Opportunities in Future Physical Implementation and Manufacturing Handoff Flows

Andrew B. Kahng

UC San Diego Depts. of CSE and ECE, La Jolla, CA 92093-0404 USA

abk@ucsd.edu , http://vlsicad.ucsd.edu/

Abstract. We discuss aspects of silicon quality and value left on the table by current physical implementation and manufacturing handoff flows. These aspects include the following. (1) Proper expectations with respect to guardbanding of process, statistical design, and gaps in nascent flows. (2) How manufacturing variability should be deal with by design flows, e.g., with approaches less dogmatic than traditional "correct by construction" (prevention) or "construct by correction" (cure). (3) Opportunities to differentiate with physical implementation 'glue' technologies in place-and-route. (4) New targets for 45nm and 32nm deployment such as stress/strain modeling, layout support for double-patterning lithography, and 'design for equipment' synergies.

1. Introduction

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 of product value. As reviewed in [1], a number of design technology gaps (electrical design for manufacturability[1], standards for modeling and characterization of process variability, etc.) must be addressed if the industry is to achieve a true "design for value" (maximizing profit per wafer) capability. The following discussion addresses (1) proper views of guardbanding, expectations from statistical design and reality checks for nascent flows; (2) 'opportunistic' mindsets for mitigation of manufacturing variability; (3) differentiating 'glue' technologies in placement and routing; and (4) directions at ≤ 45nm with respect to stress/strain modeling, double-patterning lithography, and design-equipment synergies.

2. On guardbands, statistics, and gaps

As physical implementation tools and manufacturing handoffs evolve for 45nm and 32nm production, the industry should adopt proper expectations regarding margin reductions and upsides of statistical design.

2.1. Time constants and the inevitability of guardbanding.

In the co-evolution of silicon technology and silicon products, basic 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

Also, a number of precedence and practical constraints apply, 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. This is a fundamental asymmetry between process and design. In this light, R&D goals for ≤ 45nm include:

- quantification of guardbanding costs, and benefits such as the potential tradeoff of guardband reduction and parametric yield loss for faster design closure and improved random defect yield [2];
- design robustness to variabilities (cf. self-compensation [3]), that include 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.[2]

2.2. Practicality and value of statistical design.

Despite a great deal of industry attention, deep challenges remain with respect to modeling, characterization and mitigation of manufacturing variations. Statistical analyses and optimizations have been rapidly conceived, but lack consensus on enablement for production flows. Silicon data at 65nm and below suggest that characterization of die-to-die variation and spatial correlations is essential to statistical design flows.[3] An even more basic issue is whether sufficient ROI can be shown for statistical design approaches; e.g., [4] showed limited impact of statistical power optimization. Intuitively, statistical design will have 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. When statistical optimization drives the design to essentially the same point as deterministic design, as appears to be the case for, e.g., timing-driven design when spatial correlations are considered, the potential differentiated benefits lie mainly in yield prediction.[4]

---

[1] "Electrical DFM" 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. Optimizations are driven by analysis engines that comprehend a full spectrum of physical and electrical implications of manufacturing. The degrees of freedom to achieve optimization goals include changes to layout, to target dimensions of features, and to reticle enhancement (e.g., OPC) treatments at multiple levels of granularity spanning netlist optimization and placement to individual diffusion regions.

[2] Moore's Law value scaling is roughly 1% per week. Even if margin in the form of guardbanded process models is inevitable, big wins are possible from, e.g., design-directed yield learning.

[3] Variability modeling, from easiest to hardest, spans (1) systematic WID (e.g., pattern-dependence of litho and CMP, process simulation, and process-aware analyses), (2) random DTD (statistical STA), (3) random WID, (4) correlated random WID, and (5) systematic DTD. Because process learning and design rules also address (1), and because effects of other variation types are substantial, the focus to date on (1) and (2) may in hindsight constitute 'looking under the lamp post'.

[4] Business frameworks for statistical design remain 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 end up being harmful in the matured process.

2.3. Gaps in nascent flows.

Electrical models (equivalent W', L') 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 presumably lie on the road toward "process-aware analysis" or "model-based signoff", which will inform signoff analyses (RCX, delay calculation, STA) with results of physical simulations of systematic ("deterministic") pattern-dependent variations. For example, the work of [5] models non-rectangular device channels with comprehension of narrow-width effect and resulting variation of V_{th} across the gate width. Future advances are needed with respect to overlay/misalignment, and at least two key areas.

- Effects of line-edge and line-width roughness (LER, LWR), which significantly affect inter-device variation and matching in ≤ 45nm nodes. Future electrical DFM support must provide accurate, model-based analysis of delay, capacitance and power variation due to LER/LWR – including awareness of mask data preparation and mask write contributions.
- Effects of diffusion rounding. Poly CD is increasingly well-controlled in modern processes, in part due to layout restrictions. By contrast, diffusion layouts are irregular, with many corners and jogs and small process windows due to corner rounding with defocus. Poly gates that are closer to diffusion edges will have larger performance variation. Simplistic adjustments to gate width [5][6] are inadequate, and since source-side and drain-side diffusion rounding behave differently from an electrical perspective, 'design-aware' modeling is likely required.

More generally, process-aware analysis flows for signoff at ≤ 45nm 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.

A final blocker in the electrical DFM roadmap is industry consensus on treatment of signoff analysis corners in the presence of process-aware electrical model corrections. For example, if a simulator indicates that a device's nominal L_{eff} should be changed from 40nm to 38nm due to pattern-specific litho variations, it is unclear today how to change 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.

3. Opportunistic mindsets

With industry demand for new design for manufacturability capabilities, the physical verification platform has taken over some functionality (e.g., via doubling, wire spreading and fattening, dummy fill insertion, etc.) from the upstream physical implementation (SP&R) platform. This has led to two distinct mindsets today.

Upstream "prevention", in the sense of "correct by construction", focuses on design rules, library design and manufacturing-awareness in the SP&R flow. This is costly in terms of buffer dimensions, modeling margins, and other guardbands; examples include guaranteed standard-cell composability for alternating-aperture PSM, and dummy gate poly at boundaries of standard cells. In some cases, the cost increases with scaling, e.g., as the number of pitches in the stepper wavelength increases, the area penalty of buffer dimensions grows rapidly. Solutions can also be too onerous for adoption, e.g., one-pitch-one-orientation poly layout is highly manufacturable but incurs unacceptable area penalties. Finally, "prevention" may mean attempting to solve problems with too little information – e.g., trying to solve litho hotspots that have timing impact during P&R, before golden wire parasitics and signal integrity reports are in hand.

Downstream "cure", in the sense of "construct by correction", is often performed at the post-layout handoff between design and manufacturing. This can suffer from shape-centricity and loss of design information, as well as separation from implementation flows. Without a grasp of electrical and performance constraints, timing slacks, slew criticality, etc., such flows cannot easily determine whether manufacturing non-idealities actually harm the design, or how to mitigate such non-idealities to maximize parametric yield. Moreover, any loop back to ECO P&R and re-signoff is costly since it has disturbed the 'golden' state of the design (with no guarantees of convergence) and directly impacts tapeout schedule.

Where "prevention" can address manufacturability too early for best results, "cure" often comes too late in the flow. It is possible to simply bolt manufacturing verification and SP&R tools together, but this is not optimal. Rather, optimizations should reach up into the implementation flow to introduce corrections at appropriate times. The mantra for ≤ 45nm might be: (i) opportunistically make changes that can only help (i.e., "do no harm"), and (ii) make changes at appropriate junctures – when enough information is available, and before doing work that will be thrown away. The remainder of this section gives two examples of such opportunism.

3.1. The "CORR" methodology.

With respect to the above 'mantra', the appropriate juncture for correction of poly-layer litho hotspots is after detailed placement, but before routing. This is the fundamental insight of the "CORR" methodology [7], which improves lithographic process window by removing forbidden pitches for sub-resolution assist feature and etch dummy insertion. This optimization is achieved by a dynamic programming algorithm that 'jiggles' an existing detailed placement to exploit whitespace while curing any forbidden or weak pitches. The 'AFCORR' and 'EtchCORR' variants are timing- and wirelength-preserving in that they can constrain allowed movement of each cell instance according to timing criticality.

In [8], pitch-specific CD variations through OPC and litho process window are mapped to systematic leakage variations of individual devices. The figure below illustrates how detailed placement choices (cell ordering, site choice, mirroring) can affect pitches of boundary gates. 'LeakageCORR' then optimizes leakage in a timing- and wirelength-preserving manner. The "CORR" concept extends to a variety of objectives at the design-manufacturing interface. Examples include phase-shift conflict resolution, litho hotspot removal, timing and leakage improvement, improvement of recommended rule coverage, proportion of cell-based OPC usage (discussed next), and stress exploitation (Section 5.2).

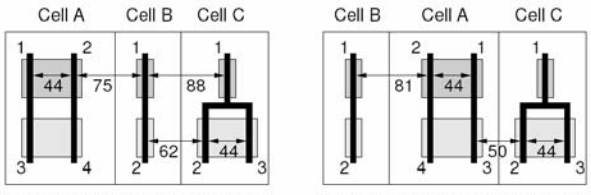

Pitch of Device 2 of Cell A: 44, 75
Pitch of Device 2 of Cell C: 62, 44
Pitch of Device 1 of Cell B: 75, 88

Pitch of Device 2 of Cell A: 81, 44
Pitch of Device 2 of Cell C: 50, 44
Pitch of Device 1 of Cell B: –, 81

3.2. Auxiliary pattern for cell-based OPC

The *auxiliary pattern* (AP) methodology of [9] is motivated by unacceptable scaling of model-based OPC (MBOPC), which is 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., [10] and [11]. COPC runs OPC once per each cell definition (i.e., per cell *master*)

rather than once per unique instantiation of each cell (i.e., per cell *instance*). Thus, 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 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 [9] *opportunistically* 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 of boundary poly features. This allows the substitution of an OPC'd cell with APs directly into the layout[5]; COPC with AP then achieves the same printability as MBOPC, but with greatly reduced OPC runtime. Here, 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. Otherwise, if there is no room for AP insertion, we leave the layout as is, and are no worse off than before. (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 'AP-CORR'. In other words, we create our own opportunities: an efficient, timing-aware dynamic programming code can maximize possible substitutions of AP cell versions and hence the runtime benefits of COPC. The resulting flow including AP-CORR is shown at right.

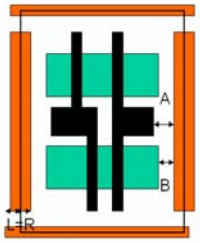

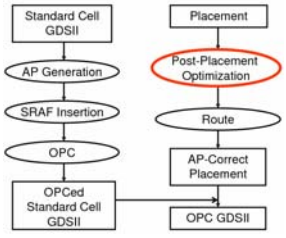

Apart from runtime improvement, AP-based OPC benefits process-aware signoff (recall Section 2.3 above). Full-chip litho simulation is implicit in such a methodology, since two instances of the same standard-cell master can print differently due to context-dependent OPC and litho variations. Since an AP 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 accuracy loss [12].

4. Glue technologies in place-and-route

As defined by G. Smith [13], a 'power user':

- has a CAD group that develops internal design tools
- is designing ICs on the latest two process nodes
- updates the design tool flow at the end of each project
- uses tools from EDA startups
- spends at least 33% more on CAD than a 'mainstream user'
- has designer productivity in line with ITRS Design Cost chart [14].

[5] As detailed in [9], 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.

By contrast, an 'upper mainstream user':

- has a CAD team but avoids internal tool development when possible
- designs ICs on processes that have been in production for ≥ 2 years
- tends to standardize its tool flow around 1+ major EDA vendors' tools
- has little or no use of startup tools
- spends less than 6% of its R&D budget on CAD
- has designer productivity that lags the ITRS Design Cost chart by ≥3 years.

According to Smith [13], the key to cost control is design productivity. A 'power user' designs the same IC as the 'upper mainstream user', but at much lower cost. Thus, the power user either prices its competition out of the market, or designs a chip with more features in the same amount of time. The catch is that EDA vendors are not supplying the tools necessary for today's designs. Hence, a large percentage of upper mainstream design groups are being forced to develop tools internally, as shown in the chart below (courtesy of GarySmithEDA; note that not all internal tool users satisfy the 'power user' criteria). The resulting trajectory choice [13]: an upper mainstream user must evolve into a power user, or else "end up doing RTL handoff". With this preamble, future differentiated, 'private-label' or internal physical design capabilities take on added significance.

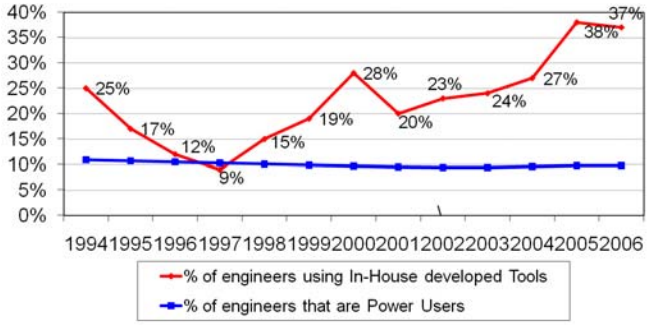

4.1. Placement opportunities.

Any near-term differentiating capability will address a subset of the following.

- Clock-to-data spatial management, and more generally, exploitation of correlated variation.
- Timing path monotonicity. There are many indications that at least one full process node of timing is left on the table by today's SP&R and optimization flows.
- Demands from 2.5-D, 3D integration. TSVs, macro placement, thermal/CTE-driven, …
- Pitch and litho hotspot management. E.g., "*CORR" methodologies; leverage fast hotspot filters
- Stress mitigation and exploitation
- Improved ECO placement. Fix-all knob: timing, SI, power, density, etc. – and closely tied to ECO routing

With respect to timing path monotonicity, the figure at right traces a typical critical path in a 90nm implementation of a JPEG encoder. Delay on the critical path is 2.796ns; delay of the path without interconnect would be 2.084ns, which is a difference of 712ps, or ~25%. The following plot shows path delays with and without interconnect for the 5000 most critical paths in the same

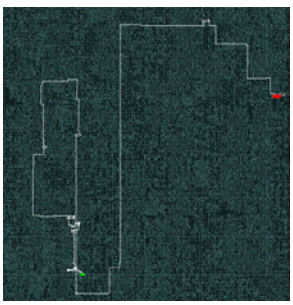

block implementation. Data such as this, along with "ASIC vs. Custom" studies by Chinnery, Keutzer and others, suggest that considerable performance – at least a full technology node – is being left on the table by today's chip implementation flows.

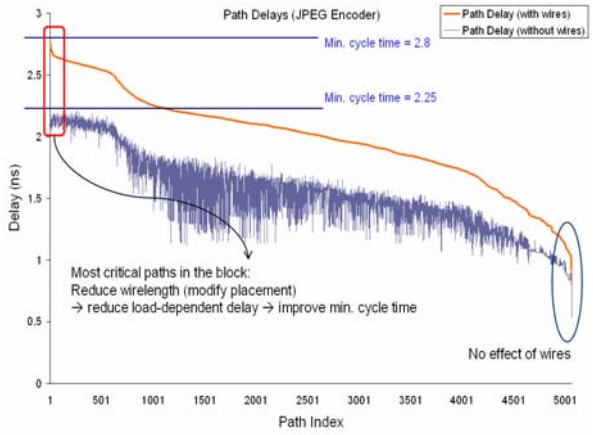

4.2. Routing opportunities.

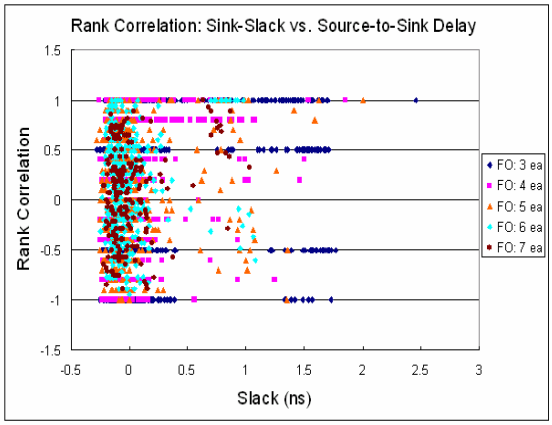

In a 2003 invited paper on the "future of routers" [15], a 'top-10' list of objectives for the industry began with the following items.

- *(0) Sensible unifications to co-optimize global signaling, manufacturability enhancement, and clock/test/power distribution*
- *(1) Fundamental new combinatorial optimization technologies (and possibly geometry engines) for future constraint-dominated layout regimes*
- *(2) New decomposition schemes for physical design*
- *(3) Global routing that is truly path-timing aware, truly combinatorial, and able to invoke "atomistic" interconnect synthesis*
- *(4) In-context layout synthesis that maximizes process window while meeting electrical (functional) spec*

Arguably, these are still at the top of today's priority list (see also [16]). To highlight the significance of item (3): a DAC-2006 work [17] demonstrates that a simple ECO routing flow can cure 'insane' topologies of timing-critical nets so as to improve clock frequency by approximately 5%. Here, an 'insane' topology is one where the rank order of sinks in terms of their slacks or required arrival times is not well-correlated with their rank order in terms of source-to-sink delay. The plot above shows rank correlation of sink slacks vs. source-to-sink delays for critical nets in a purportedly timing-optimized block. It is evident that the correlations are not uniformly positive – let alone equal to 1 – as we would hope.[6]

Opportunities for differentiated routing capability also include the following.

- Support for double-patterning lithography (see below), ranging from layout decomposition to auto-fixing of coloring conflicts, etc.
- Support for performance-, variability- and reliability-driven interconnect redundancy, including 100% redundant vias
- Overhaul of global and detailed routers for restricted pitches, simultaneous performance-aware routing and fill, support for rich library design, and ownership of swapping and final placement.
- Overhaul of clock distribution to support non-tree network synthesis and adaptive link insertion for variation-robustness
- Cleanup of 'old problems', such as 'combinatorial' timing- and SI-driven routing; assignment and ownership of pins, buffering resources, and wiring planes in hierarchical design; and early routability analyses up to floorplanning.

5. Emerging directions at ≤ 45nm

5.1. Double-patterning lithography.

Double-patterning lithography (DPL) involves partitioning dense circuit patterns into two separate masking steps, so that decreased pattern density can improve resolution and depth of focus (DOF). DPL is a candidate mainstream technology for 32nm lithography [18]. A key problem in DPL is the decomposition of layout for multiple exposure steps. This recalls strong (alternating-aperture) PSM coloring formulations, along with automatic phase conflict detection and resolution methods (see, e.g., [19], 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, with a solution being to split one feature into two. Fundamental issues for DPL are: (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. The EDA industry must rapidly bring to market tools for layout perturbation and layout decomposition to minimize the number of created line-ends, and for introduction of layout redundancy that reduces functional failures due to line-end shortening. Lithographic hotspot finding and fixing with overlay error simulation is another enabler of DPL.

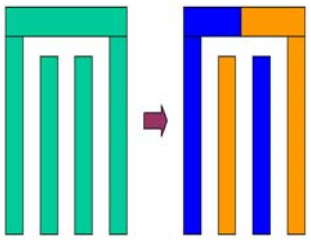

5.2. Stress modeling and exploitation

Engineering of stress and strain is the key means of achieving mobility enhancement, starting with the 65nm node. Systematic, layout-dependent impacts of stress must be modeled and exploited wherever possible to optimize the performance-power envelope of the design. As an example, [20] analyzes and exploits STI (shallow-trench isolation) compressive stress along the device channel, which typically enhances PMOS mobility while degrading NMOS mobility. STI-induced stress on a given device depends on

[6] In the flow of [17], 'virtual pin' and subnet constructs were used to force the ECO routing to deliver specific timing-driven Steiner topologies.

the device location in the diffusion region, and the width of the STI on both sides of the diffusion region. The BSIM stress model [21] accounts for STI stress as a function of device location in the diffusion region (cf. SA, SB, LOD parameters), but not as a function of STI width. On the other hand, TCAD-based studies reported in [20] show that STI width modeling can affect critical-path timing by up to 6%. Since STI width is determined by cell placement, this naturally recalls the concepts of opportunism and the use of "CORR" placement to manage deterministic variations.

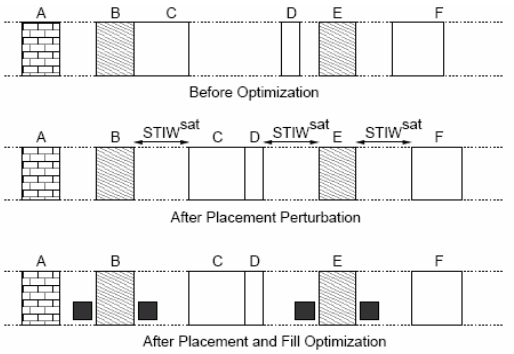

Indeed, [20] enhances the performance of standard-cell blocks by using detailed placement to modulate STI width, and by inserting active-layer dummy shapes. Additional spacing between timing-critical cells (1) increases STI width and hence speed for PMOS devices, and (2) creates space for insertion of active-layer fill next to NMOS diffusion so as to improve NMOS speed. The figure above shows 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 lines patterns are timing critical. "Don't-touch" cells with brick pattern cannot move in the placement optimization. As reported in [20], over 5% in (STI width-aware) SPICE-computed path delay can be achieved by the combined placement and active-layer fill optimization. The figure below shows delay histograms of the 100 most critical paths in a small testcase, before and after the optimization.

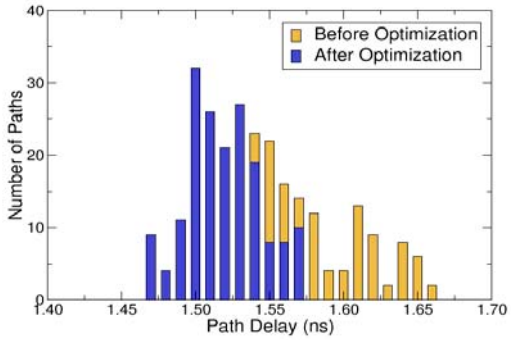

5.3. Design for equipment.

A wide range of equipment improvements (hooks to 'smart inspection', dynamic control of dose [22], various forms of adaptive process control, etc.) continually afford opportunities to leverage design information for cost and turnaround time improvements. For example, in a regime of rising mask costs, design-aware modulation of mask complexity [23], or design-aware inspection and defect disposition, could respectively reduce mask write times or increase tolerances for functionally insignificant features. As another example, ASML's Dose Mapper technology has been extensively used within the automatic process control context to improve global CD uniformity. It is a logical next step to explore whether Dose Mapper can be used (either for a fixed tapeout or in synergy with floorplanning and place-and-route) to improve design parametric yield. Finally, with equipment overlay and misalignment ever more significant within overall CD error budgets, the industry may require new misalignment-tolerant layout styles, as well as design-driven alignment targets during manufacturing.

6. Acknowledgments

A great deal of content is shared with [1], and/or represents work done in ABKGroup at the UCSD VLSI CAD Laboratory. Many thanks to Puneet Gupta, Chul-Hong Park, Kwangok Jeong, Puneet Sharma, Swamy Muddu, Rasit Topaloglu and Kambiz Samadi.

References

[1] A. B. Kahng, "Key Directions and a Roadmap for Electrical Design for Manufacturability", *Proc. ESSCIRC/ESSDERC*, Sept. 2007.
[2] K. Jeong, A. B. Kahng and K. Samadi, "Impact of Guardband Reduction on Design Process Outcomes", *manuscript,* April 2007.
[3] P. Gupta, A. B. Kahng, Y. Kim, and D. Sylvester, "Self-Compensating Design for Focus Variation", *Proc. ACM/IEEE Design Automation Conference*, 2005, pp. 365-368.
[4] 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.
[5] 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).
[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 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.
[8] A. B. Kahng, S. Muddu and P. Sharma, "Detailed Placement for Leakage Reduction Using Systematic Through-Pitch Variation", *Proc. International Symposium on Low Power Electronics and Design*, 2007.
[9] 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.
[10] 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.
[11] 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.
[12] A. B. Kahng, S. Muddu and C.-H. Park, "Auxiliary Pattern-Based OPC for Better Printability, Timing and Leakage Control", to appear in *IEEE Trans. on Semiconductor Manufacturing*.
[13] G. Smith, *personal communication*, August 2007. See http://garysmitheda.com/
[14] *International Technology Roadmap for Semiconductors*, http://public.itrs.net/
[15] A. B. Kahng, "Research Directions for Coevolution of Rules and Routers", *Proc. ACM/IEEE Intl. Symp. on Physical Design*, 2003, pp. 122-125.
[16] R. Brashears and A. B. Kahng, "Advanced Routing for Deep Submicron Technologies", *Computer Design*, May 1997.
[17] C. J. Alpert, A. B. Kahng, C. N. Sze and Q. Wang, "Timing-Driven Steiner Trees are (Practically) Free", *Proc. ACM/IEEE Design Automation Conference*, 2006, pp. 389-392.
[18] http://www.edn.com/blog/450000245/post/70007207.html?nid=3389
[19] 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.
[20] A. B. Kahng, P. Sharma and R. Topaloglu, "Exploiting STI Stress for Performance", to appear in *Proc. IEEE Intl. Conf. on CAD,* Nov. 2007.
[21] http://www-device.eecs.berkeley.edu/~bsim3/
[22] 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.
[23] 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.