Subwavelength Lithography and its Potential Impact on Design and EDA

Andrew B. Kahng and Y. C. Pati[†]

UCLA Department of Computer Science, Los Angeles, CA 90095-1596 USA
[†]Numerical Technologies, Inc., Santa Clara, CA 95051 USA

Abstract

This tutorial paper surveys the potential implications of subwavelength optical lithography for new tools and flows in the interface between layout design and manufacturability. We review control of optical process effects by optical proximity correction (OPC) and phase-shifting masks (PSM), then focus on the implications of OPC and PSM for layout synthesis and verification methodologies. Our discussion addresses the necessary changes in the design-to-manufacturing flow, including infrastructure development in the mask and process communities, evolution of design methodology, and opportunities for research and development in the physical layout and verification areas of EDA.

1 Introduction

With the advance of CMOS technology according to the International Technology Roadmap for Semiconductors [30], manufacturing cost increasingly drives design [19]. It is critical for process engineers to achieve predictability and uniformity of manufactured device and interconnect attributes, e.g., dopant concentrations, channel lengths, interconnect dimensions, contact shapes and parasitics, and interlayer dielectric thicknesses. To achieve a design solution at a reasonable point on the price-performance curve, a total *variability budget* for the design must be distributed among such attributes. In very deep submicron technologies, attaining large process windows and uniform manufacturing while bounding variability is difficult [6] [27] [19] [3]. Hence, the manufacturing process has an increasingly constraining effect on physical layout design and verification. Many physical design and physical verification methods have been proposed to address such manufacturing issues as registration errors, photolithographic random effects, random spot defects, plasma and charging effects ("antenna effect"), etc.; see such works as [18] [3] for reviews.

The heightened interdependencies between design and manufacturing are due in part to a fundamental *crossover point* in the evolution of VLSI technology. This crossover point occurs when minimum feature dimensions and spacings decrease below the wavelength of the light source. Pattern fidelity deteriorates markedly in this *subwavelength lithography* regime, leading to the use of compensation mechanisms [11] that either perturb the shape (via *optical proximity correction* (OPC)) or the phase (via *phase-shifting masks* (PSM)) of transmitting apertures in the reticle. As can be seen in Figure 1, at least the next several process generations are likely to rely on subwavelength lithography.[1]

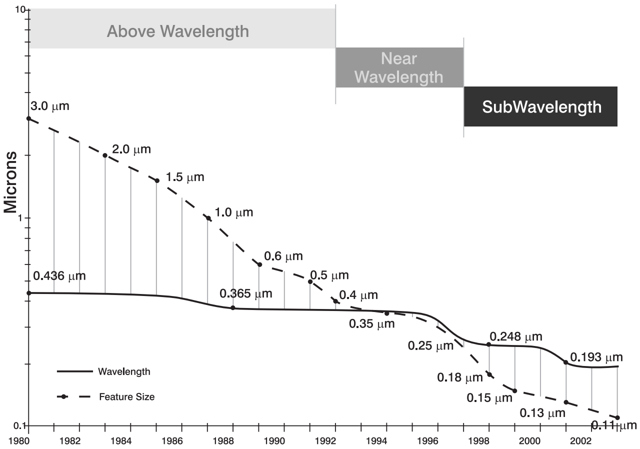

Figure 1: Shift to subwavelength optical lithography since the 0.35-micron process generation.

Available compensation mechanisms for subwavelength optical lithography have an unfortunate effect: the layout geometries being optimized in a polygon layout tool (e.g., place-and-route or custom layout) are no longer consistent with the actual mask geometries, and these in turn are no longer consistent with actual geometries on fabricated silicon. Line end shortening, corner rounding, and local context-dependent linewidth variations are all fundamentally consequences of subwavelength lithography. With the disappearance of the "WYSIWYG" regime (Figure 2), new challenges for verification, and new constraints on layout design, must be recognized and addressed.

In this paper, we assess the prospects for new tools and flows in the interface between layout design and manufacturability, focusing on layout design and verification for OPC and PSM. We will highlight necessary changes in the design-to-manufacturing flow, including infrastructure development in the mask and process communities, and opportunities for research and development in physical layout and verification.

2 Optical Proximity Correction

Optical proximity correction perturbs the shapes of transmitting apertures in the mask to address optical lithography distortions. The technique dates back to the early 1970s; see [13, 28, 2, 20] for reviews. The goal of OPC is to produce smaller features in an IC using a given equipment set, by enhancing lithographic resolution.

[1] A second crossover point in process evolution occurs when interconnect delays dominate device switching delays in deep-submicron CMOS technology [30]. Interconnect process optimization must achieve more delicate balances, e.g., affording simultaneous distribution of signal, clock and power with adequate performance and reliability while minimizing die area. Also, more interconnect layers are required at each successive node in the technology roadmap [30, 8, 33], leading to a strong requirement for, e.g., *planarized* interconnect processes that rely on chemical-mechanical polishing (CMP). Manufacturing steps involving CMP have varying effects on device and interconnect features, depending on local characteristics of the layout. This link between layout and manufacturability has grown in importance with the move to very deep-submicron (especially shallow-trench isolation and inlaid-metal) processes [9] [31] [32] [1].

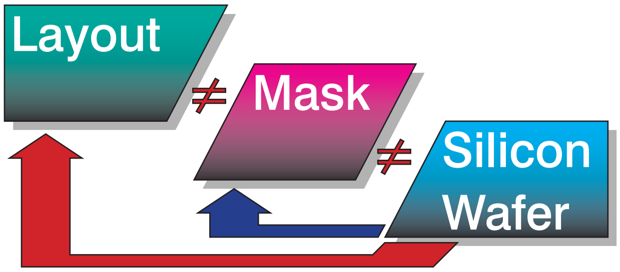

Figure 2: The disappearance of the "WYSIWYG" regime: polygon layouts no longer reflect actual mask geometries, which in turn no longer reflect actual fabricated silicon geometries.

OPC is based on systematic corrections that compensate for the nonlinear feature distortions arising from optical diffraction and resist process effects; typically, these corrections are made according to a predetermined rule set ("rule-based OPC") or else according to the results of lithography simulations that are iterated within the correction algorithm ("model-based OPC"). The OPC corrections themselves can be of several forms, including (i) serifs and hammerheads to eliminate corner rounding and line-end shortening; (ii) notches to control linewidth in the face of iso-dense effects; and (iii) subresolution assist features ("outriggers", or "scattering bars" [2]) for narrow gate geometries.[2] Figure 3 conveys the flavor of a layout after OPC has been applied. We observe that determining the optimal type, location, size and (a)symmetry of the corrections is highly complex and context-dependent. Furthermore, after OPC the "number of features" is no longer correlated to the "number of connected shapes", and the complexity of the geometry description is greatly increased.

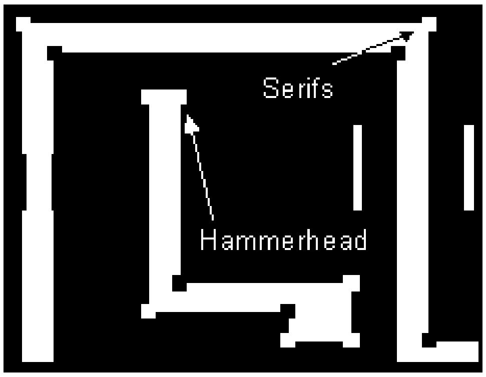

Figure 3: Example of aggressive OPC, showing serif, hammerhead and "outrigger" (subresolution gate assist) features.

OPC is very much a fact of life in deep-submicron (subwavelength) lithography, both today and into the future. OPC is also somewhat more mature than phase-assignment in terms of available software solutions. At the same time, OPC is properly viewed more as a "corrector" than an "enabler". In particular, the ability of OPC to extend the life of optical lithography equipment is limited: (i) the technique only compensates for distortions of features that can be printed (and if a feature does not print, as typically results when one attempts to manufacture gates below 180nm using 248nm equipment, it cannot be corrected), and (ii) an additional challenge to use of OPC below 180nm lies in actually making the mask (e.g., with current mask manufacturing equipment, tolerances are not sufficient to reliably create sub-180nm OPC masks). Therefore, OPC-related developments (in contrast to PSM-related developments) are aimed at disconnects within an existing and comparatively well-understood infrastructure. The remainder of this section highlights three of the most prominent disconnects: (i) application of OPC in hierarchical and reuse-centric methodologies, (ii) application of OPC without regard to functional requirements implicit in a feature, and (iii) application of OPC without regard to verifiability, e.g., at the mask writing step.

2.1 OPC for Hierarchical and Reuse-Centric Methodologies

In hierarchical (e.g., cell-based) design methodologies, the layout context for any given instance is not known *a priori*. If OPC introduces corrections (e.g., subresolution assist features) that are outside the original cell layout, instances may not be freely composable and the key assumption of the hierarchical methodology becomes invalid. Furthermore, if the OPC applied to a given feature depends on characteristics of features in a local neighborhood that possibly extends outside the cell boundary, the notion of a single "master" that can be instantiated in arbibrary contexts again becomes invalid. Within a reuse-centric methodology, the key concern is that the layout design must be amenable to OPC insertion in a variety of target processes.

2.2 Integration of Functional Knowledge

A standard measure of the cost of optical proximity correction is *data volume*, i.e., the number of edges in the corrected layout (versus the number of edges in the original layout). Data volume impacts the transmission and manipulation of correct layout data, as well as the time to write a mask and the ease of verifying the mask. We note that the true purpose of OPC insertion is not to make the manufactured structure "look like" the on-screen geometry in the layout editor. Rather, the purpose of OPC is to preserve a *functional correspondence* between the designed circuit and the manufactured circuit. The complexity of the inserted OPC should be as small as possible, consistent with this purpose. We identify three major avenues of tool development.

- The first type of new tool that must be developed will integrate "silicon-level" modeling of the fabricated geometry, for purposes of *analysis and verification* of function. In other words, silicon-level layout parasitic extraction, layout-vs.-schematic verification, and design-rule correctness verification will be added into current design flows. An example output of such a tool is shown in Figure 4, where the layout designer may be presented with an image showing the differences (subtractions and additions) of fabricated versus laid-out geometries.
- To compensate for the extra burden that such an added level of detail imposes, a second avenue of tool development will seek means of passing *functional intent* down to OPC insertion. The goal should be for OPC insertion to make only those proximity corrections that actually reduce the cost of the design – i.e., in the sense of reducing performance variation and the amount of guardbanding needed. The concept is similar to that of "filtering" in parasitic extraction for performance verification (PV) flows. For example, critical dimension (CD) control of an individual gate or individual wire jog may not be important if the gate or wire is not in any timing-critical path. On the other hand, CD control of devices and interconnects in timing-critical paths is extremely important. Such methods for passing functional intent will be applicable in any

[2] A *serif* is a small L-shaped geometry added to (subtracted from) a convex (concave) corner to compensate for rounding; a *hammerhead* is a U or inverted-U geometry that compensates for line-end shortening; a *notch* is a local thinning of a feature to compensate for linewidth variation; and an *outrigger* is a disconnected, non-printing geometry that uses diffraction effects to enhance linewidth control.

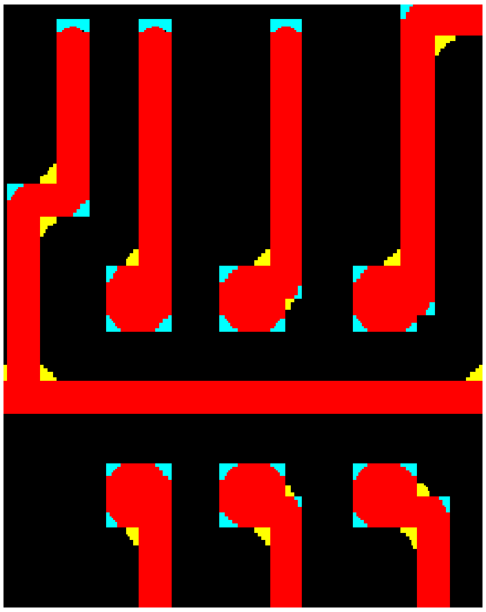

Figure 4: Output from "silicon-level" process modeling, showing the difference (subtractions, additions) between fabricated and laid-out geometries.

design flow, including today's flows where OPC is a layout post-processing step that is performed in physical verification or in mask processing.

- A third avenue for development is related to the question of how layout should best model the cost of the OPC insertion process. For example, it is not yet understood how a given geometric configuration affects the cost of the OPC needed to reliably yield a given functionality. Further study is also needed to understand how breaking hierarchy in the layout (or, in the OPC insertion) can affect data volume and verification costs at other stages of the design process.

2.3 Integration of Mask Verifiability

With the long write times of complex masks, the cost of discarding a faulty mask (or, repairing the mask) can be substantial. Furthermore, highly contorted shapes on the mask can be difficult to inspect and verify (the inspection process itself is subject to optical distortions, increased runtime due to mask complexity, etc.). Hence, it is imperative that we investigate and understand the relationship between the type of OPC applied (e.g., serif, notch, subresolution scattering bar) and the verifiability and repairability of the mask. An initial goal should be to develop new abstractions of the limits of mask verification. These would guide OPC insertion: no correction should be made that cannot be manufactured or verified. Eventually, tools must abstract mask verification up to the layout design and performance optimization stages: performance-driven layout design should not create situations where very aggressive, difficult-to-verify OPC is required to save the functionality of the circuit.

3 Phase-Shifting Masks

Phase-shifting mask (PSM) technology enables the clear regions of a mask to transmit light with prescribed phase shift. Figure 5 illustrates a layout with two closely spaced featuers: the conventional "binary" mask with no phase-shifting, but in the phase-shifting mask the two adjacent clear regions have respective phase shifts of 0 and 180 degrees. In the phase-shifting mask, light diffracted into the nominally dark region between the clear regions will interfere destructively; the improved image contrast leads to better resolution and depth of focus. All PSM variants employ this basic concept, which was proposed by Levenson et al. [12] in 1982. See [6] [27] [15] [29] [14] [17] for reviews of PSM technologies.

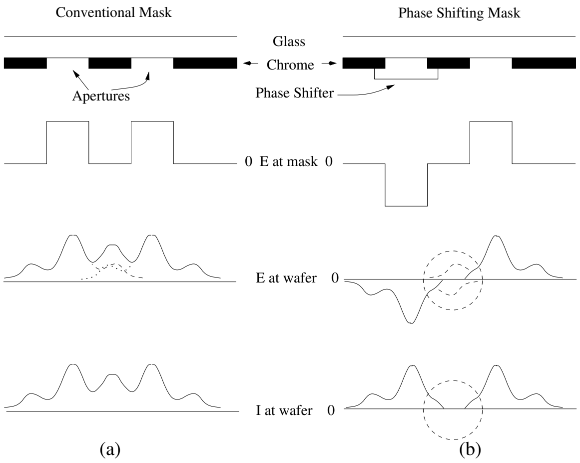

Figure 5: Comparison of diffraction optics of conventional and phase-shifting masks. E denotes electric field and I denotes intensity. With the conventional mask (a) light diffracted by two adjacent apertures constructively interferes, increasing the light intensity in the dark area of the wafer between the apertures. With the (alternating) phase-shifting mask (b), the phase shifter reverses the sign of the electric field, and destructive interference minimizes light intensity at the wafer in the dark area between apertures.

Two positive constants $b < B$ define a simplified relationship between printability and the distance between two clear regions [24]. The distance between any two features cannot be smaller than b without violating the minimum spacing design rule. If the distance between two features is at least b but smaller than B, the features are in *phase conflict*. Phase conflict can be resolved by assigning opposite phases to the conflicting features.

The Phase Assignment Problem: Assign phases to all features of a given layout such that no two conflicting features are assigned the same phase.[3]

Phase conflict in PSM layout design can occur in distinct contexts. Two classes of photolithographic masks are used to transfer circuit patterns onto silicon: bright field masks, and dark field masks. On a bright field mask the background (substrate) is transparent and the pattern is defined in chrome. Hence, the image projected onto silicon by a bright field mask defines circuit patterns through the formation of dark (unexposed) regions of the photoresist, a photosensitive material with which the silicon wafer is coated prior to exposure. Photoresists also come in two flavors: positive and negative. With positive photoresists, the development process following exposure removes photoresist material from all (exposed) regions that have been exposed with sufficient energy. For negative resists, the development process removes photoresist material from all unexposed areas. Today, positive photoresist is

[3] The Phase Assignment Problem is often stated in the context of the *phase conflict graph*, which is constructed by defining a vertex for each feature and introducing an edge between two vertices exactly when the corresponding features are in phase conflict. All phase conflicts are resolvable if and only if the vertices of G can be 2-colored with phase 0 and phase 180, which is possible if and only if G is bipartite (i.e., has no odd cycles).

the primary vehicle for lithographic pattern transfer due to superior performance and a more advanced stage of development. The majority of critical circuit layers are imaged onto positive photoresist using bright field masks; this includes polysilicon, metal, and active layers. Dark field masks and positive photoresists are primarily used for contact and via layers.

Much of the early work on PSM design was in fact performed for dark field masks. Although this approach called for the use of negative photoresists, it was widely held that both layout design and mask manufacturing issues could be more readily solved in this case. On the other hand, most commercial applications of phase shifting, which use positive resists, have been based on bright field mask applications [5]. Although such methods have been applied in volume production, they continue to pose mask manufacturing problems that are yet unresolved. Thus, a very significant recent advance has been the commercially viable application of double-exposure phase shifting with a combination of (i) a dark field phase shifting mask with positive photoresist, along with (ii) a bright field binary mask to protect critically-sized features (e.g., poly gates) and also define non-critical features [15] [16]. This technique essentially uses phase-shifting only for the critical features in the layout, leaving non-critical features for the traditional binary mask type. Significant manufacturing-related benefits result, including simplicity and ease of verification for the more complex phase-shifting mask, and simplicity and low cost for the binary mask used in the second exposure [15] [16]. Figures 6 and 7 show the effects of double-exposure dark field phase-shifting in the manufacture of an SRAM cell in 180nm process technology.

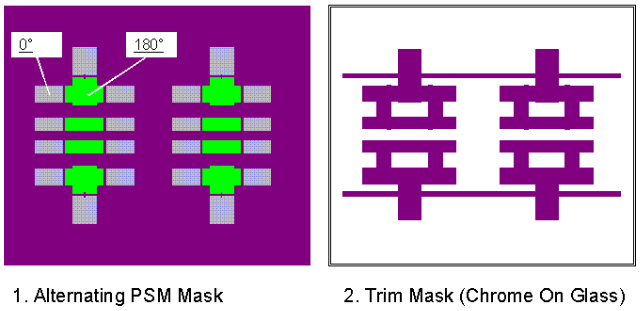

Figure 6: Double-exposure dark field phase shifting mask design for SRAM cell in 180nm process technology.

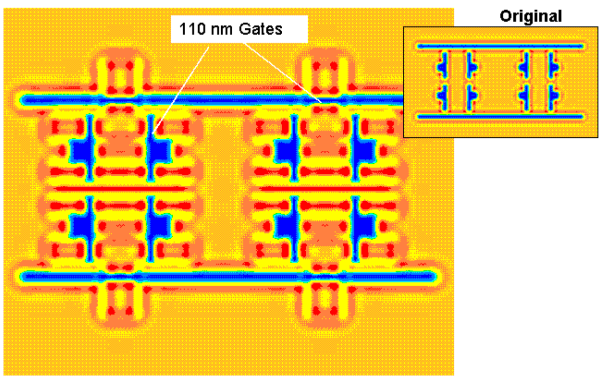

Figure 7: Effect of applying double-exposure dark-field phase-shifting mask to manufacture of SRAM cell in 180nm process technology.

From the perspective of designing phase-shifted layouts, bright field phase-shifting designs pose algorithmic problems that are substantially different from those encountered in dark-field phase shifting layout design. In both cases, the notion of phase conflicts calls for early resolution of such conflicts by introducing phase conflict verification throughout the physical design flow.

PSM Issues

The benefits of PSM include reduced gate lengths, as well as better CD (critical-dimension) control for gate lengths (see Figure 8). This results in higher-performance, lower-power circuits. Applied on the full-chip level, both performance and area gains can be realized since minimum feature spacings are reduced with appropriate phase assignments. On the other hand, the complexity of layout design and verification may increase substantially. In particular, since consistency of the phase assignment is a *global phenomenon* in the layout (as opposed to a local phenomenon such as a DRC tool might check), it is important to reconcile freedom in the (full-chip) layout design with algorithm complexity in the layout verification, or composability of instances in hierarchical design methodologies. Cost and complexity issues arise with respect to mask manufacturing and verification as well. The remainder of this section highlights two key issues: (i) approaches to PSM layout design, particularly in hierarchical and reuse-centric methodologies, and (ii) PSM verification challenges.

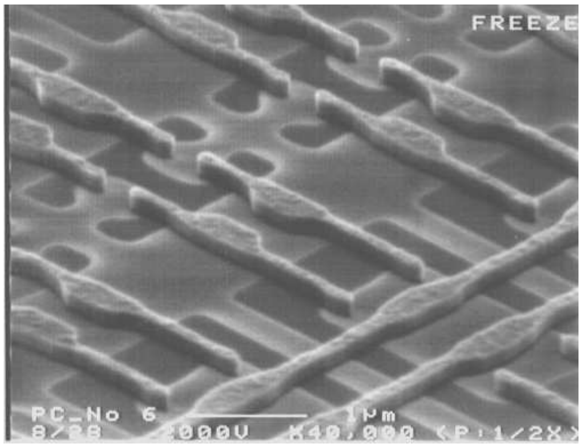

Figure 8: SEM micrograph (courtesy of Motorola) of poly gates fabricated with alternating PSM technology. Gate lengths are 90 nm.

3.1 Challenges for PSM Layout Design and Reuse-Centric Methodologies

Layout design for phase-assignability poses a variety of challenges.

- Early works in the literature center on coloring analysis of the *phase conflict graph* defined above, e.g., if the conflict graph cannot be two-colored, then compaction or other layout modifications are applied so that the modified layout induces a two-colorable conflict graph. With such approaches, a major algorithmic challenge is the computation of a *minimum-cost* set of perturbations that leads to two-colorability. Ideally, such a computation could applied to full-chip layouts.

- Layout perturbations for phase-assignability may explicitly include changes to circuit performance, either making a device slower (increasing gate length to save the cost of introducing phase shifters) or changing a given wire width or spacing (again, saving phase shifters). In this regime, the phase-assignability analysis must be tightly linked with iterative gate-level performance analysis and optimization. The problem becomes one of finding the minimum-cost set of perturbations that is consistent with the circuit performance and area requirement – as well as with the physical layout reoptimization tool (e.g., incremental cell-based or device-level placement and routing, or a pure compaction capability). We also note that there are potentially major implications for how EDA should approach performance-convergent iterations between layout synthesis and logic synthesis. For example, the fact that gates can be selectively phase-shifted or not phase-shifted means that synthesis and mapping may have a richer tradeoff space between switching speed, area, power dissipation and drive strength.

- Phase-assignability may additionally invoke such techniques as "splitting" of large features into independently phase-assignable portions (perhaps via "partial shifters" used for bright-field PSM with positive photoresists [23]), introduction of spacing to decouple sub-areas of the layout in a divide-and-conquer strategy, etc.
- When PSM is applied to local interconnect layers, reassignment of features to alternate layers (re-routing of interconnects) – beyond the sizing and spacing noted above – becomes a viable solution approach to phase conflict. This would imply tight integration of phase-assignability analyses with detailed routing. And in general, all of the foreseeable approaches (for poly routing and metal) require a tight integration of polygon-level layout syntheses with PSM-assignability analyses.[4]

Within hierarchical or reuse-based (e.g., cell-based) design methodologies, recall that the layout context for any given instance is not known *a priori*. With PSM layouts, this is a particularly difficult issue since a phase assignment solution for one cell instance may be incompatible with that of an abutted cell instance. An interesting research and development goal consists of methods to verify the composability of PSM layouts in a hierarchical methodology, as well as methods for hierarchical combination of alternative phase-shifting solutions (e.g., for standard-cell placement). With reuse-centric methodologies, it will be necessary to design layouts that can be phase-assigned to meet performance and area footprint constraints across multiple technologies and migrations. To this end, appropriate design rules (e.g., no T configurations, no uneven-length transistor fingers, limit of certain wrong-way (same-layer jogging) configurations, etc.) that enhance migratability would be beneficial. Finally, if multiple resist technologies should become simultaneous options for fabrication, then exploiting the near-duality of the respective bright- and dark-field design problems (see Figure 9) could be of interest.[5] It is possible that, despite the obvious differences between the two types of technology, there are ways in which design and verification of alternating PSM may addressed independent of the bright or dark field perspective.

3.2 PSM Verification Challenges

As is well known, three basic activities make up today's deep-submicron physical verification flow:

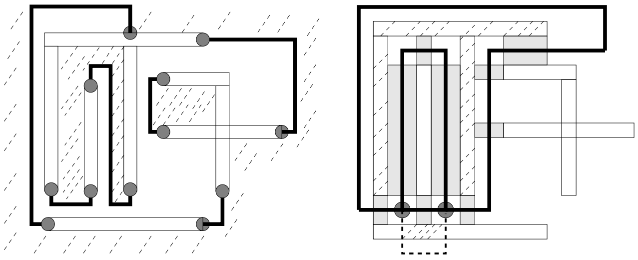

(a) (b)

Figure 9: (a) Bright-field alternating PSM (positive resist) for a layout consisting of four features (eight rectangles). Crosshatched areas are 0 phase; other areas are 180 phase (features are defined by edges between 0- and 180-phase regions). Dark lines indicate phase edges (spurious features) that must be exposed away with a trim mask. (b) Dark-field alternating PSM (negative resist). "Routes" between 0- and 180-phase regions (features) must cover all separations that are less than B, i.e., phase conflicts. The solution for the feature at the bottom of the figure assumes a hypothetical partial-shifter ("vertex splitting") approach to resolve the odd cycle in the conflict graph without changing feature placement.

- design rule checking (DRC), where the layout is checked to ensure compliance with a set of manufacturing process rules such as minimum widths or spacings,
- parasitic extraction, where layout parasitics such as line capacitance and resistance are calculated,[6] and
- layout-versus-schematic checking (LVS), where the original netlist is compared to a netlist derived from the layout.

Each of these steps uses the original physical layout as input, and each of these steps is based on a rules-based process description that assumes that the physical layout will represent the end silicon. However, with OPC and phase-shifting modifications, the original layout is no longer representative of the final silicon pattern; hence, a physical verification tools set that is based solely on original layout information cannot assure silicon results.

As discussed briefly in Section 2.2, the introduction of OPC creates a data explosion that makes the application of OPC difficult. OPC is not easily verified using rules to describe the physical relationships: the rules are too complex and too numerous to be considered by a verification tool without causing huge data volumes and intolerable runtimes. With the introduction of phase-shifting into the design process, rules-based verification becomes even more difficult, e.g., (i) features are introduced into the mask that do not print, and (ii) the layout may be split into two masks, so that a combination of two mask layouts must be verified against the original layout. To solve this problem, a method is needed to verify that the layout source for verification tools will adequately represent the end silicon, either through silicon simulations, altered process models, or a new verification paradigm.

4 Flow Changes and Futures

At the design-manufacturing interface, tools for design synthesis, analysis and verification must work together to enable the tremendous growth in "silicon complexity", design complexity and system complexity that is implied by the prevailing industry roadmaps.

[4] The method of [15] is a notable exception, being applicable today in a largely post-layout regime for critical gate length reduction and CD control.

[5] As noted above, several potentially viable double-exposure solutions for bright-field alternating PSM have been proposed [34, 4, 15]: (i) *phase edges* between 0- and 180-phase regions define thin features; (ii) a second *trim mask* exposure is used to erase unwanted phase edges (i.e., spurious features); and (iii) the key problem in layout design becomes one of "routing" the phase edges so that the trim mask is as simple and as tolerant to registration errors as possible. Figure 9 portrays the interesting *near-duality* of the dark field and bright field regimes.

[6] This is followed by performance analyses (PA) such as delay calculation, static timing, noise and delay uncertainty evaluation, etc.

Tools and methodologies will therefore rely on the following precepts in order to achieve rapid design convergence.

- Upstream tools must pass their constraints and assumptions to downstream tools, and downstream tools must pass failure diagnoses back to upstream tools. (More generally, tools must exploit all available knowledge and all available context, whenever possible.)
- Macromodels for analysis and verification must be abstracted for use as synthesis objectives. (This enables a *prevention-centric* mindset, which is an essential companion to the "checking-centric" mindset that has dominated deep-submicron design practice.)

In the context of subwavelength optical lithography, the above precepts highlight several unnatural aspects of today's separation between "ECAD" design syntheses and "TCAD" manufacturability verifications. With respect to OPC and PSM technologies, many optimizations for manufacturability are quite naturally handled as syntheses (where tools traditionally create the layout), rather than as verifications (where tools traditionally comment on, but are not empowered to change, their inputs). Thus, abstraction and understanding of manufacturing issues should be shifted up: (i) OPC- and PSM-related design rules will move up into global and detailed routing; (ii) PSM phase assignability checks and iterations with compaction will move into detailed routing; (iii) final PSM phase assignment will move up before traditional performance and physical verification; (iv) full-chip OPC insertion, full-chip aerial intensity mapping, "silicon-level" DRC/LVS/PA, and eventually function-centric DRC/LVS/PA will be added into the design flow; etc. At the same time, improved forward annotation of functional intent will ease the burden on verification tools for both layout geometry and mask geometry. Creating these new unifications and flow changes is an important challenge for the EDA industry as well as for the research and development community.

References

[1] A. Chatterjee, I. Ali, K. Joyner, D. Mercer, et al., *Integration of Unit Processes in a Shallow Trench Isolation Module for a 0.25 μm Complementary Metal-Oxide Semiconductor Technology*, Journal of Vacuum Science and Technology B, 15 (1997), pp. 1936–1942.

[2] J. F. Chen, T. Laidig, K. E. Wampler, and R. Caldwell, *Practical Method for Full-Chip Optical Proximity Correction*, in SPIE, vol. 3051, 1997, pp. 790–803.

[3] V. K. R. Chiluvuri and I. Koren, *Layout-Synthesis Techniques for Yield Enhancement*, IEEE Trans. Semiconductor Manufacturing, 8 (1995), pp. 178–187.

[4] G. Galan, F. Lalanne, M. Tissier, and M. Belleville, *Alternating phase shift generation for complex circuit designs*, in SPIE 16th Annual BACUS Symposium on Photomask Technology, vol. SPIE 2884, 1996, pp. 508–519.

[5] P. Gilbert et al., *A High Performance 1.5V, 0.10um Gate Length CMOS Technology with Scaled Copper Metalization*, IEDM 1998, pp. 1013-1016.

[6] W. B. Glendinning and J. N. Helbert, *Handbook of VLSI Microlithography: Principles, Technology, and Applications*, Noyes Publications, 1991.

[7] F. O. Hadlock, *Finding a Maximum Cut of a Planar Graph in Polynomial Time*, SIAM J. Computing, 4 (1975), pp. 221–225.

[8] A. B. Kahng, S. Muddu, E. Sarto, and R. Sharma, *Interconnect Tuning Strategies for High-Performance ICs*, in Proc. Conference on Design Automation and Test in Europe, February 1998.

[9] A. B. Kahng, G. Robins, A. Singh, H. Wang, and A. Zelikovsky, *Filling and Slotting : Analysis and Algorithms*, in Proc. International Symposium on Physical Design, 1998, pp. 95–102.

[10] A. B. Kahng, H. Wang, and A. Zelikovsky, *Automated Layout and Phase Assignment Techniques for Dark Field Alternating PSM*, in Proc. SPIE 18th Annual BACUS Symposium on Photomask Technology, 1998.

[11] M. D. Levenson, *Wavefront engineering from 500 nm to 100 nm CD*, in Proceedings of the SPIE - The International Society for Optical Engineering, vol. 3049, 1997, pp. 2–13.

[12] M. D. Levenson, N. S. Viswanathan, and R. A. Simpson, *Improving Resolution in Photolithography with a Phase-Shifting Mask*, IEEE Trans. on Electron Devices, ED-29 (1982), pp. 1828–1836.

[13] L. Liebmann, A. Molless, R. Ferguson, A. Wong, and S. Mansfield, *Understanding Across Chip Line Width Variation: The First Step Toward Optical Proximity Correction*, in SPIE, vol. 3051, 1997, pp. 124–136.

[14] L. W. Liebmann, T. H. Newman, R. A. Ferguson, R. M. Martino, A. F. Molless, M. O. Neisser, and J. T. Weed, *A Comprehensive Evaluation of Major Phase Shift Mask Technologies for Isolated Gate Structures in Logic Designs*, in SPIE, vol. 2197, 1994, pp. 612–623.

[15] H.-Y. Liu, L. Karklin, Y.-T. Wang, and Y. C. Pati, *The Application of Alternating Phase-Shifting Masks to 140 nm Gate Patterning (II): Mask Design and Manufacturing Tolerances*, in SPIE Optical Microlithography XI, vol. 3334, Feb. 1998, pp. 1-14.

[16] H.-Y. Liu, L. Karklin, Y.-T. Wang, and Y. C. Pati, *The Application of Alternating Phase-Shifting Masks to 140 nm Gate Patterning: Line Width Control Improvements and Design Optimization*, in SPIE 17th Annual BACUS Symposium on Photomask Technology, vol. SPIE 3236, 1998, pp. 328–337.

[17] Y. Liu, A. Zakhor, and M. A. Zuniga, *Computer-Aided Phase Shift Mask Design with Reduced Complexity*, IEEE Transactions on Semiconductor Manufacturing, 9 (1996), pp. 170–181.

[18] W. Maly, *Computer-aided design for VLSI circuit manufacturability*, Proceedings of IEEE, 78 (1990), pp. 356–392.

[19] W. Maly, *Moore's Law and Physical Design of ICs*, in Proc. International Symposium on Physical Design, Monterey, California, April 1998. special address.

[20] A. Misaka, A. Goda, K. Matsuoka, H. Umimoto, and S. Odanaka, *Optical Proximity Correction in DRAM Cell Using a New Statistical Methodology*, in SPIE, vol. 3051, 1997, pp. 763–773.

[21] A. Moniwa, T. Terasawa, N. Hasegawa, and S. Okazaki, *Algorithm for Phase-Shift Mask Design with Priority on Shifter Placement*, Jpn. J. Appl. Phys., 32 (1993), pp. 5874–5879.

[22] A. Moniwa, T. Terasawa, K. Nakajo, J. Sakemi, and S. Okazaki, *Heuristic Method for Phase-Conflict Minimization in Automatic Phase-Shift Mask Design*, Jpn. J. Appl. Phys., 34 (1995), pp. 6584–6589.

[23] J. Nistler, G. Hughes, A. Muray, and J. Wiley, *Issues Associated with the Commercialization of Phase Shift Masks*, in SPIE 11th Annual BACUS Symposium on Photomask Technology, vol. SPIE 1604, 1991, pp. 236–264.

[24] K. Ooi, S. Hara, and K. Koyama, *Computer Aided Design Software for Designing Phase-Shifting Masks*, Jpn. J. Appl. Phys., 32 (1993), pp. 5887–5891.

[25] K. Ooi, K. Koyama, and M. Kiryu, *Method of Designing Phase-Shifting Masks Utilizing a Compactor*, Jpn. J. Appl. Phys., 33 (1993), pp. 6774–6778.

[26] G. I. Orlova and Y. G. Dorfman, *Finding the Maximum Cut in a Graph*, Engr. Cybernetics, 10 (1972), pp. 502–506.

[27] P. Rai-Choudhury, *Handbook of Microlithography, Micromachining, and Microfabrication, vol. 1: Microlithography*, SPIE Optical Engineering Press, Bellingham, 1997.

[28] F. M. Schellenberg, H. Zhang, and J. Morrow, *Evaluation of OPC Efficacy*, in Proc. Intl. Symp. on Aerospace/Defense Sensing and Dual-Use Photonics, vol. 2726, 1996, pp. 680–688.

[29] SEMATECH, *Workshop Notes*, in 3rd SEMATECH Litho-Design Workshop, Skamania Lodge, February 1996.

[30] SIA, *The National Technology Roadmap for Semiconductors*, Semiconductor Industry Association, December 1997.

[31] B. E. Stine, D. S. Boning, J. E. Chung, and L. Camilletti, *The Physical and Electrical Effects of Metal-fill Patterning Practices for Oxide Chemical-Mechanical Polishing Processes*, IEEE Transactions on Electron Devices, 45 (1998), pp. 665–679.

[32] B. E. Stine, V. Mehrotra, D. S. Boning, J. E. Chung, and D. J. Ciplickas, *A Simulation Methodology for Assessing the Impact of Spatial/Pattern Dependent Interconnect Parameter Variation on Circuit Performance*, in IEDM Technical Digest, 1997, pp. 133–136.

[33] D. Sylvester and K. Keutzer, *Getting to the Bottom of Deep-Submicron*, in Proc. IEEE Intl. Conf. Computer-Aided Design (to appear), November 1998.

[34] T. Waas, H. Hartmann, and W. Henke, *Automatic Generation of Phase Shift Mask Layouts*, Microelectronic Engineering, 23 (1994), pp. 139–142.