



BY RON WILSON • EXECUTIVE EDITOR

Is chip design **DIFFERENT** after 90 nm?

DESIGN TEAMS FIND THAT SUCCESS REQUIRES SOME FUNDAMENTAL CHANGES IN THINKING AND IN TEAM STRUCTURE IN THE 90- AND 65-NM PROCESSES.

At every new process node, IC design becomes more difficult. But, as design teams contemplate the move to 90- and then 65-nm-process nodes, many are asking whether the increased difficulty is still just a matter of degree or whether something fundamental is changing. Does a successful 90-nm-design team differ in some way from a successful 130-nm-design team? If so, is the change a one-time thing, or are the differences even greater at 65 nm? The only way to find out is to talk to successful design teams.

Big differences exist in the degree of complexity between the 90- and 65-nm processes, especially for teams performing cell-based design. A top-level diagram of the 90-nm flow that fabless-ASIC vendor Open-Silicon uses shows new branches (**Figure 1**). Jay Jayaprakash, ASIC-design manager at Open-Silicon, groups the new challenges at the 90-nm-process node into power, signal integrity, DFM (design for manufacturability), and design for test. In the power and signal-integrity areas, increased complexity for the design team is arguably a matter of degree. Power management is more critical at the 90-nm node, and it requires the concerted effort of cell designers, tool designers, logic designers, and architects. But this situation does not fundamentally differ from the situation at the 130-nm node. Higher leakage currents at the smaller node exacerbate the problem. Similarly, signal-integrity analysis is more demanding; teams must use the available tools and pay heed to their results.

In addition, Jayaprakash says, HSpice engineers are becoming more central to the digital-design flow at the 90-nm node. Teams must use Spice for their clock nets—something many have been doing for several generations—and this sort of detailed analysis is also becoming important on critical signal paths. In the signal-path area, you see the first glimmer of the need for new skills involving digital-timing analysis. “STA [static-timing analysis] is ballooning out of control on us,” Jayaprakash says. “There are different process corners for different operating modes, and there are different corners for the parasitic-extraction tools that are providing the base data for your STA. By the time you account for all the corners, you are doing 50 or 60 runs, and then what do you do with all the data?”

As STA becomes more complex, some designers are looking at IBM’s SSTA (statistical-STA) tool. Instead of running a single STA at each combination of the process and extraction corners, SSTA attempts to express each delay as a probability distribution and then to compose these distributions as they move through the nets to produce outputs that are

AT A GLANCE

■ Guardbanding at 90 nm results in chips that are too big or too slow.

■ DFM (design-for-manufacturability) tools still cannot accurately predict circuit performance or yield.

■ The choice of which rules and guidelines to obey comes down to a judgment call that requires physical-design people to be process-knowledgeable.

■ The need for process knowledge is spreading upward to architecture.

themselves statistical distributions. It appears to have many more skeptics than fans among design teams. Part of the skepticism comes from the huge amount of process-related data that is necessary to

produce accurate delay distributions in the first place. SSTA may be practical only in an integrated-device manufacturer or a fab-owning ASIC company simply because the tool must work so closely with the process models. But another issue looms even larger in the minds of many designers: What do you do with the output? What does it mean to learn that the timing slack on a net is statistically distributed around -12 psec with a sigma of 24 psec? “We have a tool, but we have to learn what to do with it,” observes Riko Radojcic, principal engineer at Qualcomm. “SSTA is a fuzzy world, not a black-and-white one.”

The question of statistical timing opens a window—albeit a cloudy one—into the most fundamental changes that are taking place in design teams as they adapt to the 90-nm process. Imagine that a timing engineer asked a process engineer what

the delay on a net is. The process engineer might say that the delay depends on the switching speed of the driver transistor. That speed, in turn, depends on the threshold voltage of the transistor, the gate leakage, the effective channel width and length, the drain resistance, and the source and drain contact resistances. Those values vary, depending on the dimensions and shapes of the source and drain implants and the gate polysilicon. They also depend on the shape and composition of the source and drain contacts, which vary considerably in this process. And those things, in turn, depend on the proximity of the transistor’s well implants, the density and pitch of the pattern of polysilicon-gate features, and the proximity of corners or width changes in the polysilicon.

After learning about these and a host of other factors, the timing engineer, who

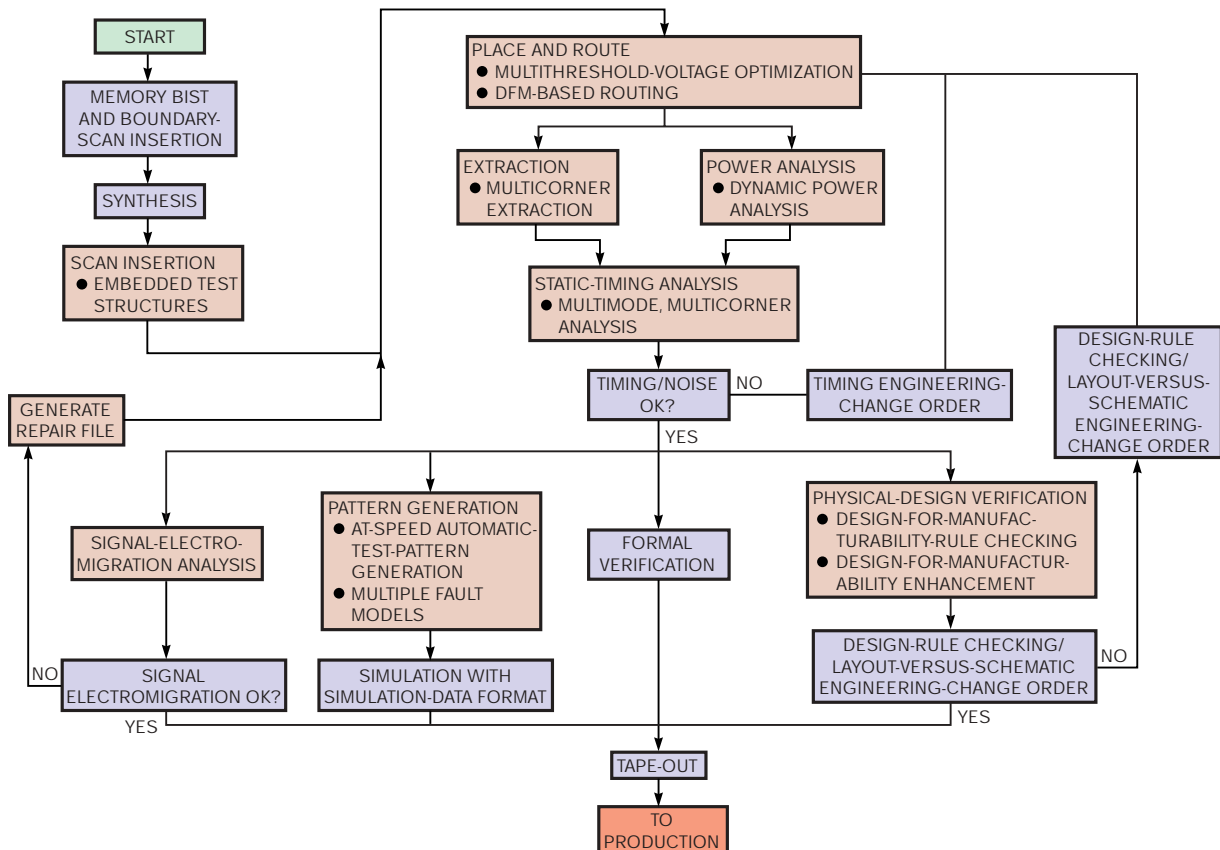


Figure 1 Open-Silicon has added or modified a significant number of stages (shaded) to its design flow to ensure manufacturability at the 90-nm process.

thought his responsibility was to achieve timing closure, suddenly becomes responsible for a range of physical-design and process variables that are beyond his experience.

This scenario is frighteningly close to reality for the early-adopter design teams—the FPGA, memory, and CPU vendors—that are the pioneers on a process node. The responsibility for all these variables does indeed fall on the design team. To succeed, these teams form intimate relationships with their process

providers, whether they are internal fabs or foundries. They also build cell-design, modeling, and process-engineering groups of their own to use the data that the fab-process engineers share with them.

However, this situation is untenable for ordinary cell-based-design teams and unthinkable for ASIC-design teams. That level of process knowledge would be impossibly expensive, even if, under some scenario, a giant foundry would share its basic data with an ordinary fabless customer. The industry is employing a num-

ber of tactics to breach this impasse. The earliest and simplest came into play at the 130-nm process and persisted into the early 90-nm designs. This era of the recommended design rule dictated that, by setting rules, the process engineers at the foundry could theoretically prevent design teams from putting anything into their physical design that would cause substantial variations in critical parameters. Then, by simply guardbanding the specs on the devices, and, at a higher level of abstraction, on the cell libraries the

DFM RAISES GLOBAL ISSUES

The global engineering community thinks that task management works without boundaries. With the possible exception of analog- and custom-design tasks, you can assign any block of a design to any available team anywhere in the world, and the results will be about the same. No interaction exists between culture and digital design.

But in the brave new world of DFM (design for manufacturability), the situation differs, say several experienced design managers. The reason hinges on one of the key differences between design at the 90-nm process and design at larger geometries. That difference is the role of judgment.

Until manufacturers arrived at the 90-nm node, design rules prohibited the features that would impact yield. You didn't put lines too close together for too long a distance, put corners of neighboring features opposite each other, and so on. If you tried to do one of

these things, the DRC (design-rule-checking) tool would flag it, and somebody would fix it.

But at the 90-nm process, many of the rules—those that forbade features that were physically possible but likely to give low yields—became guidelines.

"There are now so many recommendations that, if you followed them all, they would forbid you from doing your design," says Kazu Yamada, NEC's vice president and general manager for custom SOC (system-on-chip) products. On the other hand, if you ignore them all, you may never see double-digit yields from your product. And no magic tool exists that can tell you the impact on yield if you comply with one group of recommendations and skip the others in a block.

So, designing critical blocks at the 90-nm node and even more at the 65-nm node is a matter of judgment. And individuals approach such matters in a culturally conditioned way.

"Japanese people grow up thinking that risk is inherently evil," Yamada says. "Kids in the United States receive rewards for taking risks. And you see that in the behavior of engineers. Japanese design teams are good at absorbing a set of guidelines and identifying the points they need to pay attention to. US teams are more skilled at ignoring a set of risks and finding ways to quantify the results. You now need both cultural attitudes in a design, because if you paid strict attention to all of the yield issues at 65 nm, you would never complete a design."

Qualcomm Principal Engineer Riko Radojcic makes similar observations. "We have design centers in the United States, India, and other countries," he says. "We have some centers that are likely to question the design guidelines and others that may sometimes go too far in implementing them." Radojcic does not ascribe the whole differ-

ence to culture, however. He says that the boldest teams are often those that have an internal group of experts whose job is to convert the foundry's guidelines into a filter that designers apply to the design data. Those design centers that lack access to a resident group of experts tend to be less comfortable with the fuzziness of 65-nm design and to want hard and fast principles.

It may well be that an inherent uncertainty exists at the 65-nm node. You never know whether you have stayed close enough to the guidelines, but you always have to ignore some of them. This uncertainty illuminates individual differences in risk-taking behavior. Some designers accept the risks, whereas others, whether for cultural or personal reasons, assign that responsibility either to a tool or to an expert. This consideration may become significant in the management of global design resources.

foundries began to provide, the process folks could ensure that the chip designers—those who followed the rules, that is—would never get any negative surprises.

However, as they carried this approach into the 90-nm-process node, a negative surprise did occur. The foundry engineers kept finding new problems with each new design and creating new design rules to prevent them. Thus, for early 90-nm-process users—and not just the pioneers—the rules kept changing rapidly. Soon, one list became two lists: mandatory rules and recommended rules. Then, teams added guidelines.

By now, design teams had three problems. First, there were so many rules that even checking for compliance was impractical. Second, complying with even a fair fraction of the rules negated most of the performance and density advantages of going to the 90-nm process in the first place. Third, under the rules, some important structures became simply impossible.

This situation will only get worse at the 65-nm node, according to Kazu Yamada, NEC Electronics America's vice president and general manager. "It's not just that we are checking many more rules," he says. "At 65 nm, adding rules doesn't solve the problem. If you take all the restrictions together, they forbid you to design your chip." So what comes after rules? An intermediate step must emerge between just telling a designer to follow the rules and making him understand the process implications of every node in the design. The simple answer is that models come after rules. In theory, you can identify all the sources of variation that create uncertainty in the timing on a net. You can then model these sources, either statistically or with physics-based models, and you can use the models to tell the designer just what the result of his choices will be. Ideally, the models would give designers data they could actually use: accurate parasitic values, delay figures, and currents.

Even better, you could incorporate the

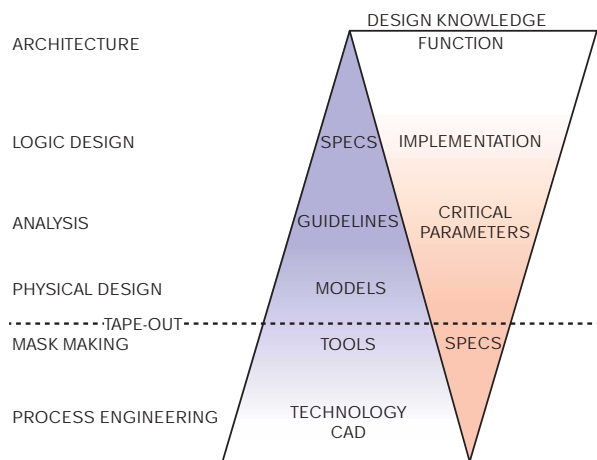


Figure 2 Process understanding is diffusing up into the chip-design team, growing more abstract as it moves up. Design understanding is diffusing down toward process engineering.

models into the placement-and-routing tools and even into the synthesis tools and floorplanners to create a fully process-aware design flow. The tools would handle all the issues that could lead to variability, so the designers could conduct their work just as they did at the 130-nm process, only with a few more tools and steps. Unfortunately, both of these scenarios are fictional. Part of the problem is the difficulty of modeling the process steps that contribute to the variability in the first place. Equipment vendors and process engineers employ PCM (process compact models), which are complex despite the term "compact." These models are too slow for chip-design teams or,

"WE HAVE TO START OUT WITH THE PCMs AND ABSTRACT JUST WHAT THE DESIGNERS NEED."

in most cases, cell designers. Also, they include many inputs that are the sole concerns of the equipment and foundry industry. A model of CMP (chemical-mechanical polishing), for instance, might include inputs for pad composition, pad pressure versus time, back-plate structure, slurry formulation, location and time of the slurry's introduction to the wafer,

rotational speed, temperature, and electric-field strength. And CMP is just one stage in the model.

Accordingly, EDA players are working—hesitantly, perhaps—with foundries and even more tentatively with equipment vendors to extract some data from the PCMs, turn them into simpler and more quickly executed models, and allow them to have inputs for things only under the control of the chip-design team. This effort is causing the growth of independent modeling within the EDA companies. "We have to start out with the PCMs and abstract just what the designers need," says Anantha Sethuraman,

vice president for DFM at Synopsys. This abstraction yields a separate model of the process that must evolve with a rapidly changing situation. "The foundries work with us to update our abstracted models, and we pass the changes on to our customers," he says. "At 90 nm now, the processes around the industry are becoming similar, and they are becoming more stable."

David Thon, group director in the nanometer-analysis and verification group at Cadence, calls this process "adaptive abstraction." He describes a two-way flow of information across the interface between the design team and the foundry (Figure 2). As foundry information flows upstream to earlier parts of the chip-design process, abstract forms become more necessary. Although cell designers may need to know how to make painful trade-offs between speed and variability on a net, architects may need only abstract information about the structures and constraints that lead to a good yield. Conversely, mask makers who run the OPC (optical-proximity-correction) tools may need to have a reasonable understanding of the role a set of polygons plays in the electrical design. Process-integration engineers may need to have only an idea of performance and yield targets.

Currently, neither these models nor PCMs can provide accurate predictions of the electrical characteristics of a net or

accurate quantitative views of the impact that a design choice will have on die yield. "We can give designers some level of comfort, some idea of the impact of the changes they are making, and some idea of what definitely won't work," Sethuraman says.

Another limitation on the EDA industry's approach to the problem is that the industry has so far built abstract models of only those process steps whose variations engineers understand well. In reality, that limitation dictates two models: one for lithography—including, depend-

ing on whom you ask, more or less provision for modeling etch processes—and one for CMP. The industry then lumps anything else that occurs during wafer processing that can contribute to variability in electrical parameters under "variation." "There are 2000 or 3000 sources of variations in an advanced process," NEC's Yamada says. "It's impossible to put models of all of them into the EDA tools, and the EDA companies cannot do their own process models. All they can do is give us faster tools for dealing with the growing number of process corners, but we can't just keep adding more corners, either."

MORE DFM COVERAGE

The July issue of *EDN's* sister publication, *Electronic Business*, explores a post-DFM (design-for-manufacturability) world in which the cost of chip design will be prohibitive. One theory posits that general-purpose chips will become the norm as high design and fab costs make custom chips uneconomical.

Whether or not you buy this scenario, there's little doubt that the relationship between design and manufacturing is changing dramatically. For more, go to www.eb-mag.com/dfm0706.

ing on whom you ask, more or less provision for modeling etch processes—and one for CMP. The industry then lumps anything else that occurs during wafer processing that can contribute to variability in electrical parameters under "variation." "There are 2000 or 3000 sources of variations in an advanced process," NEC's Yamada says. "It's impossible to put models of all of them into the EDA tools, and the EDA companies cannot do their own process models. All they can do is give us faster tools for dealing with the growing number of process corners, but we can't just keep adding more corners, either."

The consensus is that even the latest yield-aware tools still can't protect the

design team from having to understand the details of the design process. There is some hope that at least these issues would affect only cell designers and custom-circuit designers. But manufacturers often can't fulfill that hope. "We tried just guardbanding at 90 nm," Yamada says. "It resulted in chips that were either too big or too slow." NEC then incorporated deep awareness of variability issues into its cell libraries, so layout designers can now choose either conservatively designed, high-yield cells or aggressively designed, high-performance cells that exhibit greater variations. The physical-design team must assess its own risk and its understanding of the original design intent to make the selection.

This process draws a stark line between ASIC and COT (customer-owned-tooling) customers. "The ASIC customers just give us the design and expect us to deal with the DFM issues," Yamada says. "The COT customers, who typically are pushing the area/performance/power envelope, want to be intimately involved in the trade-offs between specs and yield at each critical spot in the design. They have to understand these issues in depth. The best of them have been through the experience before."

NEC's experience is that how far the need for process expertise penetrates into the design team depends on how hard the design pushes the envelope. If the point of using a 90-nm process is simply to make the die smaller and decrease dynamic power, a conventional team can turn over the RTL (register-transfer-level) logic with appropriate timing files and test cases to an ASIC vendor. If the point is to get a chip that is smaller, faster, and less power-hungry, chip designers must intuitively grasp how much risk they are buying in exchange for a more aggressive circuit design.

For example, Qualcomm uses an almost purely cell-based design flow for the 90-nm process, according to Radojic. Many DFM guidelines existed for physical designers, but they handled DFM issues in the same manner as they performed DRC (design-rule checking)—that is, by running DFM-screening tools after layout. The company develops its

own cell libraries, and its foundry runs internal inspection tools on the cell designs and suggests changes. But only the cell designers were privy to this information. "The foundry didn't require our design team to become lithography or CMP experts," Radojic says.

The only big change at the 65-nm process was that the foundries gave their DFM-scoring tools to Qualcomm to run itself. Qualcomm used these tools primarily on individual cell designs, except in the case of full-custom blocks. The design teams did not mandate that a block had to achieve a particular score on the tools. In fact, designers often ran the tools in their spare time after tape-out. "We neither identified a yield issue using the tools, nor found the 'end of the world,' as some people have been predicting. We have had a good yield ramp on our 65-nm design," Radojic says.

To deal with variations, the Qualcomm team stayed with traditional corner-case analysis. "In the brilliance of hindsight, we might have left the margins too big," Radojic says. But that mistake may be an unavoidable cost of a design that depends not on process modeling or analytical tools that don't yet exist but on the skill and experience of master designers. "There is simply no way to relate design choices and DFM-driven modifications to changes in yield," he says. "In a way, the return on investment for all these tools, beyond what you get from the experience of your design team, is faith-based. Now, we are using our 65-nm design flow to evaluate changes to make DFM more formal at 45 nm." He says that there are two schools of thought on the subject: One says to go to SSTA and try to understand how to use it, and the other says to account for as many of the variations as

FOR MORE INFORMATION

Altera
www.altera.com

Cadence
www.cadence.com

IBM
www.ibm.com

NEC Electronics
America
www.necel.com

Open-Silicon
www.open-silicon.com

Qualcomm
www.qualcomm.com

Synopsys
www.synopsys.com

TSMC
www.tsmc.com

MORE AT EDN.COM

Go to www.edn.com/070606df1 and click on Feedback Loop to post a comment on this article.

you can and accurately characterize them in the models.

When designers encounter a problematic structure, they must be able to do something about it. In this case, "something" does not mean tearing up the entire layer and doing a new layout based on handcrafted changes in the one problematic location. That approach would most likely lead only to the emergence of more problems elsewhere. So, EDA vendors such as Cadence have been working on incremental tools that will allow adjustment to a local portion of a layout; rescreening of the new layout for lithography, CMP, and DRC; and incremental modifications to the extraction files. Cadence believes that this approach requires a modification to the underlying structure of the design database, so that tools can identify objects in the neighborhood of a feature without traversing the entire database.

Ultimately, success depends upon the ability of master designers to take in data from a growing Tower of Babel of compliance checkers, rule checkers, and analysis tools and to base design decisions not on yield predictions, but on their own experience. TSMC (Taiwan Semiconductor Manufacturing Co), for its part, is trying to reduce the dissonance. For its 65-nm reference flow, the company has created a common data format for the eight DFM partners it has included in the flow. Using this format, it will put critical-area analysis, lithography-process checking, virtual-CMP modeling, and, in a later release, parametric modeling in the hands of chip designers. "We can't claim to estimate yields," says Edward Wan, TSMC's design-services senior director of

You can reach
Executive Editor
Ron Wilson
at 1-408-345-4427,
ronald.wilson@reedbusiness.com.



product marketing. "But we can predict lithography and CMP hot spots in the design. We will base the sensitivity levels of the tools on our internal variation data, so that the tools can flag things that we know will not provide good yields. This approach will put fixing the problems within the range of back-end chip-design people."

But how much the designers have to know still comes down to how hard they are pushing the envelope. Bradley Howe, vice president of design engineering at early adopter and custom designer Altera, is more careful in his assessment of how well models and tools can shield designers. Howe believes that the company's experience may transfer to design teams that later join with more cell-based flows. "You probably can't just let the cell libraries deal with the problems, at least at 65 nm," he says. "I'm skeptical that you can encapsulate the knowledge of process variations into libraries or EDA tools without leaving big guardbands. That may be OK for some designs, but not for others. Granted we are a full-custom flow. But our experience has been that you can't limit the impact of DFM to some level of abstraction. It permeates all the way up to the architectural level. We found, for instance, that when we were developing the routing architecture for our 65-nm generation, we ended up knee-deep in transistor-dynamics issues. Everybody involved had to have some level of understanding."

This belief appears to be the fundamental change on the horizon for design teams at the 90- and 65-nm-process nodes. If you want to get what a process has to offer, you need individuals who can, based on their experience, their knowledge of the sources of variation, their control, and their impact on electrical characteristics, take calculated risks on a circuit-by-circuit basis (see sidebar "DFM raises global issues").

"My advice to a design team just starting out at 90 nm is: Don't wait. Bite it off now," says Howe. "Either hire the expertise you need or help your people develop it. You have to build this infrastructure inside your team, because, sooner or later, no tool or third party can make the problems go away." EDN