

AMD CHIPLLET ECOSYSTEM

WHITEPAPER | DECEMBER 09, 2024

This white paper provides an overview of AMD's implementation of chiplet technology. This document also provides a high-level overview of chiplet technology, its advantages, challenges, and solutions for chiplet integration. It also discusses the chiplet ecosystem, the role of UCle, and the interim solutions and scenarios for chiplet integration.

TABLE OF CONTENTS

What Are Chiplets?	3
AMD's Chiplet Experience.....	4
Building Chiplet-Based SoCs	4
Chiplet Ecosystem	5
Bridging the Gap, Interim Solutions for the Chiplet Ecosystem	8
Conclusion	9

“It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected. The availability of large functions, combined with functional design and construction, should allow the manufacturer of large systems to design and construct a considerable variety of equipment both rapidly and economically.”

- Gordon E. Moore, “Cramming more components onto integrated circuits,” *Electronics*, Volume 38, Number 8, April 19, 1965

WHAT ARE CHIPLETS?

Traditional monolithic silicon design and manufacturing is being impacted by ever increasing cost and complexity. An obvious solution is to break up the monolithic silicon into smaller functions (specialized chips). The idea of breaking up monolithic silicon into smaller functions has been around for many years. As early as 1965, Gordon Moore theorized that systems would be built from specialized chips. While Moore was focused on board-level integration with separately packaged devices, his idea is easily extended into the chiplet era. Technology is readily available that allows smaller functions using specialized chiplets that can be co packaged.

The functional scope of a chiplet can vary. Chiplets can consist of a highly tuned and complex building block (for example, an AMD “Zen” CPU) or a discrete group of functions that allow for SoC differentiation. These modular chiplets provide designers and manufacturers with flexible and scalable design capabilities to meet the design requirements of modern SoCs. For over a decade, AMD has successfully addressed the challenges and limitations of traditional monolithic design by designing and manufacturing chiplet-based System on Chip (SoC) solutions.

Using chiplets offers several advantages over monolithic silicon:

- **Cost-effective:** Smaller chiplets have higher manufacturing yields, reduce waste, and lower costs. Chiplets can be produced in the most suitable silicon process for their function enabling design solutions that include chiplets produced by multiple silicon processes.
- **Scalability and Flexibility:** Chiplets are modular and can be mixed and matched for specific needs, allowing for easier upgrades and customization.
- **Faster Innovation:** Separate chiplet development allows for parallel development and faster innovation cycles.
- **Performance:** Chiplets provide scalability and flexibility that allows for specialized designs, potentially leading to better overall performance in some cases.

Using chiplets rather than monolithic SoCs has the following challenges:

- **Increased Design Complexity:** Designing the interfaces and connections between multiple chiplets requires more complex engineering compared to a monolithic chip.
- **Communication Challenges:** Data transfer between chiplets can be slower and require more power compared to on-chip communication within a monolithic design. This can impact performance for latency-sensitive tasks.
- **Packaging Costs:** Assembling and packaging multiple chiplets adds complexity and can increase overall costs compared to a simpler monolithic package.
- **Yield:** While individual chiplet yields might be higher, overall system yield can be lower due to the additional challenges of packaging and interconnection.
- **Limited Ecosystem and Standards:** Chiplet technology is still evolving. Standardized interfaces and innovative design tools are emerging.
- **Physical Limits:** Putting chiplets together results in more power per unit area. Chiplets have restrictions on their size and shape.

AMD is committed to capitalizing on the advantages of chiplets, while minimizing the disadvantages.

AMD'S CHIPLET EXPERIENCE

AMD has 10 years of innovation in chiplet architecture. In 2019, AMD's 2.5D chiplet technology was introduced with the AMD Ryzen and AMD EPYC processors. In 2023, AMD released the Instinct MI300X AI accelerator that incorporates the latest 2.5D and 3D technology.

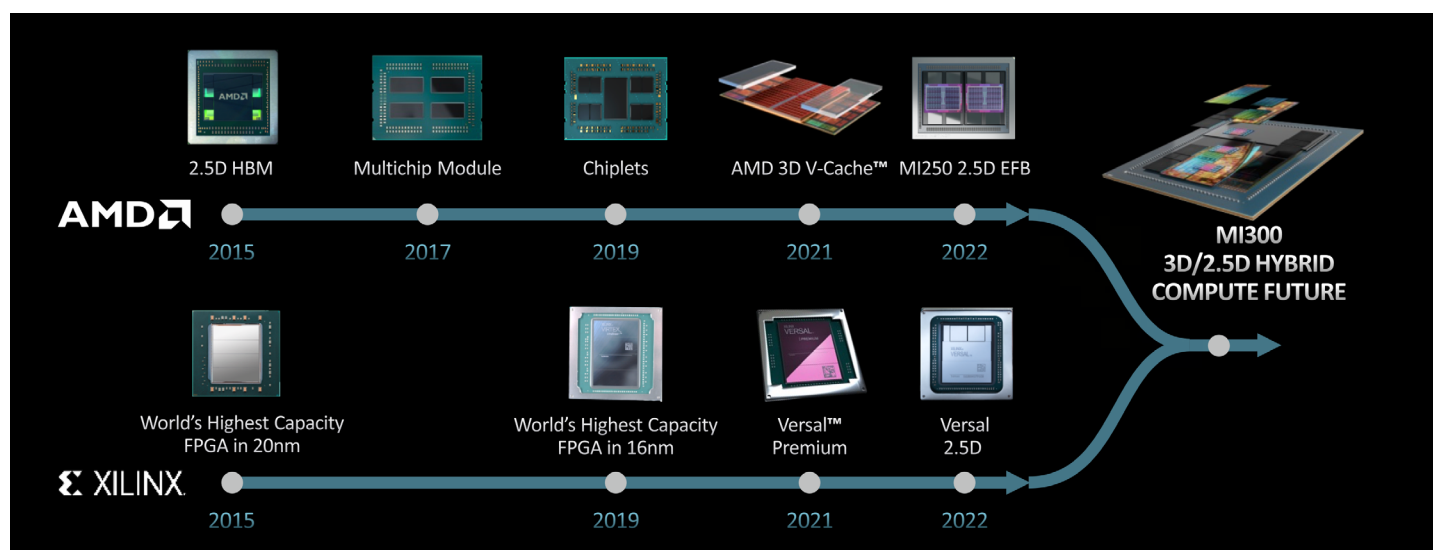


Figure 1, AMD Chiplet Innovations

AMD has been leading the way in chiplet architecture innovation. The future roadmap broadens potential solutions by integrating chiplets created by both customers and partners. These application-specific chiplets, when used with a standardized interface, can provide custom solutions which are especially valuable given the diverse range of computing tasks and workloads. New developments can be brought to market more swiftly and with substantially lower investment. This flexibility and speed will increase market-wide adoption of chiplet design and manufacturing.

BUILDING CHIPLET-BASED SOCS

An SoC (System-on-Chip) built with chiplets relies on a central chiplet called an anchor. This anchor orchestrates the essential system level functions:

- **Power Management:** Ensuring efficient and stable power delivery for all chiplets within the SoC. Including thermal management for the entire SoC.
- **Security:** Implementing security controls and protocols to safeguard the entire system.
- **Reliability, Availability, and Serviceability (RAS):** Handling error reporting, error isolation, and error recovery from issues, and employing preventative measures to maintain system health.
- **Interconnectivity:** The anchor chiplet might also manage communication between different chiplets, external memory resources and system level resources/controllers

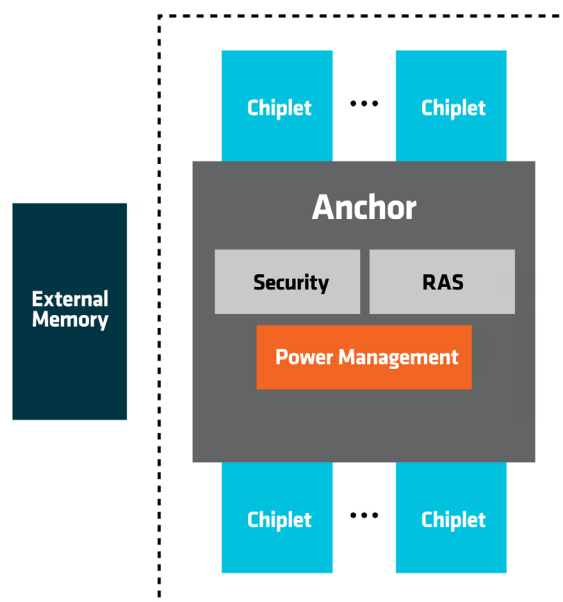


Figure 2, SoC Integration Using Chiplets

AMD classifies chiplets into three categories:

- **Internal Chiplets:** These chiplets, designed by AMD for use within their SoCs, can leverage a mix of standard interfaces and proprietary interfaces for specific functionalities.
- **Third-Party Die (TPD):** These chiplets come from external vendors and require standard interfaces and protocols. This allows seamless integration despite different protocols.
- **Third-Party Adapted die (TPA):** TPA combines AMD's intellectual property with third-party IP, allowing for advanced features and compliance with anchor services, regardless of industry standards. AMD offers a chiplet communication infrastructure to facilitate seamless integration of third-party IP with AMD anchors. TPA leverages AMD's proprietary services, which are built on their extensive experience catering to both consumer and enterprise markets.

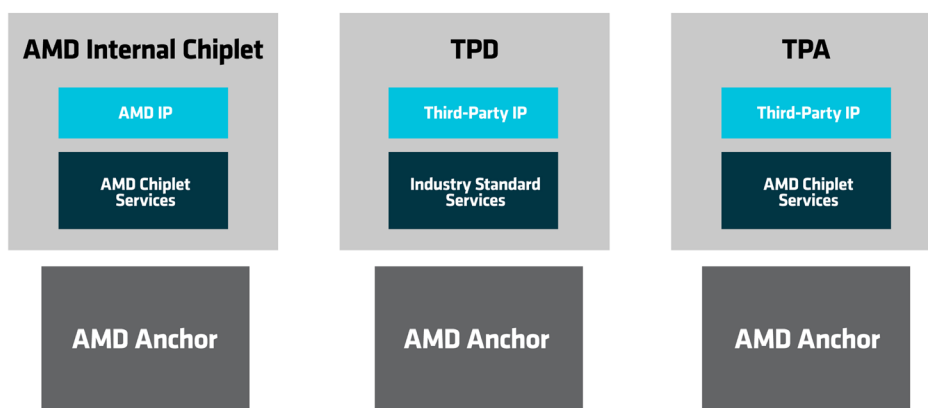


Figure 3, AMD, TPD, and TPA Chiplet Integration Diagrams

CHIPLET ECOSYSTEM

A chiplet ecosystem is required to unlock the full potential of chiplet technology and address current limitations. Having this ecosystem provides the following benefits:

- **Innovation Acceleration:** An open ecosystem allows diverse players and non-tradition SoC developers to contribute chiplet designs, fostering faster innovation and a wider range of specialized options.
- **Reduced Costs:** Standardized interfaces and design tools can streamline chiplet development and integration, leading to lower costs for both chiplet manufacturers and end-users.
- **Flexibility and Choice:** A robust ecosystem provides a broader selection of chiplets, allowing designers to mix and match for optimal performance and cost depending on the application.
- **Security and Reliability:** Shared standards and best practices can ensure consistent quality, security, and interoperability across chiplets.

Without a strong ecosystem, chiplet development might become fragmented, limiting innovation and potentially increasing costs. By working together, the industry can establish a foundation for chiplet technology to truly thrive.

A thriving chiplet ecosystem requires three main ingredients:

1. **Open Standards:** Universal interfaces, like the Universal Chiplet Interconnect Express™ (UCIe™), enable chiplets from different vendors to work together as interchangeable modules.
2. **Mature Design Tools:** EDA (Electronic Design Automation) tools need to streamline chiplet design and verification, making the process faster and more efficient.
3. **Manufacturing Capability:** Increased capacity for both chiplet production and advanced packaging is crucial to support the growing demand for this modular approach.

Chiplet ecosystem relies on a comprehensive set of technologies, forming the foundation for seamless communication and functionality across different chiplet types. This “tech stack” covers the entire development process, ensuring interoperability:

- **Interconnect Technologies:** These define the physical connections between chiplets, essentially the communication highways within the SoC. Standardized interconnect technologies robust die to die communication.
- **Interoperable Protocols:** These function as a common language, establishing how chiplets communicate and share information. Standardized protocols guarantee all chiplets understand each other, regardless of their origin or vendor.
- **Firmware and Configuration Data:** Boot and runtime configuration and loading, delivery, and validation of the on-die FW for runtime execution.
- **Command Processing:** Often implemented in firmware, command processing represents the control environment and a scheduler for chiplet tasks.
- **Software Stack, Runtime Services and Drivers:** These provide interfaces to end-user applications and services like data flow management and security enforcement.

CHIPLET CAPABILITIES

A robust chiplet ecosystem needs to support a diverse range of chiplet types, each offering unique functionalities within an SoC:

I/O Capable Chiplets: These chiplets function as the interface between the SoC and the external world. They provide essential functionalities like memory controllers, analog interfaces, and high-speed communication interfaces.

Accelerator Chiplets: These specialized chiplets offload specific tasks from the main processing cores, improving overall SoC performance. Examples include:

- **Custom Arithmetic Engines:** Tailor-made computations designed to accelerate arithmetic-heavy algorithms. These can include specialized cryptographic calculations or financial computations.
- **Compression/encoding engines:** Hardware-assisted compression and decompression of data for efficient storage and transmission.
- **Networking engines:** Optimized chiplets for handling network traffic and communication protocols.

Compute Cores: These chiplets form the heart of the SoC’s processing power. These cores include:

- **CPU cores:** Central processing units responsible for executing general-purpose instructions.
- **GPU cores:** Graphics processing units designed for high-performance computations and tasks requiring massive parallelism.

CHIPLET INTEGRATION

The technologies described above can be further detailed as *services* for chiplet integration:

Data Path Communication: Main services for management of application data. Chiplet integration requires communication services for:

- **Memory access:** Efficiency access to internal and external memory. This includes support for virtual memory.
- **Caching and Coherency:** Support to ensure all processors and accelerators in the SoC see the same view of memory.

Power Management:

- **Simple Static Power States:** Chiplets that require continuous power delivery and lack the capability for active power management.
- **Dynamic Power Management:** For more power management aware chiplets, the system can monitor overall power consumption and dynamically adjust power delivery to individual chiplets based on real-time requirements.
- **Chiplet-Level Power Management:** Highly specialized chiplets might have dedicated power budgets and manage their consumption independently, still notifying anchor about their power needs. Chiplet is also allowed to manage its low power state based on whether Anchor can tolerate this power state, and specifically the latency associated with the wakeup.

Thermal Management:

- **Chiplet level Thermal Management:** Allows for gradual regulating of the operational conditions to safely prevent thermal overages.
- **SoC Thermal Failure Management:** Catastrophic thermal event management. Depending on the policy or SoC topology this action may be limited to single chiplet or group of chiplets sharing the same power supply.

Reliability, Availability, and Serviceability (RAS):

- **Error Reporting:** If a chiplet encounters an issue, it is reported and logged for future troubleshooting.
- **Isolation and Recovery:** In case of a problem, the system needs to isolate the problematic chiplet to prevent an error or fatal event to propagate through the system. Additionally, there should be mechanisms for recovery and repair. Once isolated an anchor may attempt to repair or correct the fatal situation, without power cycling entire system.
- **Monitoring and Preventive Mitigations:** The ideal scenario is to prevent problems before they occur. This involves constantly monitoring chiplet health and taking preventative measures to avoid potential issues.

Security:

- **Centralized Security Controls:** The anchor plays a crucial role in overall security, serving as Root of Trust (RoT) for the SoC.
- **Chiplet-Level Security:** Individual chiplets can also contribute to security by handling tasks like secure boot, secure debugging, chiplet identification, and attestation. This works in conjunction with the central security services for a layered defense.

INTEGRATION EXAMPLE

Imagine integrating a cutting-edge accelerator from an external vendor into an AMD “Zen”-based SoC. In this scenario, AMD’s TPA integration model is used to incorporate the high-performance accelerator within the SoC. It interfaces with the AMD SoC through AMD’s Chiplet Communications, providing premium access to AMD’s datacenter-class Anchor infrastructure, which includes these features:

- **High-Speed Interconnects (Infinity Fabric™):** AMD’s proven Infinity Fabric™ technology ensures seamless SoC components. This high-performance fabric guarantees efficient data transfer for optimal system performance.
- **Advanced Memory Interfaces:** AMD’s expertise in high-performance memory interfaces like DDR and HBM allows the entire SoC, including the TPA, to access data efficiently. This is crucial for handling demanding workloads in datacenter environments.
- **Datacenter Reliability and Services:** AMD’s established infrastructure for datacenter reliability and services ensures the TPA and the entire SoC operate with exceptional uptime, stability, and error correction capabilities.

AMD’s TPA integration model supports all the services described in the previous section.

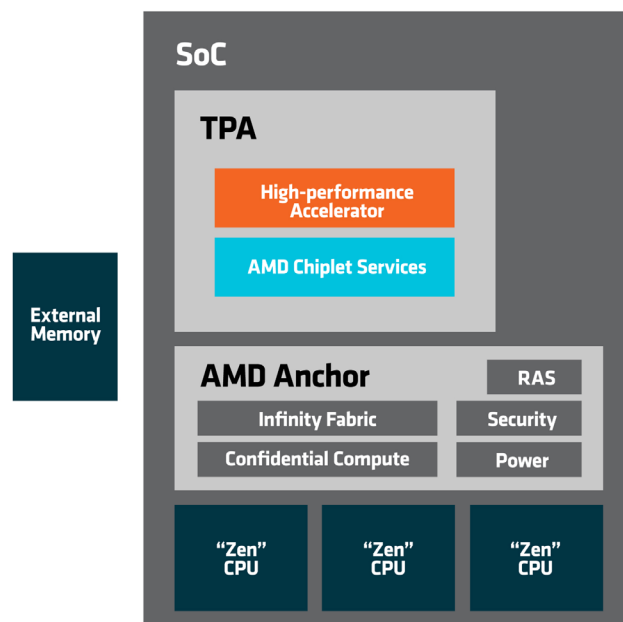


Figure 4, Enhancing “Zen”-Based SoCs with Third-Party Acceleration

BRIDGING THE GAP, INTERIM SOLUTIONS FOR THE CHIPLLET ECOSYSTEM

AMD is committed to fostering a robust chiplet ecosystem that unlocks the full potential of heterogeneous integration. While a comprehensive set of standards is still under development, AMD understands the need for immediate solutions to enable seamless integration of third-party chiplets into our SoCs.

AMD is bridging the gap with a roadmap of interim solutions:

- **Intermediate third-party die (iTPD):** This approach focuses on I/O subsystem chiplets from external vendors. These chiplets can seamlessly integrate with our anchor chiplet, ensuring compatibility within the existing infrastructure. This allows us to leverage the expertise of external vendors for specialized I/O subsystem functionality while maintaining a unified system. The iTPD can be UCle 1.1 or UCle 2.0 compatible.
- **Third-party adapted die (TPA):** As outlined earlier, the TPA feature enables customers to enhance both functionality and performance by integrating advanced accelerators from third-party vendors, all while maintaining the integrity of the overall system.

AMD EMBRACES UCIE™—THE FUTURE OF SEAMLESS CHIPLLET INTEGRATION

AMD is dedicated to being at the forefront of chiplet technology. As the UCle specification evolves, AMD is committed to fully support UCle-based Third-Party Die (TPD) for seamless integration within our SoCs:

- **Standardized Services:** Full support for UCle Management Transport Protocol (MTP) and its associated services like security, power management, RAS (Reliability, Availability, Serviceability), boot, and more.
- **Enhanced Communication:** For efficient data exchange, adoption of new protocols conforming to UCle flit formats, including CXL/PCIe, CHI/C2C, and novel streaming formats.
- **Simplified User Experience:** Development of drivers and applications ensuring effortless utilization of TPDs within our SoCs.

Partnering with AMD for Custom Chiplet Development

AMD has a history of tailoring solutions to meet customer needs, with high-performance game consoles being the most notable examples. Beyond custom SoC solutions, AMD's design services enable customers to develop unique chiplets, facilitating the rapid introduction of innovative products. Additionally, partners can work with AMD and customers to incorporate their Intermediate Third-Party Die (iTPD) and Third-Party Adapted Die (TPA) into AMD's superior-quality SoCs.

CHIPLLET DEVELOPMENT

Chiplet development involves various stages to ensure the success and reliability of System on Chips (SoCs). Key components of this process include verification, emulation, validation, and physical design. Each stage plays a role in confirming that chiplets meet design specifications, function properly under realistic conditions, and integrate seamlessly into the larger system.

Verification checks if the chiplet design meets specifications through simulation, formal methods, or assertion-based verification. AMD's strategy focuses on chiplet-level coverage, reducing the need for extensive SoC-level verification.

Emulation uses hardware platforms like FPGAs to evaluate chiplet functionality and performance under realistic scenarios, faster than simulations.

Validation ensures that each chiplet operates correctly within its intended environment before full-scale production. This process involves rigorous testing on actual silicon across various fabrication stages, confirming that the chiplet's performance, functionality, and reliability meet the requirements. Validation entails a series of evaluations, including electrical testing, performance benchmarking, and stress testing under operational conditions.

Physical design for chiplets requires joint planning between developers and integrators, affecting footprint, density, performance, cost, reliability, and bandwidth. Key factors include chiplet shoreline, I/O routing, and power/thermal considerations. As the ecosystem evolves, standardization and automation in physical design will facilitate quicker and easier chiplet integration.

AMD collaborates with chiplet and EDA vendors to provide a unified environment for verification, emulation, and validation, including compatible verification environments, well-defined interfaces, and comprehensive validation frameworks. This unified framework accelerates time-to-market and reduces chiplet integration risks.

CONCLUSION

Building a Thriving Chiplet Ecosystem

At AMD, we are dedicated to cultivating a strong chiplet ecosystem that maximizes the benefits of heterogeneous integration. We understand the importance of being flexible and working together to realize this goal.

Our interim solutions like the AMD iTPD and TPA integration models are steppingstones on the way to the standardized chiplet integration with UCle. These solutions facilitate the smooth incorporation of third-party capabilities; bringing in external proficiency while ensuring compatibility, performance, and reliability within our SoCs.

The advancement of chiplet technology depends on cooperation and open standards. We actively support the UCle standard and are committed to its development to achieve seamless communication and compatibility among various providers and chip types. This collaborative spirit enables us to build a vibrant chiplet ecosystem that propels processing power to new heights and fuels innovation throughout the sector.

DISCLAIMERS

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18

COPYRIGHT NOTICE

© 2024 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD AIDay, AMD Pensando, AMD Virtualization, AMD-V, Infinity Fabric, PowerPlay, Vari-Bright, and combinations thereof are trademarks of Advanced Micro Devices, Inc. PCIe is a registered trademark of PCI-SIG Universal Chiplet, Universal Chiplet Interconnect Express, and UCle are trademarks of UCIE Corp. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.