The Road Ahead

Design for manufacturability: Then and now

Andrew B. Kahng
University of California, San Diego

In the long-ago past, whatever the IC design process produced, manufacturing could deliver. The design rule manual and the Spice model together formed the contract between design and manufacturing. And if something didn't work, the design was at fault.

Fast-forward to the present millennium, when the contract started to unravel. Manufacturing could no longer guarantee all the shapes permitted by the design rule manual: some freedoms had to be curtailed. But manufacturing did not insist on the restrictive design rules (RDRs) that would restore a binary, pass/fail modality. Rather, manufacturers sought an unspecified cooperation from designers, called "design for manufacturability." And it was manufacturing that would decide whether the DFM went far enough to merit a high-yield promise.

In retrospect, this informal arrangement was all quite vague, but neither design nor EDA protested. Compared to RDRs, a few extra checks (litho hotspot, critical area, density, and so on) were more friendly to so-called "essential" designer freedoms. And from virtual fabs in simulation, to C-to-OPC flows in implementation, many adders to traditional EDA flows and markets could be envisioned. In the mid-2000s, the future of the design-manufacturing interface was symmetric, and all the fruit was low-hanging. At times, it seemed that the biggest issue was what to call the fruit—DFM or MAD (manufacturing-aware design)? MFD (manufacturing for design) or DAM (design-aware manufacturing)?

The recent past (say, since 2003 or 2004) taught us a number of lessons.

- *Time scales matter.* The Spice model is constant over quarters; the design is constant over months; the process is constant over weeks. To tweak a process takes weeks, but to tweak a design takes months. It is easier for the fab side to maintain flexibility (via guard bands) with respect to what the design side creates—not the other way around.
- *First-order methods work well.* Regularity, relaxation, and redundancy have been the simplest and most effective DFM approaches.
- *DFM solves "leftover" problems.* They may be fundamental and high-impact problems, yes, but leftovers nonetheless. Only after equipment providers, process engineers, failure analysis engineers, design rule changes, model guard-band changes, smarter circuit and layout designers, and so on have all failed does the product team look to DFM as the savior of yield.
- *Abstractions and proper separations of concerns must be preserved.* For example, chip implementation teams are unlikely to ever embrace "virtual process" in any of the design loops.
- *Low-order bits get dropped.* For example, modeling of chemical-mechanical planarization (CMP) in the back end of the line, to feed "CMP-aware RC extraction and timing analysis," simply wasn't needed.

Today, DFM (MAD) has come to dominate MFD (DAM). There are fewer fabs and toolsets, and the concepts of process debugging and design-friendliness have been supplanted by product development that increasingly relies on multiproduct shuttles. In a regime where manufacturing has great freedom to reject any design, product companies must first run test cases, making product introductions miserably long. It might be argued that DAM is a reality only for high-value customers or IP core

0740-7475/11/$26.00 © 2011 IEEE Copublished by the IEEE CS and the IEEE CASS

and cell providers, with the rest of the world seeing a commodity process. The simplistic explanation: maturation and consolidation in the IC industry, with wars now fought more at the level of SoC platforms and embedded software.

In the near future, two obvious trends will affect the nature of DFM. First, *double-patterning lithography* (DPL) is now the plan of record at or just below the 20-nm logic node, due to the lateness of extreme-ultraviolet lithography tools, and will bring subtle trade-offs of area, variability, and turnaround time. The levers that achieve such trade-offs span design rules, library design, detailed routing, and physical verification. Just two to three years from production, DPL-specific design rules are still an open issue: tighter rules will require smarter "coloring" in routing or verification, while more-moderate rules will incur area penalties. The two most promising patterning approaches—"litho-etch, litho-etch" and "sidewall image transfer"—are very different in terms of pitch and layout density, design flexibility, reliability, and wafer cost.

A second trend is that *continuous shrinks* are becoming the norm. Regularity, relaxation, and redundancy continue to gain in prominence. Because these techniques have such an area impact, they are attractive early in a technology node but are backed out once the technology is mature. From an area standpoint, this will tend to make the shrink look more continuous—and from a design technology standpoint, just as with double-patterning, we will need to learn how to correctly make very complex trade-offs.

Finally, in the long term, there are several looming game-changers at the design-manufacturing interface.

- *Design rule manuals will give way to pattern-matching libraries*. Even today, design rules are no longer rules: they come in flavors that are "required," "strongly recommended," "mildly recommended," and so on, causing rampant designer confusion. As the radius of pattern-dependent stress and litho effects spans ever-more layout pitches, the notion that local "rules" can guarantee yield becomes ever more outdated. Thus, when pattern interactions surpass a certain level of complexity, design rules will be replaced by libraries of yield-impacting patterns that must be avoided by place-and-route tools. (An alternative future could see the rise of "equation-based design rules" as well.)
- *The scope of DFM will broaden tremendously.* The organization charts of leading-edge companies show "design for manufacturability" steadily being renamed "design for X." Indeed, the trajectory within the *International Technology Roadmap for Semiconductors* is to generalize "manufacturability" to "variability," then "reliability," and then "resilience." Design for cost and value will also enter the picture, leading to a future where picoseconds, microwatts, millivolts (of IR drop), megahertz, sigmas (of both process variation and design guard band), mask levels, dollars, and schedule can and must be freely traded off by product teams and their design tools.
- *CMOS scaling will inevitably slow, and More than Moore will rise*. Cost, power, heterogeneity, and productivity all point to future needs in DFM for 3D integration. Beyond the obvious "port today's 2D flows into 3D," this means foundational enablement (simulation models, variability models, reliability models, predictive models, etc.) so that designers can optimally apply *future* beyond-CMOS technologies. It also means new enablement (cost of ownership models, variability scaling and process scaling models, etc.) so that designers can optimally apply *past* technologies within a 3D integrated product. To a great extent, this is uncharted territory today.

PRODUCT SUCCESS INCREASINGLY depends on the right investments, partnerships, and technologies at the design-manufacturing interface. Thus, whatever the road ahead, it seems certain that DFM must continue to solve very real problems that impact the bottom line. Last but not least: a very warm thank-you to friends and colleagues Artur Balasinski (Cypress), Puneet Gupta (UCLA), Sani Nassif (IBM), Chul-Hong Park (Samsung), Puneet Sharma (Freescale), and Rasit Topaloglu (GlobalFoundries) for their inputs to this column. Until next time, on The Road Ahead. ■

■ Direct questions and comments about this column to Andrew B. Kahng, University of California at San Diego, Dept. of Computer Science and Engineering, 9500 Gilman Dr., MC-0404, La Jolla, CA 92093-0404; abk@ucsd.edu.