4 Port Fibre Channel Host Bus Adapter (HBA) design

By Mohamad Tisani

1.0 Introduction

The PCI-X bus has been used for many years and is still running strong. This interface is being developed in Servers, Datacom and Telecom Systems. With many different PCI-X devices and peripherals being placed in systems today, the need for PCI-X Bridges becomes essential. The evolutionary PCI-X Architecture enhances system performance with better efficiency. It provides up to eight times better performance then the PCI bus. The PCI-X bus pushes the speed to 133 MHZ and adds the split transaction, which makes the utilization of the bus much more efficient. Pericom Semiconductor Corp. has a broad selection of PCI and PCI-X bridges including the PI7C21P100, PI7C8150B, and the PI7C8154B. This application note discusses the use of the PI7C21P100 PCI-X to PCI-X Bridge when designing Fibre Channel Add-in cards.

2.0 Fibre Channel

The explosion of the network traffic and the demands for low cost solution makes the multiple ports Fibre channel card ideal for servers and high bandwidth connectivity requirements. Fibre channel is an industry standard high performance interface designed to bring speed and flexibility to multiple networked storage systems. It is also hot pluggable which allows the user to remove or install the drives while the host system is still operational. Fibre channel also provides high data integrity in multiple drives systems, including Fibre Channel RAID. It can also maintain data integrity through very long cables. Optical fiber can be as long as 10 Kilometers.

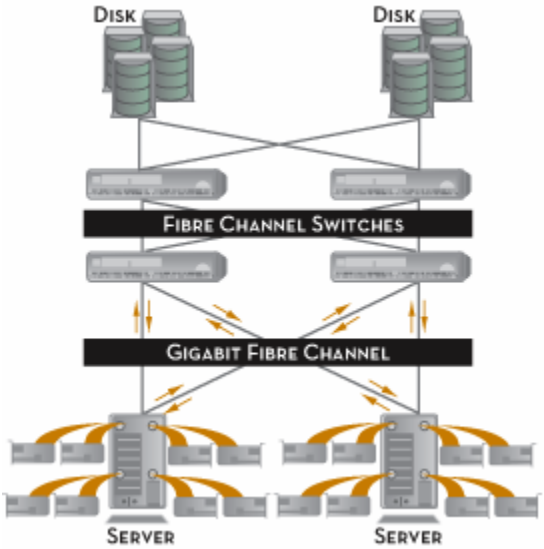

Figure 1

3.0 System Block Diagram

Normally the PCI-X Fibre Channel HBA is plugged into the PCI-X slot on the server. Servers usually provide anywhere from two to six PCI-X slots. A block Diagram of a server is shown in figure 2. Currently the PCI-X is the I/O expansion bus of choice. Most if not all of the servers provide PICX slots for I/O expansion.

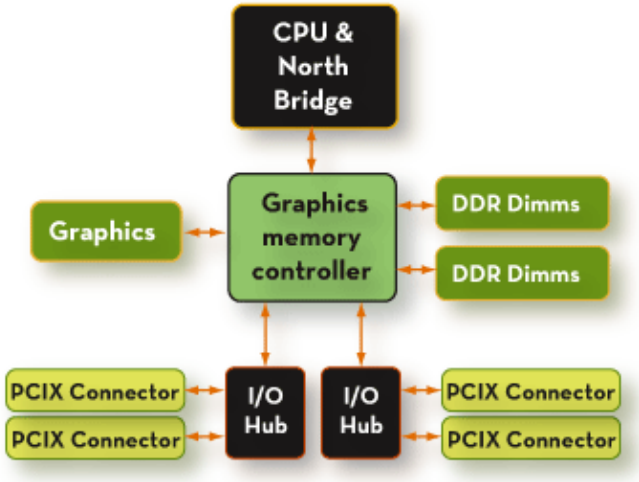

4.0 Application Block Diagram

Please refer to figure 3 for the Block diagram of the Four Port gigabit Fibre Channel HBA. The main components for the Add-In card are:

1. PCI-X to PCI-X bridge
2. 2 Ports Fibre Channel controllers
3. Clock oscillator for the PCI-X bus
4. Clock oscillator for the Fibre channel
5. PLL clock buffer for the 133 MHz PCI-X clock

All of these components are required and these are the bare minimum chips needed.

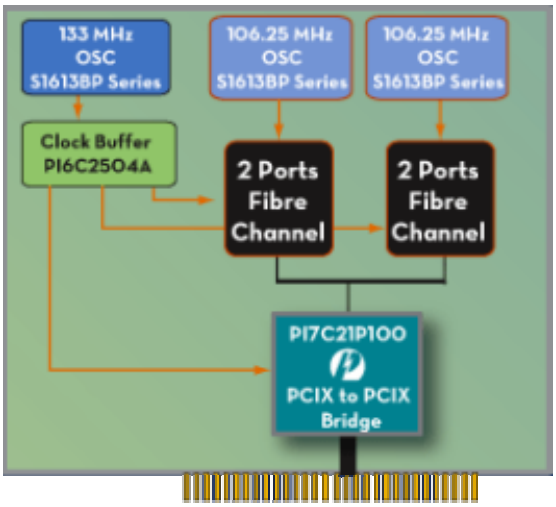

Figure 3

5.0 Unique features

Apart from some common features the PI7C21P100B device has some unique features that are used for this application:

5.1 High timing margin Output buffers

The PCI-X specification states that for the PCI-X bus the maximum loading is one load at 133 MHZ or two loads running at 100 MHz. The PI7C21P100B adheres to that specification but also allows the designer to have more loads. For example on the secondary bus we were able to have three loads running at 133 MHZ. This feature enabled the designer to have a six ports Fibre channel HBA for high integration. This feature is extremely helpful since many of the competing PCI-X to PCI-X bridges fail when used in this type of application. They fail simply because they do not have enough timing margin to work in this high-speed high loading environment. The PI7C21P100B has plenty of timing margins in order to allow the designer to have three chips load running at 133 MHz speed.

5.2 Optimum FIFO size

Pericom has experimented with many different number of write and read FIFO's. The analysis was done on many different applications. Our analysis indicated that we need the following FIFO size to achieve the best performance:

- 4 K bytes delayed read transaction FIFO's for both the primary and the secondary sides.
- 2 K bytes of posted write buffers for both the primary and the secondary sides.

The data indicated that bigger FIFO's are not necessarily better. You have to choose the right size in order to service both the write and the read transactions.

Aside from having separate read and write FIFOs the bridge can split the read or the write FIFO as 4 different buffers. This way the bridge can have up to 4 different read or write threads if needed. This is important for the six ports Gigabit Fibre Channel card design. Since you have three different Gigabit FC controllers on the secondary, in order to have them all active at the same time you will need this feature. This way all of the three controllers will be transferring data at the same.

This Architecture proved to be the optimum based on the performance report, which is also provided on the web site for your convenience.

5.3 Programmable FIFO Space

As you know all memory transaction data going in or out of the bridge comes from, or goes through the FIFO memory inside the bridge. The FIFO is described in the section above. The architecture and the size of the FIFO are very crucial to the Gigabit FC card performance since the Bridge will control all of the data throttling. In order to enhance and make unique your application, the PI7C21P100 bridge provides a register that enables you to control the size of the data moving through the bridge. The user can customize the application by programming the bridge to accept the transaction based on how much available space is in the FIFO. This available FIFO space can be programmed to 128, 256, or 512 Bytes. So the bridge will ensure that we have enough entries in the FIFO before accepting the transaction. This will help the hardware blend well with the software driver.

Lets say that the Gigabit FC Controller will have its best performance when it is transmitting 512 Bytes of data. In this case it is best to program the bridge free space to be 512 Bytes. This way the data transmitted in a single transaction with no reties or disconnects. Otherwise the chip will have to try to transfer the 512 Bytes soon to be stooped by the bridge may be at the 128 bytes boundary. Then it will have to retry again and transfer another 128 Bytes. Every time there is a retry to retransmit there is a loss of about 10 clocks for ending the current transfer giving up the bus and then rearbitrate again to get a grant and try to transmit again. This is a unique feature that the competing products do not have.

6.0 Conclusion

The PI7C21P100 PCI-X to PCI-X bridge device is well suited for the 4 ports and 6 Ports Gigabit Fibre Channel Host Bus Adapter applications. It has many unique features like: enhanced Timing margin output buffers, Optimum FIFOs sizes and up to 512 bytes of programmable FIFO free Space for a single transaction. These features will facilitate the design of the Multiple Ports Gigabit Fibre Channel Card application and it will also provide a better performance, which is critical for this application.