

“FIBRE CHANNEL NETWORK TECHNOLOGY APPLIED TO ADVANCED DSP SYSTEMS”

Curtiss-Wright Controls Embedded Computing
Data Communications

4126 Linden Ave.

Dayton, OH 45432-3068 USA

Phone: 937-252-5601 or FAX: 937-258-2729

E-mail: info@systran.com or WWW: www.cwembedded.com

ABSTRACT

New and emerging Digital Signal Processor (DSP)-based digital imaging systems, especially those with sophisticated surveillance and multi-spectral capabilities, must handle data at enormous rates. System architectures have evolved over time to adapt to these demanding data acquisition, data movement, data processing, and data storage requirements. The trend has been from dedicated, through federated, to distributed systems. A key technology that will support the continued movement to distributed systems is the new ANSI Fibre Channel (FC) standard. This paper presents a typical DSP system architecture, provides a technical overview of Fibre Channel, then specifically addresses its application in advanced radar, sonar, medical and other DSP-based imaging systems.

THE GENERIC IMAGING SYSTEM

A generic image processing system is illustrated in Figure 1. This figure is not meant to represent any particular type of imaging system, but is representative of many systems being developed today.

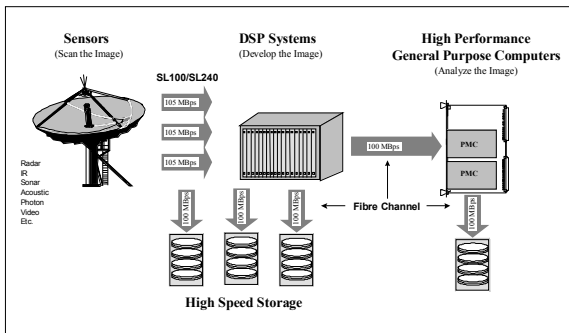


Fig. 1 - Typical Image Processing System

There are four primary elements in this generic image processing system. The Sensor Subsystem scans its environment and outputs unrelenting streams of high speed digital data. The DSP Subsystem processes these data streams and outputs usable image data. The General Purpose Computers display or further analyze the images. Finally, the High Speed Mass Storage Subsystems provide archive data storage at one or more intermediate points along the data path. The ANSI standard Fibre Channel (an efficient, high speed, serial data communications technology) allows these primary elements of the image processing system to be truly distributed with distances up to 10 kilometers.

FIBRE CHANNEL BACKGROUND

In 1988, the ANSI-sanctioned X3T11 committee was formed to develop an ultra high-speed transport for mass storage, peripheral I/O, and network communications. The X3T11 Fibre Channel committee had a vision that next-generation LAN and peripheral I/O communication requirements could be handled with a serial gigabit data stream through a single physical connector. To realize this vision, the committee borrowed the best elements of I/O channel technology and merged them with the best elements of networking technology. The objective was to develop a serial communication link that supports the bandwidth and data reliability needed by I/O channels and the flexibility, connectivity, and distance of networking technologies. The result is “Fibre Channel,” a universal carrier or transporter of data, capable of simultaneously handling networking and I/O channel protocols.

Before the first Fibre Channel standard was approved by ANSI, a group of interested companies joined together to form the Fibre Channel Industry Association (FCIA). The FCIA has two major objectives. First, its members conduct joint market ventures to raise public awareness of Fibre Channel as a next-generation technology and to cooperatively work toward delivering a broad base of Fibre Channel infrastructure to support a wide array of industry applications within the mass storage and IT-based arenas. Second, its technical working groups focus on specific aspects of the technology that target both vertical and horizontal markets, including storage, video, networking, and SAN management. Recently, the Fibre Channel standards work group FC-SW-2 unanimously voted on one interswitch routing protocol proposal, Fabric Shortest Path First (FSPF) to be submitted to the National Committee for Information Technology Standards (NCITS) T11 technical Committee for approval as a standard by ANSI. FSPF specifies a common method for routing and moving data among Fiber Channel switches and will allow switch vendors to build heterogeneous equipment based on the new protocol.

The latest market study of Fibre Channel’s market potential forecasts revenues in excess of \$1.7 billion by 2003. Fibre Channel is rapidly becoming the preferred external server storage interconnect for high-end storage solutions. There are now over 150 companies developing Fibre Channel products. The FCIA includes manufacturers, systems integrators, developers, and systems vendors. Systran Corporation, an FCIA principal member, is working with the other FCIA members to help ensure the success of Fibre Channel in all of its target applications and markets.

FIBRE CHANNEL TECHNICAL OVERVIEW

Fibre Channel is flexible, powerful, and extendible without undue complexity. This section reviews the basic concepts underlying Fibre Channel technology by layers, as defined by the standard and implemented by the community of Fibre Channel suppliers.

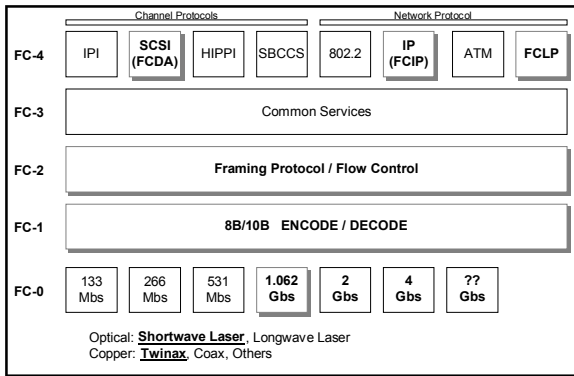


Fig.2 - Fibre Channel Implementation Layers

Fibre Channel defines several implementation layers, as illustrated in Figure 2. Layers FC-0, FC-1, and FC-2 are specified in the FC-PH standard. Layer FC-3 and FC-4 are standards still in development or specified in FCIA Profiles such as FC-PLDA or FC-IP. The next revision of FC-PH, FC-PH-2, is nearing completion of the approval process. It further refines, clarifies, and extends the existing FC-PH standard.

The FC-0 layer defines the physical layer consisting of supported baud rates, media, and connectors. The serial baud rates include 133 Mbps, 266 Mbps, 531 Mbps, and 1.0625 Gbps. Additionally, work has begun to define 2 and 4 gigabit serial baud rates. The market today is dominated by products which implement 266 Mbps, referred to as “quarter speed” and 1.062 Gbps, referred to as “full speed.” Because most Fibre Channel components are rated for full speed operation there is little cost advantage to quarter speed products. However, because higher speed signals do not travel as far, quarter speed is useful where longer distance on a particular media is the driving factor over bandwidth.

In spite of the name “Fibre Channel,” FC-0 defines both optic and copper media. A common copper implementation is on twin-ax with DB-9 or 1x3 differential connectors. The most common optical implementation is short wave laser (SWL) over 50/125 micron multi-mode fiber with dual SC connectors. Long wave laser (LWL) over 9/125 single-mode fiber with dual SC connectors, will soon be an available option. Typical distances at full speed over copper, SWL, and LWL are 30 meters, 300 meters, and 10,000 meters respectively.

These optical distances are for nodes without Open Fibre Control (OFC). OFC is a safety circuit which shuts down laser power if no optical power is being received. Redundant OFC circuits are required to allow higher optical power and remain a Safety Class 1 laser product and OFC does require point-to-point connection of optical ports to operate properly.

Due to the additional circuitry and the fact that a loop topology does not meet the point-to-point connection requirement of OFC, many FC product suppliers now provide non-OFC optical ports that use slightly lower power lasers to meet the

Safety Class 1 laser requirements. Beware, however, that OFC and non-OFC ports are not interoperable at the FC-0 level.

The FC-1 layer defines the encoding scheme, ordered sets, frame delimiters, and other signaling primitives implemented in the Fibre Channel encoder/decoder. The 8B/10B encoding used is an adaptive code designed to provide data reliability, DC balance, and word alignment. Data and special transmission characters and certain combinations of these are referred to as “Ordered Sets.” Ordered Sets are used to identify frame boundaries, transmit primitive function requests, and maintain proper link transmission characteristics.

The FC-2 layer is the signaling and framing protocol layer for Fibre Channel. FC-2 specifies the addressing, flow control mechanisms, sequence management, topology, and classes of service. The data payload, used by the FC-3 and above layers, is transparent within the FC-2 layer. This is an important feature, for this allows many new and existing Upper Layer Protocols (ULP) to be mapped and simultaneously operate over a Fibre Channel network. Fundamentally, FC-2 specifies a robust set of topologies and interconnection options referred to as the Fibre Channel “fabric.”

Fibre Channel may be viewed as a high-performance bi-directional connection between any two ports. The Fibre Channel interconnection fabric connects node ports, called “NL_Ports,” which can be as diverse as peripherals such as hard drives, RAID subsystems, and sensors; or computing systems such as DSPs, workstations, mainframes, and supercomputers. The Fibre Channel standard specifies how an NL_Port will login and communicate over a fabric to another NL_Port, but it does not specify how the fabric actually performs and manages the communication. In this way, Fibre Channel fabrics may be developed to specific applications making use of the latest technology and the most appropriate topology. By doing so, they will not be burdened with an outdated standard or forced into fitting old technology into new and evolving problems.

A Fibre Channel fabric may consist of any combination of three topologies: point-to-point, switched, and arbitrated loop as depicted in Figure 3. This set of topologies allows the system integrator a great deal of architectural flexibility. The system integrator can mix and match the best architecture for the application. For instance, a full speed arbitrated loop has a total bandwidth of 100 MB/s that is time-shared by the attached ports. This is enough bandwidth for many applications and alleviates the extra cost and complexity of a fabric switch.

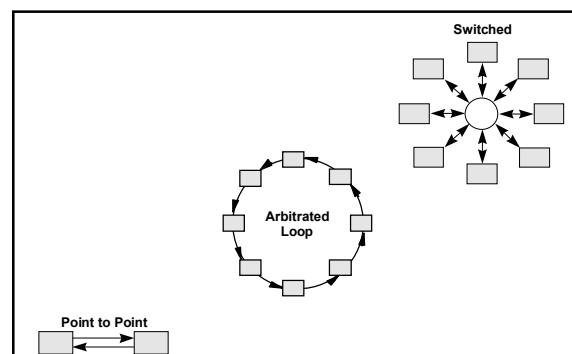


Fig. 3 - FC-2 Topologies

FC-2 also defines how NL_Ports communicate over the fabric using a high level construct called an “exchange.” An exchange is a single operation composed of one or more “sequences.” In turn, each sequence is composed of one or more “frames.” Each frame is composed of:

- a start of frame (SOF) indicator (4 Bytes)
- a frame header (24 Bytes)
- an optional header (up to 64 Bytes)
- a payload (up to 2,048 Bytes)
- CRC error checking (4 Bytes)
- an end of frame (EOF) indicator (4 Bytes)

FC-2 defines three classes of service as highlighted in Figure 4. Classes of service are not topology-dependent, and therefore can be implemented over any topology. The class of service used is a function of specific hardware and software driver implementation.

<p><u>Class 1 - Circuit Switched</u></p> <ul style="list-style-type: none"> • Connection Service—Circuit Switched • End-to-End Flow Control • Guaranteed Bandwidth and Delivery - ACK • In Order Delivery <p><u>Class 2 - Frame Switched</u></p> <ul style="list-style-type: none"> • Connectionless Service - Frame Switched • Buffer-to-Buffer Flow Control • End-to-End Flow Control • Guaranteed Bandwidth and Delivery - ACK <p><u>Class 3 - Datagram</u></p> <ul style="list-style-type: none"> • Connectionless Service - Frame Switched • Buffer-to-Buffer Flow Control • Best Effort Delivery - No ACK - Datagrams • Broadcast
--

Fig. 4 - FC-2 Classes of Service

Class 1 is a circuit-switched, connection-oriented service that guarantees bandwidth and delivery of data. Once a dedicated Class 1 connection is made, through a Class 1 capable fabric, it is maintained until broken by one of the connected NL_Ports. For the duration of the connection, Class 1 offers full bandwidth and in-order delivery of data because the connection is maintained and not broken or changed between frames. Latency is minimized because once a connection is made, the packet headers are not looked at for destination addressing. Class 1 also includes end-to-end flow control to prevent data loss.

Class 2 is a connectionless, frame-switched service where frames are multiplexed on frame boundaries within the fabric. In-order delivery of data is not guaranteed, unless the fabric happens to be an arbitrated loop where only one path for data flow exists. Also, because incoming frames for different ports are multiplexed, the bandwidth to a single port is not guaranteed, even though buffer-to-buffer flow control and data acknowledge (ACK) mechanisms still guarantee data delivery.

Class 3 is a connectionless, datagram service with buffer-to-buffer flow control, but no ACK to guarantee data delivery. In Class 3, the fabric is expected to make a best effort to deliver the data. Any acknowledgment of Class 3 frames is left to the FC-4 Upper Layer Protocols (ULP), implemented in software.

The FC-3 layer is intended to define a set of services that are common across multiple ports within a single node. No standards have been completed for the FC-3 layer at this time.

The FC-4 layer defines the mapping of ULPs such as HIPPI, SCSI, IP, and a native Lightweight Protocol to the lower layers of Fibre Channel. FC-4 is generally implemented in driver software on the host computer. At the time this paper was submitted, all of the profiles and standards for the FC-4 mappings were still in draft form and in various stages of development within the ANSI standards and FCIA technical working groups.

Some of the current implementations use the same peripheral communications model of the FC-PLDA; and other implementations, such as Systran’s FibreXpress FXLP, use a standard networking model.

APPLICATION TO IMAGE PROCESSING

For a number of reasons, Fibre Channel is an excellent technology for use in distributed image processing applications. First, very high bandwidth products, such as Systran’s FibreXpress® Network, are available today, operating at a gigabit per second and delivering sustained throughputs of 50-80 MB/s. Second, it is a very robust technology and offers tremendous flexibility in topology and configuration so that performance and cost can be matched to a particular system requirement. Third, there is enormous commercial market acceptance of Fibre Channel providing the volumes required to attain true commercial-off-the-shelf (COTS) pricing and product longevity advantages.

The first point of application for Fibre Channel in the image processing system is the connection between the sensor and the DSP. The connection to the sensor is usually very high bandwidth, located at a point downstream from a high-speed A/D converter.

In many image processing systems, the sensor subsystem is completely implemented in pipeline hardware. A full Fibre Channel NL_Port connection requires a CPU to manage the protocol chip and to execute FC-4 ULP driver software. However, a CPU, and all the hardware support that goes with it, is an unnecessary expense in this type of sensor subsystem.

In this case, a very simple simplex data link based on the Fibre Channel FC-1 layer encoder/decoder and some FIFOs, is all that is necessary to transfer data from the sensors to the DSP. Systran’s FibreXtreme™ Serial Front Panel Data Port (FPDP) data link provides this type of connection.

Frequently, a CPU is present in the sensor subsystem and the lightweight protocol (FXLP) works well for this type of connection. The advantage of using a full Fibre Channel FC-4 ULP at this connection is that multiple sensor ports can be connected through an FC fabric to a single port on one or more DSPs. The FC-4 ULP also provides two-way communication between the sensor subsystem and DSP which allows synchronization and control information to be easily exchanged.

The second point of application for Fibre Channel in the image processing system is the connection between the DSP and one or more general purpose (GP) computers. These CISC or RISC based computers perform functions such as image display, or further analysis such as automatic target recognition algorithms. This connection will generally use the fast, efficient FXLP lightweight protocol to transport the image data.

The third point of application for Fibre Channel in the image processing system is the connection between the DSP or GP computers and mass storage arrays. The Fibre Channel replacement of SCSI in storage arrays provides the possibility of connecting very large, high-speed mass storage systems through a single lightweight serial data link.

Properly designed into the image processing system, FC-based mass storage can archive the data stream at any point along the image processing path, from the raw sensor data to processed images, and any intermediate point in between.

At the heart of the system, implementation of the Fibre Channel interface within the DSP system is a very important task. To simplify this task, all of the Fibre Channel connections (simplex and FXLP to sensors, raw data and image data to mass storage, and connections to GP workstations) should share a common implementation methodology.

Most high-end DSP systems today are fundamentally VME64-based systems that implement a secondary high-speed internal data path. In the past, the VSB bus was used; more recently, cross-point switched interconnects have evolved such as the RACEway from Mercury Computer Systems.

The parsing and routing algorithms of the input, output, and storage data streams are generally unique to each image processing system. In addition, DSP processors do not generally have operating systems and software development tools that readily support off-the-shelf communication drivers. Therefore, the use of a VME SBC acting as an I/O Processor (IOP) provides the basis for a general solution.

A new wave of VME/PMC SBC products now support multiple PMC slots for multiple communication hardware options within a single slot. They allow the use of common real-time operating systems that have COTS ULP driver support, and a software development environment that can be readily programmed to perform the unique pre-processing, data parsing, and routing functions that may be required.

In addition, there are many CPU, SBC, and OS suppliers from which to choose. This allows the integrator to easily select an already familiar and comfortable programming environment which can greatly reduce the time and S/W tooling cost of image processing system development efforts.

FibreXtreme Serial FPDP data links may also be connected directly to the I/O bus of workstations and single board computers by replacing the FPDP interface with a PMC interface. The PMC versions are completely compatible with the FPDP versions and can be mixed as needed by the image processing system. Dual DMA controllers with DMA chain list capability are used to move data to and from the PCI bus to the FibreXtreme card. Setup and control of the PCI interface is a relatively simple driver for most CISC and RISC based SBCs, or can easily be provided as library functions for direct connection to DSP based SBCs with PMC sites.

Both FibreXtreme and FibreXpress can also take advantage of the LinkXchange Crossbar Switch (LX2500). LinkXchange is a full crossbar switch that can copy the signal at any input port to any or all output ports. The switch can be configured with up to eight port cards of four ports each. Thus, it can be configured in

size from a 4x4 up to a full 32x32 crossbar switch. The LinkXchange can also be cascaded up to five times. This allows a single input to be replicated and simultaneously routed to up to 32 million destinations over 50 kilometers apart for very distributed parallel processing!!!

The switch in the system not only allows real-time parallel processing of the data stream, it allows the data stream to be copied to a high speed recording device. Once recorded, it can reconfigure the system to allow the recording device to playback the data stream into the DSP system for post run processing and refinement of the processing algorithms.

SUMMARY

Fibre Channel allows all the major components of the image processing system (the sensor, the DSP, the GP workstations, and the storage) to be separated by distances up to 10 kilometers. It is available in a variety of speeds, media, and topologies so that performance and cost requirements can be easily matched. It is scaleable, providing a growth path for systems that start small, but must expand to meet the growing requirements of an overall system.

ANSI standard Fibre Channel products and technology provide a robust, available, cost-effective solution to the data communication needs of today's distributed image processing systems.

ADDITIONAL INFORMATION

Additional information about the Fibre Channel protocol, standards, working groups, and related issues can be obtained from:

Fibre Channel Industry Association
404 Balboa Street
San Francisco, CA 94118
(415) 750-8355
Email: info@fibrenchannel.com
Web Site: www.fibrenchannel.com

Additional information about the Serial FPDP protocol, standards, working groups, and related issues can be obtained from:

The VMEbus International Trade Association
7825 E. Gelding Dr.
Ste. 104
Scottsdale, AZ 85260
Phone: (480) 951-8866
Email: info@vita.com
Web site: www.vita.com



Product specifications mentioned herein are subject to change without notice. FibreXtreme, FibreXpress and LinkXchange are registered trademarks of Systran Corporation. RACE is a trademark of Mercury Computer Systems. All other trademarks or registered trademarks mentioned herein are the sole property of their respective owners. © 2004, Curtiss-Wright Controls Embedded Computing. All Rights Reserved. Doc. 550-C (07/19/04)