Shifting Methods: Adopting a Design for Manufacture Flow

John Ferguson, Ph. D.
Mentor Graphics Corp
Wilsonville, Oregon
john_ferguson@mentor.com

Abstract

Unlike traditional design rule checks (DRCs), which have a clear pass or fail definition, yield issues, such as density and antenna checks, are dependent on a number of variables, resulting in a yes, no or maybe paradigm. These issues, which have a major impact on total chip yield, have typically been identified as design constraints and embedded in DRC rule files. But there is little in the way of information on how a change in layout relates to overall improvement in yield. With advanced processes of 130nm and smaller, designs that are verified DRC clean can still result in poorly yielding or even non-functioning silicon. For this reason, a new method of communication is needed between designer and manufacturer for determining yield issues. Rather than being provided information on a simple pass/fail basis, designers need immediate access from the manufacturer to how various layout characteristics impact chip yield. This new communication loop is the first and necessary step in adopting and implementing a design for manufacture (DFM) flow. Such an approach identifies trouble spots and provides the important data that allows the designer to determine a cost/yield analysis. Designers will be deciding on and implementing a "fix or fab" methodology, resulting in greater yield and control.

1. Introduction

In the not so distant past, integrated circuit designs could be manufactured as drawn and predicting yield was a simple matter: if each structure could be manufactured, then the whole chip could be manufactured. All it took to successfully hand-off a layout to the foundry was to produce a clean DRC (design rule check) and LVS (layout vs. schematic) run based on the foundry's yes/no or pass/fail constraints. These rule files help the designer understand the limits imposed on them by the manufacturing process. Most of these constraints represent true process limitations, which, if not followed, will guarantee non-functioning silicon.

Unlike traditional design rule checks, which have a clear yes/no or pass/fail definition, yield issues are more difficult to pin down, because they are dependent on a number of variables. Traditionally, yield related checks have been confined to planarity issues (the difference in oxide heights for given region on a design, typically identified in the form of density and metal slotting rules), and antenna effects (charge accumulation of interconnect components). [1]

For most IC manufacturing processes, the only insight into yield issues appears in these design rules or guidelines. But with advanced processes of 130nm, 90nm and smaller, designs that are verified DRC clean can still result in poorly yielding or even non-functioning silicon. As designs grow more complex, process technologies become smaller and geometry counts increase, the work required to achieve acceptable yield becomes increasingly demanding and difficult. And given that mask costs are now in the $1 million range, the price of failing at first silicon is high. It is therefore important to develop methods of analyzing cost/yield issues.

2. Changing traditions

Historically, yield limitations and defects were non-parametric and random in nature. Despite the best clean room efforts, particles still found a way to land on chips, creating shorts or opens. To reduce particle problems, manufacturers developed a critical area analysis that estimated the probability of defects in relation to particle size and number of spacings. Recommendations could then be made about limiting spacings in problem areas. This was the beginning of the definition of design rules.

With the advent of smaller geometries and copper processes, random defects and particle contamination were no longer the major issue. Instead, non-random events became major sources of yield limits. (Fig. 1)

Parametric issues relating to the way designs were being created became concerns. These issues could affect an entire chip:

- Copper processes that are more susceptible to planarity issues
- Via creation that is more susceptible to opens
- Layout configurations that impede or prohibit the use of required RETs

0-7695-2093-6/04 $20.00 © 2004 IEEE

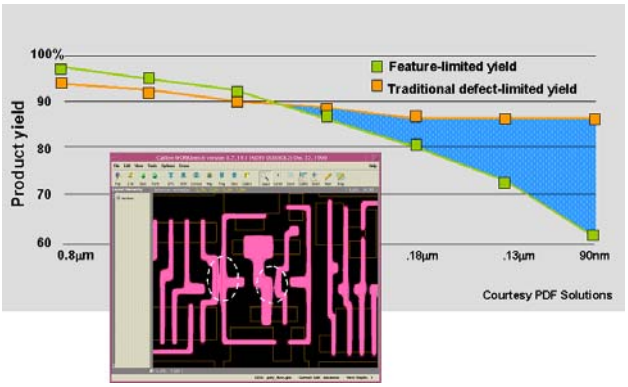

Fig. 1. Particle defect-driven vs. feature-driven yield. Failure to form features replaces particle defects as the major defect problem.

Copper processing is a less manageable manufacturing process than the former, well-known aluminum process. Via opens, due to the copper cladding process, are a result of voids traveling down into the via. Some success can be gained by populating the area with multiple vias so that if one via or area of via failed, compensation could occur.

Density issues also arose. Dramatic variations in cross-sectional thicknesses, or "dishing", due to the CMP process, had pronounced effects on yield. Some relief is found in slotting and filling to compensate for the dishing.

At 180nm, aggressive resolution enhancement technologies (RET) were introduced in order to produce designs as drawn. But the ability to apply RETs appropriately was, and remains, dramatically limited by the design layout itself. Layout designers, traditionally given minimum specifications, will design to those specs as a matter of best practices. However, minimum specs may not allow RET. If specs are extended beyond the minimum, the likelihood of failure decreases. For instance, poly end caps "pull back" of the gate when printed, making RET application difficult. Extending the end cap to compensate for pull back results in a layout that is more RET-compliant, thereby dramatically improving the overall yields that may be achievable for the design. (Fig. 2)

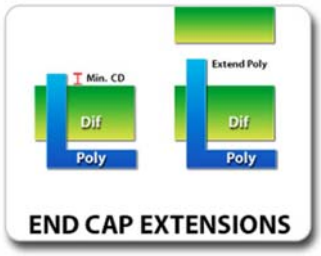

Fig. 2. In DFM, the optimal layout for highest yield must be implemented in a RET-aware environment, as with the extension of poly end caps.

Traditionally, foundries have managed yield, identifying issues in the form of yes/no or pass/fail based on design constraints embedded in DRC rule files. Current methods for IC design and manufacturing assume an acceptable level of yield. But the number of yield limiting issues is growing dramatically and can no longer be confined to a small set of constraints, as with historical examples, such as planarity and antennas.

Unlike traditional DRCs, however, where a specific instance can be defined for what can be physically manufactured and what cannot, yield related rules are given a value based on what is believed will result in acceptable yield. This is often based heavily on experimentation and experience in manufacturing, and takes into consideration the post-layout application of yield-enhancing techniques such as rule-based OPC (optical and process correction), scattering bars, antenna violation detection and planarization fill. (Fig. 3)

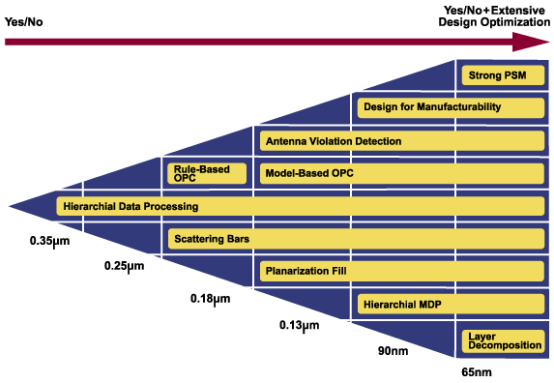

Fig. 3. Yield issues have traditionally been identified through design constraints (yes/no) embedded in DRC rule files. With advanced nanometer technology, extensive analysis of yield issues is required.

However, with advanced processes of 130nm and smaller, designs that are verified DRC clean can still result in poorly yielding or even non-functioning silicon. This is because even though a particular design element or configuration in and of itself can be physically manufactured, in the context of a specific full-chip layout, the same manufacturable element may dramatically increase the probability of failure. For example, many designs today are automatically generated specifically to meet the minimum design constraints, with minimum width and spacings. This is believed to help keep the total design size as small as possible. But although a specific placement of a minimum-spec configuration may be manufacturable in and of itself, having many such situations dramatically increases the chances of failure on that chip. By knowing how many such situations exist on a given layout, and whether any can be extended beyond the minimum configuration to a preferred configuration, may allow for a higher overall yield without dramatically effecting design size or functionality.

3. Cost/yield analysis: from yes/no to fix/fab

A layout characteristic given a "no" or "fail" may, in reality, result in a layout that achieves adequate yield, or with very little effort, result in a "yes" or "pass." Inversely, a characteristic given a "pass" may be so close to failure that it could cause the entire chip to fail within a short time. But in order to analyze the issues and make determinations, designers need to be fully informed about layout characteristics. Currently, however, designers have little or no information about just how important a given issue is for the specific design yield, or about how a change in the layout relates to overall improvement in yield. And in order to develop new approaches in achieving acceptable yield, a new method of communication is required for defining and relaying manufacturing constraints, verifying IC layouts and addressing manufacturing related issues to the designer. This requires a shift in methodology.

The possible improvement in manufacturability must be weighed against the effects in performance and functionality of a design, designers will need to be involved in analyzing the variables. *This means that manufacturability starts at the chip design level.* In a DFM flow, the designer will determine and control which issues to "fix or fab." (Fig. 4)

Fig. 4. Armed with complete layout characteristics, designers can weigh the work required to fix the yield-limiting issue vs. the benefits realized in yield.

4. DFM a full-chip issue

Like that of traditional DRC manufacturing methods, DFM is largely a full-chip problem: *data must be made available in its full context*. This means having access to DFM yield limiting issues in a cross-layer and cross-hierarchical sense. Having the ability to look across hierarchical boundaries to see how the data in one cell interacts with data outside the cell is essential. It may be possible to improve the manufacturability of one layer by manipulating another. Similarly, a cell with no known manufacturability issues may significantly impact the manufacturability of a full-chip when it is placed into context.

To implement analysis, a method of defining levels of severity must be in place. This requires a method by which the author of a configuration file can define the issues of interest and associate each of these issues with a quantifiable level of impact. For example, by calculating the number of metal transitions with one via divided by all the metal transitions in the chip, a designer can have a better feel for the total impact this issue may have on their chip yield, as pertains to an acceptable limit set by manufacturing. Similarly, this issue can be weighted in merit against other DFM related issues, resulting in a total 'grade', representing how well the design layout can be manufactured.

Naturally, the knowledge of how to determine the appropriate grading levels must come from experience with a given process. In a historical model, test structures were generated in the early life span of a process. These structures were fabricated and analyzed. Those structures with high rates of manufacturing issues resulted in the quantification of design rule constraints. As the process matured, and production designs were fabricated, this level of knowledge was improved, resulting in new design rule constraints.

For a design for yield paradigm, it is expected that the same mechanism of quantification would be expanded.

Each particular issue must now be quantified to a greater extent to determine the likelihood of failure for a given geometrical pattern. This is not a trivial task and will require new models and tools to support this effort.

Like the historic case it is expected that this grading level will also evolve and improve over the life of a process. Data collected, post-test, from manufactured silicon lots, must be analyzed to further determine the most common modes of failure. Having the ability to map design failures back to geometric layout issues, and quantify these issues by severity will also require new methodologies and supporting tool infrastructures.

With such a grading mechanism in place, rather than simply getting information on a yes/no basis, designers have access to how various layout characteristics impact the chip yield. The grading of design elements gives CAD managers a way to determine a statistical impact to overall achievable yield and control over the business decision of whether or not to "fix or fab." This results in a yield greater than that created by simply meeting the design rules and guidelines.

5. Accounting for post-layout applications

Resolution enhancement techniques (RET) cannot be considered independently from yield-limiting considerations. Certain forms of RET, such as phase shift mask (PSM) and off-axis illumination (OAI), have requirements on pitches. (A pitch is essentially half the width plus the spacing of the polygon in question.) Today, this is most problematic for poly gate transistors. With transistors of many different spacings placed throughout the design, manufacturing can be difficult, if not impossible, as adequate RETs are significantly constrained. But if a design has only a handful of pitches, then it becomes much easier to manufacture in a manner that results in acceptable yields.

With this information available, the design team can determine the best course of action by gauging the amount of work required to implement a change versus the improvement in yield they will achieve by doing so. Getting this information should be part of a *feedback loop from the manufacture*. However, in order for the manufacturer's feedback to be useful, it must also be intuitive to the designer within the context of the original design environment. Manufacturers work in a post-layout environment, analyzing the layout's manufacturing issues in GDSII format. GDSII is where RET, metal fill, slotting, redundant via insertion, and other post-layout modifications are made.

Most IC layout creation tools operate in design frameworks, which can contain significantly more information than is required in a manufacturing GDSII hand-off. In DFM methodology, the designer will need to understand the impact that the manufacturer's additions have on the original design specifications, as cost/yield decisions must be made within the context of the original design environment. If the designer's original environment is not GDSII, an integration solution is needed that bridges post-layout GSDII back to the layout design environment. Such a solution would have to have the capability to read/write in the design database format, annotate edits/changes back to the design environment, and allow access to design analysis tools.

Within such a bridge, the designer would be able to implement post-layout changes within the original layout environment. (Fig. 5) This would give the designer and CAD manager the ability to re-analyze and verify the impacts of changes created. It also provides a method of retaining IP that is more yield compliant. All of these requirements can be achieved today, due to recent efforts to open the once proprietary design environment databases.

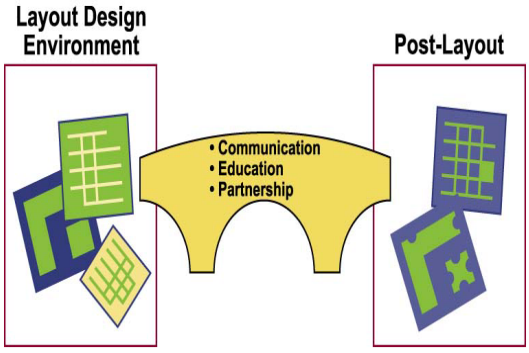

Fig. 5. A communication loop between manufacturer and designer that is intuitive to the design environment must be implemented in order for successful DFM methodology to take place.

6. Summary: The future of DFM

Because of uncertainty of yield, companies are reluctant to manufacture at smaller process technologies, such as 90nm. Most are still manufacturing at 180nm, where yield and cost is relatively predictable and controllable. As companies ramp up to 130nm and smaller, the need for managing yield issues is greater than ever. Underlying this issue is the need to maximize yield without increasing manufacturing costs. It is essential to have enough yield issue information to make a business decision about what is a cost-effective solution.

Nanometer technology is forcing a change in the way fables companies and foundries communicate. Hard and fast assumptions about what is manufacturable, and what is not, may no longer be valid. Designing for manufacture will require a new level of communication, education and partnership between design companies and foundries. The new method will definitely alter, or even completely

0-7695-2093-6/04 $20.00 © 2004 IEEE

depart, from the current design and manufacturing flow. To not shift methods will have undesirable consequences: acceptable yield in advanced nanometer technologies will not be achievable.

Fortunately, the core technologies required to build this design-for-yield flow exist now. Tools capable of reading and analyzing layout topography, in a manner that preserve hierarchy for upstream and downstream analysis are currently available, with integrated links to many existing design environments. The only new requirement, underway at advanced semiconductor companies, is a more robust communication mechanism that allows designers to be fully informed in order to make decisions about yield.

7. References

[1] J. Ferguson, "Turning up the Yield",
IEE Electronics Systems and Software, June/July 2003