Dynamic Clamping: On-Chip Dynamic Shielding and Termination for High-Speed RLC Buses

Kanak Agarwal, Dennis Sylvester, and David Blaauw
University of Michigan, Ann Arbor, MI 48109, USA
{agarwalk,dmcs,blaauw}@umich.edu

Abstract

This paper presents a novel approach called dynamic clamping for minimizing crosstalk noise and inductive effects in global buses. A simple circuit is shown that can be used to dynamically shield and terminate high-speed RLC buses. Unlike traditional passive shielding and parallel termination, dynamic clamping has no area overhead and no static power dissipation. Dynamic clamping enables significant reductions in noise (~35%) and inductive overshoot (~90%) with a small delay penalty (~10%). We also propose using Bus-Invert coding with our approach as dynamic clamping is seen to give excellent results for low to moderate bus activity.

1. Introduction

Trends in high performance digital ICs include higher clocking frequencies, longer and wider global interconnects, tighter metal pitches and faster signal rise times – these conspire to make on-chip signal integrity a primary issue. The problem is exacerbated by the increasing significance of on-chip interconnect inductance, which can lead to a number of effects such as signal overshoot, ringing, reflections, and inductive crosstalk. To control these signal integrity problems, novel design techniques should be developed which increase the robustness of the design without sacrificing the performance.

The most common technique used to reduce coupling noise and high-frequency inductive effects is *shielding*, where dedicated lines connected to power/ground are inserted between signal wires, particularly in global buses. These dedicated lines isolate signal wires from adjacent lines and can significantly reduce coupling effects. Shields also help in reducing inductance by providing good nearby current return paths. However, it has been shown [1, 2] that shielding is not completely effective for self and mutual inductance due to the long range of current return paths. Also, shields do not eliminate reflections, complicate routing, and can cause significant routing area overhead (e.g., [3] proposes using one dedicated shield for every two signal wires).

Another technique that can be used to suppress reflections and other inductive effects is *termination*. If the load impedance of a transmission line does not match its characteristic impedance, the signal is reflected at the load end. These reflections translate into undesirable effects such as ringing, overshoot, and stair-stepping [6]. These effects can be eliminated by properly terminating the line at the far-end. Figure 1 shows a basic far-end parallel termination scheme. If R matches Z_0, then the reflection coefficient at the load end becomes zero and reflections can be eliminated [4]. However, parallel termination is not used for on-chip interconnects because it results in high static power dissipation in the terminating resistor. The other disadvantage of using far- end termination is that it results in worsened signal slew rates and increased delay.

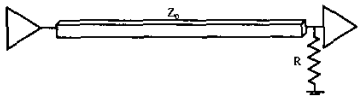

Fig 1. Far-end Termination

In this paper, we propose a technique called *dynamic clamping* that provides both shielding and termination without the area and static power dissipation overhead associated with these standard practices. Shielding is obtained by dynamically creating a low impedance path between power/ground and a quiet line at the sink point of the line. Termination is obtained in a similar way by dynamically adding a termination resistor at the far-end of the line. Both shielding and termination require adding a resistance at the far-end of the line; hence they can be obtained by using the same circuit (to be discussed in Section 2). Also since these resistances are added dynamically, static power dissipation and area penalties are eliminated.

The paper is organized as follows. We begin by describing our dynamic clamping approach in the following section. Section 3 discusses experimental results. In Section 4, we propose using Bus-Invert coding with our approach before concluding in Section 5.

2. Dynamic Clamping Circuit

In this section, we propose a simple circuit that can be used in RLC buses for reducing noise and inductive effects. This circuit performs a function called dynamic clamping and is shown in Figure 2. It consists of an inverter driving a pull-up and a pull-down transistor as shown in the figure. This circuit is added at the far-end of all the lines in a bus.

First we discuss how this circuit can be effective in providing shielding capability. In a bus, generally many lines do not switch during a given clock cycle. If these non-switching lines are tied to the power rails through a low resistance path, they can then act as shield wires. In the circuit shown in Figure 2, when a line is quiet the inverter

0-7803-8160-2/03/$17.00 ©2003 IEEE.

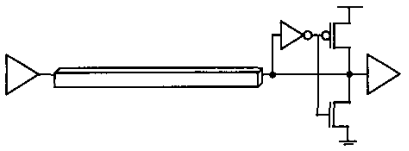

Fig 2. Dynamic Clamping Circuit

enables either the pull-up or the pull-down transistor, creating a low resistance path between the line and the power supply. All quiet lines tied to the power rail in this way effectively act as shields for switching lines. If the size of the clamping transistor is increased, the effective resistance decreases and the shielding capability of the quiet lines approaches that provided by standard passive shields. Another advantage of clamping quiet lines in this way is that it results in significant reduction in capacitive and inductive coupling noise because the quite line is tied strongly to the power supply. Dynamic clamping of quiet lines also reduces the self inductance seen by the switching lines in the bus by providing nearby low resistance current return paths.

The same circuit can be effective in suppressing reflections in switching lines by providing optimal termination as seen in Figure 1. However, unlike the traditional far-end termination of Figure 1, this approach has no static power dissipation because the termination resistor is added dynamically. When a line is switching, based on the output of the inverter one of the two transistors is active – if its resistance is matched to the characteristic impedance of the line, then the multiple reflections can be eliminated. Static power dissipation is eliminated because the state of the line is sensed and accordingly either the pull-up PMOS or pull-down NMOS is enabled, suppressing any direct path between power and ground. However, dynamic termination results in a delay penalty because of the response delay of the inverter. For a line switching low to high, the pull-down clamping transistor stays on for some time during transition resulting in an increased delay. The delay overhead depends on the sizes of the pull up and pull down transistors and they should be carefully chosen to balance the trade-off between delay overhead and noise and overshoot reduction due to clamping.

Due to the delay overhead associated with this approach, the dynamic clamping technique is effective only in high-speed buses driven by very fast drivers. For lines driven by weak drivers, the delay penalty due to dynamic clamping can be quite significant. Also for such lines, noise and overshoots are not serious concerns and hence dynamic clamping need not be used. However, for high-speed buses, on-chip interconnects act as transmission lines and the propagation delay is less sensitive to the load at the far-end of the line. In these cases, the propagation delay is dominated by the transmission line time of flight delay and hence the delay penalty incurred due to dynamic clamping is small. Also in these cases, effects like ringing and noise are heightened and dynamic clamping technique can be very effective in controlling these effects.

3. Experimental Results

In this section, we show how dynamic clamping can be effective in improving signal integrity in high-speed buses.

3.1 Experimental Setup

The experimental setup we used in our simulations is shown in Figure 3. We used an 8-bit bus topology. The line length, width, and spacing were chosen to be 2mm, 1.2μm, and 0.6μm respectively. Line parasitics were extracted from commercial extraction tool Raphael. A two-dimensional power grid of 50μm spacing and 10μm width was used in extraction. The driver size was 75X in a 0.13μm technology[1].

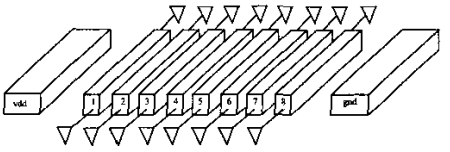

Fig 3. Experimental Setup

3.2 Sizing Pull-up/Pull-down Transistors

As mentioned in Section 2, the sizes of the pull-up and pull-down transistors are important in the trade-off between noise and overshoot reduction and delay penalty due to dynamic clamping of the lines. If the sizes of the pull-up and pull-down transistors are increased, the shielding action provided by clamping improves but it also increases the delay penalty in the switching lines. To test this trade-off, we considered the experimental setup of Figure 3. All the lines in the bus contained a dynamic clamping circuit at the far-end of the line. The dynamic clamping circuit contained a minimum sized inverter, pull-down NMOS, and a complementary PMOS (twice the size of NMOS) as shown in Figure 2. We considered a switching pattern where lines 1, 3, 5, and 7 were switching and lines 2, 4, 6, and 8 were quiet. Figure 4 shows the reduction in noise and overshoot (measured at the far-end of lines 4 and 5 respectively) as a function of pull-down transistor size. The figure also shows the delay penalty in the switching lines due to clamping.

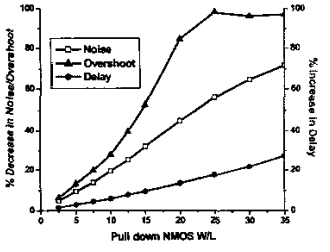

Fig 4. Noise/overshoot reduction and delay increase vs. device sizing with dynamic clamping

[1] Driver size 75X means the NMOS width in the inverter is 75 times the minimum width (=2*L_{min}=0.26μ). PMOS is twice as wide as NMOS.

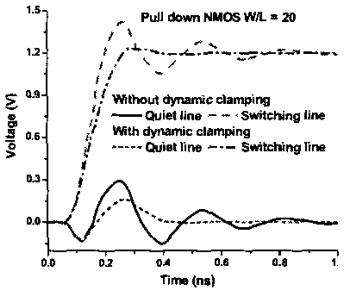

Fig 5. Waveform comparison with and without clamping

It is clear from the figure that by choosing pull down transistor size carefully, substantial reductions in noise and overshoot can be obtained without significant delay penalty. For example, for a pull-down NMOS of W/L=20, the overshoot is nearly eliminated (85% reduction) and noise is reduced by 45% with only 13% delay penalty. Figure 5 compares the waveforms with and without dynamic clamping for this case.

3.3 Sizing Sense Inverter

Next we investigate the effect of sizing the inverter in the dynamic clamping circuit. It may seem that this inverter should be made as fast as possible to reduce delay overhead in the switching lines. For fast inverters, the clamping transistor opposing the transition will get disabled faster and hence the delay penalty will be reduced. However, as discussed in Section 2, dynamic clamping is used for interconnects that behave like transmission lines and whose delays are not highly sensitive to the load at the far-end. In such cases, a minimum size inverter should be used because it gives the best results in reducing inductive effects without affecting delay significantly. This is because a weak inverter helps in reducing the loop inductance of the line by providing good current return path during transition. For example, if the line is switching from low to high, the weak inverter switches from high to low slowly and hence it keeps the pull-down NMOS enabled during most of the transition. This provides a good current return path at the far-end, thus reducing the loop inductance of the line. Figure 6 shows the waveforms with dynamic clamping for two different inverter sizes. It is clear form the figure that weak inverter gives better results in reducing noise and overshoot.

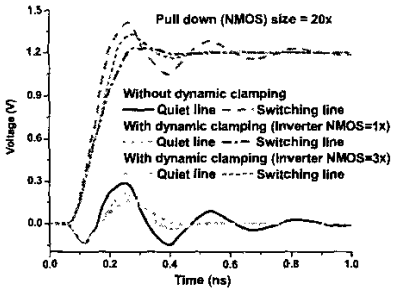

Fig 6. Waveform comparison for different inverter sizes

Table I. Comparison with driver downsizing

	Downsizing	Clamping
% Decrease in noise	24.2%	45.1%
% Decrease in overshoot	71%	85.1%

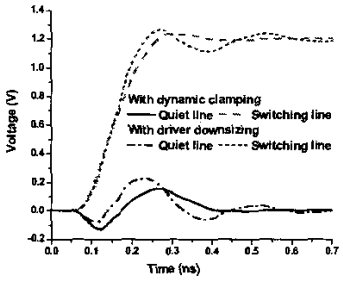

Fig 7. Waveform comparison of dynamic clamping with driver downsizing

3.4 Comparison with Driver Downsizing

Inductive effects and crosstalk noise are significant primarily in buses driven by fast drivers. Hence, downsizing drivers can improve signal integrity, but it increases line delay making buses slower. Since dynamic clamping also has a delay penalty, we compare our results with the simple downsizing of the driver. For the same setup and switching pattern as discussed earlier, we downsize the drivers such that the delay increase due to downsizing matches the delay penalty from dynamic clamping. Table I compares the reduction in noise and overshoot by dynamic clamping and driver downsizing. It is clear from the table that dynamic clamping is more effective than driver downsizing for both noise as well as overshoot. Figure 7 compares the waveforms for dynamic clamping and driver downsizing.

3.4 Comparison with Passive Shields

Now we compare our results with a traditional passive shielding approach. For the setup in Figure 3, we apply random input switching patterns (with between one and seven lines switching) and measure worst-case overshoot and noise in all cases. Table II shows results of this comparison. For the traditional shielding approach, shielding frequencies of eight, four, and two are considered (shield frequency of two indicates a shield every two wires). Shield frequencies of two, four, and eight have 50%, 25% and 12.5% routing area overhead respectively. Table II shows results for two cases – when $\leq 1/2$ the lines are switching and when $> 1/2$ the lines are switching. Dynamic clamping gives better results when no more than half the lines in the bus are switching. This is because dynamic clamping causes each quiet line to act as a shield and hence better results are obtained with higher number of quiet lines. The table also shows that when $\leq 1/2$ the lines are switching, the noise reduction from dynamic clamping is comparable to that obtained with a shield frequency of four. For the same case, dynamic clamping is more effective at reducing overshoot. This is because unlike clamping, shielding does not provide termination and hence it does not eliminate multiple reflections.

Table II. Comparison with passive shields

Less than or equal to 4 lines switching				
Shield Frequency	**Two**	**Four**	**Eight**	**Dynamic**
Avg. % decrease in noise	48.7	33.6	21	34.3
Avg. % decrease in overshoot	84.7	80.3	67.6	90.6
More than 4 lines switching				
Shield Frequency	**Two**	**Four**	**Eight**	**Dynamic**
Avg. % decrease in noise	57.4	42.7	24.6	26.7
Avg. % decrease in overshoot	82.2	69.7	50.3	62.7

3.5 Dynamic Shielding with no Delay Penalty

Dynamic clamping concept can be used without any delay overhead if the switching pattern is known *a priori*. If we know that certain lines will not switch during a clock cycle, then we can clamp these lines while disabling the clamping transistors of the switching lines. A circuit similar to the dynamic clamping (but with an external enable) can be used and is shown in Figure 8. With this configuration, we tie only the quiet lines to the power rails and hence there is no delay penalty in the switching lines. We call this technique dynamic shielding because it provides only shielding as compared to the dynamic clamping that provides both shielding and termination. Figure 9 shows the reduction in noise and overshoot as a function of pull down transistor size due to dynamic shielding. Dynamic shielding approach is as effective as dynamic clamping for reducing coupling noise. It also performs well in reducing inductive overshoot by providing good nearby current return paths. However, with this approach, the reduction in ringing is not as significant as in dynamic clamping because the switching lines are not terminated optimally.

4. Dynamic Clamping With Bus Encoding

It was shown in Section 3.4 that performance of dynamic clamping improves significantly when less number of lines are switching. In this section, we propose that dynamic clamping can be used with a bus activity-lowering scheme for best results. Lowering bus activity is also useful from power perspective and hence various techniques have been developed for minimizing switching in a bus. One such technique is *Bus-Invert* coding as described in [5]. This scheme proposes a method to ensure that the maximum number of lines switching during a clock cycle is not more than half the bus width. This ensures that, while using dynamic clamping approach, at least half the lines in the bus will act as shields to provide good noise immunity and improved current return paths. For the example in Table II, when less than or equal to half the lines were switching, then noise reduction due to dynamic clamping was comparable to that obtained with a shield frequency of four. Hence, by using dynamic clamping with the proposed bus-encoding scheme, the routing area can be reduced without violating noise and overshoot constraints.

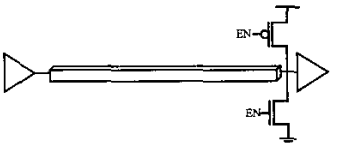

Fig 8. Dynamic shielding with an external enable

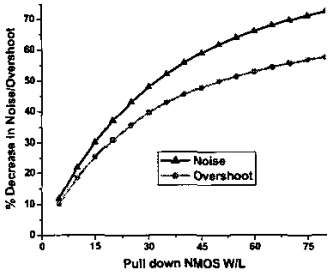

Fig 9. Noise/overshoot reduction vs. device sizing with dynamic shielding

5. Conclusion

We described a technique called dynamic clamping that is effective in reducing noise and inductive effects in high-speed RLC global buses. The technique provides shielding and termination capability without area and static power dissipation overhead. The results show that with careful sizing of clamping transistors, a significant reduction in noise (~35%) and inductive overshoot (~90%) can be obtained with small delay penalty (~10%). Using this technique with Bus-Invert coding can be very effective in reducing shield frequency.

Acknowledgments

The authors would like to acknowledge helpful discussions with Yu (Kevin) Cao of UC-Berkeley regarding the concept and implementation of dynamic shielding.

References

[1] Y. Massoud, J. Kawa, D. Macmillen and J. White, "Modeling and analysis of differential signaling for minimizing inductive crosstalk", *DAC*, pp. 804-809, 2001

[2] K. Lepak, I. Luwandi and L. He, "Simultaneous shield insertion and net ordering under explicit RLC noise constraint", *DAC*, pp. 199-202, 2001

[3] S. Morton, "On-chip Signaling", *ISSCC*, 2002

[4] B. Young, *Digital Signal Integrity*, Prentice Hall, pp. 394-398

[5] M. Stan and W. Burleson, "Bus-invert coding for low power I/O", *IEEE Trans. on VLSI*, vol. 3, 1995, pp. 49-58

[6] H.B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*, Reading, MA, Addison-Wesley, 1990.