Reed Electronics Group
THE SEMICONDUCTOR INDUSTRY'S TECHNICAL AUTHORITY™
Vol. 28, No. 6

June 2005

semiconductor INTERNATIONAL®

www.semiconductor.net

DFM: Worlds Collide, Then Cooperate

Inspecting Pattern-Related Defects

Damage to Porous Low-ks

MEMS Manufacturing Update

Look Online
www.semiconductor.net/packaging

DFM: Worlds Collide,

Since the 130 nm node, design complexity has been on a collision course with deep-submicron manufacturing issues, eventually leading to multiple respins and dramatic yield losses that culminated into the latest new field — design-for-manufacturing (DFM). But, of course, DFM is nothing new. Submicron circuits have always been designed *for* manufacturing. What is new today is the mismatch between the precisely constructed and followed design rules and layouts and the chip yield on printed silicon. The main culprit? Subwavelength lithography. But it's not the only one.

Changing defect mechanisms, shrinking process windows, lower supply voltages and power integrity issues each had their hand in necessitating a better relationship between the design and manufacturing worlds. "From a designer's perspective, things are getting more difficult because the process windows they're getting back from manufacturing are so tight that they're having a hard time getting the design methodologies to work. I saw one example where the design tolerances were so big that it didn't make any difference if the company went from 90 nm to 65 nm because they completely lost the performance benefit," said Dave Campbell, general manager of Knights Technology, a division of FEI Co. (Hillsboro, Ore.).

Mark Miller, vice president of marketing and business development for Cadence Design Systems (San Jose), added, "Actual yields were not approximating the process nominal yield from the test chips. So design teams suddenly had to take ownership for a whole new set of issues surrounding attributes like power integrity, leakage current, electromigration issues and the modeling of lower supply voltages they didn't have to worry about before."

The other key reason for DFM methodologies is the dominance of systematic mechanism-limited yield loss. This trend

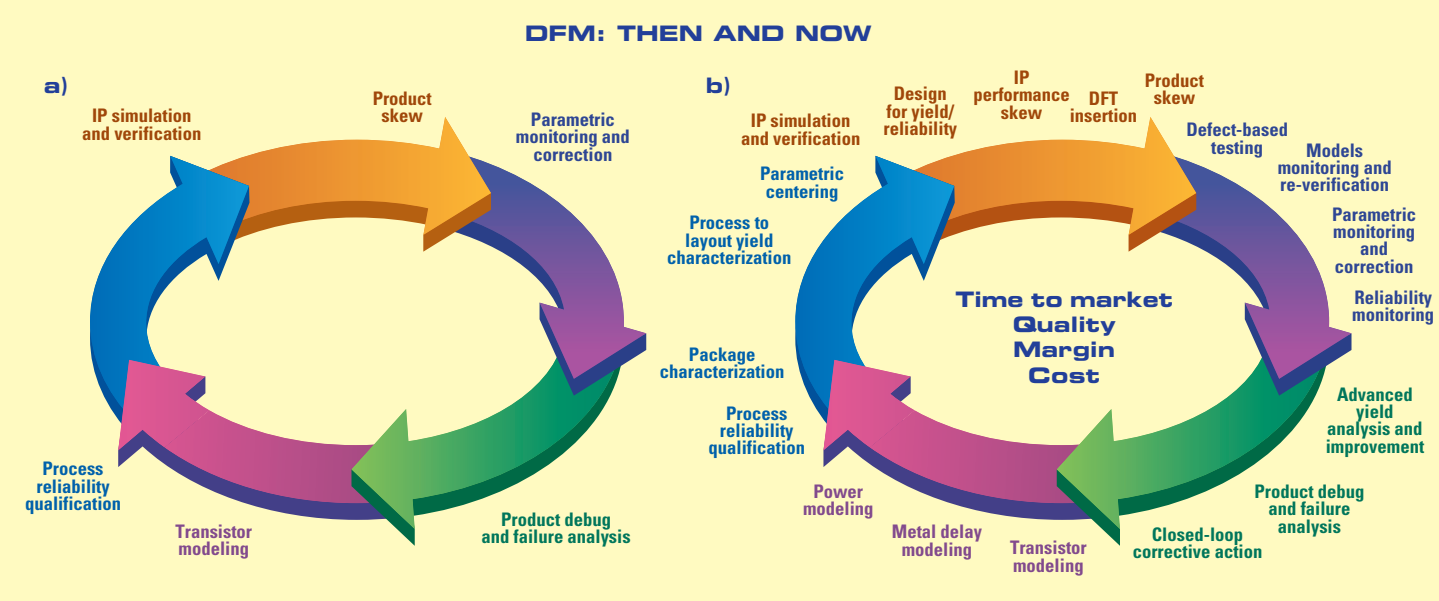

Then Cooperate

Laura Peters
Senior Editor

At a Glance

Some say the hype surrounding design-for-manufacturing has reached cosmic proportions. On earth, IC design has always been for manufacturing. But cooperation between the two worlds is elevated because of serious yield, timing and performance issues.

Artistic representation of optical proximity correction. The application of resolution enhancement techniques is imperative to ensure pattern fidelity at sub-180 nm nodes. (Source: Mentor Graphics)

leads to lower initial yields in manufacturing and lower mature yields, which tend to worsen with each process generation despite ongoing advances in metrology and yield management.[1] "Systematic, pattern-dependent yield loss, which is also called 'feature-limited' yield loss, is directly attributed to design layout, and began to exceed traditional defect-limited yield loss at 0.25 µm," said Harold Lehon, program manager for DesignScan product and the Reticle and Photomask Inspection (RAPID) division at KLA-Tencor (San Jose). "In 90 nm lithography, feature-limited yield loss is 3x the amount of defect-limited yield loss. This is such a significant issue that we must knock it out before we can address the parametric yield issues."

Lehon said what he often hears from customers is, "I need the ability to know my designs will yield before I print glass [reticles and wafers]." Engineers can't continue to troubleshoot design issues in manufacturing. He added, "We hear a lot about moving electrical design intent information downstream into manufacturing, but I think we are pretty far away from effectively implementing that vision. You must address functional yield first and then move on to improving parametric yield."

Meanwhile, design costs have escalated. Non-recurring engineering costs for 90 nm devices are ~$1.5M relative to $4.0M at the 65 nm node, according to Bob Madge, director of technology marketing at LSI Logic Corp. (Milpitas, Calif.). He estimates total design costs including verification to be $30M at 90 nm and the jury is still out on 65 nm. So for fabless companies that bring new designs to foundries, the key issue is not just low yields and respins, which are costing the industry millions of dollars in profits. For smaller companies, respins can threaten company survival.

In discussions of what DFM ultimately will be, it is a revolving supply chain — with multiple feedback loops (Fig. 1). "Traditionally, the DFM supply chain included transistor modeling, process reliability qualification, IP simulation and verification, product skew analysis, parametric monitoring and product debug and FA," Madge explained. "But things have changed. All these building blocks in the design and manufacturing flow — metal delay modeling and power modeling, package characterization and process-to-layout yield characterization, design for yield and reliability, defect-based testing, reliability monitoring and advanced yield analysis — are all needed to over-

1. Design, manufacturing and debug used to be a fairly simple and straightforward process (a). Now, every aspect of design, manufacturing, testing and failure analysis has itself become more complex (b). (Source: LSI Logic Corp.)

come the DFM challenges going forward and to meet time-to-market, cost and quality goals."

Today, the feeback loops are more heavily loaded on the front end (i.e., in lithography) than back-end test. However, these ties are nonetheless relevant to ultimate yield performance in the fab, and eventually DFM will get there.

The issue on the front burner is getting subresolution features to print correctly on the silicon; this process takes a cooperative effort among designers, EDA tool manufacturers, OPC designers, scanner suppliers and mask suppliers. Because the features printed on the wafer are smaller than the wavelength of light used to expose them, "we have to draw patterns that are different from the ones we print," said Thomas Blaesi, vice president of marketing and business development, SIGMA-C GmbH (Munich, Germany). "We apply OPC, go to mask, print on the wafer, and starting at 180 nm and now to 65 nm, we're getting more failures due to lithography problems and yield indications and a lack of good expectations of how those OPC structures we apply to the patterns are actually printing."

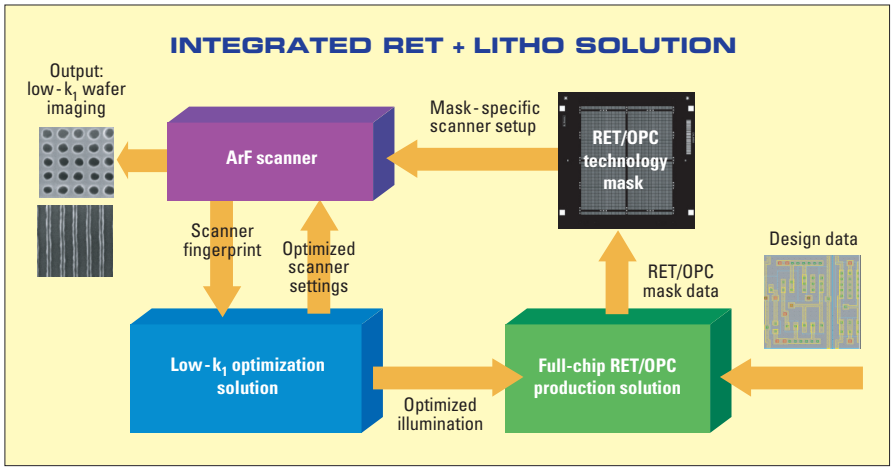

2. Characteristics of the scanner are fed into the scanner simulation and optimization solution, which produces optimized scanner settings and helps optimize the illumination conditions. These are fed into the full-chip mask data prep solution with the design data and all the RETs are applied, then fractures the information into a specific mask data format so maskmakers can make the reticles. (Source: ASML MaskTools)

Process variability

If it weren't for process variability, DFM probably wouldn't exist. However, shrinking process windows, the mismatch associated with OPC and phase-shift masks (PSMs), and the variability caused by new processes and materials has led to an overwhelming need for DFM. Solutions include the integrated mask-to-wafer infrastructure, such as that offered by ASML, ASML MaskTools and Cadence (Fig. 2).

Madge discussed the effect that process variability is having on device performance: "Litho effects are causing gate length variation and drive current variation, such that if we defocus certain parts of the design, it can cause significant parametric variation across the chip or from transistor to transistor and chip to chip. The increasing parametric variability is expected to get to the 40% level this year and increase beyond that. CMP requires tight density control to prevent intrachip density variations, and now there are increasingly stringent requirements for dummy fill.

"We also see variation in power density across the chip and large interchip variation for temperature and heat flux, which can lead to serious problems where temperature control is critical to maintaining reliability and burn-in effectiveness." Interestingly, he said that failure rates, for instance, for via stress voiding, can show unexpected behavior depending on very small changes in design rule. "Fail rate can go from very low to very high over a gradual change in design situation, so a good understanding of failure rate for different design rule situations is critical to understanding the reliability and yieldability of your design."

The first approach dealing with process variability was to do so-called "corners." "We assumed the best case, the worst case and typical in order to bracket the effects due to these process variations. This bracketing approach may be manageable at the 90 nm node, but at 65 nm, the space of the bracketing becomes multidimensional, so we had to find a probabilistic way to determine how the chip will work, taking into consideration all these process effects on resistances and capacitances. This resulted in statistical extraction," explained Rachid Salik, product manager for extraction technology at Cadence Design Systems. "A statistical approach will take process parameter variations, such as metal and ILD thicknesses,

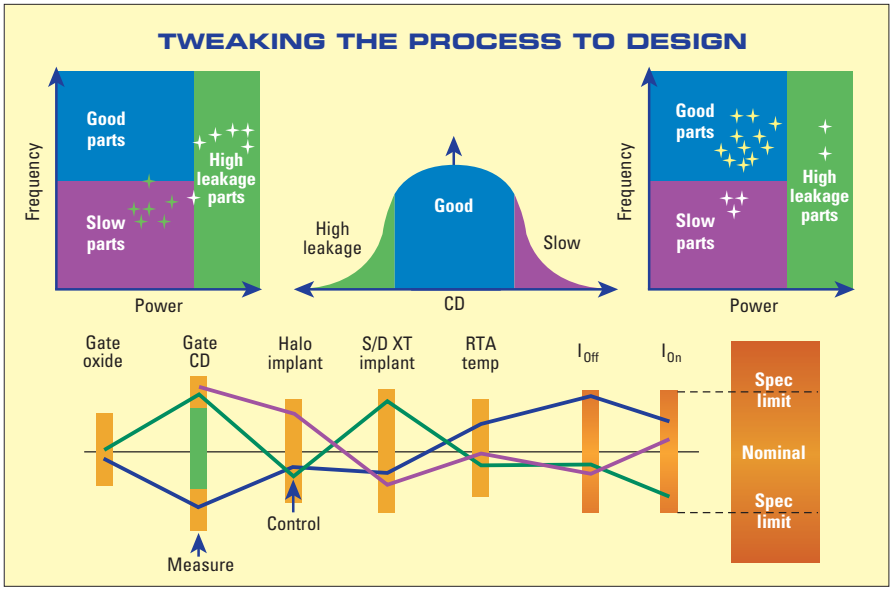

3. With minor variations in critical dimension, knowing the I_{on}/I_{off} targets, other process steps such as the HALO implant can be adjusted to provide the desired electrical results. (Source: Synopsys)

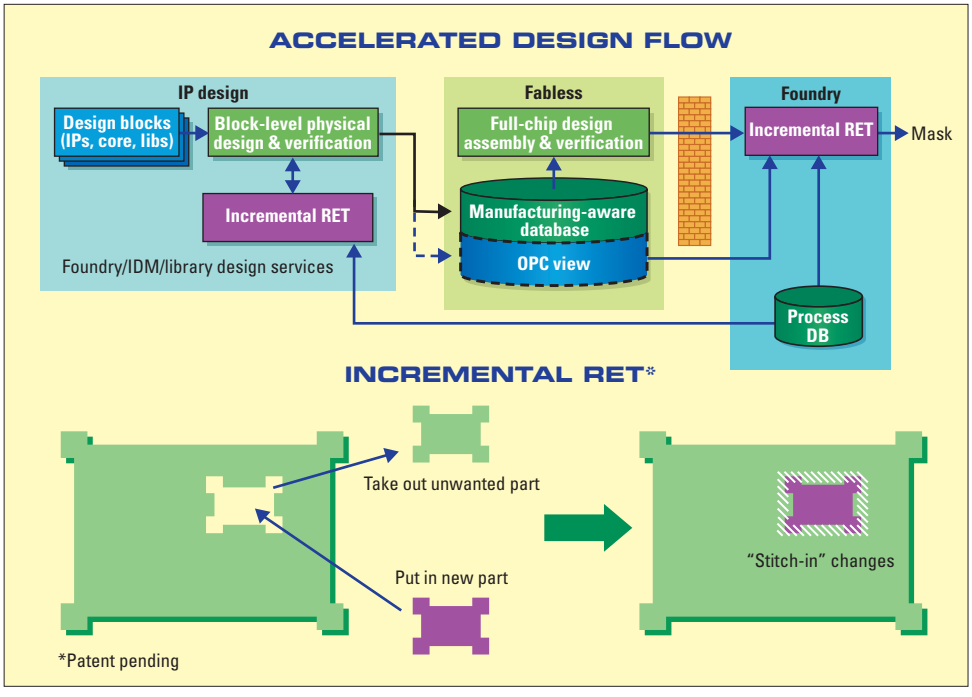

4. Incremental RET works by replacing only a segment of the layout as needed. Local fixes and reuse of existing, good OPC information can drastically reduce OPC turnaround time. (Source: Aprio Technologies)

and translate that, along with other specifications on test chips, to electrical parameters (i.e., resistance and capacitance) average and standard deviation values, for instance. Timing and signal integrity analysis tools will have to process the statistical extracted data to perform this analysis."

A key component in getting a handle on the impact of process variability on a design is calibrating and modeling the manufacturing process and translating that information to the design side. Lehon said, "We're able to inspect the design data with our models to determine how patterns will be formed throughout the lithography process window. Because we have CD-SEM expertise, we can perform the lithography calibration ourselves and directly provide calibrated simulation models to the design side." He added that most EDA tools today require that the design side of the house do the calibration, relying on information from the manufacturing side. "Our observation has been the quality of information coming up from the manufacturing side is highly questionable unless you have deep insight into lithography. This results in a lot of false starts and numerous iterations to calibrate litho models used for simulation from the design side."

Next, Lehon talked about the eventual move to design-aware process control. "People talk about hot spot identification, which primarily comes out of the OPC tools. These hot spot features are typically thousands of points that were identified by rules violations and are fed downstream for monitoring. However, if you have insight into which features are going to fail, not based on rules but lithography modeling, you can limit your sampling and target process control tools to areas that matter most."

Dipu Pramanik, group director of TCAD DFM solutions at Synopsys (Mountain View, Calif.) said, "What's coming down the road at 65 nm and 45 nm is variability, both from a random process or equipment point of view as well as from the layout, affecting the overall performance of the circuit. So we need a true two-way street and simulation to tailor the process slightly from design to design to guarantee yield and also to relate design characteristics and the critical aspects of design to the manufacturers to ensure yield."

"The performance of the device, I_{on} and I_{off}, is the main concern, and if you look at what affects these parameters, it's several process steps, not just the CD," Pramanik described (Fig. 3). "With a consideration of the specification limits of all steps simultaneously, using simulation models, the parametric yield of the devices can be optimized and the process steps can be used in control mode." In this way, the process engineer can increase the number of yielding parts inline.

Wing Leung, CTO of MoSys Inc. (Sunnyvale, Calif.), said, "In the world of embedded high-density memory, design for functional correctness is not the only consideration; memory IP must yield well across all process corners and be reliable in operation across the specification parameters. We use transparent error correction (TEC) with every memory macro to correct corrupted data caused by manufacturing defects and early life failures."

Joe Sawicki, vice president and general manager of Mentor Graphics' Design to Silicon Division (Wilsonville, Ore.), said, "We take three approaches to process variability: remove it, model it or monitor it. For instance, we're allowing designers to actually see the variability in the devices due to lithography effects, so they can minimize it by laying out the cells in a different manner. From the modeling side, people can produce parasitic models that incorporate the process variability inside them, so that that gets fed into the statistical timing in a more rigorous manner. Finally, it's possi-

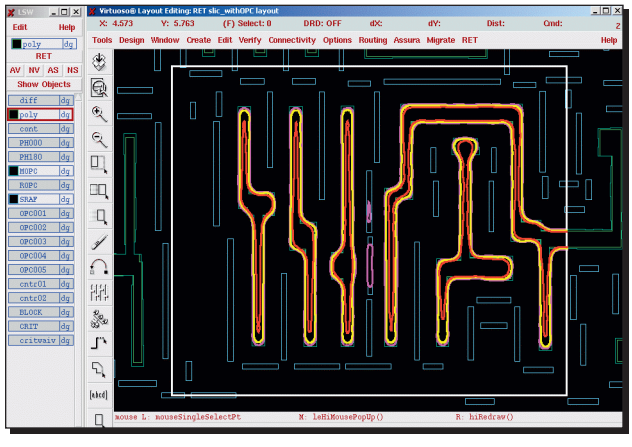

5. After OPC is applied, the green lines show the original layout, the blue lines show the subresolution assist features and the center ovals are the residual of the scattering bars. The designer can visualize this potential bridging problem. (Source: Cadence Design Systems)

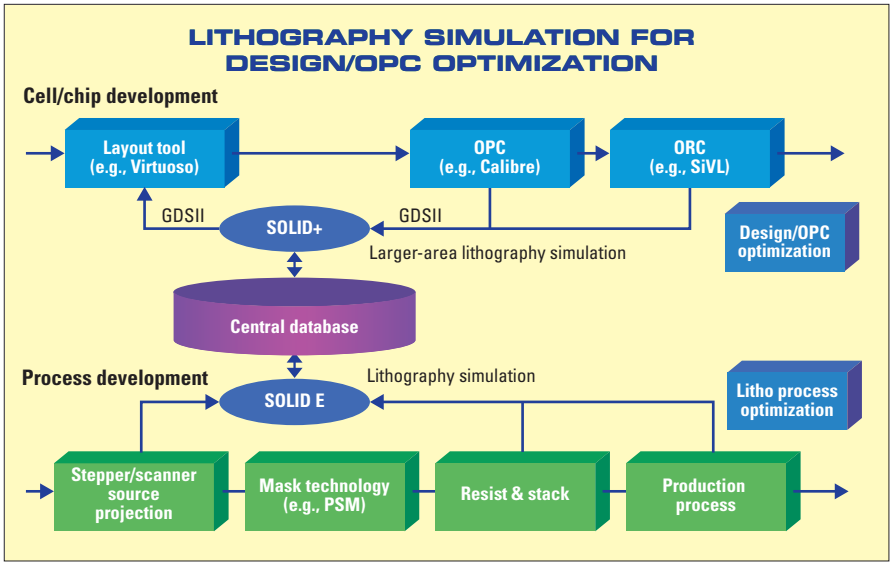

6. The lithography process is modeled for process optimization (bottom) and the data is stored in the database. Then it is made available to design/OPC engineers who perform large-area simulation after OPC. (Source: SIGMA-C)

ble to tag regions that are more susceptible to variability and monitor them on the fab line."

Finally, you have DFM tying failure analysis to design. Knights Technology's Campbell said, "There's a lot of issues because of process variation, so it's important to have methods in place in the back end to look at large amounts of test data and gather information to see where some of the issues are coming from. This way, the fab can see what process area they should be watching out for and then feed that back to the designers as a way to improve the models."

OPC

Optical proximity correction (OPC) has evolved from rule-based to model-based corrections to accommodate the complexity of design. "Now we have to use much more aggressive OPC correction. Whether or not the corrections are accurate or good enough is one set of problems, and how long they take to get there is also becoming an increasing problem," described Randy Smith, vice president of marketing and sales at Aprio Technologies Inc. (Santa Clara, Calif.). The move to largely model-based OPC at 90 nm means that run times, including multiple reticle enhancement technique (RET) correction iterations (simulating, moving edges, simulating again, etc.), can take up to three weeks from the time a design enters the mask data prep area to the begining of 90 nm manufacturing (including RET corrections and verifications of those corrections). "What's really scary is what happens at 65 nm, the total time is six weeks, and with the computational workload, perhaps 150 CPUs will be needed, vs. around 50 CPUs at 90 nm," Smith said.

One way to shorten this time considerably is through incremental RET. "Normally, if anything changes on the OPC layout, the entire mask layer has to be redone. We use a reconfigurable OPC technology to locally fix changes so that a designer can reuse OPC information that already exists," Smith said. The technique substitutes the unwanted area and also heals the HALO region surrounding the replaced area (Fig. 4). Applications of this approach include mask changes (engineering change orders, or ECOs), fab line retargeting, and verify and fix. "With verify and fix, designers will process the OPC layer, run the verification tool and find problems they want to correct. Today, they change the global settings and run the entire layer again, but in the process of fixing problems, new errors can be generated. Instead, you can do a local fix." Fab line retargeting occurs when the same chip is being brought up on a different fab line. Instead of simply copying the reticle set, the OPC can be tailored to the specific fab line, accounting for differences in scanner lenses, among others. Incremental RET software can also be used on existing OPC tools from other manufacturers without the need for substantial additional qualification.

Once certain types of corrections have been repeated several times, they can be considered recipes, and that information can be used to change the original OPC settings. "The problem people run into today is when they run through a design and have problems and start changing the OPC settings, they are changing the recipe file they use on the next design coming through the fab. So the OPC recipes are constantly changing, and you can have a design with all kinds of unique problems. So, if you put a design through mask data prep in March and the same design through in November, the results will never be the same," explained Smith.

Traditionally, a major limitation of OPC tools is that they perform corrections at the best focus and exposure point in the lithography process, when features need to print correctly across a process window. "Instead, the industry needs verification tools that look at CD uniformity, pinching and bridging problems through the process window and perform a special correction to get them back in spec," Sawicki said. Lehon explained, "We

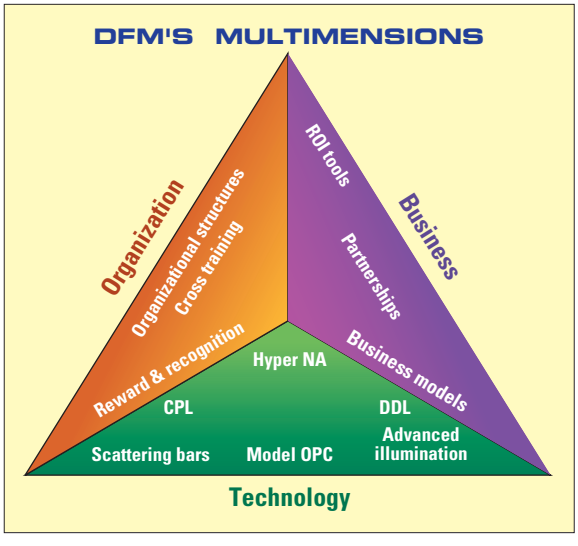

7. DFM is a multidimensional problem that requires the right business approach (partnerships, incentive), organizational structure (cross-training and encouragement/reward for bridging different domains), and, of course, technical solutions. (Source: ASML MaskTools)

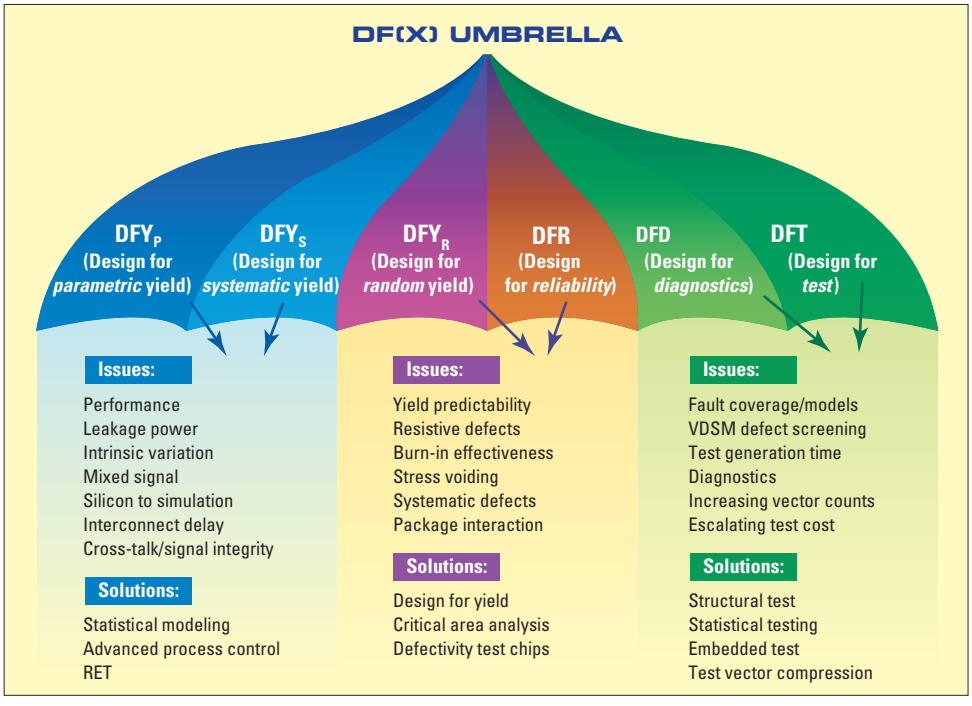

8. DFM includes design for parametric yield, systematic yield and random yield, as well as design for reliability, test and diagnostics. This chart summarizes many of the solutions on hand today. (Source: LSI Logic)

take the output from an OPC tool, run that through our system and simulate the best focus and exposure, then simulate the process window and go through 35 points where we identify features in the design that won't print correctly. This gives customers the option to go back and modify the OPC so the features print correctly through the process window."

Madge pointed out another limitation of OPC models: They typically do not account for reticle noise and dose variations. Miller said, "There's a set of issues that manifest themselves through the OPC process that cause distortions in interconnect structures, so in and of themselves and the way those distortions are used inside a stepper can be one of the causes of process variation chip to chip and wafer to wafer."

Another phenomenon arises from the way in which OPC tools work. "While the RET/OPC tool has complete freedom to break edges and model shapes to different positions, mask constraints will stop you from moving a shape where you want to move it for OPC, so being able to verify where those areas are and what they're going to cost you in terms of CD variability or process window shrinkage is very important to 65 nm customers," Sawicki said.

Figure 5 shows an example of what happens to the original layout (green) when subresolution assist features are applied (blue) and the effect. The designer is left with residual features (pink, center), which should be flagged as potential bridging problems. "What the designer doesn't know is that there's a forbidden pitch, and right in the center of the viewgraph, you have the residual of the scattering bar. That means if we are doing rule-based placement, the designer needs to see they are creating a problem for themselves, creating an area with a potential for bridging and put a warning flag in there. That flag should come when they are laying out the 65 and 45 nm libraries, not at the point of tapeout, which is typically a year and a half after they've laid out the cells," said Wolf Staud, senior technical marketing manager of DFM at Cadence Design Systems.

Kamal Aggarwal, vice president of marketing and strategy for Softjin (Bangalore, India), noted that many DFM EDA startups are developing post-layout manufacturability-related analysis tools, which each requires development of underlying software to handle complex geometric operations. Choosing a data organization that meets their tool requirements may reduce product development costs.

The closer the design can be tied to the actual electrical performance of the device, the better. Cary Vandenberg, president and CEO of HPL Technologies Inc. (San Jose), described a manufacturability simulator that takes the layout for a cell or chip, performs a lithography simulation of the layout and anticipates the effects, for instance, of misalignment, defocus parameters, etch parameters, etc. "This information is extrapolated into a transistor netlist where you can run a transistor-level simulation on the actual silicon netlist. Then there's the ability to apply OPC techniques to the layout and assess different OPC strategies to try and control the device's electrical performance."

A further issue with OPC has been verification of small patterns on the wafer. Blaesi said, "In order to accelerate process development, we have simulation technology that enables designers to verify patterns on the order of 20 x 20 µm; or if you're talking aerial image, image of the pattern without the resist, 100 x 100 µm." Such a product, in conjunction with a silicon-accurate large-area lithography simulation package (Fig. 6), can be used to provide verification of other OPC designs. "It's a verification of the verification," he added.

The DFM communication gap

Because of the "wall" that has traditionally separated design and manufacturing communities, the organizational and business aspects of DFM are perhaps more challenging than the technical ones (Fig. 7). "The organizational piece shows up in the different acronyms used by chip designs, maskmakers and process engineers — RTL, DRC, OPC, NRE, MEEF, CDU, DOF — each set of engineers has their own language. To bring design and manufacturing closer together, we need to think about new organizational structures with cross-training and innovative reward and recognition schemes for engineers," said Dinesh Bettadapur, president and CEO of ASML MaskTools (Santa Clara, Calif.).

The "silos" that currently exist between the design house, mask shop and wafer fab will be bridged by pioneering engineers whose skill sets are broader. "There is awareness, but progress in this area is in the early stages," Bettadapur said. From the business side, he talked about business models with proportional risk and reward sharing between customers and ven-

dors. Importantly, this sharing must take place throughout the development phases of a project, not just the production phases. A successful business approach, according to Bettadapur, also includes integrated financial planning and ROI tools.

Madge summarizes the business challenges of DFM as market window, reducing NRE costs, design tool costs and design resources, creating a roadmap for lower cost at high volumes and avoiding respins. He highlights the most important technical challenges as differentiating product, integrating multiple functions, timing closure, meeting power budget, product reliability and achieving all these within the target lifetime of the product.

"As more companies start working together across these domains, one of the natural things that will come out of that is a new set of standards. But there are areas where it's not so clear whether you need to define a new standard where there was none before. For example, the type of manufacturing information that should be transferred to designs, such that they become more manufacturable, will perhaps remain proprietary to a foundry and its customers," Bettadapur said.

Interestingly, with estimates of around 50 or more startup companies in the DFM space, it could be called the most attractive semiconductor market of 2005. But ultimately, how much of the DFM market will be independent of existing EDA and manufacturing markets? "It is still unclear at this evolving stage whether we are creating a brand new market or taking an existing market and transferring money from one side to another without necessarily growing the pie independently. But what is clear is that DFM does have significant market potential since it can lead to solutions for some particularly challenging design and process-related problems for 65 nm and beyond," Bettadapur said.

DFM solutions

DFM includes design for parametric yield, systematic yield and random yield, as well as design for reliability, test and diagnostics. "Each category is a function of the fab defectivity, design effectiveness and test effectiveness," Madge said. While it was beyond the scope of this article to explore every aspect of the issues under the DFM umbrella, Figure 8 summarizes many of the solutions on hand today. •

References

1. L. Peters, "Demystifying Design-for-Yield," *Semiconductor International,* July 2004, p. 42.

When you contact any of the following manufacturers directly,please let them know you read about them in Semiconductor International.

Aprio Technologies	www.aprio.com
ASM MaskTools	www.asml.com
Cadence Design Systems	www.cadence.com
FEI Company	www.feicompany.com
HPL Technologies	www.hpl.com
KLA-Tencor	www.kla-tencor.com
Knights Technology	www.knights.com
LSI Logic	www.lsil.com
Mentor Graphics	www.mentor.com
MoSys	www.mosys.com
SIGMA-C	www.sigma-c.com
Softjin	www.softjin.com
Synopsys	www.synopsys.com