We Haven’t Survived 65 nm – We’re Just in the Eye of the Storm!

by Mitch Heins, Vice President of Engineering, Pyxis Technology

Introduction

The phrase "eye of the storm" refers to the relatively calm center of a hurricane, where winds are light and the skies are only slightly cloudy, or even clear. If the eye of the hurricane passes over during the daytime, one might see sunny skies and even enjoy a rise in temperature. Observers sometimes mistake the arrival of the eye as a sign that the storm is over; but, at the end of the eye's passage, the storm returns at full force with a deluge of rain and violent winds blowing in the opposite direction to that of the storm’s leading edge.

In the case of digital integrated circuits (ICs including ASICs, ASSP, and systems-on-chip), many people seem to have the impression that the transition to the 65 nm technology node is proving to be "not as bad as expected." In reality, however, we are in the eye of the storm. So far, only a small number of chips have taped out with apparent success. However, there's a gap between tape-out and production. Reports are now coming back that yields are lower than even the pessimistic expected.

Why are Manufacturability and Yield Important?

In the context of digital ICs, the phrase design-for-manufacturing (DFM) refers to a variety of techniques used during the process of implementing the design to ensure that it can be manufactured correctly. Meanwhile, the term yield refers to the number of die that work as a percentage of the total number of die on the silicon wafer. Hence the phrase design for yield (DFY) refers to any techniques used to improve the yield of a particular device. In reality, these techniques are so intertwined that it is becoming common to consider them as being a single entity: DFM/DFY.

Yield is a function of the device's manufacturability and there are three main "buckets" into which yield-related problems may be categorized. These buckets are commonly referred to as *Random Yield* (sometimes called *Statistical Yield*), *Systematic Yield*, and *Parametric Yield*.

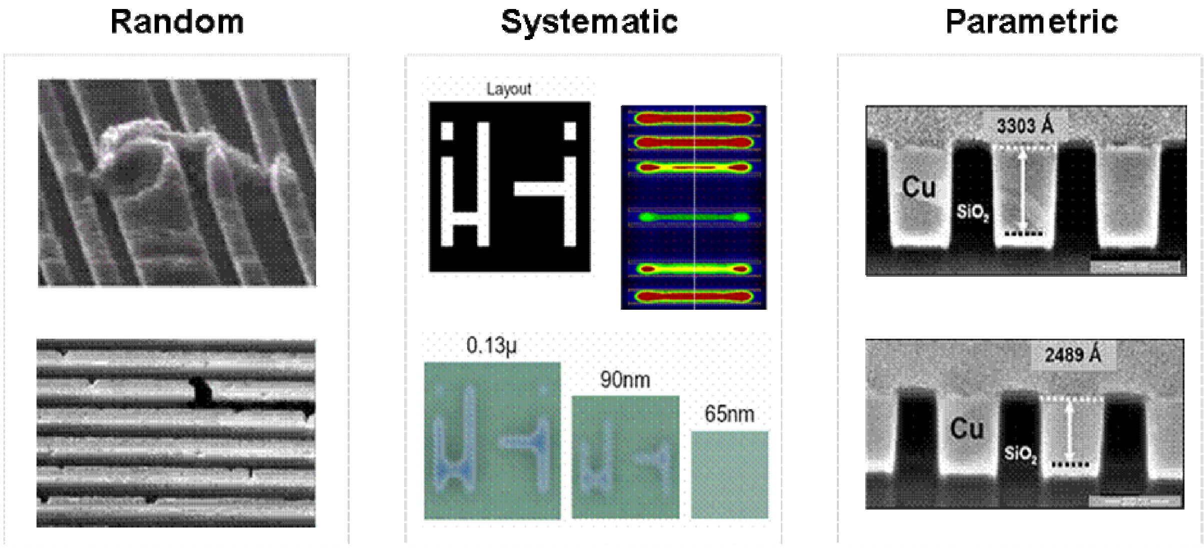

Figure 1. Yield loss due to random, statistical and parametric effects

Random (Statistical) Yield

As its name suggests, this form of yield is a function of random effects that occur during the manufacturing process. For example, no matter how clean the wafer manufacturing environment, there are always some number of small particles in the atmosphere that may land on the surface of the chip.

Such particles may cause catastrophic faults in the form of open or short circuits. Alternatively, in some cases they may cause parametric variations. For example, a particle may land on a non-critical area of a particular layer and may cause a non-planar feature (bump) in subsequent layers. In turn, this bump may end up varying the width or thickness of a wire on a higher layer, changing the electrical characteristics of that wire and resulting in a parametric yield failure (as discussed below.)

By their very nature, random defects are difficult to control. However, it is possible to create the design in such a way as to minimize their effects on final yield.

Systematic Yield (Including Printability Issues)

The term "systematic" encompasses the concepts of "logical," "methodical," and "ordered." Thus systematic yield refers to a class of manufacturability issues that are the result of some combination and interactions of events. These issues can be identified and addressed in a systematic way.

Many systematic yield issues are design-dependent. For example, some designs may have high densities (concentrations) of wires in certain areas and low densities in others. Such density variations can affect the amount of etching that takes place in the various regions. Similarly, in the case of process steps like chemical mechanical polishing (CMP), variations in wire density can cause differences in the effectiveness of the polishing process, which can result in areas where some wires are thinner than others. In turn, this affects the resistance and capacitance values associated with these wires, which can modify the power and performance (timing) of the design.

By understanding systematic effects during the design implementation process it is also possible to create a design in such a way as to minimize their effects on yield.

Parametric Yield (Including Variability Issues)

The concept of parametric yield refers to the fact that a chip may perform its logical function correctly ("stimulus X returns response Y"), but variations in the device's parameters may mean that it does not achieve its specified performance goals. If transistor channels aren't formed quite as expected, for example, the result may be lower drive capabilities, increased leakage current and greater power consumption, increased resistance and capacitance (RC) time constants, and slower chips.

Alternatively, issues in the etching and CMP processes may cause non-planarity in the surface of the chip, which, in turn, can cause wires to have higher resistances and/or capacitances than expected, which will result in the device's speed falling and its power consumption rising.

One aspect of parametric yield that is becoming extremely significant is that of variation or variability. There has always been an issue with regard to inter-wafer variation, which refers to slight differences between wafers in a lot. In the case of today's technology nodes, there can be significant variations between different areas on the same wafer (intra-wafer variation) and even the same die (OCV or on-chip-variance).

By understanding parametric effects during the design implementation process it is possible to create designs that minimize loss in chip performance and yield.

So Why Are Yield and Manufacturability Important?

The reasons yield and manufacturability are important may be summarized as follows:

- The chips (and associated products) may completely miss the market window.
- The chips (and associated products) may hit the market window, but the chips may cost too much to make the products economically viable.
- The chips may not perform at required level; that is, they still may function, but not at the required speed.
- The chips appear to be reliable after volume production, but may suffer catastrophic failures in the field earlier than their expected lifecycle.

The bottom line is that if DFM/DFY issues are not addressed, it may simply not be possible to achieve economically viable yields in the forthcoming technology nodes.

How Can We Address DFM/DFY Issues?

Until recently, design engineers had to concern themselves very little with manufacturability. As long as the design met a few simple rules – such as wires meeting minimum width and spacing values – it was assumed that the device could be manufactured.

Similarly, with the exception of specialist teams working on extremely high volume products such as SRAM devices, design engineers did not concern themselves with yield issues, which were considered to fall wholly in the fab's domain. Once the device was in production, it was the fab's responsibility to analyze and modify the process so as to bring up the yield.

Furthermore, yield issues were not significantly design-dependent. With the introduction of a new technology node, assuming that Design #1 had been brought up and the process flow had been tuned for maximum yield, the fab's engineers were reasonably confident that subsequent designs could be fabricated with minimal problems. In the case of today's technology nodes, however, such assumptions no longer hold true. For example, it is now possible to bring Design #1 up and tune the process flow. When Design #2 is introduced to the fab, however, the yield may fall dramatically or the device may fail in its entirety. In many cases, this is because the netlist structural characteristics of the second design impact the way in which it was laid out, interfering with the manufacturing process in a non-friendly way.

Over the last few years, DFM/DFY has been receiving a lot of attention. The problem with the early technologies is that they were applied post-layout, which often resulted in negative impacts on timing, power, and signal integrity. More recently, DFM/DFY analysis tools started to move upstream into the physical design portion of the flow. Although these analysis tools help designers measure new effects, which facilitate the creation of better designs, until recently the actual implementation process has been by hand.

History has shown that there has to be sufficient pain before a new technology takes hold. When timing analysis tools first emerged in the design flows, for example, they were certainly useful, but they weren't commonly used until they were tightly integrated with the logic synthesis and physical layout engines. Similarly, signal integrity analysis is useful in its own right, but it didn't take hold until it was tightly integrated with the implementation and optimization engines.

This is where DFM/DFY is today. Point analysis tools have proven themselves to be extremely useful, but they cannot reach their full potential until they are tightly integrated with the physical layout implementation and optimization engines. By integrating physical design tools with DFM/DFY analysis and simulation tools, manufacturing effects can be *during* the design implementation flow, concurrently with timing, power and signal integrity issues.

Rules-based versus Model-Based

A variety of techniques are currently employed to increase manufacturability and yield. These approaches are generally considered to be *rules-based* or *model-based*.

Rule-based Techniques

The first DFM/DFY tools were rules-based. The term “design rules” refers to a collection of rules that must be met by the physical design engineers and their tools. Examples of these rules would be the minimum width of wires and the minimum spacing between wires. One problem is that the number of such rules is increasing dramatically with each new technology node. In the case of the 180 nanometer node, for example, there were typically only a few dozen such rules, while today's 65 nanometer node can have several thousands of rules.

In many cases the rules are so restrictive that the result is to guard-band the design, leaving a significant amount of performance on the table. In some cases, the design ends up being so guard-banded that it is impossible to achieve its original performance goals. Even worse, the complex relationships between different manufacturability and yield mechanisms means that in many cases it is simply not possible to actually formulate an appropriate rule in such a way as to be meaningful to the design tool.

Model-based Techniques

Recently, DFM/DFY applications have started to apply model-based techniques. This may include, for example, modeling the way in which light will pass through the photomasks and any lenses; how it will react with the chemicals on the surface of the silicon chip; and how the resulting structures will be created.

The Best of Both Worlds

In reality, DFM/DFY tools need to use a mixture of rules-based and model-based techniques, as appropriate. Using model-based techniques for certain tasks allows the number of rules to be reduced and the remaining rules to be simplified, while still returning a much higher quality of results in the final chip layout.

Summary: The Solution is in the Routing

A very important point that is often overlooked is that the "D" in both "DFM" and "DFY" stands for "Design." On this basis, post-processing the GDSII files to correct problems introduced by the upstream design tools cannot truly be considered to be DFM/DFY.

The solution is to bring DFM/DFY upstream into the design process; to create a design that is correct by construction; and to hand-off a design that is as manufacturing-friendly and yield-friendly as possible. The obvious candidate to subsume DFM/DFY analysis and implementation is the routing engine, the "front door" into the manufacturing process. For ASICs, routing currently accounts for approximately 60% of design delays.

Conventional routing engines use only rules-based techniques. However, the number of design rules and recommended rules is increasing so dramatically with each new technology node – and the rules themselves are becoming so complex – that these engines are choking on the sheer number of rules. The answer is to bring model-based techniques into the routing domain and to use both rules-based and model-based techniques, as appropriate.

Conventional routing engines, unfortunately, simply are not architected to be capable of addressing these issues. What is required is a completely new routing architecture. A routing engine based on this new underlying architecture should be capable of performing as many DFM/DFY-related actions as possible. These actions include, but are not limited to, wire widening, wire spreading, redundant via insertion, minimizing jogs, and automatic metal fill.

But even this isn't enough. It is not sufficient to treat each layer in isolation with regard to techniques such as wire spreading and automatic metal fill. Inserting fill on one layer may address the problems of CPM and unequal etching effects for that layer, but fill areas on one layer may combine with fill areas on adjacent layers to act as capacitors, thereby impacting timing, power, and signal integrity.

Thus, the router should be capable of performing 3D wire spreading (distributing the wires evenly across each layer and across all layers) and 3D fill insertion (inserting fill using a multi-layer-aware algorithms). While performing any of its activities, such a next-generation router must take full account of timing and signal integrity effects like noise and crosstalk. Such a router must be capable of performing multi-variable, multi-value optimizations, and it must also be capable of making decisions such as when to use minimum width wires or when not to introduce redundant vias in certain cases. And, in all cases, such decisions must be made in the context of desired yield.

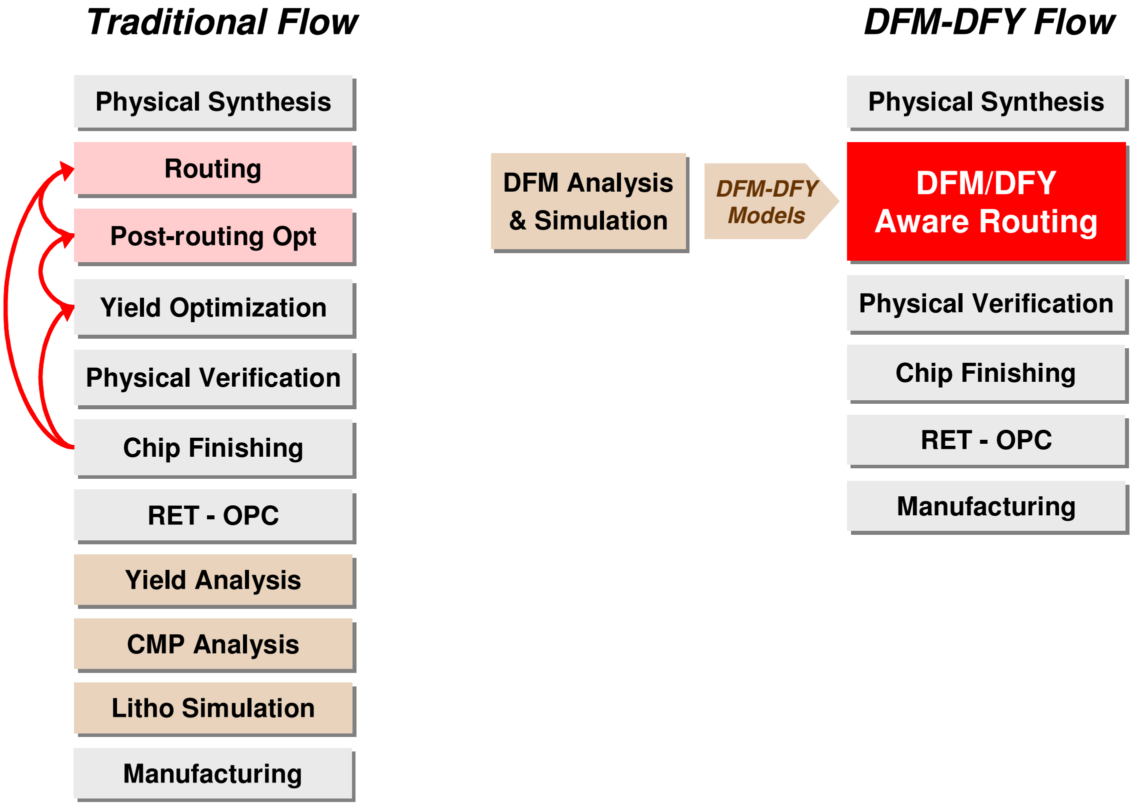

Figure 2. Traditional flow compared to new DFM/DFY aware flow

This new routing engine must be able to use all of the above techniques and be able to tune differently for different portions of the chip so that the difference in structural characteristics among different blocks in the chip may be addressed effectively and so that variability and yield loss among different designs can be minimized.

Equipping designers with such a DFM/DFY-aware routing engine will ensure that, when we encounter calm weather, it is because DFM/ DFY issues have been addressed, not because we are in the eye of the storm.

About the author

Mitch Heins, vice president of engineering at Pyxis Technology, has worked in the semiconductor and EDA industries for over 24 years, at companies such as HPL Technologies, Inc., Petersen Advanced Lithography, Cadence Design Systems, Ambit Design Systems, the CAD Framework Initiative, Inc., (a consortium focused on software standards for plug-and-play EDA tools, now called the Si2, the Silicon Integration Initiative) and Texas Instruments.