

MIPS I-Class I6500-F Processor Core

CPU IP Designed for Safety Critical Systems in an Autonomous Age

Artificial Intelligence (AI) techniques are increasingly used to give products greater awareness of their environment and, from this, a greater ability to control and automate functionality. We are seeing the emergence of artificial neural networks across a wide range of systems, from consumer products to ADAS and autonomous driving to industrial applications and beyond.

With the adoption of AI techniques, the required level of computing performance is much higher, and this is being addressed by a combination of dedicated accelerators and general purpose CPU-based compute capability. High-performance multiprocessor systems are a must.

In areas such as the automotive and industrial markets, Functional safety is also critical. These systems must be designed from the system level with a high degree of redundancy. The first step is to ensure that the product implements the correct behavior by design, which is managed by rigorous QMS processes. Safety-critical products must also detect and respond to errors that can occur during operation. Systems must be designed to adhere to industry standards for functional safety including ISO 26262 for automotive and IEC 61508 for industrial applications.

The MIPS I6500-F is the newest IP core in MIPS CPU product line, extending the variety and scalability of “off-the-shelf” licensable cores based on the proven and respected MIPS64 architecture to address the functional safety and performance requirements of emerging autonomous applications.

MIPS I6500-F as a SEooC at ASIL B(D)

The MIPS I6500-F is developed as a Safety Element out of Context (SEooC) IP core which enables a range of customers to incorporate this IP within their systems supported by a comprehensive documentation package as part of their system-level safety package. Not only does this significantly shorten the system development

process and reduce risk, it improves the efficiency as MIPS is able to optimize the additional functional redundancy to meet the required safety level while minimizing the additional overhead.

The MIPS I6500-F SEooC package is designed for systems requiring the highest level of functional safety: ASIL D. To achieve this level, not every IP core needs to reach the ASIL D. Instead the system is decomposed into IP that each reach a level similar to ASIL B but with enhanced fault detection time and fault coverage.

The additional functional redundancy built into the MIPS I6500-F includes:

- ECC across memories
- Parity protection for buses
- Time-out protection for interfaces
- Support for logic BIST (LBIST) operation on reset and periodically during operation

Heterogeneous Inside and Out

In embedded applications the computational workload is naturally multi-threaded but each thread will require a specific performance level. With the emergence of AI, many of these computational threads are sufficiently specialized to benefit from dedicated computational elements or accelerators. The MIPS I6500-F is well positioned to enable these complex systems.

Multi-threading

The performance of a CPU depends on minimizing the latency to the system memory. Even with a cache hierarchy, the CPU still stalls while waiting for data and this is where multi-threading offers significant performance improvements by running additional threads during these times.

Multi-threading is a more area-efficient alternative to the use of additional cores and offers a typical 40% performance boost for the execution of two threads simultaneously instead of sequentially.

Heterogeneous Combinations of Cores

Threads with high performance requirements can be run as a single thread on a core and this core can be optimized to maximize this single-threaded performance with increased resources such as level 1 cache size and FPU/SIMD capabilities. Other threads can share other cores for greater efficiency while lower-performance threads can be run on cores that are optimized for low power consumption with independently controlled clock frequency and voltage.

The I6500-F allows any combination of core configurations within a single cluster to optimally align to the system needs.

Accelerator Integration

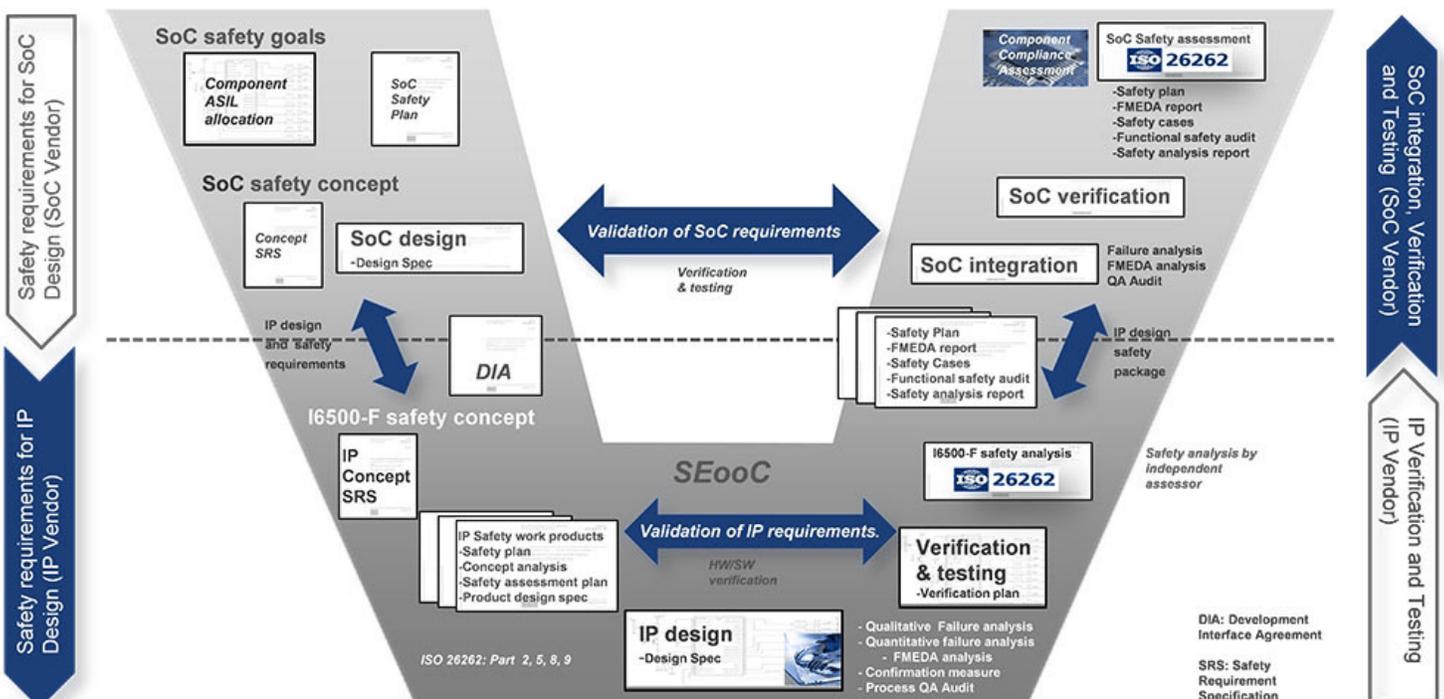
Specialist computational tasks such as artificial neural networks for AI achieve the highest performance with the greatest efficiency as dedicated accelerators. These accelerators need to work closely with the general purpose CPU cores to achieve the combination of flexibility from the CPU cores and efficient performance from the accelerators.

To enable this, the I6500-F provides very low communication latency through:

- Dedicated AXI auxiliary ports for direct communication to the accelerator control registers
- Shared Virtual Memory (SVM) with the accelerators so that data can be passed as pointers rather than through copying
- Low latency coherent access to memory via the I6500-F cluster level 2 cache using the IOCU ports
- Hardware cache coherency at the system level to allow high bandwidth traffic from the accelerators to directly access the system bus to maintain the performance of the CPU cores while retaining the benefits of SVM
- Multi-threading to enable threads to be dedicated to managing the operation of accelerators offering high efficiency with a zero context switch overhead

Performance Scaling

Hardware cache coherency at the system level allows combinations of heterogeneous I6500-F clusters and accelerators or other specialist computational IP to be integrated together to achieve whatever performance level is necessary for each system.



MIPS I-Class I6500-F Series Key Features/Benefits

- **SEooC ASIL B(D) package:** Rigorous QMS processes addressing systematic errors with optimized functional redundancy to meet system level ASIL D functional safety standards.
- **Additional functional safety packages:** Support for other functional safety market segments with IEC 61508 for industrial to follow.
- **Heterogeneous Inside:** In a single cluster, designers can optimize power consumption with the ability to configure each CPU with different combinations of threads, different cache sizes, different frequencies, and even different voltage levels. Optimized, low-latency shared virtual memory (SVM) operations with accelerators can be implemented through connecting via IOCU ports.
- **Heterogeneous Outside:** The latest MIPS Coherence Manager with an AMBA[®] ACE interface to popular ACE coherent fabric solutions such as those from Arteris and Netspeed lets designers mix on a chip configurations of processing clusters – including high bandwidth accelerator ports – for high system efficiency.
- **Simultaneous Multi-threading (SMT):** Based on a superscalar dual issue design implemented across generations of MIPS CPUs, this proven feature enables execution of multiple instructions from multiple threads every clock cycle, providing higher utilization and CPU efficiency.
- **Hardware virtualization (VZ):** The I6500-F builds on the real time hardware virtualization capability pioneered in the MIPS I6400 core. Designers can save costs by safely and securely consolidating multiple CPU cores with a single core, save power where multiple cores are required, and dynamically and deterministically allocate CPU bandwidth per application.
- **SMT + VZ:** The combination of SMT with VZ in the I6500 offers “zero context switching” for applications requiring real-time response; alongside the provision of scratchpad memory, this makes the I6500 ideal for applications which require deterministic code execution.
- **Ideal for compute intensive, data processing and networking applications:** The I6500 is designed for high-performance/high-efficiency data transfers to localized compute resources with data scratchpad memories per CPU, and features for fast path message/data passing between threads and cores.
- **Trusted:** MIPS multi-domain security technology enables isolation of applications in trusted environments, providing a foundation for security by separation.
- **Straightforward software development:** The I6500-F is based on the mature MIPS ISA which is broadly supported in the development ecosystem by multiple vendors including a wide choice of compilers, debuggers, operating systems, hypervisors and application software all optimized for the MIPS ISA.

MIPS I-class I6500 Base Core Features

- 64-bit MIPS64[®] Release 6 Instruction Set Architecture
 - Proven, successful, well supported 64-bit architecture
 - Superset of MIPS32 – runs MIPS32 software directly
- Balanced, 9-stage, dual-issue pipeline with Simultaneous Multi-Threading (SMT)
 - Superscalar on a single thread or two threads simultaneously per cycle
 - Up to four threads per core
 - Instruction bonding – merges sequential integer or floating point loads or stores into one operation for up to 2x increase on memory-intensive data movement routines
- High-performance dual-issue FPU/SIMD Unit – optional
 - 32 x 128-bit register set, 128-bit loads/stores to/from SIMD unit
 - Native data types:
 - 8-/16-/32-/64-bit integer and fixed point, 16-/32-/64-bit floating point
 - IEEE-754 2008 compliant
- Full hardware virtualization
 - Provides root and guest privilege levels for kernel and user space
 - Supports multiple guests, with full virtual CPU per guest = guest OSs run unmodified
 - Separate TLBs, COP0 contexts for root and guests → full isolation, fast context switching, exception and interrupt handling by root
 - Complete SoC virtualization support (IOMMU and interrupt handling – see multi-core features)
- L1 cache.
 - Instruction and Data of 32 KB or 64 KB each with ECC, 4-way set associative
- Data ScratchPad RAM (D-SPRAM)
 - Up to 1 MB with ECC, for deterministic low latency access and/or high performance data processing and movement outside of standard cached memory hierarchy (e.g. DMA directly into a core’s local D-SPRAM)
- Programmable Memory Management Unit (MMU)
 - First and second level TLBs with arrays for variable and fixed page size support

MIPS I-class I6500 Series Multi-Core & Multi-Cluster Features

- Coherent multi-core and multi-cluster platform, providing extensible implementations in support of both homogeneous and heterogeneous computing applications
 - Flexibility on the mix of cores and I/O coherency unit (IOCU) ports enables compute and throughput optimization to deliver better heterogeneous performance to application needs
 - Support for multi-cluster implementations of up to 64 compute clusters
 - IP available as:
 - Single cluster IP deliverable for use in combination with coherent fabric alternatives (ACE-compatible) for multi-cluster scalability, or
 - Complete multi-cluster sub-system deliverable
 - Per cluster multi-core system designed for maximum cluster-level bandwidth
 - Coherence Manager (CMv3.5)
 - Extensible to coherent multi-cluster implementations
 - Within a single cluster, supports multi-port configurations of up to:
 - Six cores in a single cluster (plus up to two hardware I/O coherency unit) IOCU ports, or
 - Eight IOCU ports for “clustering” hardware accelerators (even without a CPU core on the same cluster)
 - New directory-based coherency scheme – improves power consumption, performance and scalability
 - High-bandwidth 256-bit internal data paths and external system interface
 - Integrated L2 cache (L2\$): 16-way set associative, up to 8MB of memory
 - Dual pipelines for maximizing bandwidth on L1\$ misses
 - ECC option on L2\$ RAM for higher data reliability
 - Configurable wait states to RAM for optimal L2\$ design
 - L2\$ hardware pre-fetch for higher throughput and performance
 - Up to four auxiliary AXI ports provide for enabling features such as:
 - Separate path for non-coherent memory transactions
 - Shared access to low latency peripherals
 - Shared access to low latency and deterministic SPRAM (within a cluster, or even across clusters)
 - Inter-Thread Communication (ITC)
 - Fast path, higher efficiency alternative for messaging/data passing between threads within a core or a cluster
 - Global interrupt controller (GIC) with 256-interrupts per cluster
 - Advanced power management
 - Core-level DVFS (dynamic voltage and frequency scaling) – each core can be run at independent clock and voltage level
 - Virtualization support at system and SoC level
 - Up to 31 guest execution environments per cluster
 - IOCU include I/O MMU; GIC has virtualized interrupts
 - Guest ID brought out on system i/f for integration into multi-cluster and virtualized SoC designs
 - Advanced debug capabilities – Debug and Trace
 - Debug unit (DBU) supporting JTAG or APB i/f for Coresight™ compatibility
 - Program and Data Trace (PDtrace™), with on-chip or off chip trace buffering

About Wave Computing

Wave Computing, Inc. is revolutionizing AI with its dataflow-based systems and solutions that deliver orders of magnitude performance improvements over legacy architectures. The company's vision is to bring deep learning to customers' data wherever it may be—from the edge to the datacenter—helping accelerate time-to-insight. Wave is powering the next generation of AI by combining its dataflow architecture with its MIPS embedded RISC multithreaded CPU cores and IP. Wave Computing has been named Frost & Sullivan's 2018 “Machine Learning Industry Technology Innovation Leader,” and recognized by CIO Application Magazine's as one of the “Top 25 Artificial Intelligence Providers.” Combined with MIPS, Wave now has over 400 granted and pending patents and hundreds of customers worldwide.

Wave Computing, the Wave Computing logo and MIPS are trademarks of Wave Computing, Inc. All other trademarks are used for identification purposes only and are the property of their respective owners.