ACHIEVING PCI-E COMPLIANCE:
GETTING IT RIGHT THE FIRST TIME

EMA Design Automation®

To remain competitive in today's high-performance industry, technology must be produced quickly and correctly. As data rate requirements continue to increase for high-speed serial links, successfully fulfilling complex requirements presents a real challenge. To meet these demands, design engineers must thoroughly analyze and test the electromagnetic effects of their high-speed PCB interconnects as well as component packages for optimal signal integrity performance. When designing to an industry interface standard, the associated specification will typically have a set of measured signal integrity "Compliance Checks" that must be met in order to guarantee the correct operation of the interface.

Compliance Checks

Interface problems can show up in many forms: system crashes, failed transfers, failed attempts to read or write from memory, or corrupted data. These are all difficult and annoying problems to try to narrow down after hardware is built. Interface compliance measurements are metrics that can be used to help avoid these costly problems. Ideally, design teams would be able to qualify their design against these measurements in software using simulation before the first prototype is ever built. Knowing what to check though can be a challenge, especially as data rates increase and interface complexity goes up. Fortunately, many of the standard interfaces have working groups which look to provide guidance on the measurements and checks needed to qualify your design properly.

PCI-SIG and PCI-e Compliance

PCI-SIG (Peripheral Component Interconnect – Special Interest Group) is the organization who owns and manages the PCI industry standards whose focus is to produce specifications that define industry-wide standards for all varieties of the PCI interface. They publish a list of PCI compliant products known as the integrators list. In order to be included into this list, member companies must pass testing during a compliance workshop for the standard at hand. In this case we will discuss the PCI-e specification and its compliance requirements.

PCI-e Overview

PCI Express (PCI-e) architecture is an industry standard high-performance, general-purpose serial I/O interconnect designed for use in embedded, mobile, desktop, and enterprise platforms, touching nearly every device for communications between system components. It is the latest generation of the PCI interface standard and is the catalyst for the development of many other technologies. PCI-e also serves as a basis to help propel serial communication technology forward.

The PCI-e protocol is complex and covers a vast amount of test areas such as electrical, configuration, link protocol, transaction protocol, and platform BIOS. This protocol not only aids in product development, but it helps minimize interoperability issues once the product is launched and provides valuable feedback on product performance. To meet PCI-e compliance for the electrical category of the Physical Layer, several tests must be run encompassing the following general categories:

- Channel Tolerance/Eye Mask
- Differential Insertion Loss
- Differential Return Loss
- Stress/Swept Jitter Test

These tests can help companies and designers improve their devices before releasing them to market as well as provide confidence it will perform as intended. In turn, consumers are given peace of mind knowing the products they are purchasing have been created and tested to ensure maximum quality and reliability.

Since attending a Compliance Workshop (events hosted a few times a year, which provide members the opportunity to test and validate products before they enter the field) can

be difficult and costly, it is important to ensure your product is as close to passing as possible. Running testing requirements in-house beforehand will diminish the repercussions of major design flaws and testing failures.

Channel Tolerance/Eye Mask Check

The first tests run involve an eye diagram; a long sequence of data bits displayed by chopping the stream into 2 UI (Unit Interval) chunks and overlaying each one on top of each other. This provides a more convenient and meaningful way of analyzing large amounts of data for proper voltage values and timing results. An eye mask is a template of sorts that shows you voltage and timing reference points marking the boundaries where your eye diagram should fall. The eye mask should be able to fit within the eye diagram, clearing the inner-most signal transitions of the eye diagram. Channel tolerancing is just another way of measuring how close the eye diagram is to the eye mask.

Why is this important?

Simply put, for a signal to be understood by the receiving device without errors, it needs to be received in the proper timing sequence and above or below a specified voltage value, both of which would be defined by your eye mask.

What measurements need to be checked?

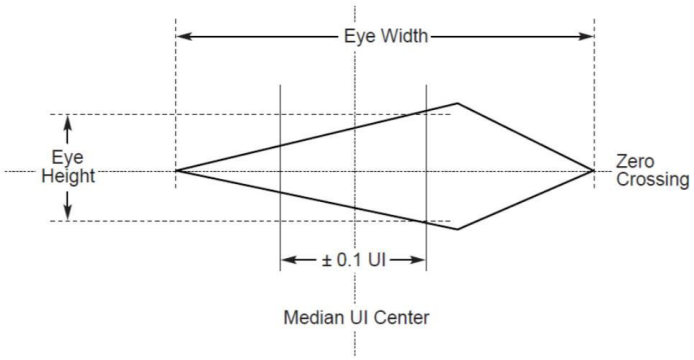

Besides the overall fit of the eye mask within your data eye, there are other common measurements that are taken right from the eye diagram, namely Eye Height and Peak EH (Eye Height) Offset from UI Center.

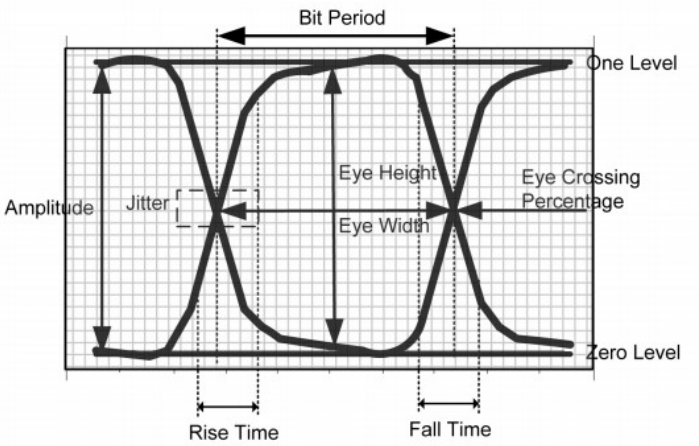

Eye height by loose definition is the maximum difference between the inner most '1' bit representation and the inner most '0' bit representation, taken in the +/- 10% vicinity of the center of the bit interval (a.k.a. the unit interval). This measurement is not always associated with the peak eye height within the bit interval. It is essentially a direct indication of your signal energy arriving at the receiving device. It can also be measured in the s-parameter domain as insertion loss. Having significant eye height, based on a requirement established in the associated interface standard provides the best probability of good signal quality which translates to receiving valid data. Jitter, noise, and signal attenuation directly affect the eye height.

Peak Eye Height (EH) Offset from UI Center is a measure of where the absolute peak eye height is in relation to the center of the bit interval. It is an indication of how much your transmitted eye is becoming distorted due to unequal attenuation characteristics of your channel and possible noise and jitter characteristics. Ideally the eye height

should peak at the center of your unit interval, but in most cases this peak is offset a bit due to the physical limits of the channel. Too much offset can result in phase and timing issues as the receiver might not have time to respond to one bit versus the next.

The electrical characteristics of your hardware channel will be limited, meaning that all compliance items for eye quality won't necessarily be met or controlled with purely hardware design techniques. **Equalization filtering** of both the transmitter and receiver data streams are employed to further compensate for hardware channel deficiencies. Equalization filters employ coefficient values that control how much attenuation is applied and at what frequency points.

For PCI-e compliance, the receiver equalization process known as Decision Feedback Equalization (DFE) is given a required range of effectiveness for eye quality by the parameter **Range for DFE d1/d2 coefficient.** These coefficients can be used to recover badly degraded signals, bringing them into conformity at the receiving end. However, equalization filter performance is limited so you still need your hardware channel design to perform with reasonable signal quality without equalizers to give your design the best chance at compliance once the equalizers are activated-

Tip: Removing stubs from vias (backdrilling), ensuring correct model extraction and connection, and an adequate range selection for the DFE coefficients can help fix failures encountered in your PCI-e design testing.

Differential Insertion Loss Check

As implied above, you channel won't be a perfect transmission medium for your signals, it will have loss. If your channel has more loss than anticipated, your signal will be harder for the receiver to correctly interpret. The amount and type of loss can be characterized in several ways.

Probably the most important loss characteristic is Differential Insertion Loss which is the amount of attenuation (loss) that your signal will experience as it travels along your channel (a differential pair net), also known as SDD21.

Why is it important?

Differential Insertion Loss is essentially the ratio of energy at your receiving end of the channel (port '2' in the S-parameter world) to the energy introduced to the channel at the transmitter end (port '1'). This measurement is assumed to be conducted using a differential transmitter and a differential receiver to mimic the conditions that your physical PCIe channel design is subjected to. Hence the SDD21 terminology:

- S => S-parameter
- D => Differential response mode (i.e. differential receiver)
- D => Differential stimulus mode (i.e. differential transmitter)
- 2 => Response port (i.e. port 2, output of network)
- 1 => Stimulus port (i.e. port 1, input of network)

What measurements need to be checked?

A general assumption is that you can expect to see about 0.1dB per inch per GHz of loss in your channel. This means that if you have a 30-inch cable at 3GHz, you should expect to see around (0.1 * 30 * 3) = 9bdB of loss at 3GHz.

The screenshot below is an example of differential insertion loss in a typical PCIe channel, which is checking SDD21 (differential insertion loss). It shows an s-parameter plot for an 8.6ns long cable which at about 5000 mil/ns is 43 inches long. Putting that into our estimation formula above, we should expect to see (0.1 * 43 * (3) = 12.9dB of loss which is quite close to the 12dB of loss that we're seeing at 3GHz.

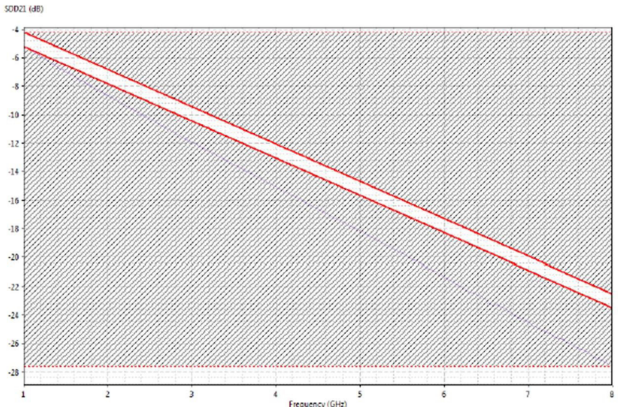

Tip: For the length of your PCI-e design channel as related to the reference calibration channel that comes with the PCI-e specification: If your design characteristics for SDD21 are above and beyond the compliance range, you're likely going to have an easier time getting things to work out. If you're below that range, you'll have a harder time. Using our estimations above, that calibration channel length is probably around 30-inches (20.5/7/0.1).

Differential Return Loss Check

The amount of signal energy reflected to your differential source when sending a signal through a channel is referred to as Differential Return Loss. It's measured as two parameters: SDD11 (input differential return loss) and SDD22 (output differential return loss). Ideally you want this value to be zero as the desired goal is to get as much signal information to the receiving end of the channel as possible. But again, no channel design is ideal, there will be loss.

Why is it important?

If too much signal is getting bounced back by the channel itself, there won't be enough signal energy making its way through the channel to be recognized at the receiver. The returning energy then will begin to interfere with the signal information being transferred, causing distortion of the signal quality.

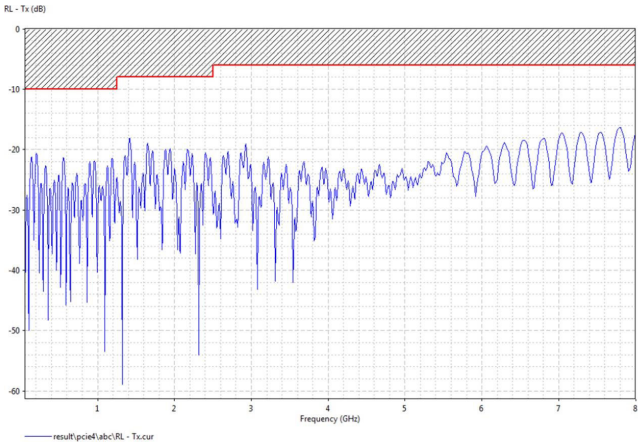

What measurements need to be checked?

Measuring the input and output differential return loss and checking against the specified limits quickly allows us to see if the channel is accepting enough of the signal without reflecting too much back. The graph above shows the differential return loss of an example channel. The goal is to have your results as far below the "shaded' or "hashed' portion of the graph as possible.

Tip: Check the models being used as well as the impedance of the channel itself to ensure you don't have impedance discontinuities. Correct or "balanced" impedance matching will achieve the most efficient energy transfer.

Stressed/Swept Jitter Test

We'll introduce the concept of the Bit Error Ratio (BER) for this next compliance item. The BER is a ratio of the number of bits received in error at the receiver, compared to the total number of bits received. Its value is affected by all the

measurable compliance items mentioned so far, including signal noise, distortion, and jitter. Depending on the design requirements, this ratio can vary as not all designs run at the same performance level. Some channel designs may run relatively slowly and would have a very low BER while other higher performance designs run much faster and would have a relatively higher BER. Whatever the value for the given channel design, there is a BER value deemed "acceptable" for the given channel performance.

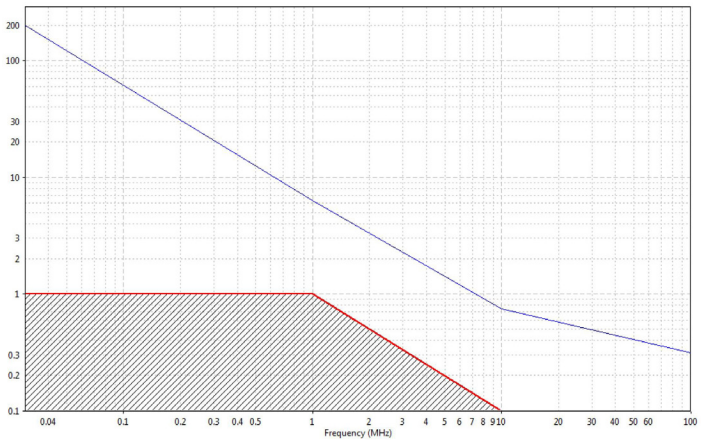

The Stressed/Swept Jitter test is a plot showing the frequency vs. amplitude of Periodic Jitter where for every frequency point, it shows the maximum Jitter amplitude allowed to still have a discernible eye at the desired BER (Bit Error Ratio). It tells you how much jitter margin (allowable jitter) you can have and still understand what is being received.

What measurements need to be checked?

At each frequency, the amplitude of the periodic jitter is increased until the eye is no longer discernible at the BER of interest (often 1e-15). The last jitter amplitude value that produced a discernible eye at the BER of interest is what's recorded in the results, giving us the blue curve (see the graph on the right) when all points are plotted.

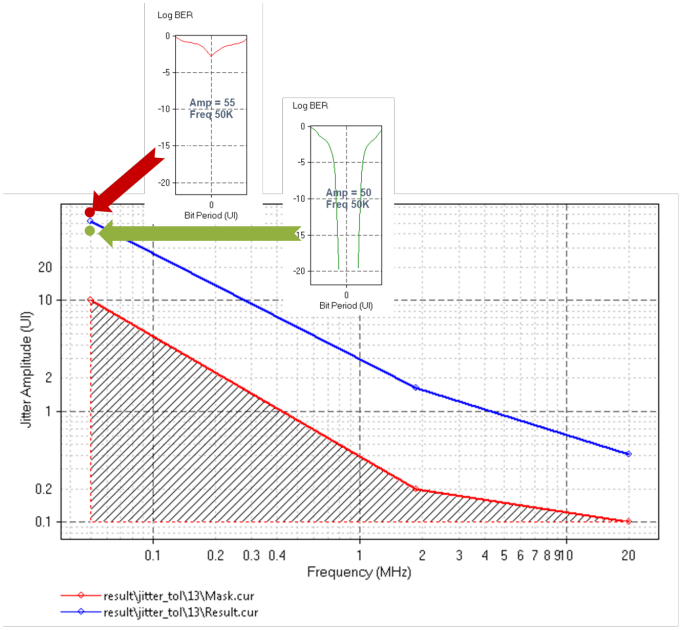

Note the red dot and arrow showing an amplitude of periodic jitter that is too high and producing a closed eye. Also note a green dot and arrow showing an amplitude of periodic jitter that still produces an open eye at the BER of interest (bathtub curves that extend down to the bottom of the graph).

Tip: If this test fails, your design is too susceptible to jitter and needs to increase its tolerance. If shortening you channel length is an option, you might want to do so. Otherwise, you'll have to adjust the configuration of the device; using AMI models can help you identify the right configuration.

Conclusion

With the design of high-speed serial data links only becoming more complex, leveraging the power of simulation to ensure designs are optimized is essential to design suc-

cess. Having a PCI-e compliant board can ensure designs perform as intended. Checking for electrical compliance for a PCIe interface can become quite involved and knowing what to check is only part of the equation. Being able to quickly analyze your interface, test alternatives, and fix issues early when the cost of change is lowest is critical.

While possible, manually completing these tests and generating compliance forms are time-consuming and can leave room for error. Since many design decisions made are based in part on other testing conducted, if the testing proves to be inaccurate it will affect the entire design, not just one piece. Automated compliance checks are a much better option as they follow a rigorous procedure as dictated in the interface standard and provide overall peace of mind for a sound, well-performing design interface.

Sigrity Can Help You Ensure PCIe Compliance

Sigrity provides advanced SI analysis for both pre- and post-layout. The compliance toolkit within Sigrity provides you with a 'push button' approach to testing your design directly against the PCI specifications—without manual calculations.

Once you run the compliance analysis in Sigrity, the software will immediately notify you of your pass/fail status. From there, you can analyze results even further with the ability to overlay results over the specifications to show you exactly how your product performance compares to the specification. This helps automate the compliance process and eliminates error prone manual checking, allowing you to quickly and accurately test before you attend a compliance workshop.

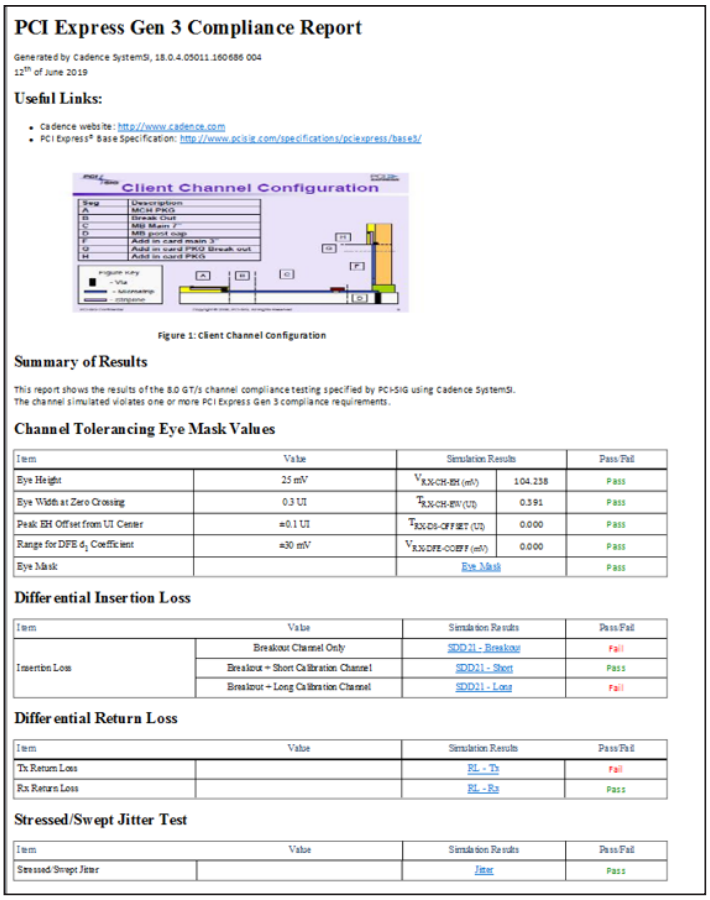

PCI Express Gen 3 Compliance Report

Generated by Cadence SystemSI, 18.0.4.05011.160686 004
12th of June 2019

Useful Links:

- Cadence website: http://www.cadence.com
- PCI Express® Base Specification: http://www.pcisig.com/specifications/pciexpress/base3/

Figure 1: Client Channel Configuration

Summary of Results

This report shows the results of the 8.0 GT/s channel compliance testing specified by PCI-SIG using Cadence SystemSI.
The channel simulated violates one or more PCI Express Gen 3 compliance requirements.

Channel Tolerancing Eye Mask Values

Item	Value	Simulation Results		Pass/Fail
Eye Height	25 mV	$V_{RX\text{-}CH\text{-}EH\ (mV)}$	104.258	Pass
Eye Width at Zero Crossing	0.3 UI	$T_{RX\text{-}CH\text{-}EW\ (UI)}$	0.391	Pass
Peak EH Offset from UI Center	±0.1 UI	$T_{RX\text{-}DS\text{-}OFFSET\ (UI)}$	0.000	Pass
Range for DFE d_1 Coefficient	±30 mV	$V_{RX\text{-}DFE\text{-}COEFF\ (mV)}$	0.000	Pass
Eye Mask		Eye Mask		Pass

Differential Insertion Loss

Item	Value	Simulation Results	Pass/Fail
Insertion Loss	Breakout Channel Only	SDD21 - Breakout	Fail
	Breakout + Short Calibration Channel	SDD21 - Short	Pass
	Breakout + Long Calibration Channel	SDD21 - Long	Fail

Differential Return Loss

Item	Value	Simulation Results	Pass/Fail
Tx Return Loss		RL - Tx	Fail
Rx Return Loss		RL - Rx	Pass

Stressed/Swept Jitter Test

Item	Value	Simulation Results	Pass/Fail
Stressed/Swept Jitter		Jitter	Pass

For more information on Sigrity, visit www.ema-eda.com.

About EMA Design Automation

EMA Design Automation, a Cadence Channel Partner, is a trailblazer in product development solutions in a complete range of EDA tools, PLM integrations, services, training, and technical support. Serving North America from Rochester, New York, EMA is a privately held corporation which develops Ultra Librarian®, TimingDesigner®, CircuitSpace®, CIP™, EDABuilder®, and a host of custom solutions to enhance the OrCAD products, and all are distributed through a worldwide network of value added resellers. www.ema-eda.com