

Reticle Enhancement Technology Trends: Resource and Manufacturability Implications for the Implementation of Physical Designs

Warren Grobman
Motorola SPS
3501 Ed Bluestein Blvd.
MD: K10
Austin, TX 78721
1-512-933-6721
warren.grobman@
motorola.com

Robert Boone
Motorola SPS
3501 Ed Bluestein Blvd
MD: K10
Austin, TX 78721
1-512-933-3559
robert.e.boone@
motorola.com

Cece Philbin
Motorola SPS
3501 Ed Bluestein Blvd
MD: K20
Austin, TX 78721
1-512-933-2627
rztw70@
email.sps.mot.com

Bob Jarvis
Motorola SPS
3501 Ed Bluestein Blvd
MD: K10
Austin, TX 78721
1-512-933-5851
ra5924@
email.sps.mot.com

ABSTRACT

In this paper, we describe the industry trends that require reticle enhancement technology during mask data preparation, as well as the diversity and complexity of solutions that are increasingly being implemented. These calculations manipulate the final physical design to create mask data that corrects for some of the resolution deficiencies of “subwavelength” lithography. Subwavelength refers to a regime in which the design critical dimension is \ll the wavelength of light divided by the stepper numerical aperture.

This portion of the design flow has been relatively unimportant in the past, but is increasingly critical in recent semiconductor generations. In the past, post-physical design (post GDSII) manipulation has been limited to trivial modifications, e.g. sizing to correct for process bias. Now the subwavelength regime is driving us to fundamental modification of the physical design prior to mask making. [1;2]

However, resources such as computational capacity and cycle time, storage, and memory are increasing by orders of magnitude as we implement these techniques. In addition, mask fabrication resources and cycle time, as well as mask yield risk and inspection issues, increase dramatically as we drive into the subwavelength regime. New post tape-out approaches will be required to address these challenges. These include both “brute force” approaches, such as methods of parallelization of mask data preparation, and more elegant approaches to these problems. One such approach involves restriction of aggressive computation to regions in which the design is especially sensitive to lithography process issues. Another powerful new paradigm would involve supplementing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISPD '01, April 1-4, 2001, Sonoma, California, USA.
Copyright 2001 ACM 1-58113-347-2/01/0004...\$5.00.

GDSII with a new data representation format capable of capturing physical data hierarchy irrespective of original design hierarchy. Finally, it is intriguing to consider a novel approach in which one partitions a flattened design, computes solutions on clustered computers, and then recaptures the physical design hierarchy afterwards for data size efficiency.

Categories and Subject Descriptors

J.6 [Computer-Aided Engineering]: Computer-aided design (CAD)

E.1 [Data Structures]: Graphs and networks

General Terms

Algorithms, Management, Measurement, Performance, Design, Economics, Reliability, Standardization, Verification.

Keywords

RET, Reticle Enhancement Technology, subwavelength lithography, mask data preparation, phase shift, OPC, tiling.

1. INTRODUCTION

As the semiconductor industry evolves toward higher densities of integrated circuits at smaller dimensions, the lithographic patterning capability becomes much more difficult to implement. In 1998 we were manufacturing 250 nm minimum feature circuits, using steppers with 248nm wavelength. Today we are about to qualify 130nm processes, while the next generation of lithography, at 193nm, is not yet mature.

One consequence of this trend is that in order to print manufacturable features on Si, we must modify physical design data so that mask data is different and corrects for lithographic deficiencies. This modification is accomplished by the application of many different types of Reticle Enhancement Technology (RET). For example, some design features are modified to form reticle features that facilitate printing (by correcting for the distortion effects of subwavelength printing), with phase-shift or scattering bar technology used for the most aggressive features (usually the gate level).

In addition, new features will be required in sparse areas to ensure better planarity, as the chemical-mechanical polish progresses at different rates in sparse and dense areas. This is true both for front-end and back-end polish. Planarity of polish achieves manufacturable results for two reasons: more uniform processing of features such as vias, and planarity during lithographic exposure. Given the severe deterioration of imaging acuity in the subwavelength regime, keeping the resist globally close to the focal plane is an important manufacturability benefit.

One major challenge, the topic of this paper, is the extreme rate of growth of resources required to transform design data into mask data. The driving forces for this resource explosion are the combination of many added shapes which are the result of RET, the growth of circuit count per square mm of Si (and consequent design size), and the deficiencies of conventional data formats (GDSII) when representing designs whose hierarchy is compromised by the application of RET.

We will describe briefly the types of RET, and will then show their proliferation as technology generations evolve. We will also quantify the experience of our and other companies with data volume and processing time issues caused by the lithographic challenge. We will also quantify the concomitant mask manufacturability issues that arise. These express themselves both in mask data size issues for the mask write process, and mask manufacturing and inspection problems. Finally, we suggest some future alternatives, both brute force (e.g. bigger computers) and elegant (smarter data representations, EDA software, etc.) which can help the industry meet this emerging challenge.

2. RET TECHNOLOGIES AND THEIR PROLIFERATION AND DESIGN DATA ISSUES

Optical Proximity Correction (OPC) is the modification of design features to correct for the “transform” applied to design shapes by the convolution of the reticle fabrication process, the imaging on resist, and the developed image transfer to Si features. (Figure 1)

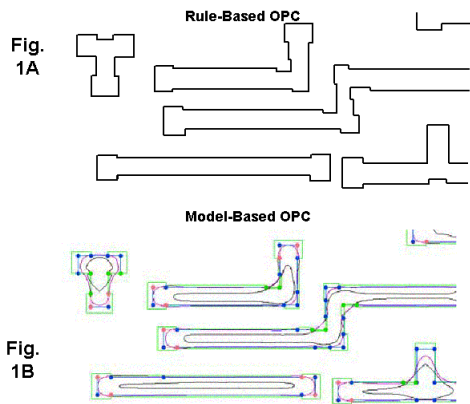


Figure 1: Two Methods for Optical Proximity Correction

Figure 1A shows rule-based modification of individual features, a method that is calibrated to the process and applies to each design shape independently. While this does add data complexity, it can also preserve hierarchy. Model-based OPC, shown in Figure 1B,

modifies edges of design features to respond to diffraction and processing effects from nearby features. It is currently implemented by a piecewise constant, or staircasing, approximation to the deformation of an edge that would respect a precise analytical model. Hierarchy information is often of necessity lost for many parts of the design, and this causes data volume issues due to both the hierarchy loss and the addition of the edge features.

Phase Shift lithography, shown in Figures 2 and 3, is another RET technique, one which increases contrast of extreme sub-wavelength features such as transistor gates.

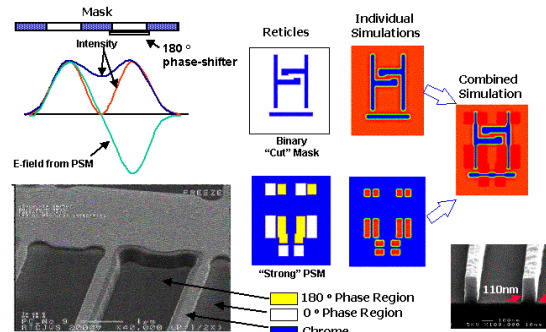


Figure 2: How Complementary PSM Works

Figure 2 illustrates the principle, in which exposure light on either side of a gate is transmitted through the mask such that on one side the light phase is not modified, but on the other it is shifted by 180 degrees. There are many strategies for the implementation of this basic idea of using destructive interference to improve contrast, manufacturability, and consequent channel length variability. The important constraint from a physical design perspective is that, given that only two “colors” are available, this “coloring” of the plane to permit gate-patterning leads to constraints in the physical design geometry. In addition, two exposures are required, one to pattern the gates, and the other to produce non-gate (e.g. routing) polysilicon lines.




Figure 3: Phase Shift – Simple Idea, Complex EDA Challenge

Figure 3 illustrates the complexity of actually creating all the shapes required to implement this concept, where the complexity

is generated not only from geometry, but also in the need to account for mask alignment tolerance during the two exposures.

Use of “scattering bars” or “assist features” is a competing technique that increases the contrast and depth-of-focus of sub-resolution gate printing by adding sub-minimum resolution features to either side of a feature. Both this and phase shift drive manufacturability and inspection challenges in mask making. Phase shift also fundamentally constrains physical design methodology [3;5].

“Tiling” to improve planarity during planarization processing is illustrated in Figure 4.




Figure 4: Model-Based Tiling

Rule-based tiling is implemented by adding new shapes to areas which at a given layer are sparsely populated, using rules to ensure that the new shapes or “tiles” do not short to design features. Model-based tiling uses an optimization technique to place varying densities of tiles in different regions to optimize modeled planarity. [6; 9] In this case, the model that drives the optimization strategy is derived from a model for the polish process. These extra shapes are typically randomly placed and consequently cause a large but necessary increase in data volume. These added data elements compromise mask inspection unnecessarily, as the manufacturability is not sensitive to small numbers of defects in the tiles.

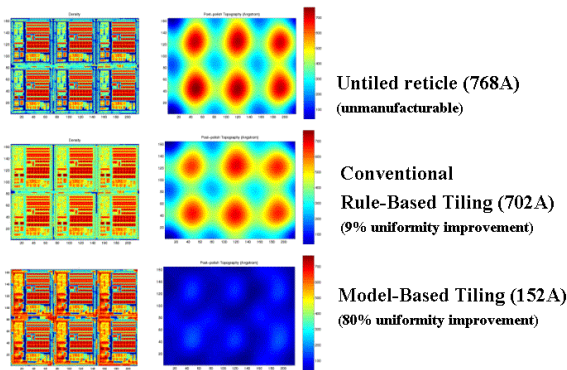


Figure 5: Rule-Based vs. Model-Based Tiling

Figure 5 shows the large benefit of model-based vs. rule-based tiling. This result is a simulation of the polish process for no tiling, rule-based tiling, and model-based tiling.

In Figure 6, we show that the implementation of RET in design results in the insertion of many model or process data driven manipulations to change design tape-out GDSII to mask shop ready GDSII. The increase in complexity of the mask data prep process can be substantial and error-prone, as different tools are often used for different variants of RET, with the mask data prep group filling the integration role.

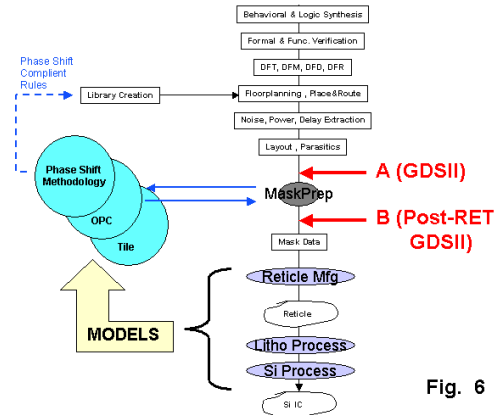


Fig. 6

Figure 6: Insertion Point for RET in Process Flow

Figure 7 illustrates that the number of these design data manipulations grows in number and innate complexity as new generations are implemented.

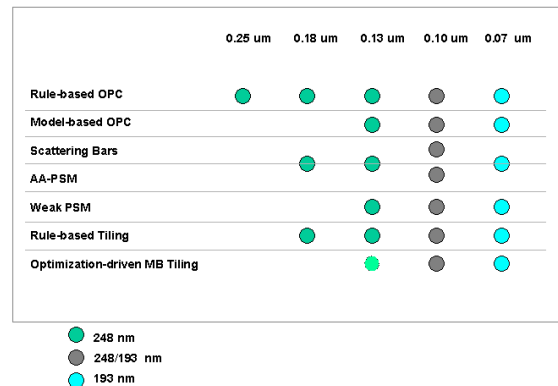


Figure 7: RET Methods Pervasively Expand as Technology Nodes Evolve “Prototypical” Scenarios

3. CONSEQUENCES OF RET IMPLEMENTATION

The improvement in manufacturability of Si in the subwavelength regime, as described above, comes at a price. Three major issues that arise are mask data preparation resources (storage, computation capacity, data preparation cycle time, etc.); mask manufacturability, including inspection; and verification technologies which guarantee design integrity between tapeout and mask data release.

Figure 8 is an estimate of several resource metrics, such as data size/layer and the number of layers subjected to RET manipulations, number of CPU's required to process data from point "A" to point "B" (shown in Figure 6), number of hours to process the mask data, and the time it takes to produce and inspect a critical gate level mask in the mask shop. These normalized resources are estimated from our own and other companies' informal communications. The main message is that 1 1/2 to two orders of magnitude expansion of resources has occurred in the last three years.




Figure 8: Mask Data Prep Resources Driven by RET

Another issue is mask manufacturability, which results from the need to manufacture on the reticle a very large number of sub-critical dimension shapes.




Figure 9: Mask Metrology Dimensions for RET Implementation Compared to ITRS Roadmap

Figure 9 is taken in part from the Semiconductor Industry Association ITRS roadmap (International Technology Roadmap for Semiconductors), and plots the mask minimum dimension vs. time, which is taken by the ITRS to correspond to about 1/2 of the CD at the mask level. Consequently, if Lpoly is 250nm on wafer, then it is 1000 nm on mask (4X stepper), and the assumption is that 500nm features on the mask can represent the small OPC features at the corners of shapes. In fact, we are discovering as an industry that another factor-of-two smaller feature is needed on the mask to fine-tune the lithographic process distortions.

Industry productivity will not continue at the historical rate unless we can reverse these trends.

4. WORKING TOWARD SOLUTIONS

Fortunately, there are several approaches to mitigate the issues that arise from RET proliferation. Some approaches are simply resource enhancement, e.g. using parallelization and more processors. Others are represented by more fundamental changes in algorithms and data representations. These strategies leverage design hierarchy or lessen the issues by restricting RET or mask inspection to critical areas only. In the following sections, we examine a few of the more promising ideas from both the mask data prep and the mask fabrication sides.

4.1 Mask Data Prep Solutions

4.1.1 GDS Replacement

GDSII has become a de facto standard data format that is widely used as a design interchange representation, but is rarely used for actual data manipulation. Its definition contains some ambiguities, and there is no authority to resolve those ambiguities and support its continued adaptation to the rapid growth and changes in the industry it serves. Most commercial EDA vendors have developed their own data formats, internal to their tools, and use ever-increasing resources for translating and storing GDSII data into these formats and translating the results back out to GDSII. While one-size-fits-all may be impractical, one-size-fits-none seems to be the trend for GDSII.

Significant reductions in data handling resources might be realized if the industry had a well thought out, widely accepted, standard for the representation of layout data. Such a representation would permit the indexing of physical design components in a manner that respects the ultimate physical – as opposed to original system – design hierarchy. Functionality and connectivity drive the latter, while final physical design geometry drives the former. Optimally, new tools would be available to transform system design hierarchy into this new physical design hierarchy.

4.1.2 Parallelization

Lacking a fundamentally new hierarchical physical design methodology, we can resort to the use of parallel computation. One approach to keeping the data bloat from significantly increasing cycle time is the use of parallel or distributed computing methods for large RET and other mask data preparation jobs.

The most straightforward parallelization strategy is one in which the area of the die is divided into grid regions, and a different node processes each region. In this case some overlap among spatial regions is required, and the overlap requires additional special processing to ensure that the region boundaries are handled properly. This simple approach however, while permitting many regions to be corrected simultaneously, has the deficiency that data size grows due to flattening of the hierarchy. Even so, by treating repeating elements such as memory cells in a special manner, this approach can be implemented in a more efficient and sophisticated manner.

Another speculative approach would be to reassemble the design after such parallel processing, combined with intelligent OPC has made compute time more efficient. If one could develop an efficient algorithm or heuristic to recapture hierarchy that remains after RET, data file sizes could be much reduced, significantly decreasing data storage investment. A high payoff research project

in hierarchy recapture would reap significant benefits in the investment needed to store product data. Such a heuristic was described in the past for solving data size issues in vector scan electron beam lithography [10], and that approach might be incorporated into such a solution. In the electron beam lithography case, compaction of orders of magnitude were achieved in affordable computation time.

Typically EDA vendors attempting to leverage parallelization have used one particular approach, tightly coupled multithreading. This approach seems chosen more to utilize the hardware most readily available to the most users (multi-processor servers with large shared memory) rather than to leverage computational strategies that are in principle most efficient.

Both loosely multithreaded and distributed processing approaches rely on separating the problem into many smaller ones that use minimal shared data or message passing architectures.

However, loosely multithreaded distributed processing that can be done on dedicated clusters of relatively inexpensive machines (or can even time-share on existing networks of workstations and desktop computers) is not available in most commercial tools, nor is it consistent with their conventional license arrangements.

4.2 Mask Fabrication Solutions

4.2.1 Mask Process Modeling

Most model-based RET tools on the market today do not explicitly capture in their models those aberrations introduced in the mask fabrication process. Those effects are folded into a generic model that proceeds directly from design to wafer. An intriguing idea for optimizing RET results is the use of rule and model-based tools to correct for sizing non-linearity, writer asymmetries, and other systematic errors in the mask-making process itself. In such a flow, tools deconvolve the two effects by including the mask-making process explicitly. This goal is illustrated in Figure 6, which shows separate models of mask writing, imaging on wafer, and Si processing as independent inputs to RET algorithms. Wafer process effects are backed out first. Following that, models and/or rules based on a characterization of the mask-making process itself extend the correction back through to design.

Current model-based lithography approaches are sufficient for capturing and correcting for local non-linearity, but if one wants to account for chrome-loading effects, writer asymmetries, and other non-local or non-uniform effects, new approaches are needed. Chrome-loading effects can be addressed using a number of process and model-based approaches and there are tools coming on the market now to address and correct for lens aberration and asymmetries in the mask writer itself.

As we move into a regime incorporating more and more process information into our models, the natural place for the characterization and modeling of mask-induced aberration is the mask shop itself. As envisioned, each mask shop would be responsible for characterizing and modeling its mask process and releasing data and models to customers, possibly using models supplied by equipment and chemical suppliers as a starting point. The cost for generating and analyzing the data and building the models could then be amortized over the customer base and passed on in cost per mask. For this to be effective, standard

characterization and model formats would need to be adopted across a sufficiently broad spectrum of the industry. Of course, even this sort of precision does not come without a price: it ties RET rather closely to a specific mask-making process within a specific mask shop. Unfortunately, that may well be the general trend for mask and wafer process alike.

4.2.2 Improvements in Die-To-Database Inspection

More precise RET is not the only benefit of modeling the mask-making process explicitly. A common approach to identification of defects during mask fabrication is die-to-database inspection, wherein the die on the mask is compared to the post-RET design data. Typically today, tools must be desensitized or put into special modes in order to ensure that intentional RET manipulations of design data, e.g. addition of small OPC features, do not get flagged as defects. This creates its own problems, because often some intentional RET still gets flagged as a defect or, conversely, true defects go undetected as a result of the desensitizing.

This situation arises because what we're comparing the mask to is not really what we ultimately want to achieve. Instead, what we really want is to ensure that the mask reflects as accurately as possible the patterning required to render the original design onto silicon. Given that, what we really should be comparing the reticle to is what we actually want to see on it. We can do this by convolving the post-RET design data with our mask model and using the result as the basis for comparison, in effect removing the mask-making component from the equation.

Such a method, proposed by Kling et al [11] allows for a more sensitive die-to-database defect inspection. This method fundamentally consists of altering the pattern generator database to create an inspection database. Many RET features are extremely small assist features, such as serifs and scattering bars. The pattern fidelity of the pattern generator process for these features is not perfect. Structures that were supposed to be sharp will become rounded and the area of the assist features may be vastly different from the design data. Using the above method, a filtering function is applied to the pattern generator database to create an inspection database that will approximate the figures expected to print on the reticle. Essentially, the difference between the reticle inspection database and the pattern generator database is the process bias and loss of data feature pattern fidelity that occurs during reticle fabrication. Because this inspection database will have features more similar to those actually printed on the reticle, there will be less false defects observed during the die to database inspection, and hence less tool desensitization required. This in turn will result in a more defect-sensitive inspection and ultimately improved reticle quality. In addition, with the reduced number of false defects, less operator time will be required for defect classification, leading to improved manufacturing cycle time.

4.2.3 Mask Layout

The layout of the mask can be optimized for manufacturability by utilizing a multi-die reticle layout, as opposed to a single die layout. Figure 10 shows an example of each type of mask layout. In the single die reticle there are no repeating or arrayed cell units in the reticle field. In the multi-die layout a single die is arrayed repeatedly in the reticle field. In the mask shop the former would require a die to database inspection, while the latter would lend itself to a die-to-die inspection mode.

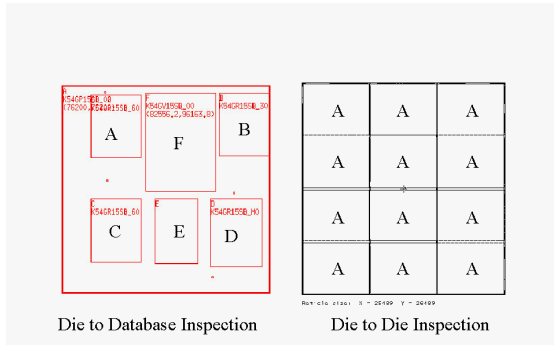


Figure 10: Die-to-Database vs. Die-to-Die Inspection




Figure 11: Global MBOPC Increases Vertex Count

Traditionally, new introductions of high-end reticle inspection tools into the reticle market are capable of operating only in die-to-die inspection mode. Implementation of die-to-database algorithms for new tools can lag anywhere from 6 months to 1 year. Mask sets consisting of a single die unit must be inspected using a previous generation inspection tool, and hence receive a less sensitive inspection. As discussed earlier, the introduction of RET into these reticle designs typically requires desensitization of the die to database inspection due to the occurrence of false or nuisance defects. In some cases the instances of false defects may in fact not simply result in a desensitized inspection, but require the pattern inspection to be waived entirely. In addition, to perform a desensitized inspection each pattern file in the single die layout must be individually prepared for the inspection tool for review. This leads to increased cycle time at the inspection stage. Given all of this, the conclusion is clear: mask layout designs utilizing a multi-die layout receive the highest quality inspection in the least amount of cycle time.

4.2.4 Mask Manufacturability

One manufacturability challenge the mask maker faces in this era of increased use of RET techniques is the sheer explosion of data volume. Write times associated with aggressive RET are often double or even triple conventional write times. Not only are write times increased, but memory allocation errors are more prevalent. The data preparation time for the inspection job setup likewise increases significantly, sometimes up to a factor of ten. Inspection verification time also goes up due to the increase in the number of false defects observed. Overall, the rise in data volume associated with RET techniques alone leads to lengthening of cycle time in the mask making process.

The increased application of RET techniques clearly points out that wafer lithographers are applying sub wavelength lithography in a regime of declining printing contrast so that reticle CD errors do not transfer linearly to the wafer. That is, errors on the mask are not reduced by the same reduction factor as in the stepper. Instead, in sub-wavelength lithography these errors are enhanced and result in a greater error contribution on the printed wafer. Hence, the CD abnormalities, missing Cr or clear defects, corner rounding, phase errors and actinic residuals once thought to be in the noise level are now becoming critical issues in both the mask shop and in wafer fabs. This is driving the mask shops to place an even greater emphasis on CD metrology, defect inspection and repair. Even now, the more stringent specifications for defects, CD, and phase control required by the customer are leading to increased cycle time and driving the mask maker to invest additional capital into the mask shop.

As stated earlier, the implementation of model-based OPC increases the number of edge features and causes loss of data hierarchy, both of which lead to increased data volumes. Smarter RET strategies are required that are more mask manufacturable and less prone to data explosion. In particular, the suppliers of model-based OPC packages may currently be too focused on achieving the perfect physical model with their algorithms and should now consider the mask manufacturability implications of their models. The users, both the mask maker and wafer lithography engineer, must answer the question: are all these added features and small jogs really necessary?

Figure 11 illustrates a typical OPC correction for a tee shape in close proximity to a larger shape.

One approach toward a smarter OPC strategy would be to use model-based OPC only when necessary/critical. For all other OPC corrections, we could still use rule-based OPC, which is better at preserving data hierarchy and often has smaller vertex count, hence reducing data volume explosion.

Figure 12 demonstrates how our previous example might look. In this figure, model-based OPC is used only where there is a minimum space between critical features. In the area where the minimum space is relaxed, rule-based OPC is sufficient.

Furthermore, rules that do not apply OPC serifs at all in non-critical regions, and use of shape extensions where allowable as opposed to hammerheads or serifs further reduce vertex count. The concurrent use of all these strategies allows for the more exhaustive checking and correction of features to be made only when necessary, with the net result of decreased data volumes. For example, in this case the corrected shape has 20 vertices in Fig. 12 vs. 44 vertices in Fig. 11 – a 50% saving.

In order to make this approach successful, the lithography engineer must verify that in fact the OPC applied improves image fidelity and leads to improved electrical data and product yields. A major point is that the optical effects of various elements of the lithography system are becoming increasingly interdependent as the process margin narrows with decreasing contrast. Mask and fab engineers, as well as software suppliers, will need to work closely going forward to develop production worthy RET techniques and processes.




Figure 12: Use of Selective MBOPC and Intelligent RBOPC Leads to Reduced Vertex Count

5. SUMMARY

Despite the growing challenges in photolithography, the industry continues to push back the limits of deep sub-wavelength photolithography through conventional and revolutionary approaches. The first method, evolutionary, focuses on optimization of current processes and wringing maximal efficiency out of today's technology. Examples of this are process fine-tuning, RET implementation, mask layout strategies that allow for die-to-die inspection, and the trend toward distributed solutions for mask data prep.

Revolutionary approaches require new design tools and data representations to make RET implementation affordable. Maintenance of system design hierarchy well into the mask data prep phase is not the correct strategy and only complicates the RET process. Another example is the need to capture the mask process explicitly. This will allow us to increase final Si pattern fidelity by considering mask-patterning as a discrete component. It will also allow us to increase the accuracy of die-to-database comparisons by providing as a basis for comparison a means of projecting how post-RET design data is *expected* to look on the mask itself.

6. ACKNOWLEDGMENTS

We acknowledge our colleagues in Motorola's DigitalDNA laboratories and in the Architectures and Systems Platforms group, who implement design tools and lithography process strategies and partner with us so that our algorithmic processes and their design standard lithographic and thin film processes work in concert. In particular, we wish to acknowledge the significant experience and learning that Clyde Browning, Bernie Roman, Kevin Lucas, Ed Travis, Tareq Bustami, and Alfred Reich have brought to us.

7. REFERENCES

- [1] A. B. Kahng, "IC Layout and Manufacturability: Critical Links and Design Flow Implications," Proc. IEEE Intl. Conf. on VLSI Design, Jan. 1999, pp. 100-105.
- [2] A. B. Kahng and Y. C. Pati, "Subwavelength Optical Lithography: Challenges and Impact on Physical Design," Proc. ACM Intl. Symp. on Physical Design, April 1999, pp. 112-119.
- [3] A. B. Kahng, H. Wang and A. Zelikovsky, "Automated Layout and Phase Assignment Techniques for Dark Field Alternating PSM," Proc. 18th BACUS Symposium on Photomask Technology and Management, September 1998, pp. 222-231.
- [4] P. Berman, A. B. Kahng, D. Vidhani, H. Wang and A. Zelikovsky, "Optimal Phase Conflict Removal for Layout of Dark Field Alternating Phase Shifting Masks," Proc. ACM Intl. Symp. on Physical Design, April 1999, pp. 121-126.
- [5] A. B. Kahng, S. Vaya and A. Zelikovsky, "New Graph Bipartitions for Double-Exposure, Bright Field Alternating Phase-Shift Mask Layout," Proc. Asia and South Pacific Design Automation Conf., Jan. 2001, pp. 133-138.
- [6] Tian, R., Wong, D. and Boone, R. Model-Based Dummy Feature Placement for Oxide Chemical-Mechanical Polishing Manufacturability. Proceedings of the Design Automation Conference, (Los Angeles CA, June 2000), ACM Press, 667-670.
- [7] Y. Chen, A. B. Kahng, G. Robins and A. Zelikovsky, "Practical Iterated Fill Synthesis for CMP Uniformity," Proc. ACM/IEEE Design Automation Conf., June 2000, pp. 671-674.
- [8] Tian, R., Tang, X., Wong, D. Dummy Feature Placement for Chemical-Mechanical Polishing Uniformity in a Shallow Trench Isolation Process. Proceedings of the Design International Symposium on Physical Design, (Sibina CA, April 2001), ACM Press.
- [9] A. B. Kahng, G. Robins, A. Singh and A. Zelikovsky, "New and Exact Filling Algorithms for Layout Density Control," Proc. IEEE Intl. Conf. on VLSI Design, Jan. 1999, pp. 106-110.
- [10] W.D. Grobman, "An Overview of Pattern Data Preparation for Vector Scan Electron Beam Lithography," J. Vac. Sci. Technol., Vol. 17, Issue 5, pp. 1156-1163, 1980
- [11] Lucas, Kevin D. ; Kling, Michael E. ; Reich, Alfred J. ; Fu, Chong-Cheng ; Morrow, James. "Process for producing and inspecting a lithographic reticle and fabricating semiconductor devices using same". U.S. Patent No. 5,849,440, 1997