Transactions on Computer Science and Intelligent Systems Research
ISSN: 2960-1800, eISSN: 2960-2238 | Volume 14
TPCEE 2026

# High-Speed Die-to-Die Interconnects and Signal Integrity in Chiplet-Based Advanced Packaging

Jiaran Yun [1, *]

[1] College of Physical Science and Engineering, Tongji University, Shanghai, China

* Corresponding Author Email: 2351060@tongji.edu.cn

**Abstract.** With the increasing cost and complexity of advanced semiconductor manufacturing, chiplet integration has become an important approach for high-performance computing and AI systems. Instead of integrating all functions on a single large die, chiplet-based systems divide different functional modules into smaller dies and connect them through package-level interconnects. In this process, advanced packaging is no longer only used for protection and external connection, but also becomes a key platform for high-density interconnection, power delivery, heat dissipation and heterogeneous integration. This paper reviews the basic concepts of chiplet architectures and advanced packaging, with emphasis on 2D, 2.5D and 3D integration schemes. Typical high-speed interconnect structures, including silicon interposers, silicon bridges, redistribution layers (RDLs), fan-out interconnects, through-silicon vias (TSVs), micro-bumps, and hybrid bonding, are discussed. Key metrics such as bandwidth density, energy efficiency, latency, scattering parameters (S-parameters), voltage transfer function (VTF), eye diagrams, and bit error rate (BER) are also introduced. Finally, signal integrity issues, optimization strategies and future trends are analyzed, including signal integrity and power integrity (SI/PI) co-design, thermal-aware design, UCIe standardization and system-level co-design.

**Keywords:** Chiplet; Advanced packaging; Die-to-die interconnect; Signal integrity; 2.5D/3D integration.

## 1. Introduction

As transistor scaling encounters diminishing returns and leading-edge process technologies become increasingly costly, monolithic system-on-chip (SoC) designs are progressively burdened by rising fabrication costs, declining manufacturing yields, and mounting design complexity.

This problem is especially pronounced in high-performance computing and AI chips, where large dies require high integration density and extensive data traffic. When all functional modules are placed on one die, manufacturing defects, power density and thermal management become serious limitations.

Chiplet integration provides another route for system design. In a chiplet-based architecture, processors, memory controllers, input/output (I/O) modules and other functional units can be divided into separate dies and then integrated through package-level interconnects. This method allows different chiplets to use different process nodes according to their functions. For example, compute dies may use advanced logic processes, while I/O or analog dies can use more mature technologies. In this way, chiplet integration can improve design flexibility and reduce manufacturing pressure.

However, the transition to a chiplet architecture changes the dominant system-level constraint, shifting the primary performance bottleneck from computation within a single die to data exchange between separate chiplets.

After partitioning, the communication between chiplets becomes a central problem. The overall system performance hinges equally on the computational capabilities of the individual chiplets and on the bandwidth, latency, energy efficiency, signal integrity and power integrity (SI/PI), and thermal characteristics of the package-level interconnects. Therefore, advanced packaging should be regarded

Content from this work may be used under the terms of CC BY-NC 4.0 licence (https://creativecommons.org/licenses/by-nc/4.0/).
Published by Warwick Evans Publishing.

as a first-order architectural determinant of chiplet-based systems rather than merely as a conventional back-end integration step.

This paper addresses chiplet packaging systems and their fundamental interconnect structures, followed by system- and channel-level performance measurements, significant SI/PI techniques, and emerging approaches in hybrid bonding, thermal-aware design, and interface standardization.

## 2. Chiplet Packaging Architectures and Die-to-Die Interconnect Performance

### 2.1. Chiplet Integration and Advanced Packaging

Fabricating a large monolithic die becomes increasingly expensive at advanced process nodes, where a defect may render the entire device unusable. A chiplet design divides the system into several smaller dies instead. Processing cores can be manufactured using an advanced logic process, while I/O and analog functions can remain on mature nodes that are more suitable and economical for those circuits. This approach allows different functions to use different technologies and may reduce the cost and yield risk associated with a large system-on-chip (SoC) [1, 2].

This partition changes the communication path within the system. Transfers that would previously have remained on the same die must cross package traces, vias, bumps, and other transitions before reaching another chiplet. These links need to provide the required bandwidth and latency without consuming excessive energy. Package design also affects power delivery and heat removal, so its influence extends beyond the electrical channel.

The physical path between two dies contains resistance, capacitance, and inductance contributed by conductors, vias, bumps, pads, and discontinuities in the reference planes. Their effects become more pronounced as the operating frequency and interconnect density increase. Poorly controlled parasitics attenuate the received signal and create reflection, crosstalk, and power-supply noise, eventually reducing timing margin and link reliability [2].

The selected packaging structure determines where fine-pitch routing is placed and how far signals must travel. A silicon interposer distributes dense wiring over a relatively large region, while a local bridge concentrates such wiring between neighboring dies. Fan-out redistribution layers (RDLs) provide another lateral routing option, and vertical bonding shortens the connection by stacking dies. Each structure has its own implications for bandwidth, package area, manufacturing cost, thermal management, and assembly reliability. Packaging decisions must consequently be made together with the chiplet architecture rather than treated only as a back-end manufacturing step [1, 2].

Advanced packaging can be classified by die placement and interconnect structure. Figure 1 compares conventional substrate routing with silicon bridges, silicon interposers, and fan-out RDLs.

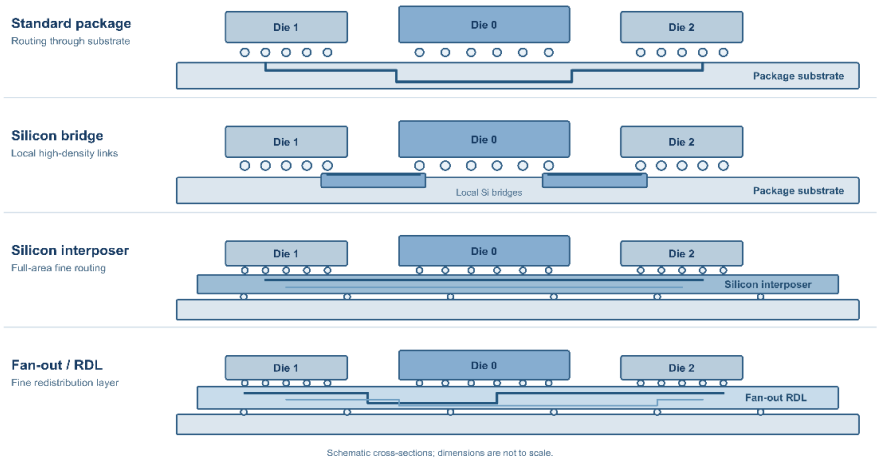


**Figure 1.** Representative package-level interconnect structures for chiplet integration (Picture credit: Original)

As shown in figure 1, placing fine routing closer to the dies shortens the signal path and increases interconnect density. A bridge covers a local region, while an interposer or fan-out RDL can provide routing across a larger area.

In 2D packaging, dies interact via an organic substrate or a fan-out RDL. The method is mature and reasonably priced, but longer pathways restrict bandwidth density. A 2.5D package includes an interposer or bridge for finer lateral routing. This supports high-bandwidth connections, but the cost and design complexity rise. Dies in 3D packaging are stacked and connected via through-silicon vias (TSVs), micro-bumps, or hybrid bonding. The routes are shorter, but power delivery, cooling, testing, yield, and dependability are more challenging [1-3]. Table 1 outlines the key distinctions between the three packing types.

**Table 1.** Comparison of 2D, 2.5D, and 3D Packaging Technologies for Chiplet Integration

| Type | Basic Structure | Typical Interconnection Form | Key Features | Key Challenges |
|---|---|---|---|---|
| 2D Packaging | Die placed directly on a packaging substrate or fan-out structure | organic substrate, fan-out RDL | Mature process, lower cost | Longer interconnection paths, limited bandwidth density |
| 2.5D Packaging | Die interconnected laterally via an interposer or bridge structure | silicon interposer, silicon bridge, RDL interposer | High interconnect density, suitable for high-bandwidth die-to-die communication | Significant issues regarding cost, area, routing and signal integrity |
| 3D | Die stacked vertically | TSV, micro-bump, hybrid bonding | Short paths, high bandwidth density, low latency | Complex thermal management, yield, testing and reliability |

The three forms meet different system needs. Their path lengths, parasitic effects, and coupling conditions also differ. Signal-integrity analysis must therefore follow the actual package structure. The main interconnect structures used in these packages are discussed next.

## 2.2. High-Speed Interconnect Structures

A silicon interposer lies between the chiplets and package substrate. Fine metal traces provide dense lateral routing, micro-bumps connect the dies, and TSVs connect the interposer to the substrate. This structure is widely used for high-bandwidth links between logic dies and high-bandwidth memory (HBM). Hou et al. demonstrated this form of logic-memory integration using second-generation CoWoS technology [4].

Silicon interposers also introduce design challenges. A large interposer increases package area, cost, and yield pressure. TSVs, bumps, pads, and metal traces add parasitic resistance, capacitance, and inductance. Routing, reference planes, via transitions, and power delivery must therefore be designed together to control loss, reflection, and crosstalk [2, 5].

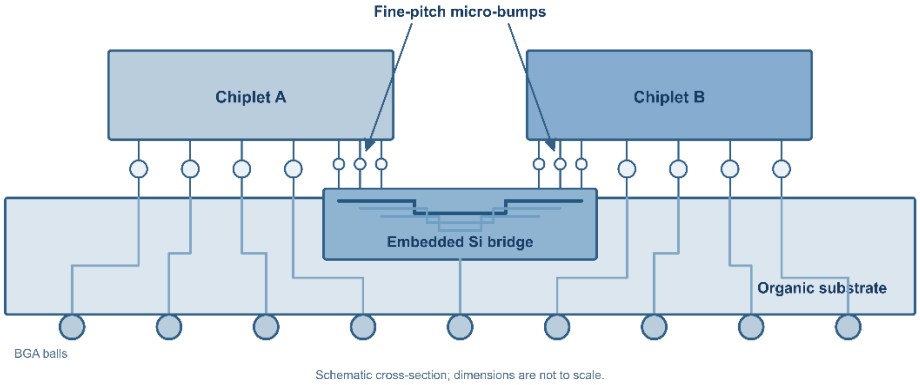


**Figure 2.** Schematic cross-section of an embedded silicon bridge for 2.5D chiplet integration (Picture credit: Original)

A silicon bridge provides dense routing only between adjacent chiplets. It is usually embedded near the surface of an organic substrate and connected through fine-pitch micro-bumps. Embedded Multi-die Interconnect Bridge (EMIB) is a representative implementation. This structure avoids the area of a full silicon interposer while retaining local high-density links [6]. Figure 2 shows its basic structure. As shown in figure 2, the bridge carries short die-to-die paths, while other signal and power connections pass through the organic substrate. It therefore combines local high-density routing with the lower area and cost of an organic package.

Fan et al. evaluated silicon bridges designed for the Universal Chiplet Interconnect Express (UCIe), using bump pitches of 45 and 55 μm. At 32 gigabits per second (Gbps), the scattering-parameter (S-parameter), voltage transfer function (VTF), and eye-diagram results satisfied the electrical limits adopted in the study [7]. Nevertheless, the routing density and transition geometry required for compact interconnection can increase electromagnetic coupling and introduce impedance discontinuities. Crosstalk suppression therefore requires coordinated optimization of return-current paths, ground shielding, and bump arrangement.

RDL and fan-out structures provide another form of lateral interconnection. An RDL moves chiplet I/O pads to positions suitable for package routing. In fan-out and some 2.5D packages, it can also connect adjacent dies. RDL-based structures offer flexible routing and may reduce package thickness and cost [1,2]. Their high-frequency performance depends on line geometry, dielectric loss, vias, reference planes, and substrate transitions. These factors must be controlled to limit loss, reflection, and crosstalk [2,5].

Unlike these lateral structures, 3D packages connect vertically stacked dies through TSVs, micro-bumps, or hybrid bonding. TSVs pass through silicon, while micro-bumps join the stacked dies. Hybrid bonding directly joins metal and dielectric surfaces, allowing a smaller pitch. Short vertical links can increase bandwidth density and reduce latency and energy use. However, stacking complicates cooling, power delivery, testing, bonding yield, and thermomechanical reliability. Electrical, thermal, mechanical, and manufacturing factors must therefore be considered together [1,3].

Each structure creates a different physical channel. A common set of system-level and electrical metrics is therefore needed to compare their performance.

### 2.3. Performance Metrics for Die-to-Die Links

The performance of a die-to-die link cannot be represented by a single parameter. From a system perspective, the line rate describes the amount of data carried by one lane, whereas the aggregate bandwidth depends on the line rate, lane count, and protocol overhead. Bandwidth density further relates the available bandwidth to the package edge length or area occupied by the interface. A wide interface consisting of many short parallel lanes can therefore provide substantial throughput without pushing each lane to an extremely high data rate. Such an arrangement, however, places tighter requirements on lane-to-lane skew, clock distribution, and package routing [2,3].

Power consumption and latency impose another set of constraints. Interconnect energy is commonly reported in picojoules per bit (pJ/bit), while latency accounts for the time required to transfer data from the transmitting die to the receiving die. Both are particularly important when processors and memory chiplets exchange data frequently. Short routing distances and compact I/O circuits may reduce energy consumption and delay, but the result still depends on the electrical properties of the package channel [2,3].

Frequency-domain results are needed to determine how the physical channel affects the transmitted signal. The forward transmission coefficient, S21, indicates the amount of signal that reaches the receiver and is used to obtain insertion loss. S11 characterizes the signal reflected at the input and is associated with return loss. Attenuation generally increases with conductor and dielectric losses,

while bumps, vias, pads, and package transitions may introduce impedance discontinuities and additional reflection [5].

These frequency-domain quantities do not, by themselves, show the voltage and timing margins available to the receiver. The VTF can be used to examine signal transfer and coupling between nearby paths. An eye diagram provides a time-domain view of the received waveform: its height indicates the available voltage margin, and its width indicates the timing margin. The bit error rate (BER) gives the probability of incorrect bit reception and provides an end-to-end measure of link reliability [5, 7].

UCIe defines a common physical layer, die-to-die adapter, and protocol framework for standard and advanced packages [3, 8]. This common framework does not make the electrical behavior of every implementation identical. Bump pitch, channel length, routing geometry, package loss, power delivery, and temperature can still affect the measured result. Consequently, compliance testing should be accompanied by package-specific evaluation using the metrics summarized in table 2.

**Table 2.** Key Performance Metrics for High-Speed Die-to-Die Interconnects

| Metric | Meaning | Relevance to chiplet interconnects |
|---|---|---|
| Data rate | Bit rate supported by one lane | Determines single-lane speed |
| Aggregate bandwidth | Combined usable data capacity of all active lanes | Represents the overall communication capacity of the interface |
| Bandwidth density | Available bandwidth normalized to package edge length or area | Allows interfaces with different physical sizes to be compared |
| Energy efficiency | Energy required to transfer one bit, usually expressed in pJ/bit | Indicates the power cost of die-to-die communication |
| Latency | Time required for data to travel from the transmitting die to the receiving die | Affects communication between closely coupled chiplets |
| S21 / insertion loss | Forward transmission response of the channel | Quantifies signal attenuation before the receiver |
| S11 / return loss | Input reflection associated with impedance mismatch | Reveals discontinuities along the signal path |
| VTF | Voltage transferred through the intended and coupled paths | Assesses signal delivery and crosstalk |
| Eye diagram | Superposition of received waveforms over multiple bit periods | Shows the available voltage and timing margins |
| BER | Probability that a transmitted bit is received incorrectly | Measures the reliability of the complete link |

These measurements address different aspects of the same link. Aggregate bandwidth, energy per bit, and latency describe its usefulness at the system level, whereas S-parameters, VTF, eye diagrams, and BER reveal the condition of the electrical channel. Package structures should therefore be compared using both groups of results rather than by quoting the maximum data rate alone.

## 3. Signal Integrity Issues and Optimization Strategies

### 3.1. Loss, Reflection and Impedance Control

Insertion loss mainly comes from conductor loss, dielectric loss, vias, bumps and package transitions. Reflection is caused by impedance discontinuity when signals pass through pads, vias, RDL transitions, silicon bridge interfaces or TSVs. In high-speed die-to-die links, these effects reduce the received signal amplitude and make the eye opening smaller [5].

Measurements of a UCIe-oriented silicon bridge at 32 Gbps show that channel performance depends on both bridge routing and die-to-package transitions [7]. The resulting loss and reflection characteristics therefore remain specific to the implemented bump pitch and transition geometry.

To reduce loss and reflection, package design usually needs shorter routing paths, smoother transitions, controlled impedance, optimized bump/via structures and continuous reference planes.

### 3.2. Crosstalk and Routing Optimization

Crosstalk becomes more pronounced as routing density increases. In silicon interposers, silicon bridges, and fan-out RDLs, capacitive and inductive coupling between closely spaced traces introduces noise into adjacent channels, reducing voltage margin and increasing timing uncertainty.

UCIe compliance alone does not ensure low crosstalk because coupling remains sensitive to trace spacing, reference conductors, and the RDL stack-up. At 32 GT/s, Wang et al. evaluated ground-shielded routing arrangements in CoWoS-R to suppress VTF crosstalk [9]. The results support package-specific optimization of signal-to-ground routing rather than transferring a single arrangement across different interconnect media.

Accordingly, crosstalk mitigation should combine adequate trace spacing, ground-signal-ground (G-S-G) routing, continuous reference planes, and short return-current paths.

### 3.3. Eye Closure, Jitter and Link Verification

Eye closure is caused by the combined effects of loss, reflection, crosstalk, jitter and skew. Eye height shows voltage margin, while eye width shows timing margin. For parallel chiplet links, skew between lanes is also important because many lanes need to be sampled together.

Fan et al. demonstrated an integrated verification flow combining S-parameter, VTF, and eye-diagram analyses at 32 Gbps, representing a comprehensive method for link-level assessment [7]. Nevertheless, its validation remains limited to a specific test vehicle, and larger multi-chiplet systems require additional system-level modeling.

For link optimization, equalization, clocking strategy, lane matching and skew control should be considered together with physical routing. In short-reach chiplet interconnects, the goal is not only higher single-lane speed, but also stable timing margin and low energy per bit.

### 3.4. SI/PI Coupling and Co-Design

Signal integrity and power integrity are closely coupled in chiplet systems. Power noise, ground bounce, resistive voltage drop and power delivery network (PDN) impedance can affect I/O drivers and receivers. Therefore, a channel may pass passive S-parameter checks but still suffer from timing or voltage noise after power delivery effects are included.

Wang et al. implemented a UCIe I/O link operating at up to 32 GT/s on CoWoS-R and reported that a bottom-side integrated passive-device (IPD) capacitor reduced PDN impedance by 23% at 100 MHz [9]. This result demonstrates that package-level decoupling directly affects the electrical environment of high-speed die-to-die links and should therefore be considered together with channel design.

UCIe provides a common interface framework, but it cannot remove package-specific SI/PI problems. Channel loss, crosstalk, bump pitch, power delivery and thermal behavior still depend on the actual 2D, 2.5D or 3D package structure [8].

## 4. Trends and Challenges

### 4.1. Higher Interconnect Density and Hybrid Bonding

The scaling of chiplet interconnects is increasingly determined by connection pitch. Micro-bump-based 3D integration shortens vertical signal paths, but further pitch reduction is limited by bump geometry, solder reliability, and assembly yield. Hybrid bonding addresses some of these limitations by enabling finer vertical connections with lower parasitic effects than conventional micro-bump interconnects [10, 11].

Lu reported hybrid-bonding pitches below 10 μm, with wafer-to-wafer demonstrations reaching below 0.5 μm [10]. Further scaling below 0.4 μm remains under investigation. These dimensions illustrate the density advantage of hybrid bonding, but achievable pitch alone does not determine whether the process is suitable for manufacturing. Overlay accuracy, pad-profile variation, surface preparation, and thin-die handling remain important constraints [11].

The practical limit is therefore not simply the smallest attainable pitch, but the pitch that can be produced with acceptable yield and reliability at manufacturing scale [11].

### 4.2. Thermal and Reliability-Aware Design

Thermal constraints become more severe as the number and power density of chiplets increase. In a 2.5D package, several high-power dies placed on the same interposer or substrate can create localized hot spots. A 3D stack introduces additional vertical thermal coupling and makes heat removal from the inner layers more difficult. These thermal effects also influence electrical behavior. Elevated temperature can alter interconnect resistance, reduce timing margin, and accelerate long-term reliability degradation.

Compared with the thermally aware TAP-2.5D baseline, analytical thermal-aware chiplet placement framework (ATPlace2.5D) reduced the maximum temperature by 5% and total wirelength by 42%, while achieving a 23-fold speedup. The results show that placement optimization can improve thermal and interconnect objectives simultaneously. However, the framework does not model detailed routing or power/ground connectivity. Its effectiveness at the package level therefore requires further validation [12].

### 4.3. Standardization and Open Chiplet Ecosystem

Standardized interfaces are necessary if chiplets from different designs or vendors are to operate within the same package. UCIe defines common requirements for the physical layer, die-to-die adapter, and protocol support, providing a technical basis for chiplet interoperability [8].

Compliance with the interface specification does not make different package implementations electrically equivalent. Channel loss, crosstalk, power delivery, and thermal behavior still depend on package geometry, bump pitch, routing structure, and material properties. A UCIe-compliant system must therefore undergo package-specific electrical and thermal verification [8].

### 4.4. Toward System-Level Co-Design

Package design increasingly involves coupled constraints rather than independent optimization tasks. Interconnect density, energy per bit, S-parameters, VTF, eye opening, crosstalk, power delivery, and temperature influence one another. Denser routing, for example, can increase crosstalk and local heat density. Conversely, a placement that improves thermal behavior may require longer signal paths and introduce additional latency.

These interactions change the nature of the design problem. The objective is not to maximize a single metric, but to evaluate trade-offs among performance,

SI/PI, thermal behavior, reliability, manufacturability, and cost. ATPlace2.5D illustrates the value of coupling placement with thermal and wirelength objectives, but it leaves detailed routing and power-distribution modeling outside the optimization process [12]. Electronic design automation (EDA) workflows that account for all of these effects within a unified design process remain underdeveloped.

## 5. Conclusion

This review shows that packaging is an integral part of chiplet-system architecture because its physical structure defines the conditions for die-to-die communication. The main integration schemes involve different trade-offs: 2D packaging favors cost and process maturity, 2.5D integration enables denser lateral interconnects with greater routing and signal-integrity complexity, and 3D stacking reduces interconnect length while increasing thermal, testing, and reliability challenges. Package selection is therefore application-dependent and must balance performance requirements with cost and manufacturing feasibility.

Reliable die-to-die communication also requires system-level and channel-level evaluation to be considered together. Bandwidth, latency, and energy efficiency characterize communication capability, whereas S-parameters, VTF, eye diagrams, and BER assess the quality of the physical channel. Because channel loss, impedance discontinuities, crosstalk, and PDN noise are coupled

through the package, routing, return-current paths, shielding, and power delivery cannot be optimized independently. UCIe provides a common interface framework, but package-specific SI/PI verification remains necessary.

Future progress will depend on both finer-pitch hybrid bonding and more integrated design methods. Hybrid bonding can increase interconnect density, but its practical use requires tighter control of alignment, surface preparation, yield, and reliability. At the same time, EDA workflows must coordinate architecture partitioning, circuit and link design, package routing, power delivery, thermal analysis, and manufacturing constraints. The long-term development of chiplet systems will therefore depend on the ability to manage these coupled electrical, thermal, and manufacturing trade-offs at the system level.

## References


[1] Lau J H. Recent Advances and Trends in Chiplet Design and Heterogeneous Integration Packaging. Journal of Electronic Packaging, 2024, 146(1): 010801.

[2] Li S, Lin M S, Chen W C, et al. High-Bandwidth Chiplet Interconnects for Advanced Packaging Technologies in AI/ML Applications: Challenges and Solutions. IEEE Open Journal of the Solid-State Circuits Society, 2024, 4: 351-364.

[3] Das Sharma D, Pasdast G, Tiagaraj S, et al. High-performance, power-efficient three-dimensional system-in-package designs with universal chiplet interconnect express. Nature Electronics, 2024, 7(3): 244-254.

[4] Hou S Y, Chen W C, Hu C, et al. Wafer-Level Integration of an Advanced Logic-Memory System Through the Second-Generation CoWoS Technology. IEEE Transactions on Electron Devices, 2017, 64(10): 4071-4077.

[5] Fan J, Ye X, Kim J, et al. Signal Integrity Design for High-Speed Digital Circuits: Progress and Directions. IEEE Transactions on Electromagnetic Compatibility, 2010, 52(2): 392-400.

[6] Mahajan R, Qian Z G, Viswanath R S, et al. Embedded Multidie Interconnect Bridge—A Localized, High-Density Multichip Packaging Interconnect. IEEE Transactions on Components, Packaging and Manufacturing Technology, 2019, 9(10): 1952–1962.

[7] Fan Y, Zhou Y, Wang Q, et al. Design and verification of silicon bridge in 2.5D advanced package based on universal chiplet interconnect express (UCIe). Microelectronics Reliability, 2025, 168: 115710.

[8] Das Sharma D, Pasdast G, Qian Z, et al. Universal Chiplet Interconnect Express (UCIe): An Open Industry Standard for Innovations with Chiplets at Package Level. IEEE Transactions on Components, Packaging and Manufacturing Technology, 2022, 12(9): 1423-1431.

[9] Wang C T, Shang S A, Hsiao Y M, et al. Signal and Power Integrity Performance of CoWoS-R in Chiplet Integration Applications. 2023 IEEE 25th Electronics Packaging Technology Conference (EPTC). Singapore: IEEE, 2023: 296-301.

[10] Lu M C. Hybrid Bonding for Ultra-High-Density Interconnect. Journal of Electronic Packaging, 2024, 146(3): 030802.

[11] Lu M C. Manufacturing Challenges of Hybrid Bonding for Chiplets Heterogenerous Integration. Journal of Electronic Packaging, 2026, 148(1): 010801.

[12] Wang Q, Li X, Jia T, et al. ATPlace2.5D: Analytical Thermal-Aware Chiplet Placement Framework for Large-Scale 2.5D-IC. Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design. New York, NY, USA: Association for Computing Machinery, 2024: 1-9.