Ensemble 7100 Platform FCN7300 Module

Exploit the Power of FPGAs in Ensemble 7100 Platforms

- Three fully integrated Virtex™-5 LX330T FPGA compute nodes
- Freescale™ MPC8640 compute node
- Flexible, high-capacity I/O connections
- Built-in IP for serial RapidIO and memory control

The FCN7300 FPGA Compute Node (FCN) module from Mercury Computer Systems provides a flexible, manageable way to exploit the power of field-programmable gate arrays (FPGAs) in Ensemble 7100 multicomputer platforms. The FCN7300 module provides massive I/O as part of a scalable system that can expand with minimal application recoding and expense. Because some algorithms can run up to 20 times faster on an FPGA than a RISC processor, developers can partition applications across FPGAs and RISC processors for maximum effectiveness.

The FCN7300 module is an M155-format board with three Xilinx® Virtex™-5 XC5VLX330T FPGAs. Each 7-million gate FPGA is connected to the RapidIO® switch fabric via an onboard serial RapidIO crossbar. Each FPGA receives support from both local QDRII SRAM and local DDR2 SDRAM to maximize its effectiveness as a compute node. Together, an FPGA and its memory chips are referred to as an FPGA compute node (FCN).

FPGA Compute Nodes

The heart of each FCN7300 module is a 7-million gate Xilinx Virtex-5 FPGA with its own memory, I/O, and RapidIO fabric connections. Each FCN7300 module includes four QDRII SRAM memory controllers, each with 9 MB of SRAM. Each controller supports simultaneous reads and writes, and has read and write datapaths that are each 18 bits wide at the physical connection between the memory and the FPGA. These 18-bit-wide datapaths operate at 400 Mb/s per wire for a total of 900 MB/s of read bandwidth and 900 MB/s of write bandwidth. The combined SRAM on the four controllers gives a total of 36 MB/s of memory with a peak bandwidth of 7.2 GB/s (4 controllers x 2 directions x 900 MB/s).

Larger datasets can be staged in the DDR2 SDRAM. Each FCN7300 module includes four DDR2 SDRAM memory controllers, each with 256 MB of SDRAM. Each controller has a bi-directional datapath that is 32 bits wide at the physical connection between the memory and the FPGA. These 32-bit-wide datapaths operate at 400 Mb/s per wire for a total of 1.6 MB/s per controller. The combined SDRAM on the four controllers gives a total of 1 GB of memory with a peak bandwidth of 6.4 GB/s.

High-bandwidth data transfers are realized through the Mercury-provided memory controller IP. The IP provides a datapath to the application that is twice the width of the physical connection: 36 bits for the SRAM and 64 bits for the SDRAM.

The three FPGA compute nodes on each FCN7300 module can work together on the same data set, communicating via LVDS mesh links. Each of these three FPGA compute nodes has four 8-bit +clock/data wide LVDS links to the other two FPGA compute nodes , with each link running at 800 Mb/s per wire, giving 800 MB/s per link. With two links in each direction, there is a total of 1.6 GB/s in each direction or a total of 3.2 GB/s. Each FCN7300 module has a full-duplex connection to the RapidIO crossbar supporting 1.25 GB/s peak bandwidth, in each direction, to the serial RapidIO fabric. A SelectMAP interface provides connection to the Configuration Manager (CM) module that is implemented in a separate FPGA.

Communication can be extended to multiple boards by creating an FPGA communication mesh. Each of three FCN7300 FPGA compute nodes provides four full-duplex 3.125 Gbaud lanes (GTPs) to front-panel copper connections and eight full-duplex 3.125 Gbaud lanes to an optional backplane mesh for board-to-board meshing. The FCN7300 FPGA compute node (#2) provides four full-duplex 3.125 Gbaud lanes (GTPs) to the other two FCN7300 FPGA compute nodes, providing a full-duplex link of 1.25 GB/s, in each direction between FCN7300 FPGA compute node (#2) and each of the other FCN7300 FPGA nodes. This link is available for providing bandwidth in addition the LVDS links.

Freescale MPC8640 Compute Node

Each M155-format board contains a Freescale MPC8640, 1.067 GHz single-core processor compute node. The compute node contains the MPC8640, with two DDR2 SDRAM memory controllers. Each memory controller has an 8-byte physical path plus bits for ECC (error checking and correction). Each memory controller has 512 MB of memory for a total of 1 GB of memory. The datapaths between memory and the MPC8640 operate at 533 Mb/s, per wire, for a total bandwidth of 4.25 GB/s per controller, yielding a total of 8.5 GB/s. There is also 128 MB of application flash. The Freescale MPC8640 compute node is connected to the serial RapidIO fabric using the processor internal serial RapidIO endpoint.

There is a 4-lane serial RapidIO connection from the MPC8640 to the onboard serial RapidIO crossbar. This enables the MPC8640 to access the FCN7300 compute nodes directly. It is intended that the MPC8640 use these connections for setting up DMA chains for the DMA controllers in each of the FCN7300 compute nodes. These serial RapidIO connections can also be used to move data between the FCN7300 compute nodes and the MPC8640 compute node, for cases where the MPC8640 processor is performing some of the DSP processing.

The application can monitor junction temperature. If the application does not keep the MPC8640 junction temperature low enough on its own, the system software shuts down the processor to prevent damage.

Massive I/O at Your Command

Sensor data can be delivered directly to the FPGAs, where it can undergo data reduction using the massively parallel algorithms particularly suited to deployment on programmable logic devices. This saves processors and inter-processor bandwidth, thereby reducing system size, cost, and complexity.

The FCN7300 compute nodes provide substantial I/O beyond the twelve 3.125 Gbaud mesh links. Two of the FCN7300 compute nodes have their own 10-Gbps, full-duplex, fiber-optic connections to the front panel, providing massive bandwidth directly to and from the FCN7300 compute nodes. There are also two groups of 8 LVDS connections on the modules front panel from each of the FCN7300 compute nodes, for a total of 48 LVDS pairs. 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.

The two independent 10 Gbps serial fiber interfaces provide high-speed streaming I/O at the front panel. Each interface has a 10-Gbps fiber-optic SFP+ transceiver, LC-style fiber-optic connectors on the module front panel, with a management interface that allows register-based management and control of the data links.

Each 10-Gbps SFP+ optical module is connected to a PHY transceiver that translates the encoded 10 Gbps signals to four 3.125 Gbaud XAUI lanes connected to an FPGA compute node. Mercury provides an end-to-end example design as a starting point for a customer-supplied 10 Gigabit Ethernet implementation. The example design uses Xilinx's XAUI core to provide four bonded channels to the internal FPGA cores.

Additionally, the Mercury example provides I2C and MDIO interfaces for control and status monitoring of the SFP+ and PHY modules. These allow the user to exploit the loopback capabilities present in the 10-GbE onboard sub-system. Because these interfaces connect directly to the design's internal status and control registers, they are accessible from any node within the multicomputer for monitoring and diagnostic purposes. The Xilinx MAC core can be purchased and inserted into the example design, so that the customer does not have to design the interface from scratch.

Delivering I/O directly to the FPGAs allows these devices to perform repetitive operations that reduce data volume before passing the data on to the balance of the system. This feature maximizes computational density without increasing the system processor count. The onboard MPC8640 has a 100-Mb Ethernet port on the front panel (available with adapter module).

On a system-wide level, FCN7300 compute nodes in one Ensemble 7100 platform can exchange data with any serial RapidIO devices on other Ensemble 7100 platform's modules via the serial RapidIO switch fabric. Third-party I/O PMCs and XMCs can also be accommodated in the serial RapidIO system via the SFM7100 modules in the Ensemble 7100 platform.

RapidIO System Connectivity

Full connectivity makes FCN7300 modules 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 module, such as an HCD7500 module, the FCN7300 module leverages the Ensemble 7100 platform's bandwidth, speed, and scalability of the serial RapidIO switch fabric communications architecture to move data quickly and efficiently. The FCN7300 module includes an 8-port serial RapidIO crossbar, with each port capable of operating at a 1.25 GB/s burst data rate in each direction. Four of the crossbar ports are connected to the module's backplane connectors: one to each of the three FPGA nodes, and one to the onboard MPC8640 processor. These ports connect the board to the system-wide fabric, which can connect dozens of simultaneous communication paths. Using the serial RapidIO switch fabric, the multiple paths between most points in the fabric greatly reduce the chance of blocking or interruption.

Real-Time Reconfiguration for Mode Changes

Mercury's FPGA technology adds the versatility of rapid reconfiguration. High-speed reconfiguration facilitates dynamic, system-level changes in mission and operating mode. The Configuration Manager (CM) provides an interface to the local MPC8640 and MMC processors, allowing users to send commands to the CM and to read status information regarding the FCN7300 module. For reconfiguring the FPGA, the CM drives the SelectMAP interface to the FPGA, so that bitstream data sent across the RapidIO network can be written into the FPGA at a peak speed of 50 MB/s.

System Management

The Module Management Controller (MMC) implements the module status and control functions common to all modules in the Ensemble 7100 platform. These control functions include various types of reset and monitors for power supplies, temperature, configuration status, and application status. The MMC communicates module status information over IPMI to a central management processor on one of the SFM7180 modules.

Applications use the multicomputer status and control (MSC) service library to retrieve module status information connected by the central management processors and to forward commands through those processors to the modules in the chassis. The MSC service is a status and control interface to the switch card modules in the Ensemble 7100 multicomputer. The MSC allows applications running on the switch card processor nodes to access the status information from FCN7300 modules and to command the FCN7300 modules to change state.

Each FCN7300 module has a set of indicator lights and onboard sensors. These include a temperature sensor for each FCN7300 compute node, as well as current and voltage sensors that provide in-system power dissipation measurements for the FCN7300 module. The measurements for these sensors are available to applications running on the SFM7180 module processors through the MSC.

A library of application programming interfaces (APIs) allows applications to interact with the MSC service. The MSC library APIs related to the use of the FCN modules include the following functions:

- Enabling FCN7300 modules in the hardware configuration
- Reading the results of power-on self-test
- Controlling the basic state of the FCN7300 modules
- Reading and controlling sensors on the FCN7300 module

Partition Applications Easily Between FPGAs and PowerPCs

FCN7300 modules can be configured in serial RapidIO systems with other Ensemble 7100 boards, including boards carrying MPC8640 compute nodes, I/O devices, and other FCN7300 modules. Because FPGAs in Mercury systems operate as a seamless element of the RapidIO environment, developers can partition their application between performance-leveraging segments that run best on the FPGA and portions that can execute on easier-to-program PowerPC microprocessors.

Algorithms such as FFTs, fast convolutions, and pulse compression on incoming data streams can run up to 20 times faster on an FPGA than on a RISC processor. However, algorithms whose functions are data-dependent and conditional in nature are more suited for implementation on a PPC. Mercury's FPGA solution implements an architecture where FPGAs and PowerPCs are combined in a serial RapidIO fabric.

An application can be partitioned across FPGAs and PPCs for maximum effectiveness. Parts of the application that are simple, fixed-point computations can run on an FPGA, saving space, power, and money. Other parts of the application can run on the MPC8640, which is easier to program, so the overall development time is kept manageable while the performance is maximized.

The onboard MPC8640 processor allows better integration, communication, and control with the onboard FCN7300 modules through the Command Ring interface without interfering with the fabric communications. For example, DMA operations can be managed from the onboard MPC8640 processor without requiring additional processor boards in the system. In addition to that, the local PPC can be used to manage the loading of the bitstreams and diagnostics of the FCN7300 modules.

Standalone Operation

The FCN7300 module takes advantage of the Ensemble 7100 platform architecture to enable more efficient application development while leveraging platform costs across multiple application development efforts. Each FCN7300 module can run in standalone mode within an Ensemble 7100 platform. In this operational mode, the onboard MPC8640 processor can act as the host processor interfacing with the FPGAs through RapidIO and with external hosts through backplane Gigabit Ethernet or front-panel 100-MB Ethernet. This allows multiple application development teams to work concurrently and completely independent of each other within the same Ensemble 7100 platform to enable faster time to market. See example in diagram in Figure 1.

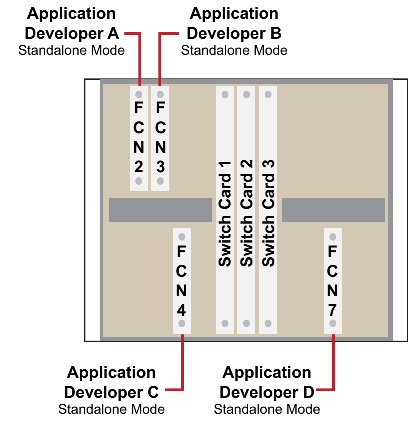

***Figure 1**. Standalone Operation*

Scalability

FPGAs in the RapidIO 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 FCN7300 modules can be deployed in a single Ensemble 7100 chassis, along with other modules carrying I/O devices and RISC processors communicating via a serial RapidIO switch fabric. FCN7300 modules can be installed in the processor slots of the Ensemble 7100 platform, including up to 24 modules in the full-size chassis.

Developers can create and test algorithms on small laboratory systems consisting of only a few processors, with the assurance that the resulting code can 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.

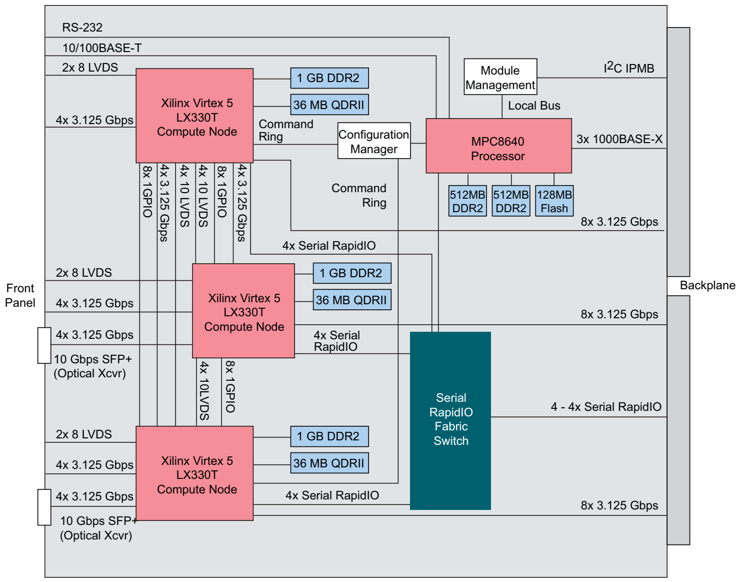

Figure 2. *FCN7300 block diagram*

Specifications

FPGA Compute Nodes

FPGA compute nodes	3
FPGA processor	Xilinx Virtex-5 XC5VLX330T
QDRII SRAM	
Capacity per FPGA	36 MB
Bandwidth per FPGA	7.2 GB/s full-duplex (peak)
DDR2 SDRAM	
Capacity per FPGA	1 GB
Bandwidth per FPGA	6.4 GB/s (peak)

Serial RapidIO ports per FPGA1

Fiber links

- 2 pairs at 10 Gbps each, full-duplex (1 on each of 2 FPGAs)
- 850 nm multi-mode fiber with LC via pluggable SFP+ module connectors

Copper serial links per FPGA

- 4 at 3.125 Gbaud from FCN2 to the other onboard FPGAs, full-duplex
- 4 at 3.125 Gbaud off board (12 total to front panel)
- 8 at 3.125 Gbaud off board (24 total to the backplane)

LVDS lines per FPGA

- 16 pairs to a front-panel connector (48 pairs total)

Freescale MPC8640 Processor

PowerPC Compute Nodes

- 1 MPC8640 processor, 1.067-GHz single-core

DRAM	
Memory capacity	1-GB DDR2 SDRAM, 533 MHz
Memory BW	8.5-GB/s peak
Flash	
Memory capacity	128-MB application flash
Ethernet port	100 MB to a front-panel connector
Serial port	RS-232 to a front-panel connector

Electrical/Mechanical

Input voltage	48.0 VDC +/-5%, main power
	5.0 VDC +/-5%, management power

Input voltages measured at the backplane pins inclusive of all ripple.
Mercury strongly recommends that system-level power designs use a ±2% margin to avoid any potential issues with respect to the system-level operating charateristics and operating environment.

Power 190W (high typical, depends largely on application IP)

Dimensions	M155 format
Slot spacing	1.057 in

Environmental

Minimum airflow (per slot at sea level)

- 28 CFM per stacked pair

Temperature*	
Operating	0°C to 35°C
Storage	-55°C to +85°C
Rate of change	5°C/min
Humidity	5-95% (non-condensing)
Vibration	
Random	0.02g²/Hz, based on 20 to 2,000 Hz, 1 hr/axis
Shock	50g in all axes, 11 ms, half-sine
Altitude*	
Operating	0 to 30,000 ft
Storage	0 to 50,000 ft
Salt/Fog	Consult factory

*As altitude increases, air density decreases and, therefore, the cooling effect of a particular number of CFM decreases. Different limits can be achieved by trading among temperature, altitude, frequency, and airflow.

2154.02E-0610-DS-fcn7300

MERCURY COMPUTER SYSTEMS™

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

Europe
Mercury Computer Systems, Ltd.
Unit 1 - Easter Park, Benyon Road • Silchester, Reading • RG7 2PQ UNITED KINGDOM
+ 44 0 1189 702050 • Fax + 44 0 1189 702321

Asia
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

Challenges Drive Innovation®