

Fibre Channel Solutions from Xilinx

Using the Virtex family and Xilinx Fibre Channel IP, you can build a complete FC solution running as fast as 4 Gbps.

by Douglas Grant
Staff Design Engineer
Xilinx, Inc.
douglas.grant@xilinx.com

The Fibre Channel (FC) standard enabled IT managers to scale compute and storage needs independently. Unlike protocols used for direct attach like SCSI, FC's switch-based architecture has enabled it to become a major player in the data center.

The advantages of FC include:

- High speed: 1 Gbps, 2 Gbps, 4 Gbps, 8 Gbps and 10 Gbps
- Low latency and high data integrity
- Upper-level protocol support: SCSI, FICON
- Distance support: >10 Km

Xilinx® technology has been used in the development of hubs, switches, and high-end enterprise storage arrays. In this article, I'll summarize FC protocol and discuss Xilinx LogiCORE™ IP solutions for Fibre Channel.

Principal Features of Fibre Channel

Figure 1 shows the relationship between Fibre Channel-defined layers FC-0 to FC-4 and the well-known ISO-OSI seven layer model. FC-0 deals with the physical transmission and reception of serial data, FC-1 refers to the encoding and synchronization, and FC-2 comprises framing and sequencing tasks and login protocols. The FC-3 layer has never been defined. FC-4 refers to the mapping of upper-layer protocols (ULPs) into Fibre Channel frames. An example of FC-4 functionality is the FCP standard, which maps SCSI commands into FC frames.




Figure 1 – OSI layers versus FC layers

Fibre Channel data is 32 bits wide at the user interface. These 32-bit words are encoded using 8b10b encoding into 40 bits in the serial domain and typically – although not exclusively – transmitted through optic fiber. Although 32-bit data words are by definition arbitrary values, Fibre Channel defines a set of so-called primitive signals, or ordered sets of specific byte values, that implement the initialization protocols. Examples of primitive signals include IDLE and RRDY.

Primitive sequences are defined as a sequence of three or more consecutive primitive signals such as LR (link reset) and LRR (link reset response). Fibre Channel also defines a number of SOF (start of frame) and EOF (end of frame) primitives that delineate the Fibre Channel frames.

Buffer-to-buffer credit is a Fibre Channel concept that allows the best use of the bandwidth available, enabling a Fibre Channel port to transmit a frame to another port only if there are receive buffers available at the other port.

Fibre Channel frames are constructed from an SOF, six header words, a payload of up to 2,112 bytes, and a cyclic redundancy check (CRC) and EOF word. Frames are separated by an inter-frame gap (IFG) of at least six so-called fill words, which are typically IDLE words but can

also include RRDYs and other credit management codes.

Virtex-II Pro and Virtex-4 FPGAs

Three main features of Virtex™-II Pro and Virtex-4 FX devices allow you to create entire Fibre Channel ports on a single die:

- Multi-gigabit transceivers (MGTs) allow the reception and transmission of serial data up to 11 Gbps. For Fibre Channel, 1 Gbps, 2 Gbps, and 4 Gbps are currently defined in the standard, with 8 Gbps just around the corner.
- The embedded PowerPC™ 405 core on both Virtex-II Pro and Virtex-4 FX devices is an ideal candidate for running Fibre Channel port software, which includes fabric and port login/logout, exchange of link parameters, and other FC-2 functions that lend themselves to implementation in firmware.
- The abundance of block RAMs on the Virtex-II Pro and Virtex-4 FX devices aids the processing of multiple frames at line rates. Any Fibre Channel port will require a number of receive buffers, and these may be implemented on-chip in block RAMs.

LogiCORE IP for Fibre Channel

Xilinx currently offers 1, 2, and 4 Gbps Fibre Channel cores in two flavors: point-to-point and arbitrated loop.

Point-to-Point LogiCORE Solutions

This core supports FC-1 and the framing and signaling part of FC-2, based on the ANSI INCITS FC-FS Final Draft v1.9 standard.

Included are the port state machine (PSM) and loss of sync state machine (LOS-FSM). All classes of Fibre Channel frames are supported. Buffer-to-buffer credit management is optionally supported, along with statistics gathering. Access to the control signals is through an optional register bank with a microprocessor-agnostic management interface or, alternatively, through dedicated configuration control and configuration status vectors.

You can use this core to create any of the following Fibre Channel ports: N, F, E, and B. Arbitrated loop support is not included in this core, nor are login protocols directly supported within – these are left for you to implement in either firmware or hardware as required.

Speed negotiation (Section 28 of FC-FS) is optionally supported in hardware for multi-speed cores. You can alternatively implement this in firmware.

Figure 2 shows the architecture of the dual-speed core.

Credit Management

The optional credit management block provides simple buffer-to-buffer credit management, keeping track of RRDY and SOF primitives received and sent and supporting BB_SCx-based credit recovery as defined in FC-FS Section 18.5.11.

MAC

The MAC block, designed to FC-FS Section 7, provides the main functionality of the core. This block comprises the port state machine (PSM) as well as the framing control and checking of the data.

Link Controller

The link controller block, designed to FC-FS Sections 5 and 6 (equivalent to FC-PH

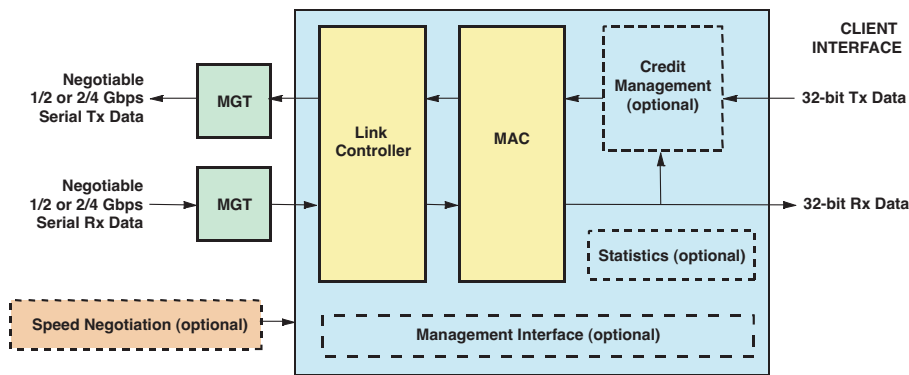


Figure 2 – Point-to-point FC LogiCORE architecture (dual speed on Virtex-II Pro FPGAs)

Parameter	Virtex-II Pro FPGA			Virtex-4 FPGA		
	Slices ¹	Total LUTs	FFs	Slices	Total LUTs	FFs
c_has_host=false						
c_has_bbcredit_mngmt=false	1,040	1,270	970	1,040	1,270	980
c_has_speed_neeg=false						
c_has_state=false						
c_has_host=true	+300	+390	+270	+300	+390	+280
c_has_stats=true ²	+220	+490	+260	+210	+470	+250
c_has_bbcredit_mngmt=true	+420	+770	+190	+400	+700	+190
c_has_speed_neg=true	+120	+210	+100	+110	+200	+100

¹ To discover the total resource requirements, begin with the numbers in the top row and add the numbers for the selected options.
² Only available if c_has_host option has also been selected.

Table 1 – Resource utilization for FC point-to-point core

Sections 11 and 12), provides word alignment and synchronization of incoming data. This block also provides CRC generation and checking on outgoing data.

Statistics

The optional statistics block collects and stores statistical information in the memory of the core. Statistics may still be collected outside of the core using the statistics vector.

Speed Negotiation

The optional speed negotiation block provides the ability for a dual-speed core to implement the FC-FS Section 28 speed negotiation algorithm.

Core Generation

The core may be generated to work at the

following speeds: 1 Gbps, 2 Gbps, 4 Gbps, 1-2 Gbps multispeed, and 2-4 Gbps multispeed. The 4 Gbps-capable cores are available only for Virtex-4 FX devices, while all other speed configurations are available for both Virtex-4 FX and Virtex-II Pro devices.

Table 1 shows the device utilization for the FC point-to-point LogiCORE IP.

Arbitrated Loop LogiCORE Solutions

This core supports FC-1 and the framing and signaling part of FC-2, based on the ANSI INCITS FC-FS Final Draft v1.9 and ISO/IEC FDIS 14165-122 FC-AL-2 standards.

Included within the core are the initialization state machine, based around a MicroBlaze™ soft processor; the arbitration state machine, known as the “old port”

state machine to allow fallback to point-to-point operation; and the loss of sync state machine (LOS-FSM). All classes of Fibre Channel frames are supported. Both alternative and buffer-to-buffer credit management are supported, along with statistics gathering through a statistics vector. Access to control signals is through a register bank with a microprocessor-agnostic management interface or, alternatively, through a dedicated configuration vector.

You can use this core to create any of the following Fibre Channel ports: NL, FL, N, F, E, and B. It should be noted that for non-arbitrated loop applications, the point-to-point core will be a much more efficient solution. No login protocols are directly supported within the core – these are left for you to implement in either firmware or hardware as required.

A fast port interface to the core allows automated arbitration/open circuit/frame transmission/close circuit processes with very little interaction from the user interface.

Figure 3 shows the main blocks within the core.

FC-1 Layer

The FC-1 layer crosses clock domains from the RocketIO™ transceiver recovered receive clock and performs clock correction on the ordered sets allowed in the FC-AL-2 standard.

CRC

The CRC block generates and verifies CRC-32 blocks on outbound and inbound frames.

Loop Initialization Controller

The loop initialization controller controls all loop initialization functions, including exchange of loop initialization frames with other ports on the loop and address assignment. It is implemented using a MicroBlaze soft processor.

Loop Port State Machine

The loop port state machine (LPSM) controls core operation such as arbitration, circuit opening, and circuit closing when the core is placed in an arbitrated loop system.

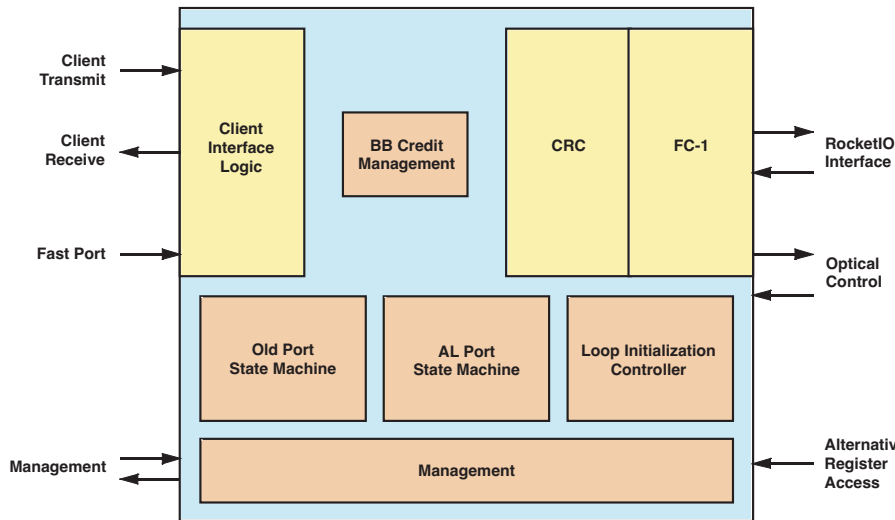


Figure 3 – FC-AL LogiCORE architecture

Old Port State Machine

Whenever the Fibre Channel arbitrated loop core fails to initialize as a loop port, it falls back to Fibre Channel point-to-point operation. The old port state machine controls core operation in point-to-point mode.

Fast Port Interface

The fast port interface manages the opening and closing of circuits in an arbitrated loop system. The core arbitrates for control of the loop by supplying the address of the destination loop port, then signaling back to the user logic when control is obtained. The availability of this fast port feature simplifies software development requirements by taking care of these functions for you.

Management Interface

Configuration of the core is performed through the management interface, a 32-bit processor-neutral control pathway independent of the data path of the core.

Credit Management

The credit management block provides both simple and alternative buffer-to-buffer credit management, keeping track of RRDY and SOF primitives received and sent. In point-to-point mode, the credit management block supports BB_SCx credit recovery as defined in FC-FS Section 18.5.11.

Core Generation

The core may be generated to work at the following speeds: 1-2 Gbps multispeed and 2-4 Gbps multispeed. The 2-4 Gbps-capable core is available only for Virtex-4 FX devices, while the former speed configuration is available for both Virtex-4 FX and Virtex-II Pro devices.

Table 2 shows the device-utilization for the FC-AL cores.

Virtex-II Pro or Virtex-4 FPGAs			
Slices	Total LUTs	FFs	Block RAMs
4,600	6,800	3,300	22

Table 2 - Resource utilization for FC-AL core

Bridge Port Reference Design

To provide an example of how the Fibre Channel cores can be integrated into a complete Fibre Channel port, Xilinx created a comprehensive Bridge (B) Port Reference Design for Virtex-II Pro devices, which is available at no cost.

The B Port Reference Design is based on an EDK project and includes:

- Fibre Channel point-to-point LogiCORE IP (v2_0)
- B Port-specific back-end hardware for automated end-end credit frame gener-

ation (ACK_1) and F_BSY/F_RJT generation

- Optional framer to encapsulate raw data in pre-programmed SOF/header and EOF
- Optional pattern generator that supports CJPAT, CRPAT, CSPAT, and a user-defined pattern
- B Port-specific back-end hardware to filter frames before login
- PPC405-based firmware, including drivers for FC core and B Port-specific hardware and ELP (exchange of link parameters) protocol implementation
- Data multiplexing to allow TX/RX without the involvement of the PPC405 subsystem
- ChipScope™ modules for enhanced debugging
- UART connecting to user-PC through RS232 to allow real-time user interaction with design
- Comprehensive documentation covering everything from register maps to adding new features

Conclusion

With Virtex-II Pro and Virtex-4 FX devices and Fibre Channel LogiCORE IP, you can quickly create an FC infrastructure including switches, bridges, and end points. Up to 4 Gbps is supported today on Virtex-4 FX FPGAs.

The cost-effective FC point-to-point and arbitrated loop cores can deliver 100% throughput. The zero-cost B Port Reference Design with an evaluation-licensed LogiCORE IP and an ML323 development platform offers a risk-free evaluation environment that you can connect to existing FC ports or networks. It also offers a strong basis for the creation of further examples of Fibre Channel ports.

For more information and technical details on Xilinx IP cores for Fibre Channel, visit www.xilinx.com/storage.