Design for manufacturability: a fabless perspective

Jason P. Cain
Advanced Micro Devices, Inc.
7171 Southwest Parkway, Austin, Texas USA 78739

ABSTRACT

Design for manufacturability (DFM) has become a key enabler of integrated circuit (IC) production over the past decade. In this paper a comprehensive DFM program for IC designs at the 28nm node and beyond is described from the perspective of a fabless design company. Challenges for future technology nodes are also explored.

Keywords: integrated circuit (IC), design for manufacturability (DFM), lithography

1. INTRODUCTION

As semiconductor manufacturing processes have continued to scale to smaller dimensions the conventional method of ensuring that an integrated circuit (IC) will be manufacturable by adhering to a set of rules set forth in the design rule manual has become insufficient. A broad range of tools and methodologies grouped under the name "design for manufacturability" (DFM) have been introduced to try to close this gap.

The semiconductor business model has also evolved. From the early days of the industry most IC makers fell into the category of integrated device manufacturers (IDMs) who handled the full product life cycle from design to manufacturing internally. The continually rising costs of building a semiconductor fabrication facility ("fab") with the large size and cutting edge equipment needed to take advantage of economies of scale created an opportunity for a new type of semiconductor company: the foundry. As a result, many existing IDM companies have divested their manufacturing operations to focus on the product design, choosing to outsource the manufacturing to one or more foundries and becoming "fabless." In addition, this has lowered the barrier of entry into the semiconductor industry, allowing new companies to enter the market without needing to build expensive fabs.

While it has long been the practice of IDM companies to co-optimize the design and the process technology (including DFM methodologies), this process looks somewhat different when the design and process development organizations reside in separate companies. This paper will provide a brief tutorial on the IC design flow, then provide an overview of how design for manufacturability is practiced inside a fabless semiconductor company. Finally, challenges for DFM at the 20nm technology node and beyond will be discussed.

2. INTEGRATED CIRCUIT DESIGN FLOW

2.1 Overview

The IC design flow is an extremely complex process consisting of many steps in which the design is refined from a set of high level specifications into physical layout data that can be used to produce lithography reticles for patterning of silicon wafers. Some of the key steps in this flow are illustrated in Figure 1 below.

The design process typically starts with a description of what specifications the IC must meet. This may include information such as power and performance constraints, instruction set to be implemented, memory type, off-chip interfaces, and other details of the architecture. This specification is then implemented as a logical design, typically in the form of a register transfer level (RTL) description. The RTL may be thought of as describing the way in which information will be processed by the chip and serves as an input for subsequent stages of the design flow.

For some design types (e.g., circuits containing analog components) an additional circuit design step is needed. The output of this step is typically a schematic of the circuit including component devices with all key parameters (e.g. transistor gate width and length) specified. The schematic may also serve as an input for additional stages later in the design flow.

Design for Manufacturability through Design-Process Integration VII, edited by Mark E. Mason, John L. Sturtevant,
Proc. of SPIE Vol. 8684, 868403 · © 2013 SPIE · CCC code: 0277-786X/13/$18 · doi: 10.1117/12.2015173

The final stage of the IC design process is the physical design step. This is the process in which the RTL and/or circuit schematics are given a physical implementation. This physical design may then be converted into a set of lithography masks to enable production of the chip. It is during the physical design stage that most design for manufacturability functions are inserted.

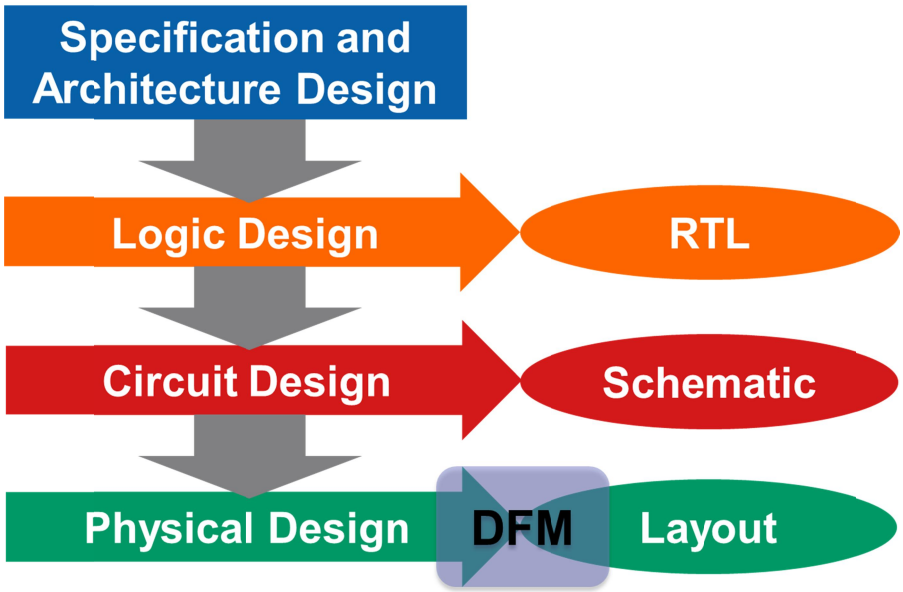

Figure 1. Semiconductor integrated circuit (IC) design flow. Enablement of design for manufacturability takes place predominantly in the physical design stage.

2.2 Physical design hierarchy

A system-on-chip (SOC) design is typically composed of many different design blocks which are created individually and assembled hierarchically to create the final chip. At the lowest levels of the hierarchy are blocks often referred to as foundation intellectual property (IP), which includes objects such as standard cells and memories. Standard cells typically implement a simple Boolean logic function (AND, OR, inverter, etc.) or a storage element (flip flop or latch). These IP elements are often used thousands or millions of times within a single product and therefore they are typically manually designed and carefully optimized. Other circuit elements may also follow a similar custom design procedure, including performance-sensitive circuits or analog and mixed-signal circuits. Digital blocks are typically designed using an automated flow (described in section 2.3 below). Finally, these design components are assembled into an SOC to implement the complete product specification. Each of these levels of the design hierarchy has different opportunities for insertion of design for manufacturability.

2.3 Digital design with automated place and route

The most common method for creating physical designs for large IP blocks implementing a digital circuit is to use a synthesis, automated place and route (SAPR) flow. This design methodology is depicted in a very simplistic way in Figure 2. The RTL from the logical design is used as the input to a software tool that performs the synthesis step. This step produces a netlist that contains the different logic gates and other digital circuit elements and also specifies the electrical connectivity between them.

Once the synthesis step is completed the netlist may be passed to an automated place and route tool along with a library of standard cells that contain the physical layout implementation of the basic logic gates, storage elements, etc. The place and route tool then creates a physical design that places the standard cells within layout and generates metal wire connections between them to match the input netlist. This process is subject to a number of constraints, including area, circuit timing, power, and physical design rules.

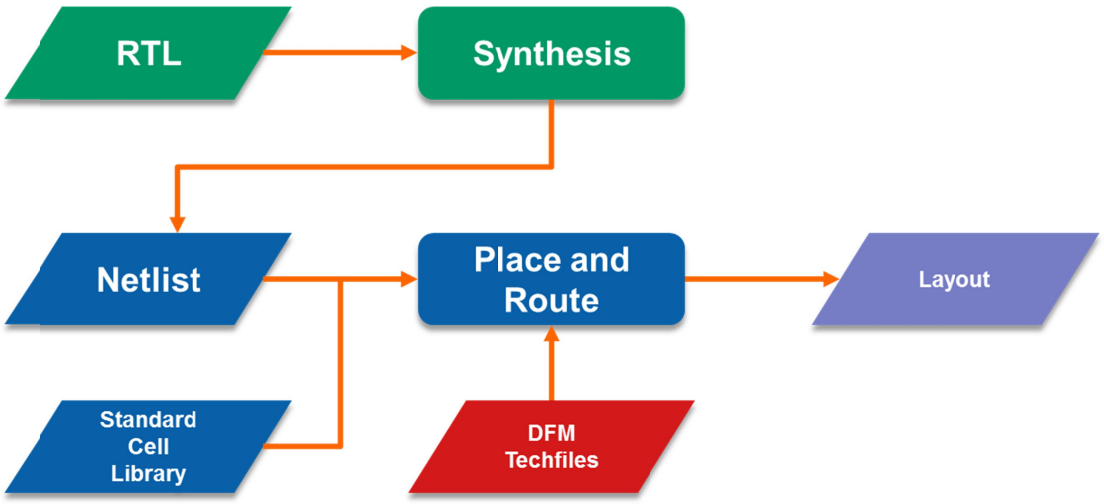

Figure 2. Automated place and route flow for digital design.

The primary opportunity for inserting DFM into the SAPR flow is via the "techfiles" that control the behavior of the place and route software. For example, the router tool can be given a prioritized list of via options for making connections between metal layers. The top priority option would be identified by the foundry as being the lowest risk with respect to manufacturability, but would also typically require the largest area. The lowest priority options would typically be the connection that contains a single, non-redundant via with the minimum amount of metal enclosure allowed by the design rules, but would also occupy the smallest area. In this way the router can make tradeoffs between area and manufacturing risk, inserting highly robust via connections where the area constraints are minimal and using the more risky vias only in areas of the layout where high metal density requires it.

The router can also be tuned to avoid known yield detractor patterns in the metal routing. By supplying the router with a database of patterns to avoid, the router can use iterative methods to identify such patterns in the metal routing and then reroute the metal wires in a different configuration to avoid the particular yield detractor pattern [1].

The placement tool can also be made "DFM-friendly" by making it aware of certain preferred standard cell placement configurations. For example, if certain standard cells are known to produce manufacturability or performance marginalities when placed next to each other the placement tool can be instructed to avoid such combinations where possible or to insert extra space between the cells.

2.4 Custom IP design

For some circuits (including performance-sensitive circuits and analog or mixed-signal blocks) the SAPR flow may not be appropriate and a custom design process is needed. In the custom design flow a layout design engineer typically creates a rough version of the layout (following a schematic) and then uses an iterative process to make adjustments to the layout to satisfy design rule checks (DRC), layout versus schematic (LVS) checks, and timing constraints. This presents an opportunity to add DFM checks into the flow in addition to the conventional physical verification checks. These could include DFM scoring (see section 3.3), yield detractor pattern checks (e.g., lithography "hotspots"), or pattern density checks among others. Enforcing requirements for each of these checks helps ensure that the design has been optimized for manufacturability from the outset before integration into the SOC.

2.5 System-on-chip (SOC) integration

As modern central processing units (CPUs) have evolved to incorporate many components that were formerly separate chips into a single system-on-chip (SOC) a design methodology has evolved to assemble and verify such designs. As more and more components have been pulled into the SOC the designs have become increasingly complex. An example SOC design implementing an accelerated processing unit (APU) including CPU cores and a full graphics processing unit (GPU) is shown in Figure 3. It is important to note that the design IP blocks included in the SOC may come from a

variety of sources: some from SAPR flows, some from custom design flows, potentially even some from external IP suppliers. This makes it critical to ensure that the complete SOC design adheres to DFM standards in order to ensure that the final product will be manufacturable.

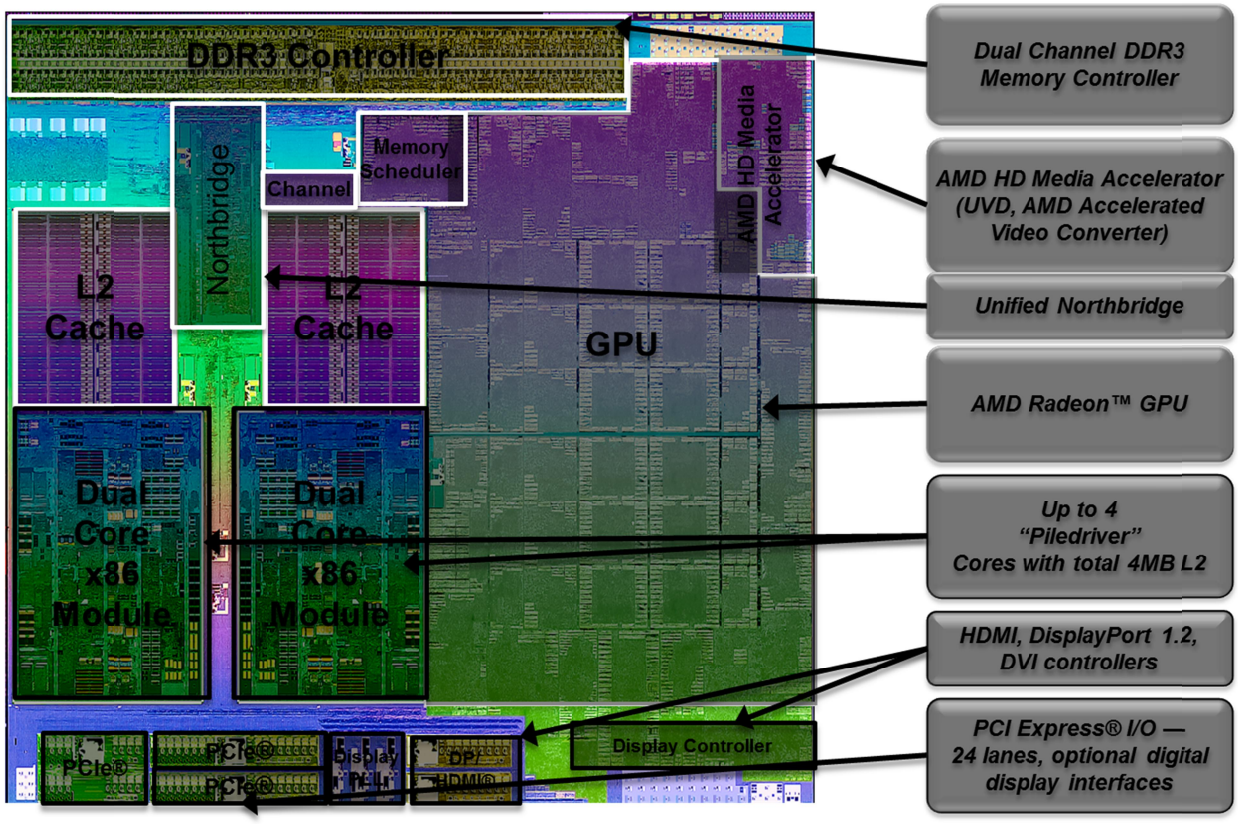

Figure 3. Die photo and floorplan block diagram of AMD "Trinity" Accelerated Processing Unit (APU) showing the complexity of modern SOC integration. This chip is manufactured in a 32nm silicon-on-insulator (SOI) process and contains over 1.3 billion transistors in an area of 246 mm^2. From [2].

2.6 IP reuse strategy

The system-on-chip design methodology involves the integration of multiple design IP components, each of which may be developed independently and reused across multiple ICs. Some commonly-used components may be designed once at the beginning of a technology node transition and reused for all products in that node. This places an extra burden to ensure that the design is manufacturable from the outset as any marginalities may persist in a large number of SOC designs.

2.7 Foundry collaboration

Collaboration between a fabless design company and the foundry partners is essential for success. For designs targeting the leading edge of a technology node the time between the availability of silicon data from test chips and the initial product design tapeout is typically very short. In order to allow for design-technology co-optimization a cycle of learning must be established in which layout blocks are sent to the foundry on a regular basis for evaluation via process simulation. This can help identify layout constructs that are problematic and either help the foundry process engineers optimize the manufacturing process to improve the manufacturability of the construct or enable the fabless designer to remove the construct from the layout. In addition, this cycle of learning can help highlight the marginal areas in the layout for further monitoring and analysis on test chips and other early silicon.

The compressed schedule between test chip silicon availability and lead product tapeout places a high degree of importance on getting feedback from the test chip quickly in order for the learning to be applied for product designs. To that end, every effort must be made to apply design for manufacturability principles to the test chip design in order to minimize the time required to obtain functional circuit learning.

One additional area in which foundry collaboration is critical is the development and implementation of tools to enable DFM in the design flow. Having early access to these tools allows for development of use models, integration into the design flow, end user training, and feedback to the foundry to improve the tools. Investment of effort early on in the development phase plays an important role in ensuring that production-ready DFM tools are available at the beginning of the technology node cycle.

3. DFM IN THE PHYSICAL DESIGN CYCLE

Successful implementation of design for manufacturability requires that it be practiced throughout the entire IC product lifecycle, from early product definition (before the actual design start) to high volume manufacturing. Figure 4 shows the various roles DFM can take in relation to the design and manufacturing cycle. These roles are explored in more detail in the following sections.

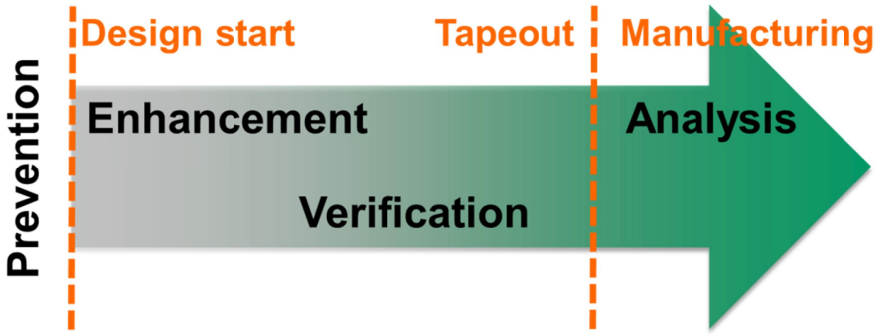

Figure 4. Roles of DFM in the physical design cycle.

3.1 Prevention of manufacturability problems

One of the most critical roles for DFM is ensuring products are manufacturable by preventing problems from being introduced into the design in the first place. The fundamental method for accomplishing this is through the design rule manual. Process technology learning from previous generations must be captured by appropriate design rules for the current generation. In addition, new manufacturability challenges for the current process generation must be anticipated and governed by new design rules.

Design rules alone are often not enough to achieve a robust design. A detailed list of layout "best practices" can help layout designers make their designs more DFM-friendly from the outset. This can include "recommended rules" from the foundry (which are not enforced as strict design rules), preferred via connection styles, and guidelines on how to use extra available space most effectively to improve manufacturability.

Finally, key layout constructs that are known to be problematic (whether due to lithography patterning problems or other process integration issues) can be identified and forbidden. Such constructs are typically identified through a combination of prior process technology learning, process simulation, and early silicon data from the current process. Often the most effective way to prevent these constructs from being present in the layout is to use a 2D pattern matching approach to identify the patterns and require designers to remove them from the layout [3].

3.2 Layout manufacturability enhancement

Once the physical design work has started DFM tools can be used to modify the layout to enhance manufacturability in both the custom layout and SAPR domains. One important example is the optimization of via connections between metal layers to minimize the risk of via failures (either open circuit or high resistivity).

In the SAPR flow this is often accomplished through "via swapping," in which the router first makes all necessary via connections using a basic minimum area via style, then makes one or more passes through the layout opportunistically swapping the basic via style for more robust connections in locations where the change will not result in an area penalty or introduce a design rule or circuit timing violation. The via swapping could include replacing a single square via with two vias to make a redundant connection or with a rectangular via to improve the patterning performance of the via. It could also include improving the enclosure of the via by the metal layers above and below it to ensure a quality connection.

In the custom layout flow an automated tool may be used to identify locations where the via connections could be improved or possibly generate layout shapes to enhance the via connections without additional intervention by the layout designer [4][5]. The automation of this process is critical as the large number of connections often makes it difficult or impossible for a layout designer to review and correct each location manually.

A second example of layout manufacturability enhancement is the automated insertion of "dummy" features to improve pattern regularity. The insertion of dummy metal shapes has been common since the advent of chemical mechanical polishing (CMP) for planarization of metal layers. Reducing the range of local pattern density variation improves the uniformity of CMP significantly. In more recent technology nodes, the need to control pattern density in front end of line (FEOL) layers in order to reduce transistor performance variation has led to extensive use of dummy features for these layers as well [6]. An example of pattern density uniformity improvement through the use of automated dummy pattern insertion for a metal layer is shown in Figure 5.

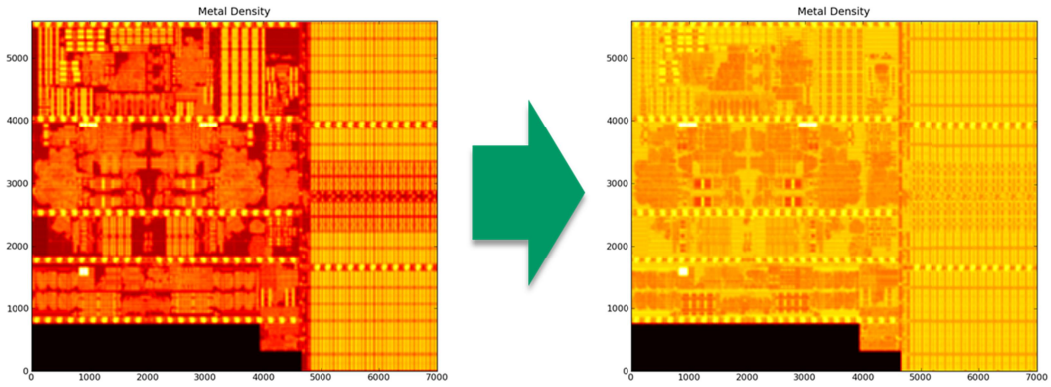

Figure 5. Metal pattern density enhancement by insertion of "dummy" metal features. The left side of the figure shows the metal pattern density variation in a block of physical design IP. The right side of the figure shows the metal pattern density of the same IP block after insertion of dummy metal features with significantly less density variation. The color scale is the same in both plots.

3.3 Manufacturability verification

As the design moves closer to tapeout (when the data is sent to the foundry for mask data preparation and mask writing) there is an increasing need to verify that it meets expected guidelines for manufacturability. However, quantifying the many aspects of manufacturability is not a trivial task. One reason is that, unlike design rule checks, there are often no clear cut criteria for determining when a manufacturable threshold has been met for DFM verification purposes.

This dilemma is illustrated in Figure 6. For an arbitrary design dimensional parameter (e.g., metal line width or space) there exists a relationship between the parameter value and the probability that the design feature will fail due to patterning problems or some other process limitation that causes the feature to be formed incorrectly. Typically a

threshold value of acceptable risk is chosen and assigned to be the design rule value. The risk could potentially be reduced further by adding additional margins to the design. However, this may come at a cost in terms of area, increased power consumption, or performance degradation. The DFM verification methodology should account for these tradeoffs.

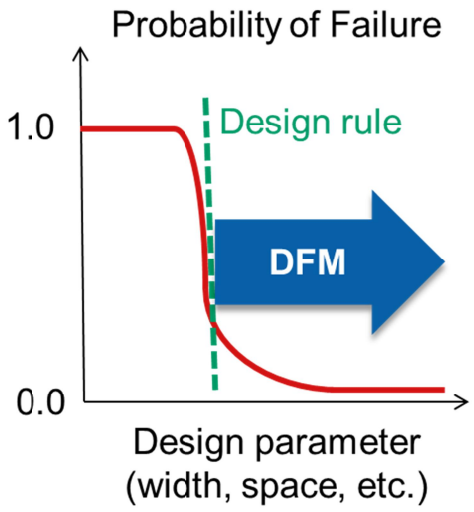

Figure 6. Relationship between an arbitrary design dimensional parameter and the probability that the design feature will fail. A certain threshold is typically assigned to be the design rule value, but improvement beyond the design rule may provide additional benefits for manufacturing robustness. One of the jobs of DFM is to find the right tradeoffs between such design parameters and yield, power/performance, and area.

One approach to DFM verification checks is to generate a list of potential yield issues that come from the design (e.g., single contacts or vias, minimum area metal or active shapes, etc.) and identify locations where each issue occurs in the design to be analyzed. This information can be used to generate a 'score' for each yield sensitivity type, and weighting can potentially be used to prioritize some issues over others (see [7], for example). These scores can be used by layout designers to focus their efforts to improve the manufacturability with a finite amount of time and resources available. Acceptable thresholds for the scores can also be determined in order to ensure that each design is robust enough against known yield detractor issues to meet manufacturability goals.

3.4 Physical design analysis

The role of DFM does not end when the design phase is completed but continues into the high volume manufacturing stage as well. When systematic yield issues are identified either by the fab or through the product debug process the DFM team is often well-positioned at the interface between design and manufacturing to help find the root cause. Once such a systematic yield detractor is identified the first question to be answered is whether it can be addressed with a design change. If so, the cost of such a fix (in terms of design resources, mask costs, etc.) must be weighed against the potential yield gain, power/performance or area penalties, and time-to-market impact.

Often a design modification alone is not optimal when weighed against the related tradeoffs and some manufacturing process improvements may also be necessary. In this case close collaboration with the foundry is needed to ensure that an optimized design/process correction is achieved. Finally, the learning from this analysis must be preserved for future designs and the same yield detractor must be prevented from appearing in future designs. This may take the form of a design rule modification or a forbidden layout pattern, for example. In this way the DFM analysis experience gained at the end of the product life cycle can be fed forward to future product designs via 'preventative' DFM (Figure 4).

4. DFM CHALLENGES AT 20NM AND BELOW

A number of new DFM challenges will be introduced at the 20nm technology node and beyond. The most significant change at the 20nm node is the use of double patterning lithography [8]. This technique requires that mask layers that rely on double patterning be decomposed into two complementary masks that combine to form the final printed pattern. This need for decomposition presents several challenges for designers. One key challenge is the manner in which the layout decomposition will be performed. One option is for the layout designers to create layouts with the patterns assigned to each mask from the outset (the "pre-coloring" approach). A second option is for the design to be created without pre-coloring and then decomposed using an automated tool to assign layout shapes to each mask. This could be done either by the fabless company prior to tapeout or by the foundry as part of the mask data preparation process. Each approach has advantages and disadvantages.

In general, circuit designers would prefer to pre-color their designs. While this places an additional burden on the designers to handle the decomposition themselves, it has the benefit of having a predictable and fixed decomposition solution. This is important in a situation in which IP blocks are reused across multiple products. The automated decomposition process is not guaranteed to arrive at the same coloring solution every time. That means that the same design block could have different coloring solutions when used across different products or even when used multiple times within the same product. This could have performance variation impact as the different instances of the same block (each with a different decomposition solution) may respond differently to variations in lithography dose, focus, or alignment. On the other hand, relying on an exclusively pre-colored layout could lead to pattern density variation between the two masks in the double exposure process that might be more efficiently handled if the full layout was decomposed at the same time. It seems likely that some compromise between the two approaches will need to be found.

Beyond the 20nm node the most significant near-term change is in the front end of line (FEOL), where a shift from planar CMOS devices to FinFETs or other related device architecture is likely. This will require enhancements to the existing design flow to account for the discrete device widths available with FinFET devices. In addition, it seems likely that FinFET devices will impose new requirements for pattern uniformity to control the fin shape and height and reduce device variability.

Finally, the lithography roadmap is unclear for process technology nodes at 10nm and below. Several candidate technologies, including extreme ultraviolet (EUV) lithography, multi-patterning (triple or quadruple exposures), directed self-assembly (DSA), or electron beam direct write (EBDW), have all been explored but so far none has emerged as a clear favorite. This creates uncertainty in the design infrastructure roadmap as well, because it is not clear what additional capabilities will be required. If EUV lithography were to be successful it might require minimal changes to the current design flow (from a design perspective it doesn't look much different from conventional 193nm lithography), while a technique like DSA might require significant modifications to enable design work due to the restrictions on allowed pattern configurations. These challenges will need to be met with close collaboration between the foundries, fabless design companies, and EDA software providers.

5. CONCLUSIONS

Design for manufacturability (DFM) spans a wide range of techniques applied throughout the physical design flow at the custom design, digital place and route, and system-on-chip (SOC) integration stages. In addition, successful DFM requires that it be practiced throughout the product life cycle, including preventative actions before design start, enhancement and verification during the design cycle, and analysis of data from manufacturing. It is a critical enabling technology for modern IC design, and the importance of DFM will only increase for future technology nodes.

ACKNOWLEDGMENTS

The author would like to thank Norma Rodriguez, Yan Wang, and Nawej Mwez of Advanced Micro Devices for their contributions to this work.

REFERENCES

[1] Jie Yang, Norma Rodriguez, Olivier Omedes, et al., "DRCPlus in a router: automatic elimination of lithography hotspots using 2D pattern detection and correction", Proceedings of SPIE Vol. 7641, 76410Q (2010).

[2] Sebstian Nussbaum, "AMD Trinity Fusion APU", *Hot Chips 24*, Cupertino, CA (2012).

[3] Vito Dai, Jie Yang, Norma Rodriguez, et al., "DRC Plus: augmenting standard DRC with pattern matching on 2D geometries", Proceedings of SPIE Vol. 6521, 65210A (2007).

[4] Ahmad Abdulghany, Rami Fathy, Luigi Capodieci, et al., "Smart double-cut via insertion flow with dynamic design-rules compliance for fast new technology adoption", Proceedings of SPIE Vol. 8327, 83270C (2012).

[5] Shobhit Malik, Sriram Madhavan, Piyush Pathak, et al., "Automated yield enhancements implementation on full 28nm chip: challenges and statistics", Proceedings of SPIE Vol. 8327, 83270X (2012).

[6] Norma Rodriguez, Jie Yang, Bill Graupp, et al., "The complexity of fill at 28nm and beyond", Proceedings of SPIE Vol. 8327, 83270Q (2012).

[7] Piyush Pathak, Sriram Madhavan, Shobhit Malik, et al., "Framework for identifying recommended rules and DFM scoring model to improve manufacturability of sub-20nm layout design", Proceedings of SPIE Vol. 8327, 83270U (2012).

[8] Jo Finders, Mircea Dusa, Bert Vleeming, et al., "Double patterning lithography for 32 nm: critical dimensions uniformity and overlay control considerations." *J. Micro/Nanolith. MEMS MOEMS* **8**(1), Jan–Mar 2009, 011002-1.