



Design Verification of Multi-Chiplet AI Accelerators: Challenges and Solutions

VenkataVaraha Chakravarthy Kanumetta*

Independent Researcher., USA

* Corresponding Author Email: vvchakravarthykanumetta@gmail.com - ORCID: 0000-0002-5247-7110

Article Info:

DOI: 10.22399/ijcesen.4357

Received : 27 August 2025

Revised : 21 November 2025

Accepted : 24 November 2025

Keywords

Multi-Chiplet Verification,
AI Accelerators,
Heterogeneous Integration,
Distributed Simulation,
Chiplet Interconnects

Abstract:

The verification of multi-chiplet AI accelerators represents one of the most formidable challenges in contemporary semiconductor design, driven by the industry's transition from monolithic architectures to heterogeneous integration strategies that combine multiple specialized dies within a single package. This paradigm shift addresses fundamental limitations in Moore's Law scaling by enabling the integration of chiplets fabricated on different process nodes, optimized for specific functions such as computation, memory management, and input-output operations. However, this architectural evolution introduces unprecedented verification complexity stemming from the integration of billions of transistors across multiple dies, each containing intricate intellectual property blocks and subsystems that must communicate through sophisticated die-to-die protocols. The verification process must address not only functional correctness but also critical non-functional properties, including thermal management across stacked dies, signal integrity through vertical interconnects, power delivery network integrity, and timing predictability across asynchronous clock domains. Traditional monolithic verification approaches prove inadequate for these distributed systems, necessitating advanced methodologies that combine hardware-accelerated emulation platforms, distributed simulation techniques employing virtual channels, formal verification methods for protocol compliance, and intelligent clock gating strategies. Despite these sophisticated approaches, significant challenges persist, including substantial infrastructure costs, extended debugging cycles for distributed simulations that can span days or weeks, a combinatorial explosion of the verification space, and the complexity of achieving adequate coverage across all chiplets and operating modes. The semiconductor industry is responding with innovative solutions, including hybrid verification frameworks that strategically combine multiple methodologies, standardized chiplet interfaces that enable reusable verification components, advanced debugging tools with unified cross-chiplet waveform databases, and emerging artificial intelligence-driven techniques for intelligent testbench generation, coverage closure prediction, and verification planning optimization. These developments, while promising enhanced productivity and effectiveness, underscore that successful multi-chiplet verification requires not only technological innovation but also substantial capital investment, specialized expertise, and close collaboration between design and verification teams to deliver high-quality, bug-free AI accelerators that meet stringent performance and reliability requirements.

1. Introduction

The fast pace of artificial intelligence and machine learning implementations has spurred semiconductor companies to design ever-more intricate AI accelerators that break the limits of computer performance. Contemporary AI accelerator architecture has moved away from monolithic designs and has instead adopted multi-chiplet configurations, where several silicon dies

are interfaced to each other using advanced die-to-die (D2D) protocols to deliver unprecedented processing power. The advent of chiplet technology is a revolutionary method in semiconductor design that provides a paradigm shift from conventional monolithic integrated circuits to heterogeneous integration techniques that support the assembly of various smaller dies manufactured through different process nodes and technologies in a common package [2]. This design innovation responds to the

intrinsic problems of Moore's Law scaling, wherein die sizes rise with exponentially decreasing manufacturing yields, causing monolithic solutions to become progressively economically unfeasible for sophisticated systems. Disaggregating huge monolithic designs into functionally specialized small chiplets allows manufacturers to produce higher yields, lower development expenditures, and facilitate mix-and-match architectures integrating compute, memory, and I/O dies optimized for their individual purposes [2].

This paradigm shift in architecture, while providing great benefits of scalability, yield, and customization, presents overwhelming design verification challenges. Design verification grows exponentially more difficult as these systems combine billions of transistors on several chiplets, each having complex intellectual property (IP) blocks and subsystems. Multi-die packages have distinct thermal properties that are very different from those of single-die packages, and thermal interactions between neighboring dies lead to hot spots and thermal coupling effects that need to be modeled and validated correctly during the design stage [1]. The thermal management problems in multi-die packaging are most critical within high-power AI accelerators, where power densities can differ widely between different chiplets, and the thermal resistance between the package substrate and dies forms intricate heat dissipation routes that affect both reliability and performance [1]. These thermal implications have a direct impact on verification approaches, given that testbenches need to consider temperature-dependent timing skews, noise coupling between chiplets caused by power, and thermal throttling characteristics that can change system behavior under stress.

Since tape-out expenses for designs so sophisticated can easily amount to millions of dollars, ensuring thorough verification before fabrication is not just prudent—it is financially essential. The chiplet design flow, though providing modularity and reusability advantages, is demanding in verifying inter-chiplet interfaces, coherency protocols, and system-level integration cases that were secondary issues in monolithic design [2]. The standardization activities surrounding chiplet interconnects, such as universal chiplet interconnect standards initiatives, are intended to simplify verification by allowing clearly defined interface specifications, but the verification problem is still large because designers have to ensure not only standalone chiplet functionality but also emergent behaviors from interactions between the chiplets under various operating conditions and corner cases [2]. This paper investigates the key challenges involved in

multi-chiplet AI accelerator verification and discusses sophisticated methodologies and techniques through which verification teams can deliver high-quality, bug-free hardware within the limitations of simulation time, system controllability, and debugging complexity.

2. Architectural Complexity of Multi-Chiplet AI Accelerators

Multi-chiplet AI accelerators are a revolutionary shift from conventional monolithic SoC designs, and they introduce novel architectural concerns that directly influence verification approaches. These systems are usually composed of several specialized chiplets, such as compute dies, memory controllers, I/O interfaces, and interconnect fabrics, that are coupled together in one package through advanced packaging technologies like silicon interposers or organic substrates. The trend toward heterogeneous integration via chiplet architectures has been necessitated by the necessity to go beyond the monolithic scaling model, whereby growing die sizes result in exponentially diminishing yields and out-of-reach production costs. Chiplet model standardization has become important to achieve interoperability with chiplets sourced from various vendors and design houses, with the latter already collaborating as industry consortia to define shared electrical, physical, and protocol-level specs that allow chiplets sourced from different sources to be integrated [3]. These standardization activities include not just the electrical and mechanical interfaces but even the behavioral models used for system-level simulation and verification, covering the issues related to power delivery network modeling, thermal interface characterization, and signal integrity analysis across chiplet interfaces [3].

Die-to-die protocols that provide inter-chiplet communication add new layers of complexity with their respective protocol stacks, physical layer definitions, and timing specifications. Individual chiplets can comprise billions of transistors divided into intricate hierarchies of IP blocks, such as tensor processing units, vector engines, multiple levels of memory hierarchies, and advanced power management subsystems. The verification complexity is enhanced by the heterogeneity of these systems, with various chiplets running at various clock domains, voltage islands, and power states, necessitating meticulous verification of clock domain crossings, power sequencing, and inter-chiplet synchronization. Mixed-criticality systems, which are becoming more important in AI accelerator designs that need to support safety-critical inference workloads and best-effort training

workloads, require strict verification approaches to guarantee that high-criticality functions preserve their timing guarantees and functional correctness even under sharing resources with low-criticality tasks [4]. The validation of such mixed-criticality multi-core systems is necessitated by techniques such as temporal and spatial partitioning verification, interference analysis to determine the effect of resource contention on worst-case execution times, and formal methods for proving isolation properties between criticality levels [4]. In addition, the sheer size of these designs—ranging typically over ten billion gates when all chiplets are counted cumulatively—places practical constraints in simulation capacity that prevent it from verifying the entire system as a monolithic unit using traditional verification techniques. The verification needs to cater to not just functional correctness but also non-functional properties like timing predictability, bounds on resource utilization, and fault tolerance mechanisms. For multi-core and multi-chiplet systems, verification methods need to cover multiple levels of abstraction, ranging from unit-level verification of individual IP blocks to integration-level verification of inter-chiplet communication protocols and system-level verification of end-to-end application scenarios [4]. The challenge is compounded by the requirement to validate emergent behavior that exists only when there are interactions among many chiplets under certain timing constraints or resource competition situations, necessitating advanced co-verification environments that can successfully model the timing relationships and cause/effect dependencies among events in various chiplets with potentially asynchronous clock domains [4].

3. Advanced Verification Methodologies

To address the challenges posed by multi-chiplet designs, verification teams employ a combination of advanced methodologies that leverage both hardware acceleration and sophisticated simulation techniques. Emulation platforms provide hardware-accelerated verification environments capable of running designs at speeds orders of magnitude faster than pure software simulation, enabling the execution of complex system-level scenarios and full software stacks that would be impractical in traditional simulators. However, even multi-rack emulation systems have capacity limitations, necessitating strategic design partitioning and selective stubbing of certain design components to fit within available resources. The architectural optimization of multi-die systems requires careful consideration of interconnect topologies and switch architectures that can efficiently route data between

chiplets while minimizing latency and power consumption. Novel switch architectures for multi-die optimization have been proposed to address the challenge of establishing efficient connections between dies, particularly in configurations where traditional mesh or crossbar topologies become impractical due to routing congestion and area constraints [5]. These switch architectures must balance competing objectives, including minimizing hop count between communicating chiplets, reducing hardware complexity to fit within inter-chiplet routing channels, and maintaining sufficient bandwidth to prevent bottlenecks during periods of high inter-die traffic [5].

In such cases, critical interfaces are replaced with behavioral models or transactors that maintain functional accuracy while reducing resource consumption. Complementing emulation, distributed simulation techniques enable verification teams to partition multi-chiplet designs across multiple processing nodes, with each chiplet or subsystem simulated independently on dedicated computational resources. These distributed simulations employ virtual channels—abstract communication interfaces that replace detailed RTL implementations of inter-chiplet buses—to connect the independently simulated components, effectively creating a distributed yet cohesive verification environment. The emergence of chiplet-based architectures for heterogeneous integration has fundamentally transformed the semiconductor landscape, enabling the combination of disparate technologies and process nodes within a single package to achieve levels of integration and customization previously unattainable with monolithic designs [6]. These architectures leverage advanced packaging technologies, including two-point-five-dimensional integration using silicon interposers and three-dimensional stacking with through-silicon vias to achieve high-density interconnections with pitch dimensions measured in tens of micrometers, far exceeding the capabilities of traditional package-level connections [6].

Additionally, intelligent clock gating strategies can significantly improve simulation performance by selectively disabling portions of the design not under active test, thereby reducing event activity and computational overhead without compromising verification coverage of the target functionality. The heterogeneous nature of chiplet-based systems allows for the strategic placement of computational elements, memory hierarchies, and I/O interfaces optimized for their specific functions, with high-performance compute chiplets fabricated on leading-edge process nodes coexisting with analog, RF, and sensor chiplets built on mature, cost-

effective technologies [6]. This heterogeneous integration capability is particularly valuable for AI accelerators, where specialized tensor processing units, high-bandwidth memory interfaces, and power management circuits can be implemented as separate chiplets, each optimized independently before integration at the package level [6]. The verification of such heterogeneous systems must account for the diverse timing characteristics, power profiles, and functional behaviors of each chiplet type, requiring verification environments that can seamlessly integrate models spanning multiple abstraction levels from register-transfer level for digital logic to analog mixed-signal models for power delivery networks and SPICE-level simulations for critical analog interfaces [5].

4. Technical Challenges and Limitations

Notwithstanding advanced methodologies for multi-chiplet verification that have been made available, numerous major issues still hinder verification effectiveness and efficiency. The cost impediment is the major concern since commercial EDA tools that provide distributed simulation capabilities have high licensing costs, and the computational infrastructure needed to support them—such as high-performance compute clusters, large memory resources, and high-bandwidth networking—has high capital outlays and operational costs. The three-dimensional integrated circuit and multi-die system verification challenges differ in essence from those experienced in the conventional two-dimensional planar design, adding new complexities associated with the vertical integration of several silicon layers and the high-density interconnects enabled by through-silicon vias and micro-bumps [7]. These three-dimensional designs allow for record levels of integration density and heterogeneity of functionality in small form factors, but they also bring verification issues related to thermal management between stacked dies, signal integrity through vertical interconnects, and power delivery network integrity between multiple tiers that need to sustain voltage regulation while having different current demands from different functional layers [7].

Debugging distributed multi-chiplet simulation failures constitutes perhaps the most challenging problem, with signal propagation and data flow paths that need to be traced across various simulation instances that could be on different machines with asynchronous execution profiles. Root cause analysis of inter-chiplet transaction failures can take days or even weeks, as there is a need for the engineers to correlate waveforms,

transaction logs, and debug traces from different sources while considering differences in timing and abstraction boundaries caused by virtual channels. With the advent of high-bandwidth chiplet interconnects, advanced machine learning and AI applications have become imperative, where huge data movement among compute chiplets, memory chiplets, and I/O interfaces has become a prime performance bottleneck that needs to be plugged through new interconnect technologies and packaging solutions [8]. Such high-bandwidth interconnect solutions present several challenges such as attaining enough bandwidth density in terabits per second per millimeter of die edge, preserving signal integrity over relatively extended interconnect paths that could extend over several millimeters within the package substrate, controlling power consumption as interconnect data rates exceed multiple gigabits per second per lane, and providing reliable operation over temperature ranges and manufacturing process corners [8].

Gate-level simulation, required for final sign-off and timing-critical path validation, introduces another level of complexity; gate-level netlist simulation of multi-chiplet designs must be scoped to individual chiplets or subsystems to control simulation times and debug tractability. Furthermore, achieving adequate verification coverage across all chiplets, operating modes, and inter-chiplet interaction scenarios requires meticulous planning and resource allocation, as the verification space expands combinatorially with the number of chiplets and their possible configurations. The verification strategy needs to solve not only the functional accuracy of discrete chiplets but also the system-wide behavior resulting from interactions among chiplets, such as cache coherency protocols that cut across multiple dies, memory consistency models in distributed memory hierarchies, quality-of-service assurances for inter-chiplet traffic under congestion, and fault tolerance procedures that need to handle gracefully link errors, chiplet failures, or thermal throttling events [7]. The intricacy is also compounded by having to validate new chiplet interconnect standards and protocols that describe the electrical, physical, and logical aspects of inter-chiplet communication, necessitating full protocol compliance testing and interoperability validation across chiplets from multiple vendors and design generations [8].

5. Emerging Solutions and Best Practices

The semiconductor industry continues to develop innovative solutions and establish best practices to mitigate the challenges associated with multi-chiplet verification. Formal verification techniques,

including formal property checking and equivalence checking, provide mathematically rigorous methods for verifying specific properties and interfaces without exhaustive simulation, particularly valuable for critical inter-chiplet protocols and safety-critical functions. Hybrid verification approaches that strategically combine emulation for system-level validation, distributed simulation for detailed functional verification, and formal methods for protocol compliance are becoming standard practice for complex designs. The application of hybrid automata for modeling and verification of multi-agent systems provides a powerful framework for reasoning about complex distributed systems where multiple autonomous agents interact through well-defined interfaces and protocols [9]. Hybrid automata combine discrete state transitions with continuous dynamics, enabling the formal modeling of systems where digital control logic interacts with analog physical phenomena such as thermal variations, power delivery dynamics, and signal propagation delays that are inherent in multi-chiplet architectures [9]. This formalism is particularly valuable for verifying inter-chiplet communication protocols where discrete protocol state machines must operate correctly despite continuous timing variations, voltage fluctuations, and environmental perturbations that characterize real-world multi-die implementations [9].

The adoption of standardized chiplet interfaces enables the development of reusable verification components and reference models that can be shared across projects and organizations, reducing verification effort and improving confidence in inter-chiplet communication. Verification IP vendors increasingly offer sophisticated models for chiplet interconnect protocols, providing pre-verified transactors, checkers, and coverage collectors that accelerate verification development. Advanced debugging methodologies, including

unified cross-chiplet waveform databases and transaction-level debugging frameworks, help verification engineers navigate the complexity of multi-chiplet interactions more effectively. Machine learning techniques are beginning to be applied to verification, with applications in intelligent testbench generation, coverage closure prediction, and anomaly detection in simulation results, promising to further enhance verification productivity in future design cycles. The emergence of AI-driven methodologies for chip design and verification represents a paradigm shift in how semiconductor development is approached, with artificial intelligence algorithms being deployed to optimize various aspects of the design flow from initial architecture exploration through physical implementation and verification [10]. AI-driven chip die size estimation frameworks leverage machine learning models trained on historical design databases containing thousands of completed chip designs with their associated area, power, and performance metrics to predict the physical characteristics of new designs early in the development cycle [10]. These predictive models can estimate die size with accuracy levels approaching ten to fifteen percent based on high-level architectural parameters such as logic gate counts, memory configurations, I/O requirements, and target frequency specifications, enabling more informed decisions about chiplet partitioning strategies and package planning before substantial engineering resources are committed to detailed implementation [10]. The application of artificial intelligence extends beyond physical design estimation to encompass verification planning, where machine learning algorithms analyze design complexity metrics, historical bug distributions, and verification resource allocation patterns to recommend optimal verification strategies that maximize coverage while minimizing schedule and cost [10].

Table 1: Multi-Chiplet System Components and Verification Requirements [3, 4]

Compute Dies	Tensor processing units, vector engines	Billions of transistors	Functional correctness, timing predictability	Resource utilization bounds, fault tolerance
Interconnect Fabrics	Die-to-die communication protocols	Variable scale	Protocol stack verification	Physical layer specifications, timing requirements

Table 2: Multi-Die Interconnect Architecture Optimization Requirements [5, 6]

Interconnect Aspect	Design Considerations	Optimization Objectives	Verification Requirements	Implementation Challenges
Switch Architecture	Topology selection (mesh vs. crossbar)	Minimize hop count between chiplets	Protocol compliance, latency validation	Routing congestion, area constraints
Hardware	Inter-chiplet routing	Reduce complexity while	Resource utilization	Fit within routing

Complexity	channels	maintaining functionality	verification	channel constraints
Bandwidth Management	Traffic handling capacity	Prevent bottlenecks during high traffic	Throughput testing, congestion analysis	Peak load handling
Latency Optimization	Data path efficiency	Minimize communication delays	Timing verification across dies	Multi-hop latency accumulation
Power Consumption	Interconnect energy efficiency	Reduce power for data movement	Power profile validation	Balance performance and efficiency
Process Node Heterogeneity	Leading-edge and mature nodes	Optimize each chiplet for function	Cross-technology verification	Interface compatibility across nodes
Packaging Technology	2.5D/3D integration methods	High-density connections (tens of μm pitch)	Electrical and thermal modeling	Signal integrity, thermal management
Chiplet Specialization	Function-specific optimization	Independent chiplet optimization	Integration-level verification	System-level behavior validation

Table 3: Key Verification Challenges in Multi-Chiplet Systems [7, 8]

Challenge Category	Primary Issues	Impact Level	Time Impact
Cost Barriers	EDA tools, infrastructure investment	High	Ongoing operational expenses
3D Integration	Thermal management, vertical interconnects	High	Complex modeling required
Debug Complexity	Multi-instance signal tracing	Very High	Days to weeks per failure
High-Bandwidth Interconnects	Bandwidth density (terabits/sec/mm)	High	Power and signal integrity issues
Gate-Level Simulation	Timing-critical path validation	Medium	Extended simulation time
Coverage Achievement	Combinatorial verification space	High	Meticulous resource planning

Table 4: AI and Machine Learning Applications in Verification [9, 10]

AI Application	Methodology	Prediction Accuracy	Input Parameters	Output/Benefits
Die Size Estimation	ML models on historical databases	10-15% accuracy	Logic gate counts, memory configs, and I/O requirements	Early chiplet partitioning decisions
Verification Planning	Design complexity analysis	Optimized strategies	Historical bug distributions, resource patterns	Maximize coverage, minimize cost
Testbench Generation	Intelligent automation	Enhanced productivity	Design specifications	Automated test creation
Coverage Closure	Predictive analytics	Targeted testing	Coverage gaps, historical data	Efficient coverage achievement
Anomaly Detection	Pattern recognition	Improved reliability	Simulation results	Early bug detection
Physical Prediction	Trained prediction models	High-level estimates	Architectural parameters, frequency specs	Package planning optimization

6. Conclusions

The verification of multi-chiplet AI accelerators represents a critical inflection point in semiconductor design methodology, where the economic and technical imperatives of heterogeneous integration collide with the practical limitations of traditional verification approaches. As this article has demonstrated, the transition from monolithic to chiplet-based architectures, while enabling unprecedented levels of functional specialization and manufacturing efficiency, fundamentally transforms the verification landscape by introducing challenges that span multiple

dimensions, including architectural complexity, simulation scalability, debugging tractability, and coverage adequacy. The sophisticated verification methodologies currently employed—encompassing hardware-accelerated emulation, distributed simulation with virtual channels, formal verification techniques, and hybrid automata modeling—represent substantial advances over conventional approaches, yet they simultaneously demand significant capital investment in computational infrastructure, premium electronic design automation tool licensing, and highly specialized engineering expertise. The debugging challenges inherent in multi-chiplet systems, where

failure analysis can extend over weeks as engineers correlate events across distributed simulation instances and abstraction boundaries, highlight the profound shift in engineering practices required for this new paradigm. The emergence of standardized chiplet interfaces and artificial intelligence-driven verification methodologies offers promising pathways toward improved efficiency and effectiveness, with machine learning algorithms demonstrating capability in predictive die size estimation, intelligent test generation, and verification strategy optimization. But the underlying problem is still that verification complexity grows non-linearly as a function of chiplet count and system heterogeneity, building verification space combinatorially and eluding exhaustive coverage by any one approach. Companies that will thrive in this climate need to adopt hybrid verification approaches that responsibly balance resources among emulation, simulation, and formal techniques and invest in state-of-the-art debugging infrastructure and developing teams with interdisciplinary skills across digital design, analog mixed-signal verification, thermal modeling, and, more recently, data science and machine learning. As the semiconductor business continues its relentless journey toward higher integration density and functional specialization through chiplet-based designs, excellence in verification will increasingly be what separates successful deployments from expensive failures, so whole-verification will no longer be a quality assurance activity but a strategic business differentiator in the competition to deliver next-generation artificial intelligence accelerators driving the computational needs of new applications.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available

on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

- [1] Andras Poppe et al., "Thermal Measurement and Modeling of Multi-Die Packages," July 2009. ResearchGate. [Online]. Available: https://www.researchgate.net/publication/224386369_Thermal_Measurement_and_Modeling_of_Multi-Die_Packages
- [2] Vivek Gujar et al., "Chiplet Technology Revolutionizing Semiconductor Design - A Review," February 2024, ResearchGate. [Online]. Available: https://www.researchgate.net/publication/378156097_Chiplet_Technology_Revolutionizing_Semiconductor_Design-A_Review
- [3] Anthony Mastroianni et al., "Proposed Standardization of Heterogeneous Integrated Chiplet Models," October 2021. ResearchGate. [Online]. Available: https://www.researchgate.net/publication/358325947_Proposed_Standardization_of_Heterogeneous_Integrated_Chiplet_Models
- [4] Samuel Pagliarini, "Verification methods, techniques, and processes for mixed-criticality applications on multi-core systems," April 2015, ResearchGate. [Online]. Available: https://www.researchgate.net/publication/339505204_Verification_methods_techniques_and_processes_for_mixed-criticality_applications_on_multi-core_systems
- [5] Feng Yu et al., "A Novel Switch Architecture for Multi-Die Optimization with Efficient Connections," August 2024, ResearchGate. [Online]. Available: https://www.researchgate.net/publication/383110922_A_Novel_Switch_Architecture_for_Multi-Die_Optimization_with_Efficient_Connections
- [6] Murali Krishna Reddy Mandalapu, "Emerging Chiplet-Based Architectures for Heterogeneous Integration," March 2025. ResearchGate. [Online]. Available: https://www.researchgate.net/publication/389855062_Emerging_Chiplet-Based_Architectures_for_Heterogeneous_Integration
- [7] Niranjana Ravishankar, "Verification Challenge in 3D Integrated Circuits IC Design," February 2020, ResearchGate. [Online]. Available: https://www.researchgate.net/publication/388763899_Verification_Challenge_in_3D_Integrated_Circuits_IC_Design
- [8] Shenggao Li et al., "High-Bandwidth Chiplet Interconnects for Advanced Packaging Technologies in AI/ML Applications: Challenges and Solutions," January 2024, ResearchGate. [Online]. Available: https://www.researchgate.net/publication/386162569_High-Bandwidth_Chiplet_Interconnects_for_Advanced_Packaging_Technologies_in_AI_ML_Applications_Challenges_and_Solutions

Bandwidth Chiplet Interconnects for Advanced Packaging Technologies in AIML Applications Challenges and Solutions

- [9] Ammar Mohammed & Ulrich Mohammed, "Multi-Agent Systems Modeling and Verification Using Hybrid Automata," May 2009, ResearchGate. [Online]. Available: https://www.researchgate.net/publication/48172941_Multi-Agent_Systems_Modeling_and_Verification_Using_Hybrid_Automata
- [10] Puneet Gupta, "AI-Driven Chip Die Size Estimation: A Technical Framework," August 2025, ResearchGate. [Online]. Available: https://www.researchgate.net/publication/394294824_AI-Driven_Chip_Die_Size_Estimation_A_Technical_Framework