

DesignCon 2017

End-to-End System-Level Simulations with Repeaters for PCIe Gen4: A How-To Guide

Yongyao Li, Huawei
liyongyao@huawei.com

Casey Morrison, Texas Instruments
cmorrison@ti.com

Fangyi Rao, Keysight
fangyi_rao@keysight.com

Cindy Cui, Keysight
cindy_cui@keysight.com

Geoff Zhang, Xilinx
geoff.zhang@xilinx.com

Abstract

The design solution space for high-speed serial links is becoming increasingly complex as data rates climb, channel topologies become more diverse, and tuning parameters for active components multiply. PCI-Express Gen-4 is a particularly relevant example of an application whose design solution space can be a daunting problem to tackle. This paper is intended to help system engineers navigate through these design challenges by providing a how-to guide for defining, executing, and analyzing system-level simulations including PCIe Gen-4 Root Complex, Repeater, and End Point.

Authors Biography

Yongyao Li is currently the principal engineer at Huawei Technologies where he is responsible for high-speed system design of leading server and storage products. He received his MSEE from University of Electronic Science and Technology of China, and has over 14 years of experience in the area of signal and power integrity.

Casey Morrison is a Systems Engineering Manager in the Data Path Solutions product line at Texas Instruments. His primary role at TI for the past eight years has been architectural definition, system simulation, and validation of high-speed interface products. His industry experience is in the area of high-speed serial communications and signal integrity spanning numerous interface standards including Ethernet, PCI-Express, SAS/SATA, and others. He received his Masters of Science in Electrical Engineering from the University of Florida.

Fangyi Rao is a master R&D engineer at Keysight Technologies. He received his Ph.D. degree in theoretical physics from Northwestern University. He joined Agilent/Keysight EEsoft in 2006 and works on Analog/RF and SI simulation technologies in ADS. From 2003 to 2006 he was with Cadence Design Systems, where he developed SpectreRF Harmonic Balance technology and perturbation analysis of nonlinear circuits. Prior to 2003 he worked in the areas of EM simulation, nonlinear device modeling, and medical imaging.

Cindy Cui is an Application Engineer with the EEsof Team at Keysight. She received her Master's degree in Micro-Electronics from Tianjin University. She has over six years of service and support experience on high-speed digital, RFIC design, and device modeling at Keysight. Her focus is mainly on DDR4 Design, PCIe, IBIS modeling and the correlation between simulation and measurement. Before Keysight, Cindy worked at Cadence for four years as the Application Engineer for Cadence simulation platform.

Geoff Zhang received his Ph.D. in 1997 in microwave engineering and signal processing from Iowa State University, Ames, Iowa. He joined Xilinx Inc. in June, 2013. Geoff is currently Distinguished Engineer and Supervisor, in transceiver architecture and modeling under SerDes Technology Group. Prior to joining Xilinx he has employment experiences with HiSilicon, Huawei Technologies, LSI, Agere Systems, Lucent Technologies, and Texas Instruments. His current interest is in transceiver architecture modeling and system level end-to-end simulation, both electrical and optical.

Introduction

The design solution space for high-speed serial links is becoming increasingly complex as data rates increase, channel topologies become more diverse, and tuning parameters for active components multiply. PCI-Express (PCIe) Gen-4 is a particularly relevant example of an application whose design solution space can be a daunting problem to tackle, especially when each link can contain three active components: Root Complex (RC), Repeater, and End Point (EP). This paper presents a how-to guide for defining, executing, and analyzing system-level simulations involving all three components.

The use of a Repeater to extend the reach between RC and EP over extremely lossy channels is a common practice and can present unique challenges in examining the design solution space. The Repeater's settings must be co-optimized together with SerDes transmitter (Tx) and receiver (Rx) settings to maximize the overall link performance.

The sections which follow present a how-to guide for performing end-to-end system-level simulation analyses involving Tx, Repeater, and Rx with a focus on PCIe Gen-4 channel topologies. The approach is based on IBIS Algorithmic Modeling Interface (IBIS-AMI) simulations. IBIS-AMI's standardized interface offers interoperability between models provided by different integrated circuit (IC) vendors. More importantly, critical component-level impairments such as non-linearity and device noise can be represented in IBIS-AMI models and reflected in the overall link performance—effects that standard s-parameter representations of equalizer components fail to capture.

For the purposes of this paper, a Xilinx FPGA SerDes is used for the RC and EP; and a Texas Instruments Linear Repeater is used in between the RC and EP to achieve reach extension. The methodology outlined here can be extended to any RC, Repeater, and EP device.

The proposed methodology for simulating the solution space of a Tx+Repeater+Rx system in the context of PCIe Gen-4 is as follows:

1. Determine if a Repeater is required
2. Define a simulation space
3. Define evaluation criteria
4. Execute the simulation matrix and analyze the results

The goal is to reach a conclusion regarding the optimum configuration of the system in an efficient and timely manner.

Step 1: Determine if a Repeater is Required

Before evaluating a Tx+Repeater+Rx link, you must first understand whether a Repeater is required for the link. There are a few ways of reaching this conclusion, and they vary in complexity.

1. Compare the end-to-end channel insertion loss (RC to EP) to the PCIe channel requirements. Section 9.4.1.2 of the PCIe Base Specification (Rev. 4.0, Version 0.7 as of this paper's publication) calls for a maximum channel loss of 28 dB at 8 GHz, including the RC and EP package loss. If the system channel loss exceeds this, then a Repeater is required.
2. Alternatively, a channel's s-parameter can be simulated in a tool such as the Statistical Eye Analysis Simulator (seasim) together with a reference transmitter, reference transmit-side package model, reference receive-side package model, and a reference receiver to determine the

post-equalized eye height (EH) and eye width (EW). If the EH and EW fail to meet the requirements set forth in Section 9.5.1.6 of the PCIe Base Specification, then a Repeater is required.

In this paper, the following two-connector PCIe channel topology is considered. Although every PCIe link has both a downstream (RC-to-EP) direction and an upstream direction (EP-to-RC), this paper focuses on analyzing the downstream direction. A similar analysis can be performed for the upstream direction.

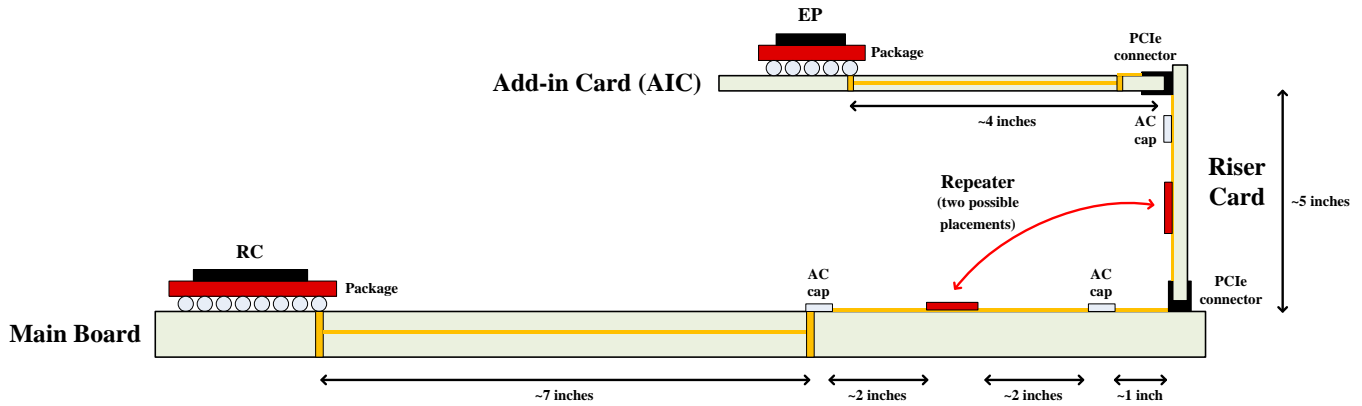


Figure 1: Two-Connector PCIe Channel Topology

The individual components of this channel topology—RC package, main board trace and via, connector, and so on—are characterized with separate s-parameters. The composite channel thru insertion loss, return loss, and crosstalk frequency responses are shown in Figure 2.




Figure 2: Composite Channel Thru-Loss, Return-Loss, and Crosstalk Frequency Response

It is immediately obvious that the total channel loss (38 dB at 8 GHz) exceeds the PCIe specifications by 10 dB. Furthermore, simulating this channel in seasim version 0.58.6 yields the following post-equalized eye.

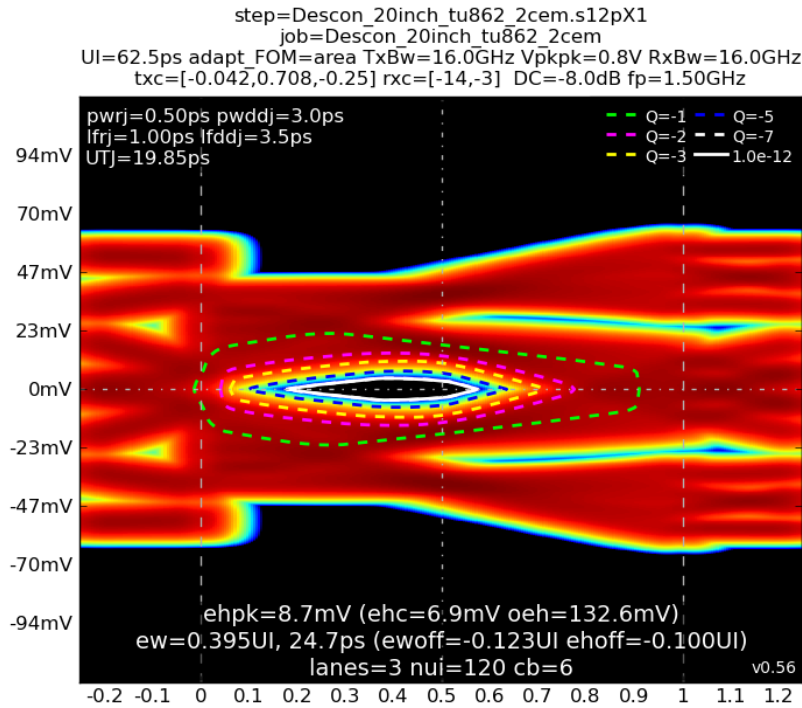


Figure 3: Seasim Simulation of Channel Topology

The conclusion from both methods of channel analysis is that a Repeater is required for the selected channel topology.

Table 1: Evaluating the Selected Channel Topology Against the PCIe Specifications

Channel analysis method		Value	PCIe Requirement	Conclusion
1	End-to-end channel insertion loss	-38.8 dB at 8 GHz	≥ -28 dB at 8 GHz	Repeater is required
2	Channel simulation with behavioral Tx, Rx, and package	EH = 8.7 mV EW = 0.395 UI	EH ≥ 15 mV EW ≥ 0.3 UI	Repeater is required

Step 2: Define a Simulation Space

With multiple active components—RC, Repeater, EP—and multiple channel topologies, the scope of the simulation task can grow exponentially. To achieve meaningful results in a reasonable timeframe, it is critical to define a narrow simulation scope, one which focuses on the parameters which are most likely to impact link performance.

The system-level simulation task can be broken down into two sequential phases:

- Phase 1. **Initial link performance analysis.** Analyze the impact of Repeater placement and Tx/Repeater/Rx settings on link performance.
- Phase 2. **Sensitivity analysis.** Quantify the sensitivity of link performance to process/voltage/temperature (PVT) variation and to variations in Repeater placement.

Table 2: Simulation Space

Simulation Phase	RC Tx Parameters	Repeater Parameters	EP Rx Parameters	Channel parameters
1	Initial link performance analysis <i>Presets:</i> 0, 1, ..., 9 <i>VOD:</i> 1000 mVppd,	<i>Boost:</i> Sweep six values <i>Wide-band gain:</i> -1 dB	<i>Rx parameters automatically adaptive</i>	<i>Channel topologies considered:</i> 1. Repeater placed on Main Board 2. Repeater placed on Riser Card
2	Sensitivity analysis <i>Presets:</i> 0, 1, ..., 9 <i>VOD:</i> 1000 mVppd	<i>Boost:</i> Optimum setting from Phase 1 <i>Wide-band gain:</i> ± 4 dB	<i>Rx parameters automatically adaptive</i>	<i>Focus on optimum topology, vary specific Repeater placement by ± 1 inch</i>

RC Transmitter Parameters

PCIe requires that RC and EP transmitters have the ability to generate a wide range of de-emphasis (post-cursor) and pre-shoot (pre-cursor) equalization voltages using a standard three-tap finite impulse response (FIR) filter.

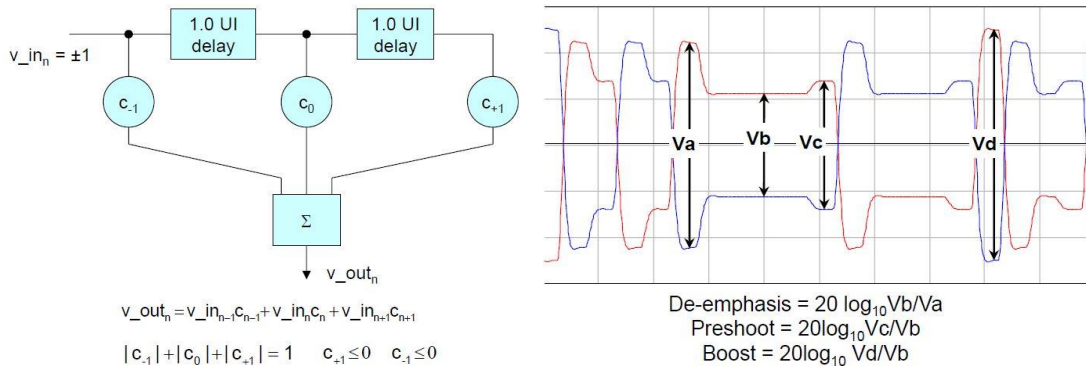


Figure 4: Tx Equalization FIR Representation (Left) and Tx Voltage Levels (Right)

To reduce the total number of simulations, the methodology outlined in this paper limits the simulated pre-shoot and de-emphasis settings to the ten Tx equalization Presets defined by PCIe. This has the added benefit of ensuring interoperability with multiple root complex and endpoint devices, since all PCIe-compliant devices are required to have the capability of generating these ten Preset equalization levels.

Table 3: Tx Preset Ratios and Corresponding Coefficient Values

Preset #	Pre-shoot (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	0.000	-0.167	1.000	0.668	0.668
P0	0.0	-6.0 ± 1.5	0.000	-0.250	1.000	0.500	0.500
P9	3.5 ± 1	0.0	-0.166	0.000	0.668	0.668	1.000
P8	3.5 ± 1	-3.5 ± 1	-0.125	-0.125	0.750	0.500	0.750

Preset #	Pre-shoot (dB)	De-emphasis (dB)	c ₋₁	c ₊₁	V _a /V _d	V _b /V _d	V _c /V _d
P7	3.5 ± 1	-6.0 ± 1.5	-0.100	-0.200	0.800	0.400	0.600
P5	1.9 ± 1	0.0	-0.100	0.000	0.800	0.800	1.000
P6	2.5 ± 1	0.0	-0.125	0.000	0.750	0.750	1.000
P3	0.0	-2.5 ± 1	0.000	-0.125	1.000	0.750	0.750
P2	0.0	-4.4 ± 1.5	0.000	-0.200	1.000	0.600	0.600

In addition to pre-shoot and de-emphasis, PCIe requires that the full-swing signal amplitude (V_d) be between 800 mV_{ppd} (min) and 1300 mV_{ppd} (max). In this analysis, 1000 mV_{ppd} was selected because it falls roughly in between the two PCIe limits.

Repeater Parameters

Linear Repeaters are commonly used in PCIe systems to provide reach extension between the RC and EP, especially when the system channel exceeds the insertion loss limits set forth in the PCIe specification. Such devices conventionally provide two mechanisms for signal conditioning:

1. *High-frequency boost*: A continuous-time linear equalizer (CTLE) amplifies high frequencies relative to DC in order to counteract the frequency-dependent loss characteristics of a passive channel.
2. *Wide-band amplitude gain/attenuation*: A wide-band amplifier stage can boost or attenuate all frequencies to adjust the overall amplitude of the signal.




Figure 5: Linear Repeater Frequency Response Showing 0 dB Wide-Band Gain (Left), Max Wide-Band Gain (Center), and Min Wide-Band Gain (Right)

Selecting the Repeater settings appropriate for the simulation analysis requires an understanding of the channel topology and the insertion loss in the pre-channel and post-channel segments. It is common practice to configure a Repeater to slightly *under-equalize* the pre-channel. This strategy allows the RC Tx and EP Rx to apply their own equalization so that the total equalization burden is shared across all components in the channel.




Figure 6: Generic Topology Including RC, Repeater, and EP

In this analysis, the pre-channel loss is approximately 19 dB. For Phase 1, the wide-band amplitude gain is set to a small value (-1 dB), and a range of eight high-frequency boost settings spanning 16 dB to 21

dB are swept. For Phase 2, the optimum boost setting from Phase 1 is selected, and the wide-band amplitude gain setting is varied by ± 4 dB to gauge the link's sensitivity to variations in peak-to-peak amplitude.

EP Receiver Parameters

In this analysis, the End Point receiver was configured to automatically adapt its CTLE and decision feedback equalizer (DFE) circuits, therefore no parameters are swept on the Rx side.

Channel Parameters

As is often the case, this analysis focuses on two options for placement of the Repeater: (1) on the Main Board, or (2) on the Riser Card.

In Phase 1 of the analysis, both placements are included in the simulation space. For each placement, an extra 1 inch of PCB trace is added to either side of the Repeater so that, in Phase 2 of the analysis, the specific Repeater placement can be shifted by ± 1 inch. This approach provides useful insight into the sensitivity of link performance to the specific placement of the Repeater.

Step 3: Define Evaluation Criteria

An important aspect of any simulation analysis is deciding how to evaluate the results to determine margin or pass/fail. Bit error rate (BER) is the ultimate gauge of link performance, but unfortunately an accurate measure of BER is not possible in relatively short, multi-million-bit simulations.

Instead, the methodology outlined in this paper uses two criteria to establish link performance:

1. A link must meet receiver's EH and EW requirements
2. A link must meet criterion 1 for all Tx Preset settings

Criterion 1 establishes that there is a viable set of settings which will result in the desired BER. The specific EH and EW required by the receiver is implementation-dependent. For example, the Xilinx Rx requirements for post-equalized EH and EW are shown in Figure 7. Criterion 2 ensures that the link has adequate margin and is not overly-sensitive to the Tx Preset setting.




Figure 7: Rx EH and EW Requirements

Step 4: Evaluate the Simulation Matrix and Analyze the Results

For this analysis, IBIS-AMI models are used for each of the active components: RC and EP SerDes from Xilinx and a Linear Repeater from Texas Instruments. Keysight ADS is the tool used to execute the IBIS-AMI simulations, measure the extrapolated EH and EW, and plot post-equalized eye. The simulation schematic is shown in Figure 8.

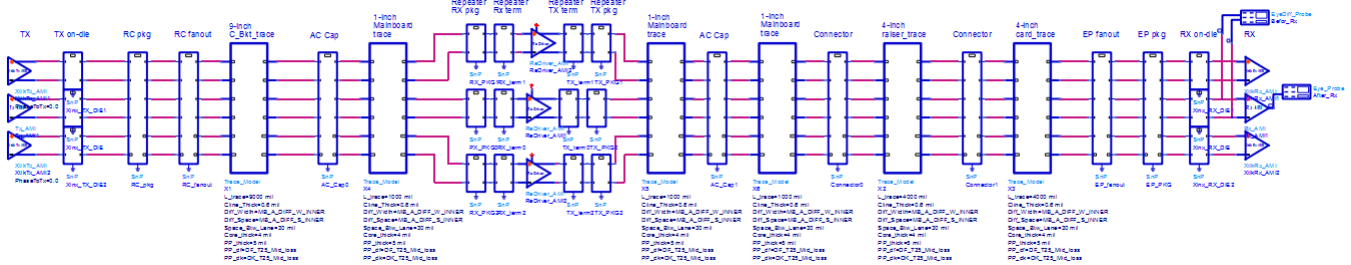


Figure 8: Keysight ADS Simulation Schematic (Main Board Repeater Placement Shown)

The simulation execution step is separated into two phases, and each phase uses the same basic simulation topology. Note that, whenever possible, details such as Rx/Repeater/Tx package, crosstalk, and other non-idealities should be included in the simulation as they have been here.

Phase 1: Initial Link Performance Analysis

The focus of Phase 1 is to run a broad set of relatively short simulations to explore the design solution space. Minimizing the simulation time for each individual simulation is crucial. To achieve this, the following simulation parameters are used:

Table 4: Simulation Parameters for Phase 1

Simulation Parameter	Value
Data Rate	16.0 GT/s
Data pattern	PRBS31
Total number of bits	1M
Ignore_Bits	500k <i>Note: This is set by the Rx model</i>
Simulation type	Time domain (a.k.a. bit-by-bit) <i>Note: Simulations will be faster running in Statistical mode, however non-linear behavior may not be adequately represented.</i>
Bit-by-bit extrapolation	Enabled <i>Note: Simulations will be faster without this mode enabled, however RJ will not be accounted for as accurately.</i>

The results are visualized in a few different ways. Figure 9 shows the EH and EW pass/fail result for each Tx Preset (horizontal axis) and Repeater boost setting (rows, in dB). This shows that there are multiple Repeater boost settings, namely the higher boost settings, which result in passing EH and EW for all Tx Presets. At lower Repeater boost settings, there are some Tx Presets which yield marginal or failing EW and EH results.

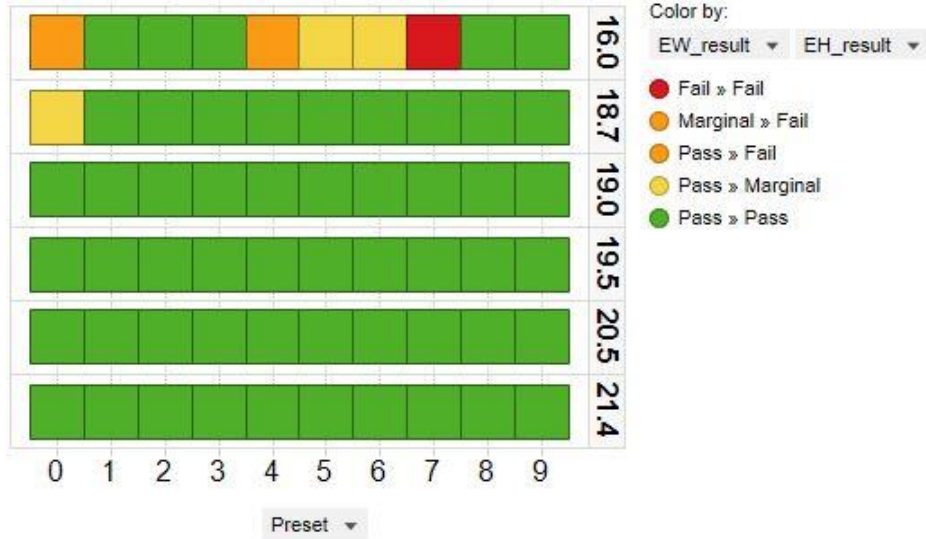



Figure 9: EH/EW Pass/Fail Result by Preset and Repeater Boost Setting (Rows, in dB)

Figure 10 shows the average EH and EW for each Repeater boost setting. Since multiple Repeater boost settings yielded passing results, this perspective helps to identify the best setting—the one which maximizes EH and EW.

Rx Post-Equalized EW (UI) vs. Repeater Boost (dB)



Rx Post-Equalized EH (mV) vs. Repeater Boost (dB)



Figure 10: Rx Post-Equalized EH and EW vs. Repeater Boost

Figure 11 shows the post-equalized eye diagram in the end point Rx for all Tx Preset settings. The optimum Repeater boost (20.5 dB) is used for these cases. Tx Preset 8 yields the largest post-equalized eye opening.

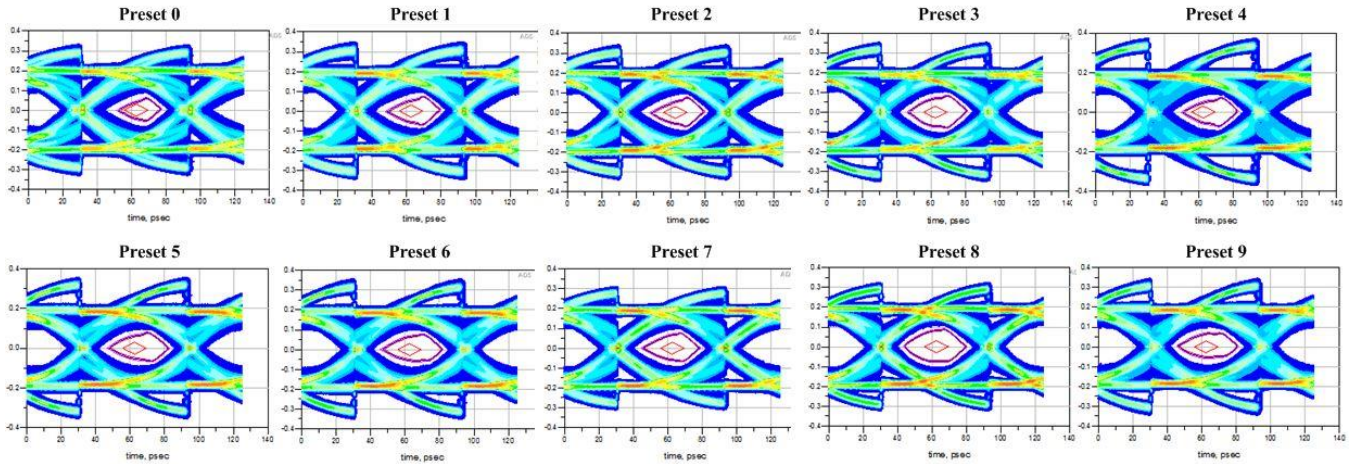


Figure 11: Rx Post-Equalized Eye Diagrams for all Tx Presets

A similar analysis is conducted for the alternate placement: Repeater on the Riser Card. In this configuration, the pre-channel loss is approximately 25 dB at 8 GHz. The results for this placement show consistently reduced performance compared to the Main Board placement of the Repeater. Figure 12 shows a comparison of the EP receiver’s post-equalized eye for the optimum Tx Preset and Repeater boost setting for the Main Board and Riser Card placements. This identifies the Main Board as the optimum Repeater placement, and this placement will therefore be used for Phase 2 of the analysis.

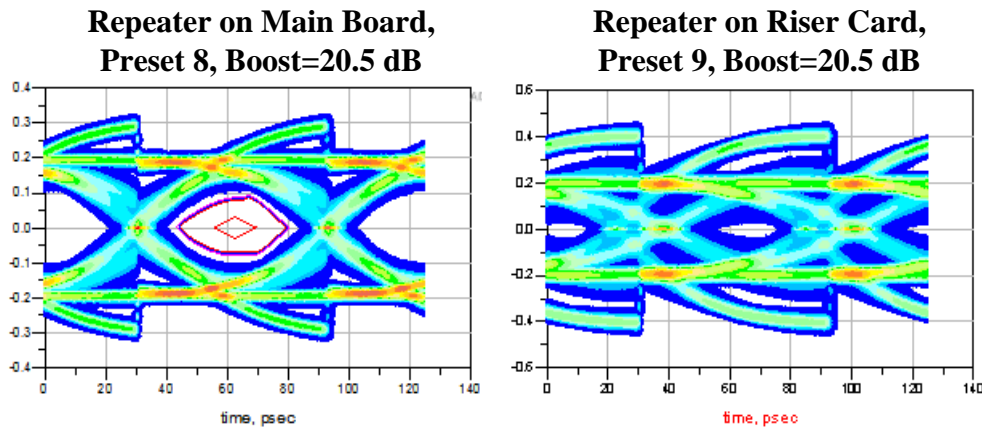


Figure 12: Comparison of Main Board vs. Riser Card Placement

Phase 2: Sensitivity Analysis

Once an optimum Repeater placement is identified and baseline performance is established, the next phase is to understand the sensitivity of link performance to common variables like process, voltage, temperature (PVT), and the specific placement of the Repeater. For simplicity sake, the sensitivity due to specific placement and PVT are analyzed separately.

To analyze the sensitivity of link performance due to the specific placement of the Repeater, simulations are run on the optimum link topology (Repeater placed on the Main Board) with a ± 1 inch variation in the placement of the Repeater.

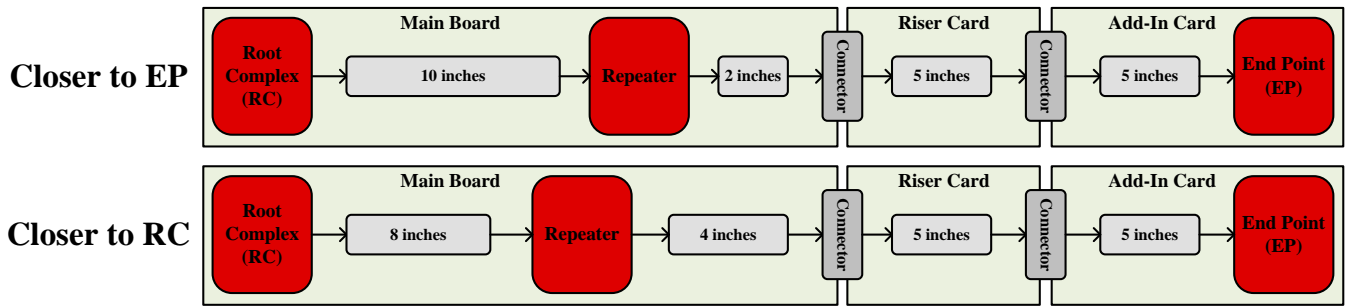


Figure 13: Block Diagram of Specific Repeater Placement Sensitivity Analysis

The results are visualized in a couple different ways. Figure 14 shows the EH and EW pass/fail result for each Tx Preset (horizontal axis) and Repeater boost setting (rows, in dB). This plot shows that the overall link performance is relatively insensitive to specific repeater placement, with a minor improvement in overall results in the case where the Repeater is shifted closer to the RC.




Figure 14: EH/EW Pass/Fail Result by Preset and Repeater Boost Setting for Two Specific Repeater Placements: 1 Inch Closer to RC, and 1 Inch Closer to EP

Another important aspect of the placement sensitivity is answering the question: does the optimum Repeater setting change with a relatively minor shift in the specific placement? Figure 15 answers this question by showing the average EH and average EW across all Tx Preset settings for each of the Repeater boost settings. This plot shows that the same Repeater boost setting (19.5 dB) is optimal even if the Repeater is shifted ± 1 inch from its original location. This gives the system designer some comfort knowing that a small shift in the specific placement of the Repeater will not result in a time-consuming settings optimization exercise.

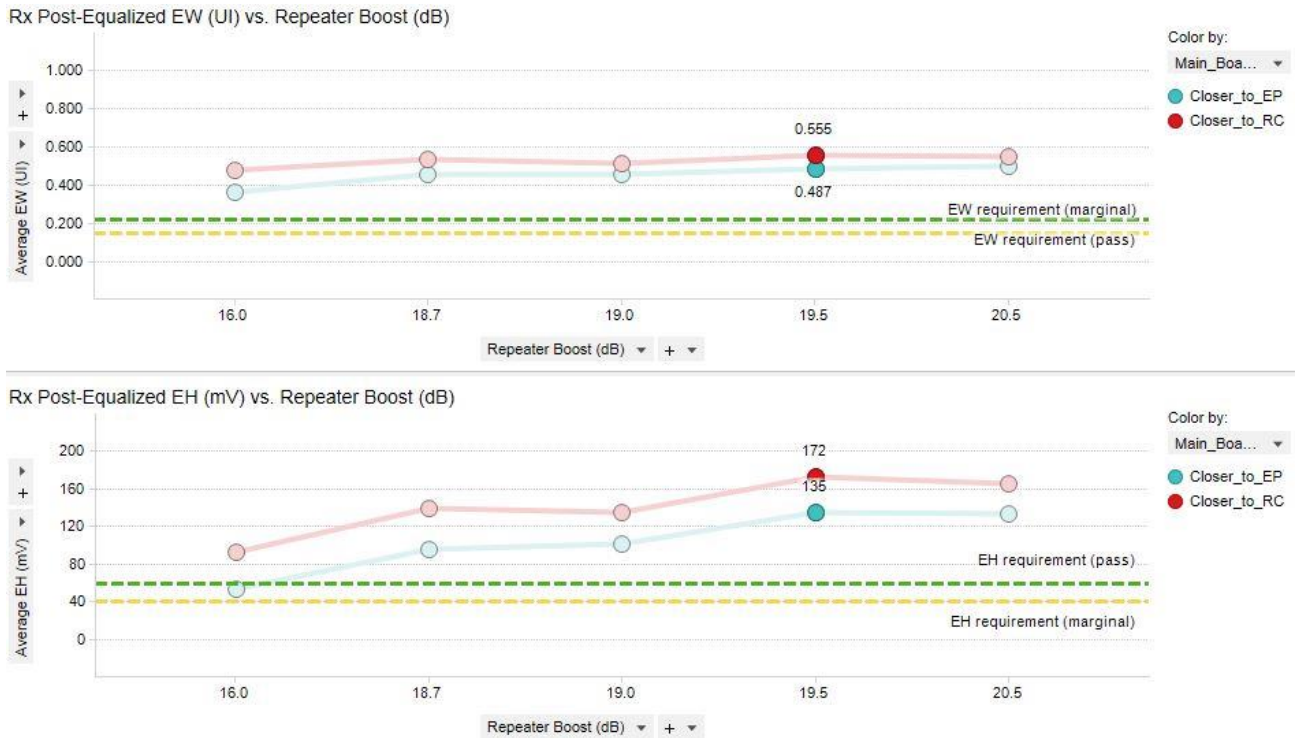


Figure 15: Rx Post-Equalized EH and EW vs. Repeater Boost for Two Specific Repeater Placements

The last part of the Phase 2 sensitivity analysis is to look at the sensitivity of link performance to process, voltage, and temperature (PVT) variations. To exacerbate the effects of PVT variation, the Repeater’s wide-band gain is also varied by approximately ± 4 dB. Figure 16 shows that the overall link performance is not affected until *both* extremes of PVT variation and wide-band gain are realized. This gives the system designer more comfort knowing that, as long as the Repeater’s wide-band gain setting is kept to a reasonable, mid-level value, the link performance will be robust across PVT corners.



Figure 16: EH/EW Pass/Fail Result Across PVT and Different Wide-Band Gain Variations

Conclusions

The methodology outlined in this paper allows system designers to quickly and effectively evaluate the end-to-end link performance of a PCI-Express Gen-4 link involving a Root Complex, a Repeater, and an End Point. By carefully selecting a simulation space, defining the evaluation criteria up front, and executing the simulations in a two-phase process, a system designer can optimize total simulation time while achieving a meaningful result which aides board design and device configuration. Simulations are not a substitute for real-world lab validation, but they can be an extremely useful tool when designing your Gen-4 PCIe links.