The Road Ahead

Shared red bricks

Andrew B. Kahng
University of California, San Diego

REQUIREMENTS SUCH AS THOSE in the *International Technology Roadmap for Semiconductors* (*ITRS*) drive the electronic design automation (EDA) industry (see http://public.itrs.net) as well as other semiconductor supplier industries, such as lithography, front-end processing, and assembly and packaging. This new department will explore the semiconductor roadmap and its implications for EDA. In particular, future columns will attempt to answer three important questions:

Where should the industry place its R&D bets?

Which design technology innovations will be most critical to the semiconductor industry's overall health? Possibilities include

- embedded software synthesis and verification;
- automatic, reusable, analog and mixed-signal design;
- modeling to enable power management at the architectural and operating-system level;
- design with mixed thresholds and supply voltages;
- new paradigms for global signaling and synchronization; and
- mitigation of increased process variability and nonrecurring costs in the mask and foundry flows.

What should the industry invest in design and test R&D?

The *ITRS* identifies at least eight distinct semiconductor supplier industries. The semiconductor industry and system houses determine the amount of research in each of these industries, through company purchases, investments, and contributions to research consortia. In this light, how are the value and cost of design technology measured in the context of the overall semiconductor industry? For example, how does the value of a new formal-verification methodology stack up against that of a new lens coating that enables 157-nm lithography? What is the value of interoperability versus that of improving a given point tool's capability? What is the cost of failing to improve analog synthesis or die-package codesign?

Why is the roadmap always wrong?

No one can predict the future perfectly. However, industry-wide roadmaps such as the *ITRS* have several built-in imperfections. First, the *ITRS* is the joint effort of competitors. This precludes going beyond a certain level of accuracy. Second, the *ITRS* is intended to drive many distinct semiconductor supplier industries toward the global goal of exponentially decreasing cost per transistor (Moore's Law). This requires careful couching of "messages" to these industries. Requirements messages must have some built-in overshoot (the industry needs capability X in year *Y*, but, just to be on the safe side, says X is required in year $Y - 1$). On the other hand, suppliers may ignore some messages if they are too premature ("all tools must interoperate on platform Z") or if their requirements are too difficult and expensive. Furthermore, even when messages are perfectly correct and timely ("we need a complete signal integrity solution by the 130-nm technology node"), suppliers may deliver capabilities late because of business dynamics (for example, market size). Third, roadmapping imperfections arise because the *ITRS* combines many different sup-

0740-7475/02/$17.00 © 2002 IEEE

plier industry perspectives. This brings us to this first column's topic: *shared red bricks*.

In *ITRS* parlance, a red brick is a "technology requirement for which no known solution exists." Solving any given red brick is expensive and requires large R&D investments. The *ITRS* is now full of red bricks, to the extent that these red bricks seem to form a *red brick wall* in the not-too-distant future. My contention is that many red bricks stem from trying to continue old ways or old trends without seeking synergy with other parts of the semiconductor supply chain. The following metaphor may help to clarify this point.

Think of the *ITRS*—the semiconductor industry's technology foundations—as a car. The supplier industries (packaging, lithography, design, and so on) are the car's parts. The car must continue along the Moore's Law road; for example, four years from now, it must reach speeds of 600 mph. It is absurd to think that super tires alone, or super seats alone, will make the car go 600 mph. However, the seat industry might specify its requirements—*and the concomitant levels of R&D investment*—from the perspective that super seats *alone* must enable the 600 mph car.

It is economically wasteful and technologically impossible for each supplier industry to attempt to continue Moore's Law all by itself. We need a more globally optimized allocation of R&D investments—that is, shared red bricks. (By the way, in this metaphor I think of design technology as both the steering wheel and the tires: Application and market drivers such as microprocessor or RF/mixed-signal design drive the car using the steering wheel, and the power generated by the lithography "engine" is transferred to the real-world road via the tires.)

Consider these four examples of potential shared red bricks:

- Can ATE ever handle at-speed test of high-speed interfaces? Is the paradigm of accelerated lifetime testing (burn-in) even scalable to lower supply voltages? Or are there built-in self-test solutions that are shared between the ATE and design technology industries?
- Must lithography, front-end processes, and interconnect technologies continue to push for 10% tolerances in critical dimensions? This would mean gate length and oxide thickness tolerances in the range of a single atomic monolayer by the end of the roadmap. Or are there design-for-variability solutions that share the red brick of variability between the design industry and these other industries? The first generation of variability-aware analysis tools is available now. However, variability-aware synthesis tools (centering for robustness under variability, or for maximum dollars/wafer) are a long way off. Appropriate (and standardized) characterizations of variability sources in manufacturing equipment and processes also appear to be a long way off. Circuit and layout techniques for high-variability regimes must also be explored.
- Should the industry build new, faster mask writers that can handle 250 Gbytes of data for a single mask layer, after optical proximity correction (OPC) and fracturing? Or, should the industry reduce data volumes and relax inspection tolerances—thus improving mask throughput, yield, and cost—by exploiting design hierarchy and an awareness of which features are functionally critical? Obviously, it is more important to apply OPC to, and verify the mask geometry for, a transistor that is in the critical path, as opposed to the company logo.
- Do researchers really need to push for dielectric permittivities below 2.0, or copper interconnect resistivities below 1.8 μΩ-cm, as specified in the *ITRS*? Is the latter even possible? Would developing better circuit and interconnect architectures and layout techniques be a more cost-effective way to share the performance and noise management red bricks between design and interconnect technologies?

THESE EXAMPLES highlight the need for deeper partnerships between design and other *ITRS* technology industries. Such partnerships can potentially resolve key red bricks at greatly reduced cost to the semiconductor industry. ■

■ Direct questions and comments about this department to Andrew B. Kahng, University of California at San Diego, Computer Science and Engineering Dept., 9500 Gilman Dr., MC 0114, La Jolla, CA 92093-0114; abk@ucsd.edu.