RACE++ SERIES

Computer Systems, Inc.
MERCURY

500 MHz PowerPC 7410 Daughtercard

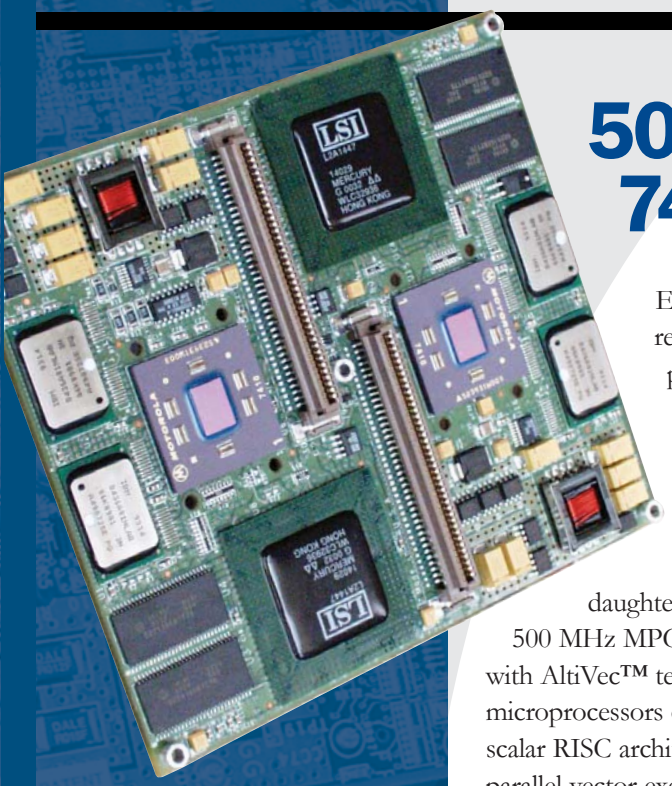

AltiVec Parallel Vector Technology

Scalable from Two to Hundreds of Processors

1K CFFT in 20 µs on Each Processor

125 MHz Memory System with Prefetch and ECC

Advanced DMA Engine for Chained Submatrix Moves

267 MB/s RACE++ Switch Fabric Interconnect

Fast L2 Cache (250 MHz)

mc.com

Embedded computing reaches a new level of performance with RACE++® Series PowerPC® 7410 daughtercards from Mercury Computer Systems.

Each PowerPC 7410 daughtercard contains two 500 MHz MPC7410 microprocessors with AltiVec™ technology. These unique microprocessors combine a modern superscalar RISC architecture with an AltiVec parallel vector execution unit.

The AltiVec vector processing unit revolutionizes the performance of computationally intensive applications such as image and signal processing. Each vector unit can operate in parallel on up to four floating-point numbers or up to sixteen 8-bit integers. This dramatically accelerates vector arithmetic and provides greater application performance on smaller, less power-hungry processors.

AltiVec technology also represents a leap in simplifying the programming required to achieve high performance. Whereas previous DSP-based systems required handcrafted assembly language code for optimal performance, easy-to-use extensions to the C language provide a direct mapping to AltiVec instructions. This permits developers to program more productively in a higher-level language, even for critical sections of code.

Optimized Performance

With the huge increase in processing performance brought by AltiVec, most applications are no longer CPU-limited. To keep the processor fed with ample data, increased emphasis is placed on the memory system and communications fabric that delivers data to the processor. Each compute node on the 500 MHz PowerPC 7410 daughtercard has a dedicated fabric interface at 267 MB/s and maximum memory speed of 125 MHz. By maximizing the performance of the memory and the fabric interface to the processor, Mercury has optimized RACE++ compute nodes for processing continuous streams of data.

AltiVec in RACE++ Computers

The computational power of RACE++ Series systems is built from compute nodes comprised of processors, memory, and interfaces to the RACE++ interconnect. PowerPC 7410 daughtercards each contain two compute nodes.

Each compute node (CN) consists of an MPC7410 microprocessor with AltiVec technology, level 2 (L2) cache, synchronous DRAM (SDRAM), and a Mercury-designed ASIC. This CN ASIC contains architectural advancements that enhance concurrency between arithmetic and I/O operations.

PowerPC 7410 Daughtercard Architecture

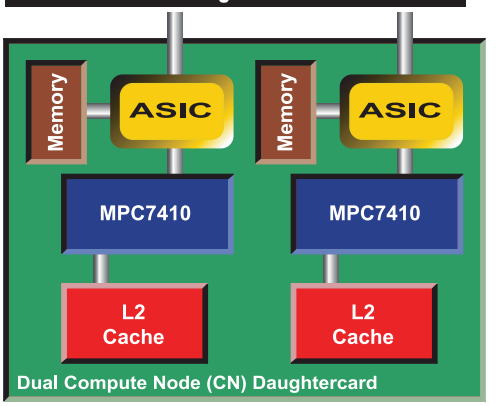

Mercury can configure systems with hundreds of compute nodes, communicating over the second-generation RACE++ switch fabric interconnect. Merging RACE++ and AltiVec technology provides embedded computers with unprecedented computational power.

AltiVec Vector Processing Unit

The AltiVec vector processing unit operates on 128 bits of data concurrently with the other PowerPC execution units. AltiVec instructions may be interleaved with other PowerPC instructions without any penalty such as a context switch. The 128-bit wide execution unit can be used to operate on four floating-point numbers, four 32-bit integers, eight 16-bit integers, or sixteen 8-bit integers simultaneously.

AltiVec instructions are carried out by one of two AltiVec sub-units. The Vector arithmetic logic unit handles the vector fixed-point and vector floating-point operations. Two floating-point operations are possible in a single cycle with the vector multiply-add instruction and the vector negative multiply-subtract instruction.

The Permute sub-unit incorporates a crossbar network to perform 16 individual byte moves in a single cycle. This capability can be used for simple tasks such as converting the "endian-ness" of data or for more complicated tasks such as byte interleaving, dynamic address alignment, or accelerating small look-up tables.

PowerPC RISC Architecture

In addition to the AltiVec execution unit, the MPC7410 contains a floating-point unit and two integer units that can operate concurrently with the AltiVec unit. Data and instructions are fed through two on-chip, 32-Kbyte, eight-way set-associative caches that enhance performance of both vector and scalar code.

Each PowerPC 7410 CN also includes a fully pipelined backside L2 cache operating at 250 MHz. This high-performance cache system provides quick access to data previously loaded from memory but too large to fit into the on-chip cache.

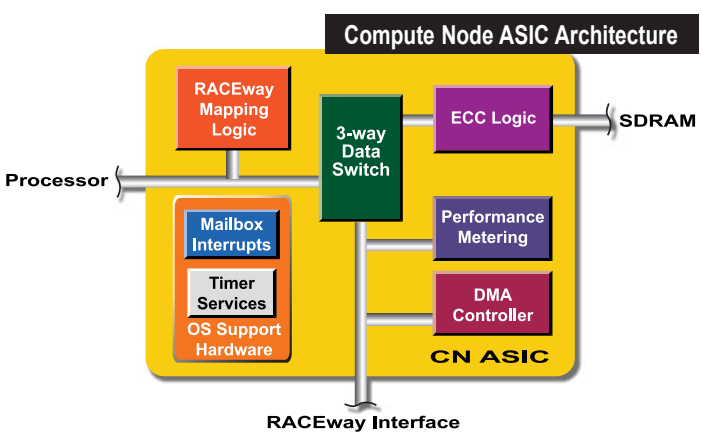

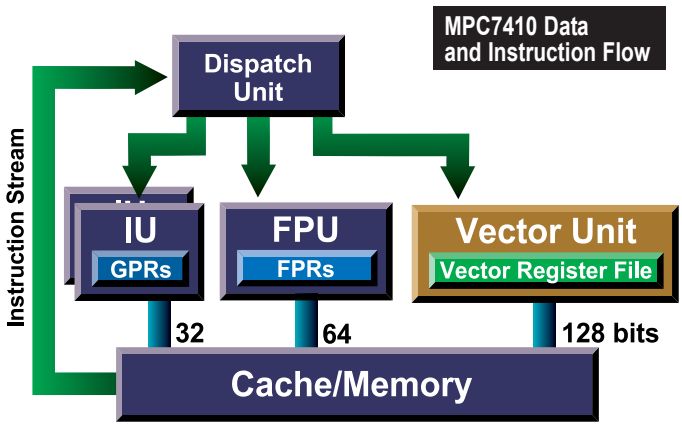

Compute Node ASIC

The CN ASIC, included in each compute node, acts as both a memory controller and as a network interface to the RACE++ switch fabric interconnect. The CN ASIC includes an enhanced DMA controller, a high-performance memory system with error checking and correcting, metering logic, and a RACE++ interface. By combining memory control and network interface into a single chip, Mercury's compute node provides the highest performance with the lowest power consumption and highest reliability.

High-Performance Memory System

Mercury's high-performance memory subsystem allows the memory to reach the intrinsic limits of its performance capability with:

125-MHz Synchronous DRAM

Prefetch Buffers: bring sequential data to the ASIC ahead of their explicit requests by the processor. These prefetch buffers greatly improve the performance of the CN in vector operations such as those used in DSP applications.

FIFO Buffers: efficiently overlap accesses to SDRAM from the local processor and the RACEway interconnect.

The PowerPC CN contains error-correcting circuitry for improved data integrity. One-bit errors are corrected on the fly, and multi-bit errors generate an interrupt error condition.

Enhanced DMA Controller

Each CN has an advanced DMA controller to support RACEway transfers at 267 MB/s with chaining and striding.

With chained DMA, the DMA controller works from a linked list in memory so that a complex chain of DMA requests requires no overhead from remote processors. DMA requests also support non-sequential access to and from local memory. This strided DMA capability enables high-performance submatrix transfers as required for distributed 2-D applications in image processing and synthetic aperture radar (SAR).

Performance Metering Logic

The CN ASIC also provides non-intrusive hardware performance metering facilities that enable application performance tuning and real-time data-flow analysis, essential for developing real-time multicomputer applications. A RACE++ node can be directed to collect performance metrics on memory bandwidth utilization, RACEway crossbar port usage and contention, or memory contention. The information can be used to locate sub-optimal code sections during execution of the application.

RACE++ Switch Fabric Interconnect

The PowerPC 7410 daughtercard is fully compatible with the RACE++, the second generation of the RACE architecture. The PowerPC CN ASIC connects the processor and memory directly to the 66.66-MHz RACEway interconnect.

As the latest version of the RACE architecture, the RACE++ interconnect represents an architectural evolution that includes increased communication speed, more richly connected topologies, and augmented adaptive routing. Together these enhancements represent a performance revolution yielding significantly higher bisection bandwidth and lower latency suitable for the most challenging real-time problems.

Crossbar-Connected Communication

All communication paths among compute nodes, memory nodes, and I/O nodes are connected using eight-port crossbars, both on-board and between boards. By providing these point-to-point connections, the RACE architecture eliminates the bottlenecks found in systems that use buses to connect processors together on a board or between boards.

Each node on a RACE++ daughtercard is connected to the RACEway switch fabric through an interface ASIC on the daughtercard. The RACE++ VME motherboard contains the first set of crossbars in the RACEway switch fabric. The second set of crossbars is contained on RACEway Interlink modules that connect multiple motherboards. This architecture assures each compute node direct and unblocked access to the RACEway network.

High Bandwidth, Low Latency

Each connection through a RACE++ crossbar can run at 66.66 MHz for a peak bandwidth of 267 MB/s. Each crossbar can connect four simultaneous communication paths for a total peak bandwidth over 1 GB/s.

Low latency is often as important, if not more important, than high bandwidth. When making a connection through a RACE++ system at 66.66 MHz, each crossbar along the data transfer path adds only 75 ns to the latency. Once the connection is established, each crossbar adds only 15 ns of latency.

Priority-Based Communication

Real-time applications demand more than speed alone. The priority-based communication of RACEway is designed to minimize the worst case latency for selected communications.

Adaptive Routing

The RACE++ crossbars adaptively route data using alternate path selections to actively avoid any congestion in the network without user intervention.

PowerPC 7410 Daughtercard

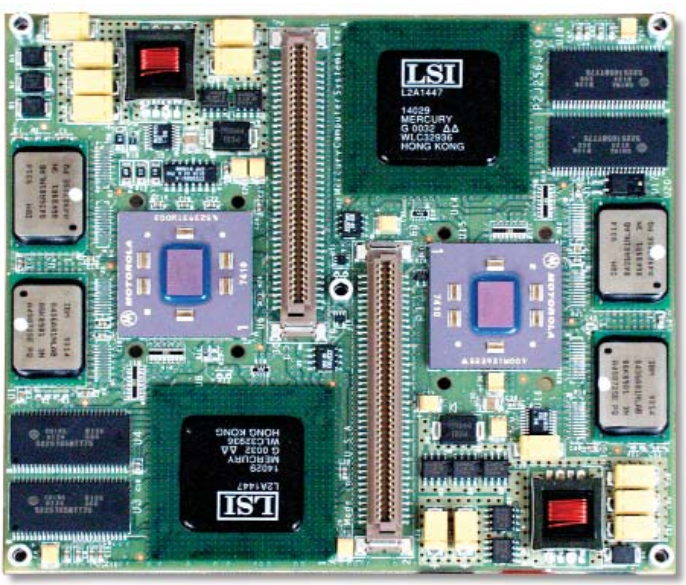

RACEway Broadcast Support

The PowerPC compute node supports RACEway crossbar broadcast capabilities for enhanced bandwidth utilization when sending the same data to multiple nodes. This hardware broadcast also reduces latency compared to software-based schemes that emulate a hardware broadcast by repetitively sending the same data in separate transfers.

As the initiator of a broadcast, either the PowerPC processor or the DMA controller may use any of the broadcast or multicast modes provided by the RACEway crossbar. As a receiver of broadcast data, the PowerPC CN has the programmable option of accepting or ignoring the data.

Software

Mercury also provides programming tools that simplify development and testing of real-time applications. The RACE++ Series MULTI® Integrated Development Environment (IDE) enables programmers to use familiar, mainstream tools to develop real-time processing routines. The MULTI IDE's integrated debug, program build, profiling, and version control facilities streamline development of complex multiprocessor software.

Mercury's Scientific Algorithm Library (SAL) provides the industry's easiest method of accessing the full power of the AltiVec vector unit. The SAL contains hundreds of image and signal processing functions that are optimized for AltiVec execution. These high-performance algorithms can also be accessed through the industry-standard VSIPL interface using Mercury's VSIPL-Lite signal processing library. For applications that require custom algorithms with demanding vector processing requirements, Mercury's C compiler includes the AltiVec C extensions that allow direct access to the AltiVec vector unit.

For large-scale programming, the Parallel Acceleration System (PAS™) supports data reorganization techniques for a wide variety of efficient and scalable programming models, by means of a high-level set of library calls. Scaling is achieved by changing runtime variables rather than recompiling the core code.

The TATL™ Trace Analysis Tool and Library from Mercury is also available to assist during software development and debugging. TATL is a minimally intrusive software logic analyzer that allows developers to visualize multiprocessor interactions. Using TATL, developers can spot bottlenecks, such as contention for shared resources or load imbalances.

Electrical/Mechanical Specifications

	P2J256J-R0C
RACEway ports	2
Processor frequency	500 MHz
Compute nodes	2
Memory frequency	125 MHz
SDRAM per CN	256 Mbytes
SDRAM per daughtercard	512 Mbytes
L2 cache frequency	250 MHz
L2 cache per CN	2 Mbytes
Weight	0.43 pounds*
Dimensions	5.0" x 4.435"
Power consumption**	20.8 W
Daughtercards per MCJ6***	2
Daughtercards per MCJ9	8

* Rugged version weighs an additional 0.01 pounds
** Maximum typical power consumption measured with concurrent FFTs and I/O
*** Requires 5-row connectors on MCJ6 and VME backplane

Commercial Environmental Specifications

Operating temperature	0°C to 40°C up to an altitude of 10,000 feet
(inlet air temperature at motherboard's recommended minimum airflow)	
Storage temperature	-40°C to 85°C
Relative humidity	10 to 90% (non-condensing)

As altitude increases, air density decreases, hence the cooling effect of a particular CFM rating decreases. Many manufacturers specify altitude and temperature ranges that are not simultaneous. Notice that the above operating temperature is specified simultaneously with an altitude. Different limits can be achieved by trading among altitude, temperature, performance, and airflow.
Contact Mercury for more information.

Computer Systems, Inc.
MERCURY
The Ultimate Performance Machine

For more information, go to www.mc.com

199 Riverneck Road
Chelmsford, MA 01824-2820 U.S.A.
978-256-1300 • Fax 978-256-3599
800-229-2006 • http://www.mc.com
NASDAQ: MRCY

 DS-6C-11

Nihon Mercury Computer Systems K.K.
No. 2 Gotanda Fujikoshi Bldg. 4F
5-23-1 Higashi Gotanda
Shinagawa-ku, Tokyo 141-0022
JAPAN
+81 3 3473 0140
Fax +81 3 3473 0141

Mercury Computer Systems, SARL
Immeuble le Montreal
19 bis, avenue du Quebec
Villebon
91951 COURTABOEUF CEDEX
FRANCE
+ 33 (0) 1 69 59 94 00
Fax + 33 (0) 69 86 10 875

Mercury Computer Systems, Ltd.
Campbell Court, Unit 19
Bramley, Tadley
HANTS RG26 5EG
UNITED KINGDOM
+ 44 1 256 880090
Fax + 44 1 25688 4004