

ARC VPX DSP IP Family

Highlights

- Four-way VLIW combining scalar and vector operations
- 128-bit, 256-bit, 512-bit, and 1024-bit vector word lengths
- 8, 16, and 32-bit integer SIMD engines
- IEEE 754-compliant vector floating point unit option offers single-precision or half-precision operations and advanced math functions
 - Dual vector floating point SIMD pipes
 - Math accelerator
 - FFT accelerator
- Optional data compression support for microscaling OCP-MX data types
- Single- and multicore-configured offerings
- MetaWare tools include an auto-vectorizing C/C++ compiler, debugger, simulator, vector DSP and vector linear algebra libraries, vision processing library, and neural network SDK

Target Applications

- Automotive driver assistance systems (ADAS), including RADAR, and LiDAR processing
- Vision processing
- Sensor fusion
- Automotive powertrain and engine management
- Voice/speech recognition
- Natural language processing
- AI/Machine learning
- Smart home / IoT

Overview

MIPS ARC® VPX DSP IP is a family of VLIW/SIMD processors targeting a broad range of signal processing applications, from always-on devices to automotive ADAS to communications and high-performance computing.

The VPX family supports multiple vector lengths and core configurations:

- 128-bit vector word-VPX2 (single core), VPX2x2 (dual core)
- 256-bit vector word-VPX3 (single core), VPX3x2 (dual core)
- 512-bit vector word-VPX5 (single core), VPX5x2 (dual core), VPX5x4 (quad core)
- 1024-bit vector word-VPX6 (single core), VPX6x2 (dual core), VPX6x4 (quad core)

The VPX Family also includes DSPs with integrated hardware safety features such as ECC protection for memories and interfaces, safety monitors and lockstep mechanisms that achieve the most stringent levels of ISO 26262 functional safety compliance.

To speed application software development, the VPX Family is supported by MIPS' ARC MetaWare tools, which provide a comprehensive and vectorlength agnostic software programming environment that enables code portability among all members of the VPX Family. The tool suite includes an optimizing C/C++ vector compiler, debugger, instruction set simulator, as well as vector-based DSP, vision processing, machine learning inference and linear algebra libraries.

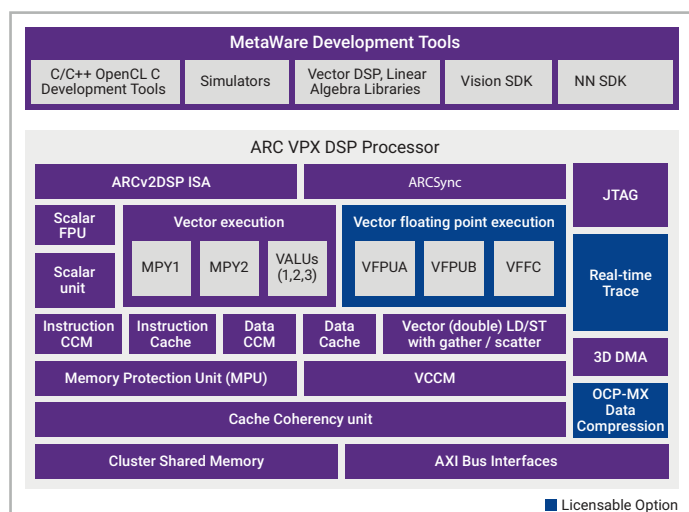


Figure 1: MIPS VPX DSP IP Block Diagram

The VPX DSP Hardware Features

VPX DSP processors feature a 4-way VLIW architecture optimally balanced to achieve high performance with low power consumption. Each VPX DSP core integrates a high-performance 32-bit scalar pipeline and a multi-slot vector processing unit supporting 8-bit, 16-bit, and 32-bit SIMD computations. Each VPX DSP core can execute one scalar and three vector instructions per cycle. The VPX DSPs are supported by configurable instruction and data caches for scalar operations and vector closely-coupled memory (VCCM) with single cycle access for vector processing. Like all ARC processors, the VPX DSPs are highly scalable and configurable, enabling users to tailor them to meet specific performance-power-area (PPA) requirements.

Each ARC VPX DSP core has a scalar register file of up to 32, 32-bit registers. This register file can be constructed from fast, single-cycle access memory or flip-flops, and supports one or two write ports and two read ports. Each VPX core also includes a vector register file with up to 64 registers of varying widths: 1024-bit for VPX6, 512-bit for VPX5, 256-bit for VPX3 or 128-bit for VPX2.

The VPX DSPs feature scalar 64-bit load double and store double instructions that can be included in the processor at build time. These are single instructions that load or store 64 bits of data to and from register pairs. There is no additional cycle penalty due to the wider and banked data closely-coupled memory (DCCM) that supports non-aligned loads and stores. In addition, the VPX DSPs support vector load double and vector store double instructions that can transmit data up to twice the vector word length.

The VPX DSPs' separate instruction and data L1 caches can be independently configured for 4K, 8K, 16K, 32K, or 64KB sizes.

The I-cache supports 2- and 4-way associativity, and a 32-, 64-, and 128-byte line size. The caches can be individually configured to support line locking and invalidate, as well as debug visibility. The data cache implements the MOESI protocol and supports cache-to-cache transfers. The ARC VPX DSP processors support 8KB to 128KB of DCCM, and 32KB to 2048KB of vector closely-coupled memory (VCCM) for the vector units. The CCMs are implemented as separate memory spaces and can be accessed every clock cycle.

The VPX DSPs provide an integrated 3-dimensional DMA for flexibly moving data between memories, including external system memory and internal memories like VCCM, DCCM, ICCM and CSM. The DMA supports 1D, 2D and 3D (block) transfers and I/O coherency, and has a configurable number of interfaces to the AXI bus. Together with an integrated coherency unit, optional clustered shared memory, and hardware support for inter-core communication (semaphores, interrupts, message passing), the DMA provides seamless support of integrated VPX multicore configurations. As a licensable option, all VPX family members can process data stored in the OCP-MX microscaling format, with data conversions handled by the DMA on the fly.

The ARC VPX vector lengths are configurable up to 1024-bit (VPX6), 512-bit (VPX5), 256-bit (VPX3) or 128-bit (VPX2) with vector support for dual issue of 8x8, 16x16 or 32x32 MAC processing. The DSP SIMD capability can deliver up to 256 8-bit, 128 16-bit, 64 32-bit MACs per cycle. The VPX DSPs support full predication and have scatter/gather instructions and hardware to maximize performance.

Each VPX DSP core can be configured to have up to three parallel floating-point processing pipelines, including two optional IEEE-754 compliant vector floating point units that supports both full- (32-bit) and half- (16-bit) precision floating point operations, with dedicated support for Radix2 and Radix4 for FFT acceleration. The VPX cores also have the option to add a dedicated vector floating point pipe that accelerates an extensive set of algebra math functions including $\text{div}(x)$, $\text{sqrt}(x)$, $1/\text{sqrt}(x)$, $\cos(x)$, $\sin(x)$, $\log_2(x)$, 2^x , e^x .

Cluster Shared Memory

Low-latency shared data memory is optionally available to support information passing and coordination between the multiple DSP cores. This memory is used as a software-managed scratch pad and is configurable from 0 to 8MB. To allow for larger sizes, the memory is internally multi-banked. This memory includes arbitration to support concurrent access from the DSP cores in a multicore configuration.

Real-Time Trace

The optional ARC Real-Time Trace (RTT) unit helps trace executed instructions or program flow and data. ARC RTT generates Nexus 5001 class 3-compliant trace messages. The RTT system can be set up in many different configurations by including the trace generator in the core and the RTT module at build time. ARC RTT can support on- and off-chip memory setups to suit application trace requirements.

Bus Interfaces

The VPX DSP processors have native support for the Arm® AMBA® AXI™ bus protocol. There are multiple AXI bus interfaces and supported bit widths to improve system throughput.

- Standard interface: 64-bit or 128-bit
- Vector direct memory interfaces (DMI): VPX2 supports 128-bit, VPX3 supports 128-bit or 256-bit, VPX5 and VPX6 support 128-bit, 256-bit or 512-bit. Up to 4 DMI interfaces are supported in parallel
- DMA interface: VPX2 supports 4x128-bit, VPX3 supports 4x128-bit or 4x256-bit, VPX5 and VPX6 support 4x128-bit, 4x256-bit or 4x512-bit. Up to 4 DMA channels are supported in parallel

SoC Integration

The VPX DSPs are designed for seamless integration into an SoC. They can be used with any host processor and operate in parallel with the host. The VPX processors include support for synchronization with the host through message passing and interrupts.

In addition, part of the VPX memory map can be made visible to other processors.

VPXxFS for Functional Safety and Security

For automotive and other applications requiring functional safety, the VPXxFS portfolio (VPX2FS, VPX3FS, VPX5FS, and VPX6FS) meet random fault detection and systematic functional safety development flow requirements to achieve ISO 26262 compliance. The ARC VPXxFS DSPs integrate hardware safety features such as ECC protection for memories and interfaces, safety monitors and lockstep mechanisms. A comprehensive set of safety documentation helps automotive designers achieve ISO 26262 functional safety certification.

To address the growing concern about malicious cyberattacks, automotive designs must comply to latest cybersesecurity regulations. VPXxFS processors are SAE / ISO 21343 Automotive CyberSecurity compliant.

For more information on VPX DSPs for functional safety and security, see the separate VPXxFS datasheet.

Comprehensive Software Environment

The ARC MetaWare Development Toolkit is a complete set of tools and libraries that provides everything needed to program the VPX DSPs. The MetaWare Toolkit includes a C/C++ compiler, a debugger, and an instruction set simulator (ISS). It comes with a rich set of compute libraries including DSP, linear algebra, vision processing and machine-learning interference. SPEED, a light-weight runtime library, provides drivers, execution services, inter processor communication support using a modular approach that allows for deterministic execution and minimal runtime overhead. All functionality is contained in one comprehensive suite of tools and libraries to deliver a high productivity software development environment and vector-length agnostic programming model allowing for smooth migration of software among members of the VPX Family.

Compilers and Debuggers	MetaWare C/C++ Compiler with auto vectorization, MetaWare Debugger	<ul style="list-style-type: none"> • Develop highly optimized code with an efficient C/C++ compiler for vector and scalar processors • Debug code with a comprehensive source-level debugger and profiler
Simulators	Fast nSIM Instruction Set Simulator	<ul style="list-style-type: none"> • Use a fast simulator to develop and debug DSP algorithm software before hardware is available
Compute Libraries	Vector DSP Library, Vector Linear Algebra Library, Vision Library	<ul style="list-style-type: none"> • Vector-length agnostic libraries provided in source code for optimization to specific core configurations • Linear algebra functions • Machine learning kernels • Computer vision kernels
Runtime Libraries	C/C++ runtime library SPEED runtime library	<ul style="list-style-type: none"> • Accelerate application software development using a modular, lightweight runtime library supporting multicore systems • Include drivers, execution services, interprocessor communications • Enable rapid setup of heterogeneous multicore systems
Neural Network SDK	Software Development Kit with Neural Network Compiler	<ul style="list-style-type: none"> • Compiler performs automatic compilation and optimization of neural network models—no manual optimization required • Supports TensorFlow framework; other frameworks (e.g., PyTorch) are supported via ONNX exchange format

Table 1: ARC MetaWare Development Toolkit components

Documentation

The following documentation is available for the ARC VPX DSPs:

- VPX Release Notes
- VPX Getting Started guide
- VPX Databook
- ARCV2 Programmers Reference Manual
- Vector DSP, Vector Linear Algebra and Vision Library Reference Manuals, and Performance Databooks

Testing, compliance, and quality verification of the ARC VPX DSPs follow 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 ARCV2 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 VPX DSPs are delivered as Verilog HDL in the ARChitect IP Library. The HDL is configured and output from the ARChitect IP Configurator tool. 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.