FIBRE CHANNEL BUS MONITORING WITH AIRBORNE DATA MULTIPLEXER / RECORDER SYSTEM

Albert Berdugo
VP and Chief Technology Officer
Eric Pesciotta
Development Manager, Network Systems
Teletronics Technology Corporation
Newtown, PA USA

ABSTRACT

Modern aircraft now employ widely accepted and standardized technology commonly found in COTS applications. One such technology, Fibre Channel, has been deployed to transport both low and high-speed measurement data. Data as varied as "command and control", "Radar Sensors" and "video" are being transmitted over fibre channel on many aircrafts. Some of these applications require data monitoring in listening mode only where transmission from the instrumentation equipment is not allowed or possible. As a result, standard off the shelf Fibre Channel devices cannot be used, and a development of a general purpose Fibre Channel monitor/ analyzer device and product is required.

This paper discusses the concept, merits, and implementation of fibre channel bus monitoring in modern data acquisition systems. Techniques for tapping into an optical fibre channel network, as well as, a recording format for IRIG106 Chapter 10 are included. An overview of fibre channel topologies and protocols is also provided.

KEY WORDS

Fibre Channel, Data Acquisition, Recorder, IRIG-106 Chapter 10

INTRODUCTION

Fibre Channel bus is widely used in the telecom and storage industries. In the past 15 years the bus has been making its way on many aircrafts applications. Original application of the Fibre Channel on aircrafts was to pump high data rate from radar sensors to radar processors. Today the Fibre Channel bus can be found as the main avionic bus, as well as a dedicated bus for some high data rate to transport video and data between units. Additionally, aircrafts using avionics fibre channel as their core bus have switch fabric to route video and data. Data monitoring of the fibre channel can be relatively simple or very complex task. That complexity depends on many factors such as fiber vs. electrical bus, receive only vs. the ability to establish fibre channel protocol link, availability of an instrumentation port on a switch fabric vs. snooping on a bus, etc.

This complexity may extend beyond the simple data collection, and may require extensive post flight data processing to assemble fibre channel data frames into data messages that can be analyzed. The messages may include video, data, and aircraft specific formatted payload.

The simple and straightforward fibre channel monitoring applications are the ones that establish complete protocol link (transmit/receive) to acquire and record fibre channel data. For this, there are many off the shelf devices that can be used to communicate over fibre or electrical interface and establish proper link to collect and record data frames and data messages. The more difficult application is to monitor fibre channel in the receive mode only. This is the topic that the paper will be discussing in detail. A primer on fibre channel standards is provided, followed by specifics of monitoring fiber channel data.

INTRODUCTION TO FIBRE CHANNEL

Since 1988, the T11 working group of the InterNational Committee for Information Technology Standards (INCITS, accredited by ANSI) has been defining standards for Fibre Channel. These standards create a class of computer communications protocols to satisfy high performance information transport. Communication speeds up to 10 Gigabits per second, varied connection topologies, copper, fiber mediums, and an abundance of application level protocols are accommodated.

Definition of the fibre channel standard is split amongst multiple documents. The standard is organized into levels, FC-0 through FC-4. The foundation levels, FC-0, FC-1, and FC-2 were originally defined by a series of documents; *ANSI X3.230-1994, Fibre Channel – Physical and Signaling Interface (FC-PH), ANSI X3.297-1997 (FC-PH-2), and ANSI X3.303:1998 (FC-PH-3).* The later two documents were supplements that refined the original definition. A fourth document was added to extend the definition, "*ANSI X3.272-1996*, *Fibre Channel Arbitrated Loop* (*FC-AL*)". In time, this collection of documents was reorganized into 2 documents covering levels FC-0 through FC-3; *INCITS 352:2002, Fibre Channel Physical Interface* (*FC-PI*) *and INCITS 373:2003, Fibre Channel Framing and Signaling* (*FC-FS*). The top-most level, FC-4 provides definitions for mapping higher level protocols into fibre channel. These are often referred to as upper layer protocols and each mapping is defined in its own document. There is a large collection of standard upper layer protocol mappings. Some of those directly applicable to the flight test community are listed below.

- SCSI-FCP ANSI X3.269-1996 Fibre Channel Protocol for SCSI
- FC-LE ANSI X3.287-1996 Fibre Channel - Link Encapsulation
- FC-AV Project 1237-D Fibre Channel - Audio-Visual (FC-AV)
- FC-AE Project 2009-DT Fibre Channel - Avionics Environment (FC-AE)
- FC-AE-1553 Project 1648-DT Fibre Channel Avionics Environment - Upper Layer Protocol MIL-STD-1553
- FC-AE-ASM Project 1649-DT Fibre Channel Avionics Environment - Anonymous Subscriber Messaging
- FC-AE-LP Project 1650-DT Fibre Channel Avionics Environment - SCSI-3 Lightweight Protocol

- FC-AE-RDMA Project 1651-DT Fibre Channel Avionics Environment - SCSI-3 Remote Direct Memory Access

FC-0: Physical Interface
The lowest level of the standard defines the physical interface. Data rates, media variants, connectors, distance limitations and signal quality are all defined. Media definitions include single and multimode fiber, copper, and an assortment of adapters and connectors.

FC-1: Encode/Decode and Link Control
Above the media level, transmission encoding and decoding (codec) is dictated at the FC-1 level. Fibre channel carries eight bit data bytes in 10 bit characters. Known as the "8b/10b" encoding scheme, it provides superior transmission characteristics with very low error rates. Building on the codec, FC-1 defines unique sets of characters called *ordered sets*. These are used to provide primitive signaling capability and link-level protocol.

FC-2: Framing Protocol
Fibre channel data is carried in frames above the link layer. A group of ordered sets are defined as *frame delimiters*. There are 22 frame delimiters in all; 12 *start-of-frame* (SOF) and 10 *end-of-frame* (EOF). The remaining ordered sets are transmitted between frames for signaling and link control. FC-2 defines the structure, format and sequencing of frames. Flow control mechanisms are defined, as well as, the classification of data. Frames are classified by service level, allowing sending ports to control the reliability and order of data delivery. This is one of the reasons why there are so many start-of-frame, end-of-frame pairs.

The combinations of flow control and transport attributes make up 4 classes of service suitable for different applications. Each *class of service* has unique characteristics. *Class 1* is connection-oriented, providing exclusive access to the full bandwidth of the link. *Class 2* is connectionless with frame delivery notification, but no bandwidth guarantee. *Class 3* is connectionless without frame delivery notifications or guaranteed bandwidth. *Class 4* provides a connection-oriented virtual circuit with frame delivery notifications plus bandwidth and latency guarantees.

Regardless of class of service, all data is transported in frames (Figure 1). Each contains a 24 byte header with various frame attributes. Optional frame headers are defined where needed and all frames include a CRC. Among other things, the mandatory header controls frame routing and sequencing information. Each frame can carry up to a maximum of 2112 bytes of data. Application level data is not limited by this restriction. Larger units of application data are fragmented into a *sequence* of frames and reassembled at the destination.

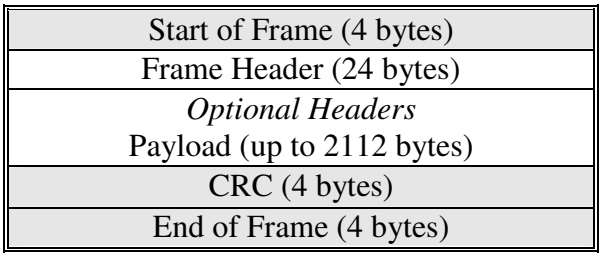

Figure 1. Fibre Channel Frame

FC-3: Common Services
The FC-3 level was loosely defined for future use. It is reserved for the role of common services spanning multiple ports.

FC-4: Protocol Mappings
FC-0 to FC-2 provide sufficient definition for effective transmission of data. However, it is the FC-4 level that makes the transport available for application protocol use. The ability to transport upper layer protocols like SCSI, or TCP/IP make fibre channel practical for application use. The encapsulation of upper layer protocol (ULP) within fibre channel is defined for multiple applications. To date this includes, but is not limited to; SCSI, IP, Link Emulation, Audio-Video, and Avionics. Each ULP mapping defines the information transferred between endpoints as *information units*. A grouping of information units transferred between a pair of ULPs is called an *exchange*. Each exchange may be transported in one or more related information units, which may be transported in one or more sequences, which may be transported in one or more frames. Throughout an exchange, it is the frame headers that contain all the information needed to make sense of the hierarchical transport.

Network Topology
The Fibre Channel standards define three topologies; point-to-point (Figure 2), arbitrated loop (Figure 3), and switched fabric (Figure 4). Each topology is defined for a particular purpose.

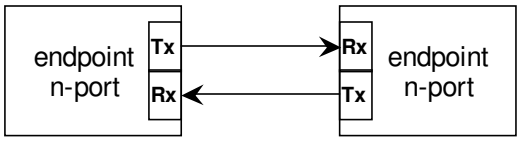

Figure 2. Fibre Channel Point-to-Point

The simplest topology, point-to-point, is designed to interconnect two endpoints. The transmitter on each device is connected directly to the others receiver. All fibre channel ports in a point-to-point network are referred to as N-Ports. N-Ports require the least amount of protocol interaction to establish a connection.

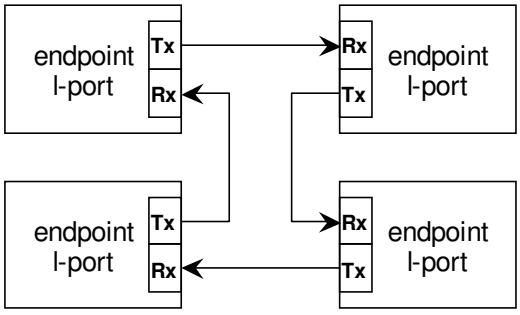

Figure 3. Fibre Channel Arbitrated Loop

The arbitrated loop topology was designed for mid-sized networks. A loop can accommodate up to 127 endpoints. Devices are directly interconnected in a ring. All ports within an arbitrated loop are referred to as L-Ports. Wiring for L-Ports is complicated by the fact that a single

endpoint requires connection to 2 others in the loop. The wiring can be dramatically simplified through the use of fibre channel hubs. Fibre channel hubs are mostly passive devices that are not considered endpoints and do not introduce additional protocol overhead. Arbitrated loops do not use a token passing scheme. All basic communications remain essentially point-to-point between endpoints. Communication on the loop is limited to a single pair at a time. Thus, additional protocol overhead is required to arbitrate access to the loop.

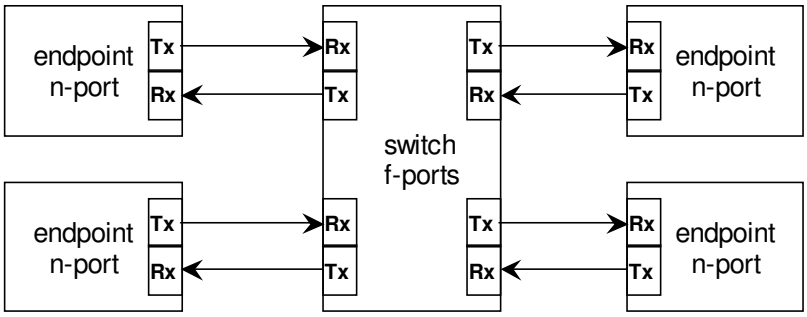

Figure 3. Fibre Channel Switched Fabric

For large scale networks, the switched fabric topology (Figure 3) can accommodate up to 2^{24} endpoints. Each endpoint in a fabric connects directly to a port on a network device, a fibre channel switch. Within a switched fabric, ports on the switch are referred to as F-Ports and those on the endpoints are N-Ports. Switches within the fabric may be interconnected to build larger and larger networks. Switches are active participants in the network and may provide high level services substituting for a server in other topologies. Switches add to the protocol overhead within a network. However, a switch enables simultaneous communications between many endpoints.

HARDWARE ARCHITECTURE

There is a need to monitor several optical fibre channels in receive mode only. Since there is no true fiber channel link, it was clear that there was no possible way to identify messages and payload protocol within those messages. If the transmit and receive bus signals are acquired in the receive mode as two separate channels, and major processing is possible within the bus monitor, then a potential frame stitching can be done to identify application based messages. In our case the monitor card received two independent channels.

It was assumed that there must be some devices or Intellectual Property (IP) based solution to monitor fibre channel data in the receive mode only. We identified several fibre channel bus analyzer companies that implemented FPGA based IP for data monitoring. It was impossible to justify their IP cost and the recurring per channel cost due to the low production quantities. We decided to develop the IP to resolve the immediate customer needs and to add in-house capabilities that can be transformed with some minor changes into opportunities in other applications.

The requirements were to monitor several optical Fibre Channel signals operating at 1 Gb/s in the receive mode only at optical wavelength of 830 to 860nm. This function was to be designed

as part of IRIG-106 Chapter 10 multiplexer / recorder system. A quick area study showed that it was possible to integrate two channels per circuit card with a single FPGA and some peripheral devices. The two channel single board design was based on several key components including the fiber connector, Optical to Electrical transceivers, FPGA, Ping-Pong storage buffers and the PCI bridge. The goal was not only to meet current customer requirements, but also to meet other potential applications that operate at 2 Gb/s.

Board Design Description

The two Fibre Channel inputs interface through Glenair Fibre D-shell ruggedized connectors. This connector introduces about 1 dBm loss and uses 62.5 / 125 um multi-mode fiber. Each channel is routed to its optical to electrical transceiver. The transceiver can operate at rates of 1 and 2 Gb/s and is rated at –17dBm to 0dBm of optical input power. The transmitting section of the transceiver is not used in the current design, but it is wired to the FPGA and the connector for future use. Received data is routed through the FPGA for processing (see FPGA Data Processing) at a preprogrammed line rate, and then sent for recording through the on board PCI Bridge. Data transfers between the FPGA and the bridge (Local Bus) occur at 32-bit 66MHz, while transfers between the bridge and the system backplane occur at 64-bit 66MHz or 64-bit 33MHz depending on the chassis type. The backplane provides time with 100ns resolution to the FPGA for Fibre Channel frame time stamp. Please refer to Figure 5 for the Board and FPGA Block Diagram.

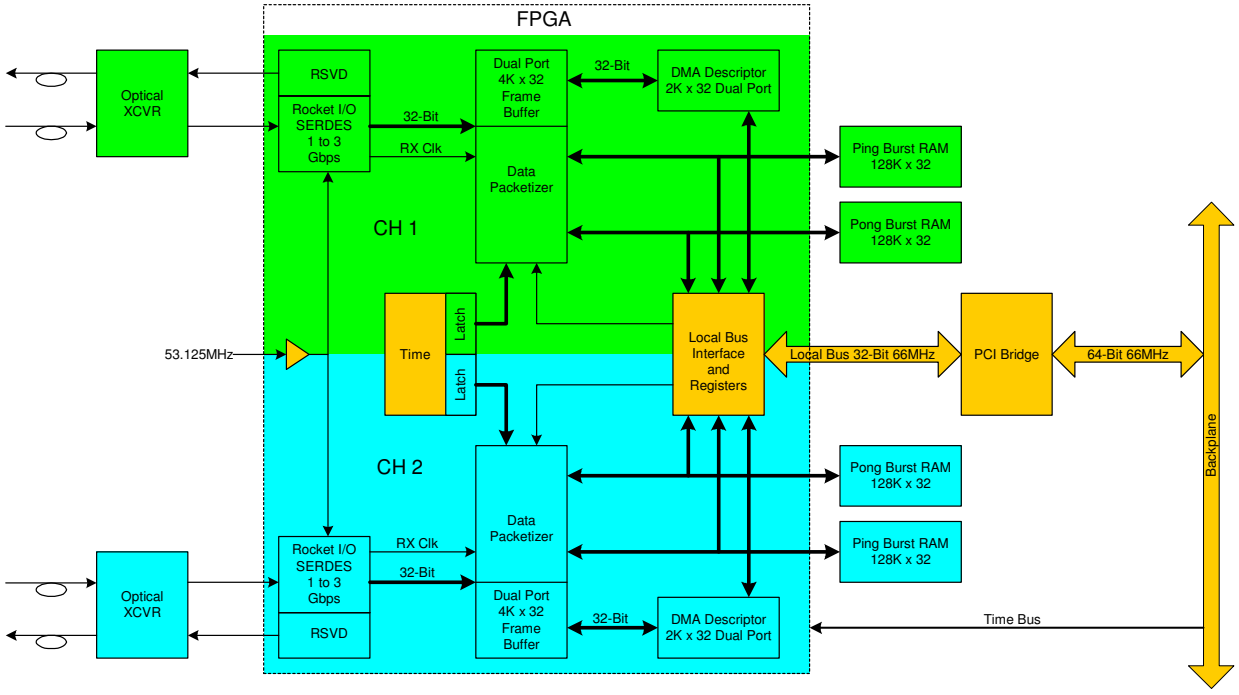

Figure 5. Board and FPGA Block Diagram

FPGA Data Processing

A single XILINX FPGA was chosen to process received data from up to two channels operating at either 1 or 2 Gb/s. The XILINX Virtex 2 Pro was the device of choice due to its Rocket I/O IP core, the device speed, and its internal memory. This device incorporated a proven reliable multi-

gigabit receiver needed to accomplish the task. Early in the design phase it was discovered that the XILINX Rocket I/O IP did not calculate the packet CRC correctly for some classes of operation. To enhance the CRC checking capabilities of the Fibre Channel interface, it was decided to perform the CRC32 in the FPGA fabric to cover the full range of classes of operation.

Each channel used an FPGA internal RAM for temporary storage of data frames before sending the data to the packetizer and descriptor memory in order to control the Ring Management Scatter Gather DMA Mode. Frames are timestamped at the SOF, and stored with the data in the dual port RAM. The packetizer retrieves the data from the dual port RAM and based on the frame size, determines if it can fit in the external 128Kx 32 Ping-Pong buffer. If it cannot fit, the frame is held for buffer swap, and then be inserted in the next buffer as the first frame. One buffer is used for the FPGA to write incoming data frames, while the other buffer is used by the DMA to transfer the data to the host memory buffer for recording. Data is transferred to the host for recording under two conditions: 1) The 128K buffer is full, and 2) buffer is not full but 90 ms time elapsed since the first data frame was written to the buffer. The second condition is based on the IRIG-106 Chapter 10 requirement that "Each Packet Must be Generated within 100 milliseconds whenever data is available to ensure packets contain less than 100 milliseconds worth of data".

Two key components were critical to the success of the Fibre Channel bus monitor card, the Rocket I/O clock oscillator and the Ping-Pong buffer devices. The clock oscillator provides LVDS interface and very low jitter, with possible operation at 1 or 2 Gb/s. The Ping-Pong buffer used for each channel is a 200MHz burst RAM that enabled the size and speed to pump the incoming data at a very high data rate to the host.

Network Access and Monitoring

The fibre channel monitor card must exists within the fibre channel network without being an active participant. There are two options for optical networks. In a switched fabric topology, a custom switch may provide a non-standard port (Figure 6). It must be a port that does not rely on protocol to operate and the switch must mirror all fibre channel traffic of interest to the port. This solution requires special circumstance and may be difficult to find. However, some modern aircraft provide such a port. Often, these are referred to as Instrumentation Ports (I-Ports).

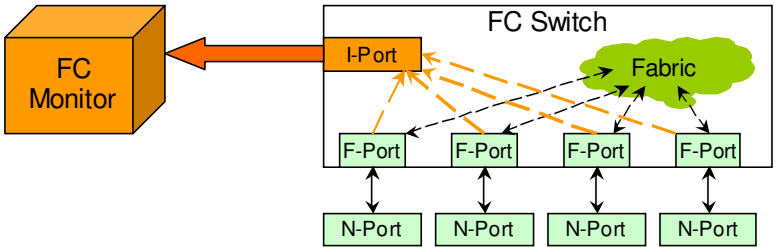

Figure 6. Fibre Channel Switch with non-standard I-Port

All topologies can accommodate monitoring through the use of an optical splitter. The optical splitter must be located on the fibre channel segment of interest. One side of the split would feed data to and from the network, while the other side feeds the monitor card.

SOFTWARE ARCHITECTURE

The fibre channel monitor card operates in a PowerPC processor based system with a PCI backplane. The host processor runs a Linux operating system with custom software. For the fibre channel monitor card a new device driver and custom application was needed. This software was designed for performance, data integrity and flexibility.

Device Driver

The device driver is a low level piece of software that interacts directly with the hardware. Most notably, the driver communicates with the cards PCI bridge (PLX9656) and FPGA (XYLINX Virtex 2 Pro). The PCI bridge is programmed by the driver to facilitate DMA transfers from the card to host memory. This involves careful setting of a myriad of registers on the bridge. The bridge supports multiple DMA strategies. Due to the large amounts of memory required to satisfy high bandwidth data capture, a scatter/gather mode was selected. In this mode, the bridge is programmed with an array of ring buffer descriptors. Each contains a source, destination address, and buffer size. Once DMA is enabled, the bridge manages transfers using the source and destination buffers in the descriptor ring. 8 KB of memory is reserved per channel, providing up to a maximum of 512 descriptors on a ring. While running, bridge registers are continuously checked for status to detect and handle any error conditions. The FPGA exposes a register set as well. Through registers, the device driver initializes and controls operational state of the FPGA. Registers are utilized to detect error conditions and govern the data rate during overload conditions.

Captured Fibre Channel Frame Format

Through an optical splitter, or I-Port, all traffic on the network is directed at the fibre channel monitor port, including; all ordered sets, framing overhead and application data. The FPGA is designed to capture all transmissions between SOF, EOF ordered sets and ignore all other ordered sets (considered protocol overhead). The monitor exclusively targets application data regardless of the ULP, class of service or frame type. Frames and all information required to allow reliable exchange reassembly (*in post processing*) are needed. Metadata is added to each frame as it is captured (Figure 7).

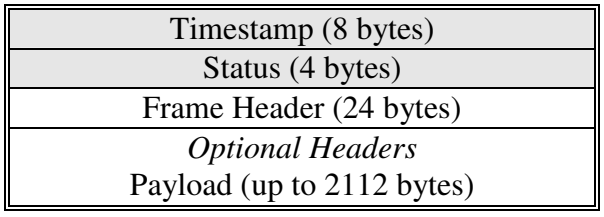

Figure 7. Captured Fibre Channel Frame

The data capture process strips the SOF, EOF, and CRC words (4 bytes each) from each frame and encodes this information in a packed format within the metadata. The metadata provides critical information for reassembly. This includes a timestamp and status word. The timestamp is a 64 bit relative time taken from the 10 Mhz oscillator (see IRIG 106-05 Chapter 10). The status word indicates errors, frame type, and frame length (Figure 8).

31	30	29	28 19	18 16	15 12	11 0
FE	CE	OE	reserved	EOF	SOF	Frame Length

Figure 8. Captured Frame Status Word

Error Handling

As a practical matter, it is important to not overlook the potential for problems in the network. The monitor must not fail in the presence of poorly implemented protocol, or faulty transmissions. To satisfy this demand, the FPGA tracks frame state. Protocol problems such as missing EOF, multiple SOF, runt frames (too short, or missing headers), and excessive frame length (greater than 2112 bytes) are handled gracefully. The status word indicates these errors with the framing error bit (FE) when they occur. Transmission problems such as bad CRC are equally handled. The metadata status word indicates these errors with the CRC error (CE) bit.

Integrity of captured data must be maintained under extreme conditions. Operating at 1 Gb/s per channel, the fibre channel monitor board must be capable of capturing and moving over 200 MB/s across its PCI backplane. Alone, this is a challenging task. Combined with the requirement to record at that rate makes the task almost impossible. Modern fibre channel based hard drives have trouble sustaining 50 MB/s. Solid state drives are generally slower. It is possible with RAID and parallel recording operations, however the average airborne system today would have trouble sustaining this rate. Moving data at such high rates also burdens the host processor. Whether the system bottleneck is at the recording interface or elsewhere, care must be taken to maintain data integrity. The fibre channel monitor board software was designed to detect overload conditions and properly manage them to protect previously captured data.

Though the on-board memory is limited to 1 MB per channel, a super buffer ring technique is used to overcome this. The device driver can allocate as much memory as needed to handle high performance applications with extreme bandwidth bursts. The design incorporates a host based ring buffer scheme overlaid on top of the PCI bridge ring. Effectively, this creates a ring of super buffers chained together. As data is gathered in the boards ping buffer the pong buffer is DMA'd to a super buffer in the ring. When the ping buffer fills, the roles reverse and the pong buffer is used to capture frames while the ping buffer is DMA'd. For each successive ping/pong transfer, an interrupt signals the host and the FPGA advances to the next super buffer in the ring. As the FPGA advances, a *Last Descriptor Written* register is updated. When the host processes an interrupt, this register provides the location of fresh frames to consume. Likewise, when the host consumes a buffer, a *Last Descriptor Read* register is updated.

Under overload conditions, capturing and routing more data only increases the system burden. The monitor board and software govern incoming data at the point of ingress to the system. The pong buffer transfer must complete before the ping buffer fills up. In case this fails, the error is signaled in the overrun error bit (OE) of the first status word in the next transfer. However, it is not likely that DMA across the PCI bus will cause a bottleneck. Overruns are far more likely to be caused by slow recording media. Also, under extreme loads, it is possible for the host to miss an interrupt. In any case, integrity of the recorded data is maintained. The *Last Descriptor* registers were designed to coordinate super buffer consumption between the host and FPGA. The FPGA is designed to not advance past the current read location. When all super buffers in the

ring are full, there is no place to transfer the content of a ping/pong buffer without corrupting previously captured data. Thus, the FPGA will stop the DMA process and start dropping incoming frames. This maintains data integrity in the super buffers and allows the host to recover from the overload. Still, it is not enough to blindly drop data. The board is designed to generate overrun reports which are transferred in-stream (for recording) as if they were a fibre channel frame. The overrun report includes the same metadata timestamp and status making it appear similar to a fibre channel frame. However, rather than a frame header and payload, the report contains the number of buffers lost, the number of frames dropped, and the number of bytes dropped. When the OE bit is set in the status word, the frame contains an overrun report (not a true fibre channel frame).

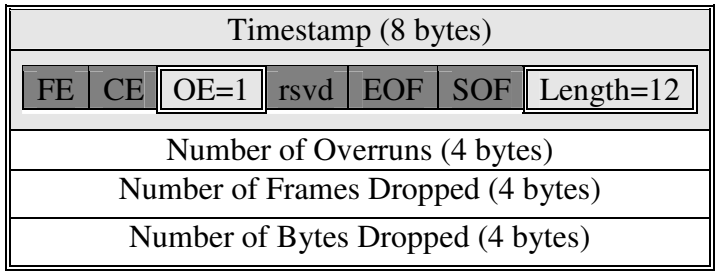

Figure 9. Buffer Overrun Report

Functional Validation

Validating the fibre channel monitoring functionality is a significant task. A multitude of network conditions must be simulated, the data captured and verified. To generate data, a commercial fibre channel simulator was used. The product provided tools to create custom frame content and generates fibre channel traffic with great accuracy and variety. To validate captured data, a custom verification tool was written. A single frame format and multiple test scripts were designed to test all of the following conditions.

- All combinations of SOF, EOF
- Ordered sets between frames (normal operation), within frames (abnormal), and unrecognized sets (abnormal)
- Framing errors
  - invalid size (>2112 bytes)
  - bad CRC
  - bad protocol (SOF followed by SOF, EOF followed by EOF, etc.)
- 100 ms timer
- Error recovery on overruns
- Timestamp accuracy
- Frame dynamics
  - Varying size frames at fixed rates
  - Varying rates of fixed size frames
  - Varying size frames at varying rates
- Line speed and throughput; 1 channel 1 Gb/s, 2 channels 1 Gb/s, 1 channel 2 Gb/s, 2 channels 2 Gb/s

Each test script controlled the rate at which frames were sent and the content of each frame. When a script finished, it was automatically rescheduled and run again. Embedded within each frame payload was a sync word, the expected status word, and an incrementing pattern indicating ping-pong buffer numbering, frame counters, and byte counters. This complex pattern was then written into a validation program that could analyze every captured frame in great detail, based primarily on the content of the frames themselves. With these tools, terabytes of data, collected for many hours on end, were scrutinized and verified to be 100% accurate.

CONCLUSION

The emergence of fibre channel in modern aircraft clearly defines the need for a reliable, high performance, highly adaptable monitoring device. The fibre channel monitor described by this paper has been proven to satisfy those needs as a result of successful design, manufacturing and validation.

REFERENCES

ANSI X3.230-1994, Fibre Channel – Physical and Signaling Interface (FC-PH)
ANSI X3.297-1997, Fibre Channel – Physical and Signaling Interface-2 (FC-PH-2)
ANSI X3.303:1998, Fibre Channel – Physical and Signaling Interface-3 (FC-PH-3)
ANSI X3.272-1996, Fibre Channel Arbitrated Loop (FC-AL)
INCITS 352:2002, Fibre Channel Physical Interface (FC-PI)
INCITS 373:2003, Fibre Channel Framing and Signaling (FC-FS)