Adopting curvilinear shapes for production ILT: challenges and opportunities

Ryan Pearman[1], Jeff Ungar[1], Nagesh Shirali[1], Abishek Shendre[1], Mairusz Niewczas[1], Leo Pang[1], Aki Fujimura[1]
[1] D2S, Inc, 4040 Moorpark Ave, #250, San Jose CA 95117

ABSTRACT

We have recently demonstrated that curvilinear shapes and multi-beam mask writing are necessary to minimize the impact of mask variability on wafer hotspots. Several key challenges and opportunities remain. We ask how we update mask-inspection rules, and how we correct for mask-process systematics for extreme ultraviolet (EUV), where the optical response must be taken into account. This paper proposes updated mask rule checks (MRC), derived from a mask variability perspective. We also demonstrate the need for MRC-aware inverse lithography technology (ILT) metrics as a pre-requisite to ensure the reticle shapes are what the wafer-side lithographer desires. Armed with a fully curvilinear ILT and mask data preparation (MDP) system, there is an opportunity to relax the restrictions on fully Manhattan designs where possible.

1. INTRODUCTION

Inverse lithography technology (ILT) has long been seen as the best way to maximize the overall process window for immersion lithography. However, since ILT natively generates purely curvilinear shapes, which until recently could not be manufactured in high volumes, the industry has searched for ways to accommodate curvilinear ILT mask shapes into a Manhattan representation. A seminal study showed that one way to approach the mask manufacturability problem was by examining how simplifying mask shapes via Manhattanization minimally affected the overall process window [1]. However, more recent studies have demonstrated that up to a 75% reduction in the overall process window can be obtained by adopting curvilinear over Manhattan shapes for ILT when mask variability is considered, for both 193i as well as EUV technologies [2,3].

Until recently, the only high-volume mask manufacturing tool has been the variable-shaped-beam (VSB) eBeam mask writer, which exposes one rectangle (or unusually, a right triangle) of a pattern at a time. To make mask shapes compatible with this system, the input data (generally the output of optical proximity correction – OPC – or ILT) must be broken into a set of rectangles ("fractured") before the data can be sent to the VSB writer. For mask patterns with curvilinear shapes, it is necessary to fracture the data into a large number of small rectangles to accurately replicate the shape. This causes a dramatic increase in the write times for curvilinear shapes. For practical mask manufacturing, mask write time needs to be limited to 24 hours or less to minimize overall defects [4]. For budgetary reasons, mask turnaround time (TAT) is optimized for 10 hours or less at the major mask shops [5]. Therefore, curvilinear masks cannot be manufactured practically in a high-volume production environment using conventional fracturing methods.

In recent years, new mask writing tools arrived on the market [6,7]. Multi-beam mask writers are now in production use and have as a key advantage the ability to write any mask shape without runtime cost. This means that curvilinear shapes are now just as manufacturable as Manhattan shapes are using multi-beam tools. These multi-beam tools rasterize an image of the desired mask pattern into a greyscale bitmap of varying doses. This bitmap is processed, and each pixel is assigned to a dose delivered by a small number of individual beamlets (between 10nm to 20nm squares in size).

Current MRC rules govern the minimum feature sizes for Manhattan shapes, but no industry standard exists for their curvilinear analog. This paper will propose a simple curvilinear MRC scheme from a mask variability perspective, based on a review of the existing Manhattan MRCs and what they really try to constrain.

As with Manhattan features, if the curvilinear features are small, they will not print to size. Therefore, mask process correction (MPC) needs to be applied to these curvilinear features. Existing methods to perform MPC for Manhattan features are well known; it is computationally more complex to perform MPC on curvilinear shapes. While addressing the computational complexity is outside of the scope of this paper, it will be shown that, if MRC constrains the curvilinear shapes properly, it improves the convergence properties of MPC significantly.

2. REVIEW OF MANHATTAN MRC

In the distant past of mask manufacturing, before OPC was commonplace, there were typically two simple rules which constrained mask shapes: the minimum feature width and minimum space between features. Since the features were large, these limits were not defined by a manufacturing limit, per se. Rather they were dictated by the limitations of the inspection tools of the time. In order to go to high-volume manufacturing (HVM) for production, one must be able to inspect and verify what one is manufacturing.

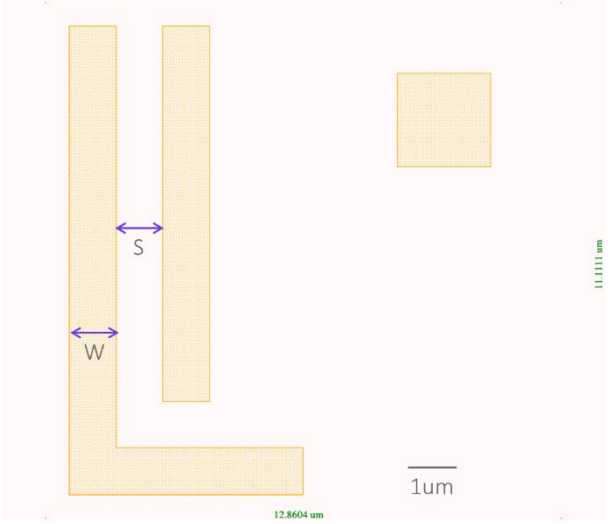

Figure 1: Simple width and space checks for large features (not shown for 2D contact)

As features shrunk over the next decade, and rule-based/model-based OPC became more common, additional line-end rules and corner-to-corner rules were adopted, with smaller constraints. This was not because tight 2D features could be manufactured more easily, rather it was a Manhattan implementation of a curvilinear mask effect; corner rounding on the mask was significantly pulling back the line ends and corners. In effect, these rules were nothing more than the width-and-space inspection rules in disguise. It was an indication that the MRCs were ill-suited, even then, to describe the real curvilinear effects of mask processing, and necessitated another MRC: that of corner rounding.

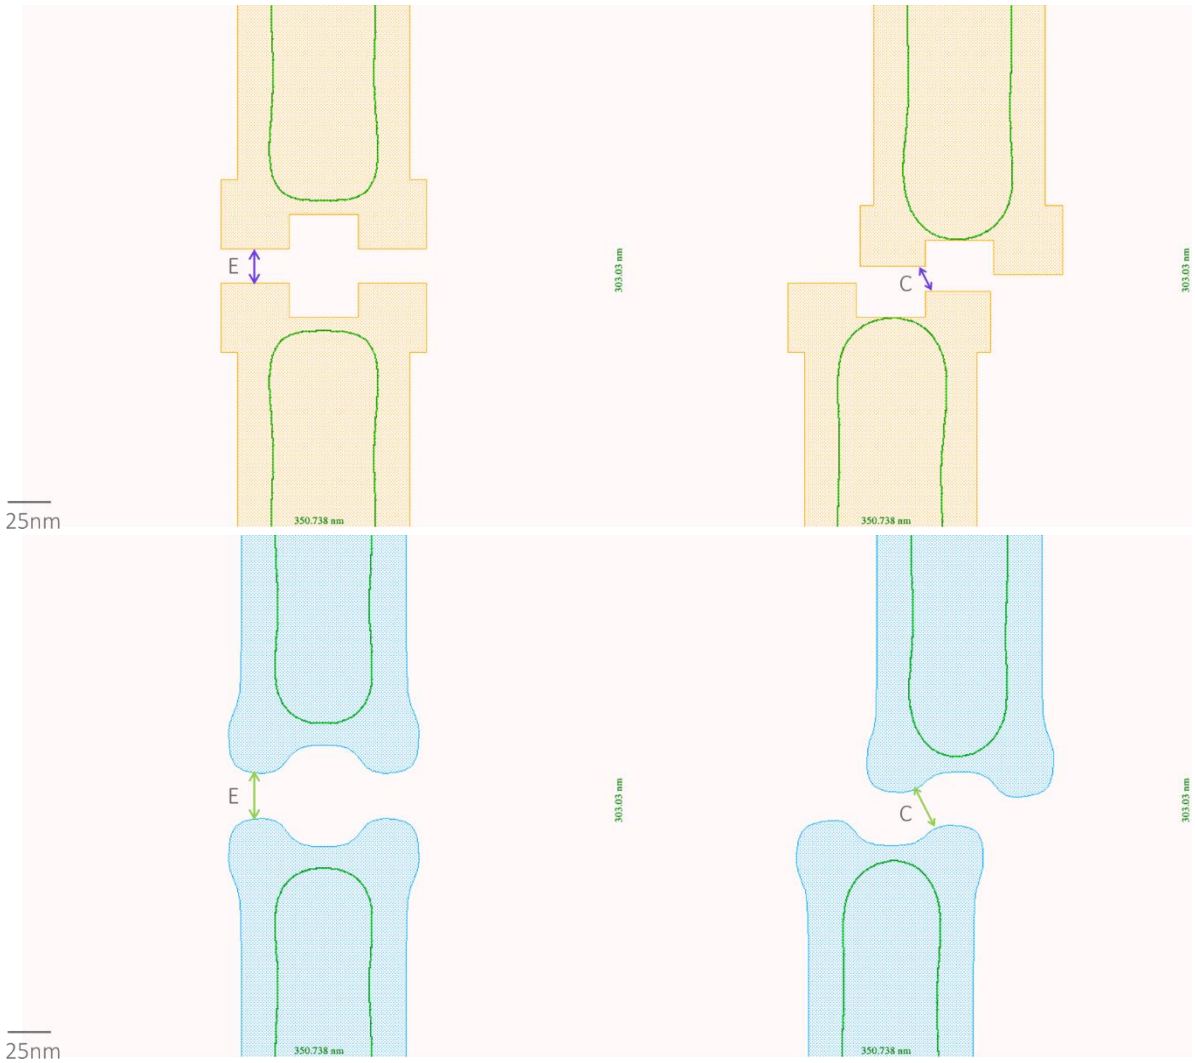

Figure 2: Top example shows end-to-end (E) and corner-to-corner (C) checks. Note these appear smaller than the width and space checks. Bottom example shows the same checks when corner rounding is taken into account. When including the pullback, the distance between features grows allowing inspectability.

As model based OPC began to draw more aggressive shapes, the corner rounding began to significantly impact the kinds of decorations placed on patterns. About this time, the resolution enhancement technology (RET) community began to take a keen interest in the effect of the mask patterning on their optimized mask shapes. While corner rounding was controlled by the mask shop, the impact of that corner rounding on wafer patterning was not well understood by the mask shop. Typically, it would result in print biases that were unaccounted for. An unofficial agreement was made between the mask shop and wafer fab to avoid features that were significantly impacted by corner rounding, but it was typically nebulous as to what those features were.

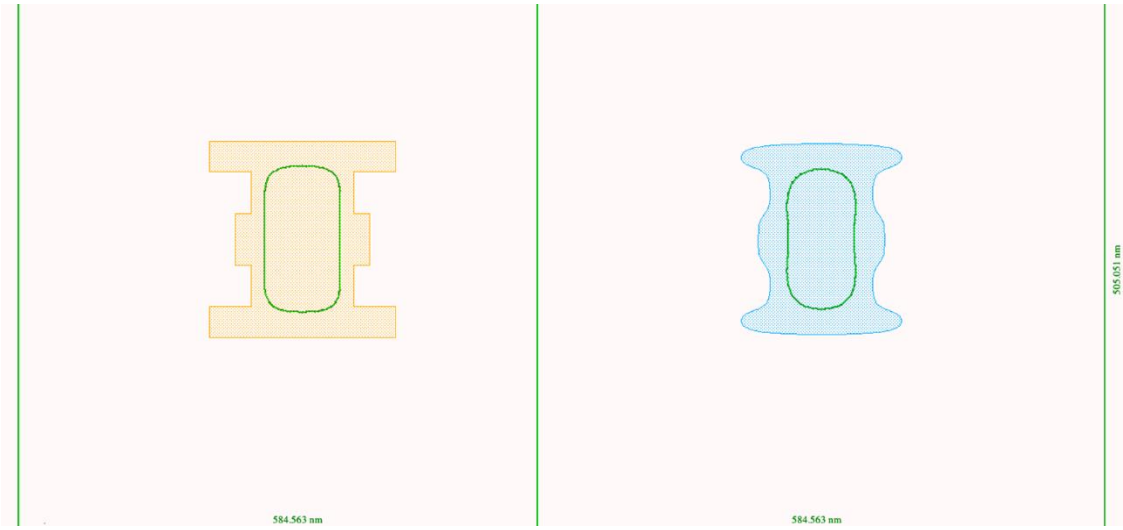

Figure 3: Left shows an aggressive RET intended to correct to a rectangular shape. Right shows the same feature after corner rounding (right). Observe that the long-aspect-ratio RET has significant pullback and the inner slots are not well resolved.

As a result, the OPC community had to account for corner rounding themselves in their process model—by explicitly rounding the corners of their mask patterns during simulation. This was a slow process, and more importantly, was only empirically matched to the real mask. For immersion lithography, such techniques were adequate to get model predictability, but they exposed a key gap in MRC formulation in the industry. Many mask shops, as a result, have gone to great lengths to create complex jog and slot and other rules in an attempt to cover all possible cases of corner rounding that could impact mask manufacturing.

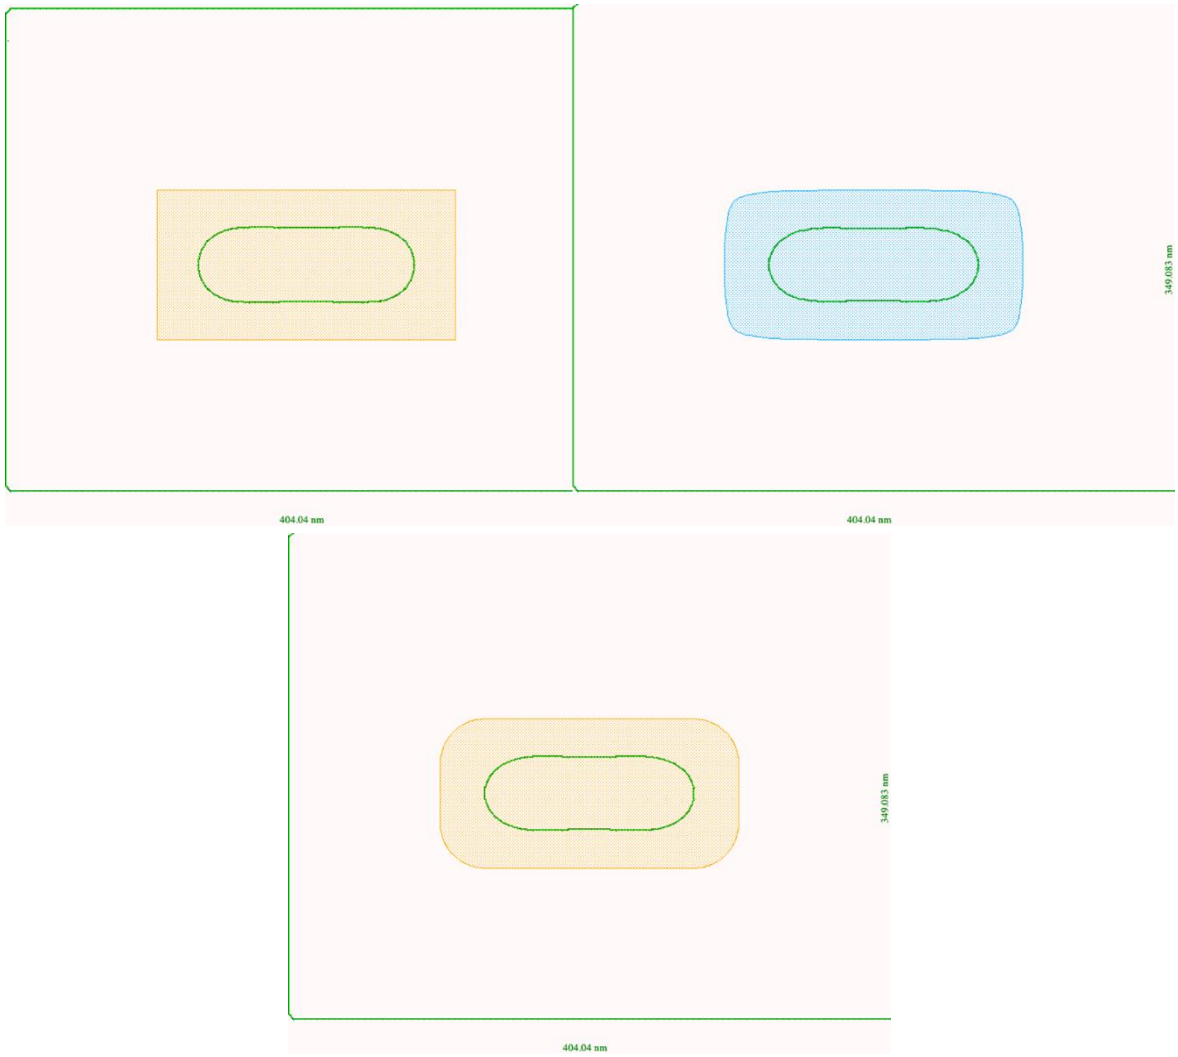

Figure 4: Top, left shows an ideal mask shape with wafer simulation (green). Top, right shows the real mask shape as printed on the reticle, with similar simulation. Notice the wafer contour is smaller. Bottom, center shows a corner-rounded mask, empirically determined to match the real mask shape as part of wafer model calibration.

As immersion lithography became the norm, sub-resolution assist features (SRAFs) were introduced, first in the glass tone, later in the chrome tone. Such features were designed not to print on wafer, but they also suffered from a linearity problem on the mask. They did not print to the desired size on the mask. Fortunately, these were mainly 1D structures, and could easily be resized. Mask process correction (MPC) was born and is still used today to resize small 1D and 2D mask features to print to the desired size.

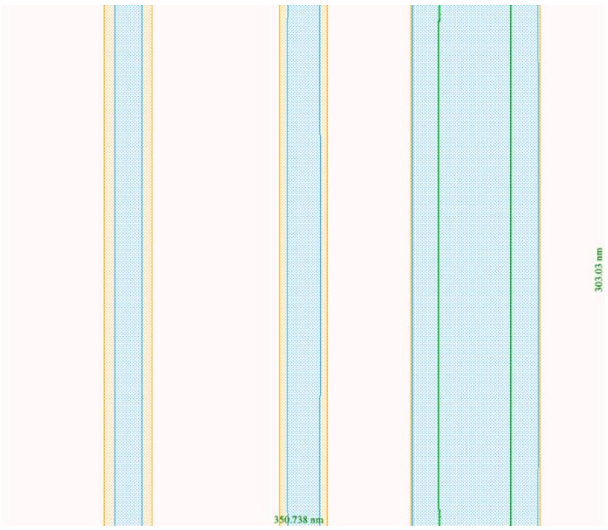

Figure 5: Print biases of assist features on mask. While the main features (right) do not have a significant print bias, the assist features do. The print bias on the assist features not only affects the process window of the main feature, but also the nominal printing size.

Unfortunately, there is a limit to what MPC can do. Wafer fabs are continually asking for smaller and smaller features, and at some point, we reach the limit of what the mask process can support in terms of feature size. Whether the features have line-edge roughness (LER) that is too high, or the features do not adhere to the mask substrate after etch, there is now an MRC for minimum feature size. What was once an inspection limit, the width and space checks have now become manufacturing limits.

This section highlights the fact that the mask rules are overly complex for Manhattan features because they are trying to account for strictly curvilinear effects in the Manhattan domain. Even with this added complexity, it is just not possible to completely describe what can and cannot be printed successfully, so the wafer fab and mask shop still rely on unofficial agreements or empirical evidence to inform the kinds of mask shapes the wafer fab asks to be printed. Therefore, moving MRC to a curvilinear domain *should* simplify the rules somewhat, and allow for a complete agreement with no ambiguity as to the kinds of masks which can be manufactured successfully.

3. MASK VARIABILITY AND CURVILINEAR MRC AND MPC

Let us first focus on the printability for 1D features. Previous studies of mask variability used Monte-Carlo methods in beam position and dose variability to determine the impact of variance on the wafer result [2,3]. The distributions used were typical of the mask manufacturers' tool specifications and were highly dominated by the variability in the dose. Therefore, for results in this paper, we will focus on dose margin (how many nanometers a mask edge moves as a function of dose change) as a proxy for overall mask variability, and we will use 0.5nm/% dose as an upper limit of manufacturability.

Figure 6 shows the 1D dose margin as a function of printed feature size, normalized to the sigma of the shortest Gaussian of the mask exposure (and resist) point-spread function (PSF). The minimum resolvable feature size is about 1.5σ, below which the variability, as observable in LER, of the feature will be too high, due to contrast reasons. It should be noted that we are ignoring effects that are second order to exposure in nature, like pattern lift-off or collapse, in this analysis. That may push the minimum feature higher. The minimum feature can be pushed lower if one increases the dose, ignoring those same lift-off effects. Of course, there is no direct 1D analog to the curvilinear domain, however, the idea of a "long," "mostly straight" edge at any angle will behave like a Manhattan 1D edge.

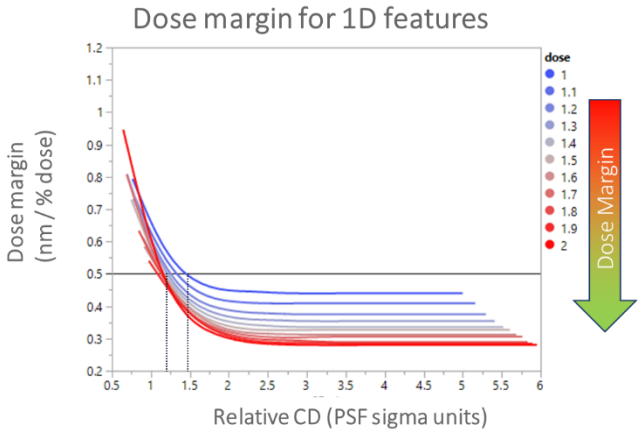

Figure 6: Dose margin for 1D features, as a function of width, in units of Gaussian sigma (PSF) of the feature. 0.5nm/% dose is defined as "acceptable". A higher dose margin than this will result in too much line edge roughness (LER). Features down to about 1.5sigma print with acceptable variability at the standard dose. Higher doses (different colors) give lower (better) dose margins, and given a dose of 1.4x, the acceptable printing limit can be extended down to around 1.25sigma.

For "short" features, the dose margin will be worse, as there is less dose (and less contrast). This means that the minimum feature in highly 2D regions should be larger than the minimum 1D feature. If we take a square, and redo the same dose margin analysis, shown in Figure 7, we see that the minimum printed "square" (which actually prints as a circle) has a minimum width of around 2.25σ. The caveat is that this is the dose margin along the width of the drawn square. As noted above, it actually prints as a circle, and if we look at the dose margin along the corner-to-corner axis, we find the dose margin *never* meets the 1D requirement as shown in Figure 7, unless the mask shop chooses to always write with a slightly larger dose. The other possible implications are that the mask shop has difficulty in measuring the variability along the 45-degree axis, and/or the mask shop allows for a larger variability in 2D regions.

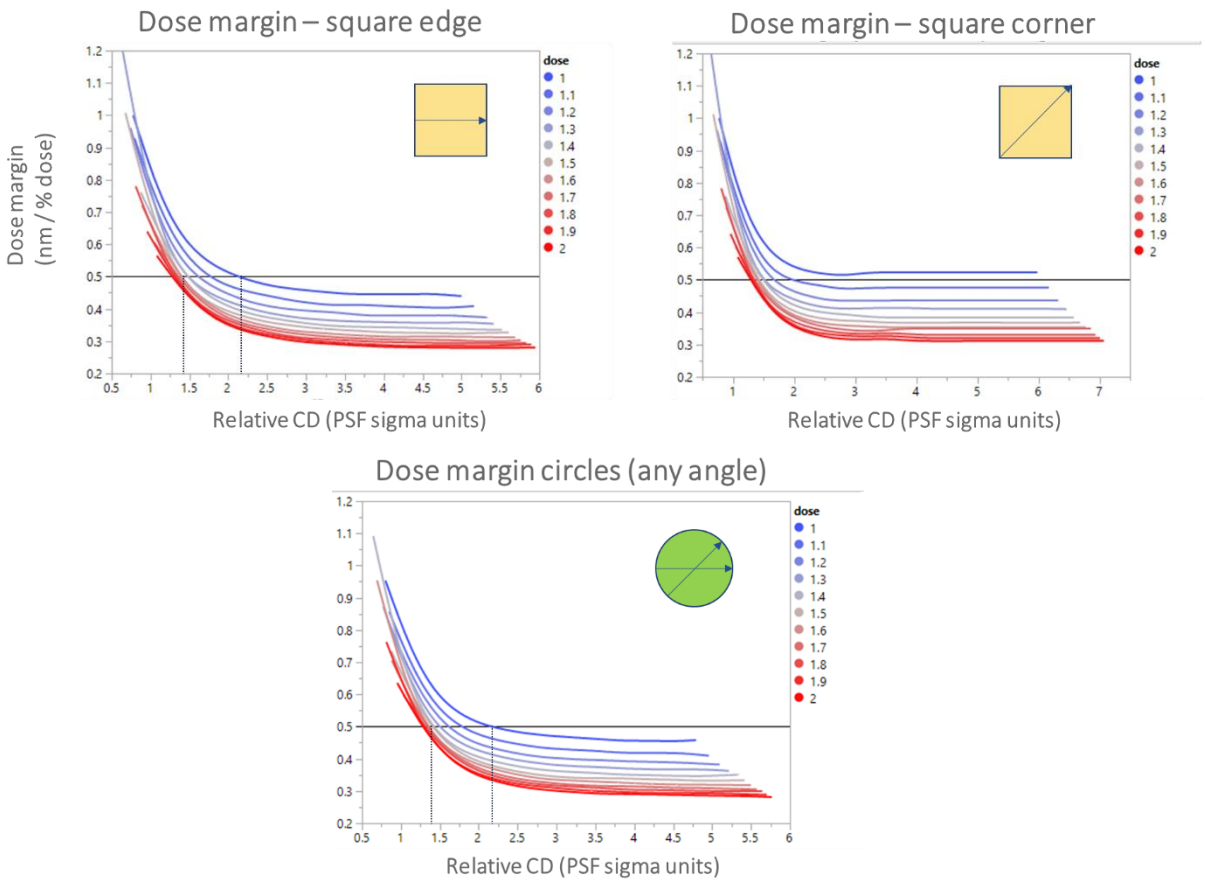

Figure 7: Top, left shows dose margin for a 2D square measured along the width that looks very similar to the dose margin for a 1D feature, shifted to the right. Minimum square is around 2.25sigma, or 1.5 sigma for a higher dose. Top, right shows dose margin for the same square, measured along the diagonal. The dose margin is significantly worse, seemingly requiring a higher dose to print reliably. Bottom, center shows dose margin for a circle is almost identical to that of the square, but for any angle.

Contrast this to using a circular shape to print directly the same size circle. The dose margin is only very slightly worse than the square along the edge axis, however, that dose margin is constant in every direction. Note that, when either using VSB or multi-beam tools to create the circle, while one must approximate the circle using square features, references [2,3] demonstrate that overall the mask variability is significantly improved. It is the more constant dose margin that is responsible for the resultant significant (20-30%) improvement in wafer process variability (PV) band.

This idea of a minimum viable circle is an important one. Not only does it define an effective minimum 2D feature area, but it implies that there is a minimum radius of curvature that the mask process can support. This maximum curvature, combined with the "1D" width and space checks, form the complete set of MRCs from a mask variability perspective. If any incoming design meets the width/space and curvature requirements, then the pattern is manufacturable, even if MPC is required to achieve the patterning.

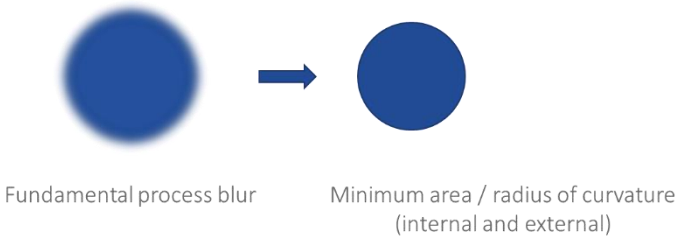

Figure 8: Schematic on how the fundamental process blur leads to the idea of a minimum viable circle, and therefore a minimum radius of curvature.

More importantly, these rules can be easily adopted to existing ILT engines. Non-MRC-aware ILT engines typically produce sharp angles, small features, and thin features, none of which can be manufactured reliably. Recall, however, that there are an infinite number of possible ILT masks which generate nearly equivalent wafer results, so if ILT can constrain the output to be MRC clean, not only will the optical process window be nearly ideal, but the mask variability will be minimized, maximizing the overall wafer process band and, ultimately, yield.

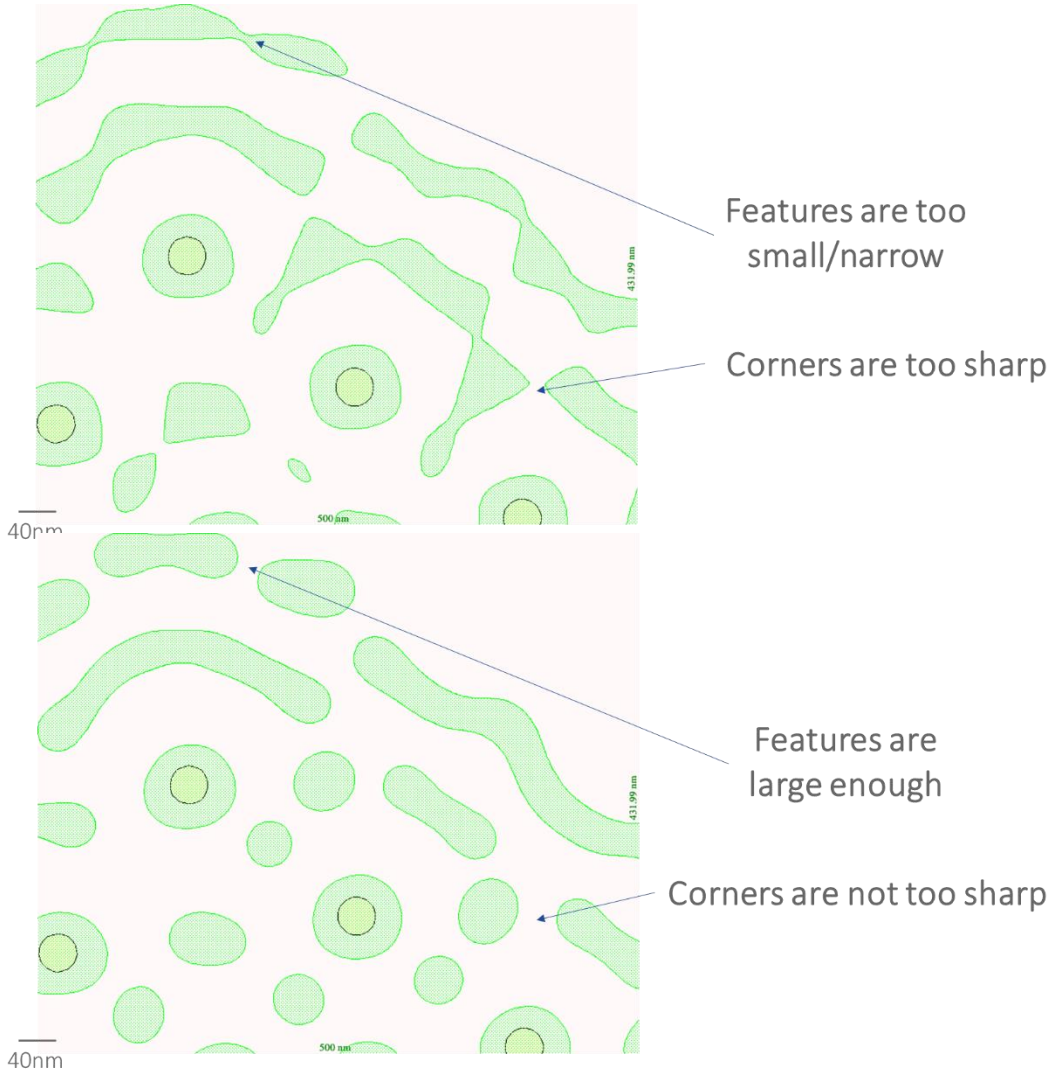

Figure 9: Top shows that unconstrained ILT output typically returns features that are not manufacturable from a mask variability perspective. Sharp corners and small features must be avoided. Bottom shows curvilinear MRC constrained ILT provides not only a similar process window, but guarantees a manufacturable mask.

4. ONE IMPLEMENTATION OF CURVILINEAR MRC

All-angle masks are best processed using GPU architecture [8], and we propose one such curvilinear MRC implementation, based on the idea of two circles. One small circle represents the minimum width/space checks, while one larger circle represents the minimum radius of curvature for the 2D areas. Some mask shops may still want to separate checks into 1D and 2D checks. 1D areas can still be defined by looking for long edges which have very low curvatures.

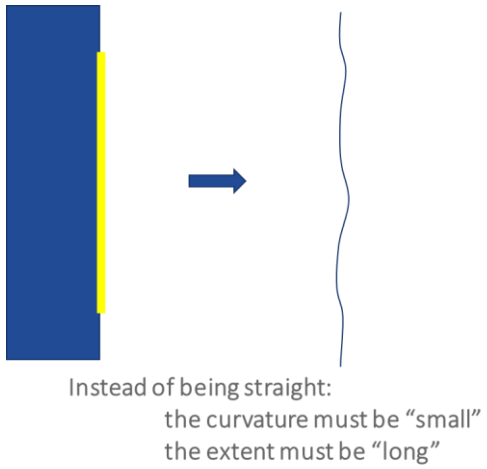

Figure 10: The analog of a 1D edge on a curvilinear mask is that of an extended run of lengths with curvature below a threshold. Not shown would be a 1D feature, which would be two parallel edges with these properties.

Width and space (or bridge and pinch) checks can be performed by simply sliding the minimum width circle throughout the entire geometry. Any places where the circle cannot traverse are violations. Typically, the minimum width and minimum space will be different sizes, due to print biases as part of the mask manufacturing process.

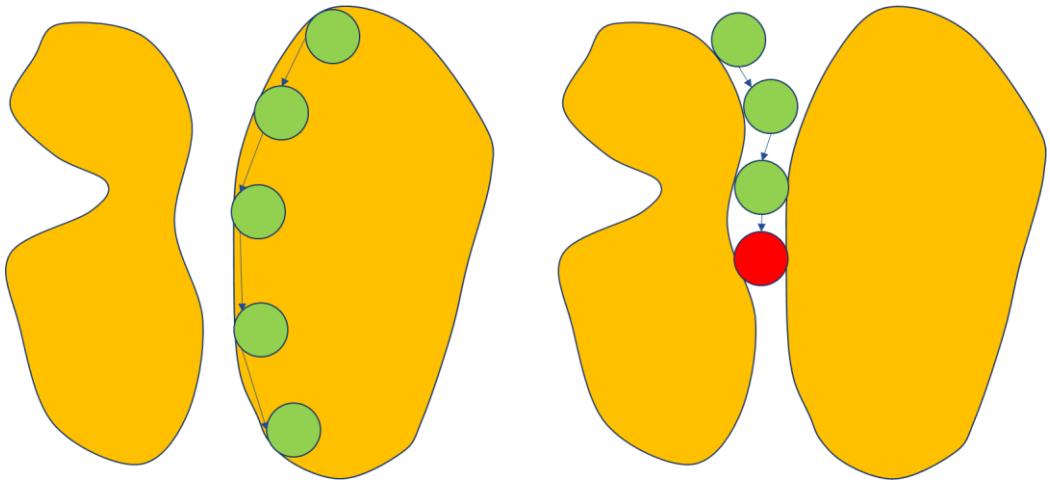

Figure 11: Left shows internal (width) checking. Right shows external (space) checking. As long as the minimum circle can slide entirely within and around the features, it should meet the minimum width and space checks. This example fails on the space check as indicated in red.

Curvature checks can be done by sliding circles around the edge of each boundary. Again, if there is any overlap between the circle and the outside of the pattern, the curvature is too large, and is not reliably manufacturable.

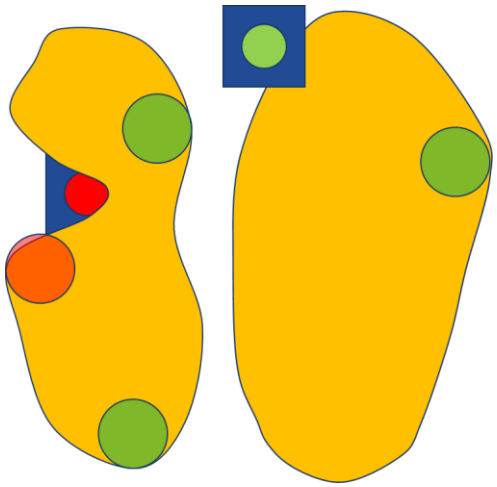

Figure 12: 2D curvature checks. In this representation the internal curvature is larger than the external curvature, which can happen depending on the resist and etch. All that is cared about here is the overlap of the circle tangential to the curve. The checks fail in red.

One comment should be made about the data volume for curvilinear features: when we talk about sliding circles, we are implicitly assuming that there is some sort of smooth (polynomial?) representation of the curves. Today, we typically use the OASIS format, which is polygon-based. When checking distances or curvatures with vertices, there will always be some unavoidable overlap; care needs to be taken to implement algorithms that avoid false positives.

5. SUMMARY

Both EUV and ILT masks will greatly benefit from moving to the natively curvilinear geometries from a process window perspective. Two simple MRC checks based on circles have been identified to constrain the width/space and curvature of the mask geometry, based on a mask variability perspective. These metrics have been created to ensure that a mask which conforms to these constraints can always be manufactured exactly, removing any guesswork as to the mask shape from the OPC model. As a result, it makes the MPC convergence easier (as the target is achievable), and the integrated process more robust.

With such an implementation, the data flow, post OPC, would now be completely curvilinear. While data volume issues still need to be addressed, no longer will mask shapes be restricted to simple geometries, Manhattan or otherwise. With work, this presents an opportunity to move the curvilinear domain up to the design community, in the form of curvilinear designs and the associated DRC rules. Not only will the Manhattan restrictions on RET and mask manufacturing lifted, but the GPU-based computational capabilities to process such geometries will have been proven.

REFERENCES

[1] Jin Choi, Sang Hee Lee, Dongseok Nam, Byung Gook Kim, Sang-Gyun Woo, Han Ku Cho, "E-beam shot count estimation at 32 nm HP and beyond," Proc. SPIE 7379 (2009)

[2] Ryan Pearman, P. Jeffrey Ungar, Nagesh Shirali, Abhishek Shendre, Mariusz Niewczas, Linyong Pang, Aki Fujimura, “Enhancing ILT process window using curvilinear mask patterning: dual mask-wafer simulation”, Proc SPIE 10961 (2019)
[3] PMJ paper
[4] Mahesh Chandramouli, Frank Abboud, Nathan Wilcox, Andrew Sowers, Damon Cole "Future mask writers requirements for the sub-10nm node era", Proc. SPIE 8522 (2012)
[5] Aki Fujimura, Jan Willis "2018 mask makers’ survey conducted by the eBeam Initiative", Proc. SPIE 10810 (2018)
[6] Christof Klein, Elmar Platzgummer "MBMW-101: World's 1st high-throughput multi-beam mask writer", Proc. SPIE 9985 (2016)
[7] Hiroshi Matsumoto, Hideo Inoue, Hiroshi Yamashita, Takao Tamura, and Kenji Ohtoshi "Multi-beam mask writer MBM-1000", Proc. SPIE 10584 (2018)
[8] E. Tzirita Zacharatou, H. Doraiswamy, A. Ailamaki, C. T. Silva, and J.Freire. “GPU Rasterization for Real-Time Spatial Aggregation over Arbi-trary Polygons”, .PVLDB, 11(3): 352, 2017