

Key Directions and a Roadmap for Electrical Design for Manufacturability

(Invited Talk)

Andrew B. Kahng

abk@ucsd.edu Depts. of CSE and ECE, University of California at San Diego, La Jolla, CA 92093-0404 USA

Abstract. Semiconductor product value increasingly depends on “equivalent scaling” achieved by design and design-for-manufacturability (DFM) techniques. This talk addresses trends and a roadmap for “equivalent scaling” innovation at the design-manufacturing interface. The first part will discuss precepts of *electrical* DFM. What are dominant aspects of manufacturing variability and design requirements? Can designs match process, or must process inevitably adapt to designs? In what sense can concepts of “virtual manufacturing” or “statistical optimization” succeed in the design flow? How should design technology balance analyses that preserve value, versus optimizations that extend value? How should we balance preventions (correct by construction), versus early interventions, versus cures (construct by correction), versus “do no harm” opportunism? Or, tools that can model and predict well, versus tools that can make upstream assumptions come true? The second part will give a roadmap for electrical DFM technologies, motivated by emerging challenges (stress/strain engineering, mask errors, double-patterning lithography, etc.) and highlighting needs for $\leq 45\text{nm}$ nodes.

1. Introduction: Scaling to $\leq 45\text{nm}$

Semiconductor manufacturing technology faces ever-greater business challenges of capital cost and risk, along with ever-greater technical challenges of pitch, mobility, variability, leakage, and reliability. To enable cost-effective continuation of the semiconductor roadmap, there is greater need for design technology to provide “equivalent scaling”, and for product-specific design innovation to provide “more than Moore” scaling.

Equivalent scaling – the introduction of design technologies that reduce power or improve density without requiring innovation in process technology – is required to continue Moore’s Law trajectories of performance, density and cost. Conservatively, half of a process node of power, a third of a node of area, and one full node of performance can be easily gained, with the only question being whether the industry will recognize and invest sufficiently in this opportunity. *Design innovation* – multi-core architecture, software support, beyond-die integration, etc. – will be the workhorse for “more than Moore” scaling that goes beyond what underlying process and design technologies can achieve. One can envision a roadmap aligned to the continued delivery of semiconductor *value*, with that value arising from a combination of manufacturing technology, design technology, and design innovation. Finally, the balance of these contributors to scaling must be determined by a new understanding of *system scaling*. In the future, Dennard’s classical scaling theory [17] no longer holds, and nor will ITRS-style scaling in terms of such parameters as ‘A factors’ or ‘FO4 delays’ or ‘CV/I metrics’. The future of semiconductor technology scaling will instead be dominated by

application- and product-specific system constraints on reliability, adaptivity, cost, reusability, software support, etc.

Technical concerns with scaling to $\leq 45\text{nm}$ will inevitably include the following.

- *Variability.* Critical dimension (CD) control and process variations will challenge both manufacturing and design until the end of the CMOS roadmap. In mature 45nm products, local pattern and pitch dependencies in resist and etch processes will be ameliorated by restricted layout rules and improved dummy structure methodologies, but these issues are still problematic today. Back end of the line (BEOL) RC performance will exhibit increased variability even as ‘percentage-wise’ control of chemical-mechanical polishing (CMP) remains constant or improves. This is due to, e.g., non-scaling of barrier thickness that magnifies RC impacts of CMP-induced thickness variation, and CD variation induced by etch and lithographic defocus [15].
- *Leakage currents* (subthreshold leakage, gate leakage, junction leakage, band-to-band tunneling, etc.) remain a dominant concern into the 45nm node. Leakage power is not only “wasted”, but also compromises achievable form factor, integration density, packaging choice, reliability, and other product metrics. Further, variability impact on leakage is substantial: subthreshold leakage scales exponentially with respect to both operating temperature¹ and transistor gate CD (i.e., channel length), and total leakage can vary 5X to 20X across chips from the same wafer lot. Mitigation of leakage through multi- V_{th} , MTCMOS, or higher-level design techniques incurs area overhead and design process complexity, along with added variability.²
- *Stress and reliability.* Today’s scaling of on-current and device speed is based on stress and strain engineering (embedded SiGe, stress memory, dual stress liner, etc.). Shallow trench isolation (STI) stress can affect device on-current by up to 40%. FEOL stress changes mobility and threshold voltage of transistors; BEOL stress affects interconnect integration and reliability. By the late 45nm node, design tools and methodologies must actively modulate stress to improve timing (mobility change) as well as leakage (threshold voltage change).

¹ Each additional °C of ambient temperature (e.g., due to global warming) will increase subthreshold leakage by roughly 5%.

² E.g., random dopant fluctuations and reduced supply voltage headroom make triple-V_t strategies less viable in future nodes.

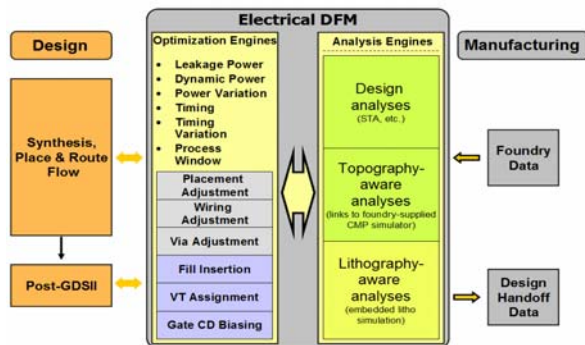
Chip design must holistically comprehend *unpredictability* from wearout (NBTI, TDDB, electromigration, etc.), parametric variation (line-edge roughness, random dopant fluctuation) and transient phenomena (particle strikes, supply noise); here, the challenge will be to minimize the cost of on-chip monitoring, redundancy, and reconfiguration structures. Of course, many other challenges will compete for attention, ranging from next-generation lithography and consensus on radical layout restrictions, to software development for the highly concurrent, multi-core SOCs to which many applications have converged.

2. Electrical Design for Manufacturability (DFM)

By $\leq 45\text{nm}$, parametric failures – chips that fail to meet power and timing requirements – will become a dominant yield-limiting mechanism. In this context, there are many opportunities for design for manufacturability (DFM) tools that bridge chip design/implementation and process/manufacturing know-how, to deliver high-value equivalent scaling advances. The precepts

- drive design requirements into manufacturing;
- bring manufacturing awareness into design; and
- work within existing design environments without requiring major changes to the design flow, the design signoff, the handoff to manufacturing, or process flow

must be kept in mind if the industry is to eventually achieve a true “design for value” (maximizing profit per wafer) capability.



Of particular interest is the notion of “electrical DFM”, which focuses on objectives that the designer or product engineer cares about: leakage power, dynamic power, timing, timing and power variability, timing, process window, and even reliability. As illustrated in the figure above [13], the drivers for such optimizations consist of analysis engines that comprehend a full spectrum of physical and electrical implications of manufacturing. The “knobs” or degrees of freedom to achieve the optimization goals include changes to placement, wiring, vias – even the dimensions of individual transistors.

Several prototype or near-production techniques exemplify how electrical DFM solutions can take into account design-specific information to improve design analyses and optimizations. These include (1) iso-dense awareness of pitch-dependent through-focus CD variation, to reduce timing guardbands and improve timing robustness [12]; (2) post-layout transistor gate-length biasing, specified at tapeout but realized in the foundry’s OPC flow, to reduce leakage and leakage variability [10]; (3) “self-compensating design” techniques that minimize the inherent sensitivity of critical paths to various sources of process variation

(dopant density, oxide thickness, Leff, etc.; see, e.g., [11]); and (4) timing- and SI-driven CMP fill that maximizes both timing robustness and post-CMP wafer uniformity [13,14]. These techniques lie along a necessary trajectory for the industry as it addresses manufacturing variability:

- first, address systematic variation (“model-predict-compensate” or “measure-model-mitigate”) as in [12];
- second, make designs robust to variation, e.g., by forcing the sum of sensitivities to a given variation source (wire thickness, defocus, V_{th} shift, etc.) to zero as in [11]; and
- third, address remaining random variations, e.g., through statistical timing and leakage optimization techniques.

3. Reality Checks

In an ideal future at $\leq 45\text{nm}$, chip designers’ power and timing requirements will be used to tailor the manufacturing line for each individual transistor of each individual design – without any changes or adjustments to fab equipment. One can envision that designers will be able to take advantage of available entitlement or process margin so that the process delivers significantly improved parametric quality of the silicon product – and that designs can be driven to a sweet spot for the process, just as the process is today driven to a sweet spot for the design. This being said, at least two important realities must be acknowledged.

Time constants and guardbanding. In the co-evolution of silicon technology and silicon products, applicable time constants range over nearly three orders of magnitude:

- O(years): technology development; application market definition; architectural and front-end design
- O(quarters): SPICE model revision; design rule manual revision; library/IP design; library/IP silicon qualification
- O(months): library/IP modeling/characterization; RTL-to-GDS implementation; reliability qualification
- O(weeks): fab latency (wafer start to wafer out); cycle of yield learning; design re-spin; OPC and mask flow
- O(days): process tweak; design ECO

A number of precedence and practical constraints also hold, e.g., the SPICE model version 1.0 must be fixed before libraries/IPs are fixed; libraries/IPs must be fixed before RTL-to-GDS physical implementation can occur; only limited changes to the SPICE model are permissible after a certain volume of library/IP/chip design activity has taken place; etc. Furthermore, even though a design change can be made in O(days), the latency for assessment in silicon must span the OPC, mask and foundry flows. Hence, (1) the process must continue to adapt to the design, as it does today; and (2) the ability of the foundry to tweak the process even when SPICE and RCX models are fixed implies that significant guardbanding, i.e., overdesign, is inherent in today’s design-foundry relationship. With this in mind, R&D foci for $\leq 45\text{nm}$ nodes must be driven by quantified assessment of guardbanding costs and benefits (cf., e.g., [18]), and also enable

- design adaptation to process in the face of significant, and possibly intentional, model-to-silicon miscorrelations; and
- more rapid process adaptation to design, e.g., through improved understanding of how parametric tests in the fab map through SPICE models to design signoff constraints.


Practicality and value of statistical design. Notwithstanding the “necessary trajectory” noted above, significant challenges lie ahead with respect to modeling and mitigation of manufacturing variation. The EDA community has rapidly developed various statistical analyses and optimizations, but it is unclear how the semiconductor industry as a whole will reach consensus on enablement of such techniques for production flows. We note that inter-die (die-to-die, or DTD) variations are easier to model in the manufacturing process, as well as in statistical design techniques. On the other hand, intra-die (within-die, or WID) variations have a significant component that is systematic and pattern-dependent. With many distinct variability phenomena and length scales in play – from wafer radial bias, reticle bending, lens aberration, CMP planarization length, flare, etc. down to mask CD and mask error enhancement factor (MEEF), etch and litho – modeling of spatial and pattern-dependent correlations is a key challenge to deployment of statistical design flows at $\leq 45\text{nm}$.³ Another key challenge is to demonstrate sufficient ROI from statistical design approaches. The work of [19] showed limited impact of statistical power optimization. Intuitively, statistical design has only limited impact with respect to “sum” objectives such as power, as opposed to “max” objectives such as timing. Impact will also be limited for phenomena such as subthreshold leakage which are exponential in most parameters (L_{eff} , temperature, etc.) and for which sensitivities and variances track nominal values.⁴

4. Elements of an Electrical DFM Roadmap

This section presents selected elements of an electrical DFM roadmap for $\leq 45\text{nm}$ process nodes.


4.1. Modeling the Electrical Impact of Variability

Electrical models of non-rectangular devices and interconnects have enjoyed recent interest as a means of assessing impact of lithographic and CMP errors on power and performance. Such models enable “process-aware analysis” or “model-based signoff”, which informs signoff analyses (RCX, delay calculation, STA) with results of physical simulations of systematic (“deterministic”) pattern-dependent variations. The work of [7] models non-rectangular device channels with comprehension of narrow-width effect and resulting variation of V_{th} across the gate width. The figure above illustrates definitions of gate width (W) and edge width (w), with edge regions shown in blue. These concepts are used in the figure below, which shows variation of V_{th} along the width of the device for different gate widths, when



³ Variability modeling, from easiest to hardest, spans (1) systematic WID (e.g., pattern-dependence of litho and CMP), (2) random DTD (“SSTA”), (3) random WID, (4) correlated random WID, and (5) systematic DTD. For example, nascent approaches to (4) still gloss over the question of how to model the fact that $\text{BUF} = \text{INV} + \text{INV}$ or $\text{AND} = \text{NAND} + \text{INV}$ [20].

⁴ Business frameworks for statistical design are still unclear. For example, it seems impractical for foundries to deliver the exact process statistics to which a design was optimized. Or, if the process evolves during the course of a given design project, optimizations targeted to early process statistics could turn out to be harmful in the matured process.



$W > 2w$ (left) and $W < 2w$ (right). Edge width is the width of the region near the boundary between poly and diffusion. Models analogous to those described in [7] are necessary to capture the electrical impacts of line-edge roughness (LER) and line-width roughness (LWR), which are likely significant contributors to inter-device variation in $\leq 45\text{nm}$ nodes. Particularly for analog and mixed-signal circuits, edge roughness can affect matching and delay requirements. While today’s design methodologies still model the effects of LER/LWR as random, future electrical DFM flows require more accurate, model-based accounting for (and bounds on) delay, capacitance and power variation with LER/LWR.

A critical extension for $\leq 45\text{nm}$ is the modeling of diffusion rounding. Poly CD is increasingly well-controlled in modern processes, in part due to layout restrictions. However, diffusion is still very irregular, resulting in imperfect printing. Although diffusion patterns have larger CD than poly patterns, corners and jogs are more prevalent in smaller technologies, and process windows are small due to significant corner rounding with defocus. Hence, poly gates placed in close proximity to diffusion edges are more likely to demonstrate larger performance variation than those away from the edges. Such diffusion patterning issues are likely to become more significant as the average gate width scales down with each technology generation.

Simple models that account for diffusion rounding by adjusting gate width (cf. [6],[7]) have unacceptable error in I_{off} predictions. Moreover, source-side diffusion rounding and drain-side diffusion rounding behave very differently from an electrical perspective, which strongly suggests that diffusion rounding modeling must be performed in a design context-aware manner. Future flows require new modeling techniques to determine equivalent L and W given both poly and diffusion patterning imperfections.

More generally, process-aware analysis flows for signoff at $\leq 45\text{nm}$ require industry consensus on “deconvolutions” to solve:

- in the FEOL, silicon-calibrated LPE (layout parasitic extraction) rule decks potentially double-count litho contour effects (LPC); and
- in the BEOL, silicon-calibrated RCX tools potentially double-count post-CMP wafer topography effects.

Another critical blocker for electrical DFM is industry consensus concerning treatment of signoff analysis corners in the presence of electrical model corrections for litho and CMP variation. For example, if a tool indicates that a device’s nominal L_{eff} should be changed from 40nm to 38nm due to pattern-specific litho variations, it is not clear today how to modulate the qualified BC/WC SPICE corners for the device. Related issues include (tractable) standardized silicon qualification of process-aware analysis, and enablement of full-chip signoff analyses in cell-based methodologies.

4.2. Opportunistic, “Do No Harm” DFM

In keeping with precepts of Section 2, electrical DFM should blend with existing chip implementation, manufacturing verification, and OPC/mask/wafer flows. We observe:

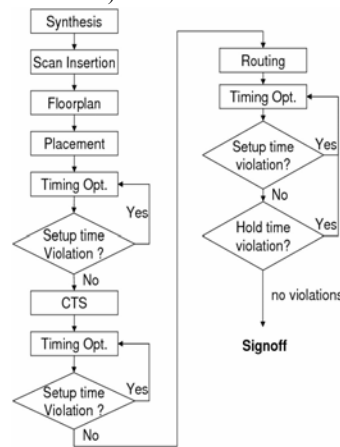
- “Prevention” in the sense of “correct by construction” can be onerous and hence difficult to adopt. For example, one-pitch, one-orientation poly layout is highly manufacturable, but is foreign to layout engineers and incurs unacceptable area penalties. Similarly, insertion of dummy devices or enforcement of phase-shift mask 2-colorability in standard-cell layouts (i.e., “composability by construction”) becomes increasingly costly as pitches decrease while stepper wavelength (and hence optical radius) remains constant.
- “Cure” in the sense of “construct by correction”, particularly at the post-layout handoff between design and manufacturing, can suffer from shape-centricity, loss of design information, and separation from implementation flows. Without understanding of electrical and performance constraints, timing slacks, etc. it is difficult to determine whether manufacturing non-idealities actually harm the design, or how to mitigate such non-idealities to maximize yield. Moreover, any loop back to ECO P&R and signoff is viewed as costly, since it occurs essentially at tapeout and has potentially disturbed the ‘golden’ state of the design.

With this in mind, a crucial mantra for $\leq 45\text{nm}$ electrical DFM is to opportunistically ‘do no harm’. Optimizations should reach up into the implementation flow (shown below) to introduce

corrections at appropriate times – e.g., the best way to correct litho hotspots on poly is after detailed placement and before routing (cf. the “Corr” methodology of [25]).


The remainder of this subsection reviews a recent example of opportunistic DFM, the *auxiliary pattern* (AP) methodology of [23]. The AP methodology is motivated by unacceptable scaling of model-based OPC (MBOPC), which has emerged as a major bottleneck for

turnaround time of IC data preparation and manufacturing. To address the OPC runtime issue, the cell-based OPC (COPC) approach has been studied by, e.g., [22] and [24]. The COPC approach runs OPC once per each cell definition (i.e., per cell *master*) rather than once per placement or unique instantiation of each cell (i.e., per cell *instance*). In other words, in the COPC approach, master cell layouts in the standard cell library are corrected before the placement step, and then placement and routing steps of IC design flow are completed with the corrected master cells; this achieves significant OPC runtime reduction over MBOPC, which is performed at the full-chip layout level for every design that uses the cells. Unfortunately, optical proximity effects (OPE) in lithography cause interaction between layout pattern geometries. Since the neighboring environment of a cell in a full-chip layout is different from the environment of an isolated cell, the COPC solution can be incorrect when instantiated in a full-chip



layout, and there can be significant CD discrepancy between COPC and MBOPC solutions.

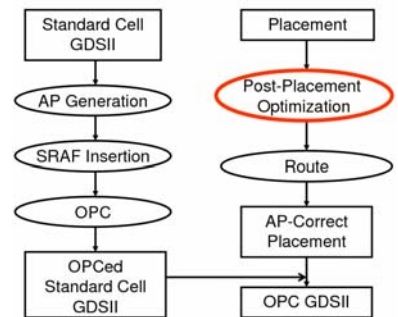
The AP technique of [23] shields poly patterns near the cell outline from the proximity effect of neighboring cells. Auxiliary patterns inserted at the cell boundary (e.g., as shown at right)



reduce discrepancy between isolated and layout-context OPC results for critical CDs (e.g., region “S” in the figure) of boundary poly features. This allows the substitution of an OPC’ed cell with APs directly into the layout⁵; then, COPC with AP can achieve the same printability as MBOPC, but with greatly reduced OPC runtime. Opportunism arises in two forms. (1) If the layout context of a standard-cell instance has room to substitute an AP version for the non-AP version, this should always be done, since it reduces OPC cost without affecting OPC quality. (2) The placement of cells in a given standard-cell block might not permit insertion of APs between certain neighboring cell instances. To maximize AP insertion in such cases, the detailed placement can be perturbed using an efficient, timing-aware dynamic programming algorithm to maximize possible substitutions of AP cell versions and hence the runtime benefits of COPC. The resulting flow is given below.

Apart from the runtime improvement, AP-based OPC benefits the process-aware signoff methodology discussed in Section 4.1 above. Full-chip litho simulation is implicit in such a methodology, since two instances of the same standard-cell master can print differently due to context-dependent OPC

and litho variations. Since an AP cell version has a pre-determined OPC solution and aerial image in litho simulation, the runtime of process-aware signoff can be substantially reduced without any loss of accuracy [26].




4.3. Layout Support for Sub-45nm Patterning


Electrical DFM at $\leq 45\text{nm}$ spans not only modeling of electrical-, design- and product-level impacts of variability, but also (1) methods for layout and manufacturing handoff for novel patterning approaches, as well as (2) quantification of design and cost tradeoffs inherent in various forms of layout regularity. This begins with an understanding of patterning options. An important option is Double Patterning Lithography (DPL), which involves partitioning of dense circuit patterns into two separate layers so that decreased pattern density can improve resolution and depth of focus (DOF). DPL is the likely mainstream technology for 32nm lithography [27], especially given that EUV appears likely be

⁵ As detailed in [23], APs consist of vertical (V-AP) and/or horizontal (H-AP) non-functional (dummy) poly lines. V-AP features are located within the same cell row and print on the wafer. H-AP features are located in the overlap region between cell rows; their width is comparable to that of sub-resolution assist features and hence they do not print on the wafer.

adopted late if at all. DPL can be performed in several different ways. However, most known double patterning techniques have relatively complex process flows, which may slow DPL use in production. A complicating factor is the use of two etch steps for the first and second exposures, as illustrated below [28]. The first etch step is necessary to transfer the pattern of the first resist layer into an underlying hard mask which is not removed during the second exposure. Photoresist is again coated on the surface of the first process for a second exposure. The second mask, having patterns separated from the first mask, is exposed – and then the flow finishes up with the hard mask and resist of second exposure.⁶



A key problem in DPL is the decomposition of layout for multiple exposure steps. This recalls strong (alternating-aperture) PSM coloring issues and automatic phase conflict detection and resolution methods (see, e.g., [4], which gave one of the earliest automated and optimal compaction-based phase conflict resolution techniques). With



DPL layout decomposition, two features must be assigned opposite colors if their spacing is less than the minimum color spacing. The figure at left shows a pattern in which features cannot all be assigned different colors. The easiest workaround is to split one

feature into two; other methods include H-V decomposition, pitch doubling, etc. However, two fundamental problems of DPL remain: (1) generation of excess line-ends, which cause yield loss due to overlay error in double-exposure, as well as line-end shortening under defocus, and (2) resulting requirements for tight overlay control, possibly beyond currently envisioned capabilities.⁷ A challenge for the industry is timely productization of techniques for layout perturbation and layout decomposition to minimize the number of created line-ends, and to introduce layout redundancy that reduces line-end shortening related patterning failures; quite possibly, such measures will be critical to the success of DPL. Lithographic hotspot finding and fixing with overlay error simulation is another potential enabler to adoption of DPL; cf. the graph-based hotspot finding approach [8], which provides 50-100X speedups of post-OPC lithography checks for both necking and bridging type hotspots.

4.4. Stress Modeling and Exploitation

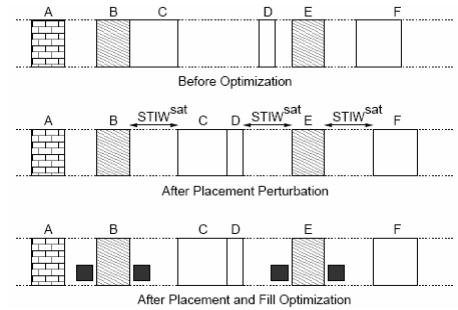
Starting at the 65nm node, stress engineering to improve performance of transistors has been a major industry focus.

⁶ SEMATECH has supported investigation of a materials solution that allows performing two exposures of the same material, i.e., double exposure and single etching.


⁷ Workarounds have been proposed, e.g., simply expanding line-ends in the mask to achieve larger overlap. Other issues include line edge errors due to overlay error, and interference mismatch between different masks.

However, even a very well-understood intrinsic stress source – shallow trench isolation (STI) – has not yet been fully exploited for circuit performance analysis and improvement. Given that process-based device scaling knobs have run out of steam, it is particularly important to model yet-unexploited stress effects, and to develop circuit analysis and optimization techniques that comprehend and *exploit* these effects. This subsection reviews one example work [29] which focuses on STI compressive stress along the device channel; such stress typically enhances PMOS mobility while degrading NMOS mobility. Stress due to STI at a device depends on the location of the device in the diffusion region and the width of the STI on both sides of the diffusion region. While present BSIM modeling accounts for STI stress due to device location in the diffusion region (using SA, SB, SC parameters), stress due to STI width is not accounted for. The work of [29] uses the Synopsys Sentaurus process simulator to simulate the STI process up to the gate deposition step, and applies a rigorous DOE across a range of STI widths, diffusion lengths, and gate to diffusion edge spacings. The analysis shows that delay of standard cells changes by 10-20% depending on STI width, which in turn depends on placement. This naturally recalls the concepts of opportunism and the use of placement to manage deterministic variations [23][25].

[29] enhances the performance of standard-cell blocks by changing the placement to modulate STI width, and by inserting active-layer dummy shapes. The goal of the placement and active-layer fill



optimization is to introduce additional spacing between timing critical cells to: (1) increase the STI width for PMOS devices and consequently improve PMOS speed, and (2) create space for fill insertion to be performed only next to NMOS diffusion so as to improve NMOS speed. Above is a standard-cell row before optimization, after placement perturbation, and after fill insertion. In the figure, $STIW^{sat}$ is the STI width beyond which stress effect saturates. Cells with diagonal line patterns are timing critical. “Don’t-touch” cells with brick pattern cannot move in the placement optimization. As reported in [29], significant reductions of up to 11% in SPICE-computed path delay can be achieved by the combined placement and active-layer fill optimization. The figure at right shows path delay histograms of the top 100 critical paths of a small testcase before and after the optimization.



Electrical DFM in ≤ 45 nm nodes requires additional methods to exploit STI width

impact. Analogous to the AP approach, it is possible to perform opportunistic, timing-driven instantiation of dummy diffusion features in cell layouts on the sides of, or around, diffusion regions

to completely mitigate layout-dependent stress effects. “Rich library” possibilities abound, e.g., cell layouts can be changed to permute devices and place them in the diffusion region to enhance cell performance over some or all timing arcs. Cell variants with different spacings of devices from the diffusion edges can afford finer-grain tradeoff of device speeds and leakage currents. Dual stress linear (DSL), a recently-introduced and very effective stress engineering technique, exhibits stronger layout dependencies; development of layout-dependent performance models for devices with DSL, along with analysis and optimization techniques that leverage such models, is a likely direction.

5. Conclusion

Many elements of the electrical DFM roadmap have been left undiscussed due to time and space constraints. A few more:

- *Cost- and design-driven DFM.* Design knowledge must be leveraged to reduce costs, particularly in mask and wafer flows (insertion points include OPC complexity [16], write optimization, as well as inspection and defect disposition).
- *Variability characterization.* Standard techniques – DOEs, simulation structures, and silicon TEGs – must afford rigorous foundations for model-based and statistical design. Prototypes might include [3] (CMP fill/extraction) and [9] (post-OPC litho contour prediction).
- *Overlay robustness.* With increased BEOL resistivities, and DPL in sight for critical layers, new design-manufacturing synergies must be developed around overlay and alignment. Possibilities include misalignment-tolerant layout styles, as well as design-driven alignment targets.
- *“Design for equipment”.* A wide range of equipment improvements (hooks to ‘smart inspection’, dynamic control of dose [21], various forms of adaptive process control, etc.) continually afford opportunities to leverage design information for cost and turnaround time improvements.

Certainly, there is no shortage of challenges and opportunities for the EDA, process, mask and design communities in $\leq 45\text{nm}$ nodes.

6. Acknowledgments

This talk draws on many joint works and discussions with students in ABKGroup at the UCSD VLSI CAD Laboratory, colleagues at Blaze DFM, Inc., and others. Thanks are especially due to Puneet Gupta, Puneet Sharma, Chul-Hong Park, Swamy Muddu, Rasit Topaloglu, Kambiz Samadi, Kwangok Jeong, Professor Dennis Sylvester, and Professor Hoong-Joo Lee.

References

[1] A.B. Kahng and R.O. Topaloglu, “Device-Level Layout Design Guidelines for Lower Leakage through Reduced Temperature”, in *submission to IEEE Trans. on VLSI*.

[2] A.B. Kahng and R.O. Topaloglu, “A TCAD-Based Study of Fill Pattern and Via Fill Impact on Low-k Dielectric Stress”, invited paper, *Proc. CMP-MIC*, March 2007.

[3] A.B. Kahng and R.O. Topaloglu, “A DOE Set for Normalization-Based Extraction of Fill Impact on Capacitances”, *Proc. ISQED*, 2007.

[4] A. B. Kahng, H. Wang and A. Zelikovsky, “Automated Layout and Phase Assignment Techniques for Dark Field Alternating PSM”, *Proc. BACUS Symposium on Photomask Technology and Management*, Sept. 1998, pp. 222-231.

[5] C. Chiang, A. B. Kahng, S. Sinha and X. Xu, “Fast and Efficient Phase Conflict Detection and Correction in Standard-Cell Layouts”, *Proc. ACM/IEEE Intl. Conf. on CAD*, Nov. 2005, pp. 149-156.

[6] P. Gupta, A. B. Kahng and P. Sharma, “Lithography Simulation-Based Full-chip Design Analyses”, *Proc. SPIE Conf. on Design and Process Integration for Microelectronic Manufacturing*, 6156, 61560T (2006).

[7] P. Gupta, A. B. Kahng, Y. Kim, S. Shah and D. Sylvester, “Modeling of Non-Uniform Device Geometries for Post-Lithography Circuit Analysis”, *Proc. SPIE Conference on Design and Process Integration for Microelectronic Manufacturing*, 6156, 61560U (2006).

[8] A. B. Kahng, C.-H. Park and X. Xu, “Fast Dual-Graph Based Hot-Spot Detection”, *Proc. 27th BACUS Symposium on Photomask Technology and Management*, SPIE vol. 6281, 2006, p. 628104.

[9] P. Gupta, A.B. Kahng, S. Muddu, S. Nakagawa, “Modeling Edge Placement Error Distribution in Standard Cell Library”, *Proc. Design and Process Integration for Microelectronic Manufacturing IV*, SPIE vol. 6156, 2006, pp. 61560S1-61560S12.

[10] P. Gupta, A. B. Kahng, P. Sharma and D. Sylvester, “Method for Correcting a Mask Layout”, *U.S. Patent No. 7,149,999*, Dec. 2006.

[11] P. Gupta, A. B. Kahng, Y. Kim, and D. Sylvester, “Self-Compensating Design for Focus Variation”, *Proc. Design Automation Conference*, June 2005, pp. 365-368.

[12] P. Gupta and F.-L. Heng, “Toward a Systematic-Variation Aware Timing Methodology”, *Proc. ACM/IEEE DAC*, 2004, pp. 321-326.

[13] <http://www.blaze-dfm.com/products/products.html>

[14] Y. Chen, P. Gupta and A.B. Kahng, “Performance-Impact Limited Area Fill Synthesis”, *Proc. ACM/IEEE DAC*, June 2003, pp. 22-27.

[15] P. Gupta, A. B. Kahng, C.-H. Park, K. Samadi and X. Xu, “Topography-Aware Optical Proximity Correction for Better DOF Margin and CD Control”, in *Proc. SPIE* 5853, 2005, pp. 844-854.

[16] P. Gupta, A.B. Kahng, D. Sylvester and J. Yang “A Cost-Driven Lithographic Correction Methodology Based on Off-the-Shelf Sizing Tools”, *Proc. ACM/IEEE DAC*, June 2003, pp. 16-21.

[17] G. Baccarani, M.R. Wordeman and R.H. Dennard, “Generalized Scaling Theory and its Application to a Micrometer MOSFET Design”, *IEEE Trans. Electron Devices* 31 (1984), pp. 452-462.

[18] K. Jeong, A. B. Kahng and K. Samadi, “Impact of Guardband Reduction on Design Process Outcomes”, *manuscript*, April 2007.

[19] S. M. Burns, M. Ketkar, N. Menezes, K. A. Bowman, J. W. Tschanz and V. De, “Comparative Analysis of Conventional and Statistical Design Techniques”, *Proc. ACM/IEEE DAC*, 2007, pp. 238-243.

[20] P. Gupta, *personal communication*, Jan. 2007.

[21] N. Jeewakhan, N. Shamma, S.-J. Choi, R. Alvarez, D. H. Son, M. Nakamura, V. Pici, J. Schreiber, W.-S. Tzeng, S. Ang and D. Park, “Application of Dosemapper for 65-nm Gate CD Control: Strategies and Results”, *Proc. SPIE* vol. 6349, 2006, pp. 6349G.

[22] P. Gupta, F.-L. Heng and M. Lavin, “Merits of Cellwise Model-Based OPC”, *Proc. SPIE Conference on Design and Process Integration for Microelectronic Manufacturing*, SPIE vol. 5379, 2004, pp.182-189.

[23] A.B. Kahng and C.-H. Park, “Auxiliary Pattern for Cell-Based OPC”, *Proc. 27th BACUS Symp. on Photomask Technology and Management*, SPIE vol. 6349, 2006, p. 63494S.

[24] X.Wang, M. Pilloff, H. Tang and C.Wu, “Exploiting Hierarchical Structure to Enhance Cell-Based RET With Localized OPC Reconfiguration”, *Proc. SPIE Conf. on Design and Process Integration for Microelectronic Manufacturing*, SPIE vol. 5756, 2005, pp. 361-367.

[25] P. Gupta, A. B. Kahng and C.-H. Park, “Detailed Placement for Improved Depth of Focus and CD Control”, *Proc. Asia and South Pacific Design Automation Conf.*, 2005, pp. 343-348.

[26] A. B. Kahng, S. Muddu and C.-H. Park, “Auxiliary Pattern-Based OPC for Better Printability, Timing and Leakage Control”, *manuscript*, 2007.

[27] <http://www.edn.com/blog/450000245/post/70007207.html?nid=3389>

[28] <http://imec.be/wwinter/mediacenter/en/SR2006/681406.html>

[29] A. B. Kahng, P. Sharma and R. Topaloglu, “Exploiting STI Stress for Performance”, to appear in *Proc. IEEE Intl. Conf. on CAD*, Nov. 2007.