Debugging PCIe Dynamic Link Behaviors with CrossSync PHY for PCIe

Application Note

October 20, 2021

Summary

This application note describes how to use CrossSync PHY for PCIe software with an oscilloscope and protocol analyzer to make measurements related to transmit-side PCIe link equalization. An example shows how to debug unexpected behavior at the end of the link-training process.

Adapted from the webinar How to Debug PCI Express® Power Management and Dynamic Link Behaviors by Patrick Connally and Gordon Getty

Introduction

With successive generations of PCI Express® operating at 8, 16 and 32 Gbps, dynamic link equalization becomes essential. Equalization involves the intentional distortion of a data signal to compensate for deficiencies in the communications channel. Those deficiencies include the link acting as a lowpass filter that attenuates key high-frequency components of the data stream. In addition, impedance discontinuities in the link caused by connectors and vias can further degrade the link performance. PCIe® equalization can be applied at the transmit side (TxEQ), the receive side (RxEQ) or both. TxEQ involves de-emphasis and pre-shoot, while RxEQ involves continuous-time linear equalization (CTLE) and decision feedback equalization (DFE).

On the transmit side, de-emphasis causes the first bit after a transition to be transmitted at full amplitude (Va). Subsequent bits of the same polarity are transmitted at a reduced, or de-emphasized, level (Vb), except for the final bit before the next transition, which is transmitted at a boosted pre-shoot level (Vc). In addition, a single bit between transitions is transmitted at a maximum boost level (Vd). The combination of de-emphasis and boost adds to the signal high-frequency content that the link would attenuate. Equalization involves a multiphase link-training sequence that can sometimes yield unexpected results. The ability to correlate protocol-layer and physical-layer traces using *Cross*Sync™ PHY for PCIe can help you isolate logical and electrical problems that can appear after link training.

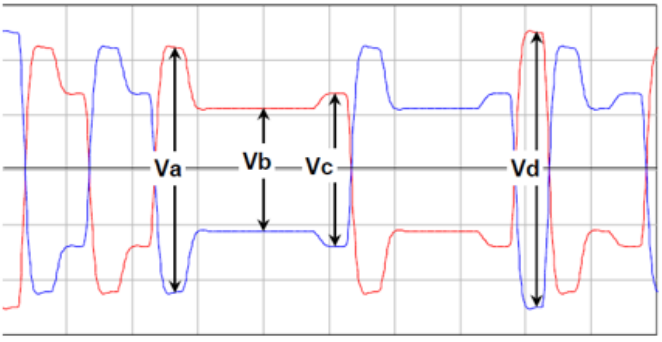

Figure 1. Transmit voltage levels and equalization ratios.
(Source: PCI Express Base Specification Revision 4.0 Version 1.0.)

Overview of the Link Training Process

For transmit-side equalization, de-emphasis, pre-shoot and boost are implemented by a three-tap finite impulse response (FIR) filter inside a PCIe system's TxEQ block. The goal of link training is to determine the optimum FIR filter coefficients, also called cursors, for a given communications link. Link training involves the exchange of ordered sets of data, including training sequence 1 (TS1) and training sequence 2 (TS2), between the downstream port and upstream port.

How PCIe Link Training Is Implemented

For example, PCIe 4.0 link training begins with a speed-change negotiation and extends from phase 0 through phase 3. In phase 0, the downstream port might send TS2 ordered sets at an 8-GT/s data rate to the upstream port, advertising a 16-GT/s maximum data rate. In phase 1, both ports exchange TS1 ordered sets, interspersing an Electrical Idle Exit Ordered Set (EIEOS) after every 32 TS1 ordered sets, to establish an operational link. The purpose of EIEOS is to guarantee that a link partner can detect the electrical idle exit state. The EIEOS packet symbols (four alternating 00 00 FF FF sequences) result in an electrical signal with regular and relatively few transitions, which can be useful for observing a signal's physical-layer properties during debug.

Presets and the Role of P10

The subsequent phases involve the exchange of data to optimize electrical performance. The PCIe standard specifies 11 predefined combinations of de-emphasis, pre-shoot and boost cursor coefficients called presets and labeled P0 through P10. During link training, a PCIe device may request either presets or cursors—the latter provide finer resolution and more setting options, while the presets provide convenience. Presets are defined in terms of voltage ratios and pre-shoot and de-emphasis coefficients in dB, with the exception of P10, which is used for transmitter boost-limit testing at full amplitude and whose boost limits are not fixed.

Preset #	Preshoot (dB)	De-emphasis (dB)	c-1	c+1	Va/Vd	Vb/Vd	Vc/Vd
P4	0.0	0.0	0.000	0.000	1.000	1.000	1.000
P1	0.0	-3.5 ± 1 dB	0.000	-0.167	1.000	0.668	0.668
P0	0.0	-6.0 ± 1.5 dB	0.000	-0.250	1.000	0.500	0.500
P9	3.5 ± 1 dB	0.0	-0.166	0.000	0.668	0.668	1.000
P8	3.5 ± 1 dB	-3.5 ± 1 dB	-0.125	-0.125	0.750	0.500	0.750
P7	3.5 ± 1 dB	-6.0 ± 1.5 dB	-0.100	-0.200	0.800	0.400	0.600
P5	1.9 ± 1 dB	0.0	-0.100	0.000	0.800	0.800	1.000
P6	2.5 ± 1 dB	0.0	-0.125	0.000	0.750	0.750	1.000
P3	0.0	-2.5 ± 1 dB	0.000	-0.125	1.000	0.750	0.750
P2	0.0	-4.4 ± 1.5 dB	0.000	-0.200	1.000	0.600	0.600
P10	0.0	Note 2.	0.000	Note 2.	1.000	Note 2.	Note 2.

Notes:

1. Reduced swing signaling must implement presets P4, P1, P9, P5, P6, and P3. Full swing signaling must implement all the above presets.
2. P10 boost limits are not fixed, since its de-emphasis level is a function of the LF level that the Tx advertises during training. P10 is used for testing the boost limit of Transmitter at full swing. P1 is used for testing the boost limit of Transmitter at reduced swing.

Figure 2. Transmit preset ratios and corresponding coefficient values. (Source: PCI Express Base Specification Revision 4.0 Version 1.0.)

In phase 2, The upstream port requests that the downstream port configure its transmitter equalization presets or cursors to compensate for the link channel deficiencies and ensure optimal performance. Phase 3 reverses the roles, with the downstream port requesting that the upstream port configure its transmitter equalization presets or cursors to compensate for the link deficiencies. After completion of equalization, the downstream port and upstream port exchange TS2 ordered sets. The link training and

status state machine (LTSSM) goes through Recovery.RcvrLock, Recovery.RcvrCfg and Recovery.Idle states, sending an EIEOS after every 32 TS1 or TS2 ordered sets before establishing the active L0 state.

Therefore, the TS2 ordered sets and EIEOS can be useful for triggering your instrumentation and zooming in on physical-layer signals to help debug link-training behavior after equalization.

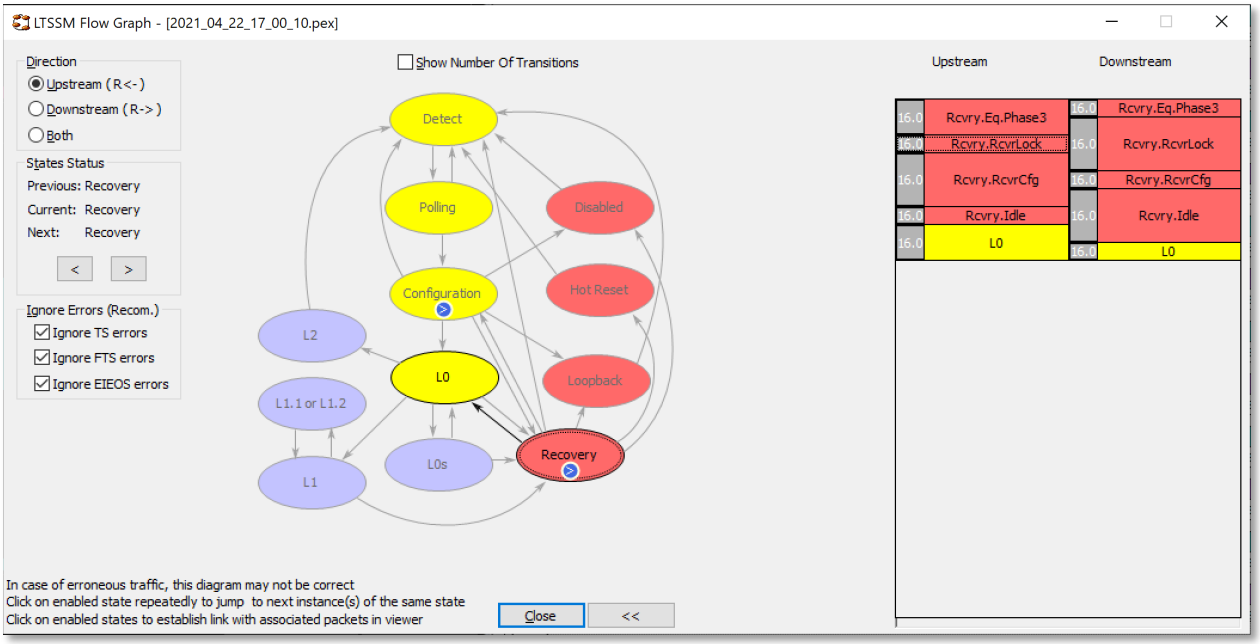

Figure 3. LTSSM establishing the active L0 state on completion of equalization.

Comparing Presets and Reported TxEQ

To validate link equalization in the real world, you can use an oscilloscope and protocol analyzer along with Teledyne LeCroy's *Cross*Sync PHY for PCIe software framework to tie the two instruments together. *Cross*Sync PHY resides on the oscilloscope and correlates data from both instruments to provide total link visibility, allowing you to view electrical waveforms from the oscilloscope correlated with protocol-layer data from the protocol analyzer. In addition, you will need a *Cross*Sync PHY-capable interposer to monitor the device under test and provide data to the protocol analyzer as well as the oscilloscope.

How to Set Up Trigger

To determine the effectiveness of the link equalization process, you will want to examine link behavior at the end of phase 3. To do that, configure the protocol analyzer to trigger on the first TS2 ordered set that occurs after the speed change to 16 GT/s, and set up the oscilloscope to capture multiple lanes of upstream traffic. This trigger setup will ensure the data is captured after the completion of the final equalization settings and the transition to the active L0 state.

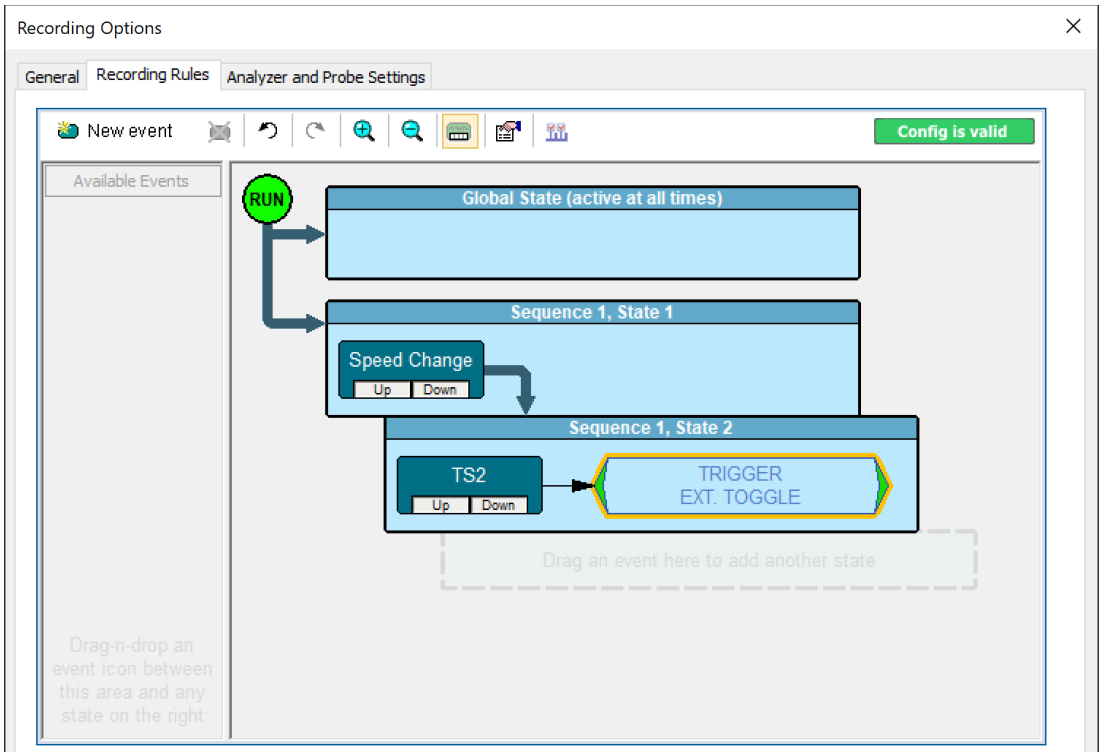

Figure 4. Protocol analyzer configured to trigger on the first TS2 ordered set that appears after the speed change to 16 GT/s.

How to Check Reported Presets

The resulting protocol trace displayed by *Cross*Sync PHY shows packet details such as packet number, ordered set, data rate and equalization control, including the preset number. *Cross*Sync PHY also displays the time-correlated oscilloscope traces, showing the electrical effects of the transmitter equalization. The oscilloscope traces in Figure 5 show a clear disparity in the electrical behavior of lanes 1 and 2 upstream signals.

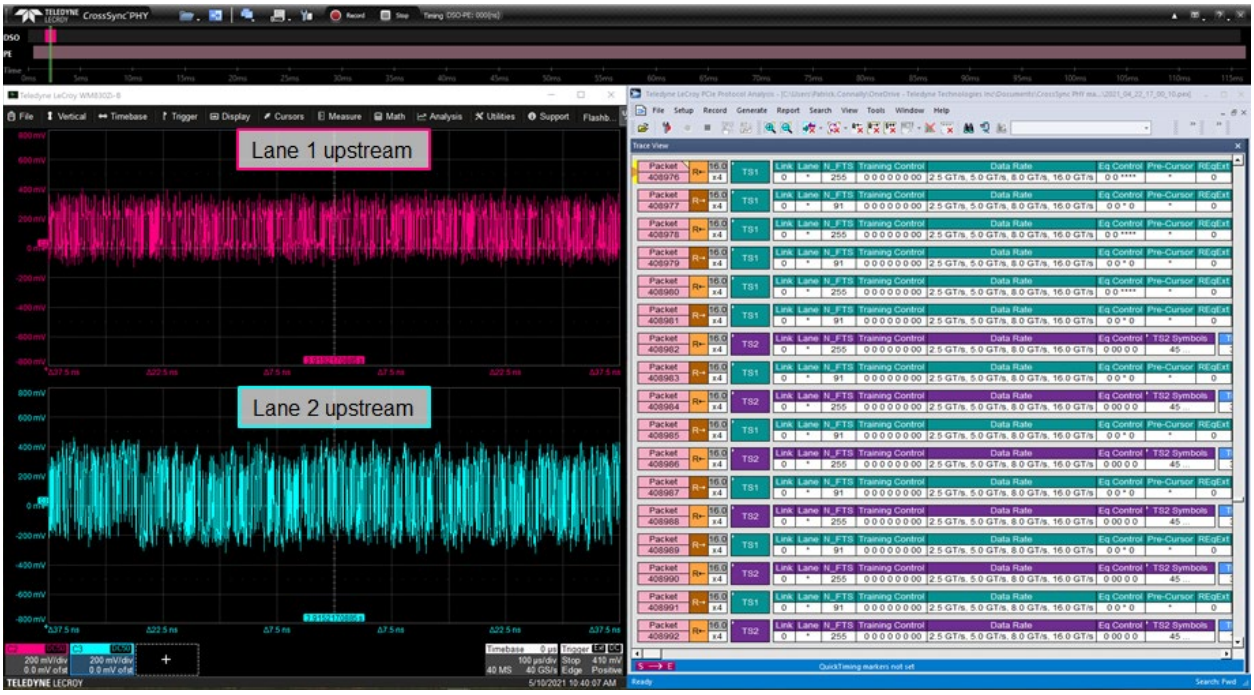

Figure 5. Oscilloscope traces showing disparity in electrical behavior between lane 1 and lane 2.

Determining Whether Problem Is Logical or Electrical

A close look at the reported TxEQ protocol-layer data at the end of phase 3 shows that lanes 0 and 2 report having trained to TxEQ preset P6, while lanes 1 and 3 report having trained to TxEQ preset P10. These results represent potentially unexpected behavior, perhaps because of one lane misreporting its status. It is not impossible for one device to train different lanes to different TxEQ presets, and P6 is a relatively common preset that many devices use during signal-quality compliance tests at 16 GT/s. However, P10 is not a preset you would expect to see being used in a live link. As mentioned previously, it exists primarily to facilitate device electrical test, and a device on the other end of the link cannot know what to expect if it requests P10.

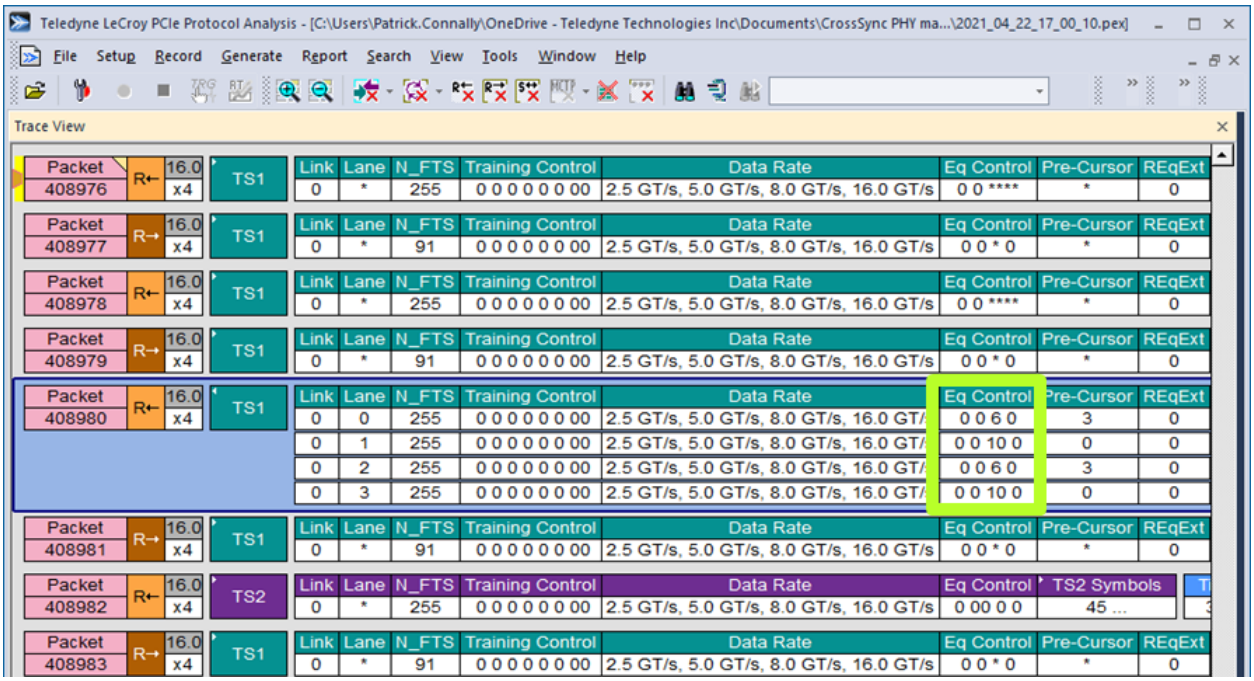

Figure 6. Lanes 0 and 2 reporting having trained to TxEQ preset P6, while lanes 1 and 3 reporting having trained to TxEQ preset P10,as highlighted by the light green rectangular outline.

Zooming Electrical Traces to Check Emphasis Levels

The question arises as to whether lane 1 is really trained to P10 or whether it is erroneously reporting that it is trained to P10. In other words, do the unexpected results indicate a purely logical problem or a logical-electrical problem? To investigate further, you can select an EIEOS packet near the end of phase 3 on the protocol trace to zoom in on the corresponding oscilloscope traces. The EIEOS packet, with its relatively few and regularly occurring transitions, lets you see on the time-domain oscilloscope traces a clear view of the differences in electrical emphasis between the two signals. As shown in the figure below, the lane reporting that it is trained to P10 shows much more emphasis placed on the signal after a transition than does the lane reporting that it is trained to P6. Further investigation would likely demonstrate that the P10 lane would have a much more closed eye than the lane trained to P6. The solution here is to examine the firmware for the logical problem that is causing the device to train to P10.

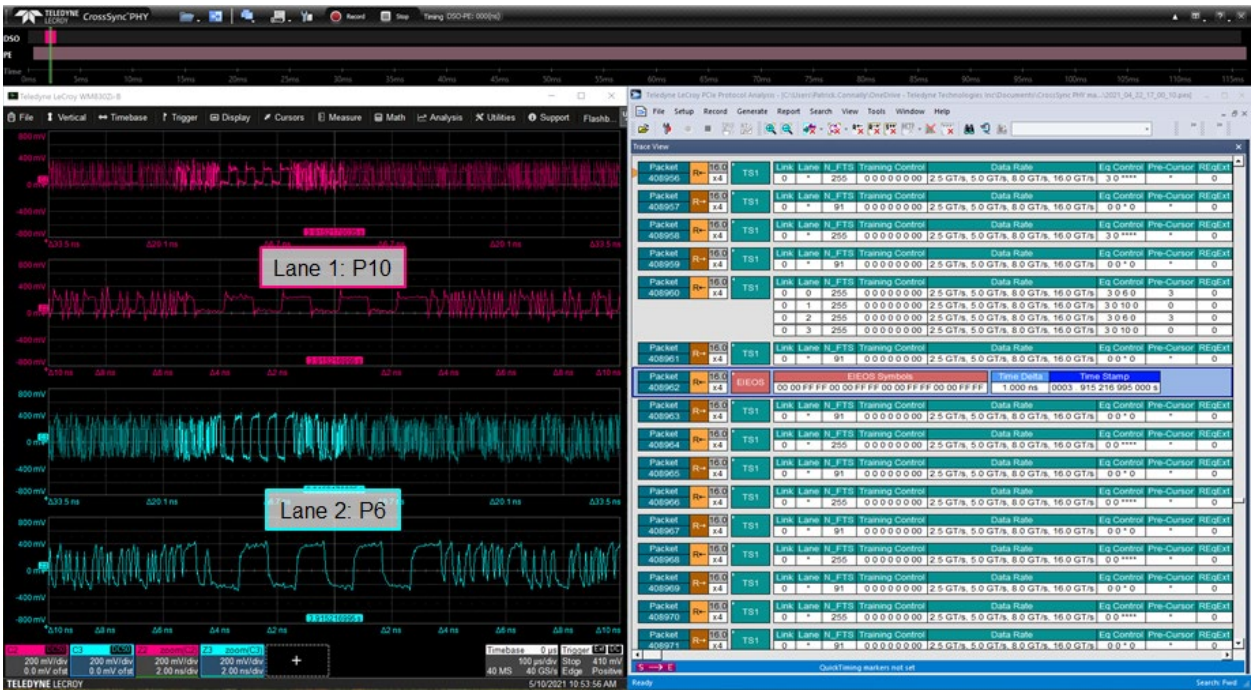

Figure 7. Zooming in on oscilloscope traces (left) correlated to an EIEOS packet in the protocol trace (right).

Conclusion

In summary, Teledyne LeCroy's *Cross*Sync PHY software framework synchronizes an oscilloscope and protocol analyzer to let you visualize, save, recall and analyze linked oscilloscope and protocol-analyzer traces to help resolve unexpected issues that can arise during the PCIe equalization process. An example related to link behavior after equalization demonstrates how to use CrossSync PHY software to debug anomolous link behavior. In addition to investigating problematic link training behaviors, the instruments and software can help characterize the entire boot sequence with visibility into sideband signals, the reference clock, data lanes and power rails. They can also help you observe speed changes in both the electrical and protocol domains.

More information about Teledyne LeCroy *Cross*Sync PHY for PCIe software can be found on our website.