ARC-V RHX-100 Series Processors

Highlights

- Dual-issue, 32-bit processors for high- performance real-time applications
- Multicore Processor versions with up to 16 CPU cores and up to 16 hardware accelerators
- Based on the RISC-V ISA, leveraging standard 32-bit protocols (and extensions)
- Real-time enhanced virtualization
- High degree of configurability
- Support for custom instructions
- Support for up to 16 MB of closely coupled memory and direct mapping of peripherals
- Optional MMU with hardware page table walk and up to 16 MB page sizes
- Optional support for RISC-V defined vector extensions (RVV)
- Real-time Trace (RTT) provides real-time trace debugging features

Target Applications

- Industrial: Robotics, Medical devices, factory automation
- Automotive: Domain controller, infotainment, safety management
- Consumer: Display control, laser printers, soft modems
- Storage: Solid state drive (SSD) controller, computational storage
- Networking: NIC, Ethernet switches, wireless access (WAP)

Overview

The MIPS ARC-V™ RHX-100 series processors feature a dual-issue, 32-bit superscalar architecture for use in applications where real-time performance is required. The cores offer outstanding performance delivering with a small area footprint and low power consumption.

The ARC-V RHX-100 processors are based on the RISC-V instruction set architecture (ISA). The processors feature a 34-bit physical address space defined by the RISC-V Sv32 MMU. For applications requiring higher performance, the multi-core RHX-105 and RHX-105V are available with up to 16 CPU cores and up to 16 hardware accelerators in the processor cluster. RISC-V vector extensions (RVV) are available in the RHX-100V (single core) and RHX-105V (multi-core) processors.

The ARC-V RHX-100 features level 1 (L1) instruction and data cache and close coupled memory (CCM) and is optimized for use in high-performance real-time embedded applications.

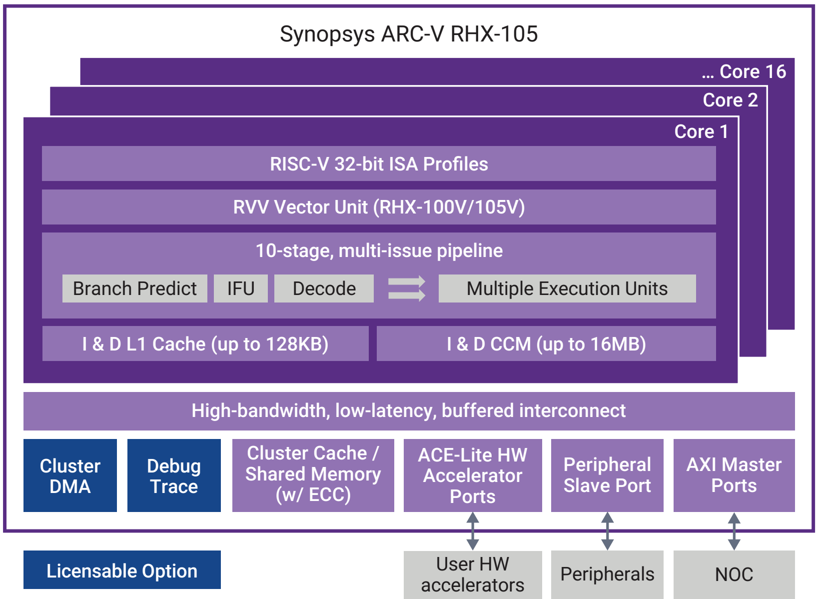

Figure 1. ARC-V RHX-105 Real-Time Processor Block Diagram

Features

- High-speed, 32-bit, dual-issue, 10-stage pipeline
- Multicore support for up to 16 CPUs and up to 16 user hardware accelerators per processor cluster
- Optional enhanced RISC-V Sv32 MMU with support for Linux and SMP Linux
- 4 KB to 128 KB instruction and data L1 cache
- Up to 16 MB cluster cache / shared memory
- Up to 16 MB instruction and data closely coupled memory (CCM)
- Real-time enhanced virtualization with 3-stage memory protection unit for OS needs (L1), for VM isolation (L2) and for system level isolation (L3). 128-bit loads and stores per clock
- Radix-4 hardware divider
- Up to 2x Advanced Platform Level Interrupt Controllers (APLIC) each supporting up to 1023 wired interrupts for a maximum of 2046 interrupts
- Native ARM® AMBA® AXI interfaces™, AHB-Lite™ and CHIinterfaces
- JTAG and Compact JTAG (cJTAG) debug interface

Dual-Issue Pipeline

The high-performance, 32-bit, dual-issue, 10-stage pipeline is optimally balanced to achieve very high embedded performance with low power consumption. The pipeline is designed to issue up to two instructions per clock (in order) and features two-cycle access to memory allowing the processor to be run at higher clock speeds and making it less sensitive to memory size. The pipeline supports out-of-order retirement, sophisticated branch prediction (with early correction of mispredicted branches) and a patented late-stage ALU that improves instruction throughput. The processors feature a 32-bit virtual address range and a 40-bit physical address range enabling the direct addressing of memories up to 4.5 Petabytes (4.5x1015) in size. Configurable support for a 32-bit hardware multiply, vector addition and subtraction, and a Radix-4 hardware divider are included. Several separately licensed options are available including: real-time trace debug, and cluster DMA.

Multicore Versions

The RHX-100 cores are available in single-core and multicore processor versions. The multicore processor versions support up to 16 CPU cores and up to 16 user hardware accelerators per processor cluster enabling very high-performance scaling for high-end embedded applications. The advanced multicore interconnect can support up to 800 GB/s aggregate bandwidth enabling the CPU cores and hardware accelerators to run at maximum throughput and speed. The CPU cores and hardware accelerators can be implemented in their own clock and power domains and can have an asynchronous clock relationship to the other cores and interconnect. The CPU cores support L1 coherency, and coherency is support between the hardware accelerators, and also between the shared cluster cache and the Network-on-Chip (NoC) if desired. Quality of Service (QoS) control is built into the multicore interconnect that enables the user to schedule bandwidth to the CPU cores and hardware accelerators to insure balanced loading and real-time operation. The shared cluster cache is configurable up to 16MB and supports the CPU cores and user hardware accelerators and can be kept coherent with memory.

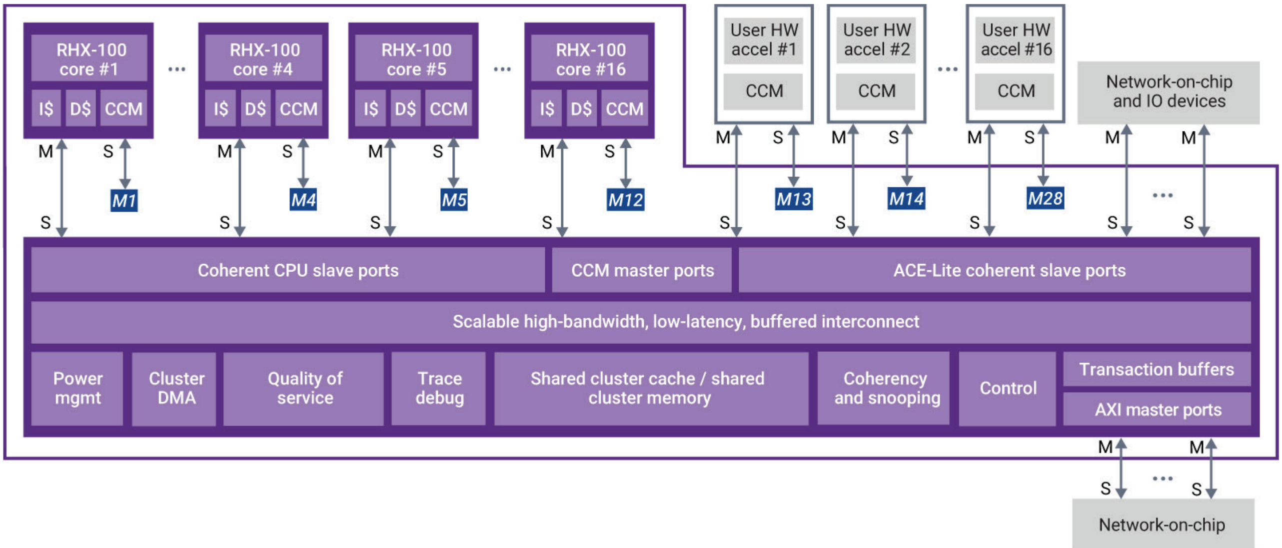

Figure 2. ARC-V RHX Architecture

Configurable Options

The processors support a broad range of configurable options, enabling optimization for a specific application's performance, power and size requirements. The included ARChitect configuration tool features a graphical interface and produces verified RTL and synthesis scripts that are compatible with industry-standard design flows.

Memory Architecture

L1 Cache Memory

The ARC-V RHX-100 processors feature separate instruction and data L1 caches that can be independently configured for 4 K, 8 K, 16 K, 32 K or 64 K size. The instruction and data cache is fixed to 2-way associativity, and a user-selectable line size of 32, 64 or 128 bytes. The caches can be individually configured to support line locking and invalidate, and to offer debug visibility. For multicore configurations the L1 data cache implements the MOESI protocol and supports coherency and cache to cache transfers.

L2 Cache Memory (optional)

The ARC-V RHX-100 has optional support for up to 16 MB of L2 cluster shared cache. The L2 cache and the shared cluster memory share the 16MB space and the size of each is under the user's control. The L2 cache can be used with a single-core or multicore processor implementations. The coherent L1 caches will work in concert with the L2 cache and all CPU cores and hardware accelerators in a multicore processor can be used with the L2 cache. This cache is designed to run at the same clock frequency as the CPU core. The L2 cache is tightly connected to the core(s) through a proprietary low latency bus. The L2 cluster shared cache feature can be separately licensed for use with the HS66.

Cluster Shared Memory

The Cluster Shared Memory (CSM) provides a high capacity, high-throughput memory shared by all cores within a multicore processor cluster. It can be used for instructions and data, cacheable and non-cacheable, and sharable and non-sharable address space. The CSM can be accessed by each core using the memory access instructions and instruction fetch operations. The CSM is located in the ARC-V RHX-100 cluster and can be mapped to the system memory space. Transactions from components within the cluster, such as processor cores, SCU, hardware accelerators and so on, access the CSM through the IBP interfaces, and transactions from external masters access the CSM through the CSM-DMI interface (if CSM-DMI is configured). The memory macro consists of up to 8 banks so multiple transactions from different interfaces can access the CSM simultaneously. The size of the CSM is build-time configurable, supporting 32 KB, 64 KB, 128 KB, 256 KB, 512 KB, 1 MB, 2 MB, 4 MB, 8MB and 168MB, and memory ECC and parity protection are supported. The CSM shares the memory space with the L2 cluster cache and the maximum size of both combined is 16MB. The CSM clock is gated to reduce dynamic power consumption, and it can be powered down independently while other components and cores are powered-up.

Closely Coupled Memories

The ARC-V RHX-100 processor supports 512 B to 16 MB of closely coupled memory (CCM) for both instruction and data. The CCM is implemented as separate memory spaces for the Instruction Closely Coupled Memory (ICCM) and Data Closely Coupled Memory (DCCM), and can be mapped in the global address space. Both ICCM and DCCM have optional support for error-correcting code (ECC) to increase application reliability.

Memory Management Unit (optional)

The ARC-V RHX-100 processor features an optional MMU, which enables the core to run sophisticated embedded operating systems like Linux and SMP Linux. The MMU implements a hardware page table walk and supports variable page sizes with concurrent support for the normal pages (up to 16 KB) as well as large pages (up to 16 MB). The primary Translation Lookaside Buffer (TLB) on the MMU has 1,024 entries and is four- way set associative. The TLB has fully associative micro-TLBs: a four-entry iTLB for instructions and an eight-entry dTLB for data. There is also a 16 entry secondary TLB for super/large pages. The MMU can be separately licensed for use with the RHX-100 series processors.

System Architecture and Interfaces

Register File and Program Counter

The RHX-100 processor cores each support 31 base integer registers with an additional 32 floating registers as a configuration option.

Bus Interfaces

The processors have native support for the AMBA AXI protocol for connection to a NoC. This is a build-time option with the AMBA AXI interface as the default selection. The AXI buses are user configurable with 64- or 128-bit widths to improve system throughput.

Interrupts and Exceptions

The RHX-100 series supports the Advanced Interrupt Architecture (AIA). Individual RHX-100 cores receive Message Signaled Interrupts (MSIs) from an Advanced Platform Level Interrupt Controller (APLIC) with real-time delivery guarantees. The RHX-100 processor supports up to two APLICs, each of which handles up to 1023 wired interrupts for a maximum of 2046 interrupts. The interrupts are serviced through jump tables and interrupts can be triggered by software as well as hardware.

Optimized and guaranteed latency of external interrupt delivery - direct or/and MSI mode. Multiple flexible external interrupt controller integration options (tied to the core, cluster level, system level), all provide the same hardware/software interface including native and virtualized environments.

Configurable number of always resident (hardware) interrupt controller contexts for virtual machines. Optimized interrupt controller emulation (trap-and-emulate) for virtual machines without a dedicated HW interrupt context. Redirection of interrupts coming to an idle state VM to hypervisor

Fast Context Switch Mechanism

Advanced gate-count efficient mechanism for saving / restoring application context, including fast interrupts handling. Provides virtually unlimited number of contexts. Software configurable handling of core registers as well as CSR's.

Real-time Enhanced Virtualization

3-stage memory protection unit for OS needs (L1), for VM isolation (L2) and for system level isolation (L3). Fast MPU context switching for accelerated and predictable VM context switching. Support of virtualized timer and extended virtual IPI (SWI).

ARC Processor EXtension (APEX) Interface (RISC-V-defined custom extensions)

The processors are designed to be extendable with the addition of custom instructions. These instruction extensions may include more processor and auxiliary registers, new instructions, and additional condition code tests. Custom instructions enable designers to efficiently add their proprietary hardware to the processor to further increase application performance. The APEX custom instruction interface is available on all RHX-100 CPU cores in a Multicore processor and is separate from and in addition to the user hardware accelerators.

Low-Latency Auxiliary Port

The low-latency auxiliary port offers fast access to on-chip peripherals and memory. The port supports single-cycle read and write register access bringing these peripherals close to the processor and significantly reducing the latency that accompanies accessing peripherals and memory over a multi-layer AMBA bus interface. In addition, system performance is improved because this processor-to- peripheral bus traffic is moved to the auxiliary bus. The auxiliary port can support all peripheral registers on an SoC.

Optional Features (Separately Licensed)

- ARC Trace I/Fprovides real-time trace debugging features for the ARC-V RHX processor
- Cluster DMA programmable memory access controller
- RISC-V vector extensions (available with RHX-100V & RHX-105V) processors

Complete Suite of Development Tools

To facilitate rapid development, the processors are supported by a complete suite of development tools. This includes the MetaWare Development Toolkit that generates performance optimized code that takes advantage of the dual-issue pipeline that is highly efficient and ideal for deeply embedded applications, the ARC nSIM simulators and the ARChitect configuration tool.

MIPS offers a suite of GNU tools (ARC GNU) for developers targeting the Linux operating system as well as bare metal systems. The ARC GNU Toolchain includes the GCC compiler and GDB debugger as well as a number of utilities and libraries that make up a complete software toolchain.
Additional development tools are available as part of the RISC-V ecosystem.

Compile	MetaWare Compiler	• Optimize your code for size and performance • Leverage core-specific features to reduce cost and increase performance • Utilize your user-defined instructions to achieve design goals
	GNU GCC Compilerr	• Freely access an open source solution with the GCC compiler
Debug	MetaWare Debugger	• Easily debug multiple targets with the same user interface • Quickly profile hotspots in your code • Use scripting to increase productivity
	JTAG Debuggers	• Efficiently bring-up hardware with tools from Ashling, Green Hills and Lauterbach
	GNU GDB Debugger	• Use the open source GDB debugger to debug real and simulated targets
Deploy	nSIM Simulator	• Develop & debug software before hardware is available • Simulate large programs with very fast ARC-V models • Quickly optimize your software with near cycle-accurate simulation
	Zephyr Real Time OS	• Open source RTOS optimized by MIPS for ARC processors
	Linux, SMP Linux	• Linux for ARC-V Processors provides all of the benefits of open source software, including complete source code and a large installed base

Documentation

The following documentation is available for the ARC RHX processors:

- ARC-V Programmers Reference
- ARC-V RHX-100 Series Technical Reference Manual
- ARC-V RHX-100 Series Integration Guide
- ARC APEX Databook

Testing, Compliance and Quality

Verification of the MIPS ARC-V RHX-100 series processors follows a bottom-up verification methodology from block level through system level. Each functional block within the product follows a functional, coverage-driven test plan.

The plan includes testing for RISC-V ISA compliance as well as state- and control- specific coverage points that have been exercised using constrained pseudo-random environments and a random instruction sequence generator.

Deliverables

The MIPS ARC-V RHX-100 processors are delivered in Verilog HDL in the ARChitect IP Library. The HDL is configurable by the user and output from the ARChitect IP Configurator tool. To test that the product performs as expected, a basic testbench of Customer Confidence Tests (CCT) is included.

About MIPS:

MIPS by GlobalFoundries delivers software to silicon with RISC-V for building physical AI platforms. MIPS delivers software-hardware co-design, optimized AI, and custom ASSP design and manufacturing. Together with ARC, MIPS delivers the open, standards-based processor IP portfolio for embedded applications. Physical AI is built on MIPS.

For more information, visit **www.mips.com/arc.**