

Reinforcement Learning for Hybrid Bonding Terminal Legalization in 3D ICs

Wanqi Ren^{*1,2}, Chengrui Gao^{*1,2}, Yunqi Shi^{1,2}, Mingzhou Fan³, Siyuan Xu³, Ke Xue^{1,2},
Chenjian Ding³, Mingxuan Yuan³, Chao Qian^{+1,2}

¹ National Key Laboratory for Novel Software Technology, Nanjing University, China

² School of Artificial Intelligence, Nanjing University, China

³ Huawei Noah's Ark Lab, China

Abstract—Hybrid bonding (HB) in 3D ICs enables scaling but introduces overlap challenges from large pitch requirements. Existing legalization methods use exhaustive sliding-window scanning, resulting in significant computational inefficiency. To address this, we propose a reinforcement learning (RL) approach that adaptively selects subregions for targeted displacement optimization. The learned policy generalizes to unseen designs without fine-tuning. Experimental results on open-source and industrial benchmarks show our method fully eliminates overlaps with minimal displacement and reduced runtime compared with baselines.

Index Terms—3D ICs, Bonding Terminal Planning, Legalization, Reinforcement Learning

I. INTRODUCTION

Face-to-Back (F2B) integration is limited by TSV area penalties [1] and MIV manufacturing constraints [2], [3]. Face-to-Face (F2F) stacking has emerged as a promising alternative. Utilizing hybrid bonding (HB) for direct metal-to-metal connections (Fig. 1(a)), this approach offers fine-pitch capabilities and high yield [4]. HB facilitates high-density heterogeneous integration, as demonstrated in recent research [5] and commercial applications [6]–[8].

Most 3D flows [9]–[11] treat terminals as vias, overlooking the disparity between large HB (1–10 μm) [12] and fine metal pitches (e.g., 0.064 μm) [13]. This mismatch causes overlaps, necessitating legalization to align terminals to a valid grid [14].

Previous works [14]–[16] formulate legalization as weighted bipartite matching. While the Hungarian algorithm guarantees optimality, its complexity is prohibitive for large designs. Heuristics like sliding windows [14] and clustering [16] reduce complexity but remain inefficient due to exhaustive, indiscriminate scanning (Fig. 1(b)). We propose an RL-based approach that adaptively selects subregions for targeted refinement (Fig. 1(c)). Initialized from a greedy solution, our agent employs a mask-based vision network and Proximal Policy Optimization (PPO) [17] to prioritize high-potential windows. By modeling window selection as a Markov Decision Process (MDP), our framework eliminates unnecessary computations and demonstrates strong generalization to unseen designs without fine-tuning. Ultimately, this approach achieves zero overlaps with substantial runtime efficiency improvements.

*Equal contribution

+Corresponding Author: qianc@nju.edu.cn




Fig. 1: (a) F2F 3D IC stacking with HB. (b) SlidingWindow [14] and (c) proposed RL method. Blue grid markers indicate legal candidate sites, while yellow dots denote ideal positions. Brown shaded regions mark the current optimization window; light brown areas indicate scanned regions.

II. PRELIMINARIES

Hybrid bonding (HB) enables fine-pitch (sub-10 μm) connections but requires strict alignment to a predefined grid [18]. However, upstream routers typically generate *ideal positions* that violate these pitch constraints [15], [16], necessitating legalization. Formally, let $\mathcal{P} = \{\mathbf{p}_i\}_{i=1}^n$ be the set of ideal HBT positions and $\mathcal{C} = \{\mathbf{c}_j\}_{j=1}^m$ be the set of legal candidate sites ($m \geq n$). The legalization problem is formulated as a minimum weight bipartite matching problem, seeking an injective mapping $f : \mathcal{P} \rightarrow \mathcal{C}$ that minimizes total Manhattan displacement:

$$\min_f \sum_{i=1}^n |\mathbf{p}_i - f(\mathbf{p}_i)|_1. \quad (1)$$

The Hungarian algorithm [19] solves Eq. (1) optimally. However, its time complexity of $O(|\mathcal{P}|^2|\mathcal{C}|)$ is prohibitive for large-scale high-density HB, requiring more efficient approaches.

III. METHODOLOGY

In this section, we introduce the proposed HBT legalization method, an iterative optimization method based on RL, which comprises two steps: (1) a greedy legalization strategy for efficiently generating a fully legal initial solution, and (2) an RL-based framework to iteratively refine the initial solution, where the RL policy is trained to identify the most promising window to minimize displacement. The algorithmic overview is presented in Algorithm 1.

REFERENCES

- [1] M. Motoyoshi, "Through-silicon via (TSV)," *Proceedings of the IEEE*, vol. 97, no. 1, pp. 43–48, 2009.
- [2] P. Batude, B. Sklenard, C. Fenouillet-Beranger, B. Previtali, C. Tabone, O. Rozeau, O. Billoint, O. Turkyilmaz, H. Sarhan, S. Thuries *et al.*, "3D sequential integration opportunities and technology optimization," in *IEEE International Interconnect Technology Conference*, 2014, pp. 373–376.
- [3] Y. Zhao, L. Zou, and B. Yu, "Physical design for advanced 3D ICs: Challenges and solutions," in *Proceedings of the International Symposium on Physical Design (ISPD)*, Austin, TX, 2025, pp. 209–216.
- [4] J. H. Lau, "Recent advances and trends in advanced packaging," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 12, no. 2, pp. 228–252, 2022.
- [5] T. Hiratani, N. Fujiwara, N. Inoue, M. Kurokawa, T. Kikuchi, T. Mitarai, T. Nitta, T. Watanabe, H. Fujikata, N. Nishiyama *et al.*, "SOA integrated InP/Si hybrid tunable laser by utilizing chip-on-wafer hydrophilic bonding," in *Proceedings of the Silicon Photonics Conference*, 2024, pp. 1–2.
- [6] J. Wu, R. Agarwal, M. Ciraula, C. Dietz, B. Johnson, D. Johnson, R. Schreiber, R. Swaminathan, W. Walker, and S. Naffziger, "3D V-Cache: the implementation of a hybrid-bonded 64MB stacked cache for a 7nm x86-64 CPU," in *Proceedings of the International Solid-State Circuits Conference (ISSCC)*, San Francisco, CA, 2022, pp. 428–429.
- [7] W. Gomes, S. Khushu, D. B. Ingerly, P. N. Stover, N. I. Chowdhury, F. O'Mahony, A. Balankutty, N. Dolev, M. G. Dixon, L. Jiang *et al.*, "8.1 Lakefield and mobility compute: A 3D stacked 10nm and 22FFL hybrid processor system in 12× 12mm 2, 1mm package-on-package," in *Proceedings of the International Solid-State Circuits Conference (ISSCC)*, San Francisco, CA, 2020, pp. 144–146.
- [8] D. B. Ingerly, S. Amin, L. Aryasomayajula, A. Balankutty, D. Borst, A. Chandra, K. Cheemalapati, C. Cook, R. Criss, K. Enamul *et al.*, "Foveros: 3d integration and the use of face-to-face chip stacking for logic devices," in *Proceedings of the International Electron Devices Meeting (IEDM)*, 2019, pp. 19–6.
- [9] L. Bamberg, A. García-Ortiz, L. Zhu, S. Pentapati, S. K. Lim *et al.*, "Macro-3D: A physical design methodology for face-to-face-stacked heterogeneous 3D ICs," in *Proceedings of Design, Automation & Test in Europe Conference*, Grenoble, France, 2020, pp. 37–42.
- [10] P. Vanna-Iampikul, C. Shao, Y.-C. Lu, S. Pentapati, and S. K. Lim, "Snap-3D: A constrained placement-driven physical design methodology for face-to-face-bonded 3d ics," in *Proceedings of the International Symposium on Physical Design (ISPD)*, New York, NY, 2021, pp. 39–46.
- [11] B. W. Ku, K. Chang, and S. K. Lim, "Compact-2D: A physical design methodology to build commercial-quality face-to-face-bonded 3D ICs," in *Proceedings of the International Symposium on Physical Design (ISPD)*, Monterey, CA, 2018, pp. 90–97.
- [12] C. Netzbund, K. Ryan, Y. Mimura, S. Ilseok, H. Aizawa, N. Ip, X. Chen, H. Fukushima, and S. Tan, "0.5 μm pitch next generation hybrid bonding with high alignment accuracy for 3D integration," in *Proceedings of the Electronic Components and Technology Conference (ECTC)*, 2023, pp. 1100–1104.
- [13] "Nangate technology node," <https://si2.org/open-cell-library/>.
- [14] S. Pentapati, A. Agnesina, M. Brunion, Y.-H. Huang, and S. K. Lim, "On legalization of die bonding bumps and pads for 3D ICs," in *Proceedings of the International Symposium on Physical Design (ISPD)*, Virtual Event, USA, 2023, pp. 62–70.
- [15] S. Liu, J. Jiang, Z. He, Z. Wang, Y. Lin, B. Yu, and M. Wong, "Routing-aware legal hybrid bonding terminal assignment for 3D face-to-face stacked ICs," in *Proceedings of the International Symposium on Physical Design (ISPD)*, Taipei, Taiwan, 2024, pp. 75–82.
- [16] G. Kim and H. Park, "Clustering-driven bonding terminal legalization with reinforcement learning for F2F 3D ICs," in *Proceedings of Asia and South Pacific Design Automation Conference (ASP-DAC)*, Tokyo, Japan, 2025, pp. 308–314.
- [17] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv:1707.06347*, 2017.
- [18] J. Derakhshandeh, L. Hou, I. De Preter, C. Gerets, S. Suhard, V. Dubey, G. Jamieson, F. Inoue, T. Webers, P. Bex, G. Capuz, E. Beyne, J. Slabbekoorn, T. Wang, A. Jourdain, G. Beyer, K. June Rebbis, and A. Miller, "Die to wafer 3D stacking for below 10um pitch microbumps," in *Proceedings of IEEE International 3D Systems Integration Conference*, San Francisco, CA, 2016, pp. 1–4.
- [19] S. Martello and P. Toth, "Linear assignment problems," in *Surveys in Combinatorial Optimization*, 1987, vol. 132, pp. 259–282.
- [20] T. Ajayi, V. A. Chhabria, M. Fogaça, S. Hashemi, A. Hosny, A. B. Kahng, M. Kim, J. Lee, U. Mallappa, M. Neseem *et al.*, "Toward an open-source digital flow: First learnings from the openroad project," in *Proceedings of Design Automation Conference (DAC)*, Las Vegas, NV, 2019, pp. 1–4.
- [21] P. Liao, Y. Zhao, D. Guo, Y. Lin, and B. Yu, "Analytical die-to-die 3-D placement with bistratal wirelength model and GPU acceleration," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 43, no. 6, pp. 1624–1637, 2024.
- [22] Y. Shi, C. Gao, W. Ren, S. Xu, K. Xue, M. Yuan, C. Qian, and Z.-H. Zhou, "Open3DBench: Open-source benchmark for 3D-IC backend implementation and PPA evaluation," *arXiv:2503.12946*, 2025.