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

Sagar Jain[a*], **Pieter Wöltgens**[b], **Puneet Gupta**[a]
[a]UCLA, ECE Department, Los Angeles, USA
[b]ASML Technology Development Center, Kapeldreef 75, Leuven 3001, Belgium

Background : High-NA EUV lithography, with a higher numerical aperture (NA = 0.55), enables finer feature printing than the previous Low-NA EUV (NA = 0.33). However, it introduces anamorphic imaging, creating a smaller exposure field (16.5 mm × 26 mm). To support larger chip designs, two half-field exposures must be stitched, which can cause shape alignment issues at the boundary, requiring exclusion zones for manufacturability.
Aim : This work aims to evaluate and reduce the performance and design overhead of implementing a physical design (PD) methodology that supports overlay error–tolerant, stitch-aware design for High-NA EUV lithography.
Approach : The proposed method uses exclusion zones to prevent stitching issues in critical High-NA EUV layers, while full-field (Low-NA or DUV) layers bridge connections across the stitch. For single-core designs, the methodology involves placing blockages on High-NA layers. For multi-core designs, a more comprehensive approach is used, involving stitch-aware floorplanning and a stitch boundary placement optimizer to minimize insertion costs such as macro redesign or floorplan changes.
Results : For single-core designs, the exclusion-zone-based methodology results in an estimated 1% performance loss. In multi-core designs, the optimized floorplanning approach effectively accommodates the exclusion zone with minimal additional design cost, demonstrating the practicality of the proposed PD methodology.
Conclusions : The proposed stitch-aware physical design methodology offers a simple and robust solution to the challenges of High-NA EUV half-field stitching. By strategically managing stitch zone placement and routing, the approach ensures manufacturability and yield without significant performance degradation, making it suitable for both single-core and multi-core chip designs in advanced nodes.

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

***Corresponding Author:** sagarjain1997@g.ucla.edu

1 Introduction

Lithography has helped drive Moore's law delivering 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 (High-NA) EUV, by a combination of stepwise reduction of the lithographic wavelength and increasing the NA to enable printing of smaller critical dimensions

(a) (b)

Fig 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 demagnification in slit-direction (x) and 8x demagnification in scan-direction (y)

A fundamental innovation in High-NA EUV lithography systems is the implementation of anamorphic imaging [1–3]. Unlike conventional lithography systems that employ a uniform 4× demagnification factor in both slit and scan directions, High-NA EUV systems utilize differential demagnification - maintaining 4× reduction in the slit direction while increasing to 8× in the scan direction. To maintain compatibility with existing mask infrastructure, the mask dimensions were preserved, resulting in an exposure field that is halved in the scan direction (16.5 mm) compared to previous systems (33 mm), as illustrated in Fig. 1. This anamorphic approach presents both opportunities and challenges for pattern fidelity and throughput.

The reduced field size (16.5 mm × 26 mm) in High-NA EUV lithography necessitates field stitching for manufacturing die that exceed the single-exposure area [4]. This stitching process

introduces critical challenges in overlay control and image continuity at stitch boundaries, which can adversely impact both device performance and manufacturing yield. Current research approaches this problem through two primary methodologies:

1. **At/Near-Resolution Stitching**: When features approach or operate at the resolution limit of the lithographic process, they may span the stitching boundary, requiring division across two half-field masks. For these features to form continuous structures post-exposure, sub-nanometer alignment accuracy and precise image matching at the boundary might be necessary. This introduces several critical challenges such as **aerial image interactions** (overlapping exposure fields must maintain consistent pattern fidelity to prevent edge placement errors or discontinuities), **mask-induced artifacts** (unwanted reflections from the mask's black border can perturb the stitched region's imaging quality)[4]. Mitigating these effects require **advanced overlay control** (to sub-1nm overlay precision to align split features), **stitch-aware resolution enhancement techniques** (to optimize OPC and SRAF placement near boundaries) and **mask-level corrections** (to compensation for edge effects and reflectivity variations at the stitching interface)[5].

2. **Exclusion Zone Stitching**: This methodology eliminates direct pattern interactions between adjacent half-field exposures by instituting an **exclusion zone** within the stitch region for all layers fabricated via High-NA EUV [5–7]. The exclusion constraint propagates hierarchically, applying to all interconnect levels below the uppermost layer patterned with High-NA EUV. Crucially, higher interconnect levels fabricated using full-field lithography (Low-NA EUV or DUV) serve as a routing bridge, interconnecting the discontinuous lower-level structures partitioned by the exclusion zone. This enables the integration of discrete half-field exposures into a unified, functionally coherent die exceeding the single-exposure field size. The approach offers **process robustness** by decoupling overlay-sensitive interactions at critical lower layers and **design simplicity** by avoiding complex

RET or mask-level corrections required for at-resolution stitching.

This paper quantitatively provides different PD (physical design) methodology techniques to implement the exclusion zone paradigm's and analyze its implications in terms of **performance metrics** (timing, area, DRC) & **design cost** (macro redesign efforts, methodology changes)

1.1 Errors in Stitched Design

A predominant yield-limiting factor in stitched designs originates from overlay-induced misalignment of interconnect segments crossing stitch boundaries. As the two design halves undergo independent fabrication processes, systematic and stochastic overlay errors frequently manifest as discontinuities in routed interconnects, significantly impacting yield. Fig. 2(a) qualitatively demonstrates this phenomenon that may lead to yield degradation due to misalignment between adjacent routing segments.

Additionally, significant yield losses emerge at vertical interconnect interfaces between High-NA and non High-NA EUV layers. As illustrated in Fig. 3(a), via enclosure failures occur when misalignments between non High-NA EUV segments (light brown) & High-NA EUV segments (yellow) results in incomplete via landing (red). This failure mode is particularly acute for directly stacked vias connecting same-orientation layers, where the overlay error compounds across multiple lithographic exposures.

[8, 9] introduces several place-and-route (PnR) methodology techniques to address exclusion zones in chip design. The authors suggest that rearranging sub-module and I/O port placements can prevent logic blocks from being fragmented across stitch regions while also minimizing unnecessary routing through these areas. However, their approach is design-specific, requiring designers to analyze the original module placement before implementing the exclusion zone. Additionally,

relocating modules may in some cases, result in performance degradation. Furthermore, if the block is integrated into a system-on-chip (SoC), modifying I/O port placements may not be feasible due to constraints imposed by the locations of interfacing blocks. The authors also propose the use of 2x-width metal layers and modified standard cells within the exclusion zone to enhance robustness, minimize design rule check (DRC) violations, and optimize performance. However, this methodology introduces significant complexity, as it necessitates the development of custom design kits (DKs) and process design kits (PDKs) tailored to the modified structures. Furthermore, interfacing dissimilar metal widths, such as those in the exclusion zone with the non-exclusion zone demands rigorous electromigration (EM) verification and additional DRC validation to ensure reliability.

[10] has also proposed PD methodology solutions to address exclusion zone constraints in integrated circuit layouts.While their approach aligns conceptually with ours [7] particularly in the strategic use of placement blockages within the exclusion zone their work lacks critical analysis of two key implications: (1) the performance degradation induced by such blockage-based methodologies, and (2) the consequent increase in design rule check (DRC) violations. Moreover, their framework does not explore compensatory design techniques to mitigate these effects, such as performance recovery mechanisms or DRC optimization strategies for exclusion-zone-affected regions.

1.2 Resolving Overlay Errors in Stitched Design

To mitigate yield losses induced by overlay errors in stitched designs, we propose a comprehensive set of design-for-manufacturing (DFM) rules that must be systematically implemented during physical design (PD). These rules establish critical methodology constraints to generate stitch-aware

layouts that inherently compensate for overlay variations. By proactively addressing lithographic challenges at the design stage, this approach ensures robust fabrication while maintaining yield targets.

1.2.1 Design Rules

1. Stitch Region Isolation: Implement route blockages on all High-NA EUV patterned metal layers within exclusion zone to restrict routing activity in critical boundary zones to minimize overlay-sensitive interactions.

2. Placement Blockage Isolation: Introducing hard placement blockages within the exclusion zone to mitigate potential pin accessibility issues for cells situated directly beneath the exclusion zone.

3. Robust Vertical Interconnect Methodology: To prevent direct via formation between High-NA and non High-NA EUV layers (Fig. 3(b)), an orientation-decoupled intermediate routing strategy is employed. This involves the use of orthogonally aligned metal layers and adherence to via enclosure rules through directional isolation. In essence, if a High-NA layer is adjacent to a non High-NA metal layer, the routing directions of these two layers must be orthogonal. Any coincidence in routing orientation combined with via interconnections may lead to via misalignment and incomplete metallization, as exemplified in Fig. 3(b). To avoid such defects, an intermediate orthogonal layer (red) can serve as a barrier layer. This orthogonal interpose prevents direct via formation between the High-NA and non High-NA layers, thereby ensuring a robust transition and adherence to manufacturability criteria. While contemporary technology nodes embed this orthogonality requirement intrinsically within design rules, older nodes lack such built-in enforcement. Consequently, in these legacy nodes, explicit directional routing constraints can be specified to

maintain orthogonality between adjacent High-NA and non High-NA layers.

(a) (b)

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

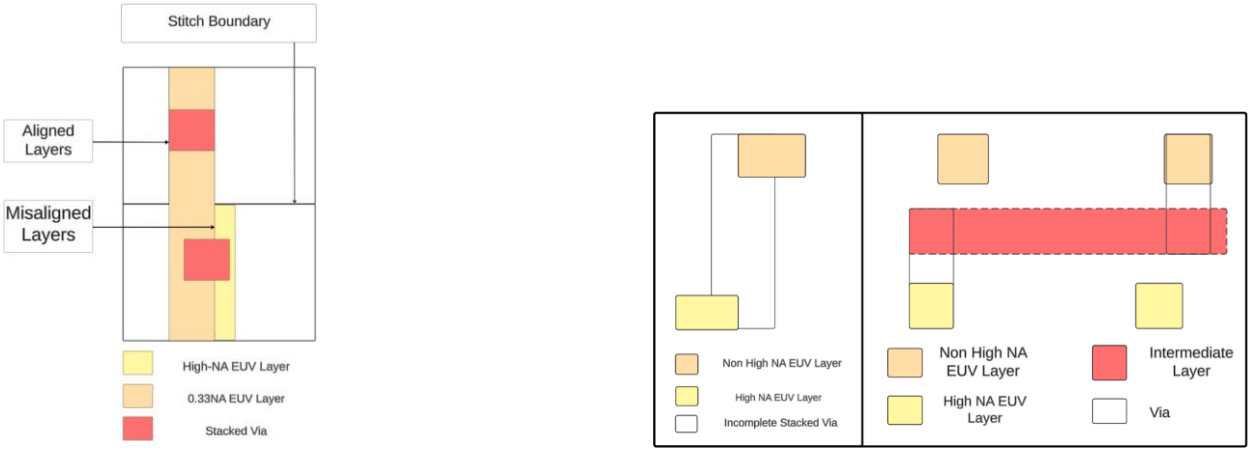

(a) (b)

Fig 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

To assess the efficacy of our overlay error-tolerant stitching methodology, we implemented the proposed design rules on two standard cell-based designs: CORTEXM3 and MEMPOOL_TILE.

Table 1: Design Parameters of CORTEXM3 & MEMPOOL_TILE

Design Property	CORTEXM3	MEMPOOL_TILE
Utilization	80%	80%
Dimensions	74.05x73.87um	151.66x151.63um
Exclusion Zone (Route Blockage Width)	2um	2um
Metal Layers Available for Routing	M1-M6	M1-M6
High-NA EUV Layers	M2 M3	M2 M3
Maximum Number of Layers in Stacked Via	2	2
Technology	ASAP7	ASAP7

Our experimental framework incorporated a 2 um exclusion zone boundary - similar to the one used in [8, 9] on M2 & M3 layers along with the full suite of stitch-aware physical design constraints detailed in Section 1.2. Table 1 summarizes all the design parameters we used for evaluating the performance. Route blockages on the High-NA EUV layers were employed to prevent ill-formed interconnections across the fabrication divide while introducing placement blockages to address pin accessibility-induced design rule violations. Further, the ASAP7 technology LEF file was also modified to limit the maximum of metal layers in a stacked via to 2.

2.1 Single Core Design

2.1.1 Performance Variation with Location

Modeling stitch boundaries as routing blockages can negatively impact circuit performance by inducing signal detours, exacerbating routing congestion, and increasing net lengths. Usage of placement blockages compounds this further by restricting the available placement area, thereby limiting opportunities for cell optimization and timing closure. This constrained approach inevitably introduces trade-offs: route blockages precipitate routing detours and increased net lengths while placement blockages increase the overall design utilization. To alleviate the performance challenges

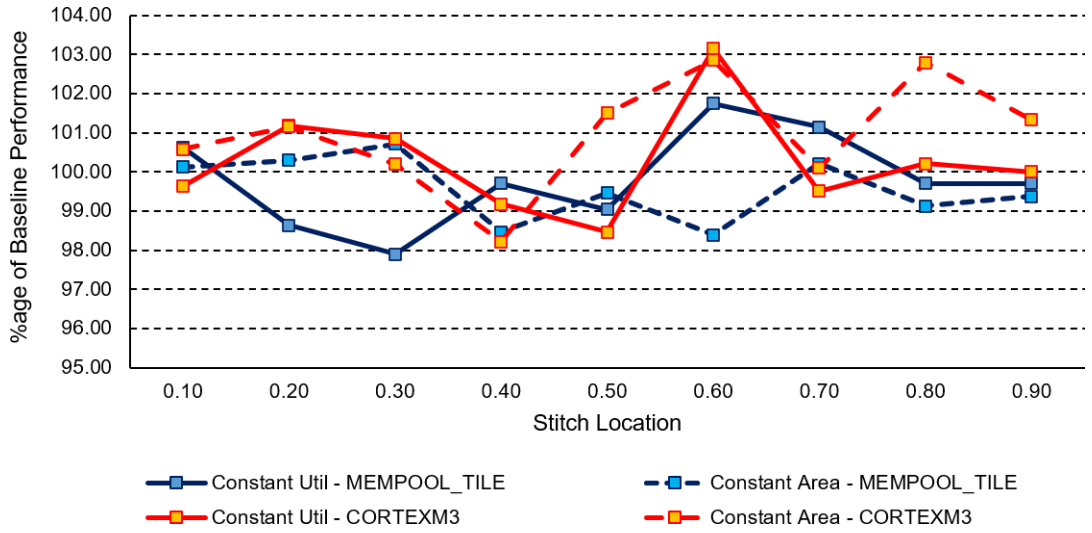

Fig 4: Variation of Performance due to 2um stitch boundary placed at different locations in ASAP7 technology

due to blockages, designers can employ constant-utilization design strategies – where the total design area is systematically scaled to maintain cell utilization with and without the exclusion zone. This approach also helps in mitigating routing congestion, ensuring minimal performance degradation despite the presence of stitch-induced blockages.

Fig. 4 shows the variation of performance [1] as the exclusion zone location is varied for ASAP7 technology node. The results show that for a 2um exclusion zone, designers can choose from either constant area [2] or constant utilization [3] approach and a maximum performance loss of about 2% for both CORTEXM3 and MEMPOOL_TILE is observed. Although the constant utilization approach results in a maximum performance loss of approximately 2% at certain stitch locations, designers can mitigate this degradation by adopting the constant area approach for those specific

[1] Performance: The maximum clock frequency at which the design can operate

$$\text{Performance} = \frac{1}{\text{Clock Period} - \text{Slack}}$$

[2] Constant Area: In these designs, the total design dimensions remains same before and after exclusion zone insertion

[3] Constant Utilization: In these designs the total design utilization remains same before and after exclusion zone insertion

regions. By strategically combining the constant area method in critical locations with the constant utilization approach elsewhere, the maximum performance loss is reduced to around 1% for the affected areas. This hybrid approach ensures improved overall efficiency while minimizing localized performance trade-offs. In [7] we have also shown that the performance of a design is also impacted by the number of High-NA EUV layer usage. As the number of High-NA EUV layers increase more signal detours and routing congestion is observed leading to increased net lengths and reduced performance.

Table 2: Performance Impact (%age of Baseline) of Stitch Width Variation

Type	Technology	0.1um	0.5um	1um	2um	5um	10um	25um
Constant Area - CORTEXM3	ASAP7	99.8	98.7	97.5	98.9	96.2	26.6	-
Constant Util - CORTEXM3	ASAP7	100.1	99.2	99.5	99.1	98.9	98.4	97.4
Constant Area – MEMPOOL_TILE	ASAP7	101.8	100.0	100.5	102.1	96.5	96.5	94.1
Constant Util – MEMPOOL_TILE	ASAP7	100.5	99.6	100.6	99.0	99.6	99.6	100.5
Constant Area - CORTEXM3	TSMC40	98.8	98.9	99.0	98.6	95.8	89.2	55.4
Constant Util - CORTEXM3	TSMC40	100.2	99.0	101.6	98.7	98.6	97.9	98.1
Constant Area – MEMPOOL_TILE	TSMC40	97.8	99.5	101.0	98.9	99.9	98.5	95.0
Constant Util – MEMPOOL_TILE	TSMC40	102.1	101.8	101.2	100.4	98.5	98.6	97.2

Table 3: Design Area

Type	Technology	Baseline Design Dimensions (um)
CORTEXM3	ASAP7	74.05 x 73.87
MEMPOOL_TILE	ASAP7	151.66 x 151.63
CORTEXM3	TSMC40	229.18 x 227.5
MEMPOOL_TILE	TSMC40	521.08 x 512.56

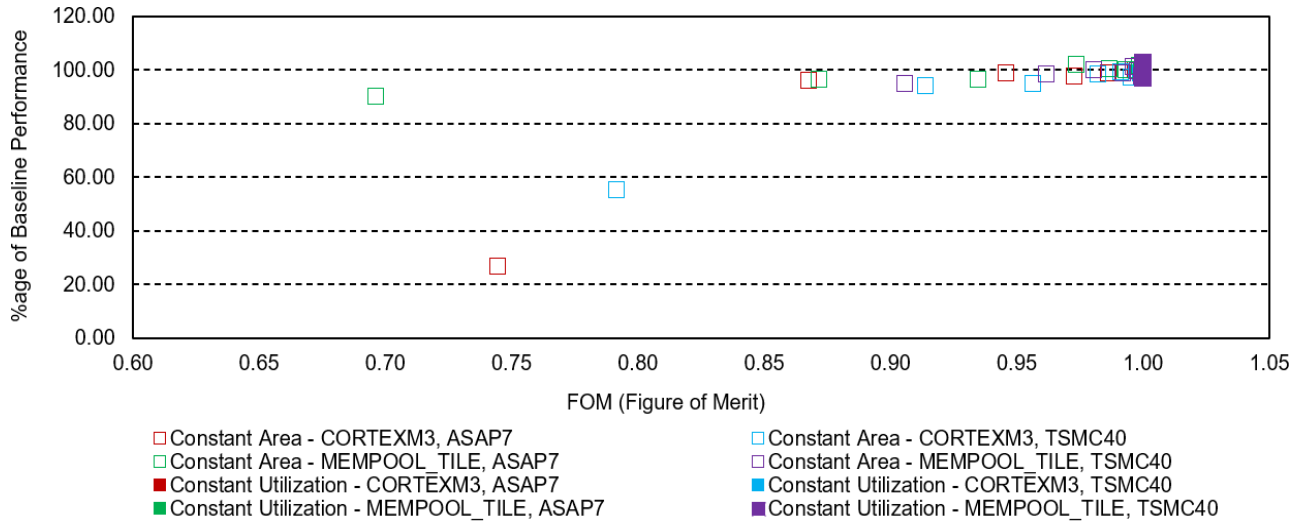

Fig 5: Variation of Performance with FOM (Figure of Merit)

2.1.2 Performance Variation with Technology

Table 2 lists the performance impact as a percentage of baseline performance in the designs at different technology nodes. It shows that as the stitch width increases, the performance drops significantly in constant area designs which can be stemmed by switching to constant utilization designs. Table 3 lists the Baseline Design Dimensions for the two designs in ASAP7 & TSMC40 technology. Based on Table 2 & Table 3 we can see that not only the exclusion zone width but the design size also plays a role in performance degradation. A 10 um exclusion zone boundary in a smaller design (ex. CORTEXM3 in ASAP7) can cause higher performance impact as compared to bigger design (ex. CORTEXM3 in TSMC40).

If all other factors (number of High NA EUV layers, exclusion zone location) are kept the same, the performance of a design post exclusion zone insertion is highly dependent on the design size as well and the exclusion zone width. Based on this we created a FOM (Figure of Merit) metric that can be used to roughly estimate the design performance post exclusion zone insertion without performing actual PnR.

$$FOM = x^2$$

$$x = \frac{Total\ Available\ Area}{Total\ Design\ Area}$$

$$x = \frac{(Total\ Design\ Area - Blocked\ Area)}{Total\ Design\ Area}$$

Here, x is a measure of how much design area is available for optimization post exclusion zone insertion. It is defined as the ratio of Total Available Area [4] and Total Design Area[5]. In case of constant area designs, x keeps on decreasing as the exclusion zone width is increased. The overall performance changes arise from multiple factors such as design area, routing resource availability, etc. Relying exclusively on x as a performance predictor would only capture the area-related component and will fail to account for other components. Accurately modeling the impact of routing resources as a deterministic function of x is inherently challenging due to the multifaceted interactions between routing topology, spatial layout, and signal integrity constraints. Consequently, we have adopted an empirical evaluation methodology, wherein we assess performance across a variety of input configurations and scenarios. It is from these results we observed that x^2 exhibits a stronger correlation with performance than x.

Fig. 5 plots the variation of performance impact as the Figure of Merit (FOM) is varied for different design configurations (ex. constant area or utilization, CORTEXM3 or MEMPOOL_TILE , technology node). For a given design configuration, multiple (x,y) values correspond to different performance impacts due to different exclusion zone widths. For instance, constant area COR-

[4]Total Available Area: Area left for placement post exclusion zone insertion
[5]Total Design Area: The total design area without any exclusion zone insertion

TEXM3 design in TSMC40 technology has an FOM of 0.7 for an exclusion zone width of 25 um while it is 0.87 for an exclusion zone width of 10 um. From the graph, we can roughly estimate that for FOM values lower than than 0.8, the performance impact is more than 5%. In Fig. 5, a single data point (hollow green box) corresponds to a 10% performance loss at a FOM of 0.58, under the configuration of MEMPOOL_TILE with a 25 um exclusion-zone width. Although this point may visually appear as an outlier, it nevertheless aligns with the broader trend: FOM values below 0.8 are consistently associated with performance degradations exceeding 5%. The comparatively reduced performance impact observed in MEMPOOL_TILE versus CORTEXM3 can be attributed to the substantially larger design area of MEMPOOL_TILE it occupies approximately four times the area of the CORTEXM3 design. This increased design size affords the optimization tool greater freedom to compensate for performance loss even after exclusion zone insertion. Consequently, despite the imposed constraint, the tool can still recover performance through optimization opportunities available in the larger area.

Overall, the performance impact of overlay error tolerant stitching methodologies manifests as a complex function of various factors such as exclusion zone width, design area, exclusion zone boundary location, quantity of blocked routing layers, etc.

2.1.3 DRC

A critical metric in the development of exclusion zone-aware physical designs is the impact on Design Rule Check (DRC) violations. The introduction of route blockages inherently reduces the number of available routing tracks, exacerbating congestion and potentially increasing DRC violations. Resolving these violations is a labor-intensive process, often requiring significant manual intervention particularly in highly congested designs where automated correction methods

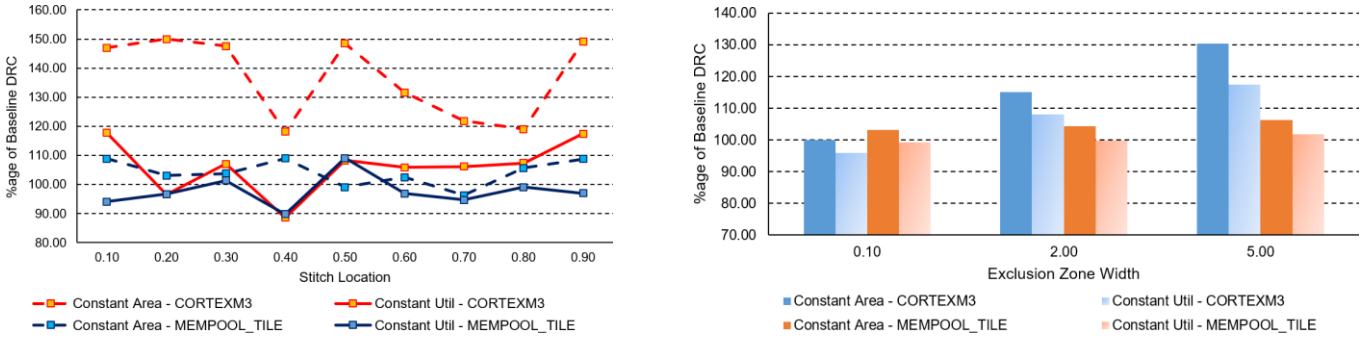

(a) (b)

Fig 6: Variation of DRC with Exclusion Zone a) Location b) Width

prove insufficient.

Fig. 6 (a) shows the variation of DRC count as the exclusion zone location is varied in constant area and utilization designs. Fig. 6 (b) shows the variation of DRC count as the exclusion zone width is varied in constant area and utilization designs. Both the graphs show that as the exclusion zone width increases, the number of DRCs increases in constant area designs. However, this increase in the DRC count is reduced as we switch to constant utilization designs.

2.1.4 Area

Area is also an important metric that is used for comparing two designs. In constant area designs, inserting an exclusion zone does not alter its total area. In constant utilization designs, the design is modified to include more area such that overall placeable area (before and after exclusion zone insertion) remains same. However, the additional area penalty is proportional to the exclusion zone width and significantly less (< 1%) of the total design area.

2.1.5 Power Dissipation

Along with area, total power dissipated is also another metric that is crucial in comparing the design performance. In [7] we have already shown that maximum variation in power dissipated due to

exclusion zone insertion is around 1-2%.

2.2 Multi Core Design

While the previous section established the performance impact of stitch-aware methodologies on single-core designs, the challenges become significantly more complex in multi-core implementations. Here, a single stitch boundary can intersect multiple macro blocks, potentially affecting the performance of each macro as well as the overall system, as illustrated in Fig. 7(a). In the figure, the colored rectangular boxes represent the different macros present in the design while the remaining area (light blue) is where stdcells and top level blockages can be placed. To address these challenges, stitch-aware methodologies for multi-core designs necessitate the exploration of tailored floorplanning strategies aimed at accommodating the exclusion zone.

1. **Floorplan Redesign**: This approach preserves the original exclusion zone boundary and the relative macro placement by expanding the overall floorplan dimensions, thereby ensuring that the exclusion zone does not intersect with any existing macros. The expanded floorplan provides additional spatial flexibility, allowing macro repositioning to accommodate an exclusion zone that avoids cutting across macros. Although this approach is conceptually straightforward, it introduces a significant area overhead and necessitates redoing of the place-and-route (PnR) process. This additional area required to accommodate the exclusion zone will henceforth be referred to as the floorplan redesign cost. Fig. 7 (b) illustrates this strategy, highlighting the trade-off between effective exclusion zone integration and increased silicon area utilization. The extent of floorplan expansion is governed by multiple design-specific factors, including macro placement, macro orientation, and the location of the exclusion zone. In the designs evaluated in Section 2.7, macros typically have heights on the order of microns. Consequently, the maximum increase in floorplan

area due to this approach is also of a similar magnitude.

2. Stitched Macro Redesign: In this approach, the exclusion zone boundary is permitted to intersect existing macros, thereby obviating the need to modify the overall floorplan. However, this necessitates macro-level redesign, since any macro traversed by the exclusion zone must be reengineered. In scenarios where multiple instances of a macro are infected i.e., uniquely affected by the exclusion region each instance may require a specialized variant to preserve design performance. The cumulative burden of developing and maintaining these macro variants will henceforth be referred to as the macro redesign cost. Fig. 7 (c) illustrates the exclusion zone situated at a fixed elevation corresponding to 50% of the total floorplan height. Under this configuration, the exclusion boundary intersects macro 2 and macro 4, both located at the mid -height mark and would consequently require redesign. This strategy sharply increases design complexity and verification effort, as multiple macro versions must be managed, verified, and integrated uniquely.

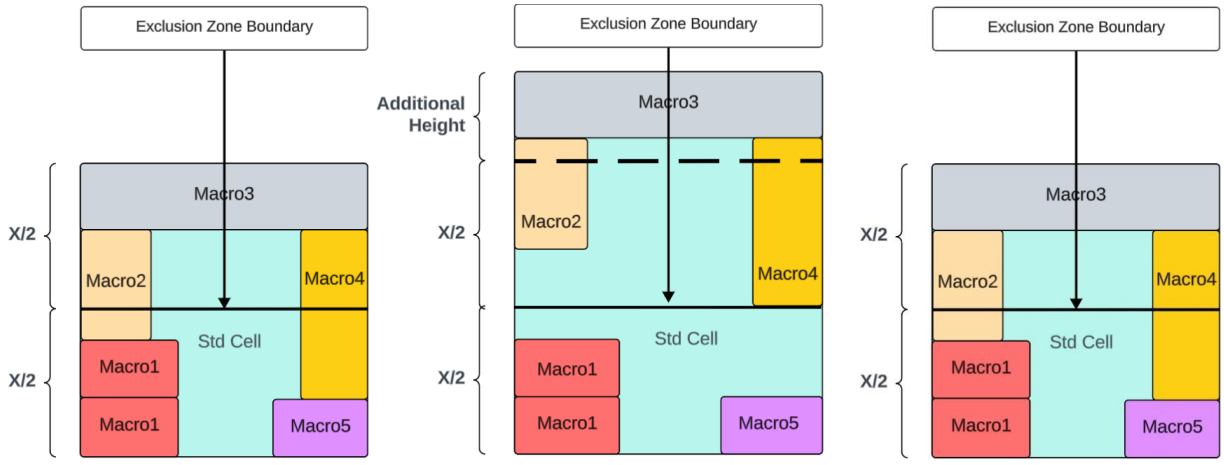

Fig 7: a) Baseline Floorplan of a Multi-Core Design with an Exclusion Zone. b) Floorplan Redesign Approach for the Multi-Core Design. c) Stitched Macro Redesign Approach for the Multi-Core Design

It should be noted that the figures presented are intended primarily to illustrate the accom-

modations of exclusion zones under different design approaches; they may not necessarily depict full-field die configurations. Fig. 7 (a) displays a multi-block layout ($x < 33mm$) through which an exclusion zone passes through. The objective here is to demonstrate the interaction between the exclusion region and internal macro placement under a standard layout. Fig. 7 (b) illustrates the floorplan redesign approach, wherein the design's height is increased and macro blocks are repositioned to avoid intersection with the exclusion zone. In the figure, macro 2 and macro 4 are shifted up so that that exclusion zone boundary can be accommodated in the design. The assumption underpinning this method is that the required height extension on the order of a few micrometers is negligible compared to the full-field scale (on the order of millimeters), and thus the modified plan remains compatible with the original full-field. Fig. 7 (c) depicts the macro-stitch redesign strategy, whereby the exclusion zone is allowed to intersect existing macros. This scheme obviates floorplan expansion but requires targeted macro-level redesign to maintain performance and physical integrity.

2.3 Optimizing Stitch Costs in Multi Core Designs

As discussed previously, the implementation of stitch-aware methodologies in multi-core architectures introduces distinct overhead profiles contingent upon the strategy chosen to accommodate the exclusion zone. The floorplan redesign approach principally incurs area-driven costs, wherein dimensional expansion to accommodate the exclusion zone directly translates to increased die area and consequent manufacturing expenditures. This spatial inflation may additionally precipitate performance degradation through suboptimal macro placement and elongated global interconnects, potentially compromising timing closure. Conversely, the stitched macro paradigm shifts the cost burden toward design complexity, requiring extensive macro-level modifications to maintain stitch

compatibility. In cases where multiple macro variants must be generated to accommodate diverse boundary intersections, the methodology imposes increased engineering effort for custom macro redesign, verification for performance validation along with resource allocation challenges in maintaining design consistency.

These cost matrices underscore the fundamental design-technology co-optimization (DTCO) challenge inherent in High-NA EUV stitching implementations. The selection of an optimal strategy must carefully balance area efficiency, performance preservation and minimizing design costs. While High-NA EUV lithography necessitates stitching for chip fabrication, the stitch boundary need not be rigidly positioned at the exact mid-point (50% height) of the design. The stringent alignment requirement is specifically applicable to designs utilizing full reticle field dimensions. This is due to the inherent limitation that neither of the half-field mask projections can extend beyond 16.5 mm, which constitutes 50% of the standard 33 mm reticle width. Consequently, regardless of whether the design comprises an even or odd number of dies spanning the full field, the exclusion zone must invariably be positioned at the 50% height mark of the design to ensure compatibility with the half-field mask constraints.

For smaller designs, the stitch location exhibits significant flexibility, potentially residing anywhere such that neither of masks size exceeds the half-field wafer (16.5mm). This positional flexibility introduces a critical degree of freedom in design optimization, as floorplan redesign (area and connectivity) & macro redesign can be made functions of the selected stitch location.

In this context we present a comprehensive PD design methodology that incorporates these multidimensional trade-offs for deriving optimal stitch locations for multi-core High-NA EUV designs.

Fig. 8 illustrates the detailed algorithm employed for the cost optimization 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 redesign costs necessary to accommodate the boundary. The redesign cost could be some function of the floorplan redesign cost and macro redesign cost for the given stitch boundary location x. The computed cost is then compared against the currently lowest redesign cost. If the newly calculated redesign cost is lower, it replaces the existing lowest redesign cost, and the associated stitch location is updated accordingly. This process is iterative and is repeated for multiple stitch boundary locations across the design space. In the end, the algorithm identifies and returns the stitch location that minimizes the overall redesign cost, thereby offering a cost-efficient solution for incorporating the stitch boundary into the physical design.

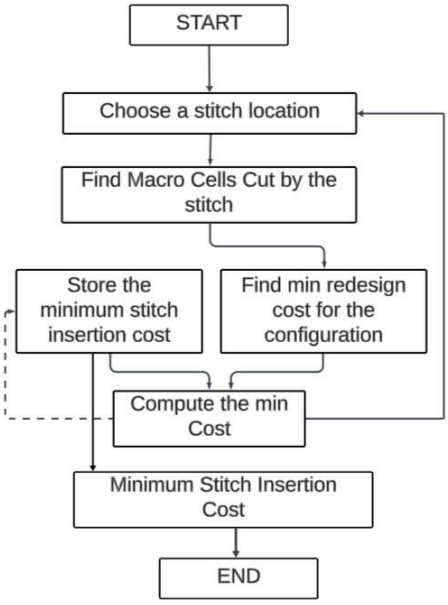

(a)

Fig 8: Cost Optimization Algorithm

2.4 Floorplan Redesign Costs

The floorplan redesign cost is primarily used for measuring the potential cost that may be incurred

due to accommodation of the exclusion zone using the floorplan redesign approach and is found by calculating the associated area & connectivity costs.

2.4.1 Area Expansion Cost

The floorplan redesign methodology involves two principal modifications to accommodate exclusion zones: (1) systematic expansion of the floorplan dimensions, and (2) strategic spatial reorganization of affected macros. This approach inherently incurs area-dependent costs proportional to the magnitude of dimensional adjustment. Crucially, macro displacement occurs exclusively along the vertical axis (Y-direction), with relocation vectors oriented either upward (positive Y) or downward (negative Y) relative to the exclusion zone boundary. Therefore, it is necessary that the exclusion zone be accommodated with minimal expansion cost.

We avoid employing the expanded area directly as the cost metric, hence we normalize the newly expanded area relative to the baseline area. This normalization yields a dimensionless cost metric that maintains proportionality with the expanded area while enabling more robust cost arithmetic within the proposed methodological framework. Such an approach enhances computational consistency and facilitates comparative analysis across varying spatial scales.

2.4.2 Connectivity Costs

The floorplan expansion involves modifying the chip dimensions and relocating macros adjacent to exclusion zones to accommodate design constraints. This spatial reorganization may displace macros vertically, potentially altering their relative position with respect to interconnected counterparts. Such topological changes can produce two opposing effects on timing characteristics (1) proximity enhancement, where reduced inter-macro distances may decrease critical path delays when previously separated macros are positioned closer together; or (2) separation-induced

degradation, where increased distances may extend signal propagation times. It is important to note that there could be a case wherein floorplan redesign reduces the macro distances. For ex. consider a baseline floorplan case as shown in Fig. 9. Due to the exclusion zone insertion, macros 2, 3 & 4 are moved up. As macro 1 is not intersecting with the exclusion zone and there is space available to place it, the designer can choose to keep it in its original location. In this case, the distance between macro 1 & 2 will decrease which could help any critical timing paths between them as shown in Fig. 9 (b) On the other hand if macro 1 has critical IO paths then it might be placed near the IO ports (at the top) and farther from macro 2 such that the separation between macro 1 & 2 increases as shown in Fig. 9 (c). Hence, proximity enhancement due to floorplan redesign is not purely speculative and can depend on various factors such as baseline floorplan, critical timing paths, etc

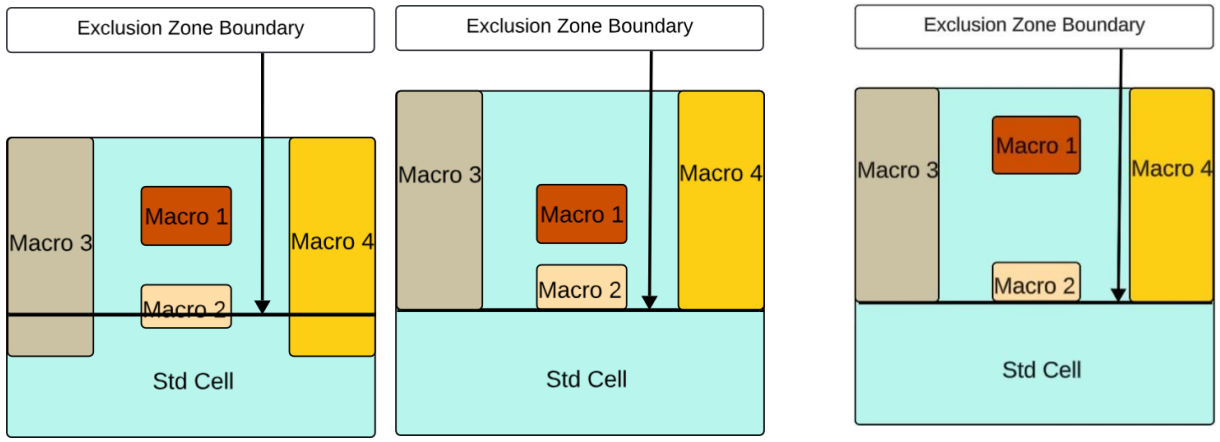

(a) (b) (c)

Fig 9: a) Baseline Floorplan b) Floorplan post-Exclusion Zone Insertion (Decreased Separation b/w macros) c) Floorplan post-Exclusion Zone Insertion (Increased Separation b/w macros)

Given the direct correlation between interconnect length and timing performance, we incorporate macro-to-macro connectivity as an additional cost metric alongside area expansion when determining the optimal vertical displacement during exclusion zone integration in the floorplan redesign approach. This metric enables systematic evaluation of whether upward or downward macro

relocation would yield net timing improvements, thereby preserving overall design performance while satisfying exclusion zone requirements.

Fig. 10 (a) presents the algorithm for determining the minimum floorplan redesign cost needed 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} and len_{DOWN}. len_{UP} & len_{DOWN} are dimensionless quantities (normalized wrt baseline height) [6] and are proportional to the macro height above or below the stitch boundary respectively.

Additionally, we also calculate the connectivity cost of each macro interacting with the exclusion zone as shown in Fig. 10 (b). For every macro interacting with the exclusion zone, we find all the macros connected to it. The flow then loops through the connected macros and creates a dictionary which stores the number of nets connecting the interacting macro and the connected macro as well as the placement coordinates of the connected macro. The placement coordinates of these connected macros are compared with the interacting macro coordinates and an upwards mobility score (UPM_i) for the i^{th} interacting macro is maintained. If the connected macro lies above the interacting macro, the number of nets connecting them is added to the UPM while if the connected macro lies below the interacting macro, the number of nets connecting them is subtracted from the UPM. This process is repeated for all the macros interacting with the exclusion zone and an upwards mobility score is calculated. This score is normalized by dividing it with the total number of nets connected to the connecting macro. The upwards mobility score is a measure of the macros tendency to move upwards based on its connectivity to other macros in the design. Higher the

[6] As width being same before and after exclusion zone insertion, normalization wrt baseline area becomes normalization wrt baseline height.

score, means that the interacting macro has more number of connections to macros placed above it and vice-versa.

As it represents the tendency to move upwards, we subtract this score directly from len_{UP} while adding it to & len_{DOWN} to arrive at the combined upward ($cost_{UP}$) or downward ($cost_{DOWN}$) movement cost respectively. If $cost_{UP} > cost_{DOWN}$ the macro is added to the set G_{DOWN} as it has a greater likelihood of being shifted downward.

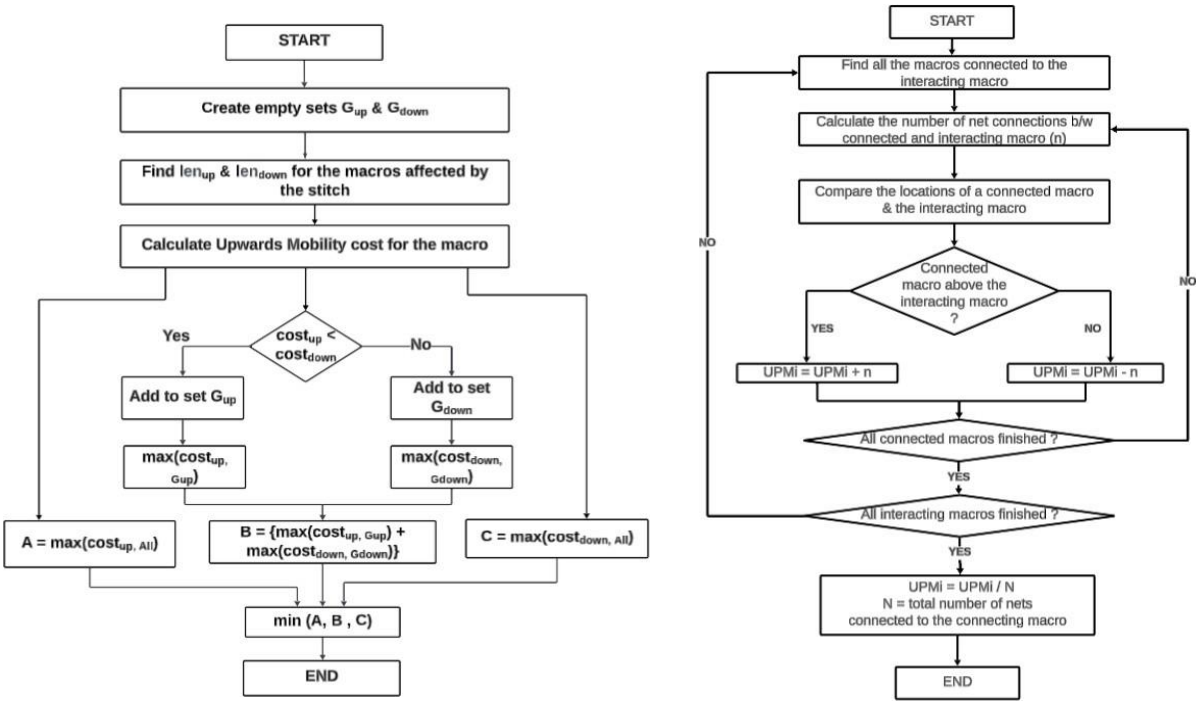

(a) (b)

Fig 10: a) Floorplan Redesign Algorithm b) Upwards Mobility Cost Algorithm

Conversely, if $cost_{UP} < cost_{DOWN}$ 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 combined costs.

The largest $cost_{UP}$ from G_{UP} & largest $cost_{DOWN}$ from G_{DOWN} are identified and denoted as as $max(cost_{DOWN,\ DOWN})$ & $max(cost_{UP,\ UP})$ respectively. These values are then summed and

represented as B such that B = max($cost_{DOWN,\ DOWN}$) + max($cost_{UP,\ UP}$).

Additionally, the maximum values of len_{UP} & len_{DOWN} across all macros intersected by the stitch boundary are computed and denoted as A = max($len_{DOWN,\ All}$) & C = max($len_{up,\ All}$) respectively. Further, we also calculate total upwards mobility cost UPM_{ALL} by summing up the individual UPM_i. UPM_{ALL} is then subtracted from A and added to B to get A_final = $cost_{UP}$ and B_final = $cost_{DOWN}$ respectively. Finally, the values A_final, B_final & C are compared, and the minimum among them determines the minimum floorplan increase required to accommodate the stitch boundary.

2.5 Macro Redesign Cost

The macro redesign methodology also introduces significant cost considerations spanning redesign efforts, sustained maintenance requirements, and performance trade-offs (spatial location, width & number of High-NA EUV layers in the exclusion zone) arising from stitch-induced design constraints. These macro-level costs are driven by multiple interdependent components that includes **redesign overhead** (design modifications to accommodate the exclusion zone boundary) , **maintenance requirements** (long term and ECO efforts due to stitch related SI) and **performance penalties** (increased net latency and congestion from detours around stitch-induced blockages).

In contrast to floorplan redesign costs, which are directly quantifiable through measurable area expansion, macro redesign costs represent a more complex category of indirect expenditures that resist straightforward quantification. These costs encompass multifaceted considerations including but not limited to engineering effort, verification overhead, and performance compromises. For the purposes of this study, we have systematically categorized and estimated these macro redesign costs in Table 4. The cost presented here highlights the design effort that went into designing these macros and hence their unit is man-hours. It is imperative to recognize that macro redesign costs are as critical as floorplan redesign costs. For instance, consider a scenario where a specific

macro is intersected by the exclusion zone at multiple unique locations one per instance. Despite each intersection individually resulting in a negligible performance impact, the necessity arises to develop and maintain distinct versions of the same macro for each instance. This requirement significantly escalates the engineering effort and verification overhead, as each macro variant must undergo comprehensive validation to ensure functional and performance integrity. Consequently, managing multiple macro versions can become a substantial bottleneck in the design process, impacting overall efficiency and time-to-market.

Table 4: Unit Level Macro Redesign Costs

Block Type	Redesign Cost
Memory Macro (ex. SRAM)	50
Core Macro	30

The total macro redesign cost for accommodating an exclusion zone boundary is given by the following:

$$Macro\ Redesign\ Cost\ = \sum_{i=1}^{n} \frac{x_i}{X_N}$$

$$x_i = i^{th}\ macro\ design\ cost$$

$$n = total\ number\ of\ macros\ interacting\ with\ the\ exclusion\ zone\ boundary$$

$$N = total\ number\ of\ macros\ in\ the\ design$$

$$X_N = Total\ Design\ Cost = \sum_{i=1}^{N} x_i$$

2.6 Joint Cost

Based on the discussion, the integration of exclusion zones in a multi-core design necessitates a balanced methodology that harmonizes the competing costs of floorplan redesign and macro

redesign. Accommodating the exclusion zone in the design incurs costs which could be in terms of area, design effort or both. Area costs are typically encountered when the design floorplan is extended to accommodate the exclusion zone. Design effort costs are typically incurred when individual macros are redesigned to accommodate the exclusion zone through them.The floorplan redesign approach causes area-driven costs, while macro redesign leads to additional overhead. We aim to design a joint cost minimization framework that strategically combines both approaches to achieve minimal joint redesign cost.

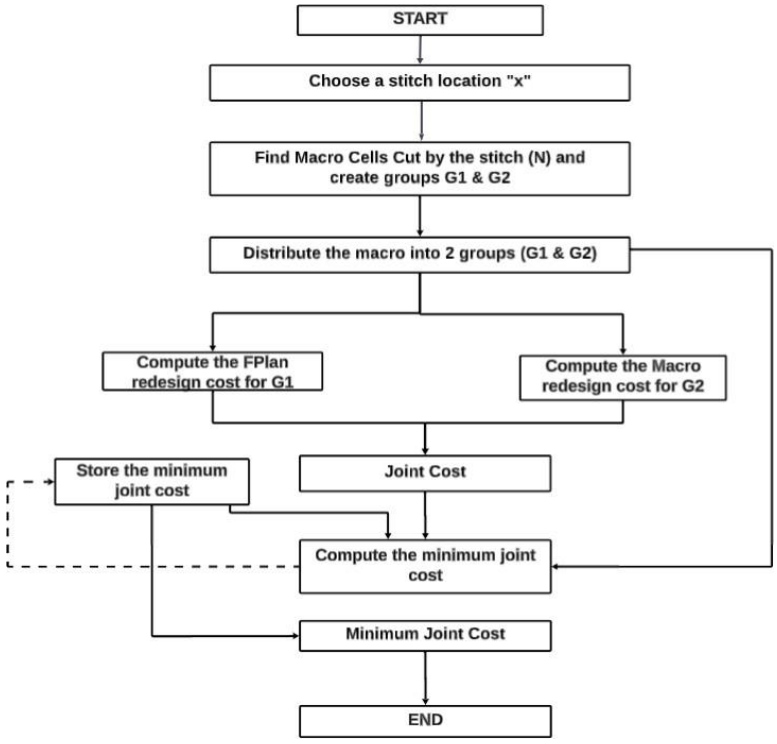

Fig 11: Joint Cost Optimization Algorithm for a given location

Fig. 11 shows the algorithm for finding the joint redesign cost for a given set of macros interacting with the exclusion zone boundary. The algorithm begins by finding all the macros (N) interacting with the exclusion zone boundary. Two groups G1 & G2 are created and all of the possible 2^N combinations of the interacting macros are created. A combination from 2^N is taken

and floorplan redesign approach is applied on G1 macros while macro redesign approach is applied on G2 macros. The two computed costs are then added together to find the joint cost which is compared with the global minimum joint cost for the interacting macro combination.

$$Joint\ Cost = \alpha * floorplan\ redesign\ cost + (1 - \alpha) * macro\ redesign\ cost$$

If the newly calculated joint cost is lower, it replaces the existing lowest joint cost. This process is iterative and is repeated for each of the 2^N combinations of the interacting macros. In the end, the algorithm identifies and returns the macro combination and the minimum joint (floorplan and macro redesign) cost for a given location.

2.7 Implementation

We test the above methodology in a modular multi-tile design (named TOP), constructed by aggregating either 14-core CORTEXM3 or MEMPOOL_GROUP tiles. Each constituent tile (Fig. 15) measures 3.2 mm × 3 mm (14 core CORTEXM3) and 3.5mm x 3mm (MEMPOOL_GROUP). Our approach systematically evaluates redesign costs across multiple scaled variants of the TOP design, with the objective of minimizing the redesign costs. The number of tiles were adjusted based on the required die size. In essence, when the die size is reduced such that a whole number of tiles can no longer span the full-field die, an entire row of tiles is removed to maintain tiling regularity and alignment within the exposure field. Fig. 12 (a) illustrates an example of a top-level design, assembled from either multiple 14-core CORTEXM3 or MEMPOOL_GROUP tiles, where the die size matches the wafer's full-field size. Similarly, Fig. 12 (b) presents another top-level design example with a die size approximately 90% of the full-field wafer size. m and n are integers can change depending upon the die size and tile size and m_n refers to either a CORTEXM3 or MEMPOOL_GROUP tile instance name.

It is important to clarify that the specified dimensions of the 14-core CORTEXM3 and MEM-

POOL_GROUP tiles are invariant with respect to die size. The design sizes selected for the purpose of this study are significantly smaller than conventional full-field dies. Due to their reduced dimensions, these standalone designs would not inherently require an exclusion zone when fabricated using High-NA EUV lithography. The primary objective of this investigation is to simulate and quantify the impact of exclusion zones such as stitch costs and floorplan redesign overheads on these smaller macro units. To this end, the full-field die is conceptually tiled with multiple instances of the 14-core CORTEXM3 or MEMPOOL_GROUP tiles. The resulting top-level structure is a synthetic aggregation of these tiles and does not represent a functional design, nor does it incorporate input/output (IO) pads or peripheral circuitry. Since the top-level assembly is not an actual design, the accommodation of an integer number of tiles within the full-field die boundaries is not a constraint. Furthermore, the floorplan and macro redesign costs are principally influenced by the exclusion zone location which is independent of the macro tile's dimension along the x-axis and thus remain unaffected by whether the tiles span the full 26 mm width of the die.

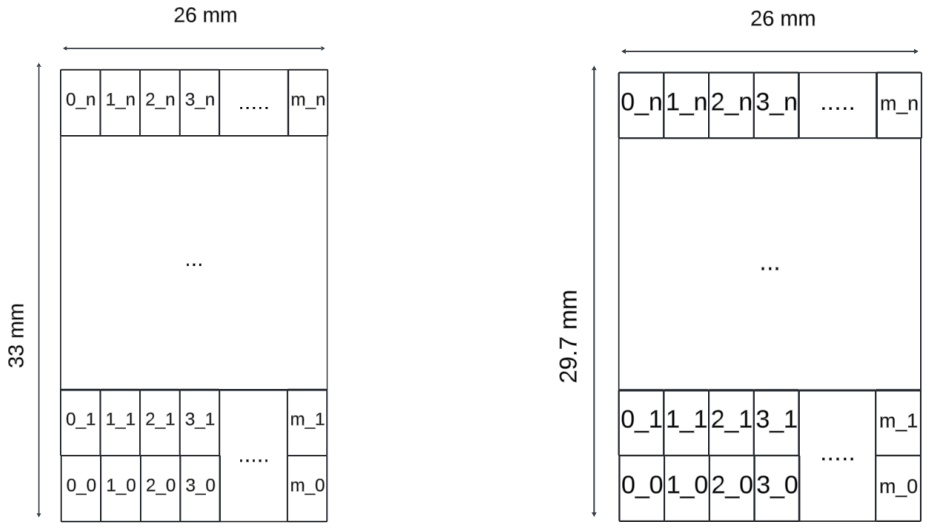

(a) (b)

Fig 12: a) Die Size is 100% FF Wafer Size b) Die Size is 90% FF Wafer Size

Fig. 13 shows the variation in different costs (floorplan redesign costs or macro redesign costs) with the die size due to insertion of an exclusion zone of width 1um. The x-axis in these graphs represents the die size as a percentage of full-field wafer size while the y-axis represents the cost (either floorplan or macro redesign) incurred for a given die size.

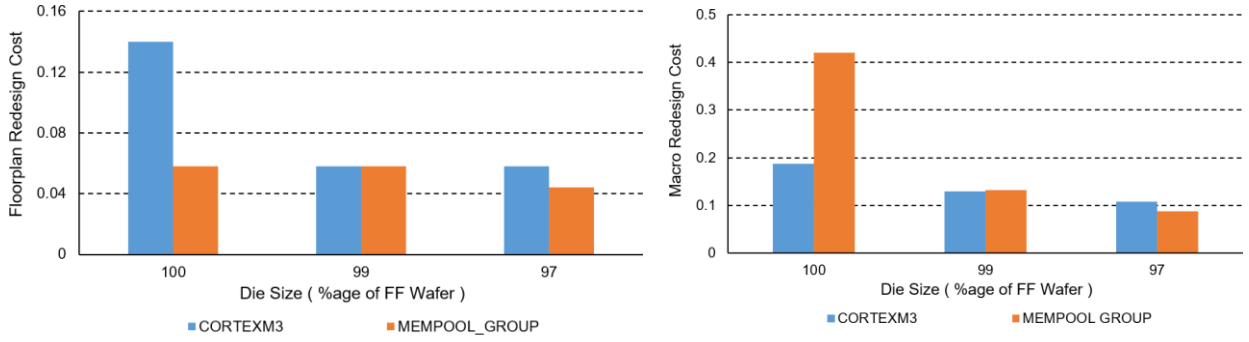

(a) (b)

Fig 13: a) Floorplan Redesign Costs b) Macro Redesign Costs

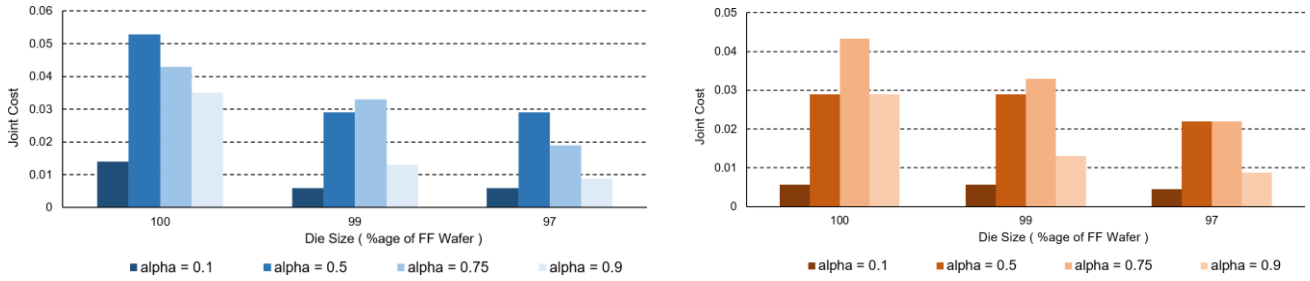

(a) (b)

Fig 14: Joint Redesign Costs for a) 14 core CORTEXM3 b) MEMPOOL GROUP

Here, floorplan redesign cost refers to normalized area expansion cost incurred by the macros due to redesigning of the floorplan. In this case only floorplan redesign approach is utilized to accommodate the exclusion zone boundary. Macro redesign refers to the normalized incremental cost incurred due to redesigning and maintenance of the macros impacted by the exclusion zone

insertion. In this approach, only the macro redesign approach is utilized to accommodate the exclusion zone boundary.

Fig. 14 shows the variation of joint cost with the die size where the x-axis represents the die size as a percentage of full-field wafer size. while the y-axis represents the joint cost incurred for a given die size. Joint cost refers to combined cost incurred when both the approaches (macro redesign for some & floorplan redesign for others) are used simultaneously. It also shows variation in joint cost with different values of alpha for the two designs. 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 exclusion zone location is fixed at a single point. As the die size decreases, the exclusion zone position becomes more flexible, resulting in lower costs. When the die size falls below a certain threshold (95% in this case), 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 in accommodating it in the design.

Looking at Fig. 13 & 14 together, one can easily see the benefits of joint cost optimization approach. For the same die size, the joint cost approach yields either lesser or similar cost than either of the floorplan redesign or macro redesign approach. In case of joint cost, the choice of approach (floorplan redesign or macro redesign) can be changed by the designer. If the designer intends to focus more on floorplan redesign approach in joint cost optimization, they can use a higher weight (alpha) value and vice-versa.

The joint redesign methodology integrates both floorplan and macro redesign costs into a unified cost function through the application of a weighting factor, denoted as α. This parameter modulates the relative contribution of each component to the total cost, thereby allowing flexibility in prioritizing either floorplan or macro redesign depending on specific design constraints. Fig. 13 (a) and 13 (b) illustrate scenarios in which only one cost component floorplan or macro redesign is

considered independently. In these cases, the design flow selects a solution optimized exclusively for the chosen cost component, without accounting for trade-offs with the other. Conversely, Fig. 14 presents outcomes under the joint redesign approach, where both cost components are considered simultaneously. It is important to emphasize that setting $\alpha = 1$ does not imply that the resulting solution consists solely of floorplan redesign. Rather, it indicates that macro redesign cost is excluded from the cost computation, although elements of macro redesign may still be present in the solution. Similarly, when $\alpha = 0$, only the macro redesign cost contributes to the total cost function, yet floorplan adjustments may still arise as part of the implementation. Since Fig. 13 (a), 13 (b), and 14 correspond to distinct optimization strategies and yield fundamentally different design solutions, the configurations represented in Figures 12 and 13 are not linearly interchangeable by simply varying the value of α between 0 and 1.

It is important to clarify that the cost analysis depicted in Fig. 13 and 14 was performed on irregular floorplan layouts as shown in Fig. 15 where macro placement does not follow a structured or ordered pattern. Fig. 15(a) & 15(b) present the initial and final floorplans for the 14-core CORTEXM3 design post-exclusion-zone insertion. Fig. 15(c) & 15(d) illustrate a similar before-and-after comparison for the MEMPOOL_GROUP architecture. Due to irregular macro placement, identifying an exclusion-zone location that avoids intersecting any macros is non-trivial. In contrast, Fig. 16(a) (14 core CORTEXM3) and Fig. 16(b) (MEMPOOL_GROUP) demonstrate orderly macro placements, where macro blocks are aligned in a structured pattern. Under such structured layouts, it becomes feasible to locate an exclusion zone that does not intersect any macros. As a result, exclusion-zone insertion may incur zero extra cost in both macro redesign and floorplan rework.

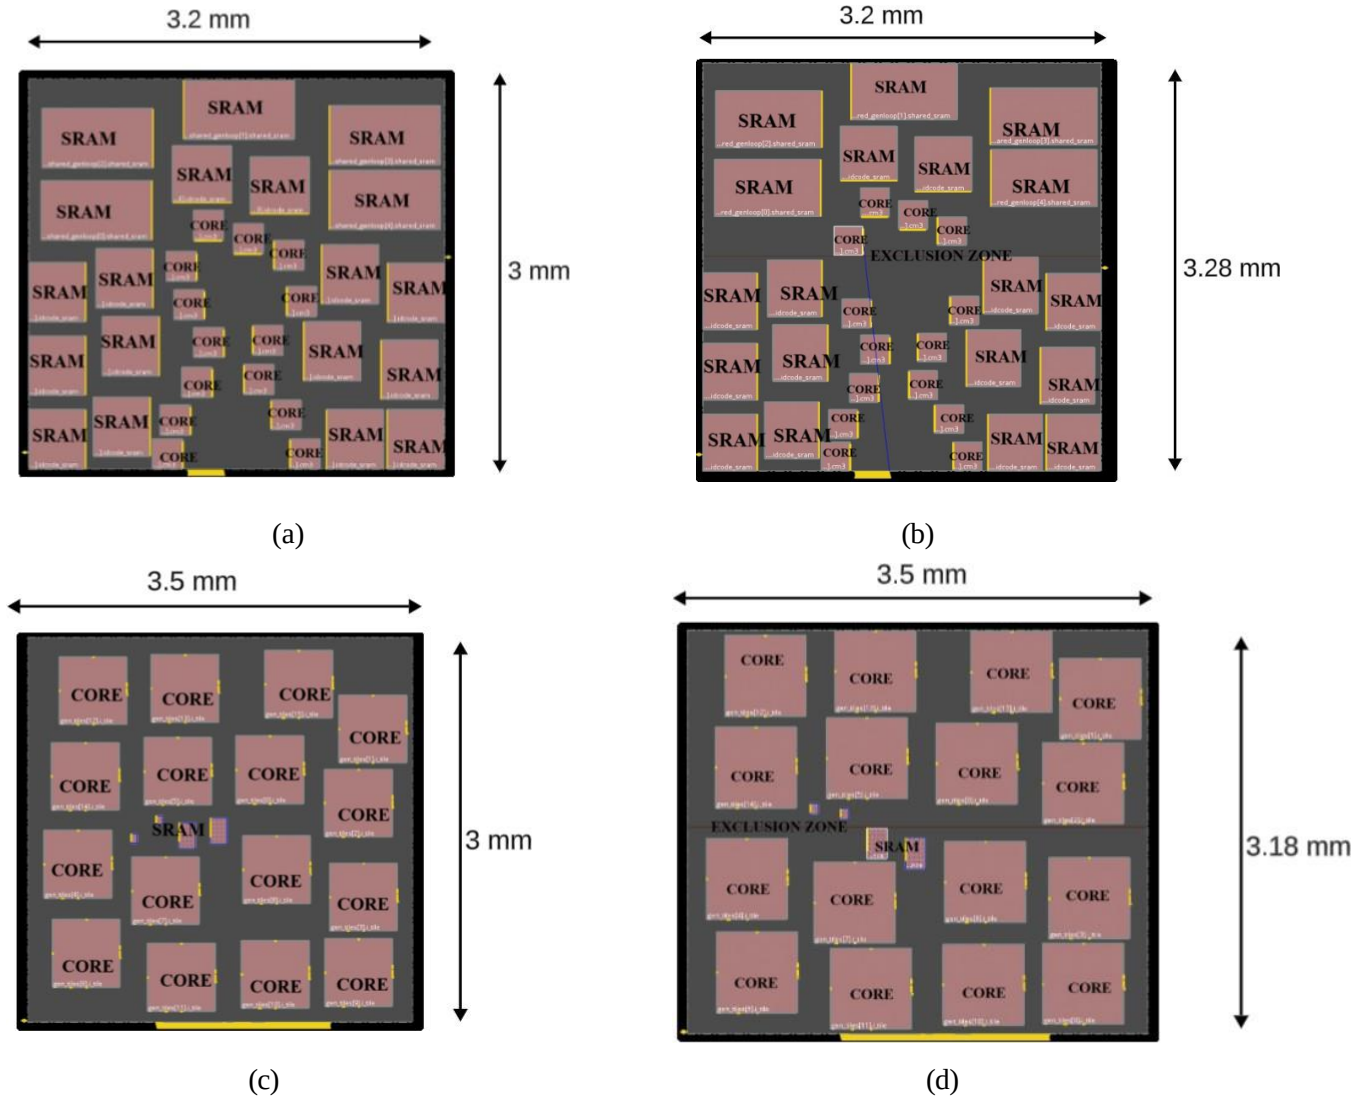

Fig 15: a) 14 core CORTEXM3 before exclusion zone insertion b) 14 core CORTEXM3 after exclusion zone insertion c) MEMPOOL_GROUP before exclusion zone insertion d) MEMPOOL GROUP after exclusion zone insertion

A notable observation from Fig. 14 is that the joint cost, defined by:

$$Joint\ Cost = \alpha * C_F + (1 - \alpha) * C_M$$

reaches its minimum at the extreme weighting values (α = 0 or 1), while peaking at intermediate values of α. Analysis of Fig. 13 (a) and 13 (b) confirm that, for the test designs under consideration, the macro redesign cost C_M is substantially larger than the floorplan redesign cost C_F. As α increases from 0, the optimization flow prioritizes minimizing the larger macro redesign cost; since

C_M remains nearly constant at this stage, even small increases in C_F cause the total joint cost to rise. This behavior produces the observed maximum in the mid-range of α. Once α surpasses a threshold dependent on the baseline floorplan and die size the floorplan cost dominates the joint function. At this juncture, the optimizer shifts its solution strategy: in order to reduce C_F it accepts higher C_M, yielding an overall decrease in C_{total} as α approaches 1. Consequently, the inverted U-shaped cost profile emerges, with the joint cost minimal at the two endpoints of α and maximal when neither cost component is fully favored.

Based on these discussions above, it is evident that the minimum cost associated with implementing an exclusion zone in a multi-core design will be influenced by three key factors: the baseline floorplan, the die size, redesign & macro redesign costs.

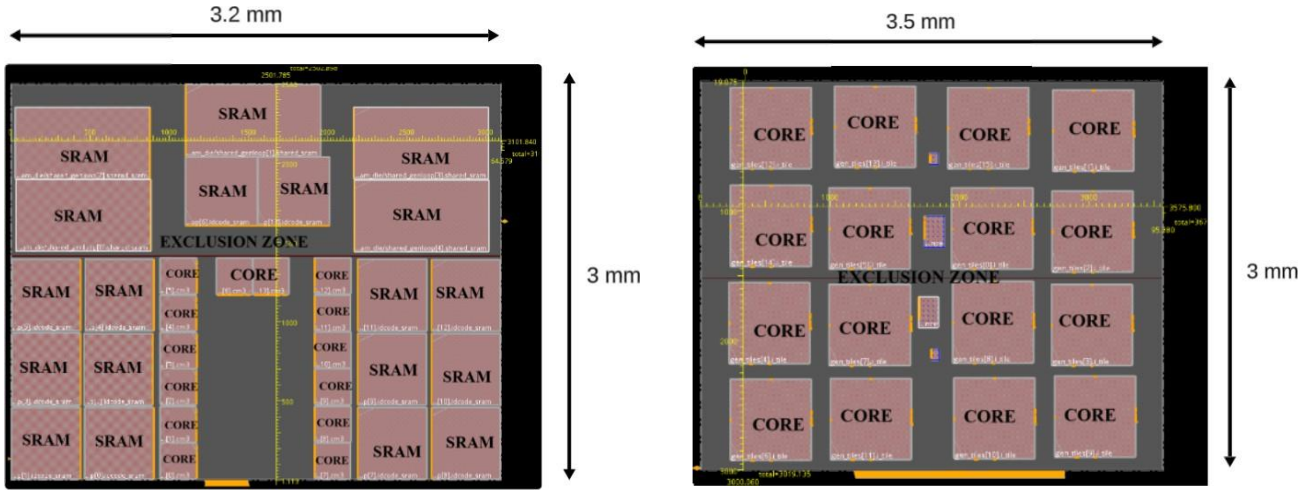

(a) (b)

Fig 16: Regularly Ordered Floorplans of a) 14 core CORTEXM3 b) MEMPOOL_GROUP Design

3 CONCLUSION

This paper introduces a stitching-aware physical design framework developed to address exclusion-zone stitching challenges in High-NA EUV lithography. At its core, the methodology defines stitch-aware design rules that balance manufacturing robustness by reducing overlay error susceptibility

with minimization of exclusion-zone insertion costs.

For single-core designs, exclusion zones are imposed via blockages on High-NA EUV layers that may cause some performance loss. However, by using the exclusion zone stitching methology solutions presented in the manuscript, the loss can be reduced to just 1% of the baseline. A Figure-of-Merit (FOM) is also proposed to estimate performance impact without requiring full place-and-route implementation .

In the context of multi-core designs, our approach incorporates various floorplanning strategies along with a stitch-boundary placement optimizer. This optimizer seeks to accommodate the exclusion zone while minimizing both floorplan redesign costs and macro redesign penalties . We evaluated three strategies: floorplan-only redesign approach, macro-stitch redesign approach, joint optimization strategy combining both cost components via a weighting factor alpha . Studies on two representative designs the 14-core CORTEXM3 and MEMPOOL_GROUP architectures demonstrate that the joint cost optimizer consistently achieves lower overall exclusion-zone insertion cost compared to either approach used individually.

4 Disclosures

The authors declare that there are no financial interests, commercial affiliations, or other potential conflicts of interest that could have influenced the objectivity of this research or the writing of this paper

5 Code, Data, and Materials Availability

Design data and TSMC40 technology library files cannot be shared. The ASAP7 technology library files are open-source and publicly available (https://github.com/The-OpenROAD-Project/asap7).

6 Acknowledgments

This work was supported by the CDEN (Center for Design-Enabled Nanofabrication) center (http://cden.ucsd.edu).

References

1 J. van Schoot, "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 W. Kaiser, "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 K. Ronse, "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 N. Davydova, L. van Look, V. Wiaux, *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 V. Wiaux, J. Bekaert, T. Kovalevich, *et al.*, "Stitching enablement for anamorphic imaging: a˜ 1um exclusion band and its implications," in *Extreme Ultraviolet Lithography 2020*, **11517**, 76–83, SPIE (2020).

6 V. Wiaux, N. Davydova, L. Van Look, *et al.*, "An experimental stitching study on the eve of high-na euv," in *Optical and EUV Nanolithography XXXVII*, **12953**, 172–182, SPIE (2024).

7 S. Jain, P. Gupta, and P. Wöltgens, "Design enablement of low-cost stitching in high-na euv patterning," in *DTCO and Computational Patterning IV*, SPIE (2025).

8 Y. J. Ban, G. Kim, and H. Shin, “Optimizing and addressing stitch challenges in high-na euv lithography for large die designs,” in *2024 21st International SoC Design Conference (ISOCC)*, 135–136, IEEE (2024).

9 Y. J. Ban, K. Cho, J. Seo, *et al.*, “Mitigating stitching-induced performance and yield losses in high- na euv lithography: Place and route implementation approaches,” in *DTCO and Computational Patterning IV*, SPIE (2025).

10 K. Miyaguchi, R.-h. Kim, A. Oak, *et al.*, “Innovative design solutions for avoiding at-resolution field stitching in high-na euv lithography,” in *DTCO and Computational Patterning IV*, SPIE (2025).

6.1 Biographies

Sagar Jain (member, SPIE) received his M.S. degree in Electrical and Computer Engineering from UCLA and is currently Physical Design Engineer at Apple Inc, USA. His current research interests include finding novel CAD solutions for chip design.

Pieter Wöltgens did his Ph.D. in solid-state physics at Utrecht University in the Netherlands. He joined the IBM T. J. Watson Research Laboratory in Yorktown Heights, NY, USA in 1994, working in VLSI design with IBM Server Division on topics like logic standard cell design, EDA, DTCO, and interfaces for 3D-stacked chips. In 2013, he joined ASML, and currently he works as assignee for ASML at imec in Leuven as well as the Netherlands, focusing on DTCO and STCO of advanced semiconductor nodes (GAA, CFET, backside power and signals), patterning with 0.33NA and 0.55NA EUV for these nodes as well as 3D chip stacking.is currently working at ASML, Netherlands. His current interests include Circuit Design and DTCO/STCO.

Puneet Gupta (Fellow, IEEE) received the Ph.D. degree from the University of California at San Diego and is currently a Professor with the ECE Department at UCLA. His current research focuses on design-technology co-optimization, system-technology co-optimization for chiplets and novel computing paradigms for hardware machine learning accelerators.

List of Figures

List of Tables