ARTERIS IP

Datasheet

CodaCache® Last-Level Cache IP

Overview

CodaCache is a configurable, standalone, non-coherent cache IP that delivers unique business value through its advanced last-level cache (LLC) architecture, improving system performance, data locality, scalability, power efficiency, application responsiveness, cost optimization, and market competitiveness.

SoC challenges.
Current challenges in SoC designs include optimizing data sharing between compute engines and accelerators, and other data processing blocks, which demand efficient data prefetching mechanisms to reduce reliance on main memory accesses. Satisfying these requirements will ensure that performance goals are met at block and SoC levels.

Arteris solution.
CodaCache Last-Level Cache IP finds practical uses in various scenarios within computer systems. Its primary purpose is to enhance system performance by reducing memory access latency. CodaCache achieves this by storing frequently accessed data closer to the processing cores, enabling faster retrieval. It acts as a shared cache, facilitating seamless communication and synchronization between components or IPs.

The CodaCache IP serves as a temporary storage and buffering mechanism for data in transit, enabling efficient data flow management. Caching frequently accessed data minimizes the need to access the slower main memory, improving overall system responsiveness. This, in turn, contributes to power efficiency by reducing memory access and associated power consumption.

Highlights

- Standalone IP.
- 1.2 GHz frequency in 16FF+TT process.
- Protocol interoperability: AMBA AXI 4.
- 64 Byte cache line size.
- Cache size up to 8MB per AXI port.
- Scratchpad RAM configurations.
- Configurable way associativity from 1 to 16.
- Way partitioning.
- Cache flushing.
- ECC protection.
- Helps meet up to ISO 26262 ASIL B and D requirements against random hardware.

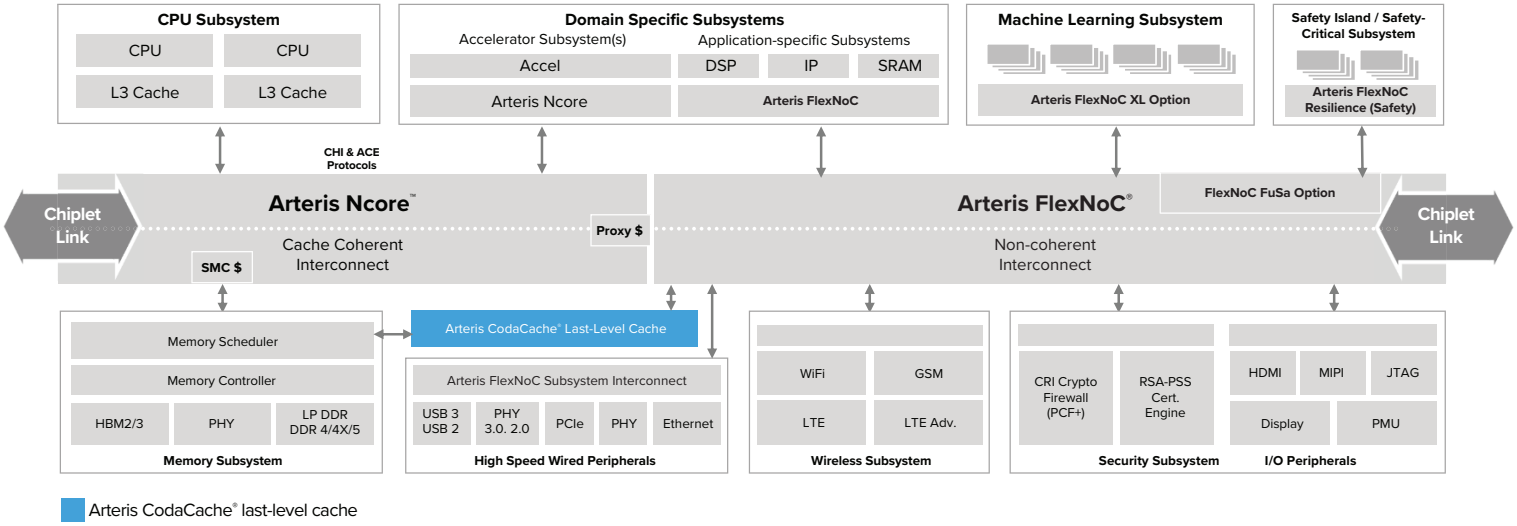

Overall, the practical uses of CodaCache IP include performance enhancement, efficient data sharing, temporary storage or buffers for data that needs to be processed or transferred between different components, and power efficiency by reducing the frequency of memory accesses. These capabilities make it a vital element for ASIC designs, supporting various applications, including high-performance computing, graphics and multimedia processing, AI, machine learning, and embedded systems.

Benefits

LLC implementation.
It provides a dedicated and configurable last-level cache that is significant and high-capacity for the SoC. The LLC acts as a buffer between the main memory and the processor cores, improving data access latency and reducing memory bandwidth requirements.

Scalability and configurability.
It is highly scalable and configurable, accommodating various SoC designs. It supports different cache sizes, associativity levels, and coherence granularities, allowing designers to tailor the cache hierarchy to their specific requirements. This flexibility ensures efficient utilization of available resources and enables scalability for diverse application scenarios.

Performance optimization.
It enhances system performance through efficient cache access and reduced memory access latency. It minimizes cache misses by providing a sizeable last-level cache closer to the processor cores, enabling faster data access, and reducing the need for frequent memory accesses. This improves overall system performance and responsiveness.

Reliability and error detection.
It includes error detection mechanisms to ensure data integrity and reliability. It incorporates error correction codes (ECC) and parity checking to detect and correct errors in cache data. These error detection capabilities enhance system robustness and reduce the risk of data corruption.

Ecosystem compatibility.
It seamlessly integrates with the broader design ecosystem, supporting industry-standard interfaces and third-party IP blocks. It ensures compatibility with various processor cores, interconnect fabrics, and verification environments, enabling designers to leverage existing IP resources and design flows, streamlining the integration process.

Features

Improved system performance.
Including a dedicated last-level cache closer to the processor cores reduces memory access latency and improves overall system performance. The larger cache capacity of the LLC allows for more data to be stored closer to the processing units, reducing the frequency of memory accesses and minimizing cache misses.

Enhanced data access latency.
It decreases data access latency by providing a high-performance cache layer between the processor cores and the main memory. This reduces the time required for fetching data from the main memory, improving the overall responsiveness of the system and accelerating application execution.

Reduced memory bandwidth requirements.
CodaCache minimizes the demand for memory bandwidth by incorporating a larger last-level cache. With frequently accessed data residing in the LLC, the need for frequent memory accesses is minimized, resulting in efficient utilization of memory resources and improved overall system efficiency.

Configurability and scalability.
It offers extensive configurability, allowing designers to tailor the cache hierarchy to meet specific requirements. It supports various cache sizes, associativity levels, and coherence granularities, enabling flexibility and scalability for diverse SoC designs. This adaptability ensures optimal resource utilization and efficient system scaling.

Reliability and data integrity.
CodaCache includes error detection mechanisms to ensure data integrity and reliability. It incorporates error correction codes and parity checking to detect and correct errors in cache data. These error detection capabilities enhance system robustness and reduce the risk of data corruption.

Safety mechanisms.
CodaCache integrates safety mechanisms to meet the requirements of ISO 26262 functional safety standards. It includes error detection and correction techniques, fault tolerance features, and redundancy mechanisms to ensure the reliable and safe operation of the cache subsystem in automotive applications.

ISO 26262 compliance.
CodaCache is designed with ISO 26262 functional safety standards in mind. It provides comprehensive safety documentation, including safety manuals, FMEDA (Failure Modes, Effects, and Diagnostic Analysis) reports, and safety verification kits. These documents and artifacts assist SoC designers in achieving compliance with ISO 26262 requirements.

Seamless integration and ecosystem compatibility.
Standalone CodaCache seamlessly integrates with existing design ecosystems, supporting industry-standard interfaces and third-party IP blocks. It ensures compatibility with various processor cores, interconnect fabrics, and verification environments, simplifying the integration process and allowing designers to leverage various IP resources.

Complementary Product

CodaCache Last-Level IP, in conjunction with FlexNoC 5 Physically Aware Interconnect IP, FlexWay 5 Core Interconnect IP, and Ncore Cache Coherent IP, form a powerful combination for system-on-chip (SoC) designs. Ncore provides a high-performance coherent interconnect fabric, and FlexNoC and FlexWay offer scalable, flexible, physically aware noncoherent NoC IP solutions. CodaCache can be used to improve IP and system performance. Together, they deliver improved performance, scalability, and efficient data sharing in complex SoC designs.

Boost Performance

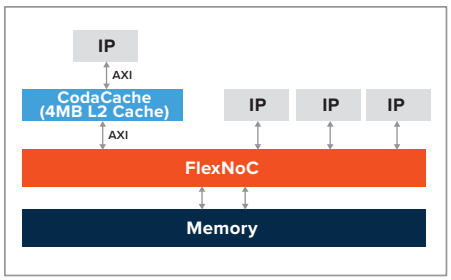

Dedicated Cache
- IP performance boost

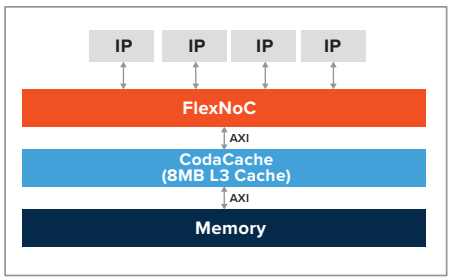

Last-Level Cache
- SoC performance boost

Reconfigure for timing closure and repartitioning for layout congestion

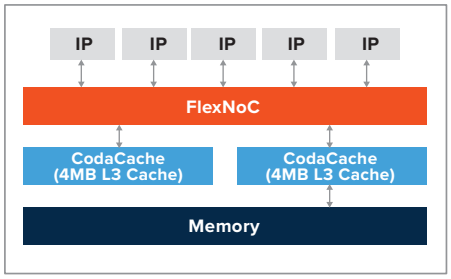

- Bigger cache size
- Higher bandwidth

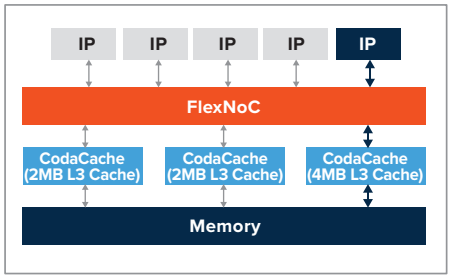

- Dedicated cache
- Address interleaved shared cache
- Smaller cache to meet frequency requirement

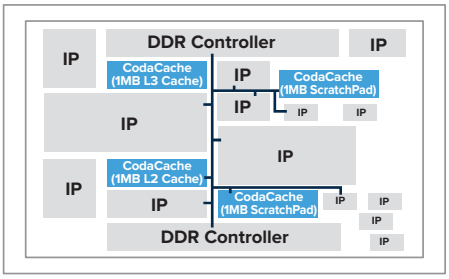

- Cache partitioning to ease the layout

Runtime configuration for unique data access patterns

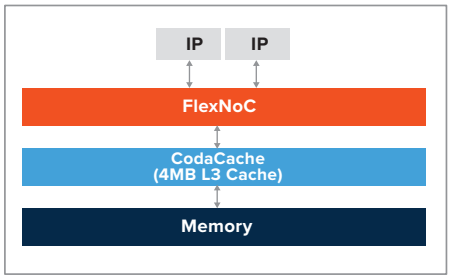

Uniform (traditional) data access

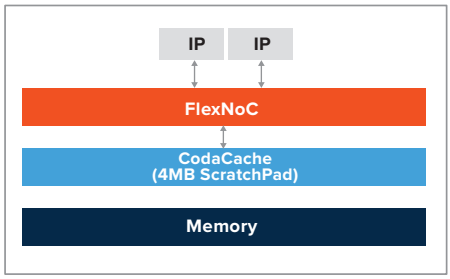

Non uniform data access – temp storage for calculation and other work in progress

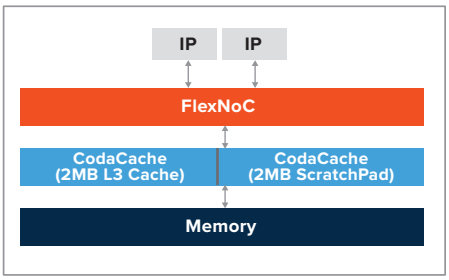

Mix (uniform and non uniform) data access

About Arteris

Arteris is a leading provider of system IP for the acceleration of system-on-chip (SoC) development across today's electronic systems. Arteris network-on-chip (NoC) interconnect IP and SoC integration technology enable higher product performance with lower power consumption and faster time to market, delivering better SoC economics so its customers can focus on dreaming up what comes next. Learn more at arteris.com.

ARTERIS IP

arteris-support@arteris.com +1 408 470 7300

Rev 2/24