

Design-Process Integration and Shared Red Bricks

Andrew B. Kahng
abk@ucsd.edu

UCSD CSE and ECE Departments, La Jolla, CA 92093-0114

ABSTRACT

This paper discusses the relationships between the *design technology* industry and other (process, manufacturing) semiconductor supplier industries, against the backdrop of the semiconductor industry's overall technology requirements (the 2001 *International Technology Roadmap for Semiconductors*). The central claim is that "design-process integration" is by no means blocked by lack of awareness of important problems; the key problems (e.g., manufacturing handoff, reticle enhancement methods, design for variability, etc.) are well known. Rather, design technology currently faces too many *more pressing* demands relative to its limited resources. This situation can improve only when the semiconductor industry and its supplier industries together adopt the perspective of "sharing red bricks", i.e., identifying the most cost-effective and synergetic solutions to technology requirements from principled analyses of costs and returns on investment.

Keywords: ITRS, technology roadmap, design-process integration, VLSI design, VLSI manufacturing

1. INTRODUCTION: THE SEMICONDUCTOR INDUSTRY ROADMAP

The 2001 International Technology Roadmap for Semiconductors (ITRS)⁴ is intended to provide a carefully balanced set of messages from the semiconductor industry to its suppliers (lithography, packaging, test, design, etc.). The messages are in the form of *technology requirements*, color-coded as white (solutions in place), yellow (solutions known and in development) or red (no known solution), with implicit promise of R&D investment and market growth to support the satisfaction of these requirements. The ITRS technology requirements, it seems, are never perfectly accurate - at least, they change rapidly with each successive edition.* Several dynamics are at work. First, the ITRS is a joint effort of competitors, and hence cannot go 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. This requires careful couching of the messages to these industries.[†] Third, roadmapping imperfections arise because the ITRS combines many different supplier industry perspectives.

Consider the following analogy. Think of the International Technology Roadmap for Semiconductors (ITRS)⁴ as a *car* that represents the entire semiconductor industry and its supplier industries (Design, Interconnect, Packaging, Test, Lithography, etc.). The supplier industries are the *parts* of the car, and the car has been driven along the "Moore's Law Road" by two drivers, MPU and DRAM.[‡] However, there have always been passengers in the car (ASIC, SOC, analog, mobile, low-power, networking, wireless, etc.) who wanted to go to different places. In the 2001 ITRS, some of these passengers have become drivers, and all drivers *must*

*It is interesting to observe the substantial accelerations and decelerations that occur between successive ITRS editions. For example, half-pitch, overlay accuracy and CD control requirements have accelerated between 1999 and 2001, while equivalent oxide thickness, metal cladding thickness, and interlayer dielectric permittivity requirements have decelerated. One unavoidable conclusion, though, is that the "red brick wall" with respect to overlay accuracy, CD control, oxide thickness, and ILD permittivity is now only four to six years away, whereas in 1999 this same wall appeared to be six to nine years away.

[†]Messages must have some built-in overshoot (the industry needs capability X in year Y, but just to be on the safe side, says that X is required in year Y-1). On the other hand, suppliers may ignore messages that are premature ("all tools must interoperate on platform Z") or contain requirements that are too difficult and expensive to satisfy. Furthermore, even when messages are perfectly correct and timely ("we need a complete signal integrity solution by the 130nm technology node"), suppliers may deliver capabilities late because of business dynamics (for example, growth of market size).

[‡]In some sense, the car has also been driven in a straight line, with the drivers navigating by the rear-view mirror.

explain more clearly who they are, and where they are going. Specifically, the 2001 ITRS has introduced a new System Drivers Chapter, which attempts to explicitly define the key IC products that drive semiconductor manufacturing and design technologies.

In this invited paper, I will move from the common backdrop of the ITRS - the System Drivers and Overall Roadmap Technology Characteristics - and develop several aspects of the interaction between design technology (DT) and other (manufacturing, process) ITRS technologies. The central claim is that “design-process integration” is by no means blocked by lack of awareness of important problems; the key problems (e.g., manufacturing handoff, reticle enhancement methods, design for variability, etc.) are well known. Rather, design technology currently faces too many *more pressing* demands relative to its limited resources. This situation can improve only when the semiconductor industry and its supplier industries together adopt the perspective of “sharing red bricks”, i.e., identifying the most cost-effective and synergetic solutions to technology requirements from principled analyses of costs and returns on investment.

2. THE ROADMAP OF DESIGN TECHNOLOGY CHALLENGES

Design technology faces challenges from many directions. The basic Overall Roadmap Technology Characteristics context - power, density, speed of devices, interconnects and systems - dictates *silicon complexity* challenges. The System Drivers context - performance and cost measures for MPUs, SOCs, mixed-signal blocks, and other consumers of silicon - dictates *system complexity* challenges. Orthogonally, design technology and the semiconductor industry as a whole face several *cross-cutting challenges*.

2.1. Silicon and System Complexity Challenges

Silicon complexity refers to impacts of process scaling and the introduction of new materials or device/interconnect architectures. Previously ignorable phenomena (implied challenges) now have greater impact on design correctness and value, including: Silicon complexity places long-standing paradigms at risk, e.g., (1) system-wide synchronization becomes infeasible due to power limits and the cost of robustness under manufacturing variability; (2) the CMOS transistor becomes subject to ever-larger statistical variabilities in its behavior; and (3) fabrication of chips with 100 percent working transistors and interconnects may become prohibitively expensive.

- nonideal scaling of device parasitics and supply/threshold voltages - leakage, power management, circuit/device innovation, current delivery;
- coupled high-frequency devices and interconnects-noise/interference, signal integrity analysis and management;
- manufacturing equipment limits - statistical process modeling, library characterization;
- scaling of global interconnect performance relative to device performance - communication, synchronization;
- decreased reliability - gate insulator tunneling and breakdown integrity, joule heating and electromigration, single-event upset, general fault tolerance;
- complexity of manufacturing handoff-reticle enhancement and the mask writing/inspection flow, NRE cost; and
- process variability - library characterization, analog and digital circuit performance, error-tolerant design, layout reuse, reliable and predictable implementation platforms.

System complexity refers to exponentially increasing transistor counts enabled by smaller feature sizes and spurred by consumer demand for increased functionality, lower cost, and shorter time to market. Implied challenges include:

- reuse - support for hierarchical design, heterogeneous SOC integration, and the modeling, simulation, verification and component block test of AMS circuits;
- verification and test - specification capture, design for verifiability, verification reuse for heterogeneous SOCs, system-level and software verification, AMS and novel device verification, test access, self-test, intelligent noise/delay fault testing, tester timing limits, test reuse;
- cost-driven design optimization - manufacturing cost modeling and analysis, quality metrics, cooptimization at die-package-system levels, optimization with respect to multiple system objectives such as fault tolerance and testability;
- embedded software design - predictable platform-based system design methodologies, codesign with hardware and for networked system environments, software verification/analysis;
- reliable implementation platforms - predictable chip implementation onto multiple circuit fabrics, higher-level handoff to implementation; and
- design process management - design team size and geographic distribution, data management, collaborative design support, “design through system” supply chain management, metrics, and continuous process improvement.

2.2. Crosscutting Challenges

The ITRS Design Chapter sets out many detailed challenges (nearly 80% of the overall text in the chapter) with respect to five traditional areas of DT: design process; system-level design; logic, circuit, and physical design; design verification; and test. However, beyond enumerating these detailed challenges, the design roadmap also identifies five crosscutting challenges that permeate relationships between electronic design automation and the other industries that support the semiconductor industry. These five crosscutting challenges are productivity, power, manufacturing integration, interference, and error tolerance.

- **Productivity.** To avoid exponentially increasing design costs, overall productivity of designed functions on chip - as well as reuse productivity (including migration) of design, verification, and test - must scale at more than two times per node. Verification has become a bottleneck that has reached crisis proportions, calling for reliable and predictable silicon implementation fabrics that support higher-level system design handoffs and, particularly in the SoC arena, automated methods for AMS synthesis, verification, and test. Reducing DT time to market requires standards that promote stability, predictability, and interoperability.
- **Power.** Nonideal scaling of planar CMOS devices, together with the Roadmap for interconnect materials and package technologies, presents a variety of power management and current delivery challenges. MPU power dissipation will exceed high-performance single-chip package power limits by 25 times at the end of the Roadmap, whereas LP-SoC PDA drivers require flat average and standby power even as logic content and throughput continue to grow exponentially. DT must address the resulting power management gap in which increasing power densities worsen thermal impact on reliability and performance and decreasing supply voltages worsen switching currents and noise. These trends stress on-chip interconnect resources, test equipment power delivery and dynamic response limits, and even current latent defect acceleration paradigms.
- **Manufacturing integration.** Feasibility of future technology nodes will depend on sharing challenges within the industry as a whole. Die-package-board cooptimization and analysis may improve system implementation cost, performance verification, and overall design turnaround time (TAT) as well as system-in-package DT. New DT for correctness under manufacturing variability - for example, variability-aware circuit design, design for regularity, timing-structure optimization, and static-performance verification - may relax critical-dimension control requirements in the lithography, process integration, devices, and structures, front-end processing, and interconnect technology areas. Finally, more intelligent interfaces that mask production and inspection flows may reduce manufacturing NRE costs.

- **Interference.** Noise and interference increasingly hamper resource-efficient communication and synchronization, which global interconnect scaling trends already challenge. Prevailing signal integrity methodologies in logical, circuit, and physical design - while apparently scalable through the 100 nm node - are reaching their limits of practicality. These methodologies include repeater insertion rules for long interconnects, slew-rate control rules, and power/ground distribution design for inductance management. Scaling and SoC integration of mixed-signal and RF components will require more flexible and powerful methodologies. Issues include noise headroom (especially in low-power devices and dynamic circuits); large numbers of capacitively and inductively coupled interconnects; supply voltage IR drop and ground bounce; thermal impact on device off-currents and interconnect resistivities; and substrate coupling. A basic DT challenge is to improve characterization, modeling, and analysis and estimation of noise and interference at all levels of design.
- **Error tolerance.** Error tolerance, correction, and self-repair could dramatically increase manufacturing yields but will require additional effort in verification and test. Technology scaling likely forces such a paradigm shift, which leads to more transient and permanent failures of signals, logic values, devices, and interconnects. Below 100 nm, single-event upsets (soft errors) severely impact both memory and logic field-level product reliability. Atomic-scale effects demand new “soft” defect criteria, such as for noncatastrophic gate oxide breakdown. In general, automatic insertion of robustness into the design will become a priority as systems become too large to functionally test at manufacturing exit. Potential measures include automatic introduction of redundant logic and on-chip reconfigurability for fault tolerance, development of adaptive and self-correcting or self-repairing circuits, and software-based fault tolerance.

3. LINKAGES AND RESOURCES

It is instructive to juxtapose the above design technology challenges against (1) specific calls to action in the realm of design-process integration, and (2) the design technology community’s resources.

3.1. Linkages: ITRS Requirements for Design-Process Integration

The following excerpts from the “logical, physical and circuit design” portion of the 2001 ITRS Design Chapter (taken from perhaps 7% of the text in the Design Chapter) indicate that the design technology community has more than passing familiarity with the mutual impacts of design and process.

3.1.1. Issues at the Design-Manufacturing Interface

The 2001 ITRS notes that as atomic-scale effects increasingly govern the statistics of many process steps, and 3-sigma variations of such fundamental parameters as L_{gate} , t_{ox} and interconnect dimensions exceed 15%, new circuit topologies and logic/layout optimizations are needed to cope with this level of variability.

Statistical timing analysis and performance verification must comprehend parasitics, delays and geometries that are parameterized by distributions. Design centering must optimize for parametric yield and revenue per wafer, rather than traditional performance metrics. (For example, consider the difference between “L” and “staircase” layouts of a global line with a single repeater, when variabilities are spatially correlated across the die.³) Manufacturing variability (and the proliferation of new materials and processes) also requires a more extensive design-manufacturing interface that supplies design rules and process abstractions to layout. Richer statistical and electrical/geometric characterization of manufacturing variability sources is needed. As power densities continue to rise, naive guardbanding against thermally induced variability will be costly. More accurate analyses and bounds for local thermal variation are needed to reduce the amount of overdesign.

Reticle enhancement technology (RET) encompasses planarization of multilayer interconnect processes (necessitating layout density control with area fill) and deep-subwavelength optical lithography (necessitating optical proximity corrections (OPC) and layout of alternating-aperture phase-shifting masks (PSM)). RET places a growing burden on physical design with respect to layout design complexity, manufacturing handoff complexity, and manufacturing (mask) NRE cost. With OPC and PSM, layout synthesis productivity is challenged by complex, context-dependent design rules. Layout verification must also handle regimes where “local design

rules” no longer exist. Physical verification must accurately understand and model, e.g., the RLC extraction impact of downstream dummy metal insertion in the post-tapeout layout database.

Indiscriminate application of RET explodes data volumes and mask write/inspection costs. RET insertion (and mask inspection) must therefore understand that only certain critical device or interconnect dimensions are worth the expense of careful enforcement, and that some enforcement mechanisms are costlier to implement and verify than others. A data volume- and cost-sensitive PD flow will enable such selectivity by passing detailed functional intent, performance analysis results, and sensitivities forward throughout the layout, verification and mask flows. These requirements have been disseminated throughout the design technology community for several years; e.g.,⁷

3.1.2. Issues Stemming From Non-Ideal Scaling

Lower supply voltages, along with larger currents stemming from increased power densities, lead to larger relative supply rail inductive noise. This is exacerbated by less aggressive scaling of bump counts and pitches. Issues such as IR drop and decoupling capacitance have been addressed in the recent literature. Near-term open issues include control of temperature variation across the die for package and performance reliability. This entails new tools spanning algorithm development, logic synthesis, and timing/layout optimization that must cooperate to manage both instantaneous and average power. The large estimated “power management gap” for microprocessor and SOC system drivers implies a continual focus on power reduction. At the physical design level, library characterization, synthesis, and layout (including power distribution design) must together deliver the roughly 5X available power reduction from fine-grain use of multiple thresholds and supplies (and oxide thicknesses, and biasing) in the same core. Potentially, design tools must automatically produce structures that enable active thermal management via OS-mediated dynamic frequency and supply scaling.

Reliability and fault-tolerance. Reliability criteria (hot-carrier effect, electromigration, joule self-heating, etc.) have been integrated into implementation flows via simple and transparent abstractions (e.g., upper bounds on gate load capacitance vs. output slew time). Such “methodological” abstractions currently permit correctness by construction with little disruption of traditional flows. However, improved abstractions and analyses that reduce guardbanding will be needed in the future. With respect to single-event upsets (SEU) caused by ionizing radiation, decreasing feature sizes lower Q_{crit} values to such levels that even the noise pulse from an alpha particle can be trapped as a logic fault. Automated methods are needed to modify logical, circuit and physical design (e.g., by automatic introduction of error correction, sizing, etc.) to prevent or manage SEU without violating design constraints.

Non-ideal scaling impacts on circuit implementation (notably from scaling supplies faster than thresholds) include higher gate and drain leakage currents, body effect (making pass gate logic less attractive), and loss of overdrive. In light of power management challenges, past tradeoffs of higher power and noise susceptibility (along with unavailability of automated tools) in return for speed become less attractive. As alternatives to static CMOS are deployed to permit overall speed/power performance gains, layout automation and physical verification (e.g., automated extraction of novel active and passive structures from layout) will be needed. For example, self-sufficient circuits such as clock-delayed domino or delayed-reset domino will become more popular with use of globally synchronous, locally asynchronous architectures. Circuit modeling must be consumable at ever-higher levels, as systems move to compiler- and operating system-based control of such parameters as body bias, clock, and supply rails. For SOI, design technology has already made progress toward necessary analyses (e.g., history-dependence of timing, coupling dependence of static power) and syntheses (e.g., planning of decoupling capacitance).

3.2. Resource Analyses

In this subsection, I will touch on three facets of design technology, in preparation for reconciling design-process integration needs versus available resources.

3.2.1. Design Technology Human Resources

According to Dataquest, the electronic design automation (EDA) industry as a whole has less than 6,000 R&D engineers worldwide. Total EDA tools revenue per IC designer has increased at 3.9% per year on average for the past decade.^{4,8} Pure research investment is low relative to other semiconductor supplier industries: a recent survey,¹⁰ made at the behest of the Semiconductor Research Corporation, suggests that the electronic design automation industry invests approximately 0.11% of total sales, or 0.5% of total R&D budget, on ITRS-timescale research. This amounts to around \$4M per year.

As noted in the ITRS document, most of today's design technology "crises" – verification, embedded software and system-level design, enabling of power management at the architectural and operating-system levels, analog/mixed-signal synthesis and reuse, design for test – lie at a "high level", away from the design-manufacturing interface. This is not surprising: the ratio of design value over effort is *perceived* to decrease as the level of abstraction moves from behavior down to layout. Such a perception creates an *a priori* bias against the portion of EDA that has most of the design-manufacturing interfaces - *physical design* - which even today is only at most one-sixth (whether by market size or by headcount) of EDA and design technology as a whole. From the design technology side, design-process integration is clearly hampered by a lack of researchers, developers and funding.

3.2.2. Demands on Design Technology

A key message in the 2001 Roadmap is that *cost of design* is the greatest threat to continuation of the semiconductor industry's phenomenal growth. Manufacturing nonrecurring engineering (NRE) costs are just reaching \$1 million (mask set and probe card), whereas design NRE costs routinely reach tens of millions of dollars. We measure manufacturing cycle times in weeks, with low uncertainty, whereas we measure design and verification cycle times in months or years, with high uncertainty. Moreover, design shortfalls are responsible for silicon respins that multiply manufacturing NRE costs. Software now routinely accounts for 80 percent of embedded-systems development cost. Test cost has grown significantly relative to total manufacturing cost. Verification engineers are twice as numerous as design engineers on microprocessor project teams - and the list goes on. These trends are illustrated in the "homunculus" of Figure 1. Traditional "hardware design" contains just a tiny fraction of the design technology challenges. If design technologists appear to be ignoring the design-process integration challenge, it is because there are bigger fish to fry.[§]

We may summarize the present state of resources in design technology as follows. (1) Despite an acknowledged design productivity gap (Figure 2 reproduces the famous Sematech analysis from 1994) by which the number of available transistors grows faster than the ability to design them meaningfully, investment in process technology has by far dominated investment in design technology. (2) The good news is that the design technology (DT) industry continues to enable semiconductor industry productivity. The estimated design cost of a leading-edge system-on-chip (SOC) - i.e., the "low-power PDA" specified as the SOC system driver - was approximately \$15 million in 2001. As shown in Figure 3, this cost would have been \$342 million if DT innovations between 1993 and 2001 had not occurred.[¶] (3) The bad news is apparent from the "homunculus": many previous design technology gaps have become crises.

[§]Consider the list of silicon and system complexity challenges given in the previous section. Design technology must spread its attention on alpha particles, real-time operating systems, and everything in between. This is one reason why only around 7% of the 2001 ITRS Design Chapter deals with issues that can be said to lie in the realm of design-process integration.

[¶]At the request of the 2001 ITRS Design working group, Gartner/Dataquest measured designer productivity at 4K gates (= 16K transistors) per year in 1990 - the year in which the so-called "Register-Transfer Level (RTL) methodology" originated - and calibrated design productivity improvements for seven major design technology (DT) innovations that have occurred or are anticipated since then. These improvements and their dates of deployment are: (1) in-house place-and-route, i.e., at the designer's site, 1993; (2) "tall-thin engineer", with small design teams able to take designs through synthesis, place and route, 1995; (3) small-block (blocks of 2,500 - 74,999 gates) reuse, 1997; (4) large-block (blocks of 75,000 - 1M gates) reuse, 1999; (5) IC implementation suite, a set of tightly integrated tools that takes an IC design from RTL synthesis through IC place-and-route and GDSII output, 2001; (6) intelligent testbench, an RTL verification automation tool ("cockpit") which partitions a system-level design description into verifiable blocks, then

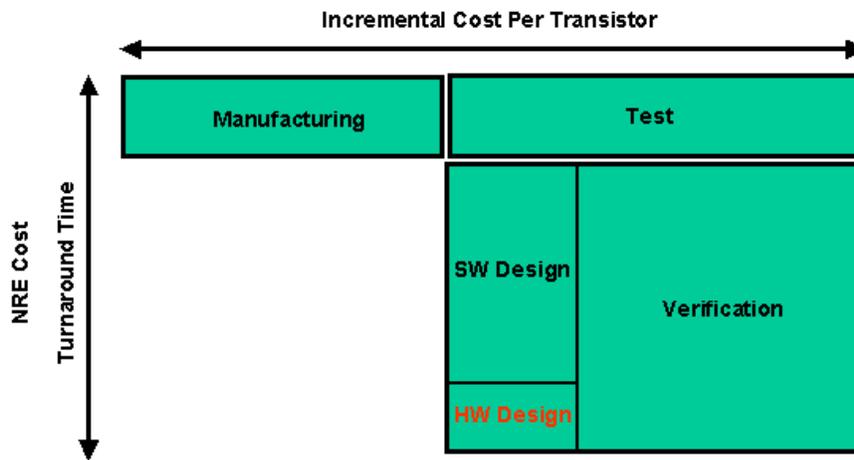


Figure 1. “Homunculus” of design technology challenges. Traditional hardware design issues occupy a very small, and decreasing, fraction of the design community’s attention.

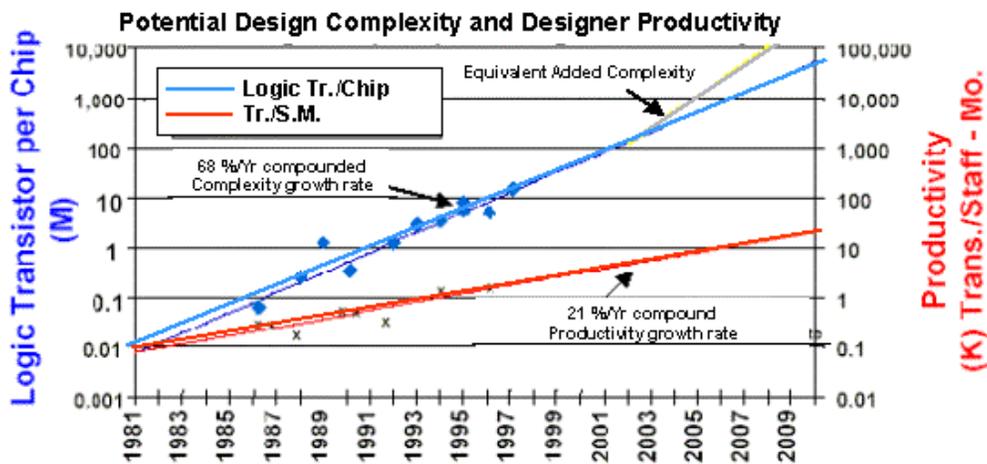


Figure 2: “Design productivity gap” published in 1994 by Sematech.

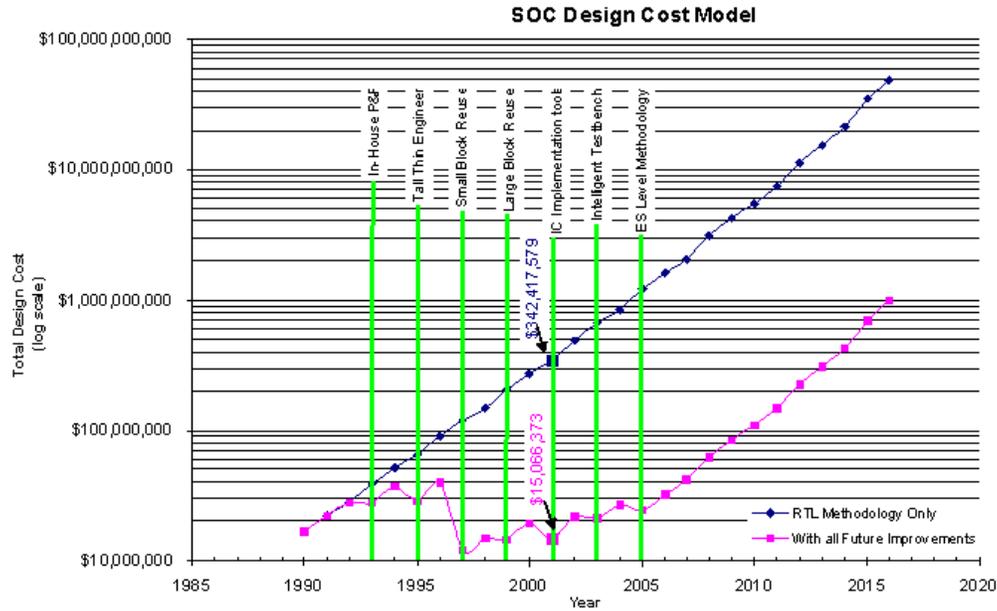


Figure 3. Analysis of design cost for the SOC Low-Power PDA design specified in the System Drivers Chapter of the 2001 ITRS. Without design technology innovations from 1993 to 2001, the design cost would have exceeded \$342M.

3.2.3. A Comment on the Design Productivity Gap

Finally, the linkage between design technology and the overall success of the semiconductor industry should not be underestimated. The role of design technology is to “fill the fab”, i.e., keep manufacturing facilities fully utilized with high-volume, *high-value* (= high-margin) parts. With any productivity failure (e.g., attributed to design technology), there will be a lack of *high-quality* designs, and hence of *high-value* designs. Figure 4 shows that quality multipliers (say, for MOPS/mW) can vary by four or more orders of magnitude between custom direct-mapped hardware and generic computing platforms (general-purpose microprocessors). Such quality is difficult to sacrifice: if there are not enough high-value designs, the semiconductor industry must look for a “workaround” (e.g., reconfigurable logic, platform-based design, embedded software, etc.). The key observation with respect to such workarounds is that they begin to extract value from somewhere other than silicon differentiation. It is not a big step to the electronics industry finding analogous workarounds (i.e., extracting value from somewhere other than silicon differentiation)

4. SHARED RED BRICKS

As noted in^{?, 5, 6} and above, the ITRS defines a “red brick” as 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

executes appropriate verification tools while tracking code coverage, 2003; and (7) “electronic system-level” (immediately above RTL, and including both hardware and software design) methodology, 2005. Figure 3 quantifies the impact of these DT innovations on total design cost for the low-power System-on-Chip (SOC-LP) PDA driver defined in the 2001 ITRS System Drivers chapter. The model sets the historical rate of increase in engineer cost at 5% per year (salary and overheads starting at \$181,568 in 1990), and the rate of increase in design tool cost at 3.9% per year (starting at \$99,301 per engineer in 1990). The number of designers per million logic gates is 250 in the year 1990 and 11 in 2001 (implying an average of 32.8% per year improvement in number of logic gates per designer-year). The low-power SOC PDA model (System Drivers chapter) has 3M logic gates in 2001, implying an SOC PDA design cost (designers + tools) of \$15.1M. Without the five major DT innovations that occurred between 1993 and 2001, the design cost for the same SOC in 2001 would be approximately \$342.4M.

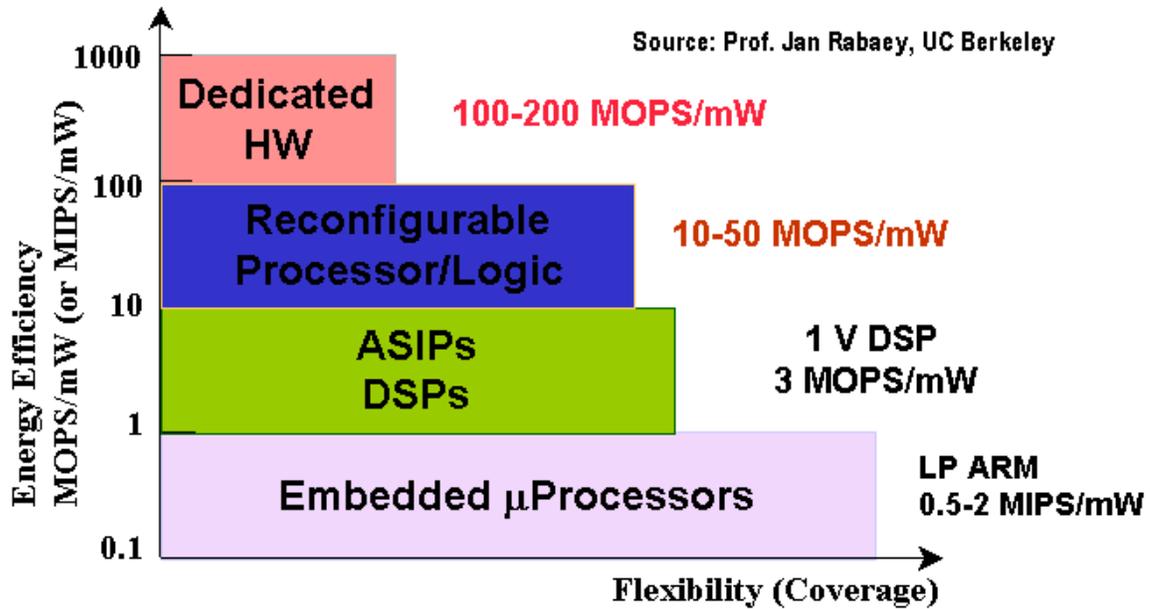
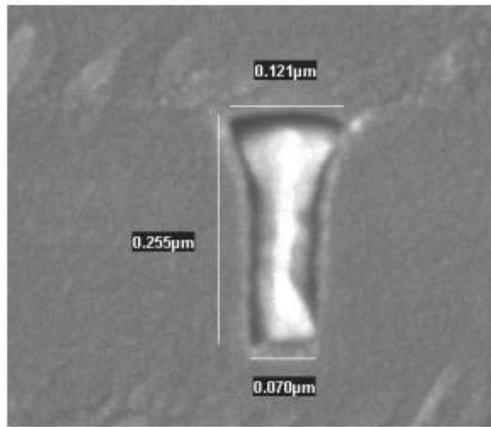


Figure 4. Quality (which influences value) can vary by four orders of magnitude or more as we traverse the range of implementation choices from direct-mapped custom designs to general-purpose computing platforms. Unavailability of custom designs due to a lack of design technology effectively reduces the efficiency of value extraction from the underlying manufacturing capability.

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. Recall the metaphor of the ITRS - the semiconductor industry's technology foundations - as a car, with the supplier industries (packaging, lithography, design, etc.) being the parts of the car. The car must continue to be driven along the Moore's Law road, e.g., if the car goes 150mph today then four years from now we require the car to reach speeds of 600mph. It is absurd to think that super tires alone, or super seats alone, will get us to 600mph. 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 600mph 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, i.e., "shared red bricks".

Three examples of potential shared red bricks are as follows. (1) 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 Design 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 \$/wafer) are a long ways off. Appropriate (and standardized) characterizations of variability sources in manufacturing equipment and processes also appear to be a long ways off. Circuit and layout techniques for high-variability regimes must also be explored. (2) 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. (3) Do we really need dielectric permittivities below 2.0 or copper interconnect resistivities below 1.8 $\mu\Omega$ -cm, as specified in the ITRS? Are performance and density gains from the former worth the huge processing and variability costs? Is the latter even possible, let alone worth it (see Figure 5)? Or would developing better



C. Case, BOC Edwards – ITRS-2001 preliminary
 Courtesy of SEMATECH

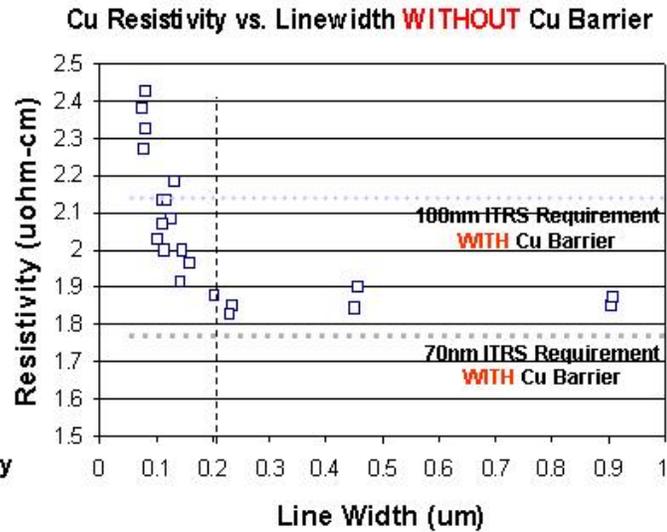


Figure 5. Copper resistivities, even without considering the barrier layers necessitated by reactivity of the material, are far above ITRS-mandated values. At some point, the semiconductor industry must ask whether it is worth pursuing impossible “solutions” to questionable “requirements”.

circuit and interconnect architectures, and better layout techniques (cf. the X Initiative¹²), more cost-effectively share performance and noise management red bricks between Design and, e.g., Interconnect technologies?

These examples highlight the potential of deeper partnerships between design technology and other ITRS technology industries. Such partnerships can potentially resolve key red bricks at greatly reduced cost to the semiconductor industry. It would be natural for the semiconductor supplier community to pursue not only shared solutions but also shared rewards from a new “virtuous cycle” of ROI, productivity and value.

ACKNOWLEDGMENTS

This is not a scientific paper - it is only “review and commentary”. References are not intended to be complete. Many passages have been excerpted from material that I have written for the Design and System Drivers chapters of *The International Technology Roadmap for Semiconductors* (2001 Edition),⁴ as well as various papers. I am grateful to my coauthors, to the many individuals who contributed to the 2001 ITRS, and to various copyright-holding entities for their indulgence.

REFERENCES

1. R. Brashears and A. B. Kahng, “Advanced Routing for Deep Submicron technologies”, *Computer Design*, May 1997. <http://www.computer-design.com/Editorial/1997/05/supplement/597suprouting.html>
2. A. E. Caldwell, A. B. Kahng and I. L. Markov, “Toward CAD-IP Reuse: A Web Bookshelf of Fundamental CAD Algorithms”, *IEEE Design and Test*, May 2002.
3. Y. Cao, P. Gupta, A. B. Kahng, D. Sylvester and J. Yang, “Toward a Framework for Assessing the Impact of Variability in Nanometer Design”, *manuscript*, 2001.
4. *The International Technology Roadmap for Semiconductors*, 2001 edition, International Sematech, Austin, Texas, December 2001. <http://public.itrs.net/>
5. A. B. Kahng, “The Road Ahead” (column), *IEEE Design and Test*, 2002.
6. A. B. Kahng, *various talks*, <http://vlsicad.ucsd.edu/~abk/TALKS/>.
7. A. B. Kahng and Y. C. Pati, “Subwavelength Optical Lithography: Challenges and Impact on Physical Design”, *Proc. Intl. Symp. on Physical Design*, April 1999, pp. 112-119. See <http://vlsicad.ucsd.edu/ISPD99TUTORIAL/ispdtutorial.ppt> (similar presentation material is in <http://vlsicad.ucsd.edu/ISQED00TUTORIAL/isqed2000tutorial-abk.ppt>).
8. A. B. Kahng and G. Smith, “A New Design Cost Model for the 2001 ITRS”, *Proc. ISQED*, March 2002.
9. M. D. Levenson, J. S. Petersen, D. G. Gerold and C. A. Mack, “Phase Phirst! An Improved Strong-PSM Paradigm”, *Proc. 20th Annual BACUS Symposium on Photomask Technology*, Monterey, CA, USA, 13-15 Sept. 2000, SPIE vol. 4186, pp. 395-404.

10. C. Nuese, "Research Needs vs. Funding: 2001 - A Report for the SRC", 2001.
11. L. Scheffer, "A Roadmap of CAD Tool Changes for Sub-Micron Interconnect Problems", *Proc. Intl. Symp. on Physical Design*, 1997, pp. 104-109.
12. <http://www.xinitiative.org>