2026 Design, Automation & Test in Europe Conference (DATE 2026)

Hetero-ChipletSim: Bridging Chiplet, Interconnect and Packaging Heterogeneity in Multi-Chiplet System Simulation

Xuguang Yuan[1], Jiangyuan Gu[1]*, Qidie Wu[1], Yang Hu[1], Shaojun Wei[1], Shouyi Yin[1,2]*
[1]School of Integrated Circuits, BNRist, Tsinghua University, Beijing, China
[2]International Innovation Center of Tsinghua University, Shanghai, China

***Abstract*—With the end of Moore's Law, multi-chiplet systems have emerged as a promising solution featuring heterogeneity across chiplets, interconnects and packaging. Existing simulators lack support for such multi-level heterogeneity, making accurate architectural exploration difficult. We propose Hetero-ChipletSim (HCS), a simulation methodology that directly integrates heterogeneous chiplet models while incorporating die-to-die(D2D) interconnect and packaging effects, enabling fast and accurate evaluation of multi-chiplet systems. Sensitivity analysis provides insights into design trade-offs under heterogeneous integration.**

***Index Terms*—Multi-Chiplet System, Heterogeneous Modeling, Modular Simulation, Packaging**

I. INTRODUCTION

Entering the post-Moore era, modern computing systems increasingly adopt chiplet-based advanced packaging to improve cost efficiency, manufacturing yield, and design flexibility [1]–[5]. Multi-chiplet systems enable the integration of dies from different process nodes and introduce multi-level heterogeneity across chiplets, die-to-die (D2D) interconnects, and packaging [6]–[8]. Such heterogeneity spans diverse chiplet models and timing abstractions, varied interconnect protocols, and advanced 2.5D/3D packaging technologies such as CoWoS, EMIB, and TSV-based 3D stacking, all of which collectively have a significant impact on system power, performance, and area (PPA) [9].

Simulators are essential for early-stage system design, but the increasing adoption of chiplet-based architectures challenges traditional, domain-specific tools [10]–[13], which lack support for heterogeneity and require costly model translation across abstraction levels. To address these issues, we propose Hetero-ChipletSim (HCS), a simulation methodology that directly accepts heterogeneous chiplet models and connects them via an interconnect layer, *LinkSim*, enabling flexible composition without translation. By modeling packaging, D2D interconnects, and packaging-induced effects within a unified framework, HCS bridges multi-level heterogeneity and provides an accurate, scalable and extensible foundation for system-level performance and power evaluation under the More-than-Moore paradigm.

II. HETERO-CHIPLETSIM METHODOLOGY

As depicted in Figure 1, the simulation framework of HCS consists of three layers: **system component layer**, **integration layer** and **evaluation layer**. In the system component layer, users select the components of a heterogeneous multi-chiplet system. The integration layer focuses on simulation-oriented modeling of the system, where a set of feasible packaging options is generated based on the chiplet connectivity. Different packaging schemes can then be selected. Finally, in the evaluation layer, actual simulations are carried out and system performance report are obtained.

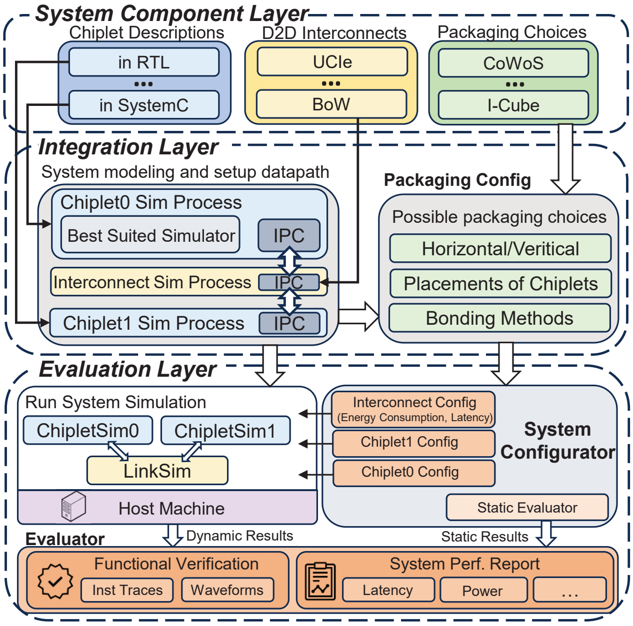

Fig. 1. Hetero-ChipletSim's modeling and evaluation flow, and different layers in the framework.

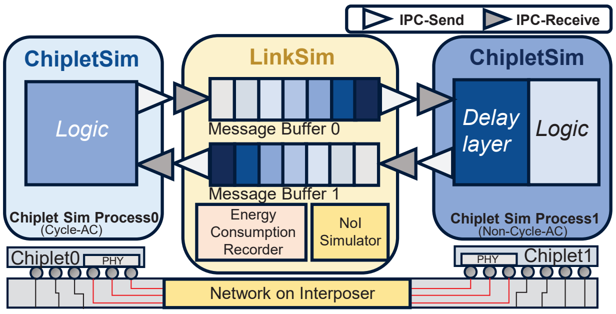

Fig. 2. Detailed HCS simulation mechanism (Top) and corresponding 2.5D packaging architecture (Bottom).

Figure 2 shows how HCS accounts for heterogeneity at different levels and the detailed simulation mechanism. **At the chiplet level**, HCS decouples chiplet modeling details from inter-chiplet communication via *LinkSim*, enabling chiplets described in different languages and timing abstractions to be simulated by their native simulators while remaining cycle-synchronized through Inter-Process-Communication. **At the interconnect level**, *LinkSim* relays D2D messages and models diverse protocols, capturing transmission latency, buffering, and energy overheads from PHYs, bumps/TSVs, and interposer

---

*Corresponding authors: {gujy, yinsy}@tsinghua.edu.cn

978-3-9826741-1-7/DATE26/©2026 EDAA

links. **At the packaging level**, heterogeneity is abstracted through extension modes, chiplet placement, and bonding methods, with a system configurator translating these choices into *LinkSim* configurations for both static analysis and dynamic simulation. Centered around *LinkSim*, HCS provides a scalable and flexible methodology for evaluating heterogeneous multi-chiplet designs.

HCS is implemented using a hybrid C++/Python framework: core simulation components such as *LinkSim* and IPC mechanisms are developed in C++, while system configuration and post-simulation analysis are handled in Python for flexibility and ease of use. HCS adopts a multi-process simulation model based on MPI for inter-chiplet communication, with deadlock avoided by alternating send/receive orders across MPI ranks. By combining MPI-based process parallelism with simulator-level parallelism (e.g., Verilator), HCS achieves good scalability and high simulation efficiency.

III. EXPERIMENTAL RESULTS

Our experimental environment runs on a dual-socket Intel Xeon E5-2620 v3 platform with Ubuntu 20.04.1, using Verilator 5.008 and OpenMPI 4.0.3. We build a multi-chiplet computing platform consisting of a CVA6-based core chiplet [14] and a cache chiplet modeled using a simulator [15]. Three representative benchmarks are selected: *multiply* (compute-intensive), *mt-matmul* (memory-intensive), and *towers* (latency-sensitive). To evaluate the sensitivity of HCS to multi-level heterogeneity, we consider different chiplet models and D2D interconnects (UCIe, BoW, and AIB) with parameters derived from public sources [6]–[8], as well as 2.5D/3D packaging configurations whose energy parameters are adopted from [16].

Figure 3(a) presents HCS simulation results with different cache chiplet models, reporting benchmark execution cycles and host simulation time. Four models are evaluated: RTL, SystemC, a C++ functional model (CPP), and a C++ functional model with an HCS delay layer (CPP-o). The results show a clear trade-off between simulation accuracy and speed. The C++ functional model achieves the fastest simulation but underestimates execution time by 30.3% on average compared to RTL. By incorporating the HCS delay layer, CPP-o significantly improves accuracy: in the *towers* benchmark, its execution time differs from RTL by only 2.6%, while still reducing host simulation time by 12%. SystemC offers a balanced compromise between accuracy and speed, demonstrating HCS's flexibility in supporting heterogeneous chiplet models.

Figure 3(b) illustrates the execution time and energy overhead of multi-chiplet designs running three benchmarks under different D2D protocols, with energy consumption broken down into bonding (micro-bump and hybrid bonding), TSVs, interposer links, and PHY-layer energy. All experiments adopt a 2.5D packaging configuration with RDL interconnects [17]. The results show that for the compute-intensive *multiply* benchmark, all three protocols exhibit similar performance, whereas UCIe's high-bandwidth capability makes it more suitable for memory-intensive workloads. BoW benefits the memory latency-sensitive *towers* benchmark due to its lightweight interface and low-latency protocol stack. Across all cases, PHY-layer energy dominates the total energy consumption, consistently accounting for more than 50%, and although UCIe incurs higher energy overhead, its performance advantages remain significant.

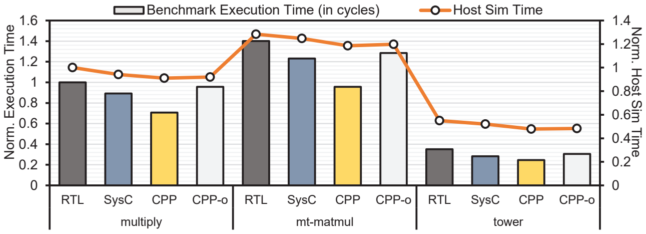

(a) Benchmark execution time and host simulation time with different cache chiplet models.

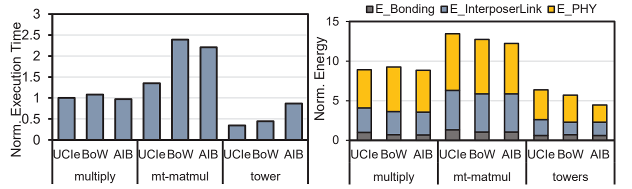

(b) Benchmark results and energy consumption breakdown with different D2D interconnects.

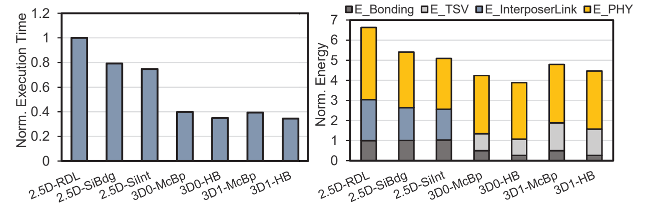

(c) Benchmark results and energy consumption breakdown with different packaging options.

Fig. 3. HCS evaluation results and host simulation time with different (a) cache chiplet models, (b) D2D interconnects and (c) packaging techniques.

Figure 3(c) shows the execution time and multi-chiplet energy overhead of the CVA6–cache system using the UCIe D2D protocol under different packaging techniques for the *mt-matmul* benchmark, with all results normalized to the 2.5D-RDL baseline. We consider three 2.5D options (RDL, silicon bridge (SiBdg) [18], and passive interposer (SiInt) [19]) and four 3D options formed by microbump (McBp) or hybrid bonding (HB) [9] with cache-on-top (3D0) or core-on-top (3D1) stacking. The results show that 3D designs achieve an average of 1.7× speedup in execution time. In terms of energy, TSV-based vertical integration is more energy-efficient than interposer-based 2.5D approaches, and hybrid bonding further reduces bonding energy compared to microbump. However, core-on-top stacking increases TSV energy by 1.65× due to the higher power demand of the core, indicating that optimized 3D stacking order is critical for reducing TSV power overhead.

IV. CONCLUSION

In this paper, we present Hetero-ChipletSim (HCS), a simulation methodology for heterogeneous multi-chiplet systems. HCS models heterogeneity across interconnects and advanced packaging, enabling direct integration of heterogeneous chiplet models for system-level evaluation. Experimental results show that HCS accepts different chiplet models, captures protocol-specific interconnect behavior, and quantifies packaging-induced performance and power impacts, providing useful insights into system-level performance–power trade-offs.

ACKNOWLEDGMENT

This work was supported by National Key R&D Program of China (2023YFB4405101), National Natural Science Foundation of China (U24B20164, U24A20234, 62125403), and Shanghai Municipal Science and Technology Major Project.

REFERENCES

[1] S. Naffziger, N. Beck, T. Burd, K. Lepak, G. H. Loh, M. Subramony, and S. White, "Pioneering chiplet technology and design for the amd epyc™ and ryzen™ processor families: Industrial product," in *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, pp. 57–70, IEEE, 2021.

[2] C.-C. Lee, C. Hung, C. Cheung, P.-F. Yang, C.-L. Kao, D.-L. Chen, M.-K. Shih, C.-L. C. Chien, Y.-H. Hsiao, L.-C. Chen, *et al.*, "An overview of the development of a gpu with integrated hbm on silicon interposer," in *2016 IEEE 66th Electronic Components and Technology Conference (ECTC)*, pp. 1439–1444, IEEE, 2016.

[3] G. Paulin, P. Scheffler, T. Benz, M. Cavalcante, T. Fischer, M. Eggimann, Y. Zhang, N. Wistoff, L. Bertaccini, L. Colagrande, *et al.*, "Occamy: A 432-core 28.1 dp-gflop/s/w 83% fpu utilization dual-chiplet, dual-hbm2e risc-v-based accelerator for stencil and sparse linear algebra computations with 8-to-64-bit floating-point support in 12nm finfet," in *2024 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits)*, pp. 1–2, IEEE, 2024.

[4] Y. S. Shao, J. Clemons, R. Venkatesan, B. Zimmer, M. Fojtik, N. Jiang, B. Keller, A. Klinefelter, N. Pinckney, P. Raina, *et al.*, "Simba: Scaling deep-learning inference with multi-chip-module-based architecture," in *Proceedings of the 52nd annual IEEE/ACM international symposium on microarchitecture*, pp. 14–27, 2019.

[5] A. Ottaviano, R. Balas, T. Fischer, T. Benz, A. Bartolini, and L. Benini, "Controlpulplet: A flexible real-time multicore risc-v controller for 2.5-d systems-in-package," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 2025.

[6] D. D. Sharma, G. Pasdast, Z. Qian, and K. Aygun, "Universal chiplet interconnect express (ucie): An open industry standard for innovations with chiplets at package level," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 12, no. 9, pp. 1423–1431, 2022.

[7] S. Ardalan, H. Cirit, R. Farjad, M. Kuemerle, K. Poulton, S. Subramanian, and B. Vinnakota, "Bunch of wires: An open die-to-die interface," in *2020 IEEE Symposium on High-Performance Interconnects (HOTI)*, pp. 9–16, IEEE, 2020.

[8] B. Wang, G. Chen, L. Y. Foo, and X. Jiang, "Signal and power integrity design of advanced interface bus (aib) for fpga packages," in *2023 IEEE 32nd Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*, pp. 1–3, IEEE, 2023.

[9] J. Kim, L. Zhu, H. M. Torun, M. Swaminathan, and S. K. Lim, "Micro-bumping, hybrid bonding, or monolithic? a ppa study for heterogeneous 3d ic options," in *2021 58th ACM/IEEE Design Automation Conference (DAC)*, pp. 1189–1194, IEEE, 2021.

[10] H. Zhi, X. Xu, W. Han, Z. Gao, X. Wang, M. Palesi, A. K. Singh, and L. Huang, "A methodology for simulating multi-chiplet systems using open-source simulators," in *Proceedings of the Eight Annual ACM International Conference on Nanoscale Computing and Communication*, pp. 1–6, 2021.

[11] G. López-Paradís, B. Li, A. Armejach, S. Wallentowitz, M. Moretó, and J. Balkind, "Fast behavioural rtl simulation of 10b transistor soc designs with metro-mpi," in *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 1–6, IEEE, 2023.

[12] P. Iff, B. Bruggmann, M. Besta, L. Benini, and T. Hoefler, "Rapidchiplet: A toolchain for rapid design space exploration of chiplet architectures," *arXiv preprint arXiv:2311.06081*, 2023.

[13] S. Herbst, N. Moroze, E. Iglesias, and A. Olofsson, "Switchboard: An open-source framework for modular simulation of large hardware systems," *arXiv preprint arXiv:2407.20537*, 2024.

[14] F. Zaruba and L. Benini, "The cost of application-class processing: Energy and performance analysis of a linux-ready 1.7-ghz 64-bit risc-v core in 22-nm fdsoi technology," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 27, pp. 2629–2640, Nov 2019.

[15] D. Million, N. Oliete-Escuín, and C. Fuguet, "Breaking the memory wall with a flexible open-source l1 data-cache," in *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 1–2, IEEE, 2024.

[16] A. Coskun, F. Eris, A. Joshi, A. B. Kahng, Y. Ma, A. Narayan, and V. Srinivas, "Cross-layer co-optimization of network design and chiplet placement in 2.5-d systems," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 39, no. 12, pp. 5183–5196, 2020.

[17] M. J. Kim, S. H. Lee, K. L. Suk, J. G. Jang, G.-J. Jeon, J.-i. Choi, H. J. Yun, J. Hong, J.-Y. Choi, W. J. Lee, *et al.*, "Novel 2.5 d rdl interposer packaging: A key enabler for the new era of heterogenous chip integration," in *2021 IEEE 71st Electronic Components and Technology Conference (ECTC)*, pp. 321–326, IEEE, 2021.

[18] L. Kang, J. Kim, B. Lee, J. Park, H. J. Yu, S. H. Lee, J. A. Lim, M. J. Lee, S. H. Han, J. K. Lee, *et al.*, "Single and multi npu chiplet heterogeneous integration packaging based on fanout rdl interposer with silicon bridge technology," in *2024 IEEE 74th Electronic Components and Technology Conference (ECTC)*, pp. 1603–1608, IEEE, 2024.

[19] M. Sunohara, T. Tokunaga, T. Kurihara, and M. Higashi, "Silicon interposer with tsvs (through silicon vias) and fine multilayer wiring," in *2008 58th Electronic Components and Technology Conference*, pp. 847–852, IEEE, 2008.