

MCJ6 FCN Module

Exploit High-Performance FPGAs for Complex Signal and Image Processing

- FPGA compute nodes fully integrated in a RACE++[®] VME system
- Improves performance by a factor of 20 over RISC processors for some algorithms
- Partition applications between FPGAs and G4 processors for configuration flexibility
- Lower size, weight, cost, and power consumption
- FDK and support services significantly reduce development time and risks



The RACE++[®] Series MCJ6 FCN board from Mercury Computer Systems provides a flexible, manageable way to exploit the power of field programmable gate arrays (FPGAs) in RACE++ multicomputer systems. The MCJ6 FCN module is a 6U VME board with two Xilinx[®] Virtex[™]-II Pro P70 FPGAs. Each seven-million gate FPGA is connected to the RACE++ switch fabric via an onboard crossbar.

Each FPGA is supported by both SRAM and DRAM to maximize its effectiveness as a compute node. Together, an FPGA and its memory chips are referred to as an FPGA compute node (FCN).

High-Capacity I/O Connections

MCJ6 FCN modules provide full-duplex fiber-optic and copper serial connections for each FCN, as well as LVDS lines for parallel I/O, giving each board over 6 GB/s of direct I/O capacity. Delivering I/O directly to the FPGAs allows these devices to perform repetitive operations that reduce data volume before passing it on to the balance of the system. This feature permits handling of more I/O without increasing the system processor count.

Partition Applications Easily between FPGAs and PowerPCs

MCJ6 FCN boards can be configured in RACE++ systems with other VME boards, including boards carrying PowerPC[®] compute nodes and I/O devices, as well as other MCJ6 FCN boards. Because FPGAs in Mercury systems operate as a seamless element of the RACE++ environment, developers can partition their applications between performance-leveraging segments that run best on the FPGA and portions that can execute on easier-to-program PowerPC microprocessors.

Real-Time Reconfiguration for Mode Changes

Mercury's FPGA technology adds the versatility of nearly instant reconfiguration. High-speed reconfiguration facilitates dynamic, system-level changes in mission and operating mode. An onboard configuration manager can load new bitstreams into either FPGA upon user or application command in less than 100 ms. Bitstreams can be obtained over the RACE++ fabric from stored files or off-board Flash, optimally staged by off-board SDRAM.

Off-the-Shelf IP for RACE++ and Memory Interfaces

Through the FPGA Development Kit (FDK), Mercury delivers the MCJ6 FCN module with intellectual property for the RACE++ fabric interface, memory transfers, and I/O management. Users need only incorporate their application-related algorithmic firmware to create complete FPGA bitstreams. Users can create their own FPGA-based solutions with Mercury's help, or contract Mercury to develop the FPGA-resident portion of an application for a complete turnkey solution.




Figure 1. Mercury's scalable solution

FPGA Compute Nodes

Each 6U VME board contains two fully connected FPGA compute nodes. The heart of each FCN is a seven-million gate Xilinx Virtex-II Pro FPGA with its own memory, I/O, and RACE++ fabric connections. The 8 MB of QDR II SRAM in each FCN provides low-latency memory with peak access rates of 6.4 GB/s. Larger datasets can be staged in the 128 MB RDRAM II of each FCN. High-bandwidth data transfers are realized through the Mercury-provided memory controller IP.

The two FPGA compute nodes on each board can work together on the same dataset, communicating together over ten 2.5 Gbps serial links. This communication can be extended to multiple boards by creating an FPGA communication mesh with the four 2.5 Gbps copper connections to each FPGA available on the front panel.

Massive I/O at Your Command

Each FPGA on Mercury's MCJ6 FCN board has six 2.5 Gbps full-duplex fiber-optic interfaces and four 2.5 Gbps full-duplex copper serial connections. Sensor data can be delivered via fiber directly to the FPGAs, where it can undergo data reduction using the repetitive algorithms particularly suited to deployment on programmable logic devices. This savings in processor and interprocessor bandwidth reduces system size, cost, and complexity.

Each FCN also has a front-panel parallel I/O interface that supports 26 pairs of general-purpose LVDS lines. Designed for application-defined communications, this interface can be used to support parallel I/O to sensors or used between boards for direct FPGA-to-FPGA connections. With a cable length of 1 meter, these 26 pairs can run at 200 MHz. Cable lengths can be extended up to 10 meters, but the longer the cable, the slower the data transfer rate.

On a system-wide level, MCJ6 FCN modules can exchange data with RACE++ Series MYRIAD™ I/O devices on other boards via the RACEway Interlink switch fabric. Third-party I/O daughtercards can also be accommodated in the RACE++ system.

RACE++ System Connectivity

Each FCN has two connections to the RACE++ switch fabric via an onboard RACE++ crossbar. Full connectivity makes MCJ6 FCN boards part of a scalable system that can expand to provide as many FPGAs and PowerPC microprocessors as changing applications demand, with minimal application recoding and redeployment expense.

When data is required to travel to another board, such as a PowerPC board, the MCJ6 FCN system leverages the bandwidth, speed, and scalability of the RACE++ switch-fabric communications architecture to move data quickly and efficiently. Each FCN has two 267 MB/s connections to the module's RACE++ crossbar. These ports connect the board to the system-wide fabric, which can connect dozens of simultaneous communication paths.

Using the RACE++ switch fabric, multiple paths between most points in the fabric greatly reduce the chance of blocking or interruption. Further, because low latency is often as important, if not more important, than high bandwidth, 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.

Application Partitioning

Algorithms such as FFTs, fast convolutions, and pulse compression on incoming data streams can run up to 20 times faster on an FPGA than a RISC processor. However, algorithms whose functions are data-dependent are not well suited to implementation on an FPGA.

Mercury's FPGA solution implements an architecture that can combine FPGAs and PowerPCs in a RACE++ fabric. Developers can partition their application across FPGAs and PowerPCs for maximum effectiveness. Parts of the application that are simple, fixed-point computations can go on an FPGA, saving space, power, and money. Other parts of the application can go on the PowerPC, which is easier to program, so that overall development time is kept manageable.




Figure 2. MCJ6 FCN board architecture

Scalability

FPGAs in the RACE++ environment become part of a scalable system that can expand to provide as many FPGAs and PowerPC microprocessors as changing applications demand, with minimal application recoding and redeployment expense. Multiple MCJ6 FCN boards can be deployed in a single VME chassis, along with other boards carrying I/O devices and RISC processors, communicating via a RACE++ switch fabric with a bisection bandwidth of 2.1 GB/s.

6U VME systems from Mercury can scale to dozens of RISC and FPGA nodes, providing ample processing power for the most demanding image and digital signal processing applications. Developers can create and test algorithms on small laboratory systems consisting of only a few processors, with the assurance that the resulting code will move seamlessly to larger deployment platforms. Additionally, as processing requirements change in future program generations, they can readily resize target platforms with minimal impact to their code.

FPGA Development Software

Mercury provides FCN Developer's Kit (FDK) software to simplify and accelerate development of FPGA-based applications. The FDK suite of development tools contains off-the-shelf, ready-to-use IP components for managing input and output dataflows, memory transfers, and the RACE++ interface in FPGA applications. Also included is the RACE-on-Chip (RoC) framework, which functions as an on-chip communication mechanism between different IP modules on an FPGA. Users can join IP modules of the FDK with their own algorithm-specific modules to create a complete FPGA application bitstream without reinventing these elements.




Figure 3. Typical FCN-based board

In effect, the intellectual property of the FDK enables each FPGA compute node to operate as a fully functional RACE++ compute node, capable of reading and writing to local and remote memory locations across the RACEway switch fabric. This functionality frees developers to concentrate on coding "inner loops" for the FPGA platform, and provides them with interfaces for connecting their computational modules with the underlying RACE++ system. See Figure 3.

To jumpstart the creation of complete application solutions, a Mercury default bitstream is part of the FDK. When loaded into the FPGA, this bitstream enables system software to monitor system health through a series of tests, including memory test, DMA test, and an I/O loopback test. The Mercury default bitstream can also serve as an application example and is fully documented.

The FDK components are built for easy integration with the leading FPGA development tools available, including the Xilinx ISE, Synplicity Synplify®, and Mentor Graphics ModelSim®. The Xilinx ChipScope™ Pro logic analyzer can also be used while developing applications for the MCJ6 FCN for improved visibility of FPGA operations during debugging. See Figure 4.

Mercury also provides a complete VHDL simulation environment or "harness" that models the FPGA compute node. This environment provides a bus functional model (BFM) for RACE++ communications, as well as for SRAM and DRAM attached to the FPGA. The simulation environment enables verification of FPGA applications prior to deployment on final-mission hardware and allows regression test suites to run with a single command.




Figure 4. FDK IP module architecture

Specifications

Module Specifications

FPGA compute nodes (FCNs)	2
FPGA processor per FCN	Xilinx XC2VP70-6
SRAM capacity per FCN	8 MB
SRAM bandwidth per FCN	3.2 GB/s full-duplex
DRAM capacity per FCN	128 MB
DRAM bandwidth per FCN	3.2 GB/s
RACE++ ports per FCN	2
Fiber links per FCN	6 at 2.5 Gbps, full-duplex
Copper serial links per FCN	4 at 2.5 Gbps off-board, full-duplex 10 at 2.5 Gbps to other onboard FCN, full-duplex

LVDS lines per FCN	26 pairs
--------------------	----------

Electrical/Mechanical Specifications

Input voltage	4.75 to 5.25 VDC 3.25 to 3.45 VDC
Power	100W max for the board 16A (80W) max at 5 VDC* 17A (56W) max at 3.3 VDC*
Dimensions	6U x 160 mm
Slot-to-slot spacing	0.8 in
*The board does not access both of these maximum values simultaneously. The actual consumption of the module depends on the FPGA application. These numbers represent the maximum consumptions for each module.	

Environmental Specifications

Minimum airflow (per slot)	12 CFM
Temperature	
Operating*	0°C to 40°C up to 10,000 ft (inlet air temperature at minimum airflow)
Storage	-40°C to +85°C
Relative humidity	10-90% (non-condensing)

*As altitude increases, air density decreases, hence the cooling effect of a particular number of CFM decreases. The operating temperature is specified simultaneously with an altitude, because different limits can be achieved by trading among altitude, temperature, performance, and airflow. Contact Mercury for more information.

RACE++ is a registered trademark, and MYRIAD and Challenges Drive Innovation are trademarks of Mercury Computer Systems, Inc. Other products mentioned may be trademarks or registered trademarks of their respective holders. Mercury Computer Systems, Inc. believes this information is accurate as of its publication date and is not responsible for any inadvertent errors. The information contained herein is subject to change without notice.

Copyright © 2006 Mercury Computer Systems, Inc.

712.00E-0606-DS-mcj6fcn



Corporate Headquarters

199 Riverneck Road
Chelmsford, MA 01824-2820 USA
+1 (978) 967-1401 • +1 (866) 627-6951
Fax +1 (978) 256-3599
www.mc.com

Worldwide Locations

Mercury Computer Systems has R&D, support and sales locations in France, Germany, Japan, the United Kingdom and the United States.

For office locations and contact information, please call the corporate headquarters or visit our Web site at www.mc.com.