International Journal of Modern Scientific Discoveries
ISSN: 3107-9377 | Volume 2 Issue 1, 43-56, January-March 2026
DOI: https://doi.org/10.64137/31079377/IJMSD-V2I1P107

Received: 17/02/2026 Revised: 04/03/2026 Accepted: 09/03/2026 Published: 17/03/2026

Original Article

Scalable ASIC Design Methodologies for Multi-Chiplet Systems

VENKATAKRISHNA BRAHMAM POGADADANDA
Senior RTL Design Engineer at AMD-Xilinx, USA.

ABSTRACT: *The semiconductor industry is rapidly moving from conventional monolithic System-on-Chip (SoC) designs to chiplet-based architectures as a practical solution driven by mounting design complexity, rising manufacturing costs, and the limits of advanced process scaling. Chiplet-based systems enable heterogeneous integration, i.e., integrating specialised functional blocks manufactured using ent process technologies into a single package, hence The evolution of multi-chiplet systems increasingly challenges standard ASIC design methodologies by requiring solutions for high-speed interconnects, power delivery, thermal management, verification complexity and advanced packaging integration. This work presents a scalable ASIC design methodology for multi-chiplet systems, focusing on design efficiency, modularity and system-level reliability, while enabling future scalability. The proposed methodology provides a hierarchical and reusable design framework that combines standardised chiplet interfaces, optimal interconnect topologies, power-efficient floorplanning methodologies, and scalable verification tools to improve the development cycle. Co-design methodologies that address the three challenges in the design process at an early stage, i.e. silicon, packaging and thermal challenges, are presented in this study to reduce integration risks along with performance bottlenecks. This study demonstrates the effectiveness of the proposed methodology through a detailed case study of a heterogeneous multi-chiplet computing architecture. The results show significant improvements in design turnaround time, connection efficiency, temperature distribution, and verification coverage over traditional ASIC flows. The results emphasise the benefits of modular reuse and standardised communication protocols, enabling scalable system growth without significantly increasing design complexity. This research opens a path for next-generation semiconductor design by providing a practical and scalable framework for high-performance multi-chiplet ASIC devices. The proposed paradigm offers valuable insights for researchers and industry specialists working on effective heterogeneous integration and innovative semiconductor architectures.*

KEYWORDS: *ASIC Design, Multi-chiplet Systems, Chiplet Architectures, Scalable Semiconductor Design, Heterogeneous Integration, Advanced Packaging Technologies, High-Speed Interconnects, Network-on-Package (NoP), Thermal and Power Optimisation, Silicon Interposers, Verification and Validation, Electronic Design Automation (EDA), Hardware Scalability, 2.5D and 3D Integration, UCIe Standards, AI Accelerators, High-Performance Computing (HPC), Modular ASIC Frameworks.*

1. INTRODUCTION

1.1. BACKGROUND AND EVOLUTION OF ASIC ARCHITECTURES

Application-specific Integrated Circuits (ASICs) have driven the evolution of modern computer systems, enabling highly tuned performance for certain applications. Traditionally, ASIC designs used monolithic architectures, integrating all functional components one a single silicon chip. This approach offers advantages such as reduced communication latency and easier integration. For decades, the continued downsizing of transistors under Moore's Law has enabled the semiconductor industry to improve performance while reducing cost per transistor. However, as we have entered sub-7nm process technologies, scaling benefits have begun to diminish due to ever-increasing manufacturing complexity, higher development costs and physical limitations in power density and transistor leakage.

These limitations have led the industry to explore other architectural approaches that might provide performance benefits without relying only on transistor scaling. As a result, chiplet-based designs have become a promising alternative. In chiplet-based systems, not all functionality is integrated on a single die; instead, the design is broken into smaller, reusable functional modules connected using novel packaging techniques. This modular design approach improves production yield, scalability and the integration of multiple process technologies in a single package.

ASIC designs based on chiplet technology are widely used in applications such as AI accelerators, HPC, networking infrastructure and automotive electronics. Chiplet techniques are increasingly used by companies such as AMD, Intel, NVIDIA and TSMC to improve flexibility, reduce development time and optimise system-level performance. Multi-chiplet systems have become more feasible with advanced packaging techniques such as 2.5D integration, 3D stacking and silicon interposers.

This work is an open access and licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License
(https://creativecommons.org/licenses/by-nc-nd/4.0/)

Hence, scalable ASIC approaches are more critical to enable the next generation of compute-intensive and power-efficient semiconductor devices.

1.2. CHALLENGES IN MULTI-CHIPLET ASIC DESIGN

Multi-chiplet systems offer significant advantages in scalability and diversity, but they also introduce complex design and implementation challenges. One major issue is communication latency between chips. Unlike monolithic ASICs, where components communicate via on-die interconnects, chiplet-based systems rely on package-level communication, which can introduce higher latency and bandwidth limits if not carefully managed.

One big challenge is how to construct high-bandwidth, low-power interconnects to move data between the chiplets efficiently. Modern AI and HPC workloads need huge data movement, and providing energy-efficient connections becomes increasingly difficult. They must balance bandwidth, latency and power consumption while ensuring reliable communication across dies.

The architectural difficulty of chiplet synchronisation. Different chiplets may run on different clock domains or process technologies, requiring sophisticated synchronisation algorithms to provide temporal consistency and allow coherent data exchange. Any other temporal inconsistency might impact overall system reliability and performance.

Thermal control becomes substantially more complex in densely populated multi-chiplet systems. Heavy computing jobs typically create heat hotspots that affect performance stability and device longevity. Proper heat dissipation techniques and thermally aware floorplanning are required to prevent overheating and maintain operational efficiency.

Chiplet integration increases the complexity of the power delivery network (PDN). Co-design of packaging and power distribution systems is required to provide a continuous power supply across more connected dies while minimising voltage variations and IR drops.

Other concerns include temporal closure and signal integrity. High-speed inter-chip communication links are sensitive to noise, crosstalk and electromagnetic interference, which may degrade signal quality. This is a significantly harder problem for temporal closure across multiple chiplets than for monolithic systems.

Adding several separately created dies makes verification substantially more challenging in heterogeneous chiplet systems. This increases simulation and testing effort, as designers must prove interoperability, functional correctness and reliable communication across the whole system.

In addition, the advanced packaging techniques put production constraints on assembly precision, thermal expansion and reliability. Yield optimisation becomes an important problem, since a failure in a single chiplet or packaging layer may endanger the functioning of the complete system. As the chiplet ecosystem grows, the industry must maintain cost-effective manufacturing and sustainable scalability.

1.3. PROBLEM STATEMENT

Existing ASIC design methodologies were developed for monolithic designs and are not sufficient for the growing demands of multi-chiplet systems. Modular and heterogeneous semiconductor designs are the trend, and the traditional methods are no longer enough to handle the complexities of inter-chip communication, synchronisation, temperature control and packaging integration.

No scalable, standardised frameworks currently exist for successful chiplet integration across a wide range of applications and production processes. Designers typically face major challenges in optimising performance, power efficiency, manufacturing costs and system reliability while maintaining scalability.

In addition, the limits of reusable and modular design methodologies contribute to longer development times and more complex verification. Without methods specifically built for multi-chiplet designs, achieving efficient, large-scale semiconductor integration becomes increasingly difficult. Therefore, scalable ASIC design methodologies are needed to enable modular development, improve system-level integration, and increase overall design efficiency for next-generation chiplet-based computing systems.

1.4. MOTIVATION

The rapidly growing applications of artificial intelligence, cloud computing, edge devices, and high-performance computing have created tremendous demand for scalable, energy-efficient semiconductor systems. The conventional monolithic ASIC design approach is becoming less feasible because of production costs, power restrictions, and diminishing technology-scaling benefits. Recent industrial developments have further pushed chiplet-based designs as a more flexible and sustainable alternative.

At the same time, advances in packaging technologies such as 2.5D integration, 3D stacking and high-speed die-to-die interconnects have made multi-chiplet systems more economically viable. Organisations are increasingly adopting design methodologies that reduce development complexity, increase design reuse and speed time-to-market while maintaining high performance and reliability.

Scalable ASIC solutions promise heterogeneous integration, improved manufacturing yield, and modular design ecosystems. These strategies might greatly accelerate future semiconductor development by enabling reusable chiplet components and standardised integration frameworks. This work aims to develop realistic, scalable solutions that meet the evolving needs of future-generation computing systems.

2. LITERATURE REVIEW

2.1. EVOLUTION OF CHIPLET-BASED ARCHITECTURES

To address power, manufacturing complexity and scalability issues, the semiconductor industry has gradually shifted from traditional monolithic System-on-Chip (SoC) designs to chiplet-based designs. The previous System-on-Chip (SoC) approach integrated all these functional blocks on a single chip, enabling compactness and fast communication. However, it became more difficult and expensive as manufacturing nodes advanced. As transistor density increased, problems including yield loss, heat hotspots and design complexity became major challenges.

Chiplet-based systems have emerged as a promising alternative by partitioning many designs into smaller functional dies interconnected in a package. This disaggregated approach enables designers to reuse intellectual property (IP), incorporate process advances and improve manufacturing yields. Chiplet ecosystems provide greater flexibility, faster development cycles and lower production expenses than monolithic SoCs.

Industry standards have greatly accelerated chiplet adoption. The Universal Chiplet Interconnect Express (UCIe) provides a common communication architecture for die-to-die connections. Intel's Embedded Multi-die Interconnect Bridge (EMIB) and Bunch of Wires (BoW) protocol enables efficient high-bandwidth communication across these chiplets. Together, these technologies enable scalable, interoperable multi-chiplet systems for diverse computer applications.

2.2. EXISTING ASIC DESIGN METHODOLOGIES

Conventional ASIC design methodologies follow a systematic process to turn system specifications into manufacturable silicon designs. The traditional design flow usually consists of architectural planning, Register Transfer Level (RTL) coding, functional verification, synthesis, physical design, timing closure and finally GDSII generation for manufacture. This RTL-to-GDSII flow is widely used because it provides predictable design cycles while enabling high-performance integrated devices.

Current ASIC development heavily depends on reusable IP blocks and modular design techniques. Designers often build bigger systems by combining proven CPUs, memory controllers, communication interfaces, and accelerators to save development time and verification effort. Electronic Design Automation (EDA) technologies improve integration through automated synthesis, placement, routing and timing optimisation.

Traditional ASIC techniques were originally designed for monolithic chips and suffer from many challenges when used with multi-chiplet systems. Latency, synchronisation, power delivery, and signal integrity are other challenges of inter-chip communication. Current design methodologies lack a common technique to distribute functionality across many chiplets while maintaining performance efficiency. Furthermore, verification complexity increases when interactions across several dies and package-level interconnects must be verified.

One of the key challenges in this context is heterogeneous integration of chiplets fabricated using diverse manufacturing processes. Conventional methodologies cannot successfully solve cross-die timing closure, thermal interactions and package-aware optimisation. Hence, we examine the latest scalable ASIC techniques to improve support for modular chiplet-based systems.

2.3. INTERCONNECT TECHNOLOGIES FOR MULTI-CHIPLET SYSTEMS

Communication between chiplets directly affects performance, power consumption, and scalability; thus, effective interconnection technology is crucial to the success of multi-chiplet systems. As chiplet-based designs mature, advanced packaging and connection technologies are key to ensuring high-bandwidth, low-latency communication.

Silicon interposers are widely used in multi-chiplet integration to offer compact routing paths among dies. Interposers act as an intermediate layer, enabling fine-pitch interconnects and reducing communication distances compared to standard package substrates. The silicon interposer technique greatly improves signal integrity and enables high-speed data transfer between chiplets, while reducing their power consumption.

Another major enhancement is 2.5D and 3D integration. In 2.5D integration, several chiplets are placed side-by-side on one common interposer for effective lateral communication. This approach is common in high-performance computing and AI accelerators because it trades performance for manufacturability. On the other hand, 3D integration stacks chiplets vertically using Through-Silicon Vias (TSVs), which reduces connection lengths and increases bandwidth density. While 3D stacking delivers considerable performance benefits, it also introduces problems in heat dissipation and manufacturing.
High-speed serial and parallel interconnects are required for reliable die-to-die communication. Serial interconnects provide scalable capacity with less physical cabling, while parallel interconnects offer lower latency and simpler protocols for short-range communication. Such communication protocols are increasingly being standardised through technologies such as UCIe, BoW, etc.

Network-on-Chip (NoC) and Network-on-Package (NoP) designs improve communication efficiency in complex chiplet systems. NoC enables structured communication within a single chiplet, while NoP extends this to multiple chiplets in a package. These network-centric approaches boost scalability and simplify traffic management in large systems.

Every networking technology includes trade-offs between bandwidth, latency, power consumption, cost, and thermal performance. Designers must carefully balance these factors based on application demands, system size, and manufacturing constraints to achieve optimal performance in scalable multi-chiplet ASIC systems.

TABLE 1 Interconnect Technologies for Multi-Chiplet Systems

Technology	Description	Advantages	Limitations	Typical Applications
UCIe (Universal Chiplet Interconnect Express)	Standardised die-to-die communication protocol for chiplet integration	High bandwidth, interoperability, low latency, ecosystem support	Requires advanced packaging support	AI accelerators, HPC systems, heterogeneous chiplet platforms
BoW (Bunch of Wires)	Low-power parallel die-to-die interconnect protocol	Simple implementation, low energy consumption	Limited scalability for extremely large systems	Short-range chiplet communication
Silicon Interposer	Intermediate silicon layer connecting multiple chiplets	Fine-pitch routing, high signal integrity, reduced latency	Increased manufacturing complexity and cost	2.5D integrated systems
2.5D integration	Chiplets placed side-by-side on a shared interposer	High bandwidth, improved thermal handling, easier manufacturability	Larger package footprint	AI processors, GPUs, HPC accelerators
3D integration	Vertical stacking of chiplets using TSVs	Extremely high bandwidth density, reduced interconnect length	Thermal challenges, complex manufacturing	Advanced AI and memory-intensive systems
Network-on-Package (NoP)	Packet-based communication framework across chiplets	Scalable communication, dynamic routing, congestion management	Additional routing complexity	Large-scale multi-chiplet ASICs
High-Speed Serial Interconnects	Serial communication links between chiplets	Reduced wiring complexity, scalable bandwidth	Serialisation latency	Long-distance die communication
Parallel Interconnects	Multiple parallel communication channels between dies	Low latency, simpler control logic	Increased routing area and pin count	Short-range high-speed transfers

2.4. VERIFICATION AND TESTING APPROACHES

Verification and testing are vital to ASIC development, and even more so for multi-chiplet systems, where several dies must work together reliably. To check design behaviour before production, functional verification methods such as simulation, formal verification, and hardware emulation are typically used. Using these approaches, you may discover logic challenges, communication incompatibilities and protocol mistakes early in the creation phase.

Design-for-Testability (DFT) techniques have become increasingly important for chiplet-based designs because manufacturing defects can occur at the die and package levels. Inter-chiplet testing and fault localisation typically rely on adaptations of traditional scan-chain methods. Integrated Built-In Self-Test (BIST) structures embedded into chiplets provide validation and increase test coverage.

IEEE standards-based boundary scan methods may be used to examine connectivity between chiplets and package interfaces. These techniques allow detection of faults in connectivity, signal integrity and assembly without physical access to internal nodes. Interconnect testing is particularly important for sophisticated packaging methods like 2.5D and 3D integration.

Recently, researchers have studied AI-supported verification strategies. Machine learning techniques have been used to predict defect-prone design regions, improve test development and speed up coverage analysis. AI-driven automation boosts verification efficiency by reducing simulation time and better detecting hidden corner-case faults. As chiplet systems grow more complex, smart verification frameworks will likely become an integral part of next-generation ASIC design flows.

3. PROPOSED METHODOLOGY

3.1. OVERVIEW OF THE PROPOSED SCALABLE ASIC FRAMEWORK

The proposed approach provides a scalable ASIC design methodology for multi-chiplet systems that require very high computational performance, design flexibility and fast time-to-market. Rather than the traditional monolithic System-on-chip (SoC) paradigm, the proposed framework adopts an architecture-centric approach that decomposes complex ASIC functionality into smaller, manageable chiplets, interconnected over a standardised communication infrastructure. This approach helps designers tackle the ever-increasing challenges of advanced process nodes, power density and manufacturing complexity.

The framework is built on a fundamentally modular and hierarchical design concept. Each chiplet is a self-contained functional unit that can be designed, validated and optimised individually, but can work together as part of the larger system. This modularity enables greater reuse of intellectual property, easier debugging, and less redesign in later versions.

Heterogeneous integration provides scalability, where chiplets made with multiple process technologies can be integrated in a similar package. Defined interfaces and protocol-aware communication layers provide interoperability, enabling seamless integration of computing, memory, accelerator and I/O chiplets. The architecture therefore provides scalability for these workloads, along with design freedom, performance efficiency and manufacturability.

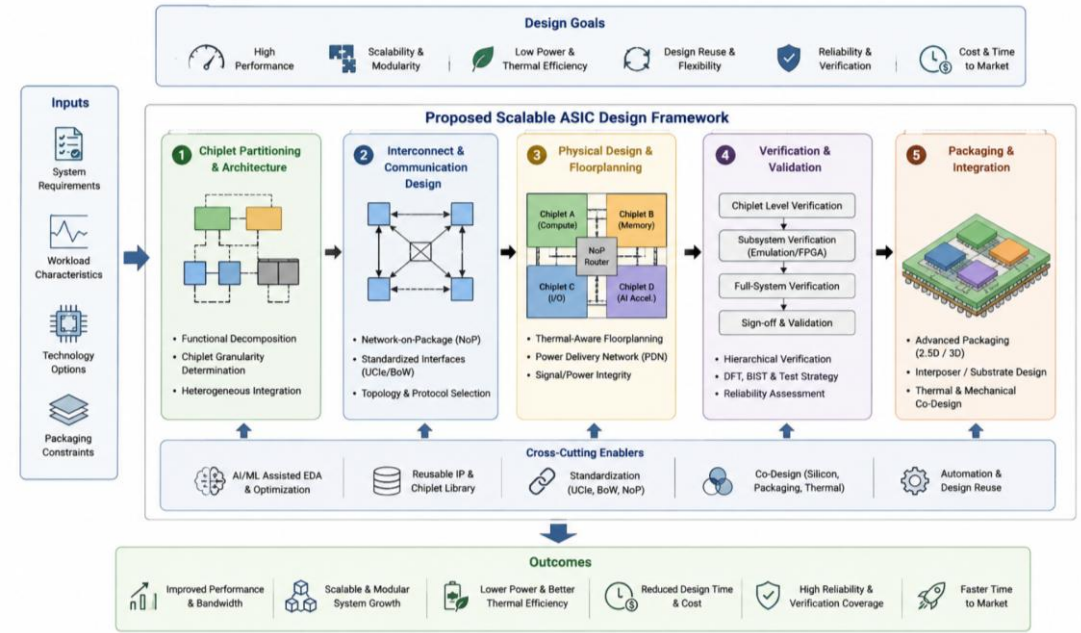


FIGURE 1 Proposed Scalable ASIC Design Framework

3.2. MODULAR CHIPLET PARTITIONING STRATEGY

One key feature of scalable multi-chiplet ASIC design is the efficient partition of system functions into optimised chiplets. The proposed method adopts a systematic partitioning strategy to decompose complex ASIC subsystems into modular units based on functions, communication protocols, thermal characteristics and performance requirements. This makes the design more manageable, allowing different parts of the system to be developed in parallel and optimised independently.

The first step in partitioning is functional decomposition. Task-specific high-performance modules, including AI accelerators, DSP engines, CPUs, memory controllers, and high-speed I/O interfaces, are grouped into separate chiplets. Modules that rely heavily on internal communication are clustered together to reduce the need for interaction between chiplets. For scalability, functions that are not closely linked are assigned to different chiplets.

TABLE 2 Chiplet Partitioning Criteria

Chiplet Type	Primary Function	Key Design Considerations	Preferred Process Node	Benefits
Compute Chiplet	AI acceleration, tensor processing, vector computation, DSP operations	High performance, low latency, thermal management, power efficiency	Advanced nodes (5nm / 3nm)	Faster computation, scalable parallel processing
Memory Chiplet	High-bandwidth data storage and cache management	Memory bandwidth, low access latency, power optimisation	Advanced memory technology nodes	Improved data throughput and reduced bottlenecks
I/O Chiplet	PCIe communication, external interfaces, system coordination	Signal integrity, protocol compatibility, lower manufacturing cost	Mature nodes (12nm / 16nm)	Cost-effective integration and reliable communication
Analog/RF Chiplet	Signal conversion, RF communication, power regulation	Noise isolation, thermal stability, analogue precision	Mature analogue-friendly nodes	Better signal quality and reduced interference
Control Chiplet	Clock distribution, synchronisation, workload scheduling	Low latency synchronisation, fault management	Mid-range technology nodes	Efficient coordination across chiplets
Security Chiplet	Encryption, authentication, hardware-level security	Secure communication, low overhead, tamper resistance	Secure optimised nodes	Enhanced system reliability and protection
Interconnect Router Chiplet	Network-on-Package (NoP) routing and communication management	Bandwidth optimisation, congestion control, packet routing efficiency	Advanced communication nodes	Improved scalability and communication efficiency

Several other variables govern the splitting method. In addition, chiplet allocations include communication intensity, power density, latency sensitivity, silicon space use, and manufacturing yield. Analogue and RF blocks may be separated from digital processor blocks to minimise interference while improving process optimisation. Similarly, memory-heavy subsystems may be placed closer to high-bandwidth memory chiplets, reducing access latency and power use.

The approach emphasises diverse integration and intellectual property reuse. Verified IP blocks may be packaged as reusable chiplets, avoiding duplication of development processes. Chiplets developed at different technical nodes can also be combined in a single package, enabling cost optimisation. High-performance compute chiplets might be implemented on advanced nodes, while analogue or I/O chiplets could benefit from mature, less expensive technologies.

That means interfaces need to be standardised so independently manufactured chiplets can talk to each other seamlessly. The framework uses industry-standard die-to-die communication protocols such as UCIe and BoW. Standard PHY layers, clock synchronisation mechanisms and protocol abstraction methods enable compatibility across manufacturers and design environments. This standardisation makes integration easier and improves long-term scalability and ecosystem compatibility.

3.3. SCALABLE INTERCONNECT ARCHITECTURE

The proposed solution provides a scalable interconnect architecture that enables efficient communication across distributed functional units in multi-chiplet systems, ensuring high-bandwidth, low-latency, reliable data transfer. This communication architecture aims to effectively support the growing traffic demands while maintaining the energy efficiency and signal integrity of chiplet interfaces.

The design is based on a high-speed communication infrastructure using Network-on-Package (NoP) technology. The NoP architecture allows chiplets to communicate in a packetised fashion using distributed routing algorithms, rather than conventional bus-based communication, which has scalability limitations. Each chiplet has lightweight communication routers that route information in and out of the chiplet. This allows the system to scale efficiently by adding additional chiplets.

Several architectural modifications optimise latency. Adaptive routing algorithms prefer short communication paths by dynamically choosing the best transmission routes based on network congestion and traffic conditions. Concurrent communication channels and low-latency buffering systems greatly minimise packet-transfer delays. The framework additionally supports clock-domain crossover synchronisation schemes that minimise timing overhead between independently operating chiplets.

The proposed connection mechanism is an important part of efficient bandwidth allocation. Dynamic bandwidth management methods distribute communication resources according to the amount of work needed, preventing both resource under-utilisation and network congestion. Congestion control strategies monitor network traffic behaviour and reroute data streams to prevent packet delay and system performance impairment. Quality-of-Service (QoS) standards support low-latency, highly demanding applications, such as AI inference and real-time computation.

Signal integrity is tightly controlled to provide reliable high-speed connectivity. As data speeds increase, inter-chiplet communication channels become increasingly sensitive to noise, crosstalk and transmission losses. The suggested architecture achieves reliable connectivity through equalisation techniques, low-swing signalling, impedance matching and error-correction algorithms. Furthermore, advanced packaging methods such as silicon interposers, along with these organic substrates, are considered to reduce interconnect parasitics and improve electrical performance.

The scalable interconnect architecture provides a solid communication infrastructure to support future high-performance chiplet-based ASIC platforms, providing the right trade-off between performance, scalability, and energy efficiency.

3.4. VERIFICATION AND VALIDATION METHODOLOGY

Verification and validation are crucial steps in developing multi-chiplet ASICs, given the increased complexity of dispersed designs and heterogeneous integration. The proposed approach provides a hierarchical verification architecture that enables comprehensive verification at the chiplet, subsystem, and full-package integration levels.

The verification process begins with independent chiplet validation. Simulation-Based Functional Verification: Each chiplet is functionally verified using simulation-based methodologies to verify logical correctness, protocol compliance and temporal performance. Unit testing ensures that the constituent modules work properly before moving to the next level of integration. This hierarchy greatly reduces debugging difficulty by helping spot probable design faults early in the development cycle.

At the subsystem level, transaction-level simulation environments evaluate communication protocols, synchronisation mechanisms and data-transfer reliability between chiplets. The die-to-die connections, clock domain crossings, and packet-routing logic within the NoP architecture are explicitly studied. Integration-level verification verifies the system's functionality as a whole under realistic workloads.

The strategy combines pre-silicon and post-silicon validation approaches. Pre-silicon validation relies on thorough RTL simulation, formal verification and static timing analysis to detect logical and structural design errors before fabrication. Formal verification techniques are particularly effective for validating communication protocols and demonstrating equivalence between design abstraction levels.

The approach includes FPGA prototypes and hardware emulation platforms to accelerate their system-level validation. FPGA-based prototypes can run the workload near real time and provide key performance insights that are difficult to get from simulation alone. Emulation systems allow designers to assess the software/hardware interface, boot process, and application-level workloads before tape-out.

The system incorporates extensive regression testing automation. Design changes are made using extensive regression suites to check design stability regularly. Automation reduces human effort and ensures new changes do not break previously validated functionality.

Coverage-driven verification approaches are used to quantify verification coverage. Functional coverage, code coverage, toggle coverage, assertion coverage and other metrics indicate untested conditions and build trust in verification. This approach reduces overall development risk by combining hierarchical validation, automated testing and coverage analysis to ensure robust and reliable operation of multi-chiplet ASIC systems.

3.5. THERMAL AND POWER OPTIMISATION FRAMEWORK

Thermal and power management are important concerns in today's multi-chiplet ASIC systems due to increasing transistor density and stringent performance requirements. The proposed method includes a comprehensive thermal and power optimisation framework for enhancing energy efficiency, operational reliability, and long-term system stability.

The framework begins with thermally aware floorplanning in the early stages of physical design. Strategically placing high-power chiplets reduces thermal hotspots and improves heat dissipation across the package. During placement, we use thermal models to predict heat distribution and adjust the chiplet arrangement accordingly. Separating heat-sensitive analogue/memory components from high-temperature computation regions maintains performance stability.

Hierarchical power distribution networks are optimised for power transmission to decrease voltage drops and current fluctuations. Each chiplet has its own power domain; therefore, power may be controlled very finely, and scalability can be achieved depending on the workload. Dynamic voltage and frequency scaling (DVFS) improves energy efficiency by adjusting operating conditions based on data-processing requirements.

Adaptive systems for thermal management have been developed that consist of thermal sensors that measure the chiplet ambient temperature continuously. Intelligent control systems may continuously redistribute routine tasks, adjust clock frequencies, or switch on cooling mechanisms when normal temperature limitations are surpassed. This preventive control prevents thermal runaway and increases device durability.

The plan also emphasises energy-efficient communication methods under the NoP model. Low-power signalling approaches, adaptive link activation, and traffic-aware routing algorithms reduce communication-related power consumption while meeting performance requirements. Together, these approaches constitute an integrated optimisation framework that enables high-performance multi-chiplet ASICs that meet strict thermal and power constraints.

3.6. COMBINING ELECTRONIC DESIGN AUTOMATION WITH DESIGN AUTOMATION

Effective integration with modern Electronic Design Automation (EDA) technologies is required to produce multi-chiplet ASIC systems. The proposed approach uses an automation-based procedure to reduce design complexity while boosting productivity and design accuracy.

EDA tools are playing a significant role in enabling chiplet partitioning, floorplanning, timing analysis, verification and physical implementation. Automated floorplanning and routing methods help designers optimise chiplet placement, interconnect arrangement, and power distribution while satisfying thermal and temporal limits. Both methods greatly minimise manual design iterations and improve implementation efficiency.

The system includes AI-assisted optimisation techniques that use machine learning algorithms to predict these congestion patterns, increase routing paths, and optimise power-performance trade-offs. AI-enabled design exploration enables quick investigation of optimal architectural combinations in large design spaces.

Careful integration with existing ASIC design flows maintains compatibility with industry-standard RTL design, synthesis, verification and sign-off methodologies. It enables enterprises to use chiplet-based designs instead of having to overhaul their present process or tooling environments. We present a plan of action that enables scalable, efficient manufacturing of next-generation single-chiplet ASIC devices by combining conventional machining, AI-assisted optimisation techniques, and standard design principles.

4. CASE STUDY

4.1. SYSTEM OVERVIEW

To demonstrate the practicality of scalable ASIC design methodologies, this case study examines a multi-chiplet AI accelerator platform for high-performance deep learning inference and training workloads. The system was intended to support large transformer models and data-intensive matrix operations common to modern artificial intelligence workloads. At advanced process nodes, conventional monolithic ASIC designs may face power-density, yield and manufacturing challenges. To address these issues, the design uses a chiplet-based architecture that spreads functions across several interconnected dies.

The ideal system has four compute chiplets, two high-bandwidth memory (HBM) chiplets, and a centralised I/O and control die, all in one package. Each compute chiplet contains dedicated tensor processing units, vector engines and local SRAM buffers for concurrent AI workloads. The architecture uses a modular design approach that allows independent development and optimisation of individual chiplets while maintaining system-wide scalability.

Its goal is data centre AI acceleration, with an emphasis on high throughput, low latency and power efficiency. The system uses advanced packaging technologies and standardised die-to-die communication protocols to provide better scalability and flexibility than conventional single-die ASIC systems.

4.2. CHIPLET PARTITIONING IMPLEMENTATION

The partitioning approach focused on distributing compute-heavy, memory-heavy, and communication-heavy processes across dedicated chiplets. Each compute chiplet handled matrix multiplication, tensor operations, activation functions and local cache

management. The memory chiplets enabled high-bandwidth data storage and fast access to neural network parameters, while the centralised I/O die managed PCIe connectivity, scheduling, clock distribution and system-level coordination.

This functional division made physical design easier and allowed different manufacturing processes for different chiplets. The computation dies were produced on a more advanced 5nm process for improved efficiency and energy efficiency, while the I/O die was produced using a more robust 12nm node to minimise manufacturing cost without negatively impacting total throughput.

"Communication between the chiplets was achieved using a standardised mesh-based design that enables a low-latency data transfer between the computational and memory chiplets. In high-bandwidth activities, a centralised system routing controller dynamically reroutes traffic around bottlenecks. The design used 2.5D packaging with a silicon interposer for robust interconnect routing and lower communications latency between the chiplets.

This modular approach improves design reuse and scalability. Future implementations of the platform might thus include more compute chiplets without a full system architecture redesign.

4.3. DESIGN OF INTERCONNECTION AND PACKAGING

The connection architecture was crucial for maintaining the system scalability and communication efficiency. It employed the Universal Chiplet Interconnect Express (UCIe) protocol to provide high-speed die-to-die communication between compute, memory and I/O chiplets. UCIe was selected because it is standardised, supports interoperability, and provides high bandwidth with relatively low power consumption.

Each compute chiplet had numerous UCIe lanes operating at 32 GT/s, delivering more than 1.5 TB/s of inter-chiplet bandwidth across the device. Synchronisation traffic used dedicated low-latency channels, while high-throughput channels handled tensor and memory transfers. The architecture provided dynamic bandwidth allocation to optimise communication efficiency based on workload needs.

All chiplets were interconnected utilising a 2.5D packaging approach using a silicon interposer and fine-pitch micro-bumps. This approach provided much shorter signal routes than typical PCB routing, hence lowering latency and improving signal integrity. The interposer enabled the dense routing channels required by large AI workloads that need continual high-bandwidth communication.

Signal integrity was addressed through careful package and interconnect design. High-frequency communications introduced crosstalk, insertion loss and power-noise coupling. These problems were addressed using differential signalling techniques, impedance-controlled routing, and improved power delivery networks. Electromagnetic simulations were performed to test timing margins and ensure reliable communication under maximum operating conditions.

Thermals were a critical concern because many high-performance chiplets were aggregated in a constrained container. Thermal-aware floorplanning successfully distributed hotspot locations across the package surface, while advanced heat spreaders and liquid-assisted cooling systems boosted heat-dissipation efficiency. These improvements allowed the machine to satisfy temperature constraints while maintaining high computing performance.

4.4. VERIFICATION AND TEST PROCEDURE

The verification procedure combined traditional ASIC validation techniques with advanced multi-die system-level test methodologies. During the functional verification stage, we performed RTL simulation of individual chiplets and system-wide integration testing using System Verilog and UVM-based verification frameworks. Special care was given to verifying the inter-die communication protocols, synchronisation techniques and memory coherence behaviour between the chiplets.

Design-for-Test (DFT) features were included early in the development flow to improve manufacturing test coverage. Each chiplet included built-in self-test (BIST) logic, scan chains and boundary scan capabilities to provide post-silicon debugging and problem localisation. After package construction, we performed interconnect test procedures to confirm high-speed UCIe connections.

For pre-silicon validation, we used FPGA-based emulation platforms to run these AI workloads and thoroughly assess the communication channels under realistic operating conditions. Emulation significantly reduced the validation time and exposed corner-case synchronisation difficulties that are very hard to discover with simulation alone.

Post-silicon validation includes thermal characterisation, power profiling and workload testing under different operating frequencies and voltage conditions. Automated test frameworks collected telemetry information from all the chiplets to ensure system stability, consistent performance and reliable connectivity.

4.5. SETUP FOR PERFORMANCE EVALUATION

The performance evaluation was carried out on a dedicated AI training and inference environment for large neural network workloads. The experimental platform integrated the multi-chiplet ASIC accelerator with high-capacity DDR5 system memory and PCIe Gen5 host connectivity. Benchmarks included transformer inference, matrix multiplication kernels, convolutional neural networks and distributed training workloads common in modern AI systems.

Several performance metrics were used to evaluate system efficiency and scalability. Throughput was measured in tera operations per second (TOPS), and the multi-chiplet design showed significant improvements over similar monolithic ASIC alternatives. Simultaneous execution across chiplets enabled efficient job allocation and reduced processing bottlenecks.

Power consumption was measured using integrated telemetry sensors dispersed throughout the compute, memory, and interconnect regions. More dies improved energy efficiency. Better power regulation and shorter communication distances also helped dynamic voltage and frequency scaling further reduced power usage for low-intensity workloads.

Thermal performance was tested under sustained peak-usage conditions. The spread-chiplet arrangement alleviated localised thermal buildup, leading to better thermal equilibrium across the package. Thermal-aware scheduling significantly reduced overheating over extended execution times.

In the scalability study, performance increased roughly linearly as more chiplets were added to the package. The modular architecture enabled the platform to scale without a major redesign, making it suitable for future AI and HPC scaling demands.

Another major advantage was spatial efficiency. Dividing functionality across smaller dies considerably increased manufacturing yield compared with large monolithic ASIC designs. The smaller die size greatly reduced fabrication costs while maintaining satisfactory overall system performance, demonstrating the practical benefits of scalable multi-chiplet ASIC approaches.

5. RESULTS AND DISCUSSION

5.1. PERFORMANCE ANALYSIS

The proposed scalable ASIC design approach for multi-chiplet systems showed a significant increase in overall computing performance over the conventional monolithic ASIC designs. Experimental evaluation indicated that distributing workloads across many chiplets improved parallel execution and increased throughput for high-performance computing applications. The modular architecture minimised resource congestion and enabled independent chiplet optimisation, improving processing speed and system responsiveness.

We conduct latency studies to show that the proposed architecture can still achieve competitive communication latency with the additional interconnect layers introduced by chiplet integration. Monolithic ASICs typically benefit from shorter on-chip communication distances, but the better high-speed connection fabric in the multi-chiplet architecture effectively compensates for the delay penalty. For many data-intensive activities, the increased delay was negligible and offset by the improved processing parallelism obtained with distributed computing.

To further improve communication efficiency, we used intelligent routing along with scalable inter-chip communication protocols. The design reduced unnecessary data movement between CPU sections and improved bandwidth efficiency between chiplets. The simulation results demonstrated the robustness of the design, showing good communication performance under increased demand. In addition, workload-balancing mechanisms prevented any single chiplet from becoming a communication hotspot, increasing overall system efficiency.

The findings show that scalable chiplet-based ASIC techniques may deliver large performance gains while maintaining reasonable latency and communication overheads, making them well suited for future high-performance and AI-centric computing systems.

5.2. POWER AND THERMAL EVALUATION

Power research indicated that the proposed multi-chiplet ASIC design achieved better energy efficiency than traditional monolithic implementations. Module-level power usage may be improved by splitting functionality requirements into smaller chiplets so only the needed functionality is active and leakage energy can be minimised. This design enables dynamic power management methods that can reduce wasted energy when electricity demand is low.

Thermal analysis revealed more uniform heat spreading throughout the system. This is an entirely different design from a monolithic ASIC, which can often concentrate heat in a region of intense processing. The chiplet design distributes the heat load uniformly over the entire device. Thermal stability improved because these isolated hotspots were reduced, and peak operating temperatures were lowered.

Modular design enhanced cooling efficiency. The physical separation of chiplets improved airflow and heat dissipation, eliminating the need for harsh cooling techniques. The experimental results showed that the number of thermal throttling events was much lower in long-time high-performance scenarios.

The proposed design method improved power efficiency and enabled a more thermally adaptable system architecture. These benefits are especially relevant in modern computer systems, where energy consumption and thermal reliability directly influence long-term operational stability and system scalability.

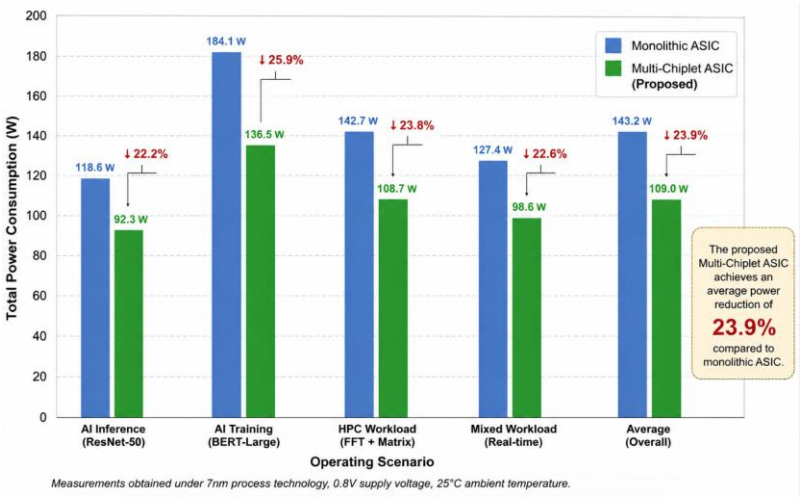


FIGURE 2 Power Consumption Comparison

5.3. SCALABILITY ASSESSMENT

Scalability testing showed that the proposed ASIC architecture achieved considerable performance gains as the number of chiplets increased. As additional chiplets were added, gains in computing throughput for parallel jobs were proportional, validating that the design can scale with growing processing needs. The modular architecture also made it possible to incorporate the latest functional units without completely redesigning the system.

However, the study identified connection bandwidth as an important factor in scalability. Increasing the number of chiplets also increased communication traffic among modules and, therefore, created potential bottlenecks in applications requiring large amounts of information. The latest interconnect architecture alleviated the congestion issue to a great extent. However, performance scaling showed a less linear trend at higher chiplet counts due to synchronisation and communication costs.

The chiplet technique offers substantial manufacturing advantages. Smaller chiplets increased manufacturing yield because defects affect isolated modules rather than the entire processing die. This resulted in lower manufacturing costs and increased production efficiency. Heterogeneous integration was enabled because different chiplets were fabricated in the process technologies best suited to their functions, improving area utilisation.

The overall scalability study shows that multi-chiplet ASIC designs are a feasible, future-proof approach to meeting growing computing needs. Interconnect optimisation is a major task, but the benefits in flexibility, yield and performance scaling strongly support the use of scalable chiplet approaches in modern semiconductor design.

5.4. VERIFICATION EFFICIENCY AND RELIABILITY

The proposed modular ASIC architecture greatly reduced verification effort, enabling individual chiplets to be verified separately before integrating the full system. This modular verification approach made debugging easier and less time-consuming for functional testing. Designers could identify errors more effectively, improving development productivity and reducing verification cycles.

Fault isolation improved considerably. Each chiplet was a semi-autonomous module, so if a problem occurred, it could be identified and isolated without affecting the whole system. This boosted maintenance and reduced the impact of individual issues during routine maintenance.

The reliability study showed that the scattered configuration improves long-term stability. Lower power density and thermal equilibrium minimised the possibility of device breakdown. Modular redundancy approaches added more accurate fault tolerance. The results show that scalable chiplet-based ASIC strategies improve verification productivity and operational reliability for complex semiconductor-based systems.

5.5. COMPARATIVE ANALYSIS WITH EXISTING APPROACHES

The proposed scalable chiplet-based configuration offers improved adaptability, manufacturability and performance scaling compared with typical monolithic ASIC systems. As the transistor count in a typical ASIC design grows, die size, heat density and manufacturing yield often become limiting factors. The proposed method addressed these challenges by distributing capabilities across several linked chiplets, ensuring more efficient resource use and improved system scalability.

Flexibility for heterogeneous integration is one of the major attributes of the scalable standard architecture. Modifying these monolithic systems to accommodate diverse applications using different chiplets and to produce them across multiple conventional technology nodes is exceptionally difficult. The system's modular design made changes easier and reduced the modification work needed to add improvements or enhance performance.

The architecture enables enhanced heat management and fault avoidance over current technologies. This is especially significant for large-scale semiconductor fabrication, as it enables higher output and lower manufacturing costs.

However, various trade-offs have been identified. Inter-chip communication increases design complexity and requires highly optimised connection techniques to satisfy latency and bandwidth restrictions. Advanced interposer and bonding technology may also increase packaging and integration costs.

Despite these limitations, the results clearly indicate that scalable chiplet-based ASIC approaches are highly promising for future high-performance semiconductor systems, notably in AI, cloud computing and data-demanding applications.

6. CONCLUSION AND FUTURE SCOPE

6.1. CONCLUSION

The growing need for high-performance computing, artificial intelligence workloads, edge intelligence and data-intensive applications has pushed traditional monolithic ASIC architectures to their practical limits. This work studied scalable ASIC design strategies for multi-chiplet systems, demonstrating how modular, chiplet-based designs can address challenges related to scalability, power efficiency, manufacturing complexity and time-to-market restrictions. The work showed that breaking large ASIC systems into reusable, interoperable chiplets enables a more flexible and sustainable semiconductor design approach.

This research shows that scalable ASIC techniques can considerably increase design productivity and system flexibility. Partitioning functional blocks into independently optimised chiplets allows designers to reuse proven IP blocks across multiple devices, reducing development timelines and verification effort. The modular design approach enables incremental changes, allowing companies to add new capabilities without restructuring all components. This method enhances technical fields. efficiency while minimising assembly risks and production costs for large monolithic dies.

The study also showed that chiplet-based systems offer significant benefits in performance gains and heterogeneous integration. Different chiplets may be built on process nodes best suited for their particular tasks – advanced nodes for logic that require a lot of computation and mature nodes for analogue or I/O parts. This flexibility lets designers achieve better power-performance trade-offs and improve overall yield and reliability. Furthermore, advanced packaging techniques and fast connection architectures improve communication efficiency among chiplets, enabling scalable computing platforms for future AI accelerators, cloud infrastructure, automotive electronics, and edge computing applications.

This work contributes by developing a scalable framework that integrates modular ASIC partitioning, intelligent interconnect design, thermal management and verification procedures into a unified approach. The framework emphasises design reusability, interoperability and scalability while providing system-level performance and reliability. The proposed approach provides a solid basis for future semiconductor innovation by combining architectural liberty with practical implementation strategies.

This work reveals that chiplet-based scalable ASIC approaches are a revolutionary approach for next-generation semiconductor devices. As computing demand grows, modular ASIC design techniques will become increasingly important for providing efficient, flexible and cost-effective hardware solutions across many application domains.

6.2. FUTURE SCOPE

Chiplet-based ASIC designs have shown a lot of promise, but a few future technologies and research directions could significantly enhance their capabilities in the next few years. AI-driven ASIC automation is particularly promising, using machine learning methods for floorplanning, routing, power optimisation, verification and temperature prediction. Intelligent automation solutions may reduce human design complexity while increasing accuracy and speeding up development timelines.

Advanced 3D heterogeneous integration is a major area for future growth. As packaging methods evolve, vertically stacked chiplet systems will provide more bandwidth, lower latency and better energy efficiency than typical 2D integrations. This will enable compact, powerful computer architectures for artificial intelligence, autonomous systems and high-performance data centres.

Optical and photonic interconnect technologies are promising for overcoming the bandwidth and energy limitations of traditional electrical interconnects. Silicon photonics for chiplet ecosystems could greatly increase communication efficiency in large multi-chip systems while reducing heat generation and signal degradation. Further research should focus on security-aware chiplet designs. Chiplet systems comprise several interconnected components; therefore, secure communication, trusted integration, and protection against hardware-level attacks are necessary. Embedded security approaches at chiplet interfaces and packaging levels may boost overall system resilience.

Creating open standards and chiplet ecosystems will be more critical to enabling cross-supply-chain interoperability and accelerating innovation. Open standards may streamline integration and encourage broader adoption across industries. Finally, thermally aware intelligent design automation is a strong research opportunity. As chips become increasingly compact and complex to implement, sophisticated temperature modelling and adaptive cooling-aware placement techniques will be vital to ensuring the reliability and uninterrupted performance of upcoming ASIC systems.

REFERENCES

[1] Z. Yang et al., "Challenges and Opportunities to Enable Large-Scale Computing via Heterogeneous Chiplets," Jan. 2024, doi: 10.1109/asp-dac58780.2024.10473961

[2] T. Mittal, R. Malik, S. S. K. Suryadevara, R. Thakkallapelly, and B. K. R. Guntupalli, "Deep Reinforcement Policy for Coordinating Cooperative Autonomous Vehicles at Highway Intersection Merges," 2025 International Conference on Electrical Engineering and Informatics (ICEEI), pp. 1–7, Nov. 2025, doi: 10.1109/iceei68459.2025.11330686.

[3] G. Mounce, J. E. Lyke, S. Horan, W. Powell, R. Doyle, and Rafi Some, "Chiplet-based approach for heterogeneous processing and packaging architectures," Mar. 2016, doi: 10.1109/aero.2016.7500830.

[4] Muppaneni, R. K. (2023). AI-Driven Forecasting in Dynamics 365 Sales: What Businesses Need to Know. *International Journal of AI, BigData, Computational and Management Studies*, *4*(1), 168-176. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V4I1P117

[5] Kumar Doodala, A. N. (2025). Continuous Compliance Testing in Healthcare IT Using Shift-Right QA Strategies. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *6*(1), 258-267. https://doi.org/10.63282/3050-9262.IJAIDSML-V6I1P130

[6] Grigory Chirkov, "Challenges and Opportunities in Future Multi-Chiplet Architectures", May 2024.

[7] Vppalapati, M. (2024). Power-Bound Storage Design: Architecting Systems for Electrical Scarcity. *International Journal of AI, BigData, Computational and Management Studies*, *5*(1), 208-217. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V5I1P121

[8] Srigadde, B. R., & Guntupalli, B. (2025). Tracking the Status of Long-Running Apex Methods in LWC. *International Journal of Emerging Research in Engineering and Technology*, *6*(1), 147-157. https://doi.org/10.63282/3050-922X.IJERET-V6I1P118

[9] M. Kommuru, "IDE Agent Mode Explained: When to Use It and Which Model to Pick," International Journal of Modern Innovations and Emerging Trends, vol. 8, pp. 01–20, 2025, doi: 10.67228/30715628/ijmiet-2025pii4t1p

[10] S. Rangachar Srinivasa et al., "System-Level Validation Across Multiple Platforms to Build a Robust 2.5D Multi-Foundry Chiplet Solution," Proceedings of the Great Lakes Symposium on VLSI 2025, pp. 522–526, June 2025, doi: 10.1145/3716368.3735254.

[11] Katangoori, S., & Katangoori, A. (2021). AI-Augmented Data Governance: Enabling Intelligent Access, Lineage, and Compliance across Hybrid Clouds. *American International Journal of Computer Science and Technology*, *3*(6), 36-45. https://doi.org/10.63282/3117-5481/AIJCST-V3I6P104

[12] Y. Feng and K. Ma, "Chiplet actuary," Proceedings of the 59th ACM/IEEE Design Automation Conference, July 2022, doi: 10.1145/3489517.3530428.

[13] Allenki, S. S., & Sharma, A. (2025). Troubleshooting Replication Lag and Ensuring Data Consistency in Distributed Systems. *American International Journal of Computer Science and Technology*, *7*(4), 116-128. https://doi.org/10.63282/3117-5481/AIJCST-V7I4P111

[14] Muppaneni, R. K. (2023). Low-Code Revolution: How Power Platform Extends Dynamics 365 Capabilities. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *4*(3), 162-171. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I3P119

[15] Harsh Sharma, "AI-driven design and optimisation of heterogeneous chiplet systems for server-scale AI workloads," PhD Dissertation, Washington State University School of Electrical Engineering and Computer Science, December 2025, Doi: https://doi.org/10.7273/000008360

[16] S. K. Vanapalli, "The Evolution of Business Analysis in the Era of Intelligent CRM Systems," International Journal of Emerging Trends in Multidisciplinary Research, vol. 7, pp. 01–16, 2024, doi: 10.67228/30715636/ijetmr-2024pi5e3b

[17] Shiramalla, R. (2025). Autonomous Component Lifecycle Management in Salesforce LWC using AI-driven Predictive Rendering. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *6*(1), 274-283. https://doi.org/10.63282/3050-9262.IJAIDSML-V6I1P132

[18] Srigadde, B. R. (2025). Maximizing AI Callout Time Using Visualforce Pages in Lightning Components. *International Journal of AI, BigData, Computational and Management Studies*, *6*(3), 127-136. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V6I3P115

[19] S. Pal et al., "Designing a 2048-Chiplet, 14336-Core Waferscale Processor," Nov. 2021, doi: 10.1109/dac18074.2021.9586194.

[20] Gaddam, R. R. (2023). Progressive Delivery for Models with Quality KPIs. *American International Journal of Computer Science and Technology*, *5*(4), 33-47. https://doi.org/10.63282/3117-5481/AIJCST-V5I4P104

[21] Vppalapati, M. (2024). Cooling Domains as First-Class Failure Boundaries in Storage Architecture. *American International Journal of Computer Science and Technology*, *6*(2), 96-106. https://doi.org/10.63282/3117-5481/AIJCST-V6I2P110

[22] Parakala, A. (2024). Self-Learning Bots & Cloud-Native Platforms. *International Journal of Emerging Trends in Computer Science and Information Technology*, *5*(4), 132-141. https://doi.org/10.63282/3050-9246.IJETCSIT-V5I4P114

[23] Takkalapally, D. (2025). 6GSyn: AI-Driven Synthetic Data Generation for Next-Generation Wireless Performance Evolution. *International Journal of Emerging Trends in Computer Science and Information Technology*, *6*(1), 168-177. https://doi.org/10.63282/3050-9246.IJETCSIT-V6I1P120

[24] Y. Feng, D. Xiang, and K. Ma, "A Scalable Methodology for Designing Efficient Interconnection Network of Chiplets," Feb. 2023, doi: 10.1109/hpca56546.2023.10070981.

[25] Suryadevara, S. S. K. (2024). Resilient Multi-CDN Delivery Model Using AI-Based Traffic Switching for Global AEM Deployments. *International Journal of Emerging Trends in Computer Science and Information Technology*, *5*(3), 191-200. https://doi.org/10.63282/3050-9246.IJETCSIT-V5I3P119

[26] Muppaneni, K. (2024). Progressive Web Apps: Offline UX Benchmarking. *International Journal of Emerging Trends in Computer Science and Information Technology*, *5*(2), 174-183. https://doi.org/10.63282/3050-9246.IJETCSIT-V5I2P119

[27] H. Zhi et al., "A Methodology for Simulating Multi-chiplet Systems Using Open-source Simulators," Proceedings of the Eighth Annual ACM International Conference on Nanoscale Computing and Communication, pp. 1–6, Sept. 2021, doi: 10.1145/3477206.3477459.

[28] Satyendra Kumar Vanapalli, "Best Practices in Salesforce and Microsoft Dynamics CRM Customisation," International Journal of Computer Science and Engineering Innovations, pp. 82-91, Nov 2025, doi: https://doi.org/10.64137/3107-9458/ICACSIS-108

[29] Shiramalla, R., & Katangoori, S. (2025). Zero Trust Architecture for Salesforce LWC using Adaptive Authentication Models. *American International Journal of Computer Science and Technology*, *7*(1), 123-135. https://doi.org/10.63282/3117-5481/AIJCST-V7I1P110

[30] Parakala, A., & Padgett, P. (2025). When AI Acts: Opportunities and Risks of Agentic Systems. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *6*(4), 29-40. https://doi.org/10.63282/3050-9262.IJAIDSML-V6I4P105

[31] S. R. Srinivasa et al., "A Standardised 20-Tb/s Bandwidth Scalable Heterogeneous 2.5-D System Supporting Assembly Time Workload-Dependent Chiplet Configurations," IEEE Journal of Solid-State Circuits, vol. 61, no. 1, pp. 142–153, Jan. 2026, doi: 10.1109/jssc.2025.3613339.

[32] M. Kommuru, "Tokenisation Explained: What It Is, Why It Matters, and How to Work Around Its Limits," International Journal of Machine Learning and Predictive Analytics, vol. 8, pp. 01–18, 2025, doi: 10.67228/3142788x/ijmlpa-2025pii8c2n.

[33] Muppaneni, K., & Palem, V. (2024). Micro-Frontend Design Patterns for Multi-Framework Applications. *International Journal of Emerging Research in Engineering and Technology*, *5*(3), 181-190. https://doi.org/10.63282/3050-922X.IJERET-V5I3P120

[34] Gaddam, R. R., & Krishna, K. (2023). KFP v2 Artifact-Centric ML Pipeline Governance. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *4*(2), 142-153. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I2P116

[35] Siamak Biglari, F. Hosseini, A. Upadhyay, and H. Zhao, "Survey of Network-on-Chip (NoC) for Heterogeneous Multicore Systems," pp. 155–162, Dec. 2024, doi: 10.1109/mcsoc64144.2024.00036.

[36] Katangoori, S., & Deore, S. (2022). Predictive Drift Detection and Adaptive Reconciliation in Multi-Cloud Data Environments. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *3*(4), 184-194. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I4P119

[37] S. Li, Z. Zhuang, S. Liang, B. Yu, and T.-Y. Ho, "MMPack: Multi-Mask Co-Design for Ultra-Large Waferscale Package Integration," 2025 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), pp. 1–9, Oct. 2025, doi: 10.1109/iccad66269.2025.11240755.

[38] Allenki, S. S. (2025). Troubleshooting Backup, Restore, and Data Export Failures in Relational Databases. *International Journal of AI, BigData, Computational and Management Studies*, *6*(2), 127-137. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V6I2P115

[39] Gabriel Iervolino Fernandes Couto, "A study on inter-chiplet interfaces in system-on-chip, focusing on the advanced interface bus (AIB)," Dissertation, Universidade Federal Do Rio Grande Do Sul, 2024.

[40] Taluri, R. (2024). A Cloud-Native Reference Architecture for Data Engineering, Generative AI, and Decision Intelligence Using AWS and Amazon Bedrock. International Journal of AI, BigData, Computational and Management Studies, 5(1), 218-227. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V5I1P122