco, CA, 2002.
[9] Y. Massoud *et al.*, "Layout techniques for minimizing on-chip interconnect self inductance," in *Proc. Design Automation Conf. (DAC)*, 1998, pp. 566–571.
[10] S. van Dijk and D. Hély, "Reduction of interconnect delay by exploiting cross-talk," in *Proc. Eur. Solid State Circuits Conf.*, 2001, pp. 316–319.
[11] R. Hossain, F. Viglione, and M. Cavalli, "Designing fast on-chip interconnects for deep submicrometer technologies," *IEEE Trans. VLSI Syst.*, vol. 11, pp. 276–280, 2003.
[12] Y. Massoud *et al.*, "Modeling and analysis of differential signaling for minimizing inductive cross-talk," in *Proc. Design Automation Conf. (DAC)*, 2001, pp. 804–809.
[13] International Technology Roadmap for Semiconductors. (2000) ITRS. [Online]. Available: http://public.itrs.net
[14] M. Kamon, M. Tsuk, and J. White, "FastHenry: A multipole-accelerated 3-D inductance extraction program," *IEEE Trans. Microwave Theory Tech.*, vol. 42, pp. 1750–1758, Sept. 1999.
[15] Synopsys, Inc. (2002). [Online] Available: http: //www.synopsys.com/products/avmrg/raphael_ds.html
[16] Synopsys, Inc. (2002). [Online] Available: http://www.synopsys.com/products/avmrg/hspice_ds.html
[17] J. Cong *et al.*, "Interconnect sizing and spacing with consideration of coupling capacitance," *IEEE Trans. Computer-Aided Design*, vol. 20, pp. 1164–1169, 2001.
[18] L. P. P. P. van Ginneken, "Buffer placement in distributed RC-tree networks for minimal Elmore delay," in *Proc. Int. Symp. Circuits and Systems*, vol. 2, 1990, pp. 865–868.
[19] C. J. Alpert, A. Devgan, and S. T. Quay, "Buffer insertion for noise and delay optimization," *IEEE Trans. Computer-Aided Design*, vol. 18, pp. 1633–1645, 1999.
[20] V. Adler and E. Friedman, "Repeater design to reduce delay and power in resistive interconnect," *IEEE Trans. Circuits Syst. II*, vol. 45, pp. 607–616, May 1998.
[21] Y. Ismail and E. Friedman, "Repeater insertion in tree structured inductive interconnect," *IEEE Trans. Circuits Syst. II*, vol. 48, pp. 471–481, May 2001.
[22] H. B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*. Reading, MA: Addison Wesley, 1990.
[23] S. Thompson *et al.*, "An enhanced 130 nm generation logic technology featuring 60 nm transistors optimized for high performance and low power at 0.7–1.4 V," in *Proc. Int. Electron Devices Meeting (IEDM)*, 2001, pp. 11.6.1–11.6.4.
[24] F. Dartu, N. Menezes, and L. T. Pileggi, "Performance computation for precharacterized CMOS gates with RC loads," *IEEE Trans. Computer-Aided Design*, vol. 15, pp. 544–553, May 1996.
[25] A. B. Kahng and S. Muddu, "New efficient algorithms for computing effective capacitance," in *Proc. Int. Symp. Physical Design*, Apr. 1998, pp. 147–151.
[26] L. Ding, P. Mazumder, and D. Blaauw, "Crosstalk noise estimation using effective coupling capacitance," in *Proc. Int. Symp. Circuits and Systems*, 2002, pp. V645–V648.
[27] P. Restle *et al.*, "A clock distribution network for microprocessors," *IEEE J. Solid-State Circuits*, vol. 36, pp. 792–799, May 2001.
[28] International Technology Roadmap for Semiconductors. (2001) ITRS. [Online]. Available: http://public.itrs.net

Himanshu Kaul (SM'00) received the B.Eng. (Hons.) degree in electrical and electronics engineering from the Birla Institute of Technology and Science, Pilani, India, in 2000. He is currently working toward the Ph.D. degree at the University of Michigan, Ann Arbor.

His research interests include on-chip signaling techniques and low-power and high-performance circuit design.

Dennis Sylvester (S'95–M'00) received the B.S. degree in electrical engineering (*summa cum laude*) from the University of Michigan, Ann Arbor, in 1995, and the M.S. and Ph.D. degrees in electrical engineering from University of California, Berkeley, in 1997 and 1999, respectively. His dissertation research was recognized with the 2000 David J. Sakrison Memorial Prize as the most outstanding research in the UC-Berkeley EECS Department.

He worked at Hewlett-Packard Laboratories, Palo Alto, CA, from 1996 to 1998. After working as a Senior R&D Engineer in the Advanced Technology Group of Synopsys, Mountain View, CA, he is now an Assistant Professor of Electrical Engineering at the University of Michigan. He has published numerous articles in his field of research, which includes the modeling, characterization, and analysis of on-chip interconnect, low-power circuit design and design automation techniques, and variability-aware circuit approaches.

Dr. Sylvester received an NSF CAREER award, the 2000 Beatrice Winner Award at ISSCC, two outstanding research presentation awards from the Semiconductor Research Corporation, and a Best Student Paper Award at the International Semiconductor Device Research Symposium. He is the recipient of the 2003 ACM SIGDA Outstanding New Faculty Award and the 1938E Award for teaching and mentoring, which is the highest award given to a junior faculty in the Michigan College of Engineering. He has served on the technical program committee of numerous design automation and circuit design conferences and was general chair for the 2003 ACM/IEEE System-Level Interconnect Prediction (SLIP) Workshop. He is an Associate Editor for IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS and is General Chair for the 2005 ACM/IEEE Workshop on Timing Issues in the Synthesis and Specification of Digital Systems (TAU). In addition, he helped define the circuit and physical design roadmap as a member of the International Technology Roadmap for Semiconductors (ITRS) U.S. Design Technology Working Group from 2001 to 2003. He is a member of IEEE, Association for Computing Machinery (ACM), American Society of Engineering Education, and Eta Kappa Nu.

David Blaauw (M'01) received the B.S. degree in physics and computer science from Duke University, Durham, NC, in 1986, the M.S. and Ph.D. degrees in computer science from the University of Illinois, Urbana, in 1988 and 1991, respectively.

He worked at the Engineering Accelerator Technology Division, IBM Corporation, Endicott, NY, as a Development Staff Member, until August 1993. From 1993 till August 2001, he worked for Motorola, Inc., Austin, TX, where he was the Manager of the High Performance Design Technology group. Since August 2001, he has been on the faculty at the University of Michigan as an Associate Professor. His work has focused on VLSI design and CAD with particular emphasis on circuit analysis and optimization problems for high performance and low power designs.

Dr. Blaauw was the Technical Program Chair and General Chair for the International Symposium on Low Power Electronic and Design in 1999 and 2000, respectively, and was the Technical Program Co-Chair and member of the Executive Committee the ACM/IEEE Design Automation Conference in 2000 and 2001.