Technical Article

3 Myths about PCIe Redrivers

Connie Lu

For nearly two decades, the Peripheral Component Interconnect Express (PCIe) specification has been the interconnect standard for both consumer and industrial applications, enabling high-speed interface between processors, solid-state drives, and various technologies or end points. You can attribute its widespread adoption to several factors:

- The PCIe ecosystem is a nonproprietary specification, with open resources that many designers are comfortable working with.
- The PCIe communication protocol for different serial peripheral interfaces has ultra-low latency and bandwidth scalability. PCIe throughput per lane has continuously doubled with each generation, with the flexibility to scale to different link widths. For example, it is possible to bifurcate or split a16-lane (x16) PCIe 5.0 configuration with 64-GBps throughput to a four 4-lane (x4) configuration with 16-GBps throughput, or a sixteen 1-lane (x1) configuration with 4-GBps throughput. This flexibility makes it possible to use PCIe across many different types of applications, ranging from a single-lane communication all the way up to 16-lane configurations.

Almost all technology that needs to process data now requires faster data rates, requiring designers to overcome signal-integrity challenges such as complex thermal designs and signal degradation before they can meet the PCIe specification. Vias, poor printed circuit board (PCB) materials, connectors, and cables have a significant effect on data margins, resulting in errors in the data received. Adding signal conditioners such as redrivers into the PCIe link helps achieve data transfers that are error-free and more reliable.

In this article, I'll explore three common misconceptions about PCIe redrivers.

Myth No. 1: PCIe Redriver Interoperability Is Challenging Given the Large Variety of Different End Points.

The Peripheral Component Interconnect Special Interest Group (PCI-SIG) governs a standardized specification for PCIe data transfers. The PCI-SIG comprises more than 1,000 different company members that specialize in developing high-speed interoperable products, and hosts compliance workshops where various companies can confirm that their products meet the most important aspects of the PCIe specification. This process provides extensive interoperability testing with other root complexes or end points so that companies can understand the full capability of their PCIe signal conditioners. For example, TI's DS320PR810 redriver passed the PCI-SIG compliance test with different root complexes and end points. While internal testing can validate key aspects of PCIe specifications, passing the PCI-SIG compliance test is more solidified proof of the product's capability.

Myth No.2: Finding the Right Setting for PCIe Redrivers Is a Very Manual Process.

Linear redrivers are simple to design because they don't require configuration or initialization on the device's physical layer, or port. Finding the right continuous time linear equalization (CTLE) setting doesn't require trying out all of the CTLE and gain settings manually. There are multiple tools and methods to help hardware designers find the most optimal setting quickly.

Designers can leverage the Input/Output Buffer Information Specification algorithmic modeling interface (IBIS-AMI) to simulate the pre- and post-channel effects on the redriver and provide the resulting eye diagram close to real link characteristics. The simulation model represents detailed analog electrical behaviors of the entire channel, including transmitters and receivers. This simulator also enables users to sweep through the entire CTLE settings and find the best setting.

Another method for tuning a redriver is to measure the insertion loss of the transmission channel at the fundamental frequency through a network analyzer and find the CTLE setting compensating for the insertion

loss measurement. For example, you may measure the insertion loss at the input of the redriver at 16 GHz when designing a PCIe 5.0 solution. If the insertion loss is 10 dB, the starting point for the CTLE setting should allow a gain of +10 dB to compensate for the insertion loss. Generally, this tuning method helps you identify appropriate initial CTLE settings quickly and tune for the best margin in the signal path.

Myth No. 3: PCIe Redrivers Should Be Placed in the Middle of the Link and Have Less Design Flexibility.

Many different factors can cause signal degradation in a system. Between various PCB materials and trace lengths, the number of vias used for signal routing, and different connector characteristics, the pre- and post-channel loss of a redriver can vary. Where you place redrivers depends on the loss profile of your system – the placement does **not** have to be in the middle of the link to achieve the best solution. Some redrivers come in different configurations with different channel counts to enable maximum flexibility in designs. Having a unidirectional configuration instead of a lane configuration enables designers to place the redriver as close to the end point as possible, achieving a longer reach.

In applications with limited PCIe slots, some redrivers have built-in multiplexing configurations that can support multiple root complexes and endpoints. For example, TI's SN75LVPE5412 linear redriver can support one root complex such as a central processing unit communicating with two solid-state drives, with SEL pin multiplex data fetching between the two end points. This configuration is popular in space-constrained PC applications. And because redrivers have small form factors and are low power, they not only address space constraints from a communication standpoint, but also remove the need for heat sinks or bulky pads, thus meeting board height requirements easily.

Additionally, because redrivers have really low latency – hundreds of picoseconds – you do not have to consider or wait for the delays between root complexes and end points, freeing up processing capacity to execute more tasks efficiently. Increased processing capacity without latency limitations from signal conditioners gives you more flexibility in your designs.

Let's look at a redriver in a real-life application.

In Figure 1a, the pre-channel loss is significant, with a closed eye diagram. This means that the root complex or end point has no way of understanding any meaningful data. Your options are to either upgrade your PCB to an ultra-low-loss material – which is more costly – and potentially open the eye, or to add a redriver into the link. After adding a redriver into the link, Figure 1b shows the open eye diagram indicating improved signal quality, with plenty of signal margin.

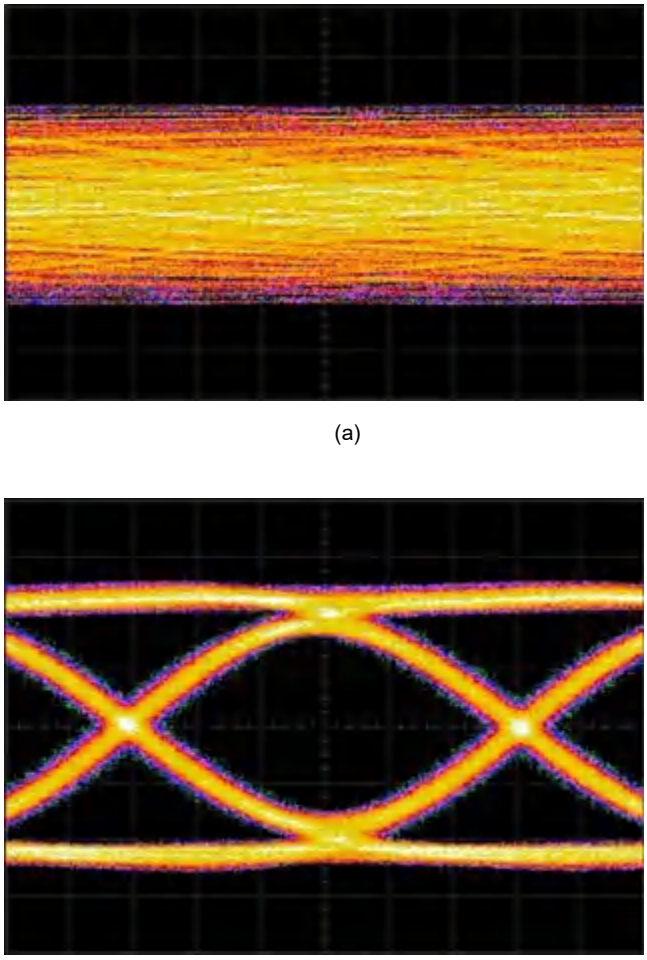

(a)

(b)

Figure 1. Pre- (a) and Post-channel (b) Signals of a PCIe 4.0 Redriver

Conclusion

Because PCIe redrivers are protocol agnostic, they can support any protocol that have PCIe based hardware such as CXL, SAS, and SATA within the data rate requirement. With multi-protocol support, redrivers can be used across many different high-speed applications that require signal conditioning. While a higher-cost PCB material with very low loss can achieve an overall improved signal across all the links on the board, you can add a redriver as-needed basis when only a few links in your system need more margin. Without having to upgrade the entire PCB material, adding a redriver can also provide more cost savings and serves as a simple alternative method to improve signal quality.

Additional Resources

Check out TI's PCIe signal conditioning solutions for more about signal degradation challenges

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATA SHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, regulatory or other requirements.

These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you will fully indemnify TI and its representatives against, any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI’s products are provided subject to TI’s Terms of Sale or other applicable terms available either on ti.com or provided in conjunction with such TI products. TI’s provision of these resources does not expand or otherwise alter TI’s applicable warranties or warranty disclaimers for TI products.

TI objects to and rejects any additional or different terms you may have proposed.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265