DFM Rules!

Chair
Naveed Sherwani
Open-Silicon

Organizer
Susan Lippincott Mack
Susan Mack Marketing

Alex Alexanian
Ponté Solutions
Mountain View, CA

Premal Buch
Magma Design Automation
Sunnyvale, CA

Carlo Guardiani
PDF Solutions
Desenzano Sul Garda, Italy

Harold Lehon
KLA-Tencor
San Jose, CA

Peter Rabkin
Xilinx
San Jose, CA

Atul Sharan
Clearshape Technologies
Sunnyvale, CA

ABSTRACT
For sub-100nm processes, predictions are putting initial process yields in the single digits. At the same time, at 130nm, we saw that two chips designed with the same methodology and same design rules could deliver completely different manufacturing yields.

This panel will discuss the reasons for these phenomena and talk about future trends in DFM that will need to be addressed for success below 100nm.

Categories and Subject Descriptors
B.7 Integrated Circuits

General Terms
Design, Economics

Keywords
Design for Manufacturability, Yield Optimization

Position Statements:
Panelists have provided their position statements (see below) outlining their opinions. In the panel, each panelist will give no more than three specific takeaway points that designers must consider in developing a design for manufacturing or design for yield solution.

1. Alex Alexanian:
There are three important points to understand in achieving high-yield chips below 100-nm. First, design rules no longer work. If design rules are enough, then we wouldn't have low yields in the first place. Two chips following same design rules may yield differently, as a result of growing gap between design and manufacturing. Statistical model-based verification solutions will complement or replace standard DRC approaches. Statistical yield models describing specific failure mechanisms can be used to characterize the future silicon at the design stage. Secondly, we require Open Yield Modeling Standards. The industry needs public and open yield models with private and closed process parameters. Open standards will allow the entire semiconductor industry to contribute to the creation of effective statistical yield models. Closed process parameters will respect the process IP of different manufacturers and keep it confidential. Finally, we need a Yield-Driven Design Methodology. Similar to timing considerations that initially started from the very back-end and grows up to the synthesis flow, yield considerations must exist at every stage of IC physical design cycle. IP Vendors must provide yield-centric variations of library cells and custom IP blocks, fabs must provide statistical yield models. Designers must be equipped with Design For Yield (DFY) solutions that provide analysis, prediction and optimization capabilities at synthesis of gate level netlist, placement, routing and full custom layout stages.

2. Premal Buch:
As semiconductor technology moves into 90nm and below processes, manufacturability and yield have become major concerns for the industry. *Manufacturing trends such as sharp reduction in fabrication yield and increasing difficulty in yield ramp-up are affecting the bottom lines of everyone involved.* Traditionally, IC design flows have been shielded from the intricacies of the fabrication process. Manufacturability is ensured by rigorously following design rules and by applying resolution enhancement techniques (RET) like optical proximity correction. With shrinking feature sizes, there has been an explosion in the number and complexity of design rules, runtime and data size for RET and the cost of mask. *Traditional techniques no longer work, and manufacturability and yield issues become uncontrollable unless they are considered early in the design flow.* Considering yield in the design flow requires concurrent optimization: techniques like wire spreading and metal fill impact timing and signal integrity; via minimization and redundant via insertion impact routability; cell yield optimization involves trade-off with area; parametric yield optimization requires modeling uncertainty in cell and wire delays. Designers require physical and interconnect synthesis solutions that will allow them to make concurrent trade-offs between these metrics to achieve complete design closure with improved yield. We believe this is the only way to make DFY a reality -*Yield needs to become another metric in RTL-to-GDSII flow along with timing, power, area and signal integrity, and an effective DFM solution needs to be tightly integrated within the design implementation flow.*

Copyright is held by the author/owner(s).
DAC 2005, June 13–17, 2005, Anaheim, California, USA.
ACM 1-59593-058-2/05/0006.

3. Carlo Guardiani:

At 130nm and below, the set of design rules required to guarantee a specific target, or "yield entitlement level" is so complicated that it can be extremely penalizing in terms of die area, product speed and power to enforce all of the rules at once. So some rules are now considered "recommendations" that designers should try to comply with as much as possible yet still meet design specs in terms of performance/die size. The responsibility for yield is now in the hands of design engineers, who need the means to quantitatively measure the yield impact of adopting a particular recommendation, or a way to trade-off a particular rule against another in case of conflicting recommendations. But yield is not a big enough concern of the designer that they'll switch or learn new tools. Giving yield information to designers through their familiar SP&R tools lets them perform tradeoffs along with timing, area, and routability. It all comes down to calibration and characterization. How do you characterize silicon? You've got to know at the microstructure level what fails and what doesn't and how many times it fails, looking at the statistical occurrence of microstructure failures versus the process variability that might come into play during the manufacturing process. The designer takes advantage of these silicon-validated characterizations by minimizing prone-to yield microstructures, or relative positions, and maximizing those microstructures that seem to have the least amount of variability. It will become increasingly necessary for design tools to understand this data so users can make informed tradeoffs to get the right mix of yield, timing and power for a given design.

4. Harold Lehon:

"Rules" will cost a bundle. Design re-spins are a significant and often uncontrollable variable to the "cost" equation. Rules are running out of steam for predicting yield. Rules-based OPC started to fail at the 130nm node, and model-based OPC is now required below 130nm. Model-based design practices that consider the transformation of the verified layout into silicon will be required. The move to model-based design inspection that considers the transformation from the RET/OPC process, reticle process, aerial image formation, resist process and etch process will be required to achieve entitlement yield sooner.

Systematic yield loss is increasing and is dominated by feature limited yield. The largest opportunity for yield improvement exists in lithography where feature size, shape, and location interact with the imaging and film processes. Variations in focus and exposure of a scanner can cause yield loss do to marginal layout or OPC which don't consider process variability. Variations in focus and exposure of all litho processes do result in yield loss when the design layout and RET/OPC have not been optimized to account for these process variations. Collapsing process windows cannot be solved with rules. "Litho Aware" design inspection will be required to model the characteristics of the reticle, the aerial image, and the resist and etch processes. Verification must also be expanded to a full process window analysis on at least nine FE points. In short, we need to achieve optimized process windows through fast model-based verification.

The design insight representations and design inspection results then need to be fed downstream into the manufacturing flow. Design-aware process control practices will close the loop while enabling increased cycles of learning. This will ensure manufacturing processes are better controlled which will result in improved device parametric yield."

5. Peter Rabkin:

To establish a strong DFM flow, we must first start from sophisticated design rules (DR) that reflect process capabilities and design specifics. We may need evolving DR sets to reflect process maturity within a technology generation development and manufacturing cycle and we may need to look at different DR sets for different applications (logic, ASICs, FPGAs, etc.) or even within a single application (e.g. logic & memory). Secondly, lithography compliance checks (LCC) for advanced technologies are a must. The litho compliance space does not always overlap with DR compliance space. In these cases, a DRC violation may be waived if/when layout is LCC clean. Finally, we must strive towards a more comprehensive implementation of DFM. DFM rules/guidelines should be a joint work by fab and a design house to reflect process and design specifics, previous experience (technology generations and/or products), analyses of test chips, anticipation of process capabilities at maturity and any known yield-limiting factors. We need integrated DRC/DFM rules and litho checks with a determination of priority for layout fixes (optimize "as you go"). Additionally, we need tools and methodologies that will enable us to perform DFM analysis of improvement opportunities for completed pieces of layout AND perform automated hierarchical layout fixes. Additionally, we must address increased process and transistor variability on the circuit and system design level to deal with parametric yield loss. Fabless companies require sophisticated DFM modeling, calibration, and simulation capabilities that work reliably under the conditions of insufficient or uncertainty of process technology data.

6. Atul Sharan:

At 90nm, designers are in deep trouble. Not only are design rules insufficient to guarantee manufacturability, but manufacturing variations cause a large performance spread. Even with the most efficient RET, manufacturers have to perform post-tapeout modifications to make designs litho-friendly. Designs do not perform as expected due to variations that were not accounted for during design. We've lost the shape closure (ability to predict the end silicon image) and the performance predictability which were the foundations of the contract between manufacturing and design. To restore shape closure, design teams need lithography-RET predictive design tools because design rules alone just won't cut it. Because of complex lithography effects, rules simply do not contain enough information to drive sub-100nm manufacturability. Bringing OPC in the design flow is just not a viable solution because RET/OPC is too computationally expensive and would disclose too much manufacturing IP. Therefore, we need new "discontinuous" technical innovation that does not disrupt the design flow. It must offer lithography-RET predictability and enable designers to control and optimize design manufacturability - without having to own RET. In order to restore performance predictability, design flows must now account for the systematic variations that are inherent in manufacturing. RET predictive tools can identify systematic variations across the process window to deliver an accurate picture. Again, we require new "non-disruptive" technologies that manage the impact of this variability on chip performance rather than expensive over-margining. Only then can we restore the contract between manufacturing and design and decrease cycle time, ensure single-pass success and regain manufacturability.