Design-Process Integration and Shared Red Bricks

Andrew B. Kahng
abk@ucsd.edu
UCSD CSE and ECE Departments, La Jolla, CA 92093-0114

ABSTRACT

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

1. INTRODUCTION

Let us begin with an analogy. Think of the International Technology Roadmap for Semiconductors (ITRS)[9] as a car that represents the entire semiconductor industry and its supplier industries (Design, Interconnect, Packaging, Test, Lithography, etc.). In the past, this car has been driven along the "Moore's Law Road" by two drivers, MPU and DRAM. (In some sense, the car has been driven in a straight line, as well.) 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* explain more clearly who they are, and where they are going. The result: a new System Drivers Chapter in the 2001 ITRS, which attempts to explicitly define the key IC products that drive semiconductor manufacturing and design technologies.

Against the common backdrop of the System Drivers and the Overall Roadmap Technology Characteristics, the years 2000-2001 saw more than 800 experts world-wide contribute to the roadmapping of eight distinct technologies in the semiconductor supply chain. This process resulted in the 2001 ITRS. In this invited paper, I will discuss some aspects of the interaction between Design Technology (DT) and manufacturing, or process, technologies.

2. THE ROADMAP OF DESIGN TECHNOLOGY CHALLENGES

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.

Despite an acknowledged design productivity gap in 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. The good news is that developers continue to make progress in design technology (DT): The estimated design cost of a low-power SoC PDA was approximately \$15 million in 2001 versus \$342 million if DT innovations had not occurred between 1993 and 2001 (see Appendix A). The bad news is that 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. In 2001, many previous design technology gaps became crises.

2.1. Complexity Challenges

Design technology faces two basic types of complexity: silicon and system.

Silicon complexity refers to the impact 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:

- 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.

Silicon complexity places long-standing paradigms at risk. System-wide synchronization becomes infeasible due to power limits and the cost of robustness under manufacturing variability. The CMOS transistor becomes subject to ever-larger statistical variabilities in its behavior. And fabrication of chips with 100 percent working transistors and interconnects may become prohibitively expensive.

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.

The following precepts govern future (physical) design technology innovation. Precepts 3 through 6 are aspects of a (top-down, iteration-free, decomposition-oriented) "correct by construction" approach. Precepts 7 and 8 are more suited to "construct by correction", where iterations are expected but made less painful. Ultimately, these lead to the future design system architecture depicted in Figure 1.

1. Exploit reuse.

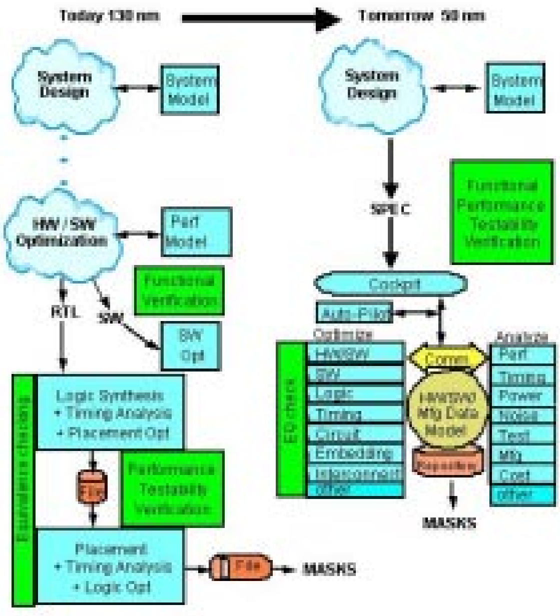

Figure 1: Slide12

2. Evolve rapidly. Typical evolutions: (a) analyses evolve into verifications, which evolve into tests, and (b) analyses and simulations evolve into models and verifications, which evolve into either objectives or constraints for synthesis and optimization. A related trend is "bottom-up commoditization" (e.g., of characterization and RLC extraction, then delay calculation, static timing / noise analyses, standard-cell placement, global routing, ...).

3. Avoid iteration. Iteration between levels of design incurs translation and other interfacing costs, and hampers predictability and reliability of the design process.

4. Replace verification by prevention. Lower-level problems (e.g., crosstalk-induced delay uncertainty) are more cheaply addressed by higher-level prevention (e.g., repeater insertion and slew rate control rules).

5. Improve predictability. Constructive estimation does not afford productivity leverage; efficient search for good design solutions requires prediction-based estimation. Better estimates enable design space exploration at higher levels.

6. Orthogonalize concerns (e.g., congestion from timing, timing from layout, computation from communication). Unrelated issues should remain separate whenever possible.

7. Expand scope. For example, we require greater integration of software and analog/mixed-signal/RF (AMSRF) design with digital flows; this must be supported by modeling, analysis and simulation at multiple levels of the system hierarchy (up to package and board, and down to mask and process).

8. Unify. Silicon complexity induces the unification of previously disparate areas, e.g., synthesis-analysis, logical-physical-timing, or even design-test. Unifications improve the downstream flow of intentions and assumptions, and the upstream flow of estimation/prediction models. Associated frameworks are successive approximation, and incremental optimization. Increasingly, unifications cross the die-package and design-manufacturing boundaries.

2.2. Crosscutting challenges.

The Roadmap sets out detailed challenges 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 2001 Roadmap also identifies five crosscutting challenges that encompass all relationships between electronic design automation and the other industries that support the semiconductor industry whose solutions are distributed across all areas of design technology. Our strong hope is that challenges impossible to solve within a single ITRS technology area are solvable with design technology partnership. For example, more rapid adoption of new fault models for crosstalk and path delay, along with corresponding automatic test pattern generation (ATPG) and BIST techniques, might address test equipment and speed limitations. The 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

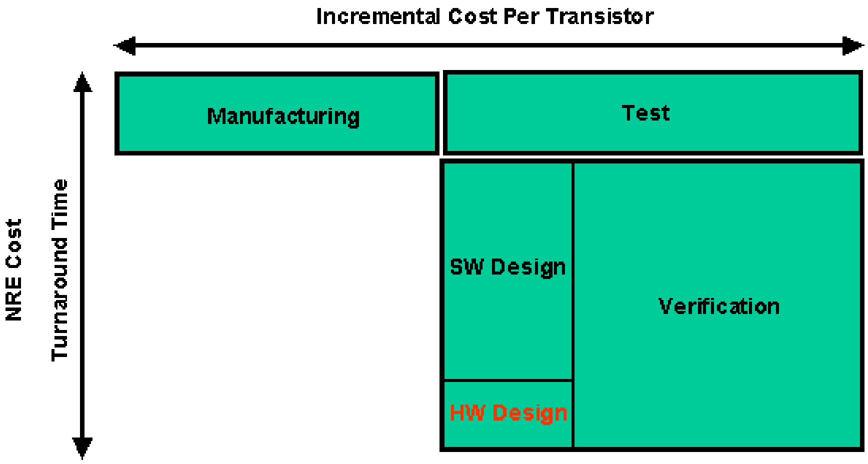

Figure 2: Slide 13

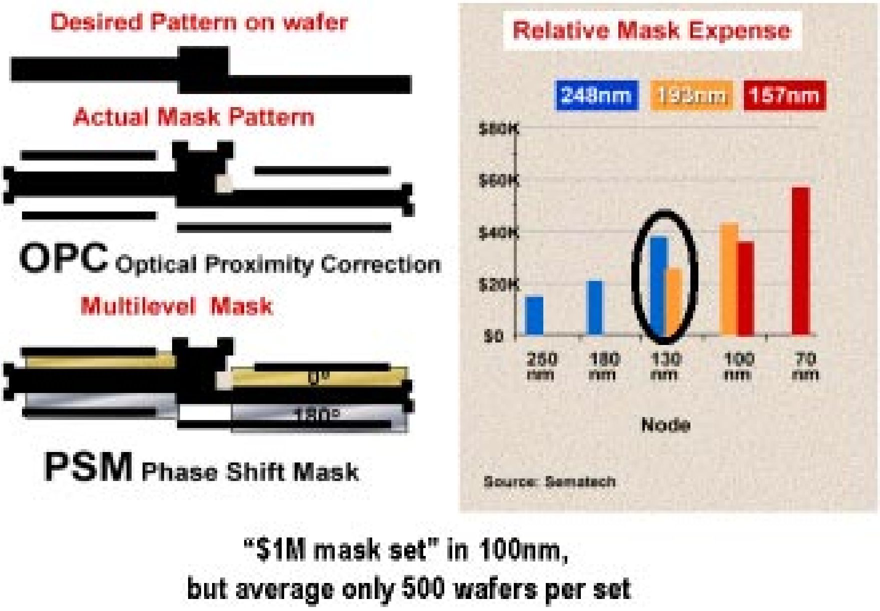

"$1M mask set" in 100nm,
but average only 500 wafers per set

Figure 3: Slide 19

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.

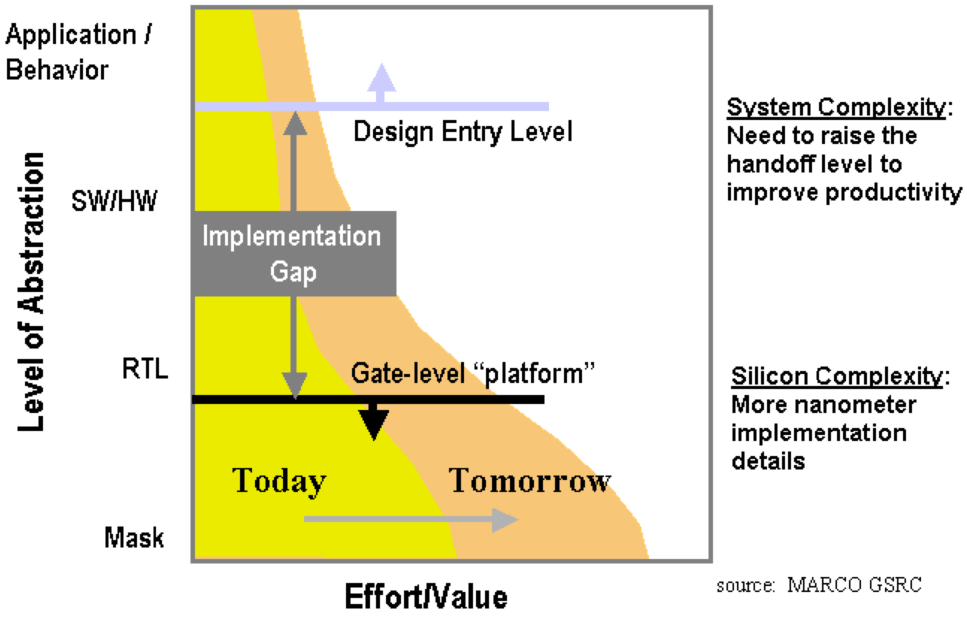

System Complexity: Need to raise the handoff level to improve productivity

Silicon Complexity: More nanometer implementation details

source: MARCO GSRC

Figure 4: Slide 20

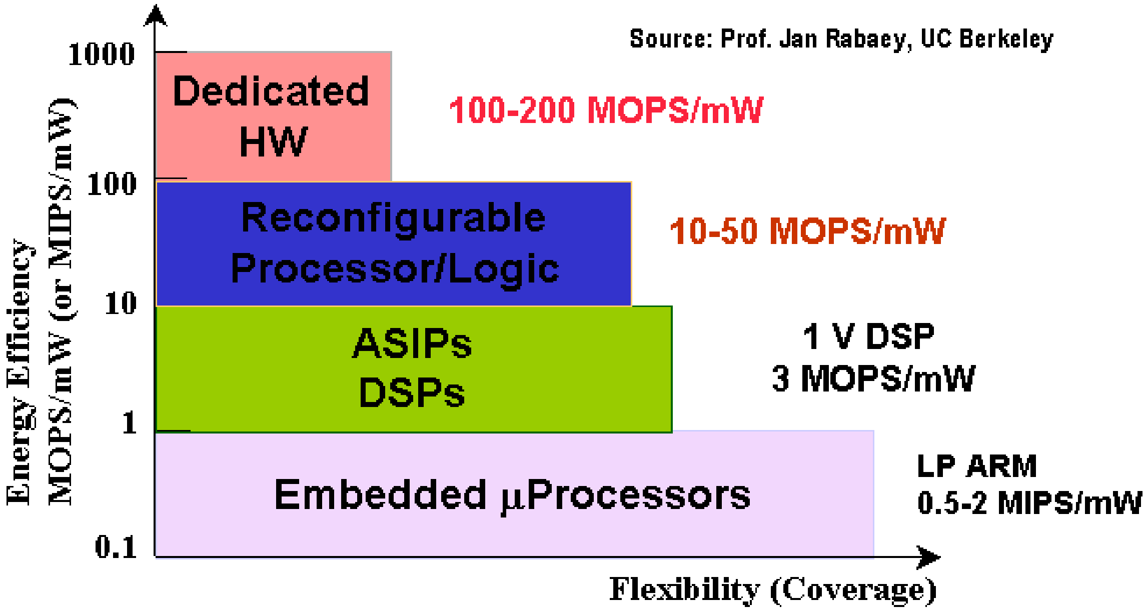

Source: Prof. Jan Rabaey, UC Berkeley

Figure 5: Slide 25

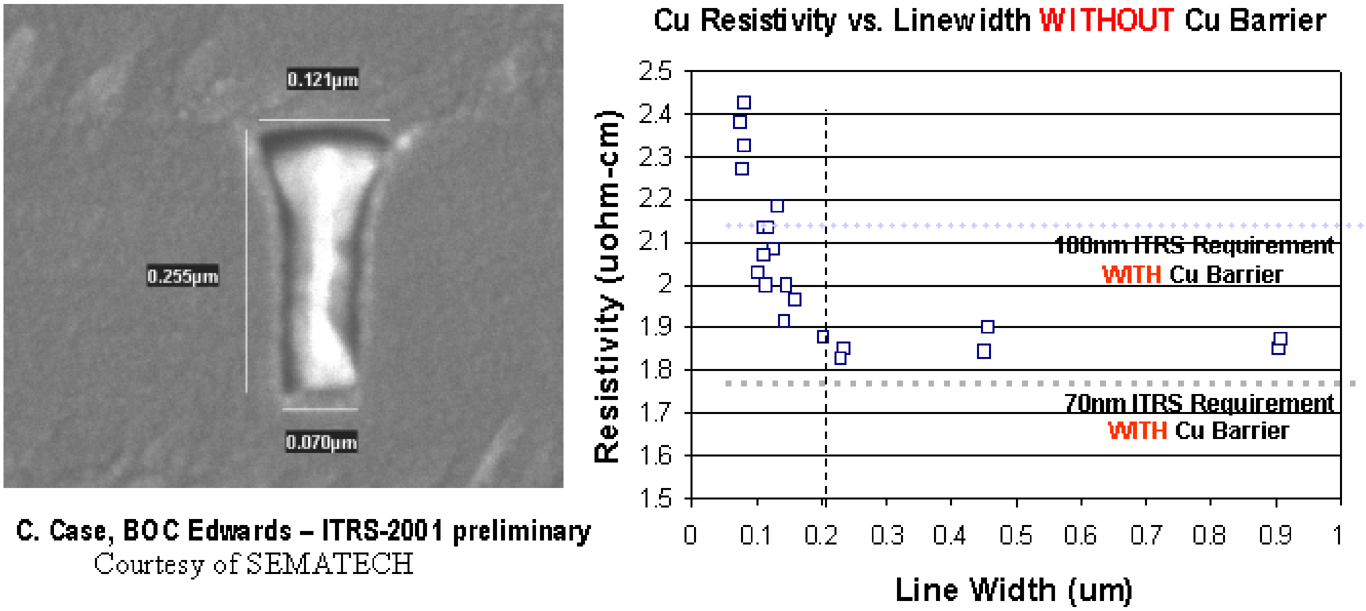

C. Case, BOC Edwards – ITRS-2001 preliminary
Courtesy of SEMATECH

Figure 6: Slide 33

3. MULTIPLE-LAYER OXIDE CMP DUMMY FILL

3.1. Density Models for Oxide CMP Dummy Fill

4. CONCLUSIONS

ACKNOWLEDGMENTS

This is not a scientific paper - rather, it is more of a “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 —em The International Technology Roadmap for Semiconductors (2001 Edition),[9] 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.

Appendix A. ITRS-2001 Design Cost Model

The vital contribution of design technology to semiconductor product realization is shown by the following analysis. At the request of the 2001 ITRS Design ITWG, 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 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 **??** 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

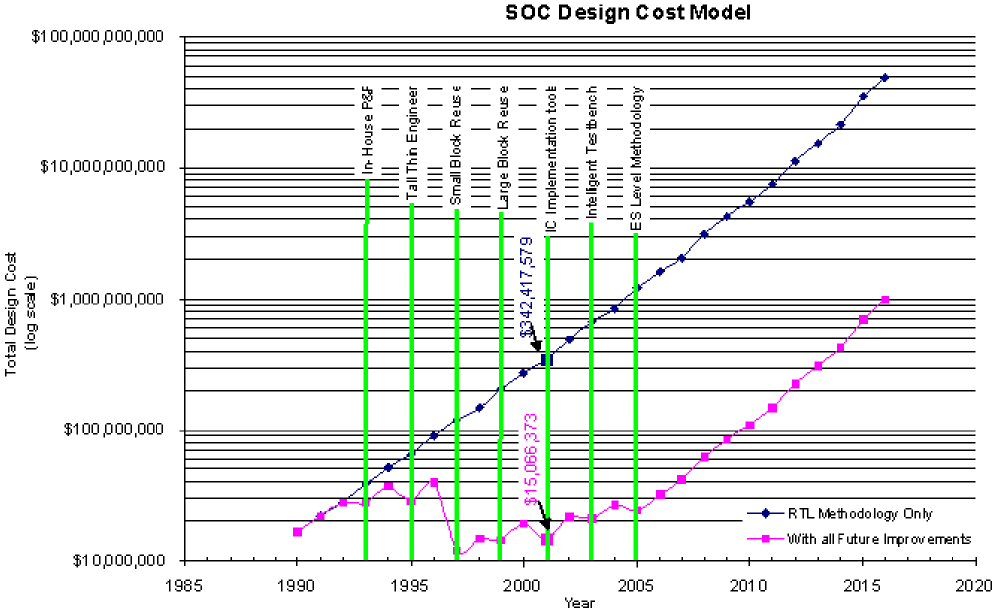

Figure 7: Slide12

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.

REFERENCES

1. C. J. Alpert, L. Hagen and A. B. Kahng, "A Hybrid Multilevel/Genetic Approach for Circuit Partitioning", *Proc. ACM SIGDA Physical Design Workshop*, April 1996, pp. 100-105.
2. C. J. Alpert, J. H. Huang and A. B. Kahng, "Multilevel Circuit Partitioning", *Proc. ACM/IEEE Design Automation Conf.*, 1997, pp. 530-533.
3. 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
4. M. Burstein and R. Pelavin, "Hierarchical Wire Routing", *IEEE Trans. CAD* 2(4) (1983), pp. 223-34.
5. A. E. Caldwell, A. B. Kahng, A. A. Kennings and I. L. Markov, "Hypergraph Partitioning for VLSI CAD: Methodology for Reporting, and New Results", *Proc. ACM/IEEE Design Automation Conf.*, June 1999, pp. 349-354.
6. 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.
7. A. E. Caldwell, A. B. Kahng and I. L. Markov, "Can Recursive Bisection Alone Produce Routable Placements?", *Proc. ACM/IEEE Design Automation Conf.*, 2000, pp. 477-482. (Cf. the "quadratic placement revisited" work as well.)
8. 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.
9. *The International Technology Roadmap for Semiconductors*, 2001 edition, International Sematech, Austin, Texas, December 2001. http://public.itrs.net/
10. A. B. Kahng, "Classical Floorplanning Harmful?", *Proc. Intl. Symp. on Physical Design*, April 2000, pp. 207-213. See talk slides at http://vlsicad.ucsd.edu/papers/slides/ispd00-cfh.ppt
11. A. B. Kahng, "The Road Ahead" (column), *IEEE Design and Test*, 2002.
12. A. B. Kahng, *various talks*, http://vlsicad.ucsd.edu/~abk/TALKS/ .
13. 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).
14. A. B. Kahng and G. Smith, "A New Design Cost Model for the 2001 ITRS", *Proc. ISQED*, March 2002.

15. G. Karypis, R. Aggarwal, V. Kumar and S. Shekhar, "Multilevel Hypergraph Partitioning: Application in VLSI Domain", *Proc. ACM/IEEE Design Automation Conference*, 1997, pp. 526-529.
16. 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.
17. R. Nair, "A Simple Yet Effective Technique for Global Wiring", *IEEE Trans. on CAD* 6(6) (1987), pp. 165-172.
18. N. R. Quinn and M. A. Breuer, "A Force-Directed Component Placement Procedure for Printed Circuit Boards", *IEEE Trans. on Circuits and Systems* 26(6) (1979), pp. 377-88.
19. S. Savage, "On the Biodiversity of SIGCOMM", *2000 ACM SIGCOMM Outrageous Opinions*, Stockholm, Sweden, August 2000.
20. S. Savage, "On the Caching and Prefetching of Program Committees", *1999 ACM SIGCOMM Outrageous Opinions*, Cambridge, MA, September 1999.
21. L. Scheffer, "We're Solving the Wrong Problems", *Proc. 6th Physical Design Workshop*, April 1996, Reston, Virginia, pp. 89-91.
22. L. Scheffer, "A Roadmap of CAD Tool Changes for Sub-Micron Interconnect Problems", *Proc. Intl. Symp. on Physical Design*, 1997, pp. 104-109.
23. B. S. Ting and B. N. Tien, "Routing Techniques for Gate Array", *IEEE Trans. on CAD* 2(4) (1983), pp. 301-312.
24. P.-S. Tseng and C. Sequin, "Codar: A Congestion-Directed General Area Router", *Proc. ICCAD*, 1988, pp. 30-33.
25. http://www.chipcenter.com/eda/dac_files/ppt/DAC99-Avant3.ppt
26. http://www.eetimes.com/dac98/news_tera.html
27. http://www.synopsys.com/products/layout/floorplan_cs.html (last modified October 13, 1997).
28. http://www.xinitiative.org
29. Workshop on Duplicating, Deconstructing, and Debunking (held in association with the Intl. Symp. on Computer Architecture), http://www.ece.wisc.edu/~wddd .