Design Enablement of Low-Cost Stitching in High-NA EUV Patterning

Sagar Jain[a], Puneet Gupta[a], and Pieter Wöltgens[b]

[a]University of California, Los Angeles, USA
[b]ASML, Netherlands

ABSTRACT

With introduction of the latest generation of High-NA EUV lithography machines with a higher Numerical Aperture (NA) of 0.55 (compared to the previous generation of 0.33NA or Low-NA EUV machines), ASML enables the semiconductor industry to print ever smaller features. However, the transition to High-NA EUV, also introduces anamorphic imaging, with a mask-to-wafer de-magnification that is different between the scan direction (8x) and the slit direction (4x). With a fixed mask size, this results in an exposure field that is half the standard size in the scan direction (16.5 mm instead of 33 mm), and full-size (26 mm) in the slit direction. To manufacture chips larger than the half-field exposure area, therefore two half-field exposures must be stitched together. One way to accomplish this is by imposing an exclusion zone at the stitch area between the two half-field 0.55NA masks to prevent shape-stitching issues for these layers. The higher-level, larger-pitch routing patterns can be made with full-field (low-NA EUV and DUV) lithography, enabling connections across the lower-level half-field stitch zone, thus offering a simple and robust solution to the half-field stitching problem. In line with this approach, we propose overlay error tolerant stitch-aware design methodology for implementing the exclusion zone in the design while mitigating any potential yield losses. This study evaluates the performance and area, and design cost impact of implementing this methodology through a comprehensive Place-and-Route (P&R) implementation of both single-core and multi-core designs. For single-core designs, our analysis indicate a 2-3% reduction in maximum frequency and a 1-3% increase in power dissipation. Additionally, our findings also reveal that the performance in single-core designs is highly dependent on factors such as exclusion zone width, the number of layers fabricated using High-NA EUV, and the exclusion zone's placement location. For multi-core designs, we explore various floorplanning strategies and introduce a stitch boundary placement optimizer designed to accommodate the exclusion zone while minimizing floorplan disruptions and macro redesign costs.

Keywords: High-NA EUV, Stitch Width, Exclusion Zone, Overlay Error, floorplan strategies, macro redesign costs

1. INTRODUCTION

Lithography has helped drive Moore's law deliver ever smaller, denser transistors by enabling printing of ever smaller features. In order to accomplish this, over time, lithography has evolved from Deep Ultra-Violet (DUV) lithography through to 0.33NA (Low-NA) EUV until the advent of 0.55NA EUV or High-NA EUV, by a combination of stepwise reduction of the lithographic wavelength and increasing the NA to enable printing of smaller critical dimensions (CD $\propto \lambda$/NA)

One of the main changes in High-NA EUV is the introduction of anamorphic imaging [1–3]: where older lithography systems use the same 4x demagnification factor from mask to exposure for both the slit direction and the scan direction, for High-NA EUV, the demagnification in the scan direction is increased to 8x, while the demagnification in the slit direction remains at 4x. To minimize the impact on the mask manufacturing, the mask size was kept the same, therefore resulting in a exposure field that has half size in the scan direction (16.5 mm) compared to that of older machines (33 mm) as shown in Fig. 1.

Thus, the half-field exposure area of high-NA EUV necessitates some form of stitching to manufacture chips that exceed the size of the half-field exposure [4]. Stitching in the design potentially introduces overlay and imaging problems between the shapes that may impact the design's yield and manufacturability.

The existing literature provides two main solutions for the High-NA stitching problem:

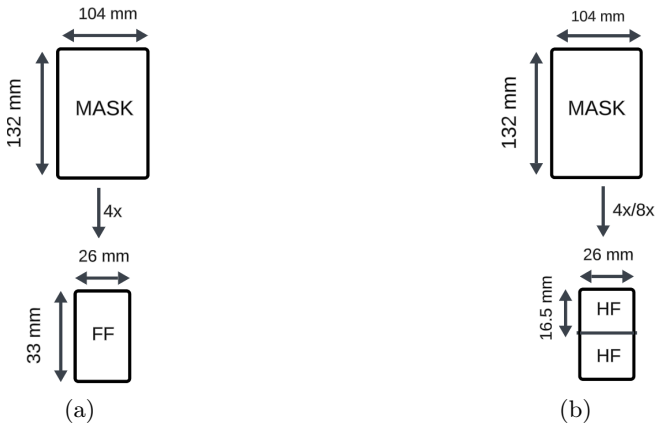

Figure 1: mask field size and exposure field size comparison for different field projections. a) for full-field DUV and Low-NA EUV lithography with 4x demagnification factor in both the slit-direction (x) and scan-direction (y). b) for High-NA EUV with 4x demagninfication in slit-direction (x) and 8x demagnification in scan-direction (y)

1. **At/Near-Resolution Stitching**: Shapes at or near the resolution limit can cross the stitching boundary: this means that these shapes are split between the two half-field masks and need to be placed and imaged perfectly to seamlessly connect : this implies very careful alignment and control of the imaging at and near the boundary between the exposures of the 2 half-field masks. Key challenges are aerial image interactions between the exposed images on both fields, and unwanted reflections from the black border around the mask [4]. Meticulous overlay management, precise imaging in the stitch area and mask resolution enhancement are needed to achieve this [5].

2. **Exclusion Zone Stitching**: This solution circumvents direct interactions between the shapes on the 2 half-field exposures, by introducing a shape exclusion zone in the stitch region between the 2 half-field exposures for layers fabricated using High-NA EUV [6, 7]. For the overall design, this exclusion zone would likely apply to shapes on all levels below the highest level that is manufactured with High-NA EUV. Levels above this level are fabricated using only full-field lithography (Low-NA EUV or DUV) and can then be used for cross-connecting the shapes of the lower level half-field exposures that are separated by the exclusion zone, thus forging the 2 lower level parts separated by the exclusion zone into a single, functional, larger than half-field chip. The key advantages of this approach are its simplicity and inherent robustness.

This paper explores the impact of the latter exclusion-zone approach on chip design in terms of performance, area and design cost.

1.1 Errors in Stitched Design

A major source of yield loss in stitched designs arises from overlay errors between routing segments crossing the stitch boundary. Since the two halves of the design are fabricated independently, misalignment between the segments of a routing shape frequently occurs, leading to yield losses. Fig. 2(a) illustrates yield losses caused by the misalignment of routing shape segments. The independent fabrication of the two halves increases the likelihood of misalignment, directly impacting yield. Yield losses can also arise when High-NA EUV routing deviates from its preferred routing directions. For instance, in scenarios where a horizontal routing layer (aligned parallel to the stitch boundary) is utilized for High-NA EUV stitching, the introduction of vertical routing segments (by the tool during optimization) within this layer may pose challenges. If these vertical segments cross the stitch boundary, they are susceptible to misalignment, akin to the misalignment observed in other High-NA EUV layers oriented perpendicularly to the stitch boundary leading to reduced yield. Another significant source of yield loss stems from incomplete via enclosure between High-NA EUV and non High-NA EUV layers. This issue is particularly more pronounced when layers with the same orientation form direct via connections, as depicted in Fig. 3(a). Misalignment between non High-NA EUV (light brown) and High-NA EUV (yellow) segments results in stacked vias (red) without complete enclosure, thereby causing yield degradation.

1.2 Resolving Overlay Errors in Stitched Design

To mitigate yield losses arising from overlay errors, we propose a set of design rules that must be adhered to during the chip design stage. These design rules translate into physical design (PD) methodology constraints that should be incorporated into the design process to create overlay error-aware, stitch-compatible layouts. By proactively addressing potential issues during the design phase, these rules aim to ensure robust and high-yield chip fabrication. The first rule aims to minimize the interaction of routing shapes within the stitch region. This is achieved by introducing route blockages on metal layers fabricated using High-NA EUV lithography, effectively restricting routing activity in the critical stitch regions. Another rule aims at enforcing higher routing effort in the preferred routing direction & prioritizing the preferred routing direction during the implementation phase to avoid the likelihood of routing shapes deviating into non-preferred directions. To address yield losses caused by incomplete via enclosure, the formation of direct vias between High-NA EUV and non High-NA EUV layers must be strictly avoided. As shown in Fig. 3(b), this measure ensures that stacked vias do not suffer from misalignment, which would otherwise result in partial or inadequate via enclosure. Instead, an intermediate layer that has a different orientation than the two is employed to ensure robust connectivity between the layers without compromising yield. By systematically applying these design rules, designers can ensure that chip layouts are optimized for fabrication under High-NA EUV lithography while minimizing overlay-related yield losses. These rules serve as a cornerstone of an efficient, cost-effective and reliable stitch-aware PD methodology.

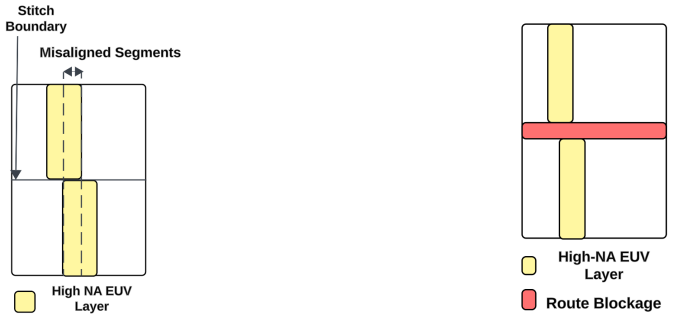

(a) (b)

Figure 2: a) Yield Loss Due to Overlay Error b) Proposed Solution - Route Blockage in the Stitch Region

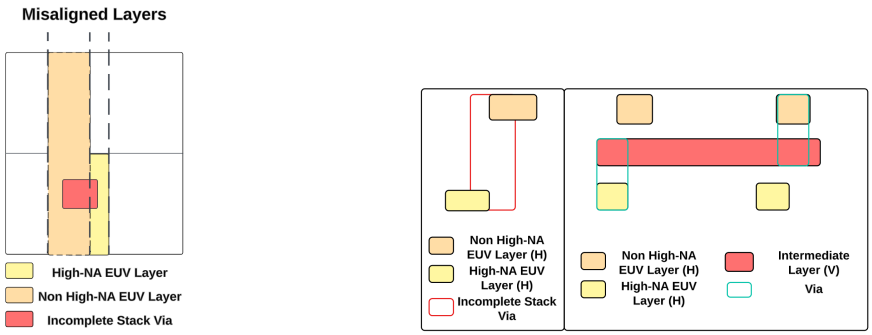

(a) (b)

Figure 3: a) Yield Loss Due to Incomplete Via Enclosure b) Proposed Solution - Preventing Direct Via Formation b/w a High-NA EUV and non High-NA EUV

2. ANALYZING OVERLAY ERROR TOLERANT STITCHED DESIGNS

The following section discusses the impact of implementing overlay error-mitigating PD solutions.

Table 1: Design Parameters of CORTEXM3 & MEMPOOL_TILE

Design Property	CORTEXM3	MEMPOOL_TILE
Utilization	80%	80%
Dimensions	227x230 um	518x521 um
Exclusion Zone (Route Blockage Width)	1um	1um
Metal Layers Available for Routing	M1-M6	M1-M6
High NA EUV Layers	M1 M2	M1 M2
Maximum Number of Layers in Stacked Via	2	2
Technology	TSMC 40	TSMC 40

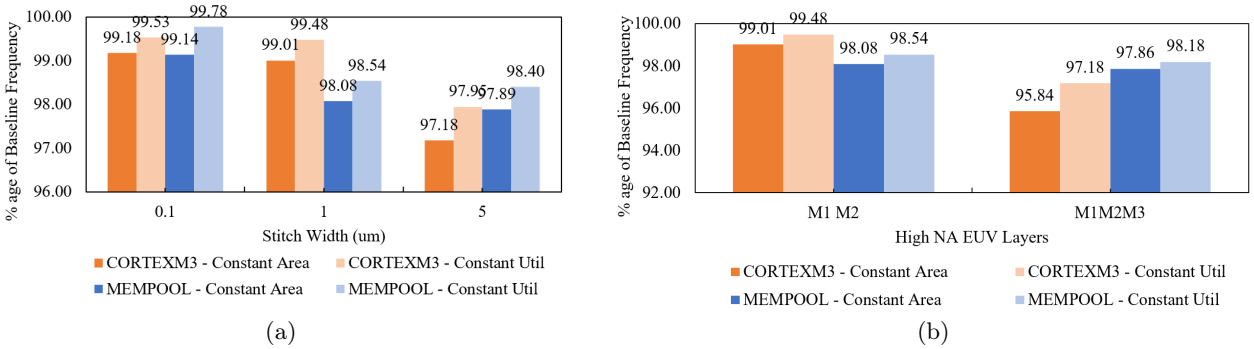

Figure 4: Variation of Performance with a) Exclusion Zone Width in Constant Area & Util Designs b) High NA EUV Layers Usage

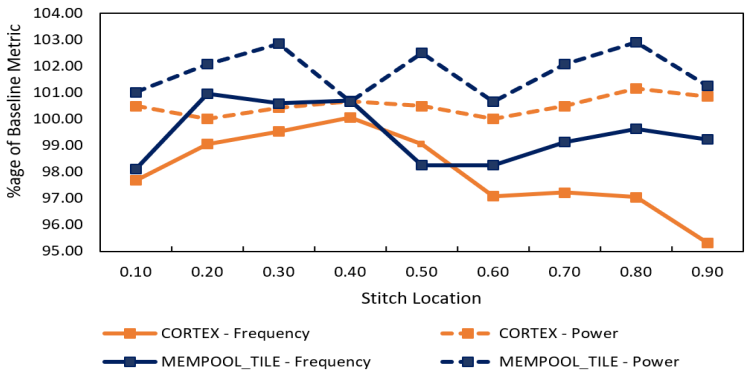

Figure 5: Variation of Performance & Power Dissipation due to 1um stitch boundary placed at different locations in Constant Area Designs

2.1 Single Core Design

To evaluate the impact of overlay error-tolerant stitched designs, the proposed design rules were applied to two standard cell-only designs: CORTEXM3 and MEMPOOL_TILE. For this analysis, a stitch boundary was implemented along with the stitched physical design (PD) rules outlined in the previous section. In [6], a 1 µm stitch boundary was shown to be sufficient for supporting a robust stitching scenario for anamorphic High NA imaging. This result informed our decision to utilize a 1 µm boundary for most of the analyses conducted in this work. Table 1 lists various design parameters that were considered during the analysis of the two designs. The stitch boundary was modelled using route blockages on High-NA EUV layers to prevent connections between independently fabricated route segments. Additionally, placement blockages were introduced in the stitch region to mitigate design rule violations (DRCs) caused by pin accessibility constraints. Stitch boundaries modeled as

route blockages can adversely affect performance by introducing detours, congestion, and increased net lengths. The degree of performance degradation is strongly influenced by the exclusion zone width, with larger widths resulting in greater detours and more significant performance loss. Placement blockages further exacerbate performance issues by reducing the effective area available for cell placement and timing optimization. To mitigate this, constant util designs—where the design area is proportionally increased to maintain a consistent utilization rate as exclusion zone width increases can be employed. Fig. 4 (a) illustrates the performance impact of stitch boundaries on CORTEXM3 and MEMPOOL_TILE designs under constant area and constant utilization conditions as exclusion zone width increases. The results reinforce the fact that performance loss (reduction in maximum frequency achieved) is more pronounced in constant area designs compared to constant utilization designs. Additionally, Fig. 4 (b) highlights the fact that as more layers are fabricated using High NA EUV, route blockage usage increases leading to increased detours, congestion hotspots in the design resulting in greater performance degradation. This effect is uniform and is seen both constant area and util designs.

The location of the stitch boundary is also a critical determinant of performance. Eventhough blockages introduced by the stitch boundary are applied prior to the place-and-route stages, but due to their small size (0.2% - 0.4% area of the whole chip) in our case, their presence does not inherently change the design tool's global optimization behavior, rather it might employ localized fixes to address them [8,9]. The tool treats these blockages as obstructions and employs localized fixes to address them and mitigate any potential performance degradation. It is due to this reason, stitch boundaries near critical paths can cause significant performance issues if limited optimization options are available to the design tool in that region. This highlights the importance of carefully selecting the stitch boundary location to minimize its impact on critical paths and maintain overall design efficiency. Fig. 5 shows the variation of performance and power dissipation in the design when a 1um stitch is placed along different locations in constant area designs. For power, maximum deviation about 1% in CORTEXM3 and 3% in MEMPOOL_TILE from baseline power is observed. In case of frequency, an average deviation of 3% in CORTEXM3 and about 2% for MEMPOOL_TILE is observed. In terms of area, the stitch boundary occupies roughly (0.2% - 0.4%) of the area of the whole design. In summary, the implementation of overlay error tolerant stitch aware design methodologies can influence design performance. The magnitude of the impact depends on multiple factors, including exclusion zone width, the number of layers blocked, and the location of the stitch boundary.

2.2 Multi Core Design

As demonstrated in the previous section, stitch-aware design methodology has some performance impact on single core designs. In case of multi-core designs, the situation becomes much more complicated as the stitch boundary may impact the performance of more than macro block (and the whole design) by intersecting multiple macro blocks as illustrated in Fig. 6 (a). In case of multi-core design, stitch-aware design methodologies would require exploring various floorplans, each tailored to address the unique challenges posed by the exclusion zone intersection.

1. **Floorplan Redesign** : In this approach, the stitch location and the relative positions of the macros remain unchanged, while the floorplan dimensions are expanded to accommodate the stitch boundary. However, this expansion results in additional costs associated with floorplan redesign. Fig. 6 (b) illustrates the floorplan redesign solution for the same.

2. **Stitched Macro Redesign** : In this approach, the stitch and macro locations are fixed, allowing the stitch boundary to intersect macros without altering the overall layout which may necessitate macro redesign. For instance, if a macro has multiple instances in a design and each instance is uniquely affected by a stitch boundary, it becomes necessary to create unique versions of the macro with their respective stitch locations. This requirement for redesigning and maintaining multiple macro variants increases the overall design complexity. Another factor to consider here is the impact of the stitch on the macro performance. Fig. 6 (c) illustrates a potential solution based on the stitched macro redesign approach.

3. **Adjusted Stitch Placement**: This approach involves repositioning the stitch boundary to a location that avoids intersection with any macros. By strategically placing the stitch and redoing the floorplan this method eliminates the need for macro redesign and floorplan extension. Fig. 6 (d) illustrates this approach,

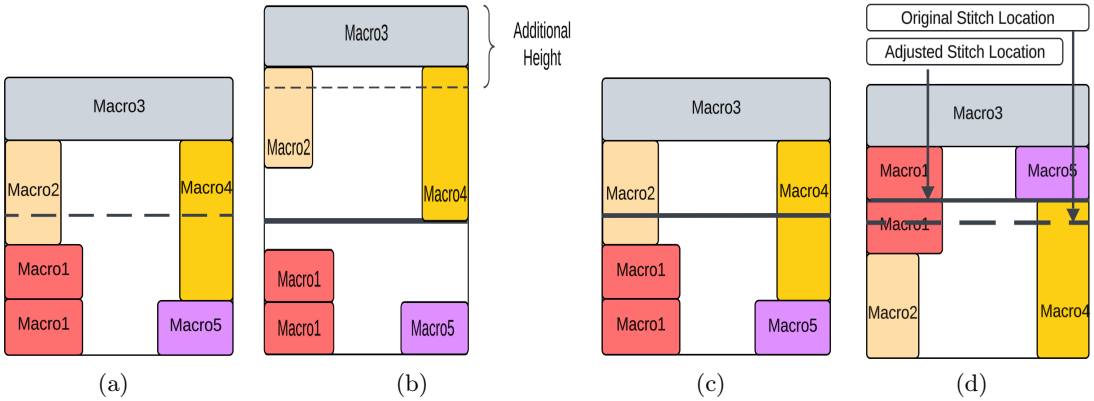

Figure 6: a) Baseline Floorplan of a Multi-Core Design with a Stitch. b) Floorplan Redesign Approach for the Multi-Core Design. c) Stitched Macro Redesign Approach for the Multi-Core Design. d) Adjusted Stitched Placement Approach for the Multi-Core Design

where the stitch is relocated and macros repositioned so as to ensure negligible impact on the existing floorplan or macro design minimizing design complexity and overhead.

2.3 Optimizing Stitch Aware Costs in Multi Core Designs

As previously discussed, stitch-aware multi-core designs incur varying overhead costs depending on the selected floorplanning strategy. In the floorplan redesign approach, the primary cost stems from the additional area required to accommodate the stitch boundary. Expanding the floorplan dimensions to integrate the stitch boundary increases the overall chip area, directly contributing to higher manufacturing expenses. Additionally, the repositioning of macros around the stitch boundary can further affect the design's overall performance due to suboptimal placement or increased interconnect lengths. In contrast, the stitched macro approach introduces a different set of overheads which includes both the macro redesign effort required to ensure compatibility with stitching and potential performance penalties. In extreme scenarios, it may be necessary to create and maintain multiple variants of a macro, each tailored to specific stitch boundary locations. This not only adds complexity to the design process but also imposes additional resource allocation challenges for design & performance validation. In the adjusted stitch placement approach, the incurred cost stems from a combination of floorplan redesign and its potential impact on performance metrics. Adjusting the stitch placement can lead to suboptimal floorplan configurations, which may affect key performance parameters, thereby necessitating additional optimization efforts. These additional costs highlight the trade-offs involved in implementing stitch-aware multi-core designs and emphasize the importance of selecting the appropriate floorplanning strategy based on design requirements and cost considerations.

Although chips fabricated using High-NA EUV lithography require the two halves of the chip to be stitched together, it is not mandatory for the stitch to be positioned at exactly 50% of the chip height. This precise placement is only required for chips with a die size equal to the reticle field size. For all other cases, the stitch location can vary and be positioned anywhere between the half-wafer field size and 50% of the chip height. The flexibility in stitch placement introduces variability in both floorplan redesign and macro redesign costs, as these costs depend on the chosen stitch location. In this work, we find optimal stitch boundary location such that redesign cost is minimized in a design named TOP - formed by joining multiple 14-core CORTEXM3 tiles together. A 14-core CORTEXM3 design, depicted in Fig. 7 (a), measures approximately 3.04 mm $\times$ 2.7 mm and consists of a single CORTEXM3 core and SRAM macros. Given that this size is significantly smaller than even the half-field wafer size (16.5mm x 26mm), multiple tiles of the 14-core design were instantiated and assembled to form a top-level design. The idea here is to add a stitch boundary at an optimal location in the TOP design such that stitching costs could be minimized. For the analysis, multiple different variants of TOP design (different sizes) were used.

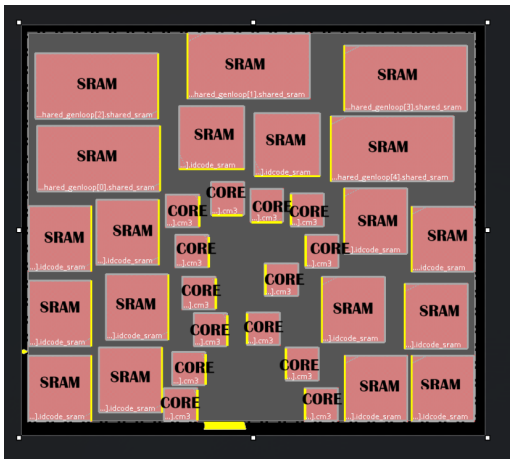

(a)

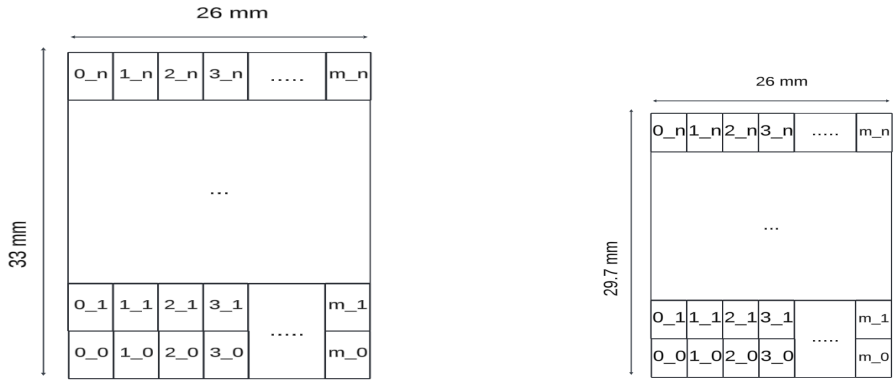

(b) (c)

Figure 7: a) 14 core CORTEXM3 Design b) Die Size is 100% FF Wafer Size c) Die Size is 90% FF Wafer Size

We find the minimum cost of implementing an exclusion zone by using a stitch boundary placement optimizer in floorplan redesign & stitched macro redesign approaches. The number of tiles were adjusted based on the required die size. Fig. 7 (b) illustrates an example of a top-level design, assembled from multiple 14-core CORTEXM3 tiles, where the die size matches the wafer's full-field size. Similarly, Fig. 7 (c) presents another top-level design example with a die size approximately 90% of the full-field wafer size. m & n are integers can change depending upon the die size & tile size and m_n refers to a CORTEXM3 tile instance name.

Fig. 8(a) illustrates the detailed algorithm employed for the floorplan redesign approach utilized in this study. The algorithm begins by selecting a candidate stitch boundary location, denoted as x. At this location, the macro cells intersected by the stitch boundary are identified and analyzed to determine the required floorplan expansion necessary to accommodate the boundary. This expansion represents the floorplan redesign cost for the given stitch boundary location x. The computed redesign cost is then compared against the current lowest floorplan redesign cost. If the newly calculated cost is lower, it replaces the existing lowest cost, and the associated floorplan configuration is updated accordingly. This process is iterative and is repeated for multiple stitch boundary locations across the design space. Ultimately, the algorithm identifies and returns the floorplan configuration that minimizes the redesign cost, thereby offering a cost-efficient solution for incorporating the stitch boundary into the physical design.

Fig. 8(b) presents the algorithm for determining the minimum floorplan increase required when the stitch boundary intersects multiple macros. The algorithm begins by initializing two empty sets. G_{up} & G_{down}. For each macro intersected by the stitch boundary, two parameters are calculated: len_{up} representing the length of

the macro above the stitch boundary, & len_{down} representing the length of the macro below the stitch boundary. The values of len_{up} & len_{down} are compared for each macro. If $len_{up} < len_{down}$ the macro is added to the set G_{down} as it has a greater likelihood of being shifted downward. Conversely, if $len_{up} > len_{down}$ the macro is added to the set G_{up} as it higher affinity for moving upward. Once all macros intersecting the stitch boundary have been categorized into G_{up} & G_{down}, the macros in each set are sorted in ascending order based on their lengths. The largest l_{down} from G_{up} & largest l_{up} from G_{down} are identified and denoted as as $max(l_{down,\ Gup})$ & $max(l_{up,\ Gdown})$ respectively. These values are then summed and represented as B such that $B = max(l_{down,\ Gup}) + max(l_{up,\ Gdown})$. Additionally, the maximum values of l_{up} & l_{down} across all macros intersected by the stitch boundary are computed and denoted as $A = max(l_{down,\ All})$ & $C = max(l_{up,\ All})$ respectively. Finally, the values A, B & C are compared, and the minimum among them determines the minimum floorplan increase required to accommodate the stitch boundary.

Fig. 9 shows two floorplan examples on which the above min floorplan increase algorithm can be applied to find minimum floorplan increase. The horizontal line represents the stitch boundary while the vertical lines represent the macro length intersected by the stitch boundary. In case of Fig. 9 (a), A = max(1,4,6) & C = max(2,3,5). $B = max(l_{down,\ Gup}) + max(l_{up,\ Gdown})$, which is equal to B = max(1) + max(2,5) = 6. Hence, min floorplan increase becomes min(6,6,5) = 5. In case of Fig. 9 (b), A = max(2,4,8) = 8 , C = max(1,3,9) = 9 & B = max(2) + max(1,3) = 5. Hence minimum floorplan increase becomes min(8,5,9) = 5. Fig. 10 shows the variation in various redesign costs with the die size. The x-axis in these graphs represents the die size as a percentage of the full-field wafer size. In Fig. 10 (a), the floorplan redesign cost reflects the increase in total die area (due to insertion of exclusion zone) as a percentage of the baseline die area (without the exclusion zone). It also shows the frequency achieved for both reg2reg and IO paths in the new floorplans as a percentage of baseline (no stitch) floorplan. Due to increase in the floorplan size, the cells are placed slightly farther (as compared to the baseline) from the ports. This causes IO paths to see larger delay leading to their reduced performance. In contrast, Fig. 10 (b) represents the cost associated with the redesign, maintenance & performance penalty of all macros impacted by the stitch, under the macro redesign approach. A detailed breakdown of the redesign and maintenance costs per macro, used to calculate the macro redesign cost is provided in Table 2.

From the graphs, it is clear that when the die size equals the full-field wafer size (i.e. 100%), maximum cost is incurred, as the stitch location is fixed at a single point. Conversely, as the die size decreases, the stitch position becomes more flexible, resulting in lower costs. When the die size falls below a certain threshold (e.g., 92% in this CORTEXM3 example), the minimum stitch boundary region can lie in the middle of two tiles. As the stitch is placed on the edge of the tile, no cost is incurred. Fig. 11 shows the exclusion zone width insertion in different size TOP designs such that the redesign cost is minimum. Based on this analysis, it is evident that the minimum costs associated with implementing a stitch boundary in a multi-core design are influenced by three key factors: the baseline floorplan, the die size, redesign & performance costs associated with macro redesign.

Table 2: Unit Level Macro Redesign Costs

Block Type	Redesign Cost
Memory Macro (ex. SRAM)	100
Core Macro	10

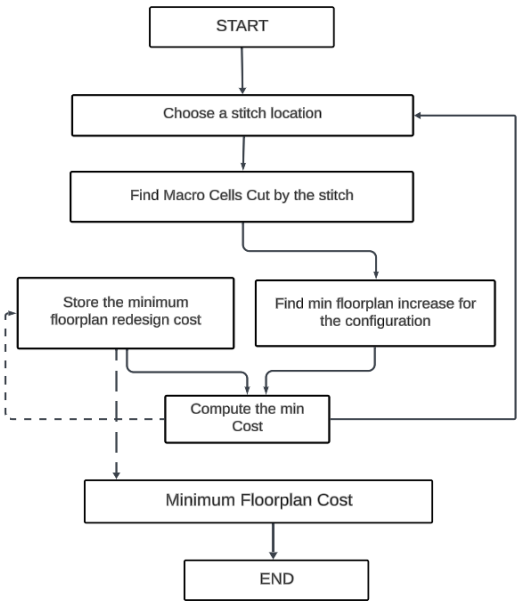

(a) Floorplan Redesign Algorithm

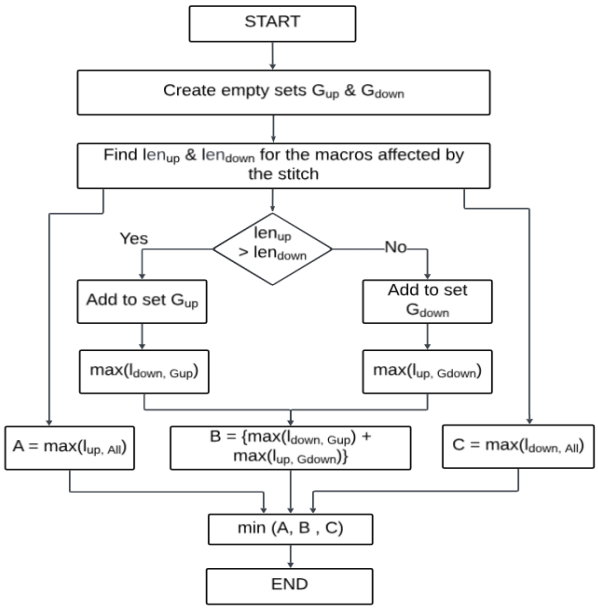

(b) Minimum Floorplan Increase Algorithm

Figure 8

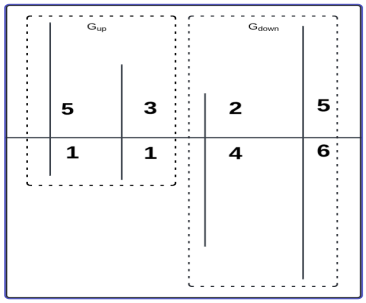

(a)

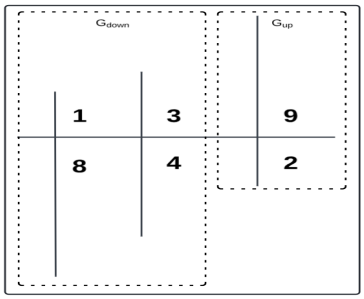

(b)

Figure 9: Minimum Floorplan Increase Algorithm Examples

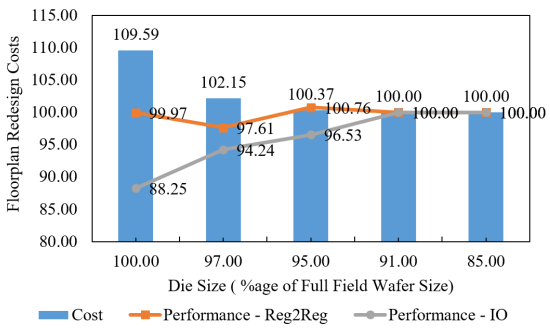

(a) Floorplan Redesign Cost & Performance as a variation of Wafer Field Size

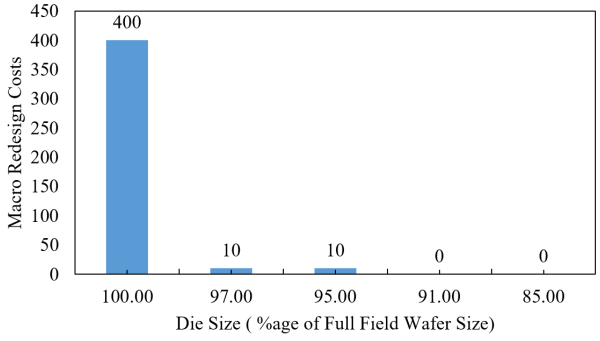

(b) Macro Redesign Cost as a Variation of Wafer Field Size

Figure 10: Redesign Costs

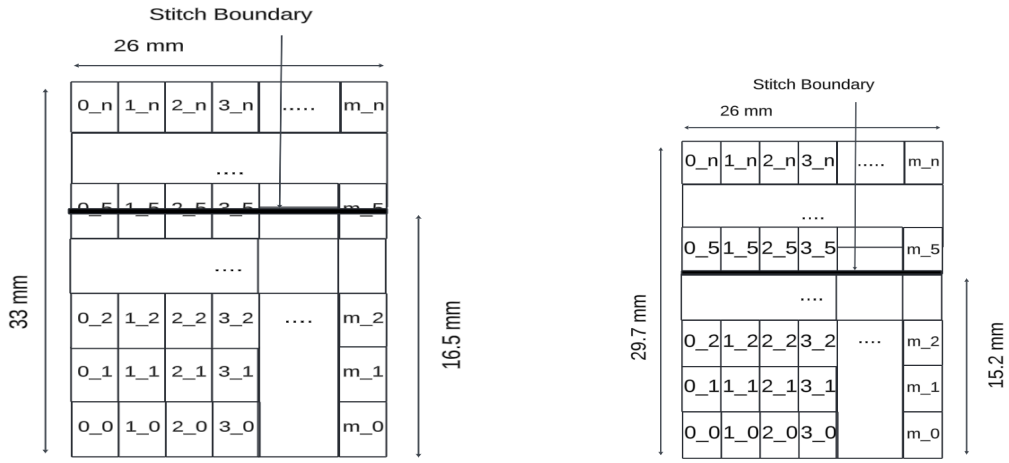

(a) (b)

Figure 11: TOP Design with exclusion zone inserted at minimum cost location when the Die Size is a) 100% FF Wafer Size b) 90% FF Wafer Size

3. CONCLUSION

This work introduces a stitch-aware physical design methodology tailored to support the exclusion zone stitching technique in designs utilizing High-NA EUV (Extreme Ultraviolet) lithography. The proposed design rules are crafted to be both robust and cost-effective, ensuring resilience against overlay errors. For single-core designs, the methodology is implemented by defining exclusion zones through blockages on the High-NA EUV layers. This approach may lead to some performance degradation, the extent of which is influenced by factors such as the width and positioning of the exclusion zone, as well as the number of High-NA EUV layers involved. For multi-core designs, the methodology encompasses various floorplanning strategies and a stitch boundary placement optimizer, which together facilitate the accommodation of the exclusion zone while minimizing the associated costs of redesign and performance penalties. Each floorplan strategy presents distinct advantages and trade-offs, with the optimal approach being contingent upon factors such as the initial floorplan, die size, and the costs related to macro redesign and performance. In summary, the proposed methodologies provide a practical and adaptable framework for effectively implementing the exclusion zone stitching technique in High-NA EUV designs.

REFERENCES

[1] Schoot, J. v., "Exposure tool development toward advanced euv lithography: A journey of 40 years driving moore's law," *IEEE Electron Devices Magazine* **2**(1), 8–22 (2024).

[2] Kaiser, W., "The evolvement of lithography optics towards advanced euv lithography: Enabling the continuation of moore's law for six decades," *IEEE Electron Devices Magazine* **2**(1), 23–34 (2024).

[3] Ronse, K., "Patterning infrastructure development for advanced euv lithography: Continuing dimensional scaling through euv lithography to support moore's law," *IEEE Electron Devices Magazine* **2**(1), 35–44 (2024).

[4] Davydova, N., van Look, L., Wiaux, V., Bekaert, J., Timmermans, F., van Setten, E., Slachter, B., Huddleston, L., van Lare, C., Zhao, R., et al., "Overview of stitching for high na: imaging and overlay experimental and simulation results," in [*Optical and EUV Nanolithography XXXVI*], **12494**, 233–251, SPIE (2023).

[5] Tabery, C. E., Hu, J., Zhao, R., Hennerkes, C., Hsu, S., Liu, Y., Davydova, N., Blanco, V., and Wiaux, V., "Computational lithography and patterning evaluation to support euv high-na stitching," in [*Optical and EUV Nanolithography XXXVII*], PC129530M, SPIE (2024).

[6] Wiaux, V., Bekaert, J., Kovalevich, T., Ryckaert, J., Hendrickx, E., Davydova, N., Woltgens, P., Tien, M.-C., de Winter, L., Maslow, M., et al., "Stitching enablement for anamorphic imaging: a˜ 1μm exclusion band and its implications," in [*Extreme Ultraviolet Lithography 2020*], **11517**, 76–83, SPIE (2020).

[7] Wiaux, V., Davydova, N., Van Look, L., Pellens, N., Weldeslassie, A., Libeert, G., Kovalevich, T., Timmermans, F., and Huddleston, L., "An experimental stitching study on the eve of high-na euv," in [*Optical and EUV Nanolithography XXXVII*], **12953**, 172–182, SPIE (2024).

[8] Sait, S. M. and Youssef, H., [*VLSI physical design automation: theory and practice*], vol. 6, World Scientific (1999).

[9] Golshan, K., [*Physical design essentials*], Springer (2007).