

ARC HS66 and HS68 Processors

Highlights

- Dual-issue, 64-bit processors for high-performance embedded applications
- 52-bit physical and 64-bit virtual addressing
- Up to 5400 DMIPS and 11,088 CoreMark per core at 1.8 GHz on 16FFC (worst case conditions, single-core configuration)
- Multicore Processor versions with up to 12 CPU cores and up to 16 hardware accelerators
- Based on advanced ARCV3 ISA
- High degree of configurability
- Enhanced MMU (HS68) with hardware page table walk and up to 16 MB page sizes
- Support for custom instructions
- Support for up to 16 MB of closely coupled memory and direct mapping of peripherals
- Floating Point Unit (FPU) supporting half, single- and double-precision IEEE 754-compliant operation
- ARC Trace Interface provides real-time trace debugging features
- Easy performance upgrade from HS3x and HS4x processors

Target Applications

- Solid state drive (SSD) controller
- Automotive systems
- Baseband control
- Home automation
- Home networking
- Data centers

The Synopsys ARC® HS66 and HS68 processors feature a dual-issue, 64-bit superscalar architecture for use in embedded applications where high-performance and high clock speed are required. The processors can be clocked at up to 1.8 GHz in 16FFC processes (worst case, single core, base configuration) and offer outstanding performance delivering 3.0 DMIPS/MHz and 6.16 CoreMark/MHz with a small area footprint and low power consumption.

The ARC HS66 and HS68 processors are based on the advanced ARCV3 instruction set architecture (ISA) and pipeline, which provides leadership power efficiency and code density. The processors feature a 52-bit physical address space and can directly address memories up to 4.5 Petabytes (4.5×10^{15}) in size. For applications requiring higher performance, Multicore Processor (MP) versions of the HS66 and HS68 are available with support for up to 12 HS CPU cores and up to 16 hardware accelerators in the processor cluster.

The ARC HS66 features level 1 (L1) instruction and data cache and closely coupled memory (CCM) and is optimized for use in high-performance real-time embedded applications. The HS68 is designed for use in applications running Linux or SMP Linux. The HS68 has all the features of the HS66 plus support for L2 cache up to 64 MB and a Memory Management Unit (MMU).

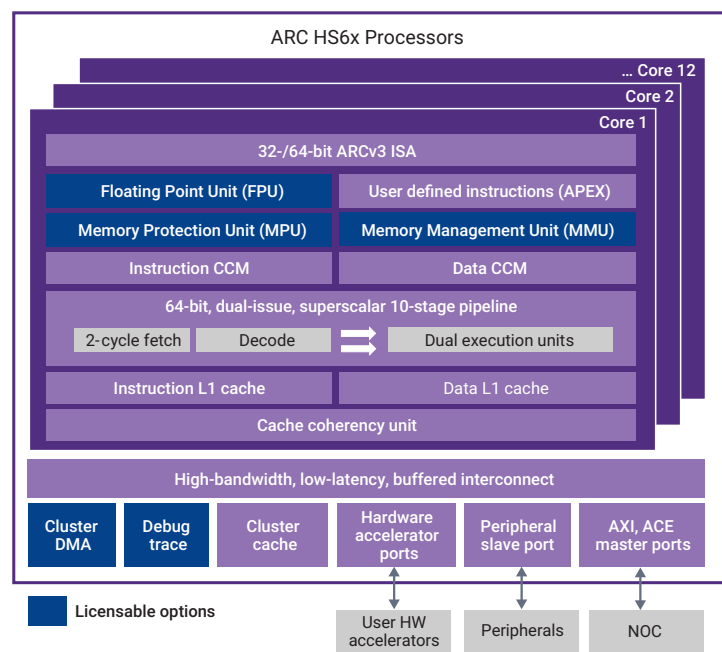


Figure 1. ARC HS6x Block Diagram

ARCV3 ISA

The ARC HS66 and HS68 processors implement the ARCV3 RISC ISA, which supports advanced processor capabilities and incorporates 16-bit, 32-bit and 64-bit instructions that improve code density compared to the previous generation ARC products. The architecture and pipeline are designed to meet the needs of next-generation system-on-chip (SoC) applications. The ARCV3 ISA enables the implementation of complex, heterogeneous SoCs with processors that are precisely targeted to meet the specific performance and power requirements for each instance on the SoC.

Features

- High-speed, 64-bit, dual-issue, 10-stage pipeline
- 64-bit virtual and 52-bit physical address space
- Multicore support for up to 12 HS CPUs and up to 16 user hardware accelerators per processor cluster
- Enhanced MMU with support for Linux and SMP Linux
- 4 KB to 64 KB instruction and data L1 cache
- Up to 64 MB L2 cache (HS68 and MP versions)
- Up to 16 MB instruction and data closely coupled memory (CCM)
- 128-bit loads and stores per clock
- Register file supports 2 write ports and 2 read ports
- Fast context switching with up to 8 register files
- 32x32 multiplier
- Radix-4 hardware divider
- Up to 240 interrupts, with up to 16 configurable preemption levels
- Native ARM® AMBA® AXI™, AHB-Lite™ and ACE interfaces
- JTAG and Compact JTAG (cJTAG) debug interface

Dual-Issue Pipeline

The high-performance, 64-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 64-bit virtual address range and a 52-bit physical address range enabling the direct addressing of memories up to 4.5 Petabytes (4.5×10^{15}) in size. Configurable support for a 32-bit hardware multiply (and multiply-accumulate), vector addition and subtraction, and a Radix-4 hardware divider are included. Several separately licensed options are available including: an IEEE 754-compliant floating point unit, real-time trace debug, and cluster DMA.

Multicore Versions

The HS66 and HS68 are available in single-core and multicore processor (MP) versions. The multicore processor versions support up to 12 HS 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 HS CPU cores support L1 coherency, and coherency is supported between the hardware accelerators, and also between the L2 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 ensure balanced loading and real-time operation. The L2 shared cluster cache is configurable up to 64 MB and supports the HS CPU cores and user hardware accelerators and can be kept coherent with memory, a host processor, or other functionality on the NoC.

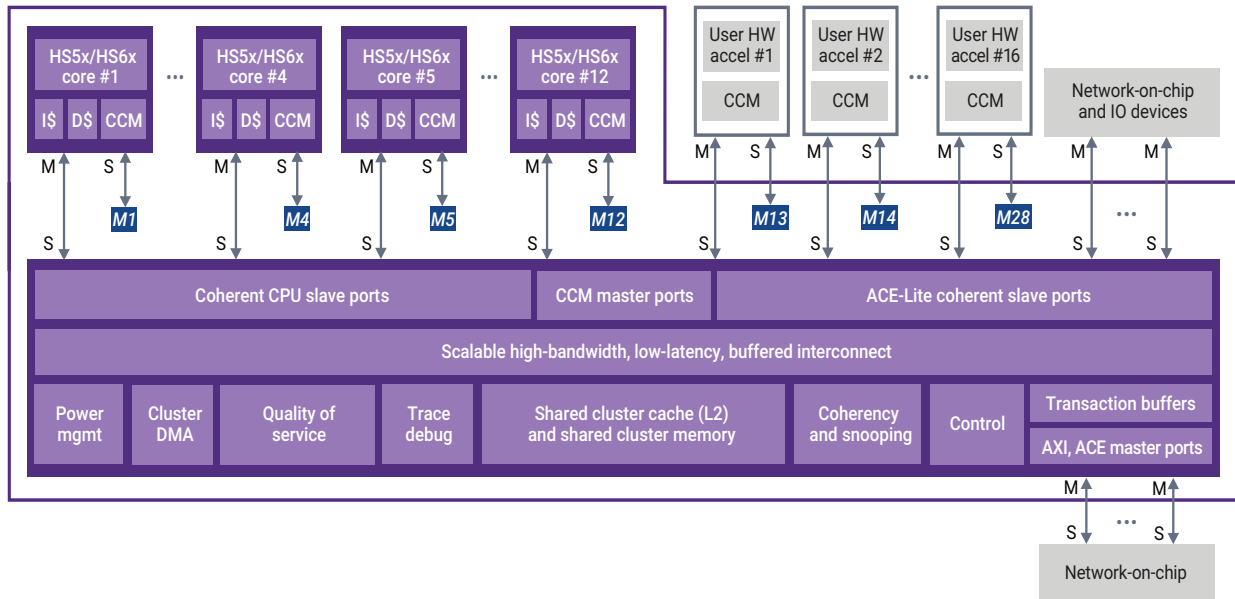


Figure 2. ARC HS6x Architecture

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 Management Unit

The ARC HS68 processor features a full 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 HS68 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 HS66.

L1 Cache Memory

The ARC HS66 and HS68 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 has 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

The ARC HS68 has support for up to 64 MB of L2 cluster shared cache. The L2 cache and the cluster shared memory share the 64 MB 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 (up to 1.8 GHz worst-case in 16FFC). 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.

Closely Coupled Memory

The ARC HS66 and HS68 support 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). Both ICCM and DCCM have optional support for error-correcting code (ECC) to increase application reliability.

Register File and Program Counter

The ARC HS66 and HS68 register file is configurable, with 16 or 32 64-bit registers. The register file can be constructed from flip-flops or SRAM memory and supports two write ports and two read ports. Some of the registers have defined purposes like the stack pointer (r28) and link registers (r29,- r31). The register file can be expanded to up to a total of 60 64-bit registers using the extension registers (r32–r59). The ARCV3 16-bit instructions can directly access registers r0–r15 and can indirectly access the remaining registers. The program counter is read only and located at r63 in the register file.

Bus Interfaces

The processors have native support for the ARM AMBA AXI or AHB-Lite bus protocols, and the multicore versions also support the ACE protocol for connection to a NoC. This is a build-time option with the ARM AMBA AXI interface as the default selection. The AXI and ARC busses are user configurable with 64- or 128-bit widths to improve system throughput.

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 and XCAM simulators and the ARChitect configuration tool.

Synopsys 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.

Linux for ARC processors (ARC Linux) allows software developers to leverage the large amount of Linux-compatible application software to quickly build complex systems for the HS68 with open source components. ARC Linux is supported for single-core HS68 processors and for multicore HS68 processor implementations (SMP Linux), and is built and tested with the ARC GNU Tool Chain.

Compile	MetaWare Compiler	<ul style="list-style-type: none"> • 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 Compiler	<ul style="list-style-type: none"> • Freely access an open source solution with the GCC compiler
Debug	MetaWare Debugger	<ul style="list-style-type: none"> • Easily debug multiple targets with the same user interface • Quickly profile hotspots in your code • Use scripting to increase productivity
	JTAG Debuggers	<ul style="list-style-type: none"> • Efficiently bring-up hardware with tools from Ashling, Green Hills and Lauterbach
	GNU GDB Debugger	<ul style="list-style-type: none"> • Use the open source GDB debugger to debug real and simulated targets
Deploy	nSIM Simulator	<ul style="list-style-type: none"> • Develop & debug software before hardware is available • Simulate large programs with very fast ARC models
	NCAM Simulator	<ul style="list-style-type: none"> • Quickly optimize your software with near cycle-accurate simulation
	Zephyr Real Time OS	<ul style="list-style-type: none"> • Open source RTOS optimized by Synopsys for ARC processors
	Linux, SMP Linux	<ul style="list-style-type: none"> • Linux for ARC Processors provides all of the benefits of open source software, including complete source code and a large installed base

Table 1: Synopsys ARC HS66 and HS68 software and development tools

Interrupts and Exceptions

The HS66 and HS68 support up to 240 interrupts and exceptions with 16 nested levels of priority. Designers can select any number of interrupts up to 240 at build time. The interrupts are serviced through a jump table that allows higher flexibility in the location and implementation of the interrupt vectors. Interrupts can be triggered by hardware or software.

128-Bit Load and Store

The processors feature 128-bit load double and store double instructions that can be optionally included in the processor at build time. These are single instructions that load or store 128-bits of data to and from register pairs. There is no additional cycle penalty due to the wider and banked DCCM that support non-aligned loads and stores.

Fast Context Switch

For applications that require a deterministic and short latency context switch, the HS66 and HS68 processors can be configured by the user with up to eight register files. This is a build-time option and allows a fast context switch without the need to save processor registers to the stack.

ARC Processor EXTension (APEX) Interface

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 HS 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.

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 HS 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 32 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, 8 MB, 16 MB, 32 MB and 64 MB, 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 64 MB. 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.

Optional Features (separately Licensed)

- ARC Trace Interface provides real-time trace debugging features for the ARC HS processor
- The FPU supports half, single- and double-precision IEEE 754-compliant math operations
- Cluster DMA programmable memory access controller

Documentation

The following documentation is available for Synopsys ARC HS6x processors:

- ARCV3 Programmers Reference
- ARC HS6x Databook
- ARC HS6x Integration Guide
- ARC APEX Databook

Testing, Compliance and Quality

Verification of ARC 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 ARCV3 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

Synopsys ARC HS66 and HS68 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 Synopsys IP

Synopsys is a leading provider of high-quality, silicon-proven semiconductor IP solutions for SoC designs. The broad Synopsys IP portfolio includes [logic libraries](#), embedded memories, [analog IP](#), wired and wireless [interface IP](#), [security IP](#), [embedded processors](#), and [subsystems](#). To accelerate IP integration, software development, and silicon bring-up, Synopsys' [IIP Accelerated initiative](#) provides architecture design expertise, pre-verified and customizable IP subsystems, hardening, and signal/power integrity analysis. Synopsys' extensive investment in IP quality, comprehensive technical support and robust IP development methodology enables designers to reduce integration risk and accelerate time-to-market.

For more information on DesignWare IP, visit [synopsys.com/ip](https://www.synopsys.com/ip).