PCI Express® Over Cable Link Tuning Essentials

By

Steve Riley, Sr. Design Engineer

PCI Express Over Cable System -- I/O Expansion Application

Introduction

The promise of simple I/O expansion in an existing host server is made an implementation reality with proper PCI Express® link tuning and design fundamentals. Engineering a successful PCI Express over cable installation requires that the PCIe link-tuning fundamentals discussed in this article be properly implemented.

These fundamentals need to be a part of any PCI Express host and target card design. PCIe link tuning parameters must be adaptable to compensate for the system variability common in today's host servers. This article will include PCI Express over cable link tuning illustrations and application examples that demonstrate effective PCIe I/O expansion chassis implementations.

The PCIe Over Cable Opportunity

PCI Express (PCIe) has evolved very quickly from a component interconnect on a PC board into the pervasive system I/O interconnect of choice in a wide variety of industrial computers and servers. There are many reasons for this, but some of the most compelling are:

- Easy migration from the previous PCI parallel interface
- Lower implementation cost compared to the PCI architecture
- Wide interface bandwidths
- Faster data transfer speeds that have scaled up over time to keep pace with board technology advancements
- Seamless backward compatibility between PCI Express iterations

The demonstrated PCIe success in connecting a system host board to the I/O card slots within an industrial rackmount computer or server has started a trend. This trend has been developing for the past few years and it involves figuring out ways to have PCI Express break out of the individual computer enclosure and function as a high-speed, wide bandwidth, scalable serial interconnect of several meters in length between control platforms and I/O, data storage or other boxes within an I/O rack . Other PCIe over cable applications include GPU Computing and video display wall controllers. Figure 1 provides a good illustration of this trend that comes under the umbrella term "PCIe Over Cable."

Figure 1 – PCIe Over Cable

Several industry standards have been developed by the PCI-SIG to cover this PCI Express over cable implementation. The base standard is known as "PCI Express External Cabling Specification," and there are several versions available covering the existing PCIe 1.x, 2.x and 3.0 variations with others in the works to cover future PCI Express versions.

PCIe Over Cable Link Tuning Parameters

The promise of PCIe over cable is a simplified way of expanding the I/O capabilities of an existing host server. To make good on this promise, proper PCI Express link tuning and design fundamentals must be followed when designing either a PCIe expansion uplink or downlink card. Regardless of PCI Express version, engineering a successful PCIe over cable installation requires that PCIe link-tuning fundamentals be properly implemented when designing a PCI Express I/O or data storage expansion system. PCIe link tuning parameters must be adaptable in order to compensate for the system variability common in today's host servers. In the following sections PCI Express over cable link tuning illustrations and application examples will be shared that demonstrate effective PCIe I/O expansion chassis implementations.

The companies making up the PCI-SIG have created the interface standards and design guidelines necessary to develop PCI Express plug-in cards that function regardless of the host system. These industry standards clearly define PCIe link parameters such as timing jitter, amplitude noise and acceptable bit error rates. The engineer's host (root complex) or device card (end point) design must adhere to these published PCIe link parameters in order for a board to function properly in a commercial motherboard or in a card slot on an industrial backplane connected via a PCIe interconnect to a system host board. This same set of link parameters applies to PCIe Over Cable designs, but with the added wrinkle of needing to compensate for cable length, cable type and PCI Express version i.e. the data rate as defined in Figure 2.

	Data Rate* (per lane/per direction)	Transfer Rate* (per lane/per direction)
PCIe 1.x	250MB/s	2.5GT/s
PCIe 2.x	500MB/s	5.0GT/s
PCIe 3.0	1.0GB/s	8.0GT/s

Figure 2 – PCI Express Data Rates

***NOTE**: The values in Figure 2 represent maximum theoretical PCI Express data speeds and data transfer rates. The PCIe 1.x and 2.x interfaces use 8b/10b coding while the PCIe 3.0 interface utilizes 128b/130b coding. This results in 20% protocol overhead with the PCIe 1.x and 2.x interfaces and a 1.5% overhead with PCIe 3.0 that will decrease the maximum practical transfer rate of any PCI Express implementation.

Each PCIe version doubles the effective data transfer rate of the preceding version. PCIe 3.0 combines protocol changes as well as transmit and receive speed enhancements to achieve this doubling of data throughput.

Proper PCIe link tuning becomes critical as PCIe data throughput speed increases. Due to the full-duplex nature of PCI Express, the desired PCIe cable length also plays a deciding roll in link tuning. Each PCI Express link over copper is made-up of a twisted wire pair. For a x16 PCIe interface between the host system and the expansion or target chassis there are 16 twisted wire pairs plus sideband signals needed for both the transmit and receive channels within the PCIe copper cable. A single x16 PCIe cable would be expensive, lack flexibility and have limited practical use due to the required overall cable diameter needed to house all 32+ wire pairs. For this reason x16 PCIe cables are made up of four x4 cables consisting of four twisted wire pairs each per direction. These channels can exhibit high frequency loss as a natural result of skin effect along with the cables' dielectric and capacitive loses.

Fortunately, there are compensation parameters that the board designer can build into the target and host card design to minimize skin effect and provide the clean interface signals needed to deliver fast and reliable communication between the remote PCIe option card and the host server.

PCIe Link Tuning Methodology

PCI Express transmit and receive compliance eye diagrams are useful to the PCI Express board and PCIe over cable system product designers. The acceptable range of various link voltages and signal frequencies as they relate to timing jitter, amplitude noise and acceptable bit error rates are defined in the PCI Express specifications. The eye diagram in Figure 3 is an example of an acceptable eye diagram for a PCIe 2.0 x16 link being driven over a 3m long cable between a Trenton PEU8039 uplink/host card installed in 1U H-P server and a Trenton PED8044 downlink/target card installed in a 14-slot industrial backplane commonly used in a 3U rackmount PCI Express I/O expansion chassis like the Trenton TTX3100 shown in Figure 4. Notice the clear area around the diamond-shaped reference mask in Figure 3. This indicates proper peak-to-peak signal voltages at the desired signal frequency.

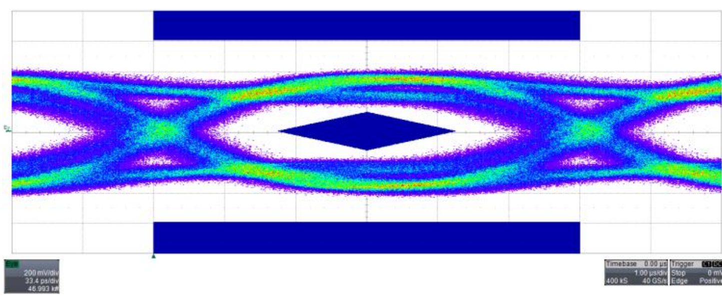

Figure 3 – Good Eye

Figure 4 – 3U Rackmount PCI Expansion Chassis

Conversely, a bad eye is one in which the link voltages and signal frequencies as they relate to timing jitter and amplitude noise fall out of spec resulting in an unacceptable bit error rate on the PCI Express link. Figure 5 is an example of a bad compliance eye diagram at a downstream or target device.

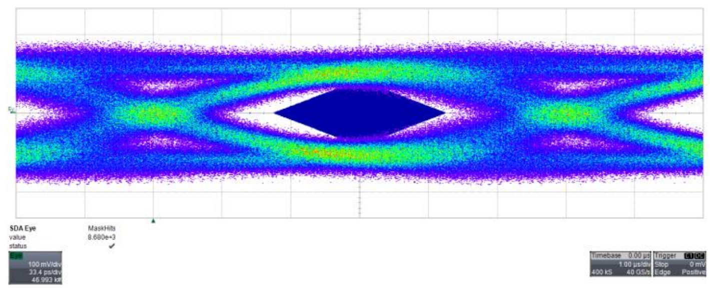

Figure 5 – Bad Eye

Notice how the PCIe signal pattern is encroaching into the reference diamond. If the devices manage to link with a compliance eye of this type the user could experience a larger number of data packet errors and link retries resulting in an unacceptable bit error rate. Slower interface performance may also occur because the PCIe link may train down to a x1 link in order to establish communications. The link may also need to resend a lot of data packets thereby degrading overall interface performance. A compliance eye that is this poor will likely not allow the host sever to establish a link with the PCI Express expansion chassis.

PCIe Link Pre-Emphasis and Equalization

There are a couple ways to compensate for this signal loss over the cable. One can amplify the high frequency component of the signal at the receiver or device target end of the cable using a process called: Equalization. However, amplifying the receive end alone runs the risk of increasing the deterministic jitter which will in turn close down the compliance eye. The process of pre-emphasis avoids this noise induction problem. Board and system designers use the process of pre-emphasis to place more emphasis on the transmit voltage signals in order to ensure robust host-to-target device communications over the length of the PCI Express cables. Pre-emphasis opens up the compliance eye up at the target device as illustrated in the transmit voltage differential signal patterns in Figure 6.

Pre-emphasis enables a more judicious use of equalization at the receiver-end of the cable. Employing pre-emphasis and equalization together opens up the compliance eye at the target device even further.

Figure 6 – Pre-Emphasis Example

Summary & Conclusion

So does a customer need to purchase an expensive PCI Express protocol analyzer and fiddle around with device parameters like pre-emphasis and equalization to get a good PCIe signal sent over a copper cable to an I/O expansion chassis? Not at all! These parameters are engineered into the PCIe components used in the target and host boards and optimized for a variety of application scenarios. There may be diagnostic LEDs and dip-switches incorporated into the PCIe host and target expansion card to enable a user to make field adjustments as necessary to decrease the rate of data re-tries, but these adjustments will only be required in extreme applications such as using maximum cable lengths with high-speed PCI Express 3.0 applications. In general the default settings of the dip switches should meet most system requirements. It's unlikely that these dip-switch adjustments will be needed with the majority of PCI Express over cable installations, but it's nice to know that they are there if needed.

PCI Express expansion chassis can be implemented quickly and effectively to expand the I/O card capabilities of an existing host server with proper understanding of the underlying engineering principles of PCI Express over cable. An I/O expansion system using PCIe over cable will extend the capabilities of an existing host server at a significant cost savings compared to installing all new host and I/O hardware and software.

For Additional Information

Contact us for more information on the Trenton PCI Express over cable expansion systems or any of our other embedded computing products or integrated computer systems. Trenton team members welcome the opportunity to work with you.

Trenton Systems

Phone: 770.287.3100
Toll Free: 800.875.6031 (U.S.)
Fax: 770.287.3150
E-mail: Info@TrentonSystems.com
Web: www.TrentonSystems.com

Follow Trenton on:

About Trenton Systems

Trenton Systems is a designer and manufacturer of PCI Express expansion systems, rackmount computers, video capture and display wall controllers, GPU computing systems, custom systems, embedded motherboards, single board computers, system host boards, and backplanes for critical embedded computing applications. All Trenton products are designed and built in the U.S. and are well suited to deliver long-life performance, precision and reliability in military, DoD, medical instrumentation, industrial control, and video capture and display wall systems.

Trenton industry partnerships with Intel, Matrox Graphics, Nvidia and other leading technology companies play an important role in the design, manufacture, integration and deployment of our high-performance system solutions. For example, Trenton is a member of the Intel® Intelligent Systems Alliance, a community of communications and embedded developers and solution providers. Trenton board level products and integrated computer systems are used to create a variety of mission critical solutions, from Government and Defense, to Industrial Automation, Virtualization, Video Processing, Medical, Communications, Energy, GPU Computing, Test & Measurement and Video Display Walls.