3 Ways Linear Redrivers For PCIe 3.0 Differ from ReDrivers

Pericom Semiconductor Corporation 1545 Barber Lane Milpitas, CA 95035

For PCIe 3.0, linear redrivers do the best job walking the line.

Designers know that redrivers are an essential part of solving signal integrity challenges in gigabit systems. But at speeds up to 8 GT/s in PCI Express 3.0 systems, *linear* redrivers are needed to also comply with PCI-SIG specs and with chipset vendors' datasheets. That is, it's not enough to open the signal eye—the redriver also needs to precisely match PCI Express Link Training waveforms while offering designers the ability to "tweak" channel characteristics (Figure 1).

OPEN EYE

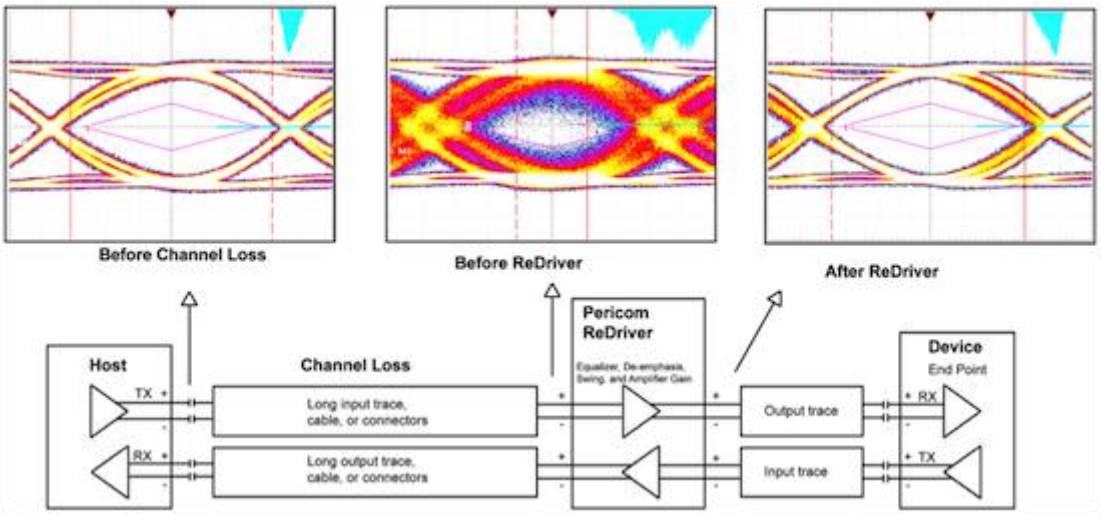

Figure 1: Redrivers amplify and shape high frequency signals to compensate for channel "challenges."

Redrivers Revisited

The continuum of super high-speed, serial communications standards is visually shown in Figure 2, a graphic from integrated circuit (IC) connectivity supplier Pericom Semiconductor. We have previously covered the issues with and mitigation strategies for dealing with signal integrity (SI) in serial standards in the article "Signal Integrity in a GHz World." The reader is encouraged to refer to that article for a refresher.

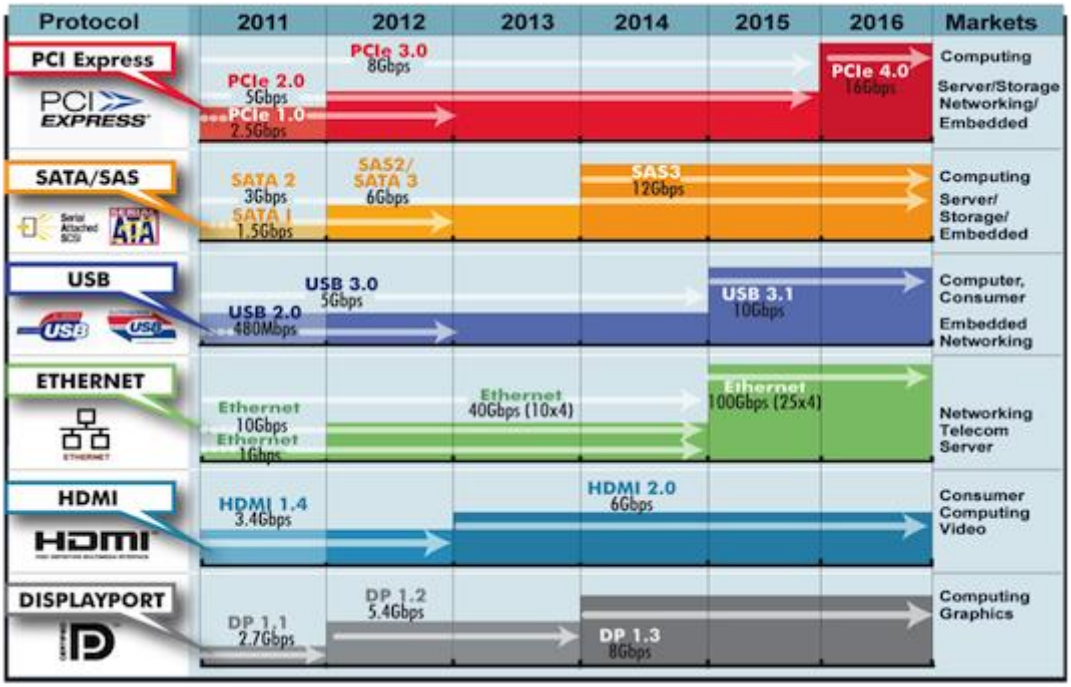

Figure 2: Serial communications standards just keep getting faster. Along with signal integrity challenges also come specification compliance issues such as PCI Express Link Training. (Courtesy: Pericom Semiconductor.)

Here, this article builds on the previous material by arguing that while redrivers *absolutely* solve the problems as described, as serial standards get faster—and IC chipsets and system design specifications evolve even more—there is a need for a new kind of redriver.

The amplifier portion of the linear redriver is designed to have a "more linear" transfer function (Vout/Vin) than the limiting redriver (Figure 3). While both limiting and linear redriver amplifiers are analog in nature, their gain slopes are very different. This is necessary to comply with *Link Training* protocols now becoming essential in PCI Express 3.0 and related Gen 3.x standards.

What is Link Training?

Driven by Intel and introduced with PCI Express 3.0, Link Training is part of the complicated in-band PCIe enumeration sequence that not only "connects" Source to Destination nodes, but also determines the ideal speed for the channel. Since it's a given that the complete channel from Source to destination will experience signal integrity (SI) issues (refer to "Signal Integrity in a GHz World"), Gen 3 seeks to optimize the transfer speed and minimize bit error rate (BER) based upon the channel's capability.

Note that "channel" refers to the sum of the complete interconnect *and boundaries* between the Source IC's pins and the Destination IC's pins or leads. For all practical purposes, this includes PCB traces and vias, board connectors, cabling, backplanes and any other electrical path for the 8 Gbps signals. In fact, the point of Link Training is to verify if the channel is capable of this speed or if the speed should be reduced to PCIe Gen 2 speed—or less.

The Link Training protocol passes a series of waveforms from the transmitter (Source) to the receiver (Destination) across the channel according to a pre-defined state machine sequence. If all of the Link Training protocol waves arrive at the Destination as expected, the channel is deemed capable of full speed. As waves degrade or are not received as expected, this infers information about the channel's speed capability for Gen 3, Gen 2, or even Gen 1 (the slowest) PCIe speed.

Link Training signals use various pre-emphasis and pre-shoot waveforms to represent the different presets. For example, Table 1 shows how the various Link Training Tx signals contain pre-shoot and de-emphasis, which translate to emphasized portions of the input waveform as seen by the linear redriver (Figure 3). Because of these intentional waveform characteristics, the linear redriver *must* reproduce these various waveforms without any changes—such as limiting (clamping) the signal—which would destroy the Link Training waveform as it passes through the redriver.

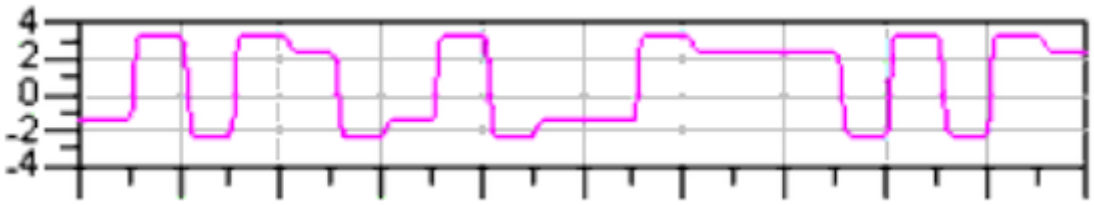

Figure 3: This is an example of a Link Training Tx waveform corresponding to preset P03 with 2.5dB of pre-emphasis.

Preset	Preshoot	Deemphasis
P0	0.0 dB	-6.0 dB
P01	0.0 dB	-3.5 dB
P02	0.0 dB	-4.4 dB
P03	0.0 dB	-2.5 dB
P04	0.0 dB	0.0 dB
P05	1.9 dB	0.0 dB
P06	2.5 dB	0.0 dB
P07	3.5 dB	-6.0 dB
P08	3.5 dB	-3.5 dB
P09	3.5 dB	0.0 dB
P10	0.0 dB	-9.5 dB

Table 1: Examples of Link Training waveforms and associated pre-shoot and de-emphasis. A linear redriver is needed to precisely reproduce the waveform characteristics. Source: PCI-SIG.

Any passive or active impedance, interference, or device in the channel has the potential of *adversely* modifying these waveforms. While the channel itself "is what it is", the redriver added to the channel to fix signal integrity must at the same time not modify the waveforms in any way!

Therefore, while non-linear redrivers *regenerate* high-speed signals to open the signal eye to compensate for channel SI effects, *linear* redrivers precisely *reproduce* high-speed signals and open the signal eye. With linear redrivers, all of the characteristics of each Link Training waveform are completely reproduced, including amplitude. Hence, the linear redriver compensates for SI challenges without affecting the Link Training protocol waveforms. The difference between linear and non-linear redrivers is shown in Figure 4.

As mentioned above and shown in Figure 4, you can see that the limiting amplifier has much higher gain so that the output signal is driven to its maximum more quickly from the input signal, resulting in cutting off the input waveform. A useful analogy is the compressor/limiter in the audio world, that clamps the output

voltage to some limit regardless of how hard the input signal is driven. The linear redriver, on the other hand, has a lower gain, no limiting threshold and therefore allows the output to track the input without clamping the signal to some limit.

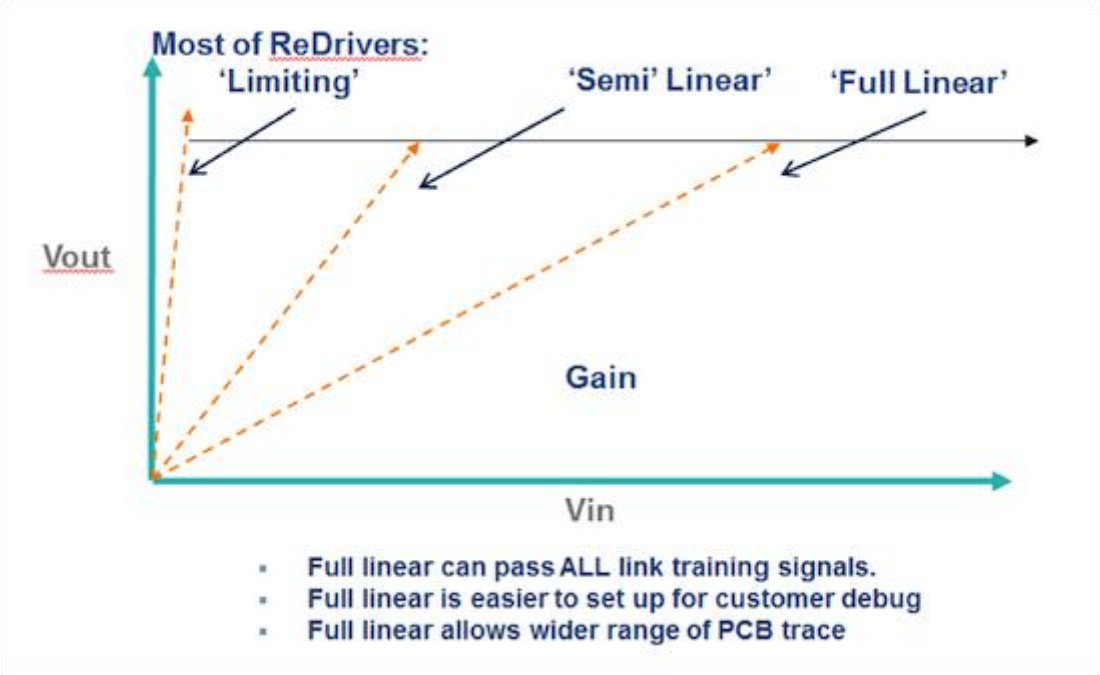

Figure 4: The difference between linear and regular (non-linear) redrivers. Non-linear redrivers compensate for signal integrity problems, but regenerate signals that may not pass PCI Express 3.0 Link Training waveforms.

Waveform Details: Comparison

It's important to understand that while redrivers will "fix" SI problems, they do so by regenerating signals as shown in the Figure 5 graphs provided by Pericom Semiconductor. Notice that the top waveforms in Figure 5 (non-linear redrivers) show some pre-emphasis waveform attenuation. This is just fine for opening the signal eye downstream; however, changing the waveform *even just a little bit* can cause PCIe Link Training to fail.

The result may be that the channel is deemed insufficient for PCI Express Gen 3 speeds (8 Gbps) and the Source/Destination agree to communicate a slower Gen 2.0 (5 Gbps) speed. Conversely, the bottom waveforms in Figure 5 show how a linear redriver precisely reproduces the input waveform by accurately tracking the source signal. The channel would pass Link Training, and the devices would communicate at the fastest possible channel speed.

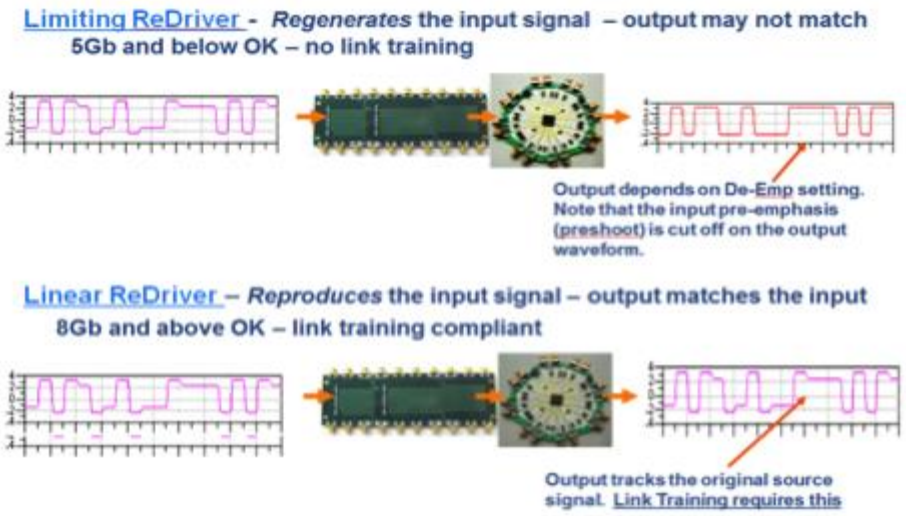

Figure 5: Redriver vs linear redriver waveforms, according to a Pericom Semiconductor lab test.

Compatibility and Interoperability Testing

The notional waveforms shown in Figure 5 illustrate the differences between a redriver—sometimes referred to as a *limiting redriver*—and a linear redriver. Figure 6 shows the logical block diagram differences. In the linear case, the final analog amplifier stage drives the differential channel directly and provides channel equalization, whereas the more common redriver adds a de-emphasis (DE) stage after the amplifier. The DE stage is needed to add back the emphasis that was limited (chopped off) at the input side of the limiting redriver.

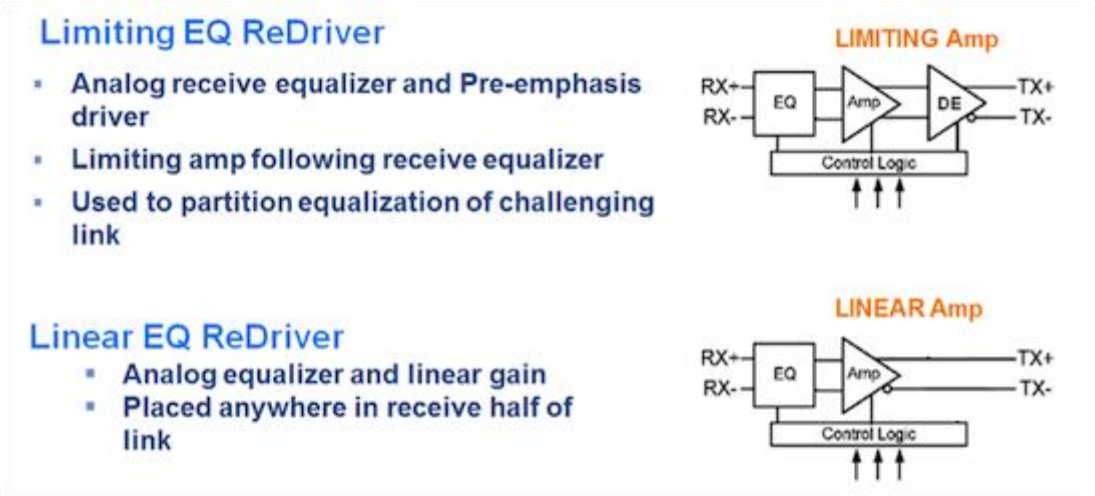

Figure 6: The block diagram difference between regular (limiting) redrivers and linear redrivers. The final analog stage makes all the difference when PCI Express Link Training is required. (Courtesy: Pericom Semiconductor.)

What's important with linear redrivers is the accuracy with which they reproduce PCI Express Link Training waveforms. This is measured by plugfests organized by the PCI-SIG and by tests conducted with major PCIe chipset vendors. The details of these tests are often not public; suppliers will only say that their linear redrivers have passed PCI-SIG workshops and that their ICs have been tested by IC vendors.

Pericom Semiconductor, a signal integrity and connectivity IC supplier, has linear redrivers designed not only for PCI Express 3.0 at 8 Gbps, but linear redrivers for 10Gbps Ethernet (KR), USB 3.1, DisplayPort 1.3,

Thunderbolt, HDMI 2.0, and SAS 3.0. In the case of PCIe 3.0—the subject of this article—Pericom's redriver family has "passed all the tests for Link Training as measured by a very large PCI Express semiconductor company," says Bill Weir, Sr. Director of Marketing for Pericom.

According to Weir, the company spent many months conducting simulations prior to going into the lab and measuring compliance to Link Training protocol waveforms. The company shares this confidential test data with customers when requested.

The company's PI3EQ8904 and PI3EQ8908A linear redrivers are PCIe 3.0 compliant at 8 Gbps and support four PCIe lanes ('8904) and eight lanes ('8908A), but can also be used at slower PCIe speeds and Gen 1/2/3 protocols when the channel so necessitates. Beyond PCI Express compliance, there are other nice-to-have attributes that designers should consider.

The "Nice-to-Haves"

Beyond the basics—improving SI and compliance with PCI Express Link Training protocol waveforms—eval boards, simulation suites, and programmable features and applications support to assist customers in the debug and validation process are very desirable.

According to Pericom's Weir, the typical customer process is:

1) Simulation of the redriver in the customer channel, typically using S Parameter models. Pericom fully supports this important phase.

2) Use an eval board or customer prototype PCB board with redrivers to do initial lab testing of the actual application.

3) Debug and validation once final customer PCB platform is back and ready to finalize redriver settings for release to production. "Settings" means programmable IC capabilities.

Eval boards allow designers to place a manufacturer's linear redriver into real world scenarios so designers can "tune" their own circuit boards, backplanes, and complete system design. Eval boards can include just the chosen redriver, or might include PCIe traces for "what if" in-system EMI and channel noise tests. The latter might be useful to see if adjacent oscillators or RF transceivers affect a PCIe channel design. Several examples of eval boards from Pericom Semiconductor are shown in Figure 7.

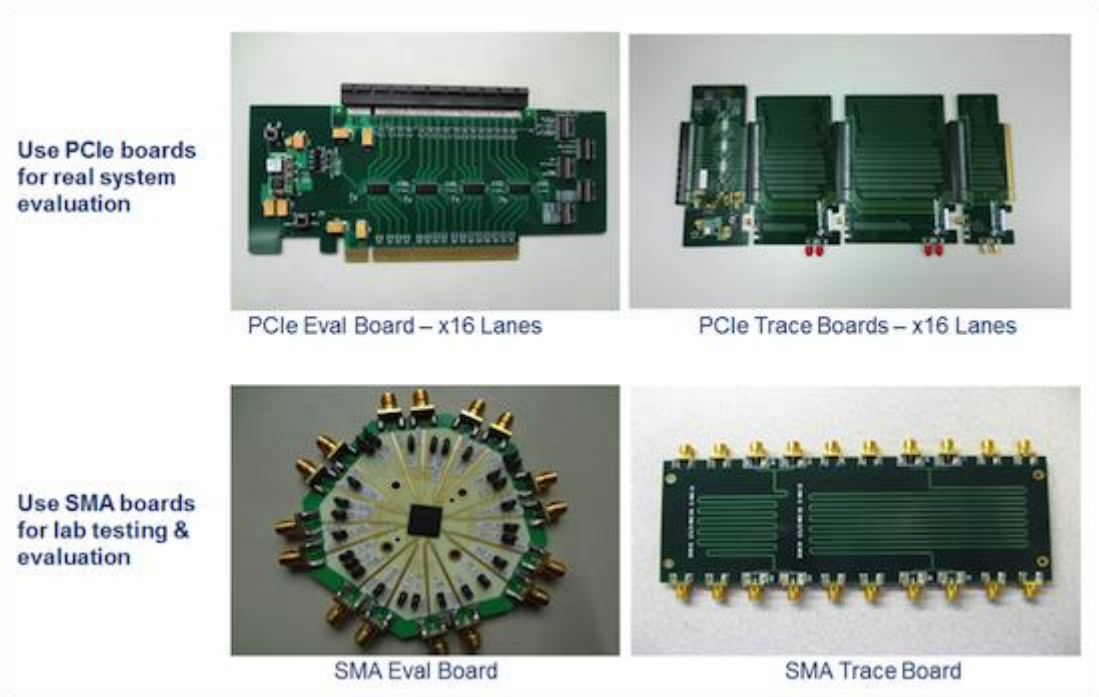

Figure 7: Examples of linear redriver and PCI Express evaluation boards. (Courtesy: Pericom Semiconductor.)

Other nice-to-haves include per-channel programmability of input/output termination, receiver equalization, output swing and overall gain. On the latter, the amplitude of the waveform can be boosted to compensate for channel losses, *but the waveform shape remains the same*. This assures that the linear redriver doesn't violate Link Training waveforms.

On the subject of programmability, a customer and cost friendly feature offered by Pericom is two-way programming. The linear redriver's attributes can be programmed via local pinstrapping, or via an I2C interface. On the latter, the redriver can be selected for either slave or master mode. In master mode, up to 16 redrivers can be programmed from a low cost EEPROM during boot up. The advantage to this programming method is that the CPU BIOS does not need to be changed in order to program the redrivers, which reduces software complexity and speeds up the design process.

As well, it's often handy to have flow-through pinouts where all inputs are on one side of the device with the outputs on the other to reduce crosstalk and also provide compatibility to double-sided PCIe connectors, such as on a riser card This allows the Rx channels and Tx channels to be routed from the connector on different sides of the PCB. Single voltage supply sources also make for easier designs and lower BOM costs by eliminating multiple DC-DC regulators. Wider temperature ranges from -40 ºC to +85 ºC are ideal for most consumer and industrial systems. For reference, all of Pericom Semiconductor's linear redrivers include these features, and more. A complete table of redrivers for PCI Express and the other high-speed channels shown in Figure 1 above can be found here.