

Using DFM Routing to Impact Design Performance and Yield

Doug Stiles^a, PT Patel^b, Mitchell Heins^b, Surbhi Agarwal^b,
Simon Favre^c, Carlo Guardiani^d

^aMicrosoft Corporation, Mountain View, CA

^bPyxis Technology, Inc., Austin, Texas

^cPonte Solutions, Inc., Mountain View, CA

^dPDF Solutions, Inc., San Jose, CA

ABSTRACT

As design teams continue to drive forward with the use of advanced process technologies there has been much discussion about the use of design-for-manufacturing (DFM) techniques and questions about the usefulness and effectiveness of DFM. This paper describes an experiment conducted to establish the quantitative performance and yield impact of proactively using DFM techniques during the routing of a design.

The design used in the experiment is an established production device for which both performance and yield are known. A methodology for seamlessly introducing DFM into the routing is presented, including the impact of DFM on the performance of the design. Also shown is a strategy for making quantitative comparisons of predicted yield between non-DFM and DFM versions of the design and the results from those comparisons for this experiment.

Final results are summarized showing better performance and a predicted yield improvement of greater than 7% by introducing DFM techniques during the routing of the design, given the same placement & timing constraints.

INTRODUCTION

There have been numerous papers written on the techniques that can be employed during integrated circuit (IC) design to achieve better overall manufacturability and yield. These design-for-manufacturing (DFM) techniques have historically been grouped by the yield limiting affects that they are meant to mitigate, such as wire spreading for particulate (or random) based yield loss, influence rules for systematic (or lithography) based yield loss and metal density control for parametric yield.

While these DFM techniques were useful, they lacked a strategy for easy deployment and acceptance by the design community. Many times, the designers were expected to understand complex relationships between different yield limiters and even if they did understand, the industry was slow to provide sufficient fab data to allow the designers to make the necessary trade-offs. When data became available the bigger problem became that there was no automation in place to help designers implement the trade-offs. Because of this much of what has been espoused as DFM had been relegated to the design of smaller IP blocks where designers could manually handle the required trade-offs.

In an effort to realize a more wide spread use of DFM techniques, Microsoft, Pyxis, PDF and Ponte conducted an experiment to establish a methodology for proactively using DFM techniques during the routing phase of the design flow and create a strategy for quantifying the performance and yield impacts of that methodology.

EXPERIMENT OVERVIEW

Experimental Test Block

The design used in the experiment is an established production device for which both performance and yield are known. To simplify the experiment one block of the design was used that represented approximately ten percent of the chip area as shown in figure 1. The chosen block is actually replicated four times throughout the chip and is made up of both standard cells and custom memories. The block area was approximately 85% utilized by the standard cells and memories and thus represented a fairly constrained environment in which to practice DFM techniques.

Basic block statistics were as follows:

- Number of cell instances: ~ 1.4 million
- Number of signal nets: 1.1 million

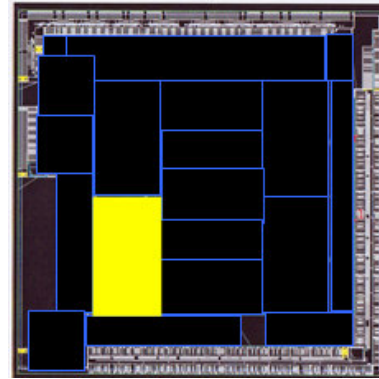


Figure 1 – Chip Floorplan

Strategy for Results Comparison

Since the design was already in production it was used to compare results against the experimental layout. A copy of the original block was made using the existing standard cell and memory placements along with the corresponding power and clocking structures and timing constraints. The rest of the routing was stripped out of the block and the block was re-routed and optimized using a correct-for-manufacturing approach with the Pyxis *NexusRoute* manufacturing-aware router.

Yield comparisons were done using yield analysis tools from Ponte Solutions, Inc. and PDF Solutions, Inc. The layout from both the taped-out block and the experimental block were run through these tools to understand the difference in the predicted yields given the original block and DFM style layouts. Ponte's Yield Analyzer (YA) tool was used to simulate yield loss due to particulates (random) and via failures (single vs double via coverage). PDF Solutions Yield Ramp Simulator (YRS) was used to simulate yield loss due to systematic via misalignment and via density failures.

Timing for the experimental block was compared to the timing of the taped-out design. Parasitics for the both blocks were calculated using the Pyxis parasitic extractor. Delay calculation and static timing analysis were performed using Synopsys' PrimeTime and PrimeTime-SI products along with the parasitics generated by the Pyxis extraction tool.

Timing slack histograms of both the taped-out design and the experimental block were compared. All timing comparisons were done inclusive of all DFM treatments.

Design & Data Flow

Data was exchanged between the existing design flow and the Pyxis *NexusRoute* product by using standard interface formats as shown in figure 2. LEF was used to pass technology data and rules; DEF was used to pass standard cell and memory IP block placements, power rails, blockages etc. Design connectivity and timing constraints were passed using Verilog, DEF and Synopsys Design Constraints (SDCs). PrimeTime and PrimeTime-SI were used for delay calculation and static timing analysis. Parasitics generated by the Pyxis extractor were passed to PrimeTime and PrimeTime-SI through the SPEF format.

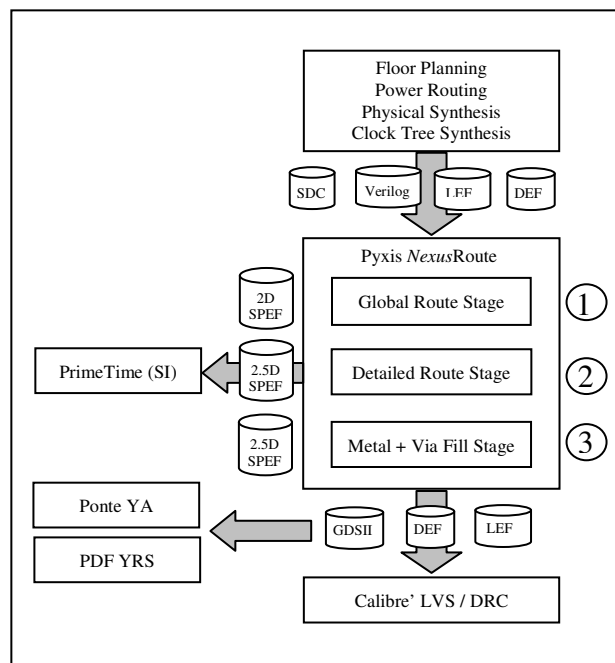


Figure 2 – Design & Data Flow

Methodology for DFM Implementation

A methodology for proactively using DFM techniques during the routing phase of the design flow was put in place using the Pyxis *NexusRoute* product. Some of the basic assumptions and tenets of the methodology were as follows:

- The routing phase should fit seamlessly into the design flow and behave like a normal design flow using standard interface formats to exchange data
- The routing phase should use the same floorplan and timing constraints as the original block flow
- DFM techniques should be introduced as early into the routing phase as possible to enable the greatest degrees of freedom for optimizing the design both in terms of performance and yield
- Use of a correct-by-construction approach whereby DFM is built into the design from the beginning as opposed to finding and fixing problems post-routing which could then affect timing
- Use of recommended design rules when ever possible but do so in an automated fashion so that the designer does not have to decide when or when not to use these rules
- Route the design in an automated fashion using DFM techniques during routing to help close timing
- Use a progressive timing closure approach to use the minimize the number of constraints put on the router to address timing and signal integrity closure

DFM Routing Methodology

Routing was done in three stages as shown in figure 2. The methodology employed works to accomplish the items detailed in each stage in figure 3. In general as much optimization as possible was done during the global route stage. These optimizations were then passed on to the detailed route stage as constraints.

Pyxis' progressive timing closure flow was used to automatically make several passes through the global router. Incremental 2D extraction, delay calculation and timing analysis using PrimeTime were performed upon iterations. Additional routing constraints were automatically generated on each pass until both timing and DFM goals were met. During the global route stage the router did a sensitivity analysis of the design's timing paths and selectively used non-default routing rules for signal integrity and cross talk avoidance. Final timing optimizations were done using netlist gate sizing techniques.

1. Global Route Stage

- a. Overall route planning with 3D wire balancing for metal density and density gradient control
- b. 2D wire spreading to reduce critical areas for shorting defects
- c. Automated progressive timing closure flow using wire and netlist optimization. Includes wire layer assignment and use of non-default rules for DFM and for signal integrity / cross talk avoidance.
- d. 2D extraction with estimations for congestion and eventual metal fill
- e. PrimeTime and PrimeTime-SI are used for delay calculation and STA

2. Detailed Route Stage

- a. Signals are routed correct-by-construction layout with incremental DRC during routing
- b. 2D wire spreading for DFM
- c. All routing is cost-based; allocates space for DFM
 - i. Preferred design rules
 - ii. Replacement of single vias with double vias
 - iii. Other via treatments if vias can not be doubled
 - iv. Widening of wires to reduce critical area for "open" type defects
 - v. Protection of lithography sensitive patterns
- d. DFM treatments are committed where the space reserved is available after detailed routing.
- e. DFM analysis is done internal to the Pyxis tool using pDfx™ models from PDF Solutions

3. Metal / Via Fill Stage

- a. Metal fill is added to bring density up to required goals. Metal fill is kept at least 3X minimum spacing away from signal wires and can be tied or left floating.
- b. A via density check is made and vias are added to reach a required density.
- c. 2.5D extraction is performed on the final layout including the metal and via fill. Delay calculation and static timing analysis is again performed using PrimeTime-SI.

Figure 3 – DFM Routing Methodology

Upon meeting timing and DFM goals at the global route stage, the design was sent through the detailed route stage. During the detailed route stage, wires were routed using a dynamic costing algorithm that reserved space for the various DFM treatments as detailed in figure 3. As part of the methodology, the router used a combination of preferred design rules and costing to reserve as much space for DFM treatments as possible. After routing, DFM treatments such as double vias, wide wires etc. were committed to the layout based upon the amount of reserved space left un-used. As before PrimeTime-SI was used to perform delay calculation and static timing analysis. After completing the detailed route stage the design was sent through the metal / via fill stage. Density goals were set for metal and via density as well as density gradients for metal using a 10um X 10um window stepped in 5um increments. A final 2.5D extraction, delay calculation and static timing analysis were done to ensure metal and via fill did not negatively affect timing.

After all routing stages were complete the experimental layout block was analyzed for errors using Mentor Graphics' Calibre' DRC and LVS tools.

During each stage of the routing process, progress towards DFM goals were checked using analysis capabilities included in the Pyxis *NexusRoute* software. The *NexusRoute* installation included fab fail rate process models from PDF Solutions which were used to analyze and Pareto the areas of highest possible yield impact for the design. Parameters of the *NexusRoute* software were adjusted to achieve optimum yield given feedback from these analysis. Final GDS data was analyzed for yield loss using Ponte Solutions' Yield Analyzer and PDF Solutions' Yield Ramp Simulators. Results from both blocks are compared below.

Yield Limiters Targeted for Optimization

Random yield

- Critical area for opens
- Critical area for shorts
- Single vs. double vias

Systematic yield

- Via protection (for litho & misalignment)
- Litho-based wire protection
- Via density
- Metal density and density gradient

EXPERIMENT RESULTS

Routing Statistics

Final routing statistics for both the original block and the experimental block are shown in tables 1 and 2. Shaded rows of the table represent data that was used in the calculation of yield impact shown in table 3.

- Table 1 represents items associated with random yield.
- Table 2 deals specifically with lithography based yield limiters.

Critical Area for Opens & Shorts

Overall wire length goes toward the calculation of critical area for both open and short type defects. In this case, the experimental block had slightly less overall wire length (~1.58% reduction) which helped to reduce the critical area for "open" type defects. Wire spreading in the experimental block resulted in approximately 53.71% of the wire length having 3X minimum spacing next to it which reduced the critical area for "short" type defects. In addition, 26.43% of the wire length of the experimental block was widened to 1.2X the minimum width which also helped to reduce the critical area for "open" type defects.

Random Yield	% Delta Experimental Block vs. Original Block
Wire Length - mm	(1.58)%
Wide Wire Length - mm	26.43%
Via Count	(7.55)%
Double Via %	18.79%
Single Vias	(53.33)%

Table 1 – Random Yield Limiters

Via Statistics

As shown in table 1, the DFM driven experimental block had 7.55% less via connections. In addition, the experimental block had 18.79% more of the vias made manufacturing robust through via doubling. Through a combination of less vias and higher via doubling the number of single vias remaining in the experimental block was cut by 53.33%

Given that vias are one of the main yield limiters in the routing layers, special attention was given to the remaining single vias to try to make them as robust as possible. The first row of table 2 shows that the experimental route reduced the number of unprotected vias by 70.98% leaving the total percentage of unprotected vias at 12.62% as opposed to the original taped-out block which had 40.2% of the vias left unprotected. In addition to protecting single vias, an effort was also made in the experimental block to ensure adequate via density to protect against layer delamination issues.

Litho	% Delta Experimental Block vs. Original Block
Unprotected Vias	(70.98%)
Line-end Spacing %	93.33%
Min Width Jogs	(96.36%)
Protected Corner %	92.09%

Table 2 – Lithography Related Yield Limiters

Litho-aware Routing

In addition to standard treatments for random and systematic yield limiters for vias, special attention was also placed on litho-aware routing in the experimental block. Some of these are listed in table 2. Special DFM treatment was also made for non-preferred direction jogs which can suffer from printability problems within the normal scanner defocus range. Greater than 90% of instances of these items were treated for better manufacturing robustness. The table rows for these items are not shaded as the existing yield simulation tools are not capable of determining their impact on a predicted yield number and as such the predicted yield improvements do not reflect the benefit of these additional DFM treatments.

CMP-aware Routing

The third stage of routing for the experimental block focused on bringing metal and via density up to fab specific goals. There was no direct comparison made against the original block. Data gathered from the experimental block confirmed that metal density goals of greater than 30% and via density goals of greater than 1.5% were met with metal density gradients of less than 9% as measured using a 10um X 10um window stepped in 5um increments. Density and density gradient goals were achieved through application of Pyxis *NexusRoute* metal and via fill routines. At this stage there was no way to judge the impact of metal fill on the parametric yield of the device and as such the predicted yield improvements do not reflect the benefit of the additional metal fill added to the design.

Predicted Yield Comparison

Both the original taped-out block and the experimental block were run through Ponte Solutions' YA and PDF Solutions' YRS tools and then compared. Partial yields were calculated for the block under comparison and scaled to be representative of the 80% of the design for which this type of routing would apply. Given the block used for the experiment was one of the highest density blocks it was felt that this would give a conservative estimate of what could be achieved on the entire design. Because of the sensitive nature of yield numbers the predicted yields have been removed from the table, leaving only the difference in the partial and final predicted yields. As shown in table 3, the experimental DFM driven route scored an

Predicted Yield Comparison	Taped-out Block Partial Yields	Experimental Block Partial Yields	Delta Yield
Shorts	Predicted Yield Numbers withheld due to proprietary nature of data.		0.70%
Opens			1.17%
Total metal yield			1.75%
Via			3.43%
Via pitch density			3.32%
Single via treatments			0.90%
Total via related yield			6.80%
Total yield			7.46%

Table 3 – Yield Comparison

improvement of 7.46% yield over the original taped-out block. It is important to remember that partial yields can not be added to get the final yield but instead the product of the partial yields must be taken.

The predicted yield improvements shown in table 3 reflect a reduction in yield limiters as highlighted by the shaded rows of tables 1 and 2. Lithography and parametric related yield improvements are not included in these calculations.

Performance Impact

Comparisons were made between the original taped-out block and the experimental block to understand the performance impact on the design when using DFM based routing.

Data for both blocks was derived using the Pyxis 2.5D extractor and PrimeTime for delay calculation and static timing analysis. Initial comparison of the Pyxis data compared to data from the EDA design flow used to tape-out the original block showed the Pyxis timing of the taped-out block to be slightly conservative, in the range of 25 picoseconds.

Analysis results for both the taped-out design and the experimental block are shown in table 4. The taped-out block had greater than 1000 paths with negative slack without taking into account cross talk induced delays. The experimental block met setup conditions with no negative slack.

It should be noted that the design team had significant difficulty in closing the taped-out block due to signal integrity issues resulting in several months of delay. Final timing still had over a 1000 paths with negative slack. Timing for the experimental block was analyzed using PrimeTime-SI using coupling capacitances. Results are shown in the far right column of table 4. The experimental design did have

277 paths with negative slack however there were a significantly smaller number of paths that were stacked up against the 0 slack-line as indicated in the last row of table 4 and figure 4.

Timing Results	Taped-out Block without SI Timing Slack in ns	Experimental Block without SI Timing Slack in ns	Experimental Block with SI Timing Slack in ns
Top path	-0.096	+0.028	-0.151
Top path -1	-0.094	+0.029	-0.136
Top path -2	-0.092	+0.032	-0.132
Top path -3	-0.085	+0.035	-0.125
Top path -4	-0.081	+0.036	-0.112
# of paths with negative slack	>1000	0	277

Table 4 – Timing Results

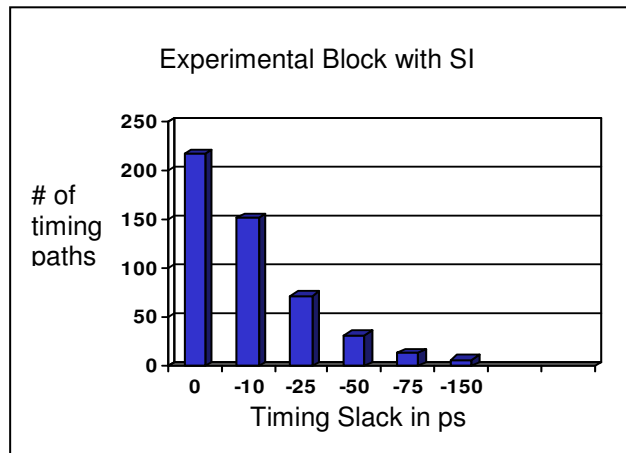


Figure 4 – Timing Slack Histogram

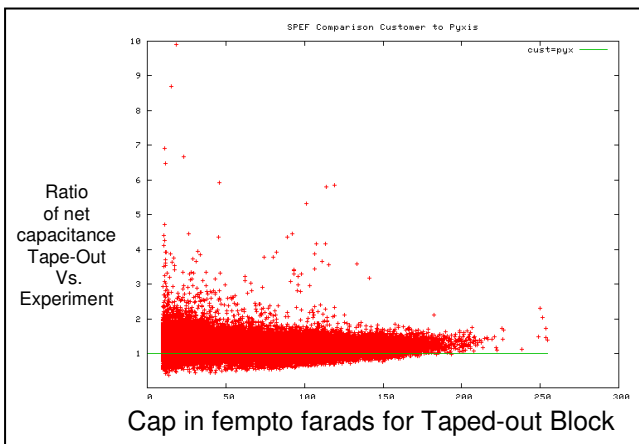


Figure 5 – Capacitance Comparison

Capacitance was also compared between the blocks. Nets with less than 10 ff were not compared. The ratio of the capacitance for the taped-out design to the experimental design is plotted in figure 5. The horizontal axis is the capacitance in fempto farads of the nets in the taped-out design. This data shows that the experimental block has less overall net capacitance.

- Number of nets of Taped-out block with cap > Experimental block cap = 216,378
- Number of nets of Taped-out block with cap < Experimental block cap = 32,288

DRCs as reported by Calibre

GDS for the experimental block was analyzed using Calibre' DRC & LVS tools. There were 0 LVS errors and 27 DRC errors reported. For the experiment this was deemed adequate given the size and complexity of the block (greater than 1 million nets, greater than 85% utilized, routed with full dfm treatments and metal and via fill).

CONCLUSIONS

The experiment established that a significant improvement in yield can be gained by proactively employing DFM techniques during the routing stage of a design. In this experiment the improvement in predicted yield was shown to be greater than 7%. In addition the experiment showed that by employing a correct-for-manufacturing approach to DFM during the routing phase that the overall performance characteristics of the design can be maintained and in some cases improved as was the case with the experimental design. This should be contrasted to existing design flows where post-route DFM techniques can many times disturb the overall performance of the chip causing design re-spins.

The experiment also showed that a DFM implementation methodology can be put into place that is non-intrusive to the physical designer's normal work flow. The experimental flow used a commercially available auto-router that interfaced to the rest of the design flow using commonly available standard interchange formats and interaction with signoff analysis tools for timing analysis and yield prediction. The experimental design was auto-routed using identical timing constraints, floorplan footprint, standard cell placements, power and clock routing structures and metal layers as were used in the original taped-out design.

Finally, the experiment put to rest the question of the real usefulness and effectiveness of Design-For-Manufacturing and showed that it can be deployed to a wide set of physical designers with minimal disruption to design flows and use models while returning a significant increase in yield and manufacturability.