POLITECNICO DI TORINO
Repository ISTITUZIONALE

Optimizing Scan Test Economics Trade-Offs in Homogeneous 3D SICs via Cost Modeling

Original
Optimizing Scan Test Economics Trade-Offs in Homogeneous 3D SICs via Cost Modeling / Iaria, G., Bernardi, P., Bertani, C., Garozzo, G., Tancorre, V.. - (2026), pp. 1-7. (2026 IEEE 32nd International Symposium on On-Line Testing and Robust System Design (IOLTS) Polignano a Mare (ITA) 01-03 July 2026) [10.1109/iolts69666.2026.11633775].

Availability:
This version is available at: 11583/3014007 since: 2026-08-05T10:05:13Z

Publisher:
IEEE

Published
DOI:10.1109/iolts69666.2026.11633775

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright
IEEE postprint/Author's Accepted Manuscript

©2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

13 November 2026

Optimizing Scan Test Economics Trade-Offs in Homogeneous 3D SICs via Cost Modeling

Giusy Iaria
Politecnico di Torino
giusy.iaria@polito.it

Paolo Bernardi
Politecnico di Torino
paolo.bernardi@polito.it

Claudia Bertani
STMicroelectronics
claudia.bertani@st.com

Giuseppe Garozzo
STMicroelectronics
giuseppe.garozzo@st.com

Vincenzo Tancorre
STMicroelectronics
vincenzo.tancorre@st.com

***Abstract*—This paper presents a scan test cost optimization model tailored for 3D stacked integrated circuits (3D SICs), designed to reduce test cost by decreasing Wafer Sort test time. The transition from traditional planar manufacturing to vertical integration introduces cost losses when test escapes are detected late in the assembly process, leading to the loss of Known Good Dies (KGDs) stacked in defective units. To address this, a scan test cost model is proposed that integrates yield, Wafer Sort, progressive partial assembly tests, and final Package Test costs. The proposed model uses a recursive cost formulation and known weighted fault metrics to evaluate if the test time saved exclusively at the Wafer Sort level justifies the extra cost of scrapped KGD and 3D packages. The proposed model is validated using six months of high volume manufacturing data from a 20 million gate industrial automotive SoC. Although the source data comes from a planar 2D process, a simulated stacking methodology is used to mathematically evaluate virtual 4-layer and 8-layer 3D configurations. Experimental results show that progressive intermediate testing successfully mitigates the KGD scrap cost, yielding maximum economic gains of up to 3.44% for 8-layer stacks. Furthermore, coupling the framework with criticality-oriented pattern set nearly doubles these benefits, achieving overall cost reductions of 6%.**

***Index Terms*—Yield, Scan Test, Test Economics, Known Good Die, 3D SICs, Chiplets**

I. Introduction

OVER the past few decades, the semiconductor industry has witnessed an exponential rise in the complexity of Systems-on-Chip (SoCs). As physical scaling approaches its limits, the demand for improved performance and advanced features has pushed traditional planar designs toward their boundary.

As a result, chiplet integration, particularly 3D-stacked integrated circuits (3D SICs), has emerged as a practical alternative. This marks a paradigm shift in integrated circuit design, introducing a modular approach to the industry. As manufacturers struggle with rising costs and physical limitations, chiplets have emerged as a promising solution to drive innovation and performance improvement beyond conventional scaling methods [1], [2].

However, this transition significantly increases the complexity and economic costs of the testing process. In traditional 2D planar manufacturing, a test escape at the Wafer Sort phase results in the packaging of a defective die, which is subsequently screened out during the Package Test. In such scenarios, the wasted value is limited to a single die and its package [3]. In contrast, 3D integration introduces the *Known Good Die* (KGD) challenge. If a defective die escapes Wafer Sort and is integrated into a multi-layer stack, the entire assembled device must be scrapped upon failure at the final Package Test or, if present, at an intermediate test step. This results in the loss of all *good* dies attached to the defective unit, as well as the costly advanced bonding efforts.

All test phases require substantial resource allocation in automated test equipment (ATE), leading chip manufacturers to search for methods for minimizing the number of devices undergoing the entire testing process. In industrial manufacturing flows, test time reduction is often achieved by pattern reordering techniques [4] or by cutting patterns from the tail of the generated set [3], [5], [6].

As part of this work, the cutting strategy is applied exclusively during the wafer selection phase and, in order to ensure the highest quality standards for the devices, the complete set of patterns is strictly maintained during all subsequent intermediate assembly checks and final package testing, as proposed in [3]. Thus, *shifting right* some test escapes.

However, the substantial difference is that while pattern reduction during Wafer Sort brings economic benefits in 2D SoC scenarios, as demonstrated in [3], adapting this practice to 3D SIC requires a rigorous reassessment due to the greater economic impact resulting from the increased cost of the stacked test escapes [2] and the difference of the manufacturing test flow.

To address this challenge, this paper introduces a cost optimization model tailored for 3D SICs. Rather than simply estimating test costs, the proposed model is designed to optimize the test flow cost by identifying the ideal point to cut patterns from the tail of the pattern set during Wafer Sort. To accurately quantify the economic impact of reducing scan patterns and the consequent increase of test escapes, the model is structured around two core pillars:

1) **Expected cost of a Known Good Die:** Adaptation of existing metrics relating to non-uniform failure distribution

and information from sacrificial batches [3], [7] to define the base cost of a *good* chiplet before it is assembled into the final or intermediate package, quantifying the yield losses resulting from a reduction of patterns.

2) **Stack-aware recursive cost model:** Use of a recursive mathematical formulation, depending on the number of stacked layers, to take into account the presence of KGD scraps in different production flows (e.g., progressive partial assembly testing versus blind assembly), ultimately identifying the optimal point for cutting the patterns during Wafer Sort.

The proposed optimization framework is validated using six months of high-volume manufacturing data from a 20M-gate industrial automotive SoC produced by STMicroelectronics. Although the source data comes from a planar 2D process, a simulated stacking methodology is used to mathematically evaluate virtual 4-layer and 8-layer 3D configurations.

The rest of the paper is organized as follows. Section II provides the technical background on the manufacturing process of 3D chiplets and state-of-the-art cost models. Section III details the proposed methodology and the mathematical formulations of the two core pillars. Section IV presents the experimental setup, the sensitivity analysis scenarios, and the results derived from the industrial case study. Section V concludes the paper with some remarks on the key findings of the study.

II. Background

This section outlines the manufacturing concepts and literature references that drive the proposed study, highlighting the critical differences between planar and stacked testing economics.

A. Manufacturing Test Flow for Chiplet-Based SoCs

With the decline of Moore's law, disaggregation involves splitting a monolithic SoC into multiple smaller dies, or *chiplets*, which are subsequently integrated into a single package. In homogeneous 3D SICs, identical dies are stacked vertically using advanced interconnection technologies such as micro-bumps, hybrid bonding, and Through-Silicon Vias (TSVs) [8].

The manufacturing test flow for such devices, illustrated in Figure 1, relies heavily on verifying dies before they are committed to an assembly. Because dies cannot be easily removed from a stack once bonded without causing irreversible damage, a single defective component forces the entire device to be discarded [2]. Therefore, achieving extremely high fault coverage at the early Wafer Sort stage is critical.

After Wafer Sort, there are two main approaches for the assembly and test process:

- **Progressive Assembly**: Intermediate test steps (partial assembly tests) are inserted into the flow to catch early failures. This prevents the waste of Known Good Die (KGD) components by identifying defective sub-stacks before they are bonded to further layers.
- **Blind Assembly**: To maximize manufacturing throughput, some industrial flows omit the partial assembly tests and proceed directly to the complete assembly. In this strategy, testing is entirely deferred until the stack is fully completed.

To accurately reflect these varying industrial constraints, the proposed cost model is designed to be highly flexible, allowing manufacturers to explicitly configure whether these partial assembly tests are executed or bypassed.

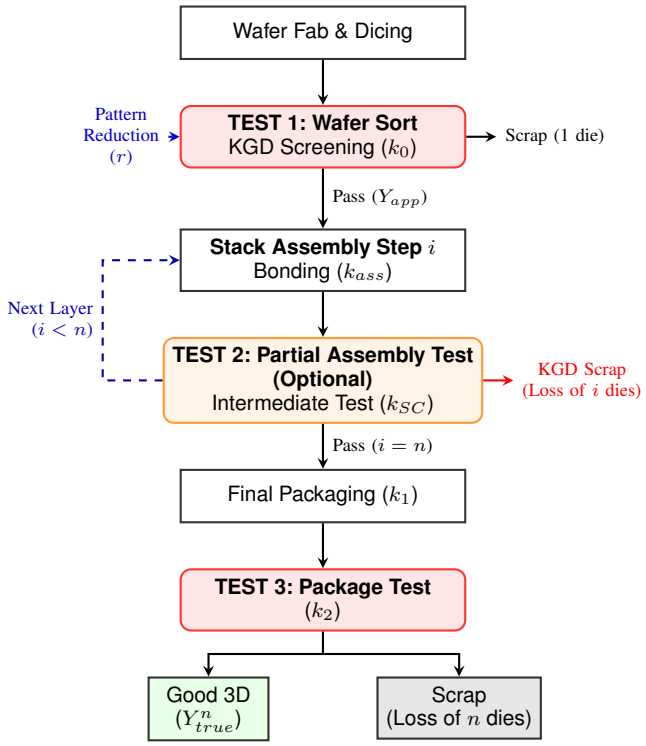


Fig. 1: Overview of a possible test flow for 3D SICs. The configurable model accounts for *Progressive* flows, where optional partial tests catch escapes early, and *Blind Assembly* flows, where these intermediate tests are omitted to maximize throughput.

B. Manufacturing Test Time Reduction and Weighted Fault Coverage

Manufacturing test execution, particularly during Wafer Sort, represents a massive portion of the overall production cost. Recent literature has demonstrated that truncating the pattern set can yield significant economic gains [3]. Since the vast majority of defects are detected by the first few patterns [6], the tail end of the pattern set covers fewer and fewer faults. By strategically pruning these last patterns, manufacturers can drastically reduce test time. If this truncation is carefully calibrated, the economic savings during the Wafer Sort will far exceed the eventual costs incurred by test escapes in the downstream assembly phases [3].

However, to accurately evaluate this economic trade off, traditional test metrics are insufficient. Standard fault coverage, derived from classical defect probability models [9], implicitly assumes that every logic gate has an equal probability of

failing. As manufacturing processes scale down, this uniform distribution model leads to severe predictive inaccuracies in complex architectures. To overcome this limitation, the concept of Weighted Fault Coverage (Ω) assigns a specific criticality weight (w_n) to each fault based on its physical and structural failure characteristics [3], [10]. For a given pattern set, it is calculated as:

$$\Omega = \frac{\sum_{n=1}^{N} w_n \cdot FD_n}{\sum_{n=1}^{N} w_n} \quad (1)$$

where N is the total number of faults, w_n represents the specific criticality weight of the nth fault, and $FD_n \in \{0, 1\}$ is a binary variable indicating whether the fault is successfully detected (1) or not (0) by the applied patterns.

A highly effective parameter for defining this criticality weight w_n is the physical layout density of the standard cells [7]. Topology information is crucial [11], as densely packed regions are statistically much more prone to manufacturing defects like critical bridging faults or electromagnetic interference issues [12]–[14]. Rather than treating all gates equally, a density metric evaluates the neighborhood of each cell within a specific distance to identify these highly critical areas [5], [7], [11]. Other potential critical areas might lie in logic gates that are difficult for the ATPG to test because of the circuit's structure [15].

By replacing standard fault coverage with Ω, manufacturers can prioritize the detection of defects located in the most critical areas. Furthermore, it is possible to exploit these criticalities to generate a set of criticality-oriented patterns in order to anticipate the coverage of these areas [3] and reduce the test cost.

C. State-of-the-Art Cost Models for Chiplet-based SoCs

For planar 2D SoCs, prior work has demonstrated that reducing scan-based patterns from the tail of an ATPG set can yield significant economic gains, provided the failure distribution is modeled using weighted defect metrics [3], [5]. As the industry shifts toward vertical integration, modeling these trade-offs becomes significantly more complex due to the stacking of multiple dies into a single device.

Recent literature has introduced various models to address the costs of 3D and heterogeneous integration. *Ahmad et al.* [8] proposed an open-source model to evaluate the business case for disaggregated systems, focusing on raw material, assembly, and general KGD costs. Similarly, *Graening et al.* introduced CATCH [16], a tool designed for early-stage co-optimization of chiplet-based systems, which looks at case studies on optimal chip size, defect density, and substrate choices.

However, as shown in Table I, while existing models provide a broad overview of 3D manufacturing test costs, they do not explicitly formulate the optimization of Wafer Sort test time through reducing applied patterns. The methodology proposed in this paper bridges this gap by providing an analytical engine that identifies the exact *cut point* for scan patterns to minimize the test costs while managing the risk of KGD scraps and their additional costs.

TABLE I: Comparison among state-of-the-art cost models for chiplet-based SoCs

Feature	**[8]**	**[16]**	**Proposed**
KGD Cost Modeling	Yes	Yes	Yes
Yield-Aware Analysis	Yes	Yes	Yes
Partial Assembly Tests	Yes	Yes	Yes
Pattern-Level Optimization	**No**	**No**	**Yes**
Optimal Cut-Point Search	**No**	**No**	**Yes**

III. The Proposed Approach

The proposed methodology introduces a comprehensive cost-benefit framework specifically tailored to optimize the test flow for 3D SICs. The fundamental motivation of this work is to reduce the number of applied patterns during Wafer Sort to optimize the overall manufacturing cost, thus *shifting right* some test escapes [3]. However, it is necessary to rigorously understand exactly where to cut the pattern set so as not to waste money due to the package of actually faulty chiplets. This precision is exceptionally crucial for 3D SICs because the cost of packaging and assembling multiple dies is significantly bigger than in standard planar scenarios. The primary objective is to balance the test time savings achieved at the Wafer Sort against the economic losses incurred when test escapes cause the scrapping of highly valuable multi-die systems [2].

To accurately reflect modern industrial manufacturing requirements, the proposed model works for both assembly flows explained in the Subsection II-A: *Progressive Assembly* and *Blind Assembly*.

The proposed model is structured around two core pillars, which directly correspond to the subsequent subsections:

- *A. Expected Cost of a Known Good Die:* Integration of previously established non-uniform defect probability metrics [3], [7] to accurately define the expected base cost of individual chiplets before they are assembled into a package, considering the potential test escapes.
- *B. Stack-Aware Recursive Cost Model:* Use of a recursive mathematical formulation to consider the cumulative cost, for each stack, of any KGD test escapes for both blind and progressive assembly scenarios, ultimately reaching the optimal cut point for the patterns to apply during Wafer Sort.

An overview of the proposed optimization workflow is illustrated in Figure 2. The subsections below explicitly detail the mathematical evolution of the model, step by step.

A. Expected Cost of a Known Good Die

To properly assess the cost risk of assembling a multi-layer stack system, it is crucial to accurately define the cost of its foundational element: the single bare die (or in this context, *chiplet*). Following standard defect probability models [9], the

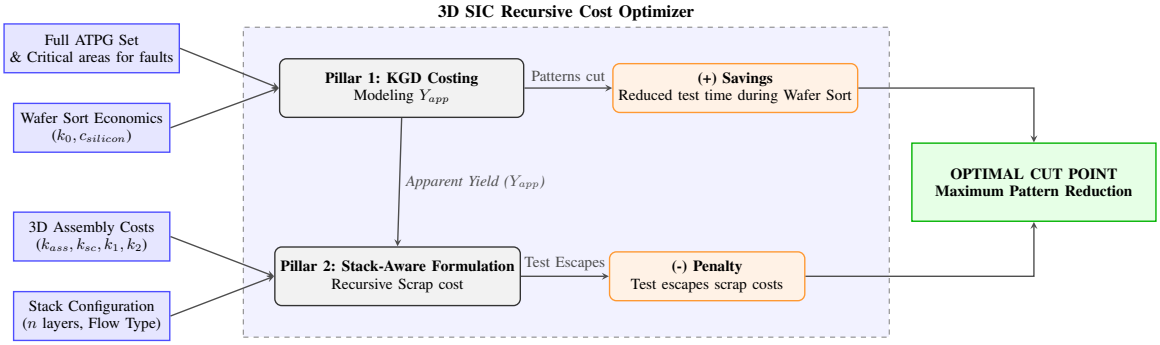


Fig. 2: Flow of the proposed stack-aware methodology for identifying the optimal Wafer Sort point where to cut patterns.

Apparent Yield (Y_{app}) of a die exiting Wafer Sort includes both truly *good* devices and test escapes [3].

Test cost models typically rely on uniform fault coverage, which assumes an equal probability for every logic fault [7]. As demonstrated in recent literature, this leads to inaccuracies in modern complex SoCs [3], [7]. To mitigate this, the proposed model leverages the Weighted Fault Coverage (Ω) methodology, which assigns a specific criticality index to each fault based on its physical probability of occurrence [3]. Consequently, the relationship between the detection of these faults and the resulting apparent yield is formulated as:

$$Y_{app} = Y_{true} + 1 - (Y_{true})^{1-\frac{\Omega(p)}{100}} \tag{2}$$

Where Y_{true} is the baseline yield. By utilizing $\Omega(p)$ instead of standard fault coverage, the calculation of Y_{app} reflects the real impact of fault detection on the number of defective dies that mistakenly enter the assembly line.

With the apparent yield established, the base cost incurred during the Wafer Sort (C_W) can be modeled dynamically based on the ratio between the applied patterns and the number of patterns in the full pattern set:

$$C_W = c_{silicon} + k_0 \cdot \frac{p_{applied}}{p_{max}} \tag{3}$$

Here, $c_{silicon}$ represents the intrinsic raw cost of the silicon. The term k_0 denotes the Wafer Sort execution cost per die. Finally, to account for the dies scrapped during Wafer Sort, the normalized expected cost of a single unit that successfully passes and is placed into the apparent KGD bin is represented as $C_P = C_W / Y_{app}$.

B. Stack-Aware Recursive Cost Model

As discussed in Subsection II-A, manufacturers may adopt different assembly strategies depending on specific throughput and quality requirements. To provide a versatile optimization framework, the proposed model formally incorporates both *Blind Assembly* and *Progressive Assembly* workflows. While the blind approach maximizes manufacturing throughput by stacking all n layers prior to a final Package Test, the progressive approach utilizes intermediate partial assembly tests to verify sub-stacks during the assembly process [2]. By accounting for both methodologies, the model allows for a comprehensive assessment of the economic trade-offs inherent in modern 3D manufacturing. Let k_{ass} represent the die-to-die bonding cost, and k_{sc} be the cost of executing a partial assembly test.

The expected cost to produce a truly *good* i-layer partial stack, denoted as $C_G^{(i)}$, must be modeled recursively to capture the growing value of the assembly. The base case for this recursion is the foundational 2-layer stack. For this initial assembly, two passed dies (C_P) are bonded and subsequently tested:

$$C_G^{(2)} = \frac{2 \cdot C_P + k_{ass} + \gamma \cdot k_{sc}}{(Y_{true}/Y_{app})^{\alpha}} \tag{4}$$

Where γ is a binary flag (1 if partial tests are performed, 0 for blind assembly) and α represents the yield loss factor. In the *Progressive Assembly* flow, $\gamma = 1$ and $\alpha = 2$ at the base stage. In the *Blind Assembly* flow, $\gamma = 0$, and costs simply accumulate. For subsequent intermediate layers ($2 < i < n$), a single passed die is bonded to the previously verified $(i-1)$-layer partial stack, and the new i-layer assembly is tested again in the progressive case:

$$C_G^{(i)} = \frac{C_G^{(i-1)} + C_P + k_{ass} + \gamma \cdot k_{sc}}{(Y_{true}/Y_{app})^{\beta}} \tag{5}$$

This recursive mathematical step is the core mechanism that captures the progressive KGD scrap cost. If the newly attached i-th die is a test escape that fails the intermediate test (k_{sc}), the manufacturer is forced to scrap the previously verified $C_G^{(i-1)}$ partial stack, losing potentially a lot of money. However, by catching the failure at layer i (where $\gamma = 1$ and $\beta = 1$), the model prevents further downstream investments in devices that are already defective. In contrast, for *Blind Assembly* ($\gamma = 0, \beta = 0$), costs accumulate linearly without filtering at intermediate stages.

The recursive assembly process continues until the final n-th layer is added to the device. At this definitive stage, instead of a simple intermediate test, the fully assembled 3D SIC undergoes final packaging (incurring cost k_1) and the exhaustive final Package Test (incurring cost k_2). Consequently, the normalized Total Manufacturing Cost ($Cost_{3D}$) function is formulated as:

$$Cost_{3D} = \frac{C_G^{(n-1)} + C_P + k_{ass} + k_1 + k_2}{(Y_{true}/Y_{app})^{m}} \cdot p_{max} \tag{6}$$

Where $m = 1$ for *Progressive Assembly* (as previous layers were already filtered) and $m = n$ for *Blind Assembly* (representing the probability that no layer contains an escape).

This independent structure forms the mathematical backbone of the proposed optimization engine. By solving for the global minimum of Equation 6, the algorithm dynamically evaluates the trade-off of the test costs. It analytically dictates exactly where to cut the patterns during Wafer Sort, determining whether the economic savings gained by reducing the applied patterns at Wafer Sort ($k_0 \cdot \frac{p_{applied}}{p_{max}}$) yield a lower overall Total Manufacturing Cost than the consequent losses triggered by allowing test escapes to reach the expensive downstream assembly and packaging tests. Normalizing the entire equation by p_{max} enables an objective evaluation of percentage gains across various stack depths, independent of the absolute manufacturing costs of the components.

IV. Experimental Results

This section validates the proposed model through experimental implementation, utilizing a sensitivity analysis across distinct industrial manufacturing flows.

A. Experimental Setup

The proposed cost optimization framework is evaluated using industrial production data provided by STMicroelectronics. The case study is an automotive SoC comprising approximately 20 million gates. The dataset encompasses six months of high-volume manufacturing, providing statistically significant tracking of defect distributions.

Exact numbers regarding the number of patterns, the yield, and the costs cannot be disclosed and percentages are reported. The results are reported for a scenario where the raw silicon and the advanced 3D packaging steps constitute the biggest part of the overall costs. The Weighted Fault Coverage (Ω) has been used for all the cost modeling formulas, as stated in Section III.

The evaluation relies on a simulated stacking methodology. Virtual 4-layer ($n = 4$) and 8-layer ($n = 8$) configurations were simulated by using the historical data for the 2D industrial case study via the recursive formulations established in Section III. To provide a rigorous comparative analysis, the total cost for each configuration is normalized independently to its own 100% baseline at the maximum numbers of patterns.

B. Sensitivity Analysis: Progressive vs. Blind Assembly

To prove the adaptability of the model and understand the impact of partial assembly testing on the optimal Wafer Sort cut point, two distinct manufacturing strategies were evaluated. Scenario A represents a *Progressive Assembly* flow, utilizing partial assembly tests to catch Wafer Sort test escapes early. Scenario B represents a *Blind Assembly* flow, where all testing is deferred until the 3D SIC device is fully assembled. Table II summarizes the results of the optimization algorithm across different stack depths, comprising also the single die 2D scenario, deeply detailed in [3].

TABLE II: Optimal pattern cut point comparison by assembly flow

Configuration	Patterns Applied	Max Cost Gain
2D Scenario (Single Die SoC)		
Classical ATPG	54%	3.12%
Criticality-oriented ATPG	8%	6.1%
Scenario A (Progressive Assembly)		
4-Layer (Classical)	55%	3.34%
4-Layer (Criticality-oriented)	10%	6.31%
8-Layer (Classical)	56%	3.44%
8-Layer (Criticality-oriented)	16%	6.08%
Scenario B (Blind Assembly)		
4-Layer (Classical)	56%	3.23%
4-Layer (Criticality-oriented)	16%	5.70%
8-Layer (Classical)	56%	3.18%
8-Layer (Criticality-oriented)	23%	5.23%

When evaluating these results, it is crucial to correctly contextualize the normalization methodology. While the maximum percentage gains may appear numerically similar between the 2D scenario (3.12%) and the 3D progressive scenarios ($\sim$3.4%), their absolute economic impact is vastly different. The 2D scenario is normalized solely against the cost of producing a single planar die. In contrast, the 3D scenarios are normalized against the total production cost of a complete multi-die 3D SIC device, which intrinsically incorporates n silicon chiplets, advanced bonding procedures, and premium packaging. Because the manufacturer successfully accumulates test time savings across all n individual dies required for the device, saving 3.44% on a highly expensive 8-layer 3D SIC yields a significantly larger absolute cost reduction than saving 3.12% on a single bare die.

Furthermore, the graphs in Figure 3 highlight a fundamental shift in the optimal test strategy for 3D SICs. For 3D SICs, the optimal cut-point is shifted to the right, toward higher number of applied patterns. This shift is driven by the KGD scrap cost; as the value of the assembled device increases, the cost of a test escape becomes so high that the model mandates a more exhaustive Wafer Sort to protect the investment.

Consequently, the window of pattern counts where reducing Wafer Sort time actually results in a net gain is much smaller for 3D SICs than for 2D. In 3D scenarios, the curve rises sharply if too many patterns are removed, indicating that the savings in ATE time are rapidly eclipsed by the cost of scrapped Known Good Dies. This proves that while 3D SICs offer higher potential gains, they also operate within a much tighter margin of error.

The results highlight a highly favorable dynamic for 3D manufacturing when progressive testing is employed. In Scenario A, because the intermediate tests successfully catch Wafer Sort test escapes before the entire stack is fully assem-

bled, the KGD scrap cost is heavily mitigated. This allows the optimization engine to safely maintain an aggressive Wafer Sort reduction, driving the maximum cost gain progressively higher as the stack grows.

Conversely, Scenario B mathematically proves the potential economic risk of blind assembly. Because a single test escape mandates the scrapping of the entire finished package and all previously bonded KGDs, the massive cost of the penalty directly erodes the Wafer Sort savings. Consequently, the maximum cost gains literally shrink as the stack grows deeper, dropping to 3.18% for the 8-layer device. To protect the highly valuable multi die assembly, the optimization algorithm is forced to become more conservative, shifting the optimal cut point further to the right to retain a higher percentage of applied patterns.

C. Impact of criticality-oriented ATPG

To further validate the economic model, the cost equation was applied using the criticality-oriented ATPG pattern set presented in [3], which detects first the most critical faults.

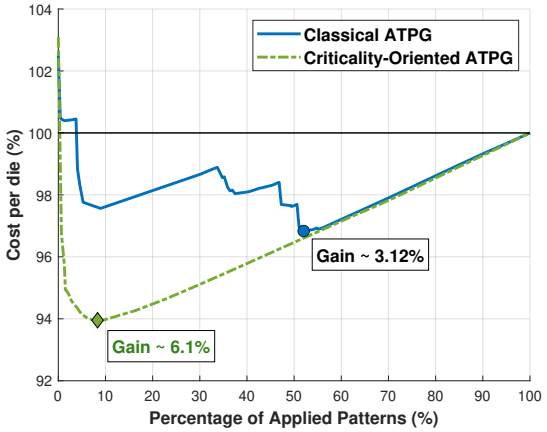
Regardless of the assembly flow chosen, criticality-oriented ATPG significantly improves the production costs of 3D SICs. By prioritizing critical failures at the beginning of the pattern set, fewer patterns can be applied while still maintaining high failure coverage. The model recognizes this behavior, reliably identifying optimal cut-off points between approximately 8% and 22% of the pattern set without triggering the KGD loss.

The use of the criticality-oriented pattern set nearly doubles the cost benefit compared to the classical uniform approach. Even in the highly restrictive 8-layer *Blind Assembly* environment, where test escapes trigger catastrophic losses, the criticality oriented pattern set yields a robust 5.23% overall cost gain. This proves that applying patterns based on defects probability distribution remains an absolute necessity for 3D test economics, compensating for the economic losses introduced by advanced packaging.

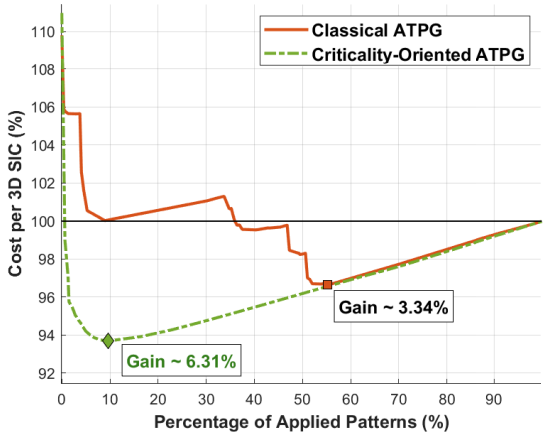
V. Conclusion

This paper introduces a stack-aware, recursive cost optimization framework tailored for 3D SICs. By evaluating both progressive and blind assembly flows, the model accurately quantifies the cost losses associated with Wafer Sort test escapes. Experimental results demonstrate that progressive intermediate testing effectively mitigates the escalating scrap losses of 3D SICs, enabling manufacturers to safely apply pattern reduction strategies to accumulate savings across the manufacturing test flow.

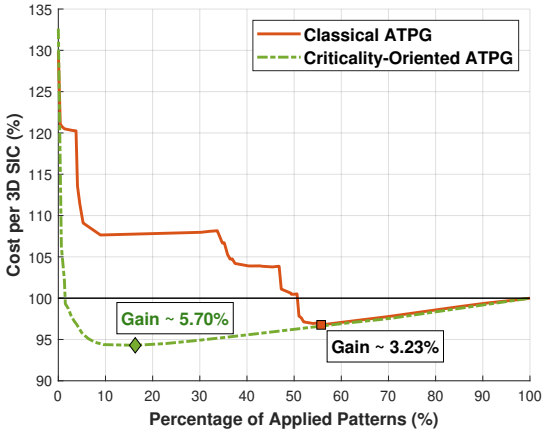
Furthermore, coupling 3D integration with criticality-oriented ATPG significantly amplifies the cost reduction compared to classical methodologies. This demonstrates that a cost modeling approach is an absolute necessity for maximizing profitability in modern advanced packaging. The proposed framework accurately identifies the optimal cut point for patterns to be applied during the Wafer Sort, minimizing the total economic cost by balancing pattern reduction with the additional costs associated with the risk of additional test escapes.



(a) 2D Scenario Test Cost per single die



(b) Scenario A: 4-Layer Progressive Assembly, Cost per 3D SIC



(c) Scenario B: 4-Layer Blind Assembly, Cost per 3D SIC

Fig. 3: Normalized Total Manufacturing Cost comparing Classical ATPG versus criticality-oriented ATPG. The optimal cut point is identified where the curves reach their global minimum.

REFERENCES

[1] H. Liu *et al.*, "Survey of chiplet technology: Soc architecture, interconnect, eda, and advanced packaging," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 15, no. 4, pp. 514–536, 2025.

[2] M. Hutner *et al.*, "Special session: Test challenges in a chiplet marketplace," in *2020 IEEE 38th VLSI Test Symposium (VTS)*, 2020, pp. 1–12.

[3] G. Iaria *et al.*, "A comprehensive scan test cost model to optimize the production of very large socs," *IEEE Transactions on Computers*, pp. 1–14, 2024.

[4] T. Song *et al.*, "Pattern reorder for test cost reduction through improved svmrank algorithm," *IEEE Access*, vol. 8, pp. 147 965–147 972, 2020.

[5] G. Iaria *et al.*, "A novel pattern selection algorithm to reduce the test cost of large automotive systems-on-chip," in *2022 IEEE 23rd Latin American Test Symposium (LATS)*, 2022, pp. 1–6.

[6] F.-F. Ferhani *et al.*, "How many test patterns are useless?" in *26th IEEE VLSI Test Symposium (vts 2008)*, 2008, pp. 23–28.

[7] P. Bernardi *et al.*, "About the correlation between logical identified faulty gates and their layout characteristics," in *2023 29th IEEE International Symposium on On-Line Testing and Robust System Design*, 2023.

[8] M. Ahmad *et al.*, "Heterogeneous integration of chiplets: Cost and yield tradeoff analysis," in *2022 23rd International Conference on Thermal, Mechanical and Multi-Physics Simulation and Experiments in Microelectronics and Microsystems (EuroSimE)*, 2022, pp. 1–9.

[9] T. Williams *et al.*, "Defect level as a function of fault coverage," *IEEE Transactions on Computers*, vol. C-30, no. 12, pp. 987–988, 1981.

[10] F. Corsi *et al.*, "Defect level as a function of fault coverage and yield," in *Proceedings ETC 93 Third European Test Conference*, 1993, pp. 507–508.

[11] W. Ruggeri *et al.*, "Innovative methods for burn-in related stress metrics computation," in *International Conference on Design Technology of Integrated Systems in Nanoscale Era*, 2021.

[12] P. Maxwell *et al.*, "Bridge over troubled waters: Critical area based pattern generation," in *IEEE European Test Symposium*, 2017.

[13] F. Hapke *et al.*, "Total critical area based testing," in *2018 IEEE International Test Conference (ITC)*, 2018, pp. 1–10.

[14] ——, "Defect-oriented test: Effectiveness in high volume manufacturing," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 40, no. 3, pp. 584–597, 2021.

[15] F. Angione *et al.*, "Automatic generation of system-level test for uncore logic of large automotive soc," *IEEE Transactions on Computers*, vol. 74, no. 9, pp. 3195–3209, 2025.

[16] A. Graening *et al.*, "Catch: A cost analysis tool for co-optimization of chiplet-based heterogeneous systems," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 45, no. 3, pp. 1091–1104, 2026.