Application Brief

Implementing the PCIe Design with Signal Conditioning Devices

David Creger, Nasser Mohammadi

Introduction

An increased Peripheral Component Interconnect Express (PCIe) bus speed is important for applications that use highly computer-intensive peripheral components that reside on the PCIe bus (for example, SSD, GPU, FPGA, and ASIC devices). To overcome design challenges of signal reach and signal quality with high bandwidths of PCIe Gen 4 at 16 GT/s and PCIe Gen 5 at 32 GT/s, PCIe signal conditioning devices are implemented in the system to reduce the design complexity.

TI offers various PCIe signal conditioning devices, including multiplexers, redrivers, and retimers. Here are the most commonly asked questions when a designer implements the PCIe design with these devices.

When Should a PCIe Retimer be Implemented?

A redriver is a simple device used to provide up to 16 dB of range extension for systems with moderate amounts of insertion loss. It utilizes a continuous time linear equalizer (CTLE) to selectively boost high frequencies of a signal which becomes attenuated due to the low pass frequency response of a transmission line. This can open an eye that is closed due to inter-symbol interference (ISI).

A retimer is a more advanced device that does the job of a redriver and more. It can be used to correct even larger amounts of insertion loss, up to 28 dB for PCIe Gen 4, as well as random jitter. These devices are used when signal integrity has been heavily degraded. In addition to a CTLE, retimers contain a clock and data recovery (CDR) unit, decision feedback equalizer (DFE), and transmit FIR filter. A recovered clock is used to completely regenerate the signal, so a retimer is able to completely reset the jitter budget of a system. Also, unlike redrivers, retimers can participate in PCIe protocol negotiation and adaptively change CTLE and FIR settings. This comes at a cost of causing potential interoperability issues, as well as increased power consumption, added thermal management, overall product reliability, and a more expensive part. Therefore, a redriver is recommended before seriously considering a retimer solution.

In summary, a redriver is a cost effective solution to deal with moderate amounts of insertion loss (up to 16 dB). For systems with more than 16 dB of insertion loss or other forms of jitter, a retimer can provide the necessary correction. However, this comes with the tradeoff of increased complexity, cost, and power consumption as noted earlier. For more information on the differences between retimers and redrivers, please reference Signal Conditioning functions go mainstream in PCI Express Gen 4.

How Should a PCIe Redriver be Tuned to Get the Best Performance?

Recent advances in modeling allow us to perform preliminary signal integrity analysis before actually completing our schematic design. Using redriver IBIS-AMI device models along with connector and transmission media S-parameter models, we can get a good overview of the signal integrity and eye margins. It is highly recommended to do system modeling to ensure signal integrity is kept within the required specifications.

Secondly, once a board is back from fabrication, there are tools such as TI Eye Scan or other IO margining utilities to check eye opening in real time – allowing further fine tuning of each link as needed.

The *How to Tune DS160PR410 for Best Signal Integrity* application report provides guidance on using the DS160PR410 for best signal integrity and maximum extension of the PCIe 4.0 links.

What is the Difference Between a Linear Redriver and a Limiting Redriver?

A redriver can be classified into two different types: linear or limiting.

For linear redrivers, the output signal amplitude is a near linear function of the input signal amplitude for a majority of the input range. This is not the case for limiting redrivers, like the DS125BR401 which include additional features like de-emphasis control that make the signal output non-linear. Figure 1 shows how de-emphasis is used to compensate for further loss on an output channel, by attenuating low frequency components of the signal.

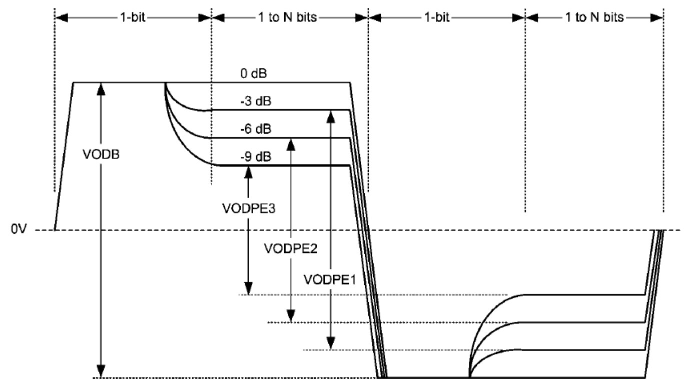

Figure 1. De-Emphasis Example

With the additional signal conditioning capabilities, a limiting redriver can help achieve compliance of industry specification even if the source signal itself is not compliant.

However, the de-emphasis provided by the limiting redriver will interfere with link training protocols. A linear redriver, on the other hand, will not interfere with link training.

Please view this TI Precision Labs video to learn more.

When Should a Passive Multiplexer or an Active Multiplexer be Used?

A multiplexer is a device that allows us to select from one of several inputs and send the data to a singular output.

A passive multiplexer, such as the TMUXHS4412, forwards the selected input data straight to the output, without any modification.

An active multiplexer, such as the DS160PR412, DS160PR421, or DS160PR822, provides the same function as the passive mux but includes a linear redriver. This allows the system to operate with larger amounts of insertion loss, but requires more power than a passive mux. For more information about redrivers, see When Should a PCIe Retimer be Implemented.

Can Multiple Redrivers or Retimers be Cascaded?

In general, it is not recommended to cascade multiple redrivers. A redriver amplifies high frequency content (including noise) and also introduces minimal random jitter. Also, it can be difficult to tune cascaded redrivers.

On the contrary, it is possible to cascade up to 2 retimers per PCIe specification. Adding additional retimers can cause jitter peaking due the PLL loop bandwidth frequency response.

The PCIe standard requires ASICs to operate on up to 28 dB of loss. A redriver can provide an additional 16 dB of reach extension, while a retimer can provide up to 28 dB of reach extension. This means a retimer will allow a system to operate with a total of 56 dB of loss. This suffices for the majority of applications. If a system includes more than 56 dB of total channel loss, it is first recommended to use a lower loss PCB material before adding an additional retimer. Only if the channel loss cannot be reduced below 56 dB should two cascaded retimers be used.

What are the Advantages of Using PCIe Instead of Ethernet Protocol in Automotive Applications?

Peripheral Component Interconnect Express (PCIe) has been used in automotive systems for many years. It offers several advantages over alternate forms of communication such as Ethernet.

First, PCIe provides ultra-low latency performance, which is an absolute must for safety critical applications in the automotive industry. By incorporating minimal data overhead, PCIe can achieve latencies on the order of tens of nanoseconds. Comparing this to Ethernet which relies on software-level handling in the TCP/IP layer and can have a latency of several microseconds, PCIe provides a huge advantage in situations where every nanosecond matters.

Another huge advantage that PCIe offers is Direct Memory Access (DMA). It allows the processor to directly access shared memory without requiring additional CPU processing resources. While other interfaces require several CPU cycles to access, copy, and buffer memory data from another domain, PCIe allows processors to directly access shared memory as if it were locally available. This further reduces the latency compared to other communication protocols.

There are several other key advantages to using PCIe in automotive applications, such as bandwidth scalability, and ecosystem breadth. Please refer to the following technical white papers titled *Connecting Zonal Automotive Architectures with PCIe* and *Implementing Native PCIe Interconnects Over Automotive Cable Channels* to learn more.