A Multi-objective Floorplanner for Shuttle Mask Optimization

Gang Xu[a], Ruiqi Tian[b], David Z. Pan[c], Martin D.F. Wong[d]

[a] CS Department, University of Texas at Austin, Austin, TX 78712, USA
[b]Freescale Semiconductor, Austin, TX 78721, USA
ECE Department, University of Texas at Austin, Austin, TX 78712, USA
[d]ECE Department, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
xugang@cs.utexas.edu, ruiqi.tian@freescale.com, dpan@ece.utexas.edu, mdfwong@uiuc.edu

ABSTRACT

Shrinkage of VLSI feature size and use of advanced Reticle Enhancement Technologies (RET) in manufacturing such as OPC and PSM have dramatically pushed up cost of mask. For example of a 130nm or 90nm mask set, the mask cost can easily reach one or two million US dollars. Shuttle mask is an effective method to share the mask cost by putting different chips on the same mask. Shuttle mask floorplanning is a key step to pack these chips according to certain objectives and constraints related to cost, yield, and manufacturability. In this paper, we present a simulated annealing based floorplanner to solve the shuttle mask floorplanning problem with multiple optimization objectives and constraints. We will consider area minimization, density optimization (for manufacturability enhancement with CMP), wafer utilization maximization, die-to-die inspection constraint, die orientation constraint and their combinations. A nice property of our floorplanner is that it can be easily adapted to different cost models of mask and wafer manufacturing. Experiments on industry data show promising results.

Keywords: Shuttle mask, floorplanning, simulated annealing, multi-objective.

1. INTRODUCTION

The past decades have witnessed VLSI feature size continuously shrinking into the deep sub-micron regime under governance of Moore's law, which brings many challenges to VLSI manufacturing. Among these challenges sub-wavelength lithography is one of the most important. Nowadays integrated circuits at 130nm and 90nm technology nodes are typically fabricated by 193nm optical lithography tools. It is necessary to use advanced Reticle Enhancement Technologies (RET) such as optical proximity correction (OPC) and phase shifting mask (PSM) to ensure the image of these circuits can be printed on the wafer as close to the layout of the original design as possible.

However, shrinking VLSI feature size and use of RET have dramatically pushed up cost of mask, because the complexity and volume of data written to the mask are quickly increasing. It needs more blank materials, longer writing time, and more accurate inspection and repair processes to make such a set of masks. As a result, a set of masks at 130nm node may cost up to one million US dollars, and the number is estimated to double at 90nm node. For prototype and low volume designs, the mask cost is prohibitively expensive because it cannot be amortized over the volume.

Shuttle mask is an effective method to share the mask cost among these prototypes and low volume designs by putting different chips on the same mask, as shown in Figure 1. Although use of the shuttle mask leads to extra overhead such as the growing size of data file, higher complexity of the mask, and extra time/expense introduced by cutting different chips from wafers, the overhead is still much lower than the total cost to make multiple sets of masks. For this reason, use of shuttle mask becomes popular. Chip designers can access the shuttle service provided by major foundries and manufacturers such as TSMC and IBM.

A key step to make a shuttle mask is to efficiently pack different chips together according to certain optimization objectives and constraints. This step is known as shuttle mask floorplanning. Unlike traditional floorplanning in circuit design whose major objective is to minimize of area and total wire length of a chip, shuttle mask floorplanning aims at multiple objectives and constraints related to cost, manufacturability and yield. These objectives and constraints include:

(1) area minimization to save mask cost; (2) die-to-die inspection constraint to improve the defect inspection; (3) density optimization to improve topography variation, which is critical to yield; (4) wafer utilization to save wafer cost and chip production time; and (5) others, for example, die orientation constraint to guarantee the manufacturability.
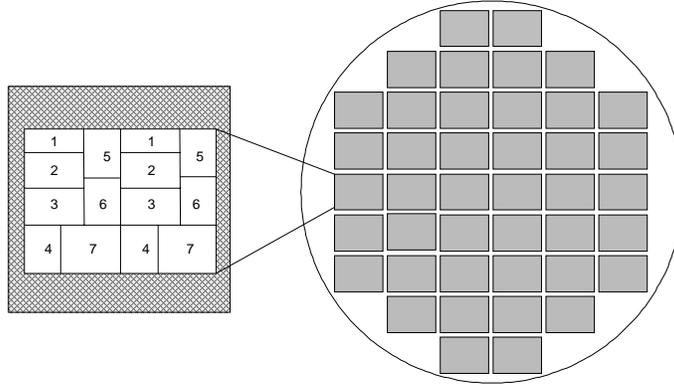
Figure 1. A shuttle mask and the projections on the wafer. Each small rectangle with a number marked represents a chip.

Because of the novelty of these objectives and constraints, shuttle mask floorplanning becomes an interesting and challenging problem and has drawn attention of researchers in EDA community [1, 2, 3, 4, 5]. To our knowledge, Chen and Lynn published the earliest paper on shuttle mask floorplanning in early 2003 [1]. They only considered the area minimization objective that was actually a simplified version of classical floorplanning problem. Later Xu et al. [2] studied the minimum area slicing floorplan problem with die-to-die inspection constraint that is important for defect inspection on mask. Around the same time appeared Andersson et al's work [3], in which they used a "grid" floorplan which tried to handle both the area minimization and wafer utilization maximization objectives. Afterwards, Kahng et al. [4] also studied simultaneous area minimization and wafer utilization maximization. They used non-slicing floorplan, and assumed the chip with varying width and height. Later, Kahng et al [5] considered another formulation of area minimization and wafer utilization maximization in which wafer utilization was represented as a constraint, rather than an objective. In this paper they revisited "grid" floorplan and claimed improvements on experimental results over their previous work.

In this paper, we present a simulated annealing based floorplanner solving the shuttle mask floorplanning problem with multiple optimization objectives and constraints. We will consider area minimization, density optimization (for manufacturability enhancement with CMP), wafer utilization maximization, die-to-die inspection constraint, die orientation constraint, and their combinations. A nice property of our floorplanner is that it can be easily adapted to different cost models of mask and wafer manufacturing. For example, different weight parameters can be assigned to these objectives and constraints to reflect preferences on area, density distribution, inspection, wafer utilization, and their combination. Experiments on industry data show promising results.

The paper is organized as follows. In Section 2, we investigate different objectives and constraints of shuttle mask floorplanning. Section 3 shows how to handle these objectives and constraints by using certain techniques and algorithms in our floorplanner. Section 4 demonstrates the experimental results. Conclusions are in Section 5.

2. OBJECTIVES AND CONSTRAINTS

In this section, we will look at the following objectives and constraints in shuttle mask floorplanning: (1) area minimization; (2) wafer utilization; (3) density optimization; and (4) die-to-die inspection. The investigation on each objective or constraint is combined with review of the related previous work. Background information such as wafer cutting and topography variation for CMP will also be presented.

2.1 Area Minimization

Area minimization is a natural and important objective for shuttle mask floorplanning. Given a set of chips, a compact shuttle mask floorplan will lead to more projections on the wafer; it also allows more chips to be put on the shuttle as long as these chips can be packed in the frame of maximum printing field. Cost is reduced in both cases,

In addition, area minimization is probably the easiest objective to optimize. If we only consider area minimization, shuttle mask floorplanning is actually a simplified version of classical floorplanning problem in circuit design, because chips are all rectangles without wire connection among them. Although optimally packing rectangles is proved to be an NP-hard problem, there have been efficient heuristics to solve it near-optimally. As mentioned in Introduction, Chen et al's work [1] mainly focuses on area minimization of shuttle mask floorplan. Their floorplan can either be slicing [6] or non-slicing [7]. A bottom-left fill algorithm is proposed to find near-optimal packing. Although they didn't report how compact the floorplan generated by their algorithm was, from the example diagram it seems the result can be improved remarkably.

2.2 Wafer Utilization Maximization

Shuttle mask floorplanning becomes novel, interesting and challenging only when the primary objective area minimization is combined with other objectives or constraints specific in manufacturing process, for example, wafer utilization. Most of the times chips on wafer are cut out by a cutting saw that traverses the whole wafer horizontally or vertically. Because chips on shuttle mask have different sizes and shapes, cutting out one chip may destroy another, as shown in Figure 2. How to pack chips on shuttle mask such that as many chips as possible can be cut out from the wafer thus becomes another important objective, i.e., wafer utilization.
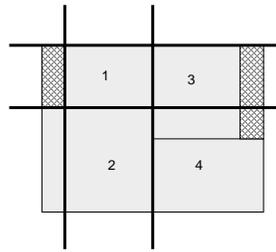
Figure 2. Wafer cutting. Cutting out chip 1 will destroy chip 2 and 3.

Obviously, the less the number of chips to be destroyed, the higher the wafer utilization. Consider two chips A and B Their positions on the shuttle mask will determine whether cutting out one will destroy another. The position of a chip can be represented by a pair of intervals: $([L_x, U_x], [L_y, U_y])$, where L_x and L_y are the coordinates of the left-down corner; U_x and U_y are the coordinates of the right-up corner. It is straightforward that A and B can be cut out simultaneously without destroying each other if and only if (1) $[L_x^A, U_x^A] \cap [L_x^B, U_x^B] = [L_y^A, U_y^A] \cap [L_y^B, U_y^B] = \varnothing$, which means the chips' projections on x-axis and y-axis are not overlapped, or (2) $[L_x^A, U_x^A] = [L_x^B, U_x^B]$, or (3) $[L_y^A, U_y^A] = [L_y^B, U_y^B]$. For any pair of chip not satisfy the above conditions, we call them are in horizontal or vertical conflict.

[3] considered a "grid" floorplan for shuttle mask that reduced the possibility of chips' overlap on x-axis or y-axisas shown in Figure 3. A shuttle mask will be partitioned into a grid first. Then each cell in the grid will be assigned to a chip. The grid structure prevents chips neither in the same row nor in the same column from conflicting. They studied the area minimization problem of such grid floorplan and its variants, and suggested a series of approximation algorithms. Their approach looks interesting theoretically. However, they didn't explicitly evaluate the wafer utilization of a grid floorplan, and no experimental results were reported to show the effectiveness of their approach.

[4] was the first paper that explicitly evaluated the wafer utilization (yield in their paper). They used the non-slicing floorplan to represent shuttle mask (multi-project reticle in their paper). Given a shuttle mask floorplan, they defined H-conflict and V-conflict graphs to indicate the conflict relation between any two chips on the mask. An example is shown

in Figure 4. A maximal independent set in H-conflict graph corresponds to a set of chips that can be horizontally cut at the same time. Assuming reticle projections were arranged as a RxT matrix, they assigned an independent set of H-conflict graph (ISH) to each row and an independent set of V-conflict graph (ISV) to each column. For the reticle projection at (i, j), the intersection of i-th row's ISH and j-th column's ISV would determine which chips to be cut out. Such an assignment of ISH and ISV was called a "wafer dicing plan", as shown in Figure 5. Cost of a dicing plan was defined as the minimum number of wafers required to get the volume of all chips. Given a shuttle mask floorplan, they proposed a non-linear programming formulation and several integer linear programming formulation to find an optimal dicing plan, and a simulated annealing heuristic to quickly find the near-optimal solution. Cost of a shuttle mask floorplan was the weighted combination of area and cost of its dicing plan.
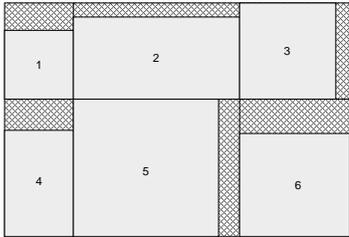
Figure 3. A grid floorplan.
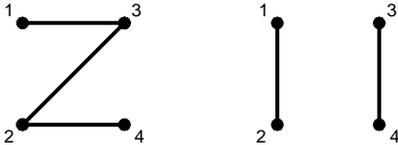
Figure 4. The H-conflict graph (left) and V-conflict graph (right) for the floorplan in Figure 2.
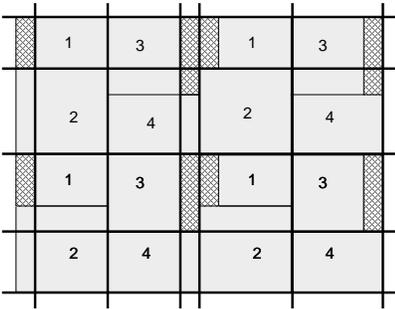
Figure 5. A wafer dicing plan. The reticle is shown in Figure 2. Assume its projections on wafer compose of a 2x2 matrix. Maximal independent sets {1,2} and {3,4} of H-conflict graph in Figure 4 are assigned to the first and the second row; .maximal independent sets {1,3} and {2,4} of V-conflict graph are assigned to the first and the second column.

However, a major problem appeared when they tried to reduce the conflict so as to improve the cost of dicing floorplan. In their paper, they made a questionable assumption that a chip could be cut out with variable margins from different reticle projections, as shown in Figure 6. Such variable margins will inevitably result in difficulties in packaging.
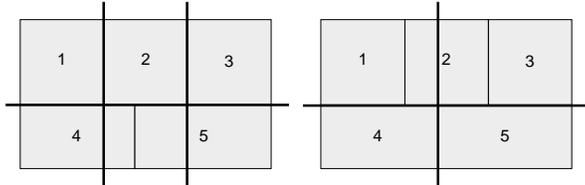
Figure 6. With variable margin assumption, chip 1 now can be cut out together with either {2, 3} or {4,5}. However, the two copies of chip 1 will have different size, for in the latter case chip 1 will have an extra margin.

Later Kahng et al revisited the grid floorplan [5]. This time they removed the assumption of variable margins. In addition, wafer utilization appeared as a constraint, instead of an objective. Their problem was formulated as finding a grid floorplan with minimum area such that the wafer utilization was no less than certain value. They used branch-and-bound search to find the optimal solution. Experimental results showed improvements on both area and wafer utilization over [4].

2.3 Density Optimization

The two objectives we have discussed so far are both about cost. As a matter of fact, different shuttle mask floorplans will also impact one of the most important manufacturability factors: chemical-mechanical polishing (CMP) for oxide planarization. CMP polishing process is required for planarization of wafer surface, as lack of planarity may lead to

insufficient depth of focus [8]. The topography variation of wafer surface after CMP is closely related to the feature density distribution [9, 10]. Given the feature density distribution, a 2-D low-pass filter model proposed by Ouma et al is widely used to estimate the post-CMP topography variation [11]. In this model, circuit layout is abstracted to a grid of small squares called cells. Each cell has a feature density defined as the total area covered by the feature in the cell divided by the area of the whole cell. With long enough polishing time, the oxide thickness distribution after CMP can be modeled as: $Z = C + k \cdot IDFT[DFT[d(i,j) \cdot DFT[f(i,j)]]]$, where $f(i,j)$ is a weighing function serving as a 2-D low-pass filter, $d(i,j)$ is the local oxide feature density, C and k are parameters independent to the feature density, and DFT and $IDFT$ refer to Discrete Fourier Transformation and its inverse operation. Tian et al [12] gave the following approximation of f(x,y): $f(x,y) \approx c_0 \exp[c_1(x^2 + y^2)^{c_2}]$, where constants c_0, c_1 and c_2 are determined by the process.

Post-CMP Topography variation can be reduced by inserting dummy features into the circuit layout to change the feature density. Besides feature density, each cell also has a capacity indicating the maximum amount of dummy feature to be added into the cell. Therefore, the feature density distribution $d(i,j)$ is actually a variable. Density optimization for minimum topography variation can be formulated as a linear programming problem and solved optimally. See [12] for more details.

As different shuttle mask floorplans result in different feature density and capacity distributions, it naturally forms a floorplanning problem how to pack chips on the shuttle mask such that the area and topography variation are minimized. We will also consider this problem in our floorplanner.

2.4 Die-to-die Inspection Constraint

Defects may appear on the mask during the process of mask making. A defect is any flaw distorting the mask image from the original design, including extra chrome region such as chrome spots and chrome bridging between geometry, or extra clear areas such as pin holes and clear extensions. In order to guarantee good yield, IC manufacturers require defects on the mask to be controlled very well. Therefore, defects on the mask must be carefully inspected and repaired before the mask set is delivered.

Die-to-die and die-to-database are two kinds of mask inspection technique. Die-to-die inspection compares two identical chip images at different positions on the mask. In contrast, die-to-database compares the chip image on the mask and the computer-generated image stored in the database. Die-to-die inspection has higher sensitivity to detect defects, as the defect is unlikely to appear twice at the same location of the chip images. However, chips under die-to-die inspection must appear pair-wise and be aligned horizontally or vertically on the mask for the sake of the requirement set by the inspection machine.

[2] studied area minimization of the shuttle mask floorplan with such a die-to-die inspection constraint, and proposed a merging method to handle the constraint. In their method, a pair of identical chips to be aligned is first merged into a superblock which has multiple shapes. The shuttle mask floorplanning problem with die-to-die inspection constraint is then transformed to an unconstrained floorplanning problem by replacing each pair of identical chips to be aligned with a superblock, as shown in Figure 7. As we mentioned in 2.1, the unconstrained floorplanning problem can be solved efficiently. In addition, by applying the merging method, the die-to-die constraint problem is easy to incorporate with other objectives besides area minimization.
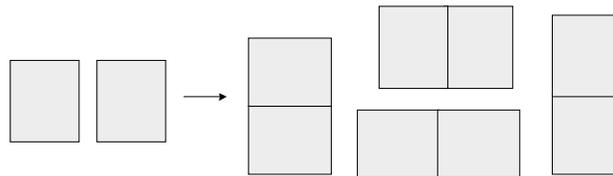
Figure 7. Multiple shapes of a super block.

As the technology continues to advance, there might be other constraints in shuttle mask floorplanning in the future. For example, the transistor orientation may become strict due to manufacturability reason [13], which will impose an

orientation constraint on each chip. Although this orientation constraint is actually a friendly constraint in that it reduces the complexity of the floorplanning by prohibit possible rotations, we cannot expect other emerging constraints always to have such a nice property. Because of the variety of objectives and constraints, we believe the shuttle mask floorplanning will be kept open for further research.

3. THE MULTI-OBJECTIVE FLOORPLANNER

We have implemented a simulated annealing (SA) based floorplanner to solve the shuttle mask floorplanning problem with multiple optimization objectives and constraints, as the SA based floorplanner is proved to be very successful in solving floorplanning problems in IC design. Figure 8 shows a generic algorithm of simulated annealing search. SA search can be approximately considered as the combination of a random search at high temperature and a greedy search at low temperature. Temperature T plays the role of control parameter for smooth transition.

```
begin
    x=initial_solution
    T=initial_temperature
    while (stopping criterion is not satisfied) do
        begin
            while (not stable)
                begin
                    x'=get_random_neighbor_solution(x);
                    delta=cost(x')-cost(x)
                    if delta<0    then accept x'
                        else accept x' with probability e^{-delta/T}
                end
            update T;
        end
    output best solution
end
```

Figure 8. A generic algorithm of simulated annealing search

We choose the slicing floorplan for the shuttle mask, which is obtained by recursively cutting a rectangle into smaller rectangles. Each chip will be put in an indivisible small rectangle called basic block. The topological structure of a slicing floorplan can be elegantly represented as rooted binary tree, or equivalently, a normalized polished expression. It has a smaller solution space than the non-slicing structure for SA search. In addition, empirical evidence shows that given a minimum area floorplanning problem, the solution of slicing structure found by SA search has close area to the one of non-slicing structure. Next, we will introduce the techniques used to handle multiple objectives and constraints we have investigated.

3.1 Die-to-die Inspection and Orientation Constraints

The orientation constraint on a chip can be explained as whether a chip with height h and width w is allowed to have multiple shapes in the floorplan: w x h, or h x w. If we take the merging method in [2] to put a pair of identical chips together so as to satisfy the die-to-die constraint, the superblock of the chip pair will also have multiple shapes.

Multiple shapes of a chip can be described by its shape curve, which is defined as feasible dimensions of the block containing this chip. In a floorplan represented as a slicing tree, each leaf node will represent a basic block. The shape curve of the leaf node is determined by the shape of the chip this leaf node refers to; the shape curve of an internal node is calculated by merging its two children nodes' shape curve. The shape curve of the root will determine the minimum area of the floorplan.

A nice property of shape curve is that, by initializing the shape curve of a basic block in different ways, we can easily guarantee that a shuttle mask is feasible with the die-to-die inspection and orientation constraints. For example, suppose

we have a chip with height h and width w. We can apply 4 possible combinations of die-to-die and orientation constraints on it: (1) the chip is under orientation constraint, but free of die-to-die inspection constraint; (2) the chip is free of both orientation and die-to-die inspection constraints; (3) the chip is under constraints of both orientation and die-to-die inspection; (4) the chip is free of orientation constraint but under die-to-die constraint. Here we refer a chip is under die-to-die inspection constraint to the scenario that it must be merged with another identical chip. Figure 9 shows how the shape curve of the chip should look like in these situations. In any case, we can always reduce the die-to-die inspection or orientation constrained floorplanning problem into a unconstrained one.
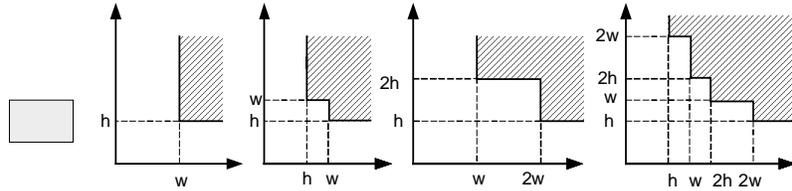
Figure 9. Different shape curves for a chip with height h and width w. From left to right, the shape curves are for single chip with and without the orientation constraint, a pair of merged chips with and without orientation constraint respectively.

3.2 Area Minimization and Wafer Utilization Maximization

Given a floorplan, our evaluation of its wafer utilization is different from that of the previous work: (1) unlike [4], we do not allow any chip to be cut out with any extra margin. For example, in Figure 6, the wafer dicing plan on the left is legal for us to cut out chip 1 while the one on the right is illegal; (2) unlike [5], we still consider wafer utilization maximization as an objective, and calculate the weighted combination of area and wafer utilization, as the combination may reflect the total cost of mask and wafer; it is a nice property of our floorplanner that it can be easily adapted to different cost models of mask and wafer manufacturing by adjusting weights of area and wafer utilization according to users' real situation; (3) unlike [4] that used the H-conflict and V-conflict graph of a floorplan separately, we use a single conflict graph which is sum of H-conflict and V-conflict graph to reflect the conflict relation among chips; (4) unlike [4] in which every wafer was cut in the same way, we may assign different dicing plans to different wafers.
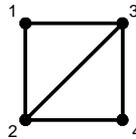
Figure 10. Conflict graph for the floorplan in Figure 2.

Consider the floorplan in Figure 2 whose H-conflict and V-conflict graph is shown in Figure 3, its conflict graph is shown in Figure 10. The optimal coloring scheme needs 3 colors, as there exist two 3-cliques in the graph but no 4-clique exists. A coloring scheme will be: {1, 4} red, {2} yellow, and {3} blue. Obviously, chips with the same color are neither H-conflict nor V-conflict. We require that for any wafer, only chips with the same color can be cut out. Three dicing plans for the floorplan in Figure 2 are shown in Figure 11.
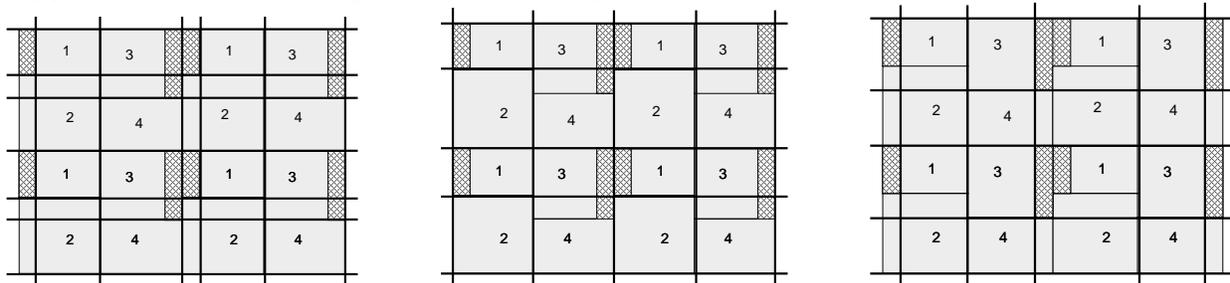
Figure 11. Dicing plans for the floorplan in Figure 2.

A quick example will show how this strategy may improve wafer utilization. Assume we use the reticle in Figure 2 to print chips on wafer. The required volume of each chip is 240. The reticle is 4x with area 100mm x 132mm. The wafer has 200mm (8-inch) diameter. These data are all typical industry value. It is simple to calculate that there are at most 4x6=24 reticle projections on a wafer. With the dicing plan in Figure 5, for each chip we can cut out 6 copies from the wafer. So we need 40 wafers to satisfy the volume requirement. However, with our dicing plans, from a wafer we can cut out 24 copies of each chip with the same color. The total number of required wafers thus becomes 240/24*3=30.

The evaluation of wafer utilization is shown in Figure 12. Given a floorplan, we construct and color the conflict graph. Then we calculate the number of reticle projections based on the sizes of the wafer and the floorplan. With the color number and the reticle projection number, we obtain the number of required wafers as wafer utilization.

```
input: a slicing floorplan F and the size of the wafer
begin
    construct the conflict graph G(F);
    color the conflict graph G(F);
    calculate the number of reticle projections on wafer;
    calculate the number of required wafers;
end
```

Figure 12. the algorithm to evaluate wafer utilization.

As graph coloring has been proved to be a NP-hard problem, we use a greedy coloring algorithms proposed by [14]. The algorithm is shown in Figure 13. Our experiment shows that this greedy algorithm is a good approximation to the optimal coloring scheme of the conflict graph.

```
input: a general graph G
begin
    sort vertices of G by degree in the descending order;
    for i = 1 to n do
        assign the lowest indexed color c to $v_i$ such that for
        any $v_j$ adjacent to $v_i$ , j<i, c is not assigned to $v_i$ yet.
end
```

Figure 13. The greedy coloring algorithm.

3.3 Area Minimization and Density Optimization

As usual, we solve area minimization and density optimization simultaneously by combining them into a single objective that is their weighted sum. The remaining question is how to evaluate the density optimality of a shuttle mask floorplan.

As we have discussed in Section 2.4, different shuttle mask floorplans may result in different density and capacity distributions that further result in different minimum topography variations (MTV). We need to find out which floorplan is optimal in topography variation after dummy features are inserted. Unfortunately, it is impractical to evaluate the MTV of a floorplan using the linear programming method in each SA search step, because a simulated annealing search may check thousands or even millions of solutions, and the linear programming method is just too costly in computation to afford. Instead, we have to take some heuristic to estimate the MTV.

A simple heuristic is to aggressively assume that the optimal amount of dummy feature to be inserted to each chip is independent of the floorplan. Before starting floorplanning, we pre-fill dummy features in each chip by using the linear programming method. With that assumption we only need two DFT and one IDFT operations in each SA search step. In this paper our floorplanner uses this heuristic, assuming each input chip has been pre-filled. Notice that this heuristic becomes accurate when the input chip has strong restriction on dummy insertion.

A more complicated heuristic is to design some predictive function to guide the floorplanner to a solution expected to have the MTV. It is critical to find a fast and accurate predictive function.

4. EXPERIMENTAL RESULTS

We implement the shuttle mask floorplanner based on the Wong-Liu floorplanner [15]. The code runs on a Linux workstation with a single P4 2.4G Hz CPU with 512K L2 cache and 1G DRAM. We derive our data set from industry shuttle masks. This data set includes 12 chips with different sizes and shapes.

Table 1 compares the quality of the floorplans found by different weighted combinations of area and wafer utilization cost. (A, W) refers to the normalized weights for area and wafer utilization. Wafer Utilization refers to the required number of wafers for each chip. # of Projections refers to the maximum number of the reticle projections on wafer. Colors refer to the number of colors to color the conflict graph of the floorplan. White Space indicates how compact the shuttle mask floorplan is. We can see the consistent trend that when the weight of wafer utilization increases, the wafer utilization is improved while the white space rate goes up. The floorplans for the best wafer utilization case and the best area case are shown in Figure 14 respectively. As the weights of area and wafer utilization are adjusted according to the different cost models of mask and wafer, the floorplan with the optimal total cost can be easily obtained from our floorplanner.

Table 1: The comparison among different weighted combinations of area and wafer utilization

(A, W)	Wafer Utilization	# of Projections	Colors	White Space
(1, 0)	60	40	5	3.74%
(1, 0.1)	50	48	5	4.84%
(1, 0.5)	40	48	4	5.12%
(1, 1)	36	40	3	9.51%
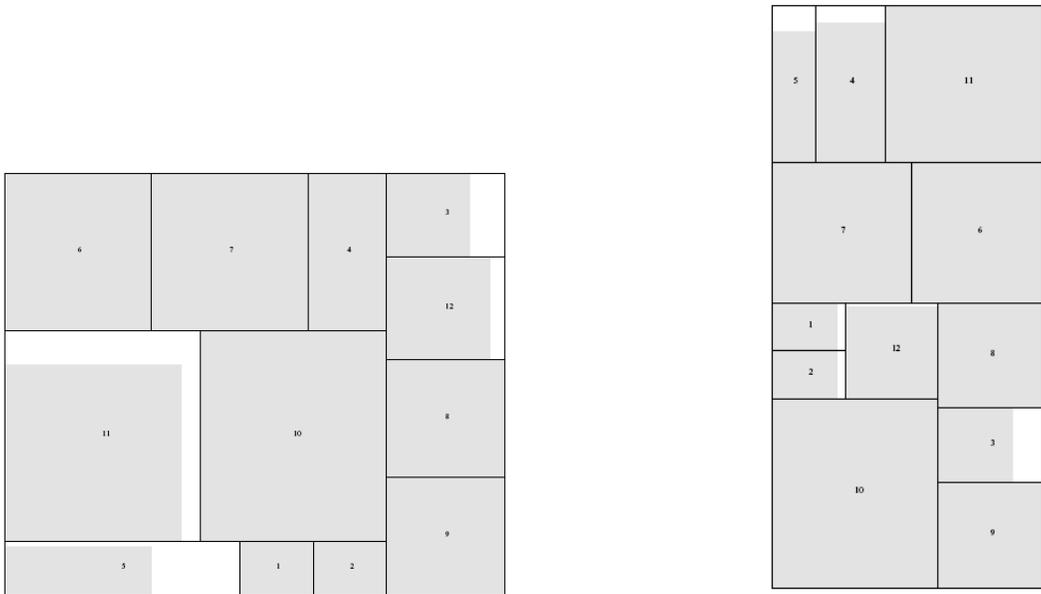
Figure 14. Floorplans for the best wafer utilization and best area.

Table 2 shows the comparison of topography variation with and without density optimization. As we can see, with consideration of density optimization, the optimal solution may have slightly larger mask area, but the post-CMP topography variation may be remarkably improved which will benefit the yield. Figure 15 shows the floorplans.

Table 2: Comparison of topography variation with and without density optimization objective

Objective	Topography variation	White space
Area+Density	1 (normalized)	5.86%
Area only	1.74	4.36%
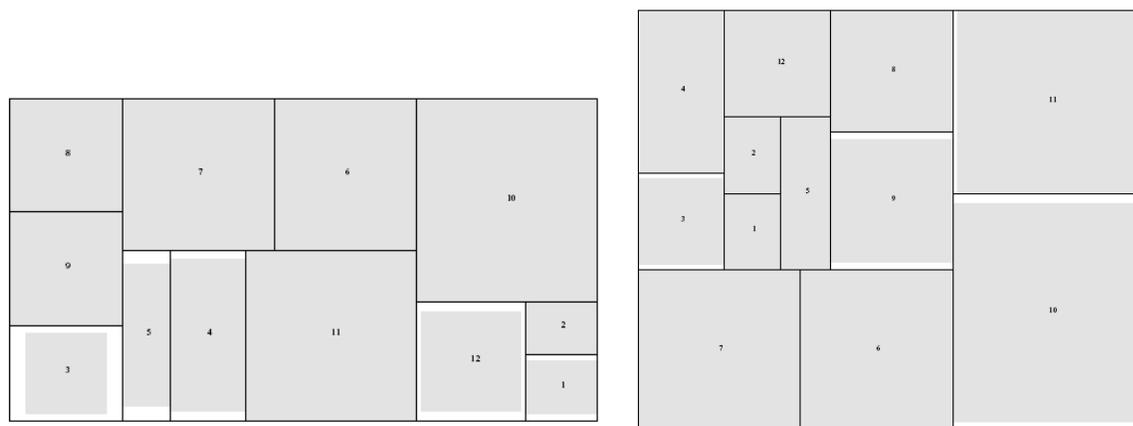
Figure 15. Floorplans with and without the objective of density optimization.

5. CONCLUSIONS

In this paper, we investigate multiple objectives and constraints related to cost, yield and manufacturability in shuttle mask floorplanning. We also present a simulated annealing based floorplanner to solve these objectives, constraints, and their combinations. Our floorplanner can be easily adapted to different cost models of mask and wafer manufacturing which may lead to different optimal solutions to different users. Experiments on industry data show promising results.

6. ACKNOWLEDGEMENT

We would like to thank Mr. Frederick Peiffer for some very helpful discussions and comments.

REFERENCE

1. S. Chen and E. C. Lynn, *Effective placement of chips on a shuttle mask*. Proceedings of SPIE, Vol 5130, pp. 681-688, 2003.
2. G. Xu, R. Tian, D. F. Wong, and A. Reich. *Shuttle mask floorplanning*. Proceedings of SPIE, Vol 5256, pp. 185-194, 2003.
3. M. Andersson, J. Gudmundsson, and C. Levcopoulos. *Chips on wafer*. Proceedings of Workshop on Algorithms and Data Structures, 2003.
4. A.B. Kahng, I. I. Mandoiu, Q. Wang, X. Xu, and A. Zelikovsky. *Multi-project reticle floorplanning and wafer dicing*. Proceedings of ISPD, 2004.
5. A.B. Kahng and S. Reda. *Reticle Floorplanning With Guaranteed Yield for Multi-Projects Wafers*. To appear in Proceedings of ICCD, 2004.
6. D.F. Wong, H.W. Leong and C.L. Liu. *Simulated annealing for VLSI design*. Kluwer Academic Publishers, 1988.
7. H. Murata, K. Fujiyoshi, S. Nakatake, and Y. Kajitani, *VLSI module placement based on rectangle-packing by the sequence pair*. IEEE Transaction on Computer-aided Design, Vol 15, No 12, pp. 1518-1524, 1996.
8. G. Nanz and L. E. Camilletti. *Modeling of Chemical-Mechanical Polishing*. IEEE Transaction on Semiconductor Manufacturing, Vol 8, No 4, pp. 382-389, 1995.

9.   S. Prasad, W. Loh, A. Kapoor, E. Chang, B. Stine, D. Boning, and J. Chung. *Statistical metrology for characterizing CMP processes*. Microelectron. Eng., Vol 33, pp. 231-240, 1997.

10. 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 Trans on Semiconductor Manufacturing, Vol 11, pp. 129-140, 1998.

11. D. Ouma, D. Boning, J. Chung, G. Shin, L. Olsen, and J. Clark. *An integrated characterization and modeling methodology for CMP dielectric planarization*. Proceedings of Int. Interconnect Technology Conf., pp. 67-69, June 1998.

12. R. Tian, D. F. Wong, and R. Boone. *Model-based dummy feature placement for oxide chemical-mechanical polishing manufacturability*. IEEE Transaction on CAD, Vol 20, pp. 902-910, 2001.

13. P. Gelsinger. *Keynote*, Proceedings of DAC, 2004.

14. D.J.A. Welsh and M.B. Powell. *An upper bound for the chromatic number of a graph and its application to timetabling problems*. Computer Journal, Vol 10, pp 85-86, 1967-68.

15. D.F. Wong and C.L. Liu, *A new algorithm for floorplan design*, Proceedings of DAC, pp. 101-107, 1986