Co-Engineering Packaging Architecture for Next Gen Data Center Xeon SoC

Sandhya Arikatla, Vishal Chandrasekhar, Judy Amanor-Boadu, Yanjie Zhu, Raul Enriquez, Arun Chandrasekhar, Amruthavalli P Alur
Intel Corporation
2200 Mission College Blvd
Santa Clara, CA 95054 USA
Ph: 408-765-8080;
Email: sandhya.arikatla@intel.com

Abstract

The rapid evolution of AI and High-Performance Computing (HPC) has driven the need for advanced modular packaging architectures (2D, 2D+ 2.5D, 3D & 3.5D). As workloads become more demanding and time-to-market (TTM) accelerates, robust, scalable, and efficient processor packages are needed. The innovations in advanced packaging architectures are intended to aid scalability and performance but also face challenges from increased SoC area, higher pin-counts, larger package form factors with more layers, higher packaging/assembly costs and lower manufacturing yields resulting from greater substrate stack complexity. This paper presents a comprehensive co-design methodology for multi-chiplet processor packaging that integrates physical design optimization, signal integrity enhancement and power delivery network improvements. Our approach addresses the fundamental challenge of maintaining performance while reducing package complexity and cost through innovative routing topologies, advanced via architectures, pinmap optimizations and optimized power distribution strategies. The integrated methodology demonstrates significant cost reduction while meeting next-generation performance targets, providing a scalable framework for future multi-chiplet processor implementations.

Key words

Advanced packaging, chiplet packaging, co-design methodology, power delivery networks, signal integrity (SI), Time-to-market (TTM)

I. Introduction

The semiconductor industry is experiencing a fundamental shift from monolithic processor designs to disaggregated chiplet architecture driven by the physical and economic limitations of Moore's Law scaling [1]. Modern Xeon processors must support increasingly demanding Artificial Intelligence (AI) and High-Performance Computing (HPC) workloads that require substantial increases in compute performance, memory bandwidth and I/O connectivity. However, traditional monolithic SoC approaches face critical constraints including reticle size limitations (typically 858mm^2 for advanced nodes), exponentially increasing manufacturing costs, and yield degradation for large die areas [2].

Chiplet-based architecture offers a promising solution by decomposing monolithic design into smaller, specialized dies that can be manufactured independently and integrated using advanced packaging technologies [3,4]. These architectures enable heterogeneous integration of compute dies, I/O dies and accelerator tiles using various packaging approaches including 2D (side-by-side placement), 2.5D (silicon interposer-based), 3D (vertical stacking), and emerging 3.5D (hybrid vertical-horizontal), integration methodologies [5,6,7]

Advanced packaging has evolved from passive integration platforms into active system enablers that directly influence critical performance parameters. Modern packaging

architecture now plays a decisive role in signal integrity (SI), power integrity (PI), thermal management, mechanical reliability and overall compute efficiency. This transformation has fundamentally changed system design approaches, requiring unprecedented collaboration between traditionally separate engineering disciplines.

The demand for higher performance has driven significant increases in package complexity: die areas exceeding reticle limits, increasing I/O pin counts and increasing substrate layer counts. These larger, more complex packages introduce strong cross domain interactions spanning physical design, signal integrity and power delivery that cannot be optimized independently.

Increasing memory channel counts and high-speed I/O requirements create routing congestion that traditionally requires additional substrate layers, significantly increasing cost and complexity. Higher pin densities and increased via stacking at C4 bump levels introduce crosstalk and signal degradation that limits achievable data rates and timing margins. Growing current demands and reduced substrate layer counts create power delivery bottlenecks that impact voltage regulation and transient response.

This paper presents an integrated co-design methodology that addresses these challenges through:

1. A novel 2x density DDR routing topology that reduces substrate layers while improving signal integrity through optimized channel isolation and via stacking reduction.
2. An interleaved high-speed serial I/O routing approach that maintains compliance with loss and crosstalk specifications in reduced layer count substrates. Pinmap optimization to provide crosstalk improvements.
3. A dual-entry power delivery architecture with strategic via placement and decoupling optimization that improves power integrity despite layer reduction.
4. A modular "copy-exact" layout methodology that accelerates time-to-market across product variants.

II. DDR Physical design techniques:

Next-generation Xeon architecture requires substantially higher memory bandwidth, leading to increased memory channels. This has significantly increased the routing density on the package. Historically, packages employed a single-channel-per-layer breakout scheme on each side of the die, as shown in "Fig 1(a) and (b)". however, scaling this approach with rising channel counts necessitates additional routing layers, resulting in significant cost and stack-up complexity. Multiple micro-vias stacked at C4 bump locations create significant crosstalk and remains a dominant limiter for DDR channel margins.

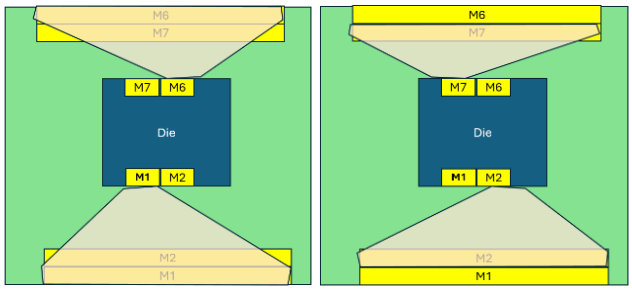

(a) (b)

Fig.1 (a) & (b) Single channel per layer per side breakout scheme

To address these scaling constraints and DDR channel margin degradation at the C4 bump level, a 2x higher density routing topology is implemented as shown in "Fig 2 (a)". The key enablers to higher density package routing are SOC floorplan alignment of memory channel stacking and a ~12mm real estate from die edge to package edge. In this approach, two different channels are routed on the same layer with signal integrity optimized channel to channel isolation using ground microvia stitching. Adjacent memory channels are fanned out on the same routing layer, and stacked memory channels are fanned out on different routing layers. A signal integrity optimized transition via pattern is created to bring the channels to be routed on the same routing layer as shown in "Fig.2(a) and 2(b)". This approach has reduced the micro via stacking at the C4 bump level, resulting in significant improvement of Far-End Crosstalk (FEXT) by ~25mV. "Fig. 3" shows the vertical crosstalk improvement that led to margin improvement to overall DDR channels.

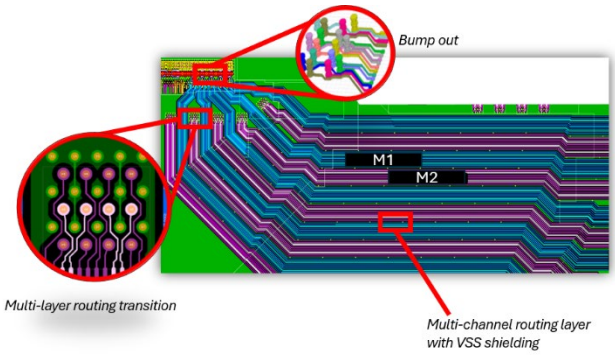

Fig. 2 (a). 2x higher density routing per layer with optimized channel to channel isolation.

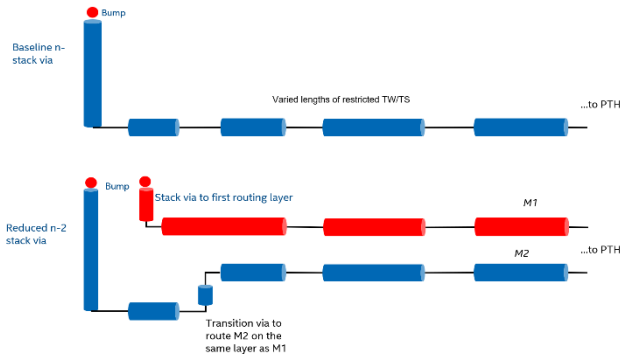

Fig. 2(b) Stacked via transition for traditional vs optimized stackup

Additionally socket pin patterns were chosen based on the aggressor count to further minimize coupling and pin count optimization. A coupled-via [8] architecture was further implemented to mitigate far-end crosstalk (FEXT) due to inductive coupling between tall socket pins. Capacitive coupling was introduced with coupled-via structures to compensate for inductive effect, yielding measurable improvements in overall DDR channel margins. "Fig 4(a)" Write (WR) shows big EH margin improvement, but slight EW degradation, still comparable margins. "Fig 4(b)" Read (RD) show both EH and EW improvement, both shows healthy margin. "Fig. 4(c)" Read (RD) with channel-to-channel Xtalk shows significant margin improvement between baseline and optimized stackup with design optimizations.

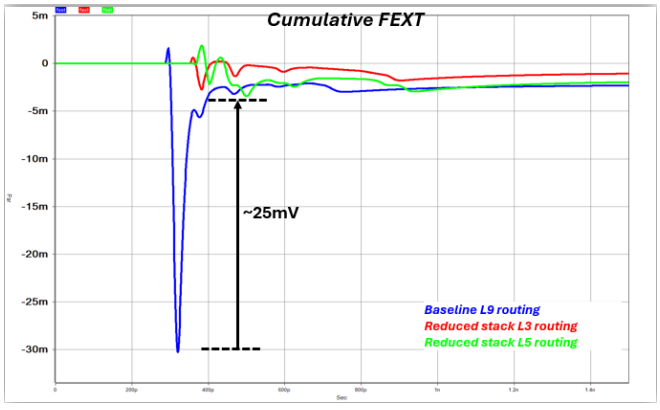

Fig. 3 Far-End Crosstalk (FEXT) for baseline vs reduced microvia stack at C4 bump level.

These optimizations collectively enabled a reduction of two-layer pairs from the package stack while sustaining SI robustness.

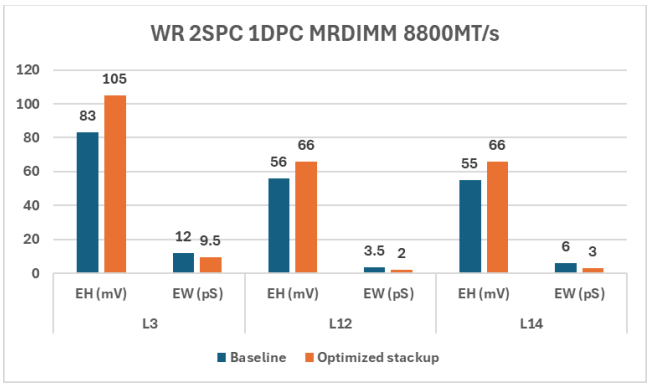

(a)

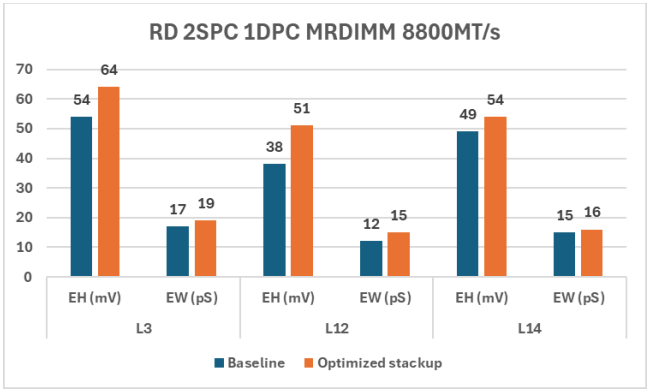

(b)

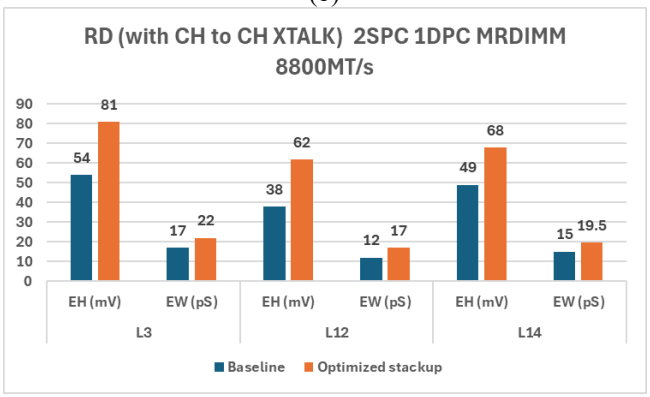

(c)

Fig. 4. EW and EH margins for WR and RD

III. High-Speed Serial IO physical design techniques:

High-Speed Serial IO is a loss sensitive interface dictating shorter package lengths. Historically the biggest challenge with the High-Seed Serial IO routing is the SOC floorplan mismatch to the socket pinmap, which is optimized for the board layer count. This drives the layer count on the package as RX and TX are routed on different layers. As new physical design methodologies in DDR enabled a reduction of two-layer pairs, innovative routing solutions for High-Speed

Serial IO are needed to meet the layer count reduction. An interleaved RX/TX routing methodology was adopted as shown in "Fig. 5" to align with the target layer count reduction.

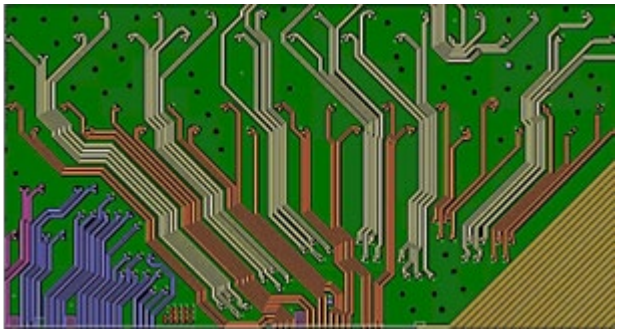

Fig. 5. High-Speed Serial IO interleaved TX/RX routing.

Interleaving minimizes the overall trace length, supports compliance with channel-loss and skew specifications. Interleaving requires enhanced crosstalk control by inserting ground isolation between RX/ TX signal pairs and ground via placement adjacent to signal vias. Signal integrity optimized PTH patterns are implemented to maintain compliance with Near-End Crosstalk (NEXT) and Far-End Crosstalk (FEXT) specifications.

Additionally, the zig-zag pin pattern "Fig. 6(a)" is an alternative package pinout arrangement designed to improve pin count efficiency and real-estate utilization in high-speed differential interfaces by controlling aggressor orientation and minimizing crosstalk. In this configuration, each differential pair is surrounded only by quasi-orthogonal aggressors, effectively eliminating worst-case coupling scenarios and reducing the dominant sources of differential crosstalk [9]. As a result, the zig-zag pattern achieves improved crosstalk performance relative to previous generation patterns. Furthermore, the zig-zag structure has been demonstrated to be practical for implementation, particularly in edge regions of the pin map, making it a viable solution for PCIe 5.0 interfaces while maintaining cost and manufacturability constraints.

In comparison, an alternative approach, commonly referred to as the twin-ax configuration "Fig. 6(b)", may offer superior crosstalk performance; however, it requires four ground pins per differential pair, whereas the zig-zag pattern requires only two. This results in a substantially higher pin count for the twin-ax approach, with approximately 35% more pins required, as illustrated in "Fig. 6(a) & (b)". Consequently, insertion loss can be reduced through shorter routing paths. These design optimizations enable compliance with signal integrity requirements without the need for additional routing layers.

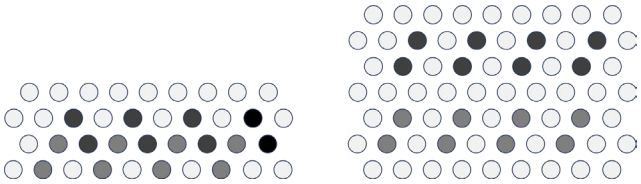

Fig. 6. Implementation of eight differential pairs in (a) zig-zag and (b) twin-ax package pinout patterns.

IV. Power delivery advancements:

The rapid growth in compute density has significantly increased power consumption [10], placing a substantial burden on power delivery networks (PDNs), particularly given limited substrate real estate constraints. Reducing layers directly impacts PDN performance by increasing impedance due to the loss of dedicated power and ground planes. There is a greater sensitivity to voltage droop and transient response, increased risk of current density problems and imbalance, and potential degradation in power integrity due to routing constraints. To mitigate this, a critical innovation is to introduce enhanced power delivery corridors that improve effective resistance paths (R_path) between voltage regulators and the package.

By rethinking package current entry points and distribution symmetry, the design reduces resistive losses despite fewer substrate layers. Power delivery efficiency in a reduced substrate count is significantly improved by balancing voltage regulator phases across the package, enabling dual entry power delivery and strategic placement of pins under dies. "Fig. 7". illustrates this power delivery design concept. The dual entry power scheme lowers the R_path, improving performance without requiring additional package layers. To further optimize PDN performance, careful construction of microvia and PTH patterns are implemented based on the transient analysis for compute and I/O workloads.

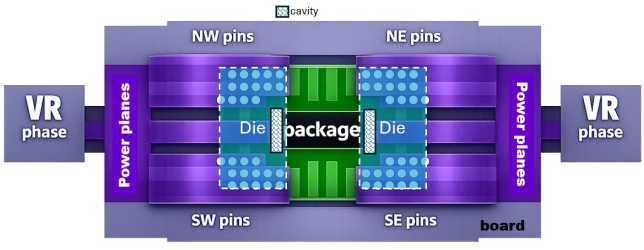

Fig. 7. Dual entry power delivery with strategically placed pins and cavity

To support higher current slew rates, package cavity size and placement is strategically designed to ensure decoupling capacitors are placed in optimized areas to support workload demands as shown in “Fig. 7”. Mechanical optimizations, such as adding notches to the Integrated Heat Spreader (IHS) to accommodate bulk capacitors are also implemented to create a shorter power delivery path from capacitors to die. This concept is illustrated in “Fig. 8”.

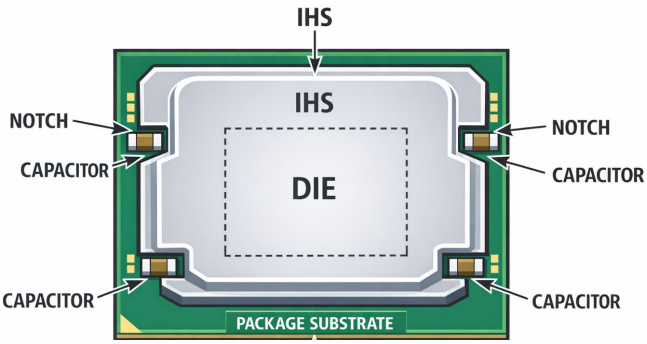

Fig. 8. Illustration of IHS notches with bulk capacitors

To further streamline time to market across products with varying power and socket pin requirements, a modular “copy exact” layout methodology was adopted to reduce validation overhead. By reusing common layout building blocks, aligning critical power structures, and making targeted power-delivery tradeoffs where needed, the methodology reduced unique validation requirements per product, shortened post-silicon validation effort, and improved resource efficiency, demonstrating how disciplined layout reuse can accelerate time to market for complex multi-die products.

V. Conclusion

This paper presents a comprehensive co-design methodology for multi-chiplet processor packaging that successfully addresses the critical challenges of routing density, signal integrity and power delivery in next-generation Xeon processors. Adoption of these methodologies and techniques results in improved performance, supporting higher TDPs, reduced package layer counts, cost optimization and meeting TTM. Reducing the micro-via stacking and implementing higher density routing enables a reduction of two-layer pairs in the package stack. The optimized routing and power delivery solutions improved SI and power performance, supporting both performance and cost objectives. This represents a significant advancement in processor package design, supporting the signal integrity and power delivery demands of next-generation enterprise and cloud workloads.

Acknowledgment

The authors wish to acknowledge Intel engineers across different disciplines for their contributions to the work and the collaborative development of the methodologies presented. Special recognition goes to signal integrity, power delivery, physical design team and assembly technology partners whose cross-functional collaboration made this integrated approach possible.

References

[1] S. S. Iyer, "Chips, Dies, Chiplets and Dielets and Heterogeneous Integration," *2022 6th IEEE Electron Devices Technology & Manufacturing Conference (EDTM)*, Oita, Japan, 2022

[2] Y. Han, H. Xu, M. Lu, H. Wang, J. Huang, Y. Wang, Y. Wang, F. Min, Q. Liu, M. Liu, Ninghui Sun, “The Big Chip: Challenge, model and architecture”, *Fundamental Research*, Volume 4, Issue 6, 2024, Pages 1431-1441.

[3] J. H. Lau, "Recent Advances and Trends in Advanced Packaging," in *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 12, no. 2, pp. 228-252, Feb. 2022.

[4] T. Li, J. Hou, J. Yan, R. Liu, and Z. Sun, “Chiplet heterogeneous integration technology–status and challenges,” *Electronics*, vol. 9, no. 4, p. 650, 2020.

[5] K. Banerjee, S. J. Souri, P. Kapur and K. C. Saraswat, "3-D ICs: a novel chip design for improving deep-submicrometer interconnect performance and systems-on-chip integration," in *Proceedings of the IEEE*, vol. 89, no. 5, pp. 602-633, May 2001

[6] P. E. Garrou, C. A. Bower, and P. Ramm, Eds., “Handbook of 3D Integration: Technology and Applications of 3D Integrated Circuits”, Weinheim, Germany: *Wiley-VCH*, 2008.

[7] R. Mahajan, R. Sankman, N. Patel, D. Kim, K. Aygun, Z. Qian, Y. Mekonnen, I. Salama, S. Sharan, D. Iyengar, and D. Mallik, "Embedded Multi-die Interconnect Bridge (EMIB) - A High Density, High Bandwidth Packaging Interconnect," *2016 IEEE 66th Electronic Components and Technology Conference (ECTC)*, Las Vegas, NV, USA, 2016, pp. 557-565

[8] Z. Zhang, Y. Mekonnnen, S. Mondal and K. Aygun, "Novel Coupled Via (CV) Feature for Far-end Crosstalk Reduction," *2024 IEEE International Symposium on Electromagnetic Compatibility, Signal & Power Integrity (EMC+SIPI)*, Phoenix, AZ, USA, 2024

[9] R. Enriquez, K. Xiao, B. Lee and M. Tlaxcalteco, "Differential symmetry principle for differential crosstalk cancellation," *2013 IEEE International Symposium on Electromagnetic Compatibility*, Denver, CO, USA, 2013, pp. 730-734.

[10] M. Swaminathan, J. Kim, I. Novak and J. P. Libous, "Power distribution networks for system-on-package: status and challenges," in *IEEE Transactions on Advanced Packaging*, vol. 27, no. 2, pp. 286-300, May 2004