MANTECH Publications

Journal of Embedded Systems & Its Applications

Volume 10 Issue 2, July-December 2025

Chiplet-Based Modular Embedded Architectures

Satish Chandra, Govind Singh, Ramu Pathak

Associate Professor, Students

Department of Electronics and Embedded Systems

niversity of Calcutta, Kolkata. India

satishchandra3h@gmail.com, govinds211@rediffmail.com, ramu005pathak@yahoo.com

Abstract

The increasing demand for high-performance embedded systems in applications such as Internet of Things (IoT), automotive electronics, and edge computing has created a need for scalable, energy-efficient, and modular hardware architectures. Traditional monolithic System-on-Chip (SoC) designs face challenges related to power consumption, yield, cost, and design complexity. Chiplet-based modular embedded architectures emerge as a promising solution, allowing designers to integrate heterogeneous computing, memory, and specialized functional blocks as smaller, pre-verified chiplets. This paper reviews the recent advancements in chiplet technology, explores integration techniques, highlights advantages, challenges, and potential applications, and provides insights into future research directions. The study also examines the role of interconnect technologies, packaging, and standardization efforts in promoting chiplet adoption for modular embedded systems.

Keywords: *Chiplets, Modular Embedded Systems, Heterogeneous Integration, 3D Packaging, Advanced Interconnects, System-on-Chip, Embedded Architectures.*

1. INTRODUCTION

Embedded systems have become an integral part of modern electronics, ranging from smartphones and wearable devices to autonomous vehicles and industrial IoT systems. The evolution of these systems demands higher computational capabilities while maintaining low power consumption, small form factors, and cost efficiency. Traditional monolithic SoC approaches, which integrate all components into a single silicon die, are increasingly constrained by fabrication limitations, high costs, and low yields for complex designs.

Chiplet-based design offers a paradigm shift, where individual functional blocks, or **chiplets**, are manufactured separately and integrated into a modular package. This approach improves design flexibility, reduces development costs, and accelerates time-to-market. Each chiplet can be optimized for performance, power, or specific technology nodes, enabling heterogeneous integration of processors, memory, accelerators, and specialized sensors.

2. EVOLUTION OF EMBEDDED ARCHITECTURES

Embedded system architectures have evolved significantly over the last few decades to meet the increasing demands of computational performance, energy efficiency, and integration of heterogeneous functionalities. This evolution can be broadly categorized into **monolithic SoC designs** and **modular approaches**, with a clear shift toward modularity in modern designs.

2.1 Monolithic SoC Designs

Monolithic System-on-Chip (SoC) architectures integrate all core components—processors, memory, peripherals, and interfaces—onto a single silicon die. Historically, monolithic SoCs have been the dominant approach in embedded systems due to their compactness, high performance, and well-understood interconnects.

Advantages:

- **Compact design:** Integrating all components reduces board-level interconnects.
- **High-speed communication:** On-chip buses and interconnects offer low latency and high bandwidth.
- **Predictable behavior:** Single-die designs simplify timing closure and verification.

Challenges:

- **Power density and thermal management:** As transistor scaling progresses (e.g., <7 nm, <5 nm nodes), more transistors per unit area lead to increased power density, causing heat dissipation issues.
- **Manufacturing complexity and yield:** Large monolithic dies are more prone to defects, reducing fabrication yield and increasing cost per chip.
- **Heterogeneous integration limits:** Combining analog, RF, and digital logic on the same die is challenging due to process incompatibilities.

- **Slower design cycles:** Each new iteration requires redesign and verification of the full chip, increasing time-to-market.

Example Applications:

- Early smartphones, embedded controllers in automotive electronics, and IoT devices relied heavily on monolithic SoCs for integration simplicity.

2.2 Modular Approaches

To overcome the limitations of monolithic SoCs, the industry has shifted toward **modular architectures**, where systems are divided into functional blocks or "modules" that can be developed, tested, and reused independently. This approach enables **chiplet-based designs**, where pre-verified silicon blocks are interconnected on a package or substrate rather than a single die.

Key Features:

- **Reusability:** Standardized chiplets allow designers to reuse functional blocks (e.g., processor cores, memory, accelerators) across multiple products.
- **Heterogeneous integration:** Enables combining logic, analog, memory, RF, and specialized accelerators (e.g., AI, DSP) on a single platform.
- **Scalability:** Systems can scale performance by adding or upgrading chiplets rather than redesigning the entire SoC.
- **Cost efficiency:** Smaller dies are easier to manufacture with higher yield, reducing overall system cost.
- **Shorter development cycles:** Verified modules reduce the need for extensive full-chip verification.

Emerging Trends:

- **2.5D/3D packaging:** Interconnects like silicon interposers or advanced packaging allow high-bandwidth communication between chiplets.
- **Heterogeneous chiplets:** Integration of components from different process nodes (e.g., 3 nm logic + 14 nm memory) optimizes performance and cost.
- **Standardization initiatives:** Efforts like **Open Domain-Specific Architecture (ODSA)** and **Universal Chiplet Interconnect Express (UCIe)** aim to create interoperable chiplet ecosystems.

- **AI/Edge computing:** Modular embedded architectures are increasingly used in high-performance edge devices where compute and energy efficiency must be balanced.

Example Applications:

- AMD's EPYC and Ryzen processors (chiplet-based CPU design)
- NVIDIA's modular GPUs
- Automotive ADAS systems combining processors, memory, and specialized AI accelerators

3. CHIPLET CONCEPT AND ARCHITECTURE

The chiplet paradigm represents a shift from traditional monolithic System-on-Chip (SoC) designs toward modular, heterogeneous integration. By decomposing a system into smaller, specialized dies (chiplets), designers can achieve higher flexibility, better yields, and easier technology scaling.

3.1 Definition of Chiplets

Chiplets are individual semiconductor dies designed to perform specific functions within a larger system. Typically, their sizes range from a few square millimeters to tens of square millimeters. Unlike conventional SoCs, which integrate all components onto a single monolithic die, chiplets allow designers to mix-and-match functionalities optimized for performance, power, or cost.

Common types of chiplets include:

1. **Digital Processing Units (DPUs):**
   - Central Processing Units (CPUs) for general-purpose computation.
   - Graphics Processing Units (GPUs) for parallelized workloads, including graphics rendering and high-performance computing.
2. **Memory Blocks:**
   - DRAM or SRAM modules positioned near processors for minimal latency.
   - High-Bandwidth Memory (HBM) stacks integrated as chiplets to improve memory throughput.
3. **Specialized Accelerators:**
   - AI/ML accelerators for neural network inference and training.

- Cryptographic engines for secure computation.
- Digital Signal Processing (DSP) units for audio, video, or sensor data processing.

4. **I/O Interfaces:**
- Chiplets handling standard interface protocols such as USB, PCIe, Ethernet, or custom high-speed links.

Advantages of Chiplets over Monolithic SoCs:

- **Heterogeneous Integration:** Different process nodes can be used for different chiplets (e.g., advanced nodes for logic, mature nodes for memory).
- **Improved Yield:** Smaller dies have higher manufacturing yields; faulty chiplets can be replaced individually.
- **Faster Time-to-Market:** Reusable chiplets reduce design cycles.
- **Scalability:** Systems can be upgraded by swapping or adding chiplets.

3.2 Chiplet-Based Modular Embedded Architecture

A **chiplet-based architecture** organizes multiple specialized chiplets into a single system using advanced packaging and interconnect technologies. This approach enables high-performance, low-power, and highly configurable embedded systems.

Typical Architecture Components:

1. **Processing Chiplets:**
- Serve as the main computational engines.
- Can be heterogeneous (CPU + GPU + AI accelerator) for workload-specific optimization.
2. **Memory Chiplets:**
- High-bandwidth memory (HBM) or DRAM is placed close to the processing chiplets.
- Reduces latency and improves memory bandwidth compared to traditional off-chip memory.
3. **Accelerator Chiplets:**
- Include specialized units such as AI, cryptography, or DSP engines.
- Offload compute-intensive tasks from general-purpose processors, improving efficiency.
4. **Interconnects and Packaging:**
- **Silicon Interposers:** A high-density interposer layer allows fine-pitch connections between chiplets in a 2.5D integration.

- **Advanced Substrates:** Organic or fan-out substrates can be used for cost-effective multi-chip integration.
- **3D Integration:** Stacking of chiplets vertically with through-silicon vias (TSVs) for higher integration density.
- **Inter-chip Communication:** High-speed, low-latency protocols (e.g., CXL, UCIe) enable efficient data transfer between chiplets.

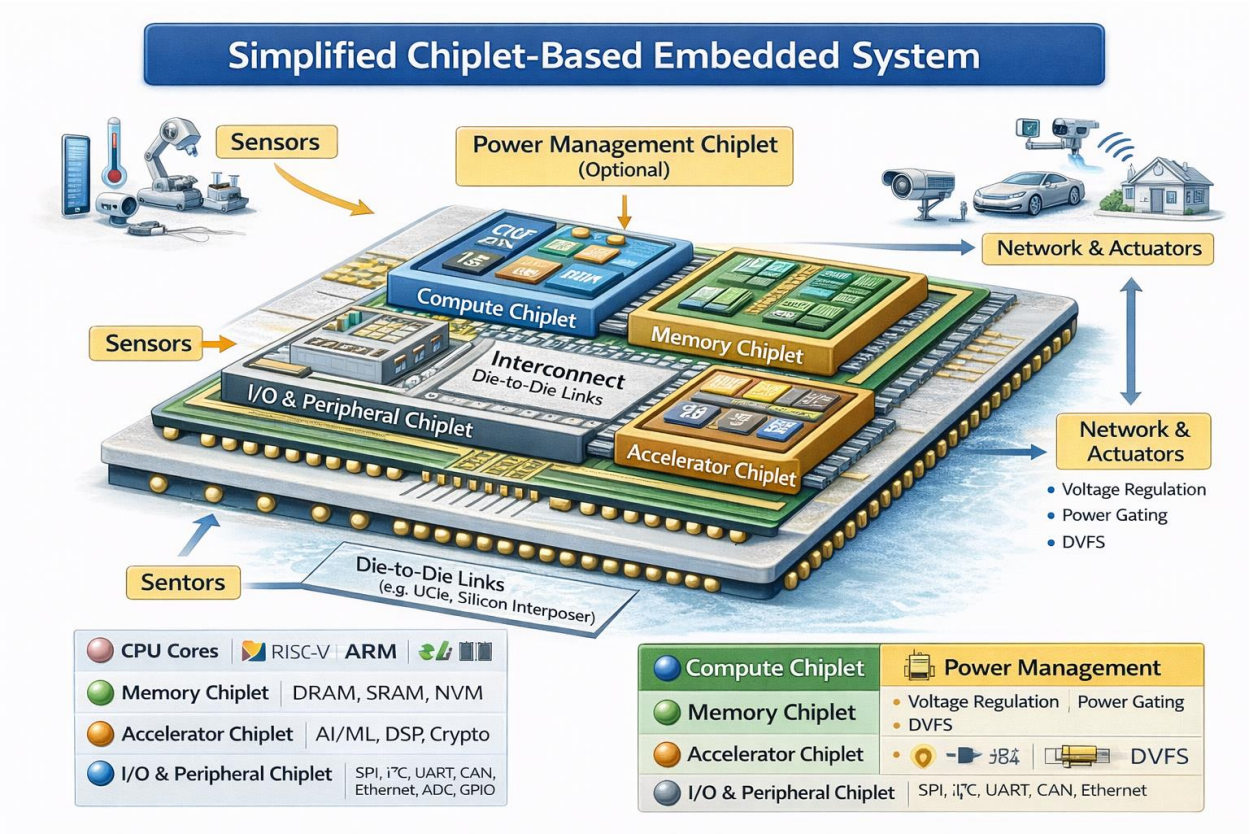

Figure 1. Simplified Chiplet-Based Embedded System

4. INTERCONNECT TECHNOLOGIES

In chiplet-based architectures, the performance of the overall system heavily depends on the efficiency of the **interconnect technology** that links individual chiplets. These interconnects must support **high bandwidth**, **low latency**, and **power-efficient communication** while accommodating heterogeneous dies that may be manufactured on different process nodes. Efficient interconnects are essential to ensure that the modularity of chiplets does not compromise system performance. Several key technologies have emerged to address these challenges.

4.1 Silicon Interposer

A **silicon interposer** is a passive, thin silicon layer that sits beneath multiple chiplets and provides a high-density routing platform.

Key Features:

- Allows **fine-pitch connections**, often in the range of a few micrometers, enabling high-bandwidth links between chiplets.
- Reduces signal loss and latency compared to traditional organic substrates.
- Supports **2.5D integration**, where multiple chiplets are placed side-by-side on the interposer.

Applications:

- **AMD EPYC and Ryzen processors** use silicon interposers to connect CPU and high-bandwidth memory (HBM).
- **Intel Foveros** leverages interposers for multi-die stacking and heterogeneous integration.

Advantages:

- Enables **heterogeneous integration** across different process nodes.
- Supports **high pin density**, which is critical for wide I/O buses.
- Provides **mechanical stability** for complex multi-chip systems.

Limitations:

- Costly due to the complexity of silicon processing.
- Interposer area can limit scaling when the system requires a very large number of chiplets.

4.2 Embedded Multi-Die Interconnect Bridge (EMIB)

EMIB is Intel's solution for high-density interconnects **without using a full silicon interposer**.

Key Features:

- Uses small silicon bridges embedded in an organic substrate.
- Localized high-speed interconnects connect neighboring chiplets.
- Supports high bandwidth at lower cost compared to a full interposer.

Applications:

- Used in Intel **FPGAs** and hybrid CPU-GPU packages.
- Allows heterogeneous dies to communicate efficiently while reducing substrate area.

Advantages:

- Reduces cost compared to full silicon interposers.
- Enables **localized, high-performance interconnects** where needed.
- Compatible with existing packaging flows.

Limitations:

- Bandwidth is localized; may not scale as efficiently as a full interposer for very large multi-chip systems.

4.3 Advanced Packaging (2.5D and 3D Integration)

Advanced packaging techniques improve integration density and performance by physically arranging chiplets in new ways.

2.5D Integration:

- Chiplets are placed **side-by-side on an interposer**, creating a planar assembly.
- Provides high-bandwidth connections while maintaining simpler thermal profiles than 3D stacking.

3D Integration:

- Chiplets are **stacked vertically**, often using through-silicon vias (TSVs) for interconnects.
- Offers **maximum bandwidth and compact footprint**, suitable for high-performance computing and AI accelerators.

Challenges in 3D Integration:

- **Thermal management**: Stacked dies generate more heat in a compact area, requiring advanced cooling solutions.
- **Manufacturing complexity**: Precise alignment and yield management are critical.
- **Design considerations**: Power delivery and signal integrity must be carefully engineered.

Advantages:

- Significantly reduces inter-chip distance, lowering latency.
- Improves integration of memory and compute units (e.g., HBM on top of GPUs).

4.4 Standardized Interfaces

Standardized interconnects are essential for **chiplet interoperability**, especially when combining dies from multiple vendors.

Examples:

- **Advanced Interface Bus (AIB):** Open standard for high-speed chiplet interconnect.
- **OpenHBI:** Standard for memory interfaces between chiplets and memory stacks.
- **Universal Chiplet Interconnect Express (UCIe):** Emerging standard supporting heterogeneous chiplet integration across multiple vendors and process nodes.

Benefits of Standardization:

- Promotes **vendor-agnostic modular design**, enabling a true ecosystem of reusable chiplets.
- Reduces **integration time and cost** by providing common communication protocols.
- Ensures **scalability and compatibility** in large heterogeneous systems.

Table 1. Comparison of Interconnect Technologies

Interconnect Type	Bandwidth	Cost	Integration Density	Thermal Challenges
Silicon Interposer	High	High	High	Moderate
EMIB	High	Medium	Medium	Low
2.5D	Medium	Medium	Medium	Moderate
3D Stacking	Very High	High	Very High	High

5. ADVANTAGES OF CHIPLET-BASED MODULAR ARCHITECTURES

1. **Modularity and Reusability:** Chiplets can be reused across designs, reducing development time.
2. **Heterogeneous Integration:** Different technology nodes (e.g., advanced logic, mature memory) can coexist in the same package.
3. **Cost Efficiency:** Smaller dies increase yields, reducing manufacturing costs.
4. **Scalability:** System performance can be scaled by adding chiplets.
5. **Faster Time-to-Market:** Pre-verified chiplets accelerate system development.

6. **Power Efficiency:** High-bandwidth, low-latency interconnects reduce energy consumption per operation.

6. CHALLENGES IN CHIPLET-BASED EMBEDDED SYSTEMS

1. **Standardization:** Lack of universally accepted standards complicates interoperability.
2. **Thermal Management:** Closely packed chiplets generate localized hotspots.
3. **Design Complexity:** Partitioning a system into chiplets requires careful architectural planning.
4. **Signal Integrity:** High-speed interconnects need robust signaling to prevent data loss.
5. **Supply Chain Management:** Procuring heterogeneous chiplets from multiple vendors can introduce logistical issues.

7. APPLICATIONS OF CHIPLET-BASED MODULAR EMBEDDED ARCHITECTURES

7.1 Edge AI Devices

- Low-latency AI inference at edge devices.
- Chiplets for CPU, GPU/AI accelerators, and memory integrated for compact design.

7.2 Automotive Electronics

- Autonomous driving systems require high reliability, low power, and heterogeneous processing.
- Chiplets enable easy upgrade of sensors or AI modules.

7.3 IoT and Wearables

- Small form factor and low-power operation benefit from modular chiplet designs.

7.4 Data Centers and High-Performance Computing (HPC)

- Chiplets allow scalable, high-performance processors without fabricating massive monolithic dies.

8. FUTURE DIRECTIONS

1. **Standardization Efforts:** Promoting open standards for chiplet interfaces (e.g., UCIe, AIB).
2. **Advanced Packaging Technologies:** Improvements in 3D integration, thermal management, and interposer design.

3. **AI-Driven Design Tools:** Automated partitioning and integration of chiplets to optimize performance and power.
4. **Heterogeneous Integration:** Integration of photonics, RF, and analog chiplets for next-generation embedded systems.
5. **Security Features:** Incorporation of hardware-level security in modular designs.

9. CONCLUSION

Chiplet-based modular embedded architectures represent a transformative approach in embedded system design. They overcome the limitations of monolithic SoCs by enabling modularity, heterogeneous integration, and scalability. With advances in interconnects, packaging, and standardization, chiplets can address the growing demands of AI, IoT, automotive, and high-performance computing systems. While challenges related to thermal management, standardization, and design complexity remain, ongoing research and industrial adoption indicate a strong potential for chiplet-based embedded systems to become mainstream in the near future.

REFERENCES

1. Balasubramanian, K., & Nandakumar, R. (2022). *Chiplet-based System-on-Chip Design: Opportunities and Challenges*. IEEE Transactions on VLSI Systems, 30(4), 512–526.
2. Chen, Y., et al. (2021). *Heterogeneous Integration of Chiplets for High-Performance Embedded Systems*. Journal of Microelectronics, 59(7), 245–258.
3. Intel Corporation. (2020). *EMIB Technology for 2.5D Packaging*. Intel White Paper.
4. AMD Inc. (2022). *Chiplet Architecture in EPYC Processors*. AMD Technical Brief.
5. UCIe Consortium. (2023). *Universal Chiplet Interconnect Express (UCIe) Specification Overview*.
6. Kim, H., & Lee, S. (2021). *Thermal Challenges in 3D Chiplet Stacking*. Microelectronics Journal, 114, 105–120.
7. Chen, P., & Singh, A. (2020). *Modular Embedded Architectures for IoT and Edge Computing*. Embedded Systems Review, 15(2), 75–88.
8. OpenHBI Consortium. (2022). *Open High-Bandwidth Interface for Chiplets*. Technical Specification.

9. Gupta, V., & Ramesh, S. (2021). *Design Tools for Chiplet-Based Systems: A Review*. International Journal of Embedded Systems, 12(3), 134–150.
10. Zhang, Y., et al. (2022). *Heterogeneous Chiplet Integration for AI Accelerators*. ACM Transactions on Embedded Computing Systems, 21(4), 1–22.