Effects of Specialized Clock Routing on Clock Tree Timing, Signal Integrity, and Routing Congestion

Jesse Craig
IBM Systems & Technology Group
jecraig@us.ibm.com

Denise Powell
Synopsys, Inc.
dpowell@synopsys.com

ABSTRACT

Signal integrity issues and their interaction with clock trees are a growing concern in deep sub-micron VLSI designs. These issues can increase skew and latency on a clock tree, generate false switching due to coupling noise, or cause pulse-width timing failures. The clock tree also acts as a strong aggressor against other signal and clock wires. Traditionally, physical designers have dealt with these issues by isolating the clock wires or shielding them with grounded metal—methods not well suited to hard IP designs where routing congestion is often extremely high and the physical size of the hard IP can decide its attractiveness to the market. This paper examines and contrasts the traditional clock routing approaches along with techniques well suited to hard IP design. The effects of each technique on skew, latency, routing congestion, and capacitive coupling are explored using a popular embedded processor design hardened with the Synopsys tool suite.

Table of Contents

Table of Figures

1.0 Introduction

With the scaling to deep sub-micron processes, 0.13µm and smaller, capacitive coupling has become a major concern. This is not only because it adds yet another design challenge to be overcome, but because its effects are so hard to predict early in the design process, when design changes are easiest to implement. Capacitive coupling not only increases the delay on a wire but can also induce coupling noise on that wire when adjacent wires switch. Unlike fringe capacitance and area capacitance, capacitive coupling is a result of the specific wire interactions found in a design and can only be measured once the design has been routed. In [1] it is shown that coupling can lead to greater wire delays along a path, excessive power consumption due to increased capacitance and glitches, and even functional failures from coupling induced noise causing false switching.

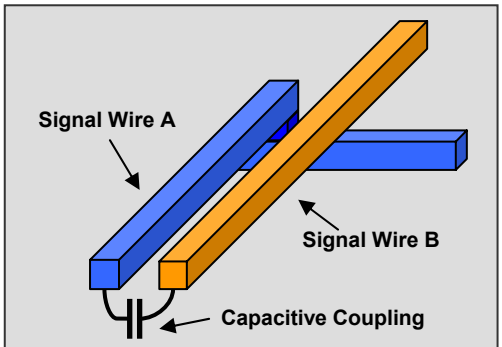

Figure 1: Capacitive Coupling

Two popular strategies have arisen to combat the problems caused by capacitive coupling: active mitigation through coupling aware routing tools, and designing around the problem using coupling avoidance techniques. Traditionally, designers have chosen to use coupling avoidance techniques on nets that are known to be strong coupling aggressors or especially sensitive to the effects of coupling. This decision can be driven by a performance requirement on a net forcing the designer to reduce the amount of capacitance as much as possible. The decision can also be driven by the characteristics of a circuit making it susceptible to coupling noise or a strong injector of coupling noise. The nets of the clock tree are a known group of strong aggressors often chosen for these techniques. The quantity of coupling noise a net creates depends on the net's switching speed, also known as the net's slew. In general, nets with faster slews induce more noise. The clock nets are most likely the fastest switching nets and they represent a non-trivial percentage of the overall nets found in a design. Paper [2] explores several strategies for avoiding coupling on the clock nets. It introduces two avoidance techniques—isolation and shielding—which have been shown to reduce the amount of coupling and increase the predictability of that coupling's effects on the nets in the design. It also combines these techniques into a strategy of first isolating and then shielding the clock nets. Of the two techniques, isolation is shown to produce the least capacitive coupling but also has the least predictability of the final amount of coupling. The shielding technique produces the greatest amount of coupling, but the majority of this coupling is between the clock wire and its shields making it a predictable value. The shielding technique also significantly decreases the amount of coupling noise injected into the shielded net. The combination of isolation and shielding produces a decrease in the amount of coupling and an increase in the predictability of that coupling, but consumes over twice the wiring resources.

While the techniques of isolation and shielding have proven effective at mitigating the effects of coupling for the clock nets in chips, they are poorly suited methods for routing clock nets in hard IP. A key factor in the marketability of hard IP is its size. Unlike the chip level, where the average circuit row utilization will rarely exceed 50%, hard IP implementations are typically pushed to achieve greater than 85% average utilization. Often the limiting factor in

achieving these high placement utilizations is the routability of the design. Hard IP pushes the edge of what is routable with today's routing tools. Techniques like isolation and shielding consume a great deal of valuable wiring resources. This paper examines the impacts of these techniques on the overall routability of a design and the quality of the clock tree routes produced. It also examines a modification of these techniques that consumes fewer wiring tracks and the latest active mitigation techniques.

2.0 Experimental Approach

To compare the clock routing techniques, each technique was used to route an identically placed and optimized clock tree from an embedded processor design. This section details the embedded processor design and the settings used for each of the clock routing experiments.

2.1 The Embedded Processor Design

The embedded processor design used for these experiments is a common piece of IP often hardened because of its high performance demands. The specific implementation used for this experiment targets a 0.13µm technology and contains only standard cells. The ASIC library used is a nine-track library with a robust assortment of cells. The design is restricted to only four levels of metal, the top level having a pitch equal to twice that of the lower levels. Approximately 0.5mm^2 of silicon area is consumed by the design, with the longest single wire being less than 1mm in length. Power straps have been inserted on the first, third, and fourth metal layers to form a fine-grained power grid. These straps are the minimum width for their respective layer and alternate between VDD and GND every ninth routing pitch. The combination of highly utilized circuit rows, only four routing layers, a robust power grid, and a double pitch fourth level of metal make routing especially challenging when implementing this design.

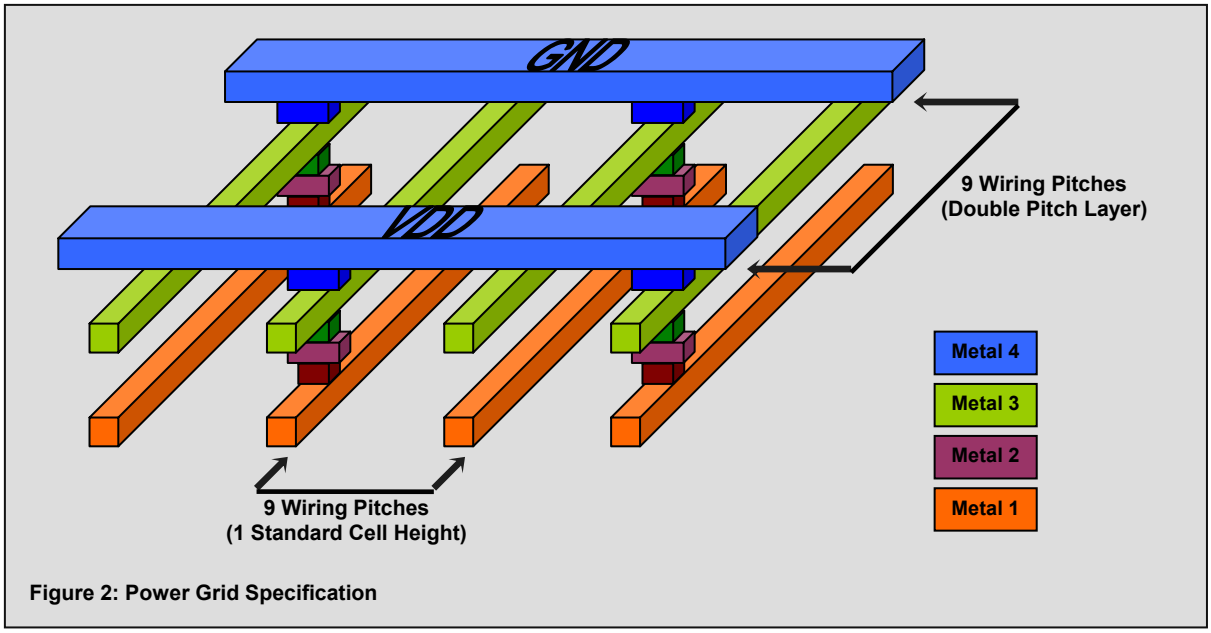

Figure 2: Power Grid Specification

The design contains only one clock domain that is driven by an external clock signal. A pin on the south edge acts as the input of the clock signal and drives a clock tree that fans out to 179 clock gates controlling 2,328 D-type flip-flops. The clock tree has been inserted post-

synthesis using Astro's clock tree synthesis feature and has been constrained to minimize both the maximum latency and skew of the clock tree. The resulting tree was implemented as an eight level deep tree composed of well-balanced standard cell buffers and inverters.

2.2 Traditional Clock Routing

Prior to deep sub-micron designs the primary distinction between clock routes and signal routes was the desire for clock routing to have zero skew across the clock tree. Specialized balanced routers have been used to equalize the wire length at any given level of the clock tree. This technique is ignorant to the effects of coupling and often yields clock routes that the routing tool incorrectly perceives to have minimal skew. Once the effects of coupling are added to the timing calculation a noticeable clock skew may be measured. The high amount of coupling may also cause the clock latency to be different from what was predicted prior to routing, leading to timing violations.

The traditional clock routing results, presented in section 3.0, were gathered using the Astro router and its balanced clock routing feature. The clock routes were constrained to prefer the third and fourth metal layers and to use a minimum width wire. The clock routing was performed prior to any signal routing so that the balanced router would not be constrained by any signal routes. After the clock routes were inserted, signal routing was performed and the experimental results were measured.

2.3 Clock Routing Isolation

Clock routing isolation is a technique that literally isolates the clock routes from any other route in the design. The tool places a routing blockage shape around the isolated routes to prevent the routing tool from using that space. Any violation of this blockage by another route is treated as a design rule violation. The quantity of isolation used to protect the clock routes is a difficult design decision. The greater the quantity of isolation used, the greater the protection from coupling, but also the more wiring resources that are consumed. Commonly designers choose to isolate one wiring pitch on either side of the clock routes. This effectively triples the amount of routing resources needed to route the clock. Also, a single pitch of isolation is insufficient to guarantee the clock routes are free from all the effects of their neighbor's coupling capacitance. Accurate capacitive coupling analysis tools generally look beyond the nearest wiring pitch when calculating the coupling capacitance. Isolating to include even the second nearest wiring pitch would use fives times the routing resources, an approach that is far too wasteful for use in most hard IP designs.

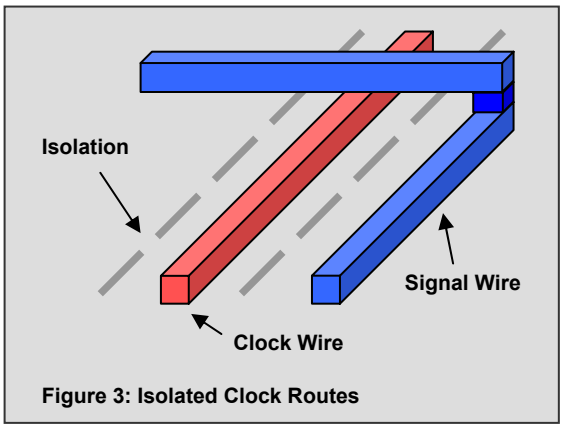

Figure 3: Isolated Clock Routes

The isolation experiment constrained the clock routes to prefer metal layers three and four and to place one wiring pitch of isolation around the clock routes on those levels only. No isolation was used for clock routes on the first and second metal layers to prevent pin access problems. The Astro router was also configured to allow power straps into the isolation area. Power straps add no coupling noise and a predictable amount of capacitive coupling that is often

considered acceptable. Using these clock routing constraints, Astro's balanced router was used to route the clock nets before a full routing of the design was performed. During the full routing of the design the router was allowed to alter the existing clock routes if either the clock route itself, or the blockage created by its isolation, prevented a legal routing solution from being found.

In the above experiment the isolation constraint was set as a strict rule. In some cases a designer will wish to make a trade-off between allowing some wires into the isolation area in exchange for being able to route the design. To measure the effects of this weakened isolation, the isolation constraint was removed from the above design and twenty iterations of Astro's search and repair were run. The localization of the search and repair's operation limits the number of wires allowed into the now freed isolation space to only those wires in areas of the design that could not be routed previously. In this way, isolation is only reduced in extremely congested areas of the design where the design could not route. Unfortunately, it almost guarantees that signal wires will be routed adjacent to the clock wires and the amount of coupling will increase. Also, the decision made by the routing tool regarding which wires to allow into this newly freed isolation area is made without respect to coupling. This combination can lead to coupling violations even when the majority of the clock routes remain protected.

2.4 Shielded Clock Routes

Shielded routing is a technique where the clock wires are surrounded on both sides by wires connected to ground or power, known as the shield wires. Instead of the clock wire interacting with an unpredictable quantity of signal wires it only couples to the two shield wires. This coupling to the shield wires is easily predicted and accounted for early in the design process, during clock tree optimization. The shield wires never switch so they do not induce coupling noise on the clock wire. In [1] it is shown that shield wires are significantly better at preventing coupling noise from neighboring signal routes than isolation for the amount of routing resources used. This increased protection is commutative; it also helps neighboring signal routes avoid the coupling noise induced by the clock routes. The greatest disadvantage of this technique is that while it decreases the amount of coupling noise and increases the predictability of the coupling capacitance, it also tends to increase that coupling capacitance. With shielded routing, a shield wire is almost always adjacent to the clock wire on both sides. Even using traditional clock routing techniques it would be highly unlikely for a signal wire to always be routed adjacent to both sides of the clock wire. The coupling between the clock wire and the shield wire will increase the latency of the clock tree and the amount of power consumed by that clock tree oscillating. The shield wires can also be difficult to insert into the design if the design's ground rules have a minimum via pitch that is greater than the wiring pitch. As the clock route vias between the routing layers, the shield routes will not be able to follow without taking special care to prevent a via-adjacency violation.

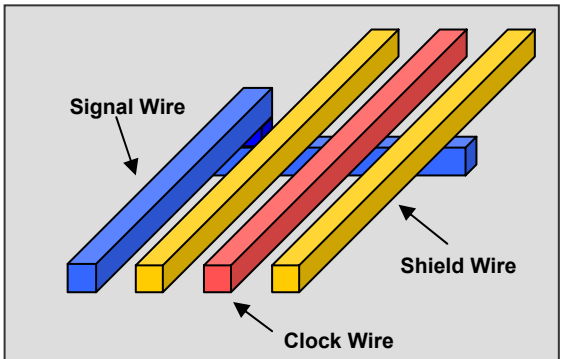

Figure 4: Shielded Clock Routes

Experimentation using the shielded clock routing technique was done by constraining the clock routes to prefer the third and fourth metal layers, leaving the first and second metal layers

for pin access to the standard cells. A minimum width shielding constraint was applied to the clock routes on the third and fourth metal layers only. The Astro router implements shielding using a two-stage approach. The user first applies a routing constraint on the shielded wires that reserves space for the shield wires during signal routing. The actual shield wires are inserted after the full signal routing is complete. Using the shielding constraint, the clock nets were first routed and then full signal routing was performed. After the design was completely routed the shield wires were inserted into the design and the results were measured. Then, as with the isolation experiment, a further experiment was performed with a weakened version of the shielding constraint. The shield metal and shielding constraint was removed from the above design and an incremental routing was performed using twenty search and repair operations. This allowed the router to use those tracks that had been reserved for shielding in areas where the routing congestion was so high that the design could not be wired. After this incremental routing was performed the shield wires were again inserted into the design. Because the router had allowed signal wires into the area where shield wires would be inserted, less shielding metal could be added to this second design.

2.5 Clock Routing Isolation Using a Fine-Grained Power Grid

This technique is a combination of isolation and shielding that takes advantage of a design's power grid. The power straps that compose the design's power grid make excellent shield wires without increasing the amount of wiring resources needed to route the design. Clock routes that are constrained to be wired adjacent to power straps will have shielding on at least one side of every wire segment. The other side of the wire segments can be protected using isolation or by adding a shield wire. The result of this technique is that the amount of wiring resources needed for preventing coupling of the clock routes becomes half that of traditional isolation or shielding. A drawback to this method is that the possible routing locations for the clock wires are limited to only those wiring tracks that are directly adjacent to the power straps. This constraint can lead to a trade-off between less coupling capacitance to the signal routes and greater timing predictability in exchange for higher clock tree skew and latency. It is also likely that there will not be a power strap near every pin in the clock tree network or on every metal layer of the design. In these cases, a small, unprotected route can be used to wire from a protected location to a pin or from one protected location to another.

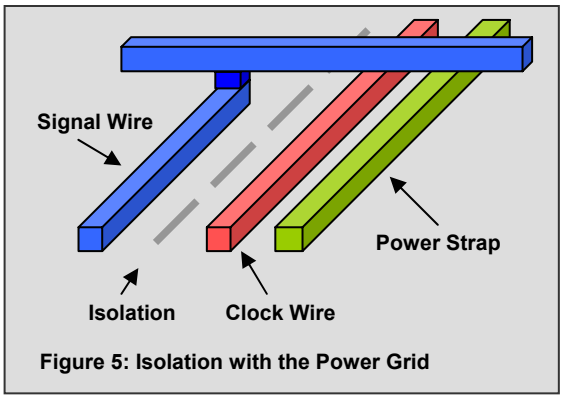

Figure 5: Isolation with the Power Grid

The experimentation for this technique is similar to the flow used in the isolation experiment. Again, the clock routes are constrained to prefer the third and fourth metal layer with one wiring pitch of isolation around the clock routes on those layers. Likewise, the isolation constraints are configured to allow power straps into the isolation area. At this point the methods for the two techniques diverge. Before the clock nets are routed the wiring tracks in the design must be reconfigured. This is done by removing all the wiring tracks on the third and fourth metal layers that are not directly adjacent to a power strap. Paper [4] covers the Astro commands to perform this operation in detail. Using only these partially shielded wiring tracks the clocks nets are routed. Next, the deleted wiring tracks are added back into the design and a full routing

is performed. During this full routing the router is permitted to move a clock route away from a power strap but at the expense of uncovering the track of isolation that was overlapping the power strap and effectively increasing the congestion in that area. From this design a second weakened experiment was performed. The isolation constraint was removed and twenty search and repair operations were run. This weakened version parallels the weakened version created in the isolation experiment. In areas where the routing congestion was too high for the design to be routed, signal wires were allowed into the area that had previously been reserved for isolation. Also, the router was allowed to move clock routes away from their shielding power strap and not suffer the penalty of exposing another track of wiring isolation.

The embedded processor design used for these experiments has a robust, fine-grained, power grid, which aids in the implementation of this technique. To understand the effects of the power grid design on this technique, the same experiment was run but with 75% of the power straps on the third and fourth metal layers deleted. During global routing and track assignment the Astro router takes each routing estimate and searches for the best possible track for routing that estimate inside a given window. This window's size is initially set to be equal to a square whose sides are the length of the circuit rows' height. Paper [7] explains how the size of this window can be controlled using the 'compactMode' parameter in Astro. This parameter scales the size of the window ranging from the default of one to a maximum of ten times the circuit rows' height. With the new power grid the spacing between the routing tracks for the clock nets increases significantly. It is often the case that there are more estimated routes in a window than the available tracks. When the Astro track assignment feature cannot find a track for a given net, the Astro router is unable to route that net. Using a default 'compactMode' setting of one, only 42% of the clock nets could be routed. Table 1 describes the effects of altering this parameter and underscores the importance of a fine-grained power structure when implementing this technique.

'compactMode' Setting	*Percent Routed*
1	41.94%
2	78.91%
3	91.94%
5	97.39%
7	98.10%
10	99.29%

Table 1: Effects of window size compensation for coarse power grids

This experiment shows how poorly suited this coupling avoidance technique is for designs with coarse-grained power grids. Even using the maximum allowed window size the Astro detailed router was unable to route all the clock nets. These nets can still be routed once the original routing tracks are added back into the design, but they will not gain the shielding from the power straps and will again consume two routing tracks for their isolation.

2.6 Active Clock Routing Mitigation with Astro XTalk

Active mitigation is a technique that focuses specifically on reducing the coupling noise aspect of the coupling capacitance problem. Coupling aware routing tools are used to avoid routing strong coupling noise aggressors near nets that would be sensitive to that coupling noise.

This is done to prevent coupling noise related false switching and to reduce the increased delays that can be caused by coupling noise. Several approaches are used to calculate when two nets should not be wired adjacent to each other. Paper [5] details an approach of ordering the nets in the design based on when they are expected to switch and in which direction they will switch. For example, two nets carrying signals that are exact negations of each other should not be routed in adjacent tracks, as they will always switch in opposition. This opposition causes the coupling noise induced by each net to delay the other. The paper also discusses a method for interleaving bus signals based on their function. Buses often switch simultaneously, and individual bits often switch in opposition to their neighboring bits. Other techniques for mitigating coupling include using parasitic extraction tools to discover strong aggressor nets or especially sensitive nets, and to keep these nets isolated from the remainder of the design. Unlike normal isolation, where the designer decides which nets to isolate, this method uses an automated technique to make its decisions. This technique can be run iteratively during the routing and timing optimization process as changes in the design introduce new candidate nets and eliminate old candidate nets.

The active mitigation experiment was performed using the Astro router in combination with the Astro XTalk capability. Astro XTalk[1] has the ability to calculate the coupling effects of the clock routes on other routes in the design. A global coupling constraint was set so that no input pin would have more than 30% of VDD worth of coupling noise injected on that pin. Using this constraint the clocks were routed, followed by a full signal routing. No further optimizations were performed on the design.

3.0 Experimental Results

The experimental results have been broken into three major sections. The first section, 3.1, discusses each technique's effect on the quality of the clock tree. The second section, 3.2, relates each technique to its effect on the overall routability of the design. The last section, 3.3, discusses how effective each method is at avoiding coupling between the clock and signal routes.

3.1 Effects on Skew and Latency

In general designers seek to have a clock tree with the least skew and the smallest latency possible. By comparing these two numbers, the quality of the clock tree routing generated by each technique can be compared. For each experiment, the routing was performed using the same standard cells and placement for the clock tree. Any difference in the skew and latency of the clock trees generated can be directly attributed to the routing technique used. Table 2 shows the results for each of the five methods. The table also includes the results from the experiment performed in section 2.5 using the sparser power grid and a 'compactMode' parameter setting of ten. The skew and latency for each technique have been measured after all signal routing was performed but before any constraints were weakened.

The results show that with the exception of the technique using isolation with a sparse power grid, all of the methods produce a clock tree with an equivalent skew. The clock skew for the sparse power grid experiment is much higher than the other methods and is caused by the router having so few tracks to use for clock routing. The same technique using the finer-grained power grid produced a clock tree with a skew that is comparable with the other techniques. This

[1] As of v2004.06-SP1

further emphasizes that the isolation technique using the power grid is only useful in designs with a fine-grained power grid. The clock tree latency value for the techniques is also similar, with the exception of the sparse power grid experiment and the shielding technique. The sparse power grid experiment again suffers from its lack of routing tracks. The router must use significantly more wire length than the ideal Steiner estimate to route the clock tree, resulting in a higher latency. The high clock tree latency with the shielding technique is caused by the capacitive coupling to the shield wires. As discussed in section 2.4 the shielding wires induce a significant amount of easily predicted coupling capacitance on the wire that they shield.

Method	*Clock Skew (ns)*	*Maximum Clock Latency (ns)*
Traditional Routes	0.134	2.128
Isolated Routes	0.151	2.109
Shielded Routes	0.190	2.230
Isolation with the Power Grid	0.142	2.117
Isolation with a Sparser Power Grid	0.265	2.291
Active Mitigation	0.140	2.132

Table 2: Effects on skew and latency

3.2 Effects on Overall Routing Congestion

For hard IP designs attempting to achieve the smallest physical size possible, routing congestion is a major concern. Three of the five techniques explored in this paper make some significant trade-off between the amount of routing resources needed to route the clocks and the amount of capacitive coupling related to those clocks' routes. Table 3 shows the amount of over-congestion in the design when using each technique and the number of routing violations left in the design after full signal routing has been run. The over-congestion was measured using Astro after all the clock routes had been inserted but before the signal nets had been detail routed. Over-congestion is the percentage of global routing tiles where the number of routes passing through that tile, either horizontally or vertically, exceeds the number of tracks available for routing. The number of routing violations shows the number of design rule violations left in the design that the router could not fix. The post-weakening violations number describes the results for methods where certain routing constraints were reduced or removed after the initial routing and an incremental routing was performed. This second routing run is detailed in the description of each technique where applicable.

Method	*Horizontal Over-Congestion*	*Vertical Over-Congestion*	*Post-Routing Violations*	*Post-Weakening Violations*
Traditional Routes	2.8%	0.4%	186	—
Isolated Routes	5.0%	8.4%	2943	179
Shielded Routes	5.1%	8.4%	8546	4942
Isolation with the Power Grid	4.8%	7.4%	945	226
Active Mitigation	2.8%	0.4%	175	—

Table 3: Effects on overall routing congestion

The isolation and shielding methods both have significant impact on the overall routability of the design. Both techniques triple the amount of routing resources needed to route the clock tree, and this is apparent in the results. It is reflected not only with the increase in the

over-congestion but also by the massive amount of routing violations left in the design after the router has run. Once the constraints are weakened, the number of routing violations decreases significantly, but it is still clear that the design using shielding will never successfully route. In contrast, the technique of using isolation in combination with the power grid has less over-congestion and far fewer routing violations. Once the constraints are weakened for this technique its routing results become comparable with those of traditional routing. The active mitigation and traditional routing techniques both show significantly less of an impact on over-congestion and produce few routing violations. Both of these methods consume minimal routing resources for coupling avoidance.

3.3 Effects on Capacitive Coupling between Clock Routes and Signal Routes

Each of the five techniques seeks to reduce the amount of coupling between the clock routes and the signal routes. It is this coupling which is so hard to predict and correct early in the design process. The amount of coupling capacitance between the signal and clock routes can only be determined after the design is completely routed. Also, coupling between the clock and signal routes involves two actively switching nets that can cause false switching and power consuming glitches on either net. While a high amount of coupling capacitance does not guarantee one of these coupling noise related failures, it is a symptom. The coupling data presented in table 4 was gathered from each experiment's fully routed design using Star-RCXT for parasitic extraction and includes only coupling between a clock and signal route. Any coupling capacitance between the clock routes and the power straps or shield wires has not been reported. Coupling capacitance that involves a clock route and a power strap or shield wire is less interesting as it is easily predicted and involves a net that never switches. In table 4, the worst signal coupling represents the greatest amount of coupling capacitance between a clock route and a signal route, regardless of whether the signal route or clock route was the aggressor. Because this worst coupling may not be representative of the entire design, all of the reported coupling events between the clock and signal routes have been summed to give the aggregate worst signal coupling.

Method	*Worst Signal Coupling (e-13 F)*	*Aggregate of Worst Signal Couplings (e-13 F)*	*Worst Signal Coupling Post-Weakening (e-13 F)*	*Aggregate of Worst Signal Couplings Post-Weakening (e-13 F)*
Traditional Routes	0.09	16.07	—	—
Isolated Routes	0.02	1.91	0.03	2.52
Shielded Routes	0.01	0.68	0.04	1.03
Isolation with the Power Grid	0.02	1.15	0.02	1.33
Active Mitigation	0.06	14.38	—	—

Table 4: Effects on capacitive coupling

The techniques of isolation, shielding, and isolation with the power grid all effectively reduce the amount of coupling capacitance between the clock and signal routes. The aggregate coupling for these techniques is an entire magnitude better than when using the traditional clock routing technique. Of the three, the shielding technique proves to be the best at protecting the clock routes, having only half the aggregate coupling of the other two. After these three techniques have their constraints weakened to improve the routability of the design, the isolation and shielding techniques both see a significant increase in their aggregate coupling capacitance.

However, all three of the weakened techniques still show less coupling than is found when using traditional routing. The active mitigation technique shows only a marginal decrease in the coupling when compared to the traditional clock routing technique. This is because the active mitigation in Astro XTalk only focuses on reducing those coupling events that exceed the coupling noise constraint and does not try to reduce overall coupling. This technique is also attempting to address coupling noise constraint violations globally and is not localized to just the clock nets – making the problem more complex.

4.0 Conclusions and Recommendations

Designers face three major problems when routing the clock tree in their design. They seek to reduce the skew and latency of the clock tree in order to minimize its impact on timing closure. They also attempt to reduce the impact of the clock tree's routing on the overall routability of their design. Lastly, the designers seek to avoid the effects of capacitive coupling on the routes of that clock tree. This includes not only the effects of coupling noise, which can cause glitches and false switching, but also the increased delay caused by capacitive coupling. Specifically, the clock tree routing should reduce the amount of coupling between the clock routes and signal routes. This coupling is difficult to predict, as it is only measurable after the clock tree has been optimized and the entire design has been routed.

Each of the five methods explored in this paper have distinct properties with respect to these three problems that make them more or less suitable for certain types of designs. The shielding technique provides the best protection from capacitive coupling but has the greatest impact on a design's routability. These traits make it an excellent technique for routing clock nets in designs where capacitive coupling is a serious concern and routing resources are plentiful. In contrast, the traditional routing technique provides little to no protection from capacitive coupling but has no adverse effect on routability. This makes traditional routing well suited to a design in earlier technologies, where capacitive coupling is not a concern. Table 5 provides a summary of each technique's properties and illustrates how no one solution is ideal at solving all three problems.

Method	*Clock Tree Quality*	*Design Routability*	*Capacitive Coupling Avoidance*
Traditional Routes	Good	Excellent	Poor
Isolated Routes	Good	Moderate	Good
Shielded Routes	Moderate	Poor	Excellent
Isolation with the Power Grid	Good	Good	Good
Active Mitigation	Good	Excellent	Moderate

Table 5: General properties of the clock routing techniques

For hard IP designs in deep sub-micron technologies it is important to use a technique that provides a good solution to all three problems. Traditional routing proved to be a technique incapable at avoiding noise. It was shown that the shielding and isolation techniques are useful at creating a reasonable quality clock tree that avoids noise but also consumes a significant amount of the routing resources. In the case of shielding, so much of the routing resources were consumed that the designer would be forced to grow the design in order to route. It was also shown that active mitigation, in its current implementation, focuses too much on the specific problem of coupling noise to be used as the sole method of capacitive coupling avoidance. Of the

five methods, only the method of using isolation in combination with the design's power grid provided a reasonable trade-off for dealing with the three problems. Unfortunately, only designs that use a fine-grained power grid can successfully implement this technique.

This paper has explored the most common techniques used for routing clock trees in deep sub-micron technologies, including specialized techniques for hard IP and newly evolving techniques that rely on active mitigation. Each of these techniques failed to completely solve the problems associated with clock tree routing. The isolation with the power grid technique, a technique combining isolation and shielding, illustrates that these techniques are not mutually exclusive. Interesting and novel combinations of these techniques, along with new methods taken from active mitigation, could yield an abundance of new design possibilities. Further exploration is needed as technologies scale smaller and the problem of capacitive coupling increases.

5.0 Acknowledgments

The authors would like to thank Ray Yock and Kevin Brelsford of the Synopsys Corporation, along with Tom Lepsic of the IBM Corporation, for their technical support with gathering and interpreting the experimental data. They would also like to thank Laura-Rose Russell for her excellent work on the design and editing of this paper.

6.0 References

[1] M. Elgamel and M. Bayoumi. *Interconnect Noise Analysis and Optimization in Deep Sub-micron Technology*, IEEE Circuits And Systems Magazine, Fourth Quarter 2003.

[2] T. Bednar, P. Buffet, R. Darden, S. Gould, P. Zuchowski. *Issues and strategies for the physical design of system-on-a-chip ASICs*, IBM Journal of Research and Development, November 2002.

[3] J. Cong. *An interconnect-centric design flow for nanometer technologies*, Proceedings of the International Symposium on VLSI Technology, Systems, and Applications, June 1999. Online at http://citeseer.ist.psu.edu/article/cong99interconnectcentric.html

[4] D. Powell. *Route Clock Nets avoiding potential X-talk violations, while minimizing consumption of routing resources*, Synopsys Solvnet Document ID 012644, December 2004. Online at http://solvnet.synopsys.com/retrieve/012644.html

[5] J. Yim and C. Kyung. *Reducing cross-coupling among interconnect wires in deep-submicron datapath design*, Proceedings of the 36th Design Automation Conference (DAC), June 1999.

[6] C. MacDonald and A. Jarrar. *Addressing Crosstalk Delay – Achieving Timing Closure*, Proceedings of the 2004 San Jose SNUG Conference, March 2004.

[7] E. Dagan. *Overflow, GRC, DSR: What's It All About?* Synopsys Solvnet Document ID 010147, January 2004. Online at http://solvnet.synopsys.com/retrieve/010147.html