

Packaging as a System-Scaling Pathway in the Post-Moore Era: Classification, Evaluation, and Application Matching

Ziyang Huang

Department of Electrical Engineering, University of Sydney, Sydney, NSW 2006, Australia

ABSTRACT

As the marginal benefit of transistor scaling decreases while lithography, yield, and design costs rise, advanced packaging has become a central route for sustaining semiconductor system performance beyond monolithic integration. This paper develops a compact but systematic framework for comparing four major packaging families: 2.5D interposer or bridge integration, 3D stacking with through-silicon vias (TSVs) or hybrid bonding, fan-out wafer/panel-level packaging, and chiplet-based heterogeneous integration. The study synthesizes academic literature, technology roadmaps, and representative industrial platforms, and evaluates each option using six criteria: interconnect density, bandwidth and latency, power efficiency, thermal reliability risk, manufacturing complexity, and cost/ecosystem maturity. The analysis shows that advanced packaging should not be understood as a single replacement for front-end scaling. Instead, it operates as a system-scaling layer in which process node, architecture, memory placement, substrate technology, thermal design, and interface standards are jointly optimized. 2.5D and 3D integration are most attractive for AI and high-performance computing, fan-out is well suited to thin mobile and edge devices, and mature-node chiplet or system-in-package solutions remain important for automotive and industrial electronics. The paper concludes with an application-matching model that supports early-stage technology selection in post-Moore chip design.

KEYWORDS

Advanced packaging; Heterogeneous integration; Chiplets; 2.5D integration; 3D integration; Fan-out wafer-level packaging

1 Introduction

For more than five decades, semiconductor progress has been driven by Moore-style transistor density scaling. Smaller devices enabled higher integration density and reduced cost per function; however, this scaling paradigm is now constrained by lithography complexity, rising mask and design costs, yield sensitivity, power density, and the physical reticle size limit^[1-2]. Although advanced lithography remains essential, front-end miniaturization alone is no longer sufficient to meet the bandwidth, energy efficiency, and time-to-market requirements of artificial intelligence, cloud computing, mobile systems, and automotive electronics.

In this context, advanced packaging has shifted from a back-end assembly step to a strategic system integration technology. Heterogeneous integration roadmaps emphasize the combination of logic, memory, analog, RF, photonic, and power devices within a single package as a key direction for future computing systems^[3].

Recent chiplet-based architectures further suggest that future systems will increasingly rely on 2.5D and 3D integration structures to connect multiple specialized dies with high bandwidth and low energy consumption^[4].

The research problem addressed in this paper is therefore not simply which packaging technology is most advanced, but rather which technology is most appropriate for a given system requirement. This paper proposes a concise classification and evaluation framework that maps packaging architectures to application scenarios. Unlike descriptive reviews, the proposed framework treats advanced packaging as a technology selection problem under multiple constraints, including performance, power, thermal behavior, reliability, cost, and ecosystem maturity.

2 Literature Review and Research Gap

Existing studies generally approach advanced packaging from three directions. The first direction focuses on device and process roadmaps. IRDS and HIR documents frame heterogeneous integration as a response to slowing device scaling and as an enabling technology for system-level density, bandwidth, and functional diversity^[2-3].

The second direction studies individual packaging technologies. Reviews of advanced packaging and fan-out packaging summarize process flows, redistribution layers, substrates, warpage, reliability, and manufacturability challenges^[4-6]. The third direction focuses on chiplet ecosystems, where die-to-die standards such as UCle aim to reduce vendor lock-in and enable reusable multi-die system design^[7].

Nonetheless, the majority of research literature classifies 2D, 3D, fan out and chiplet packaging systems in their own

unique separate technology narrative so that early designers and decision makers will find it hard to evaluate the various available package options in an efficient or systematic way.

For a given target application, the key challenge is to identify the packaging route that best balances interconnect performance, thermal feasibility, yield, cost, and ecosystem maturity.

Because “advanced packaging” includes fundamentally different types of physical architectures and business models, the gap between advanced packaging and legacy packaging is critically important. For example, a silicon-interposer GPU platform, a mobile fan-out package, and an automotive system-in-package solution will be governed by different limitations and will not solve the same types of optimization problems.

This paper addresses this gap by organizing the discussion around a multi-criteria technology-selection framework. The objective is not to rank packaging technologies on a single performance scale, but to identify their boundary conditions, trade-offs, and most defensible application scenarios.

3 Methodology

3.1 Comparative Evaluation Framework

The paper adopts a qualitative comparative methodology based on peer-reviewed literature, technology roadmaps, international standards, and representative industrial platforms. Each packaging family is evaluated using six criteria: (1) interconnect density, defined as the number and pitch of die-to-die connections; (2) bandwidth and latency, which determine communication efficiency between memory and compute units; (3) power efficiency, including energy consumption per bit for off-die communication; (4) thermal and reliability risks, including heat dissipation, warpage, mechanical stress, and testability; (5) manufacturing complexity and yield; and (6) cost and ecosystem maturity.

These criteria are selected because they capture the primary trade-offs in post-Moore system design. For example, 3D hybrid bonding maximizes vertical interconnect density but introduces significant thermal and testing challenges. Fan-out packaging enables thin form factors and favorable electrical performance for mobile systems but suffers from die-size and warpage constraints in large-area applications. Chiplet-based integration improves yield and design reuse, but requires mature interface standards, supply-chain coordination, and advanced design automation tools.

3.2 Technology Scope

Four packaging families are included in this study. 2.5D integration uses a silicon interposer, organic RDL interposer, or embedded bridge to place multiple dies side by side. 3D integration stacks dies vertically using through-silicon vias (TSVs), micro-bumps, or hybrid bonding. Fan-out packaging embeds dies in a reconstituted wafer or panel, enabling high-density redistribution layers without a conventional substrate in certain implementations. Chiplet-based integration is an architectural and ecosystem-level approach in which a system is partitioned into smaller functional dies and reassembled through standardized die-to-die interconnects.

4 Technology Classification and Comparative Analysis

4.1 2.5D Interposer and Bridge Integration

2.5D integration places multiple dies on an interposer or bridge-based structure, enabling high-bandwidth communication between logic and memory through dense lateral routing while avoiding full vertical stacking. Industry platforms such as TSMC CoWoS-S demonstrate silicon-interposer-based integration for high-performance computing and AI workloads, combining logic chiplets with HBM stacks on a large interposer [9]. Bridge-based approaches such as Intel EMIB reduce reliance on full silicon interposers by embedding local silicon bridges within the substrate, enabling a trade-off between routing density, package size, and cost^[10].

The major benefit of 2.5D integration is that it can provide extremely high memory bandwidth using mature thermal management and assembly processes. However, there are limitations to 2.5D integration in terms of interposer size, substrate complexity, and yield sensitivity which limit large scale multi die systems. Regardless of the existing issues presented by these limitations, 2.5D remains the premier option for AI accelerators and GPUs due to the memory bandwidth being the limiting factor in those systems.

4.2 3D Stacking and Hybrid Bonding

3D integration reduces interconnect length by stacking dies vertically, significantly increasing interconnect density and improving latency and energy efficiency. Early implementations relied on through-silicon vias (TSVs) and micro-

bumps, while recent high-end roadmaps increasingly emphasize hybrid bonding technologies. Intel Foveros Direct, for example, adopts direct copper-to-copper bonding to eliminate solder micro-bumps and achieve sub-10 μm interconnect pitch, enabling higher density and improved form-factor efficiency^[10].

The primary advantage of 3D integration is vertical density, which enables high-performance logic-to-memory or logic-to-logic communication with minimal interconnect overhead. However, the approach introduces significant thermal concentration, testing difficulty, and yield sensitivity, since stacking increases power density and requires known-good-die assumptions. As a result, 3D integration is most effective when performance gains from proximity outweigh thermal and manufacturing risks.

4.3 Fan-Out Wafer/Panel-Level Packaging

Fan-out packaging uses redistribution layers (RDLs) to extend I/O beyond die boundaries and achieve thin, substrate-free package profiles in certain implementations. TSMC InFO is a representative platform that enables high-density routing using advanced RDL and through-mold interconnect structures for mobile and high-performance applications^[9]. Peer-reviewed studies further highlight fan-out as a key heterogeneous integration approach for compact system design due to its ability to integrate multiple dies and passive components within a single reconstituted package^[6-7].

Fan-out packaging has three main advantages (form factor, electrical performance and manufacturability); this allows a good balance of these requirements and suits mobile devices as well as edge applications. With large dies, there are issues associated with scaling due to the problems of warping, shifting of dies and routing due to restrictions on RDL, limiting the use of fan-out packaging for high power applications in large die areas.

4.4 Chiplet-Based Heterogeneous Integration

Chiplets are simply a type of packaging that incorporate all functional circuits into one or more chips through an interconnect medium. Each chip has its own physical architecture, but they use a common electrical interface to connect together. Heterogeneous architecture allows separate processing technologies (i. e., multiple process node die technologies) to be combined, resulting in improved manufacturing yield, ability to share or reuse designs across multiple chips, and maximizing use of design methodologies (i.e., provide for different optimization of logic, memory, analog, I/O, and power).

However, chiplet systems require standardized die-to-die interfaces, advanced EDA tool support, power delivery coordination, and robust supply-chain integration. The emergence of standards such as UCle addresses part of this challenge by defining physical, protocol, and software layers for interoperable multi-vendor chiplet ecosystems^[7]. Therefore, chiplet adoption depends not only on packaging density but also on ecosystem maturity and standardization readiness.

Table 1 Application-oriented comparison of major advanced packaging routes

Route	Technical strengths	Key limitations	Representative applications
2.5D interposer/bridge integration	High-bandwidth connectivity to HBM; mature integration for large systems; effective lateral heat dissipation	Large package footprint; interposer/bridge cost; assembly yield sensitivity	AI accelerators; GPUs; HPC systems; supercomputing
3D stacking / hybrid bonding	Shortest interconnect paths; highest vertical density; strong integration form-factor efficiency	Thermal concentration; difficult testing and repair; high bonding precision requirements	Logic-memory stacks; cache stacking; high-performance AI/HPC modules
Fan-out WLP/PLP	Thin package profile; high-density redistribution layers (RDL); favorable electrical performance; scalable cost structure	Warpage; die shift; RDL routing limitations for large die sizes	Mobile SoCs; RF modules; edge devices; compact systems
Chiplet-based heterogeneous integration	Yield improvement via die partitioning; design reuse; process-node optimization; system-level flexibility	Dependence on interface standards; EDA complexity; testing and supply-chain coordination	Modular SoCs; automotive platforms; data-center processors

5 Application Matching and Strategic Implications

The comparison indicates that technology selection should be driven by workload characteristics rather than packaging type. AI and HPC systems are typically bandwidth-limited and power-dense; therefore, 2.5D integration with HBM remains a practical baseline solution. In contrast, 3D stacking becomes advantageous when latency reduction and form-factor constraints justify higher thermal and testing complexity.

Mobile and edge systems are more sensitive to package thickness, board area, and cost, making fan-out technologies more suitable. Automotive and industrial applications require long qualification cycles, high reliability, and stable supply chains; therefore, mature-node chiplet-based and system-in-package solutions are often preferred over the most aggressive density-oriented architectures.

Advanced packaging thus creates a new design space between front-end process scaling and board-level system integration. System-level performance depends on co-optimization of packaging, substrate, memory, cooling, test strategy, and EDA tools. A high-performance package may fail commercially if thermal design, yield learning, or supply-chain capacity is insufficient, while a less aggressive solution may achieve better overall system efficiency by reducing risk and accelerating time-to-market.

For emerging semiconductor ecosystems, the most actionable opportunity lies in co-optimization across the value chain. Rather than competing solely at advanced lithography nodes, competitiveness can be built through innovation in substrates, redistribution layer (RDL) processing, bonding technologies, thermal materials, known-good-die testing, package-aware EDA, and chiplet interface validation. These capabilities are critical enablers for reliable and economically viable post-Moore system scaling.

6 Conclusion

The central idea of this paper is to define Advanced Packaging as a System / Scaling Pathway in the "Post-Moore" Era as opposed to simply describing isolated Packaging Technologies. We conclude that no single Packaging Technology can be characterized as "Best".

Types of Integration Options Relative To Application Characteristics:

2.5D Integration: Highly effective for memory bandwidth intensive applications.

3D Stacking: Provides the highest levels of interconnect density and proximity, but requires careful design of thermal and test processes.

Fan Out: Compact and manufacturable for mobile/edge devices.

Chiplet Based Integration: Facilitates modularity, reuse of design and specialization of processing nodes, thereby redefining how systems are designed.

The evaluation framework being proposed, has been designed to correlate all of the available packaging options against six specific parameters, which include: Interconnect density, Bandwidth/Latency, Power consumption, Thermal/Reliability risk, Manufacturing complexity and Cost/ecosystem maturity. By establishing an association between architectural decisions and system-level constraints, the proposed framework is able to eliminate the potential for oversimplified assessments of "advanced" implying globally better performance.

The research offers a structured approach towards selecting packaging for application-driven systems following Moore's Law that considers at the design phase an integrated evaluation between disparate integration methods as opposed to evaluating each separately based on technical metrics.

The research presented here is mostly a literature synthesis and data from public sources. Future efforts could enhance this framework through quantitative benchmarking of bandwidth density, energy per bit, thermal resistive properties, yields, and costs retrieved from actual products. In addition to what was done in this research, improved modelling of chiplet placement within the die footprint (i.e., floorplanning), chiplet-aware electronic design automation (EDA) tools and thermal-mechanical reliability would increase the capability to predict application-matching decisions.

References

- [1] G. E. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, no. 8, pp. 114–117, 1965.
- [2] IEEE International Roadmap for Devices and Systems (IRDS), "2023 Edition," IEEE, 2023. [Online]. Available: <https://irds.ieee.org/>.
- [3] IEEE Electronics Packaging Society, "Heterogeneous Integration Roadmap," 2023. [Online]. Available: <https://eps.ieee.org/technology/heterogeneous-integration-roadmap/>.
- [4] D. Sharma and R. V. Mahajan, "Advanced packaging of chiplets for future computing needs," *Nature Electronics*, vol. 7, pp. 425–427, 2024. doi: 10.1038/s41928-024-01175-3.
- [5] 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.
- [6] J. H. Lau, "Recent advances and trends in fan-out wafer/panel-level packaging," *Journal of Electronic Packaging*, vol. 141, no. 4, 040801, 2019.
- [7] T. Braun et al., "Fan-out wafer and panel level packaging platform for heterogeneous integration," *Micromachines*, vol. 10, no. 5, 342, 2019.
- [8] Universal Chiplet Interconnect Express Consortium, "UCIe 2.0 Specification," 2024. [Online]. Available: <https://www.uciexpress.org/>.
- [9] Taiwan Semiconductor Manufacturing Company (TSMC), "3DFabric Technologies: CoWoS and InFO," 2025. [Online]. Available: <https://3dfabric.tsmc.com/>.
- [10] Intel Corporation, "Foveros Direct 3D Technology Brief," 2025. [Online]. Available: <https://www.intel.com/>.