cādence®

Taming the Challenges of 20nm Custom/Analog Design

Custom and analog designers will lay the foundation for 20nm IC design. However, they face many challenges that arise from manufacturing complexity. The solution lies not just in improving individual tools, but in a new design methodology that allows rapid layout prototyping, in-design signoff, and close collaboration between schematic and layout designers.

Contents

Introduction................................1
Double Patterning........................3
Layout-Dependent Effects............5
New Interconnect Layers.............6
New Design Rules 6
Device Complexity and Variation..7
A New Custom Methodology.......7
Conclusion..................................9
References 9

Introduction

For many electronics OEMs, particularly those working with mobile applications, the move to the 20nm process node will be irresistible. Early estimates point to a potential 30-50% performance gain, 30% dynamic power savings, and 50% area reduction compared to the 28nm node. Chip complexity may range up to 8-12 billion transistors. With all its benefits, the 20nm node will open the door to a new generation of smaller, faster, more differentiated devices.

The 20nm process node, however, comes with many challenges, and most of the discussion thus far has concentrated on the challenges faced by digital designers[1]. This white paper focuses on the custom and analog designers who will lay the foundation for 20nm design. Custom designers will produce the standard cells, memories, and I/Os that digital designers will assemble into systems-on-chip (SoCs). Analog designers will create the IP blocks that will be integrated into 20nm SoCs, nearly all of which will be mixed-signal.

All designers face several "big picture" challenges at 20nm. One is simply the investment that's required. The research firm IBS predicts substantial increases in fab costs, process R&D, mask costs, and design costs with the move from 28nm to 20nm (Figure 1). Profitability may require shipments of 60-100 million units. Finances will thus mandate careful risk considerations.

	28nm	20nm
Fab Costs	$3B	$4B - $7B
Process R&D	$1.2B	$2.1B - $3B
Mask Costs	$2M - $3M	$5M - $8M
Design Costs	$50M - $90M	$120M - $500M

Figure 1: Fab, process, mask, and design costs are much higher at 20nm (IBS, May 2011)

Another challenge is design enablement, which mostly represents a worsening of existing concerns. These concerns include time to market, profitability, predictability, low power, complexity, and cost. At 20nm, an increased amount of silicon IP must be obtained from multiple sources, and there will be increasing concerns about mixed-signal design, integration, and verification. Mixed-signal interactions will increase as more and more digital control circuitry is used, and as analog and digital components come into close proximity.

Of most concern to custom/analog designers— and the main topic of this white paper—are the challenges that arise from manufacturing complexity. What is unique about 20nm is the deep and complex interdependency of manufacturing and variability, on top of increasing timing, power, and area challenges. Concerns include the following:

- The use of double patterning (with extra mask layers) so 193nm photolithography equipment can print features at 20nm
- Layout-dependent effects (LDE) in which the layout context—what is placed near to a device—can impact device performance by as much as 30%
- New local interconnect layers
- More than 5,000 design rules, including some new and difficult ones
- Device variation and sensitivity
- A new type of transistor, the FinFET

Concerns about manufacturability issues vary according to design style (Figure 2). For example, analog and I/O designers are most concerned about LDE, circuit specifications, and area vs. performance tradeoffs. Memory and standard cell designers are very concerned about density, and as such, double patterning and local interconnect are key concerns.

Common Design Styles	Multi-Patterned (MPT)	Local Interconnect	Layout-Dependent Effects (LDE)
Analog	SHOULD	SHOULD	MUST
I/O's	SHOULD	SHOULD	MUST
SRAM/Memory	MUST	MUST	SHOULD
Standard Cell	MUST	MUST	SHOULD
Custom Digital	SHOULD	NO	COULD
Chip Assembly	COULD	NO	COULD

Figure 2: Key manufacturability "care abouts" differ according to design style

The solution to 20nm custom/analog challenges lies not only in new point tools, but in a new custom design methodology. In this methodology, circuit (schematic) and layout designers will work in close collaboration, and will have the ability to rapidly exchange information. Circuit designers will be able to obtain early parasitic estimates before the layout is completed. The flow will use in-design, signoff-quality engines as opposed to attempting to fix everything during the final signoff stage. And all tools will be "double patterning-aware" and ready for 20nm.

Double Patterning

The most-discussed manufacturability issue at 20nm is double patterning. This technology splits a layer into two separate masks so that 193nm lithography can print structures that are too close together to resolve with a single mask. This technology is needed to get current 193nm lithography equipment to print correctly when metal pitches are below 80nm, which will be the case for at least some of the metal layers for almost any 20nm design.

When double patterning is used, each mask is exposed separately, and the exposures overlap to create features that are half the pitch that would otherwise be printable with 193nm lithography (Figure 3).

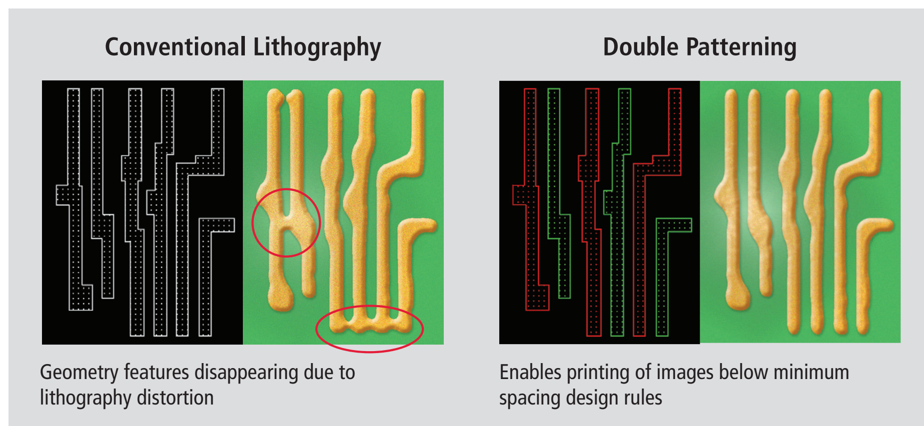

Figure 3: Double patterning makes it possible to print features that could not be printed with 193nm lithography

The concept may be simple, but managing double patterning is difficult. It requires a two-color layout decomposition process in which alternate colors (such as red and green) are used to indicate which features will be placed on which mask. This results in added design rules that restrict the placement and proximity of layout features. For example, traces that are the same color can't be placed too closely together.

As shown in Figure 4, it is very easy to create a design-rule checking (DRC) "loop," which is a coloring conflict that cannot converge on a solution that works. And in many cases, it will be necessary to trace back a number of steps to unravel how the loop was created.

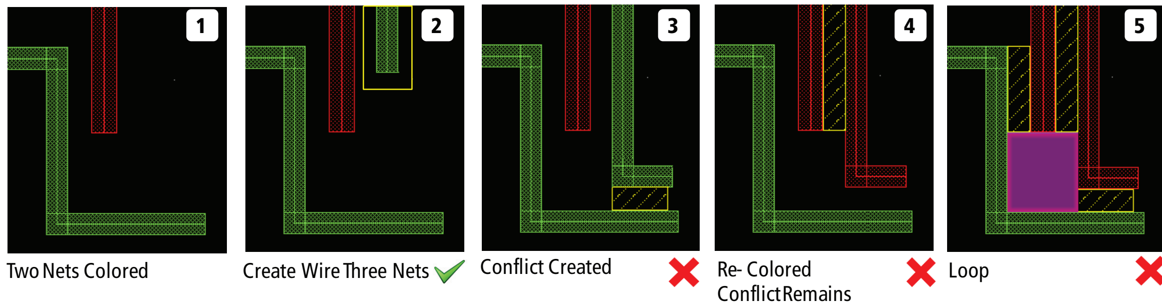

Figure 4: Double-patterning loops are easy to create

Handling double patterning properly is a big concern for custom designers of standard cells, memories, and I/Os. These designers must be cognizant of coloring as they create layouts to optimize area. It is difficult to achieve a high density while making the design decomposable. According to some reports, standard cells that previously took four hours to lay out sometimes take a week at 20nm, because designers have to keep re-running verification as they try to pack decomposable cells as tightly as possible.

Analog designers are concerned about the mismatches that an additional mask can cause. Double patterning impacts electrical performance because different masks on a given layer will shift during the manufacturing process (Figure 5). This mask shift causes variations that have a direct impact on RC and the interconnect. As a result, parasitic matching can become very challenging.

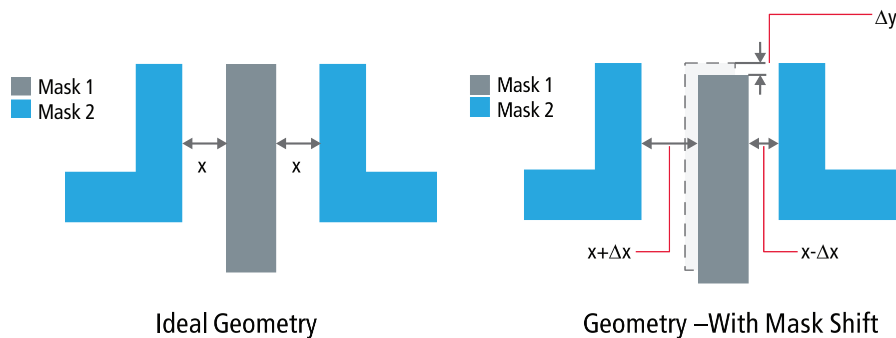

Figure 5: Mask shift occurs with double patterning

EDA tools can help automate the colorized decomposition process and can help ensure correctness. In a 20nm-aware toolset, all physical design tools should be double patterning aware, including placement, routing, extraction, and physical verification. For example, extraction must be able to predict the capacitance variation resulting from mask shift.

Another capability that's needed is automated color-aware layout. Here, color conflicts are avoided as the layout designer draws shapes and places cells. Once a shape is dropped into place, the tool automatically makes any needed coloring changes. Locking color choices down in the design phase sets a constraint for the manufacturer and helps to ensure that matching is correct by placing pairs of nets or devices on the same mask.

Instead of running signoff verification once every four hours, a quick double patterning check should be run after every editing command. In this way errors can be fixed quickly, and designers don't end up with DRC loops that may take many steps to unwind.

Layout-Dependent Effects

At 20nm, it's not enough to model the performance of a transistor or cell in isolation—where a device is placed in a layout, and what is near to it, can change the behavior of the device. This is called layout-dependent effect (LDE), and it has a big impact on performance and power. While LDE was an emerging problem at 28nm, it is significantly worse at 20nm, where cells are much closer together. At 20nm, up to 30% of device performance can be attributed to the layout "context." That is, the neighborhood in which a device is placed. Figure 6 shows how voltage threshold can change according to "well proximity effect," or how close a device is placed to a well.

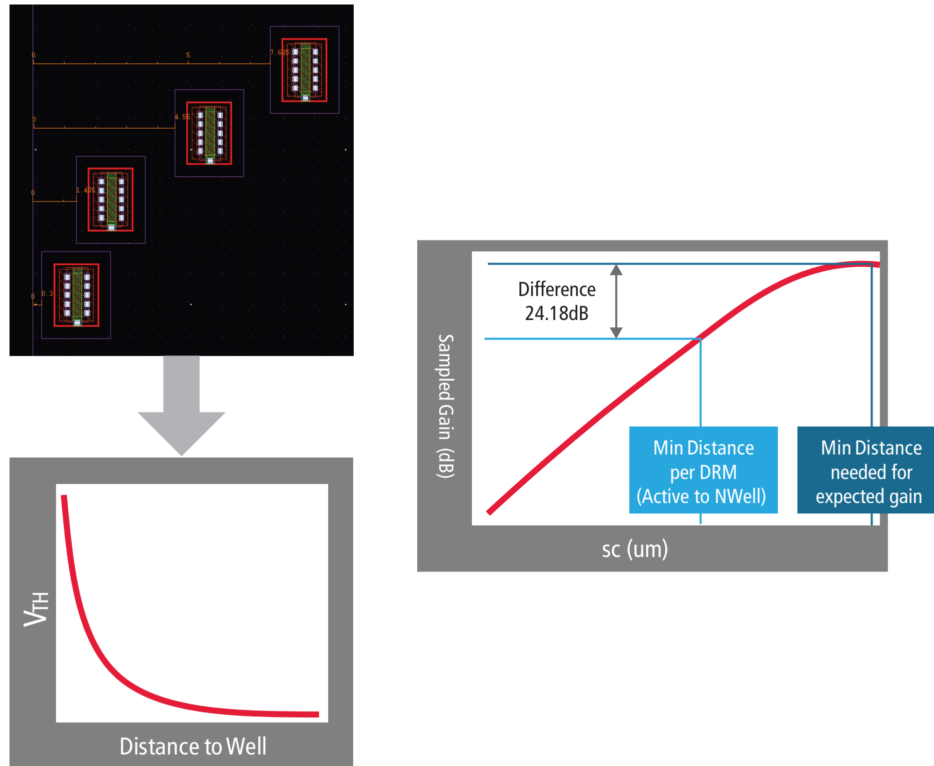

Figure 6: Well proximity effect is a source of LDE, and it impacts voltage thresholds

As shown in Figure 7, there are many potential sources of LDE. While a few of these effects emerged at 40nm and above, most are far more problematic at 20nm. For example, the distance between gates including dummy poly has a direct effect on the drain current of the transistor (poly spacing effect). Length of diffusion (LOD) is the distance from an n-channel or p-channel to the shallow trench isolation (STI) oxide edge. Oxide diffusion (OD)-to-OD spacing is active-to-active spacing.

Layout-Dependent Effects	Prior to 40 nm	At 40nm	28nm and Beyond
Well Proximity Effect (WPE)	X	X	X
Poly Spacing Effect (PSE)		X	X
Length of Diffusion (LOD)	X	X	X
OD to OD Spacing Effect (OSE)		X	X

Figure 7: Many layout-dependent effects first appear at 40nm (black X) and become severe at 28nm and below, with limited workarounds (red X)

A major cause of LDE is mechanical stress, which is often intentionally induced to improve CMOS transistor performance[2]. For example, a dual stress liner is a silicon nitride (SiN) capping layer that is intentionally deposited to be compressive on PMOS and tensile on NMOS—improving the performance of both. Nonetheless, it results in variability that may make it difficult to close timing. Stress is also unintentionally induced through technologies such as shallow trench isolation, which isolates transistors and determines active-to-active spacing.

LDE cannot be modeled in Pcells or device models. It is no longer enough to create a schematic, pick a topology, run a simulation, and throw it over the wall to a layout designer. At 20nm, circuit designers have to consider layout context as well as device topology, and they need to simulate with layout effects prior to layout completion. While this may sound paradoxical, pre-layout sensitivity analysis tools can identify devices that are sensitive to LDE. Circuit designers can also make use of LDE simulation using partial layouts and LDE-aware layout module generators. Layout engineers can use LDE hotspot detection and fixing.

What's needed is context-driven placement and optimization that can determine how different cells are going to interact and how the layout context affects timing and power. The design tool should take care of LDE during both schematic and layout phases. But it's not just about tools. Circuit and layout designers have to learn to work together in new ways, with a much higher level of cooperation.

New Interconnect Layers

For many custom designers, 20nm is all about density. Local interconnect (LI) layers—also called middle-of-line (MOL) layers—offer one way to achieve very dense local routing below the first metal layer (Figure 8). There may be several of these layers, and most don't use contacts; instead, they connect by shape overlap without any need for a cut layer. Getting rid of contacts makes routing denser because contacts are bigger than nets, and can't be placed too close to nets.

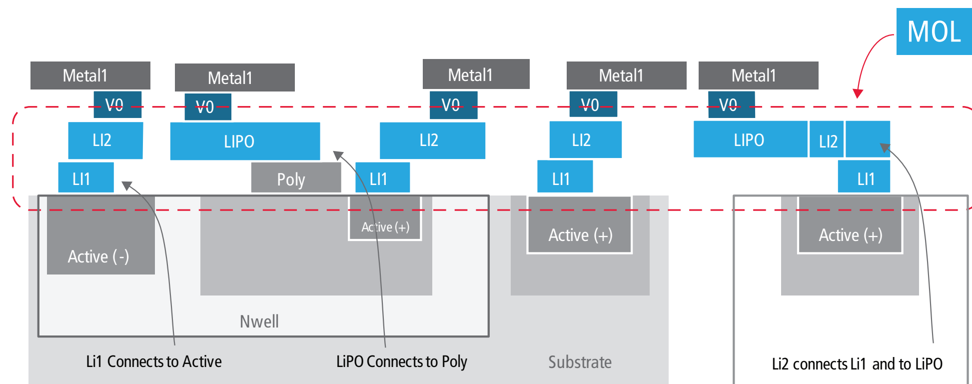

Figure 8: Local interconnect layers provide additional routing density

However, designers need to be aware that LI layers have their own restrictive design rules. For example, LI shapes can only be rectangles, and they often have fixed directions with length constraints.

New Design Rules

At 20nm, there may be more than 5,000 design rule checks. There are more than 1,000 new rules since the 90nm process node, and double patterning alone requires 30-40 checks. The 20nm node adds another 400 new advanced rules such as wrong width and wrong spacing, discrete width and spacing, and special rules for pin access on cells. Designers will face directional orientation rules, specific rules regarding length/width and transistor proximity, and new rules governing legal inter-digitation patterns.

An important point to remember is that 20nm doesn't just bring in more rules; it brings in more complex rules. Simple rules of thumb or memorization won't work any more. What's needed is a design system that can check design rules as the design progresses, rather than waiting for a final signoff check.

Device Complexity and Variation

Transistors at 20nm are very small and very fast, and variation is a constant challenge. Transistors are sensitive to channel length and channel doping, and transistor behavior is subject to short-channel effects. Custom designers must minimize leakage, ensure reliability, and achieve reasonable yields.

The reality is that 20nm transistors were designed for achieving high densities in digital design, not for optimizing leakage or gain in custom/analog design. A limited set of device sizes are available for design. Width and length parameters are limited to a small set of values, and with fewer choices, manually tuning transistors to meet specs (such as gain) is difficult.

Some 20nm designs—and many, if not most, designs at 14nm and below—will use a new type of transistor called a FinFET (or tri-gate in Intel's terminology). In a FinFET, the FET gate wraps around three sides of the transistor's elevated channel, or "fin" (Figure 9). This forms conducting channels on three sides of the vertical fin structure, providing much more control over current than planar transistors. Multiple fins can be used to provide additional current.

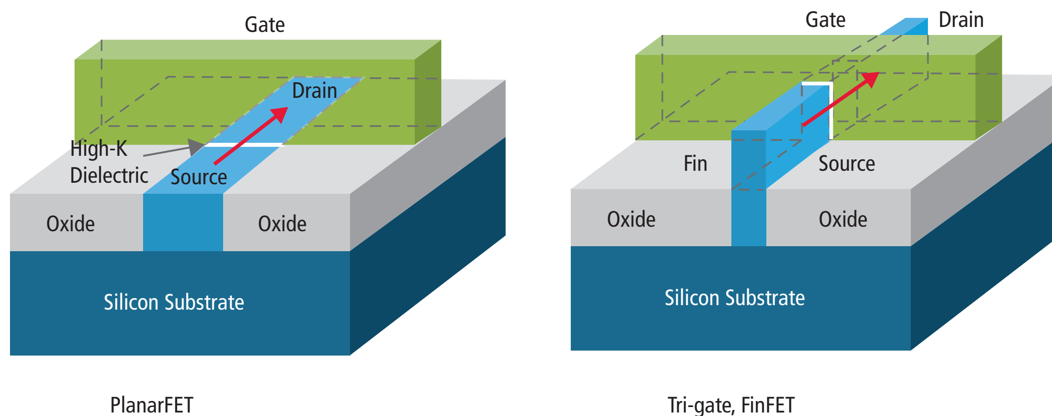

Figure 9: A FinFET (right) wraps a gate around three sides of an elevated channel

FinFETs promise greatly reduced power at a given level of performance. According to Intel[3] (which is using tri-gate transistors in its 22nm "Ivy Bridge" chip), 22nm tri-gate transistors provide a 37% performance increase and use 50% less power at the same performance than 32nm planar transistors, for an added wafer cost of only 2-3%.

However, FinFETs raise some challenges for custom/analog designers. One constraint is that all the fins on the transistors on a given chip must be the same width and height. Designers can add fins to increase the width, but this can only be done in discrete increments—you can add 2 fins or 3 fins, but not 2.75 fins. In contrast, planar transistors can be adjusted to any channel width in order to meet specifications.

Other challenges include additional design rules, manufacturing variations in the width and height of fins, and metal and via resistance. SPICE models with additional parameters will be needed for the FinFETs, and simulators must be able to interpret them. Extraction tools must be aware of the capacitance and resistance that arises from 3D transistor structures. Layout tools will have to be optimized to handle FinFETs. And like any new technology, FinFETS will require ecosystem support including EDA tools, process design kits (PDKs), physical IP, and silicon-proven manufacturing processes.

A New Custom Methodology

To resolve the challenges described in the above sections, every tool in the custom/analog flow needs to be aware of the changes that 20nm brings, including double patterning, LDE, local interconnect, complex design rules, device variation, and FinFETs. But it is not enough to just improve tools. What is needed is a new methodology that provides a higher level of automation than existing flows. In this methodology, circuit and layout designers will exchange and share information, layout prototyping will provide early estimates of parasitics and LDE, and an in-design signoff approach will greatly shorten final signoff runs.

The traditional custom/analog flow is a manual, "throw it over the wall" approach. Circuit designers do schematic entry and run an ideal simulation without layout parasitics. The design is tossed to layout designers who handle device creation, manual placement, and manual routing. Next comes physical verification, extraction, and a final simulation. It's a time-consuming, serial methodology in which issues are exposed very late in the design process, and many design iterations may occur. No wonder blocks that took four hours to lay out at 28nm might take a week at 20nm!

Figure 10 depicts a more automated and collaborative methodology. Here, circuit designers draw schematics, just like they always have. But they also pass constraints to the layout designers, and run a pre-layout parasitic and LDE estimation. On the right side of the diagram, both circuit designers and layout designers can use Modgens (a Cadence® term for automatic module generators) to quickly generate layouts for structures such as differential pairs, current mirrors, and resistor arrays. While not a final "DRC-clean" layout, these automatically generated layouts allow accurate physical effects to be extracted, analyzed, and simulated. Modules can then be fed into an analog placer and assembled into a floorplan.

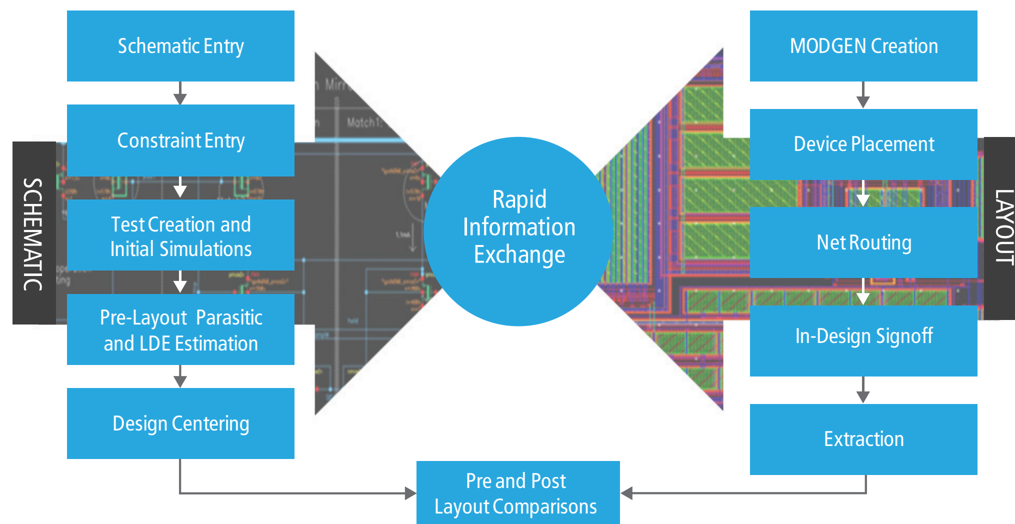

Figure 10: A new custom design methodology allows rapid information exchange between schematic and layout designers

Layout engineers then perform device placement, routing, in-design signoff, and extraction. In short, the basic roles of circuit designers and layout designers remain the same, but there is an ongoing and rapid exchange of information and a high degree of collaboration. Constraints flow easily between the schematic and layout environments.

Automatic module generators help enable "rapid layout prototyping," which is the ability to quickly generate an extracted layout view so circuit designers can run simulations with real parasitic and LDE information. In this way, the electrical and mismatch problems that might be caused by layout effects can be spotted and remedied early in the design cycle. The module generators, however, must be ready for 20nm, with place-and-route engine support for 20nm design rules, complex abutment support, array-based FinFET configurations, coloring for double patterning, and LDE awareness.

Another approach for bringing physical effects into initial simulations is incremental design. Here, designers lay out the pieces they care about with assisted or full automation, and gather as much physical and electrical information as they can. This results in a partial layout extraction. The emphasis is on placement rather than routing. The point is that designers are not taking the time to do a full layout, just what's necessary to generate the desired parasitic information.

Finally, in-design signoff is a must at 20nm. If a designer makes a mistake during layout, the feedback should come immediately, not four hours or a month later. Otherwise many design iterations may be needed for each block, with each iteration taking longer than it took to do the original design. There are two types of interactive editing checks that can help avoid those iterations. One is "in-edit checking," which warns of errors while geometry is being created. Another is "post-edit signoff-quality checking," which exercises a more robust check after each edit is completed.

The key to in-design signoff is having "signoff-quality" engines that run during the design flow. This does not remove the need for a final signoff check, but it greatly reduces the amount of time and potentially the number of licenses that a final check might consume.

Conclusion

There is no doubt that 20nm is coming. The performance, power, and density advantages of the 20nm node will provide a competitive advantage to those who adopt it. While there are a number of design and manufacturing challenges, the good news is that the challenges are manageable—but only if we look beyond individual tools and rethink the ways that custom and analog circuits are designed.

It's certainly true that every tool in the custom/analog flow must be 20nm aware, and able to handle such challenges as double patterning, LDE, and complex design rules. But improving individual tools is not enough. What's needed is a new custom/analog methodology in which schematic and layout designers work in close cooperation, schematic designers can run prototype layouts to gather parasitic and LDE information, and signoff-quality engines permit "in-design signoff" that catches the vast majority of errors before the final signoff phase.

While there's been much discussion of the problems facing digital designers at 20nm, it's the custom and analog designers who will lay the foundation for this process node. Digital design cannot proceed without standard cells, memories, and I/Os, and SoC design cannot move ahead without analog/mixed-signal IP. A complete 20nm solution must therefore include custom/analog and digital design, and allow a close interaction between these domains, preferably using a common database such as OpenAccess. Cadence Design Systems offers such a solution today.

References

[1] A Call to Action: How 20nm Will Change IC Design, white paper, cadence.com

[2] Modeling Stress-Induced Variability Optimizes IC Timing Performance, white paper, cadence.com

[3] Intel Reinvents Transistors Using New 3-D Structure, press release, intel.com

Cadence is transforming the global electronics industry through a vision called EDA360.
With an application-driven approach to design, our software, hardware, IP, and services help customers realize silicon, SoCs, and complete systems efficiently and profitably. **www.cadence.com**

© 2012 Cadence Design Systems, Inc. All rights reserved. Cadence and the Cadence logo are registered trademarks of Cadence Design Systems, Inc. All others are properties of their respective holders. 22970 09/12 MK/DM/PDF