AI Guided Multi-Objective Heterogeneous Chiplet Placement for Advanced Packaging

Partha Pratim Kundu [a], Sezin Kircali Ata[a], Furen Zhuang [a], Yubo Hou [a], Mihai Dragos Rotaru [b], Rahul Dutta [b], Ashish James [a]

[a] *Institute for Infocomm Research (I^{2}R), A*STAR*, Singapore kundupp@i2r.a-star.edu.sg

[b] *Institute for Microelectronics (IME), A*STAR, Singapore*

1. Introduction

The rising demand for green computing together with applications of high-performance, low-power, and compact systems like machine learning and autonomous vehicles, has elevated the need for optimal integration of chiplets through advanced packaging techniques [1]. Generic rule-based placement or random stochastic methods take long to converge in scenarios with multi-objective optimization (MOO) and large search space. Navigation through search space must be driven by active learning using multi-objective agents, where one agent performs chiplet placement actions and the other critic agent feedback to the actor on the placement wirelength and package temperature profile. This multi-objective RL approach intelligently arranges chiplets to reduce wirelength and improve thermal management, resulting in more efficient, reliable next-generation electronic systems that outperform traditional placement algorithms.

2. Methodology

The methodology of the proposed multi-objective Actor-Critic framework is shown in Fig. 1. The process begins by drawing batches of input states from a proportional replay buffer. These states are processed by either a GCN (Graph Convolutional Neural Network) (A2C-GCN) or NN (Neural Network) (A2C-NN) based agent. Additionally, when using GCN as the RL agent, the framework effectively incorporates interconnectivity between chiplets by treating both node and edge features. The node features provide a detailed representation of each macro's spatial characteristics, such as orientation, relative positioning within the grid, and spatial interactions with other macros whereas edge features are derived from the connectivity configuration between macros, offering a more comprehensive understanding of how each macro relates to the overall system's thermal and connectivity optimization. The Actor Local agent predicts potential chiplet placements and rotations and interacts with the interposer environment to compute wirelength and temperature metrics. The Mask module handles infeasible placement choices. A Differential Reward with a Potential mechanism evaluates the performance based on wirelength and thermal profiles, guiding the system to minimize hotspots and interconnect lengths. Simultaneously, the Critic agent, aided by the Information Gain Computation Unit, evaluates the predicted reward against actual outcomes to feedback on the Actor's loss. This feedback is crucial for improving the Actor's placement predictions. The system utilizes gradient extraction and the N-Step Reward Predictor to effectively adjust future actions. This adaptive approach extracts actionable insights from chiplet states and connections resulting in performance improvement over traditional placement methods like simulated annealing (SA). This iterative process continues as both the actor and critic refine their performance, achieving optimized chiplet placements while addressing thermal and connectivity challenges dynamically.

3. Result and Discussion

The proposed methodology's effectiveness is validated through extensive simulations and comparisons with traditional methods like SA across the three systems described in [2]. However, due to space limitations, only a specific case study featuring a multi-GPU system with 2 GPUs, 1 CPU, and 3 high-bandwidth memory (HBM) modules on a 50 mm x 50 mm interposer are benchmarked here. The aim is to optimize the placement of all chiplets to minimize wirelength, boost signal integrity, reduce latency, and manage thermal distribution for improved performance. The optimized placements determined by A2C-NN RL for these systems are shown in Fig. 2. The performance comparison with SA, A2C-GCN and A2C-NN are tabulated in Table 1. A2C-GCN optimized wirelength, achieving 92.4 m, outperforming SA by 10% and A2C-NN by 8% in multi-GPU system. However, it raised the temperature to 91.4°C, compared to 89.6°C for SA and 89.9°C for A2C-NN, showing trade-offs in thermal management. These performances showcase A2C's ability to balance performance trade-offs for optimizing complex chiplet placement, particularly in minimizing wirelength for better system performance.

4. Conclusion

The Actor-Critic reinforcement learning framework presented here leverages both GCNs and NNs to address key multi-objective challenges in chiplet placement by optimizing thermal management and reducing wirelength. This impacts the downstream cooling solution requirements and will be able to reduce the power consumption of heterogeneous compute infrastructure. GCNs excel in modelling spatial relationships and chiplet connectivity, outperforming traditional approaches. Extensive simulations confirm that the framework effectively balances trade-offs between wirelength and temperature, offering a highly efficient solution for chiplet placement in high-performance, thermally constrained systems and this approach is extendible to other types of advanced packages such as wirelength reduction from chip-to-wafer hybrid bonding and to other mechanical or electrical constraints through multi-agent actor-critic frameworks.

Acknowledgments

This research is supported by Agency for Science, Technology and Research (A*STAR) under its MTC Programmatic Funds (Grant No. M23M3b0064).

AI Guided Multi-Objective Heterogeneous Chiplet Placement for Advanced Packaging

Partha Pratim Kundu [a], **Sezin Kircali Ata**[a], **Furen Zhuang** [a], **Yubo Hou** [a], **Mihai Dragos Rotaru** [b], **Rahul Dutta** [b], **Ashish James** [a]

[a] *Institute for Infocomm Research (I^2R), A*STAR*, Singapore kundupp@i2r.a-star.edu.sg

[b] *Institute for Microelectronics (IME), A*STAR, Singapore*

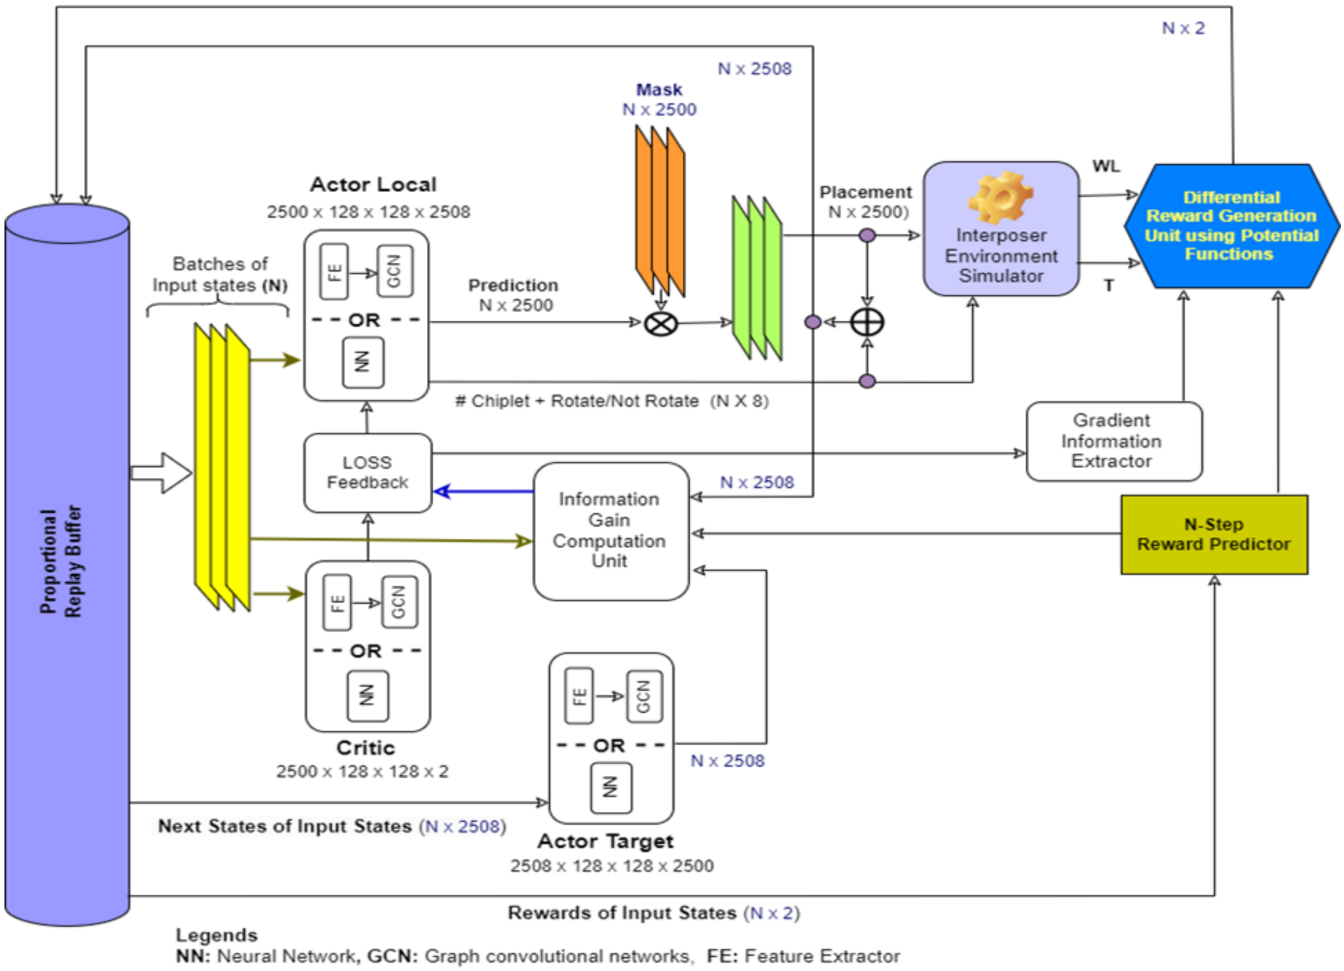

Fig. 1: A2C RL framework that leverages connectivity and thermal behavior to optimize chiplet placement in heterogeneous 2.5D systems.

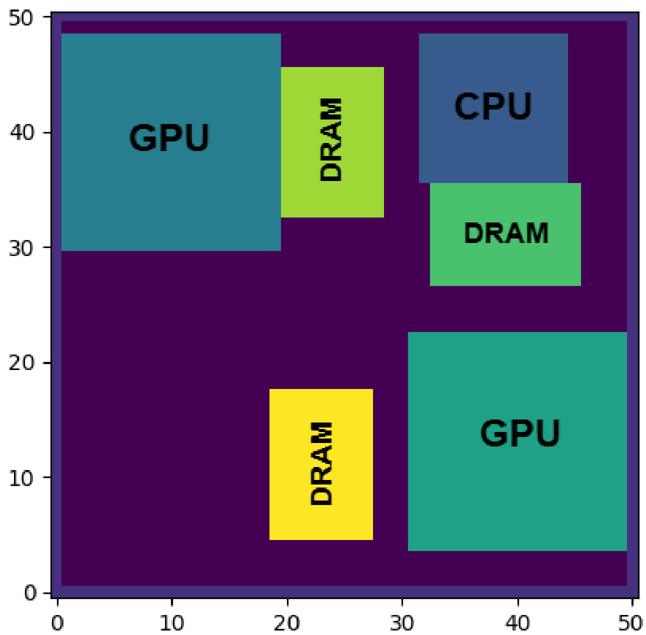

Fig. 2: Optimized chiplet placement determined by A2C-NN RL framework for multi-GPU system

Table 1: Performance comparison of both RL frameworks with SA.

Case Study		SA	A2C-GCN	A2C-NN
Multi-GPU	WL (m)	101.8	**92.4**	97.9
	T (°C)	**89.6**	91.3	89.9

References

[1] Weidong Cao et al., “Domain knowledge-infused deep learning for automated analog/radio-frequency circuit parameter optimization. In Proc. 59th ACM/IEEE DAC'22, NY-USA 1015–1020.

[2] Y. Ma et al., “TAP-2.5D: A Thermally Aware Chiplet Placement Methodology for 2.5D Systems,” in Proc. DATE, France, 2021, pp. 1246-1251.