

IC Layout and Manufacturability: Critical Links and Design Flow Implications*

Andrew B. Kahng

UCLA Department of Computer Science, Los Angeles, CA 90095-1596 USA

Abstract

We assess the prospects for new tools and flows in the interface between layout design and manufacturability. We begin with a review of classic elements of this interface, then focus on more recently critical issues: (i) layout design for reduced CMP variability; (ii) layout design for PSM; and (iii) layout design for OPC. Our discussion highlights the many ways in which layout affords effective means of optimizing manufacturability, as well as opportunities for research and development.

1 Introduction

As CMOS technology advances according to the SIA NTRS [42], manufacturing cost increasingly drives design [26]. Process engineers must 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. A total *variability budget* for the design is distributed among such attributes. In very deep submicron technologies, attaining large process windows and uniform manufacturing is difficult [1] [10] [38] [18] [26] [6], and the manufacturing process has an increasingly constraining effect on physical layout design and verification.

Several aspects of the heightened impact of design-manufacturing links can be traced to two recent crossover points in the evolution of VLSI technology. The first crossover occurs when interconnect delays dominate device switching delays in deep-submicron CMOS technology [42]. Interconnect process optimization must achieve more delicate balances, e.g., affording simultaneous distribution of signal, clock and power with adequate performance while minimizing die area. Also, more interconnect layers are required at each successive node in the technology roadmap [42, 11, 13, 47], leading to a strong requirement for *planarized* interconnect processes that rely on *chemical-mechanical polishing* (CMP). The second crossover occurs when minimum feature dimensions and spacings decrease below the wavelength of the light source [39] [36]. As reviewed in [39], pattern fidelity deteriorates markedly in this *subwavelength lithography* regime, leading to compensation mechanisms [19] that either perturb the shape (via *optical proximity correction* (OPC)) or the phase (via *phase-shifting masks* (PSM)) of transmitting apertures in the reticle.

This paper will assess the prospects for new tools and flows in the interface between layout design and manufacturability. We begin with a review of classic elements of this interface, then focus on more recently critical issues: (i) layout design for reduced CMP variability; (ii) layout design for PSM; and (iii) layout design for OPC. Our discussion will highlight the many ways in which layout affords effective means of optimizing manufacturability, as well as opportunities for research and development.

* This research was supported by a contract and grant from Cadence Design Systems, Inc.

2 Classic Design-Manufacturing Links

Many physical design methods have been proposed to address various manufacturing issues such as registration errors, photolithographic random effects, random spot defects, plasma and charging effects (“antenna effect”), etc.; see such works as [25] [6] for reviews. In this section, we sketch two of the many classic issues that are in some sense “solved”.

2.1 Antenna Effect

Many processing steps use plasmas and charged particles, during which charge can collect on poly and metal surfaces. Through capacitive coupling, large electric fields may develop over gate oxides, resulting in stresses that cause oxide breakdown and shifts in threshold voltage V_t . The area of the *antenna* connected to a gate is the sum of all poly and metal antenna areas; such area for, e.g., the i^{th} metal layer M_i is the total area on that layer that is electrically connected to the gate without using layer M_{i+1} , but that does not have any connection to an active via lower levels of interconnect. Typically, foundries impose *antenna rules* that limit the ratio of antenna area to gate oxide area for any individual gate. As discussed in such works as [51, 5], there are two ways to control antenna effects (see Figure 1): (i) insertion of protection diodes, and (ii) layer changes, or *bridging*. The former solution is expensive in terms of area and leakage power, and hence the routing-based bridging solution is preferable.




Figure 1: Control of antenna effect by diode insertion or bridging (routing).

2.2 Critical Area

Modern manufacturing processes have a high susceptibility to *defect-related yield loss* [25], particularly in metallization stages. The failure mechanisms of interest are due to extra or missing material in the routing layers, e.g., extra material can short two lines together, or missing material can create an open. Layout tools do not account for defect-related yield loss well: they fail to understand that design rule constraints reflect minimum values, as opposed to target values. For a given defect radius r , consider the set of all points x in the layout at which such a radius- r defect, if centered at x , would cause a failure. The total layout area covered by such points is a function of the defect radius r , and is called the *critical area* [6] for that defect radius. A number of works, e.g., [6, 49], show that critical area can be effectively minimized

by layer assignment or by spacing, as illustrated in Figure 2. The former technique requires a router, while the latter requires either a router or a compaction tool.




Figure 2: Reduction of critical area by layer assignment (routing) and spacing (compaction).

3 Layout Design for Reduced CMP Variability

Chemical-mechanical polishing (CMP) is the procedure by which wafers are polished using a rotating pad and slurry to achieve the planarized surfaces on which succeeding processing steps can build [18] [31] [48]. In very deep-submicron VLSI, manufacturing steps involving *chemical-mechanical polishing* (CMP) have varying effects on device and interconnect features, depending on local characteristics of the layout.¹ Recent work in statistical metrology shows that CMP variation is controlled if the *local feature density* is controlled [37] [46].² Post-CMP oxide thickness is monotone in feature density, so variation of local feature density must be minimized. This is achieved by inserting “fill” geometries into the layout.

The Filling Problem: Given a design rule-correct *layout* geometry of k disjoint rectilinear rectangles in an $n \times n$ layout region, minimum feature size c , window size $w < n$, buffer distance B , and area (or perimeter) density lower bound L and upper bound U , add fill geometries to create a *filled layout* that satisfies the following conditions: (1) circuit function and design rule-correctness are preserved; (2) no fill geometry is within distance B of any layout feature; (3) no fill is added into any $w \times w$ window that has density $\geq U$ in the original layout; (4) for any $w \times w$ window that has density $< U$ in the original layout, the filled layout density is $\geq L$ and $\leq U$; and (5) the minimum window density in the filled layout is maximized.

Here, U and L are real numbers between 0 and 1, i.e., each $w \times w$ region of the layout must contain total area of features $L \cdot w^2 \leq \text{area} \leq U \cdot w^2$. Condition (5) establishes the “minimum-variation” objective.

Today’s industry practice enforces feature density bounds only within a fixed set of $w \times w$ windows in the layout. Define a *fixed dissection* of the layout region to be a decomposition into $(\frac{n}{w})^2$ non-overlapping $w \times w$ windows. Since bounding the density in windows of a dissection can

¹In oxide polishing of interlayer dielectrics, the pad conforms to local topography and overpolishes empty oxide areas that have no underlying metal features (“dishing”); areas with dense underlying metal features are underpolished. The typical oxide thickness variation of up to 4000 angstroms uses up a large fraction of the die’s variability budget [14] [44] [45]. CMP variation is even more critical with the adoption of shallow-trench isolation and inlaid-metal (e.g., damascene copper) processes [3] [30] [43] [2].

²The definition of “local” is determined by the *length scale* at which feature density impacts oxide thickness, and defines the “window size” within which feature density must be controlled. For oxide CMP, this length scale has been estimated to be on the order of 1-3mm [7] [37] [46].

incur error (i.e., other windows not in the dissection can violate the density bound), a common practice is to enforce density bounds in r^2 dissections simultaneously, where r determines the “phase shift” w/r by which the dissections are offset from each other. That is, density bounds are enforced only for windows of a *fixed r -dissection*, which is the set of $w \times w$ windows having bottom-left corners at points $(i \cdot \frac{w}{r}, j \cdot \frac{w}{r})$, for $i, j = 0, 1, \dots, r(\frac{n}{w}) - 1$, where r is an integer divisor of w . A fixed r -dissection divides the layout into $\frac{nr}{w} \times \frac{nr}{w}$ *tiles*, each of size $\frac{w}{r} \times \frac{w}{r}$ (see Figure 3(a)).

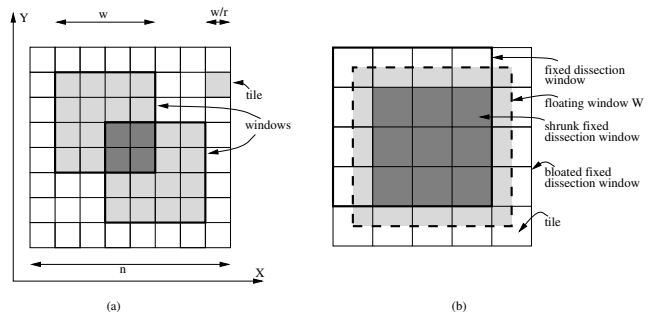


Figure 3: (a) The layout is partitioned by r^2 ($r = 4$) distinct dissections, each dissection having window size $w \times w$, into $\frac{nr}{w} \times \frac{nr}{w}$ tiles. Each $w \times w$ window (dark) consists of r^2 tiles. A pair of windows from different dissections may overlap. (b) An arbitrary $w \times w$ floating window W always contains an $(r-1) \times (r-1)$ shrunk window of a fixed r -dissection, and is always covered by an $(r+1) \times (r+1)$ bloated window of the fixed r -dissection.

The industry state of the art in density control for CMP has three weaknesses.³ (1) Many foundry density rules still constrain only the average *overall* feature density on a given layer, rather than local variation in feature density. (2) Current approaches to analysis of layout density do not find true extremal window densities in the layout. Rather, they find the extremal window densities over a fixed set of window positions within the “fixed r -dissection” approach. (3) Current methods for inserting fill geometries into the layout do not actually minimize the maximum variation in layout density between windows of the layout. Rather, simple Boolean layer processing techniques are applied to insert fill patterns into any empty region that is sufficiently large. Our problem statement above addresses the first weakness; we now address solutions for the second and third weaknesses.

3.1 Exact Analysis of Extremal Window Density

A fundamental problem in density analysis is to find the window with the maximum (or minimum) feature density. The fixed-dissection approach can incur substantial error [14] and yet is used in practice because optimal algorithms are slow. A promising new approach [16] is based on: (see Figure 3(b)):

Observation: Given a fixed r -dissection, any arbitrary $w \times w$ floating window of the layout will contain some $w(1 - 1/r) \times w(1 - 1/r)$ shrunk window of the fixed r -dissection, and will be contained in some $w(1 + 1/r) \times w(1 + 1/r)$ bloated window of the fixed r -dissection.

³These may be attributed to the genealogy of tools for layout density control. Such tools are typically evolved from physical verification and mask processing tools, where the mindset is one of verification rather than data modification, local rules rather than global rules, and Boolean rules rather than context-dependent rules.

This observation suggests an *exact, multilevel* density analysis approach. First, we can find densities of bloated windows in the fixed r -dissection, and this will bound the densities of all $w \times w$ floating windows. Second, the fixed-dissection density analysis for any given $r = r_0$ can be sped up by (i) starting with $r = 1$ and eliminating tiles that do not belong to any bloated window that can possibly contain high-density floating windows, then (ii) recursively subdividing the remaining tiles into 4 subtiles (i.e., multiply r by 2) until the necessary $r = r_0$ is reached. Finally, the recursive subdivision can be continued until the number of rectangles left in tiles is sufficiently small to run optimal (i.e., exact) density analysis, or when some user-defined accuracy tolerance ϵ is reached. This approach can be implemented to have worst-case runtime of $O((\frac{r}{w} \log \frac{w}{\epsilon})^2)$; it is faster and more accurate than fixed r -dissection analyses, especially for large designs [16].

Future research issues include the impact of emerging process technologies (notably involving reverse active-area masks) that require simultaneous perimeter- and area-density based analysis criteria. Density control requirements may also be defined at multiple length scales simultaneously, e.g., due to various loading effects. This is an interesting challenge for density analysis as well as for the fill synthesis approaches discussed in the next subsection. Yet another refinement stems from the proper window shapes being circular, rather than square. Finally, integration of density control for CMP with layout synthesis will require new, dynamic versions of the algorithms for density analysis, i.e., for efficient maintenance of min/max density/perimeter information as the layout is modified.

3.2 Synthesis of Optimum Fill Amounts

A second fundamental issue is how to decide what amounts of fill should be inserted into the layout, and where these geometries need to be added. Suppose that we know, for each tile T_{ij} ($i, j = 1, \dots, \frac{w}{r}$) in a fixed r -dissection, the total feature area in the tile, $area(T_{ij})$. Also suppose that we know the *slack* of the tile, $slack(T_{ij})$, defined as the maximum fill area that can be added to T_{ij} without violating the density upper bound of any window of the fixed r -dissection that contains T_{ij} . (The area and slack of all tiles can be computed efficiently as a byproduct of density analysis.) Then, for each tile T_{ij} , the total fill pattern area $p_{ij} = p(T_{ij})$ to be added to T_{ij} must satisfy $0 \leq p_{ij} \leq slack(T_{ij})$ and the constraint that no window can have density more than U after filling unless it was overfilled initially: Given this notation, the essence of the Filling Problem becomes:

$$\text{maximize} \quad \left(\min_{ij} (area(T_{ij}) + p_{ij}) \right)$$

i.e., we seek to maximize the minimum window density in the filled layout. As described in [15], this can be exactly solved by *linear programming*.

Other important issues include the support of hierarchical designs, so that layouts with large hierarchy factors will not explode when we apply filling analysis and synthesis. In performing fill synthesis under simultaneous area and perimeter density bounds, even the question of fill *pattern* synthesis becomes quite interesting. Finally, new methods will be needed (integrating, e.g., special-net routers) to synthesize *grounded* fill effectively.

4 Layout Design for Phase-Shifting Masks

Phase-shifting mask (PSM) technology enables the clear regions of a mask to transmit light with prescribed phase shift.

Consider two adjacent clear regions with small separation, and respective phase shifts of 0 and 180 degrees (see Figure 4). 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. [20] in 1982. See [10] [38] [23] [41] [22] [24] for reviews of PSM technologies.




Figure 4: 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 [33]. 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.

The Phase Assignment Problem may be stated in the context of the *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). Hence, if G is not bipartite, the Phase Assignment Problem requires us to *delete* enough edges such that no odd cycles exist in the remaining modified conflict graph. Edge deletion in the conflict graph is achieved by changing the placement of layout features so that they no longer conflict (see Figure 5).

Early methods for edge deletion in the conflict graph are due to [28, 33]. The heuristic of [29] (i) constructs the conflict graph G , (ii) creates a list of all odd cycles in G using an enumerative approach, and (iii) iteratively finds and deletes




Figure 5: Edges of an odd cycle in the conflict graph (here, a triangle induced by conflicts among three features) can be removed by changing the placement of a feature.

the edge that is in the greatest number of minimum-length odd cycles. Deletion is accomplished by increasing the lower bound on separation between the corresponding features, and then applying a compactor to perturb the shape or position of these features. This approach does not scale to large instances since the number of odd cycles can be exponential in the size of the layout. Moreover, it does not necessarily delete the minimum number of edges, nor will it necessarily select edges whose deletion will have minimum impact. The method of [34] (i) produces a symbolic layout from the mask layout; (ii) performs phase assignment in the symbolic layout; and (iii) compacts the symbolic layout using minimum spacing design rules consistent with the phase assignment. However, this method can significantly increase the layout area.

4.1 Approaches to Layout Modification for PSM

As discussed in [17], it is natural to seek phase-assignability using a layout modification step after detailed routing. For example, the approaches of [34, 29] can be generalized as follows. First, initially constrain the layout only by the minimum-spacing design rule, i.e., no two features can be less than distance b apart. Then, *iteratively* apply the following three steps until the conflict graph G becomes bipartite: (i) compact the layout and find the conflict graph G ; (ii) find the minimum-cost set of edges whose deletion makes the conflict graph G 2-colorable; and (iii) add a new compaction constraint for each edge in this minimum set, such that the pair of features connected by this edge must be separated by distance at least B . This approach not only requires integration of device-level compaction capability, but also requires an optimal algorithm to make the conflict graph bipartite. I.e., given a planar graph $G = (V, E)$ with weighted edges, we seek the minimum-weight edge set M such that the graph $(V, E - M)$ contains no odd cycles. An exact algorithm for this problem can be traced to Hadlock [12] and Orlova et al. [35];⁴ a faster algorithm was presented in [17]. We observe that the weight of an edge in the conflict graph should reflect the ease of perturbing the layout to break that particular conflict. For example, integration with a compactor allows us to use slack to determine edge weight.

A second layout modification approach is to use “splitting” of sufficiently long features into several parts. This is equivalent to splitting one vertex into several vertices in

⁴To eliminate all odd cycles it is sufficient to eliminate odd faces of the planar graph G . Consider a graph D that is the geometric dual of G : all faces of G are vertices of D (note that there is a vertex of D that corresponds to the “outer face” of G), and each edge of D intersects with exactly one edge of G . The odd faces of G correspond to odd-degree vertices of D . Any edge elimination in G corresponds to edge contraction in D ; in particular, we may remove a pair of odd faces from G by contracting edges of a (minimum-weight) path between the corresponding odd-degree vertices in D . Hence, the minimum-weight set of edges to break in the conflict graph corresponds to a minimum-weight matching of odd-degree vertices in D , which can be found in polynomial time.

the conflict graph, and allows assignment of different phases to neighboring parts within the same feature. Introducing *partial shifters* (used for bright field PSM with positive photoresists [32]) between these parts allows gradual transitions between 0-phase and 180-phase with no dark edge being formed. The splitting technique requires a more complicated mask, but has two advantages: (i) it can only decrease the number of odd cycles, and (ii) it does not perturb the layout geometry.

Finally, it is possible to exploit *layer reassignment* to remove phase conflicts. This degree of freedom is natural for a routing tool, but unnatural for current mask optimization tools that view a given layer’s geometries as immutable. With layer assignment, a problematic *wire* feature can be replaced by two vias and transferred to another layer (the approach does not apply to device layers). Such an approach will require the integration of device-level routing capability in a tight loop with the bicolorability analysis.

4.2 Related Directions

Our discussion has been in the context of *dark field* (*alternating* or *Levenson-type*) phase-shifting mask technology with *negative photoresist*, following, e.g., [28, 29, 33, 34]. While the resulting single-exposure regime is attractive from a cost perspective, it is not yet known whether negative or positive resists will afford the most contrast for future process generations. We observe that a number of double-exposure solutions for bright-field alternating PSM (with positive resists) have been proposed [50, 9, 23]: (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 6 portrays the interesting *near-duality* of the dark field and bright field regimes. A valuable contribution would be a unified theory for alternating PSM that is independent of the bright or dark field (positive or negative resist) perspective. Also of interest are design rules (e.g., no T configurations, no uneven-length transistor fingers) as well as methods for hierarchical combination of phase-shifting solutions (e.g., for standard-cell placement).

5 Layout Design for Optical Proximity Correction

As noted above, *optical proximity correction* perturbs the shapes of transmitting apertures in the mask (see [21, 40, 4, 8, 27] for reviews). These changes 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” [4]) for narrow gate geometries. Figure 7 conveys the flavor of a layout after reasonably aggressive OPC has been applied.

OPC is very much a fact of life in deep-submicron (subwavelength) lithography, both today and into the future. OPC is also more “mature” than filling and phase-assignment, in terms of available software solutions. Thus, in contrast to the more “pioneering” CMP- and PSM-related needs, OPC-related research must be aimed more at disconnects within an existing infrastructure. We highlight two disconnects in particular: (i) the application of OPC without regard to functional requirements implicit in a feature, and (ii) the application of OPC techniques without regard to verifiability (e.g., at the mask writing step).




Figure 6: (a) Bright-field alternating PSM 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. “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 uses the partial-shifter (“vertex splitting”) approach to resolve the odd cycle in the conflict graph without changing feature placement.

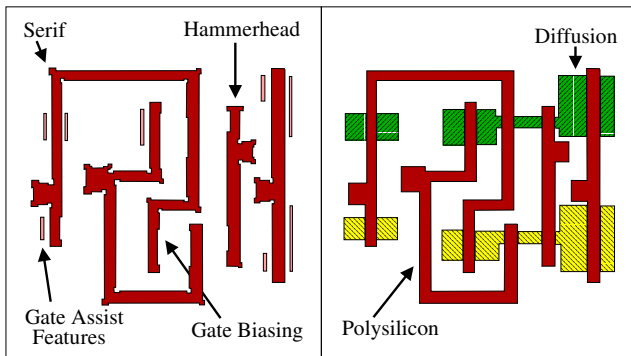


Figure 7: Aggressive OPC with subresolution gate assist features (“outriggers”). Figure is redrawn from the OPC Technology web page (www.opctech.com).

5.1 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.

New methods must be developed for 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. (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 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 related 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.

5.2 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 reparability of the mask. A modest first goal would be to develop new methods of abstracting the limitations of mask verification up to the OPC insertion stage: OPC insertion should not make corrections 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.

6 Flow Changes and Futures

ECAD and TCAD must work together to enable the tremendous growth in “silicon complexity”, design complexity and system complexity that is implied by the SIA NTRS. On the ECAD side, tools and methodologies will most likely 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.)

The discussion above has pointed out several unnatural aspects of the separation between ECAD’s design syntheses and TCAD’s manufacturability verifications. Particularly in the three critical areas of CMP, PSM and OPC, 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) CMP fill awareness will move up as early as post-placement performance analysis (PA); (ii) OPC- and PSM-related design rules,

and fill pattern models, will move up into global and detailed routing; (iii) PSM phase assignability checks and iterations with compaction will move into detailed routing; (iv) fill insertion and final PSM phase assignment will move up before traditional performance and physical verification; (v) full-chip OPC insertion, full-chip aerial intensity mapping, "silicon-level" DRC/LVS/PA, and eventually function-centric DRC/LVS/PA will be added as subsequent steps in the ECAD flow. At the same time, improved forward annotation of functional intent will ease the burden on TCAD verification tools. Creating these new unifications of ECAD with TCAD is an important challenge for the research and development community.

References

- [1] L. E. CAMILLETTI, *Implementation of CMP-based Design Rules and Patterning Practices*, in 1995 IEEE/SEMI Advanced Semiconductor Manufacturing Conference, 1995, pp. 2-4.
- [2] L. E. CAMILLETTI, personal communication, (1998).
- [3] 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.
- [4] 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.
- [5] Z. CHEN AND I. KOREN, *Layer reassignment for antenna effect minimization in 3-layer channel routing*, in Proceedings IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems, 1996, pp. 77-85.
- [6] V. K. R. CHILUVURI AND I. KOREN, *Layout-Synthesis Techniques for Yield Enhancement*, IEEE Trans. Semiconductor Manufacturing, 8 (1995), pp. 178-187.
- [7] R. R. DIVECHA, B. E. STINE, D. O. OUMA, J. U. YOON, D. S. BONING, J. E. CHUNG, O. S. NAKAGAWA, AND S. Y. OH, *Effect of Fine-line Density and Pitch on Interconnect ILD Thickness Variation in Oxide CMP Process*, in Proc. CMP-MIC, Santa Clara, February 1998.
- [8] D. DOLAINSKY AND W. MAURER, *Application of a Simple Resist Model to Fast Optical Proximity Correction*, in SPIE, vol. 3051, 1997, pp. 774-780.
- [9] 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.
- [10] W. B. GLENDINNING AND J. N. HELBERT, *Handbook of VLSI Microlithography: Principles, Technology, and Applications*, Noyes Publications, 1991.
- [11] L. GWENNAP, *IC Vendors Prepare for 0.25-Micron Leap*, Microprocessor Report, September 16 (1996), pp. 11-15.
- [12] F. O. HADLOCK, *Finding a Maximum Cut of a Planar Graph in Polynomial Time*, SIAM J. Computing, 4 (1975), pp. 221-225.
- [13] 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.
- [14] 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.
- [15] A. B. KAHNG, G. ROBINS, A. SINGH, AND A. ZELIKOVSKY, *New and Exact Filling Algorithms for Layout Density Control*, in Proceedings of the 12th International Conference on VLSI Design, 1999.
- [16] A. B. KAHNG, G. ROBINS, A. SINGH, AND A. ZELIKOVSKY, *New Multi-Level and Hierarchical Algorithms for Layout Density Control*, in Proceedings of the Asia and South Pacific Design Automation Conference, 1999.
- [17] 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.
- [18] H. LANDIS, P. BURKE, W. COTE, W. HILL, C. HOFFMAN, C. KAANTA, C. KOBURGER, W. LANGE, M. LEACH, AND S. LUCE, *Integration of Chemical-Mechanical Polishing into CMOS Integrated Circuit Manufacturing*, Thin Solid Films, 220 (1992), pp. 1-7.
- [19] 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.
- [20] 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.
- [21] L. LIEBMAN, 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.
- [22] L. W. LIEBMAN, 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.
- [23] 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.
- [24] 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.
- [25] W. MALY, *Computer-aided design for VLSI circuit manufacturability*, Proceedings of IEEE, 78 (1990), pp. 356-392.
- [26] W. MALY, *Moore's Law and Physical Design of ICs*, in Proc. International Symposium on Physical Design, Monterey, California, April 1998. special address.
- [27] 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.
- [28] 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.
- [29] 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.
- [30] S. NAG AND A. CHATTERJEE, *Shallow Trench Isolation for Sub-0.25- μm IC Technologies*, Solid State Technology, 40 (1997), pp. 129-130, 132, 134, 136.
- [31] G. NANZ AND L. E. CAMILLETTI, *Modeling of Chemical-Mechanical Polishing: A Review*, IEEE Trans. on Semiconductor Manufacturing, 8 (1995), pp. 382-389.
- [32] 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.
- [33] 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.
- [34] 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.
- [35] G. I. ORLOVA AND Y. G. DORFMAN, *Finding the Maximum Cut in a Graph*, Engr. Cybernetics, 10 (1972), pp. 502-506.
- [36] Y. C. PATI, personal communication, (1998).
- [37] S. PRASAD, W. LOH, A. KAPOOR, E. CHANG, B. STINE, D. BONING, AND J. CHUNG, *Statistical metrology for characterizing CMP processes*, Microelectronic Engineering, 33 (1997), pp. 231-240.
- [38] P. RAI-CHOUDHURY, *Handbook of Microlithography, Micromachining, and Microfabrication, vol. 1: Microlithography*, SPIE Optical Engineering Press, Bellingham, 1997.
- [39] F. SCHELLENBERG, *Design for Manufacturing in the Semiconductor Industry: The Litho/Design Workshops*, unpublished manuscript, (1998).
- [40] 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.
- [41] SEMATECH, *Workshop Notes*, in 3rd SEMATECH Litho-Design Workshop, Skamania Lodge, February 1996.
- [42] SIA, *The National Technology Roadmap for Semiconductors*, Semiconductor Industry Association, December 1997.
- [43] K. SMEKALIN, *CMP Dishing Effects in Shallow Trench Isolation*, Solid State Technology, 40 (1997), pp. 187-188, 190, 192, 194.
- [44] 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.
- [45] B. E. STINE, V. MEHROTRA, D. S. BONING, J. E. CHUNG, AND D. J. CIPICKAS, *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.
- [46] B. E. STINE, D. O. OUMA, R. R. DIVECHA, D. S. BONING, J. E. CHUNG, D. L. HETHERINGTON, C. R. HARWOOD, O. S. NAKAGAWA, AND S.-Y. OH, *Rapid characterization and modeling of pattern-dependent variation in chemical-mechanical polishing*, IEEE Transactions on Semiconductor Manufacturing, 11 (1998), pp. 129-140.
- [47] D. SYLVESTER AND K. KEUTZER, *Getting to the Bottom of Deep-Submicron*, in Proc. IEEE Intl. Conf. Computer-Aided Design (to appear), November 1998.
- [48] M. TOMOZAWA, *Oxide CMP Mechanisms*, Solid State Technology, (1997), pp. 169-175.
- [49] A. VENKATARAMAN, H. H. CHEN AND I. KOREN, *Yield-enhanced routing for high-performance VLSI designs*, Proc. Microelectronic Manufacturing Yield, Reliability, and Failure Analysis III, 1997, pp. 51-60.
- [50] T. WAAS, H. HARTMANN, AND W. HENKE, *Automatic Generation of Phase Shift Mask Layouts*, Microelectronic Engineering, 23 (1994), pp. 139-142.
- [51] K. P. WANG, M. MAREK-SADOWSKA, AND W. MALY, *Layout Design for Yield and Reliability*, in Proc. ACM/SIGDA Physical Design Workshop, 1996, pp. 190-197.