

An Introduction to PCI Express™ Measurements




Table of Contents

Validation, Debug, and Compliance of PCI Express Designs	3
Architecture and Specification Overview	4
Characteristics of the Physical Layer	5
Physical Layer Compliance Testing: An Analog View Of A Digital World	8
Validation and Compliance Measurements	8
Amplitude Measurements	9
Timing Measurements.....	9
Jitter Measurements	10
Real Time Eye Diagrams and Mask Testing	14
What about Noise and its contribution to Jitter?	14
Frequency Domain Measurements.....	15
Receiver Sensitivity and Jitter Tolerance Measurements	16
Insuring Signal Fidelity in Analog Measurements.....	17
Connecting to the Serial Device Under Test	18
Analog Acquisition Considerations.....	22
Debug the Link using Pinpoint® Triggering and Data Decoding Tools	25
Locate and Trigger on a Bit Error	25
Validate and Trigger on 8b/10b Symbols	26
Digital Validation and Debug	27
Probing Means Making the Right Connections.....	28
Triggering Tools Boost Troubleshooting Efficiency	30
Analyzing the Results	30
Looking at the “Big Picture” of a Multibus System	31
Capturing Power Management Transistions	32
Overview of PCI Express Measurement Solutions	34
Oscilloscopes	34
Real Time (RT) Oscilloscopes	34
Equivalent Time (ET) Sampling Oscilloscopes	34
Signal Generators	35
Logic Analyzers	35
Summary	35

Validation, Debug, and Compliance of PCI Express Designs

As first generation (Gen1) PCI Express deployment is well underway, the industry has been successful in implementing components, systems, and cards. There are now motherboards and laptop computers on the market with PCI Express slots and a selection of Graphical Processor Units (GPUs) and network cards that use this new expansion slot. The successful deployment of Gen1 devices in just a few short years can be attributed partly to the rigor put into the testing of devices both in vendor's labs and at PCI-SIG (PCI Special Interest Group) sponsored plug fests. Compliance and interoperability testing for Rev1.0a of the PCI Express Card Electrical Mechanical (CEM) has been successfully implemented by the PCI-SIG and several vendors have qualified for the SIG's Integrator's list. During the deployment of Rev1.0a, it became clear that some changes to the specification were necessary (primarily in the physical layer) to insure more reliable communication between different PCI-E designs. These changes have now been rolled into Rev1.1 of the Base and CEM Specifications and the SIG has begun testing to the Rev1.1 specification at workshops. The PCI-SIG has also released the Base and CEM Specifications for Gen2 as well as the first Cabling Specification to enable peripheral expansion of the bus. Gen2 doubles the base signaling frequency from 2.5 Gb/s to 5 Gb/s and the addition of a cable to system design adds an additional level of complexity. This will surely introduce some new challenges in validation and compliance testing of devices.

Beyond just meeting the pass/fail criteria of the PCI-SIG for integrator list inclusion, it is also necessary for vendors to validate their designs in many conditions. For instance, your design may pass a test at a workshop under room temperature under power supply conditions of the test fixture, but what happens as you vary temperature and power supply conditions? What happens as you mate your device with another vendor's in a system and crosstalk from another bus causes your device to fail?

This primer deals with all aspects of PCI Express testing including validation, debug, and compliance as well as basic principles, measurement techniques, and the tools needed to insure successful designs. Whether you've successfully tackled Gen1 and are starting to test Gen2 devices or you are just getting started in PCI Express, this primer is designed to help you understand the PCI Express architecture, specifications, and measurement solutions.


Some Useful links for PCI Express Developers

Tektronix PCI Express Home:


http://www.tek.com/Masurement/applications/serial_data/pci_express.html

PCI-SIG Compliance Testing Home:

<http://www.pcisig.com/specifications/pciexpress/compliance/>



► Figure 1. Layers of a PCI Express Link.



► Figure 2. PCI Express PHY link implementation.

Raw Bit rate (Gb/s) by Lane Width

Signaling Rate	X1	X4	X8	X16	X32
Gen1 – 2.5 Gb/s per lane	2.5	10	20	40	80
Gen2 – 5 Gb/s per lane	5	20	40	80	160

► Raw bit rate.

Architecture and Specification Overview

As in any serial data standard, PCI Express can be viewed as “stack of layers,” as shown in Figure 1. Layers communicate with each other while at the same time buffering each other’s operations from adjacent levels.


The stack includes the physical layer in which electronic signals pass through transmission media; a logical layer in which these signals are interpreted as meaningful data; a transaction layer, and more. Each layer has its own applicable standards and compliance procedures.

Figure 1 shows the physical layer (PHY) partitioning of a serial data link. The PHY provides isolation from the higher layers of the protocol stack, and encompasses two layers: logical and electrical. The Electrical section of the PHY handles the high speed serial packet exchange and power management mechanisms.

The Logical layer of the PHY handles reset, initialization, encoding, and decoding. Both the electrical and logical sub-blocks also may incorporate standard-specific features.

Each of the two blocks composing the PHY has unique test requirements. Analog waveform characteristics of the data eye diagram and system reference clock are a priority when making electrical interface measurements. In the logical layer, digital packets must be interpreted, embedded clocks extracted, and so forth.

Figure 2 shows a typical physical layer topology of a 4X PCI Express link. The PCI Express link is made up of a collection of dual-simplex transmission scheme known as lanes. Each lane has a transmit and receive differential pair for a total of four copper traces per lane.



▶ **Figure 3.** The layers of a PCI Express Packet.

Each lane transmits at 2.5Gb/s (Gen1) or 5Gb/s (Gen2). The specification supports X1, X4, X8, X16, and X32 lane width configurations. Raw bit rate for the link is simply a multiple of the base bit rate per lane.

While data is transmitted by embedding the clock, both transmit and receive circuits share a Reference Clock (RefClk) so that their PLL circuits can track when Spread Spectrum Clocking (SSC) is used, as shown in Figure 2. SSC is used to minimize Electro-Magnetic emissions from the computer mother board. The reference clock is modulated between 100 MHz and 99.5 MHz at a frequency of 30-33 kHz.

Packetized Data

PCI Express data is packetized with contributions from several layers, as shown in Figure 3.

- ▶ The logical sub-block of the PHY adds framing to signal the beginning and end of each packet. The sub-block is responsible for symbol transmission and alignment.
- ▶ The data link layer provides error checking and re-tries services. Packet include ACK (Acknowledge), power management information, and more.
- ▶ The transaction layer handles initialization, instruction generation and processing, and flow control.

Characteristics of the Physical Layer

The physical layer is the carrier of the packetized differential signals just described. From the functional perspective, the physical layer mechanics—circuit board traces, connectors, and cables—are simply a path for data expressed in the form of binary signals. These binary signals are the subject of your physical layer measurements for debug, validation, and compliance tests with an oscilloscope.

Within PCI Express, several different bus configurations can be defined. Figures 4, 5, and 6 depict three typical mechanical implementations.

Every individual lane of the PHY consists of transmit and receive differential pairs. The transmit pair and its transmission medium are often called a channel. Within a channel, signals may traverse three basic types of copper paths.

Chip-to-Chip (Figure 4)

A chip-to-chip lane normally resides on an etched circuit board (ECB), which serves as the transmission medium. A typical application would be a PCI Express bus on a system motherboard where large amounts of data need to be transmitted from one device to another.

Card-to-Card (Figure 5)


When a connector is part of the lane, as in an Add-In card application, the specification requires AC coupling capacitors on the transmit side of the link. This eliminates potential common-mode bias mismatches between transmit and receive devices. Low-cost links may employ long runs on FR4 board and inexpensive connectors, both of which can add jitter, crosstalk, and potential imbalances due to layout. The test points for this configuration are found in the CEM specifications for desktop and server applications and in the PCMCIA (Personal Computer Memory Card International Association) ExpressCard document for mobile such as laptop computers.

Card-to-Cable (Figure 6)


Introducing a cable connection into the lane adds yet another source of loss and jitter. Cable-connected PCI Express is in early development. Yet, there are defined test points for the transmit and receive side of the cable in the draft PCI Express Cable Specification. In this topology, the reference clock and three other sideband signals are shipped down the center pairs of the cable along with up to X16 transmit and receiver pairs.

Don't Forget RefClk


Beyond the high speed serial data signal path, it is critical to test the reference clock (RefClk) signal integrity as well. Reference clock parameters are found in the CEM specification for Gen1 and are part of the main body of the Base Specification for Gen2. Changes in the Gen1 specification from Rev1.0a and Rev1.1 were



► Figure 4. Chip-to-Chip Lane.




► Figure 5. Connector part of the lane.



► Figure 6. Cable connection part of the lane.

largely driven by the affects of reference clock jitter on data eye closure. The Gen2 CEM system test requires that both data and RefClk are captured simultaneously to eliminate common mode jitter. This test methodology is referred to as the “Dual Port” method.



► **Figure 7.** Test points in PCI Express.

Compliance and Interoperability

Compliance and interoperability testing is key to the success of any standard. This is especially true in the early stages of a standard's lifecycle, when details are still in flux and interpretations may vary.

The PCI Express standards address the important issue of transmit and receive loss budgets. They also define compliance test points at which system-level testing must be performed. Essentially, compliance points are those at which system components (usually from different vendors) need to interoperate. Figure 7 summarizes some typical interoperability points in a complete system made up of interconnected elements. In PCI Express TP1 and TP4 are defined in the Base Specification and TP2 and TP3 are defined in the CEM and Cabling Specifications depending on the interface configuration.

These are shown as test points because they are specifically called out in the standard as the probing attachment points for test instrumentation. For a Card-to-Card or Card-to-Cable connection, it is important to understand whether the compliance point includes or excludes the connector from the measurement. Note that the compliance test point can not reside in the middle of the connector, since the mated and unmated connector performance differs quite substantially. In addition any ECB channel length or cable loss in your measurement can affect results. For Gen1, the PCI Express specifications (and PCI-SIG test procedures) for the most part neglected (or included in the test margins) the loss profiles of test fixtures and instrument cables. For 5Gb Gen2, consideration of these issues becomes a specification requirement. Another trend in PCI Express is to more clearly define procedures for receiver sensitivity testing at the system level. Receiver sensitivity test definitions for Gen1 are primarily left up to interpretation of the developer.

Physical Layer Compliance Testing: An Analog View Of A Digital World

PCI Express operates in the digital realm. But much of the compliance testing task consists of analog measurements. Why? Isn't this all about digital data?

The answer is a qualified “yes.” There are important exceptions. Digital signals exist in a world of distributed capacitance, noise, power supply variations, crosstalk, and other imperfections. Each of these phenomena detracts from “ideal” digital signals, sometimes to the point of compromising their ability to carry data. As a result, the specifications set limits on signal distortions and degradation. The device under test must meet these limits in order to be considered “in compliance.” Test conditions and test points are explained in detail in the standard. The PCI-SIG provides detailed test procedures for compliance testing at their compliance testing home page. These procedures include specific instruments down to make and model number. Also provided by the PCI-SIG is signal quality test (SigTest) software. SigTest performs the compulsory signal integrity measurements on waveform data captured from a real time oscilloscope for System and Add-In Card compliance points.


Figure 8 summarizes a typical compliance test process for PCI Express.

SigTest or RT-Eye® Software?

A common question that is asked by PCI Express developers is: If SigTest is free and tells me if I pass or fail, why do I need additional analysis tools in my oscilloscope? The answer lies in how much confidence you have in the design margins. Tektronix signal quality software tools such as RT-Eye, TDSJIT3, and JNB provide a broad set of capabilities for the validation, debug, and compliance of your design. This will become apparent in the subsequent sections.

Test Fixtures for Compliance Testing

Once test points are defined, test fixtures must be developed to provide probing attachment points. The PCI-SIG offers a Compliance Load Board (CLB) for



► Figure 8. PCI-SIG Compliance and Interoperability Process.

system testing and a Compliance Base Board (CBB) for Add-In card testing. The PCMCIA offers a system transmitter compliance card for notebook system transmitter testing.

Validation and Compliance Measurements

The PCI Express Base specifications includes amplitude, timing, jitter, and eye diagram measurements for the transmitter and receiver. These measurements are commonly known as signal integrity measurements. Insertion loss, return loss, frequency domain crosstalk are also specified for the channel. The CEM and Cabling Specifications define jitter, and eye diagram margins at the system level compliance points and represent the tests performed at plug-fests. Automated measurement/analysis tools are commonly used to speed the selection and application of these tests. The following list encapsulates some of the key measurements required at plug-fests (to prove compliance and interoperability) as well as measurements needed to validate system components to insure they meet the specification.

- Amplitude Measurements
- Timing Measurements
- Jitter Measurements
- RT-Eye Diagrams and Mask Testing
- Noise and Jitter
- Frequency Domain Measurements
- Receiver Sensitivity and Jitter Tolerance Measurements

Amplitude Measurements

Does the signal achieve the voltage levels and stability to reliably pass through the transmission medium and communicate a proper “one” or “zero” to the receiving circuit? These tests ensure that the signal has enough amplitude tolerance to do its job under worst case conditions.

Differential Voltage – A peak-to-peak differential voltage specification is the fundamental specification to guarantee that transceiver is transmitting the correct voltage levels. The minimum transmitter differential voltage is specified such that a stated minimum differential voltage will arrive at the receiver under worst-case media conditions (maximum loss). This ensures proper data transfer. The specification for PCI Express defines differential peak voltage as follows:

$$V_{TX-DIFFp-p} = 2 * |V_{TX-D+} - V_{TX-D-}|;$$

as measured from the waveform data after the clock is recovered from the data signal.

Note that the differential peak-to-peak voltage is not equivalent to the traditional pk-pk measurement you would find in an oscilloscope menu. This is often a source of confusion. A PCI Express specific algorithm is needed for this measurement.

De-Emphasis – This is the ratio of the amplitude of the second and subsequent bits after a “transition bit” to the amplitude of the “transition bit.” Other terms for de-emphasis are pre-emphasis and equalization. De-emphasis is used in serial data transmission systems to compensate for the frequency characteristics of “lossy” media such as the low-cost FR4 boards and connectors used in desktop computers. By making the transition bit higher in amplitude than the subsequent bits, the signal will arrive at the receiver pins with an “open eye.”

Common Mode Voltage Measurements (AC, DC) – common mode imbalance and noise on the transmitter can create undesirable effects in the differential signal. It is often useful to break apart the differential signal into its single-ended components to troubleshoot such issues. This technique also uncovers crosstalk and noise effects that may be coupling into one side of the differential pair and not the other.

Waveform Eye Height – Eye height is the data eye opening in the amplitude domain. It provides a very valuable measurement, because it represents the actual sample point of the receiver circuit. It is measured at the .5 Unit Interval (UI) point, where the UI timing reference is defined by the recovered clock. For more details about eye diagrams see the Eye Diagram section.

Timing Measurements

Is the signal free from timing variations, and are its transitions fast enough to preserve the critical data values the signal is meant to deliver? These tests, which require uncompromised performance on the part of the measurement toolset, detect aberrations and signal degradation that arise from distributed capacitance, crosstalk, and more.

Unit Interval and Bit Rate – Variations in the embedded clock frequency can be measured by looking at the mean measurement of the embedded clock over a large number of consecutive cycles. In PCI Express, if the mean measurement is more than 300 parts per million (PPM) away from the specified value (without SSC), then the device is not compliant.

Rise/Fall Transition Time – Tools for measuring accurate transition times are most important when making transmitter measurements. The minimum allowable transmitter risetime measurement for Gen2 PCI Express is 30 ps (measured with 20-80% reference levels). Risetimes that are too fast can cause EMI issues, while those that are too slow rise can cause data errors.

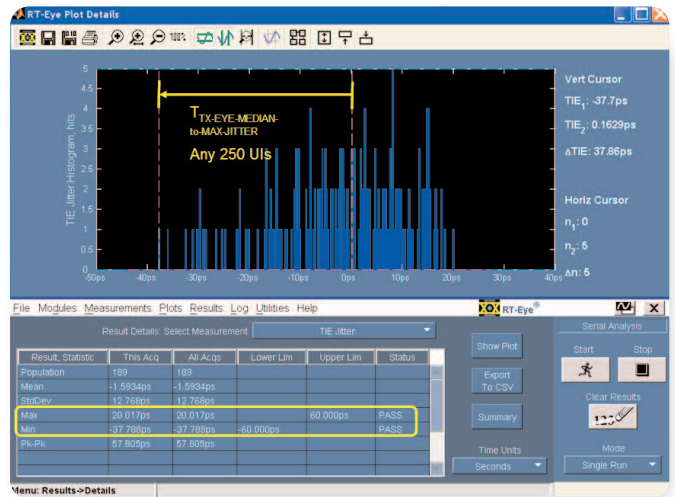
Waveform Eye Width – The waveform eye width is defined by one interval less the measured composite jitter; where the jitter is measured at the 0V differential reference level. For more details about eye diagrams see the Eye Diagram section.

Jitter Measurements


Jitter measurement methodology has been a major topic of discussion amongst PCI Express developers and authors of the specification. This section discusses the evolution of Jitter analysis in PCI Express as it applies to the measurements in the specification.

Median-Max-Jitter - The original Gen1 PCI Express specification (Rev1.0a) specified that “Median-Max-Jitter” was to be measured over any 250 consecutive bits. The Median-Max-Jitter is the maximum deviation in the positive or negative direction of the Time Interval Error (TIE) using the median as the timing reference. Figure 9 shows how the Median-Max-Jitter is measured from a histogram of the TIE over 250 unit intervals. Measurements limits (highlighted in Figure 9) are placed around the Min and Max values of the TIE which represent the Median-Max-Jitter specification.

The clock recovery method involves taking the mean of 3500 consecutive bits and then applying a 250 bit analysis window at the center of the clock recovery window. This method of analysis provides a good model for emulating a phase interpolator receiver behavior. The 3500:250 method approximates a third order filter function of the jitter (-60 dB/decade attenuation with a cutoff frequency of 1.5 MHz). It filters out low frequency SSC which enables the technique to be used on systems whether or not SSC is turned on.

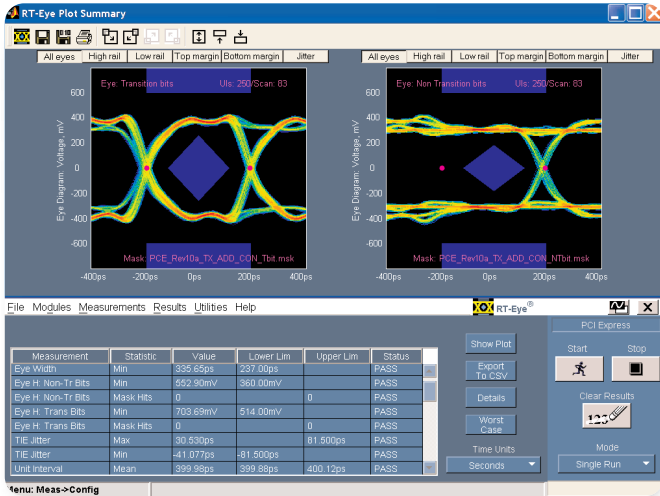


▶ **Figure 9.** Rev1.0a Median-Max-Outlier jitter over any 250 UIs.




▶ **Figure 10.** SmartGating used for 3500:250 windowing.

Tektronix has generalized this method in its RT-Eye software with the SmartGating feature. SmartGating (Figure 10) allows the user to specify clock recovery and analysis windows to represent different effective filter functions on the time interval error measurement. The ClockRecovery:Analysis window can also be set to scan the waveform with a user defined increment. Setting up a 3500:250 window that scans the waveform in increments of 100UI provides results that correlate to the PCI-SIG’s SigTest software.



▶ **Figure 11a.** Add-In Card test using Gen1 (Rev1.0a) method.


After experience with testing real world devices, it was determined that a more robust specification for jitter was needed to insure interoperability. The new specification (Rev1.1) measures transmitter jitter independently from the reference clock phase jitter. Clock recovery is achieved by applying a first order (-20 dB/decade attenuation at 1.5 MHz corner frequency) filter function to the jitter data and measuring it with a “clean clock” applied to the RefClk input of the transmitter device. Jitter in Rev1.1 is specified to be measured over a population of one million unit intervals, where the 1.0a specification was ambiguous in this regard specifying jitter “over any 250 unit intervals”. For Transmitter Component and Add-In card testing, the new Rev1.1 transmitter jitter



▶ **Figure 11b.** Add-In Card test using Gen1 (Rev1.1) method.

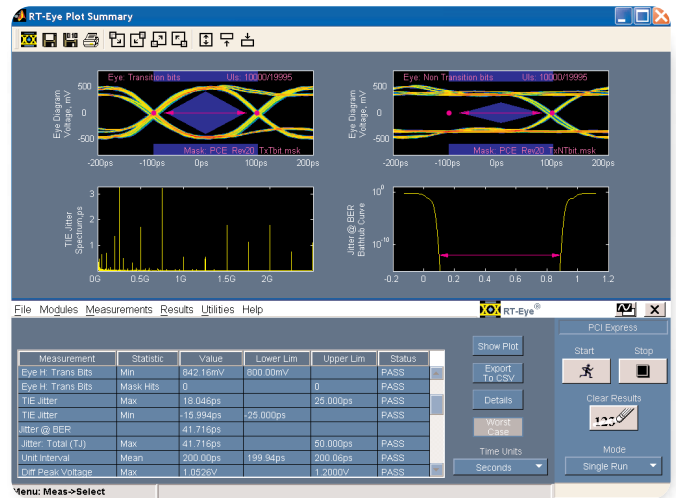
specification is to be used for compliance. For Systems with SSC turned on, the original third order function from the Rev1.0a specification is observed for compliance. Figure 11 show an example of an Add-In card being tested using the two specifications. Figure 11a is tested with a Rev1.0a Compliance Base Board (CBB) using the 3500:250 method, with scan mode enabled. Figure 11b is the same device tested with a Rev1.1 CBB using a first order software PLL. Observe that Rev1.1 has tighter jitter tolerances (wider mask), but since a clean clock is applied to the DUT, the measured TIE jitter is much smaller (Median-Max-Jitter is 27.6 ps vs. 41.07 ps). The end result is that the device passes under both methodologies.

**Assume bi-modal distribution (dual-Dirac), measure Tj at two BER
 Fit curve to points, slope is Rj, Intercept is Dj**




▶ **Figure 12a.** Dual-Dirac model of jitter estimation.

Dual-Dirac Jitter – As described in the Median-Max Jitter test, the Gen1 PCI Express specifications rely on TIE histogram approaches for measuring jitter. The Gen2 specification goes a step further in specifying jitter by using the Dual–Dirac model. This method estimates Total Jitter (Tj), Random Jitter (Rj), and Deterministic Jitter (Dj) as shown in Figure12a. Multiple acquisition and analysis techniques have been developed to derive the Dual-Dirac curve. One method, used primarily in Time Interval Analyzers (TIA) is to build a histogram of the Time Interval Error (TIE) over an extended period of time. Tj at two different BER levels can be measured; for example 10⁻⁴ and 10⁻⁷ BER. Once these values are measured, a straight line can be drawn to estimate the Tj at 10⁻¹² BER. The slope and intercept can then be drawn to derive the Rj and Dj values. In contrast, a real time (RT) oscilloscope using TDSJIT3 or RT-Eye estimates Tj at 10⁻¹² BER using the “spectrum approach to jitter measurement”. Tj is then measured at 10⁻⁹ using the same approach. Once the Tj at these two levels are



▶ **Figure 12b.** Transmitter test using Gen2 Dual-Dirac method.

determined, the same straight line can be established and Rj and Dj estimated. The spectrum approach to jitter (as developed by Tektronix) includes a patented technique for separating random and deterministic components.



▶ **Figure 13.** Differential RefClk Compliance test.

The spectrum approach has been correlated to BERT scan methods for measuring jitter. It has been widely adopted by serial data developers because it provides accurate jitter estimation in a fraction of the time of histogram based techniques. Figure 12b shows the results of a Gen2 PCI Express transmitter being tested using this technique.

Reference Clock (RefClk) Jitter – As discussed in the previous section, the Rev1.1 specification requires that transmitter jitter and RefClk jitter be measured independently. PCI Express transmitters and receivers will exhibit differing phase jitter tracking behavior due to variations in the transfer functions of their respective PLLs, differences in transport delay, and the tracking bandwidth of the CDR located in the receiver. It is necessary to specify the reference clock in terms of the amount of jitter that a worst case combination of transmitter and receiver will propagate and filter. This may be done by means of a jitter mask function that is specified in the Rev1.1 CEM and Gen2 Base Specifications. Also specified is the allowed Jitter at 10^{-6} and 10^{-12} BER. Figure 13 is the result of applying the required Rev1.1 CEM Specification to a RefClk device using the RT-Eye software. In addition to Total Jitter, the tool also



▶ **Figure 14.** Acquisition of RefClk and Data using Dual Port method.

measures parameters such as Slew Rate, High and Low Voltages, and Duty Cycle of the clock. For Gen2, the Dual Port method is used for data jitter compliance. With this method, the RefClk and data signals are captured simultaneously using a four channel oscilloscope with 50 GS/s sample rate on each channel. 10Meg record length is used to capture 1Million UI per the specification. In Figure 14, the Clock signal M1 is the result of Ch1-Ch2 and the data signal M2 is the result of Ch3-Ch4. By acquiring RefClk Data simultaneously, “common mode jitter” on the RefClk can be eliminated from the data jitter analysis.

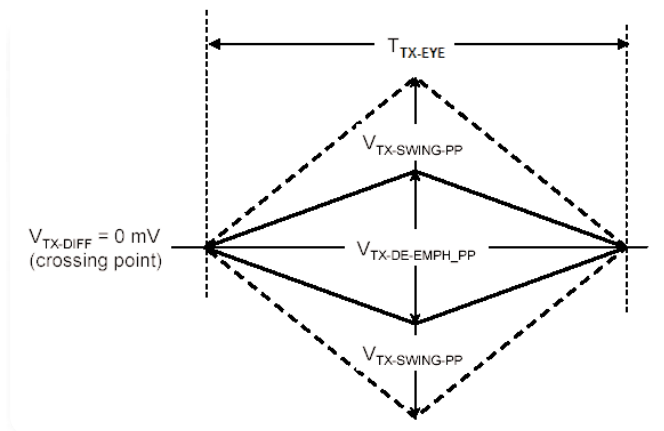
SSC profiling – In addition to signal integrity of the clock, the PCI Express specification specifies that the clock frequency be spread $-0.5%$ spread from the base data rate at a modulation frequency of 30-33 kHz. This will result in a spread of the data period from 400 ps to 402 ps. In order to achieve a profile of the SSC from the data signal, the data period must be measured and filtered. TDSJIT3 Advanced Jitter Analysis software can be used to analyze the spread spectrum profile on a long record from a high performance real time oscilloscope.

Real Time Eye Diagrams and Mask Testing

The “eye diagram”—an oscilloscope display that results from the overlap of a series of data waveforms—is an important tool for establishing the quality of the PCI Express transmitter signal. To produce an eye diagram, it is necessary to recover the clock from the data in order to synchronize the data stream carrying random, pseudorandom, or compliance pattern bits. The Tektronix patented technique of overlaying real time waveform data with a software recovered clock is the bases for the RT-Eye (Real Time Eye) software that runs on Tektronix real time oscilloscopes. As discussed in the previous section, there are multiple techniques for recovering the clock, depending on the device under test and the specification version. In addition to recovering the clock, the transmitter and interconnect test points require that transition bits and de-emphasized bits be separated and mask testing be performed on each bit type. Figure 15 shows the definition of the mask geometries as they appear in the specifications. Figures 11 and 12 show how the resulting eye diagrams are displayed simultaneously using the RT-Eye software. As discussed earlier, The T_{TX-EYE} specification for Gen1 is a ‘waveform eye mask’ defined as one unit interval less the composite TIE jitter. In Gen2, T_{TX-EYE} is specified as a ‘jitter mask’ that is defined as one unit interval less the total jitter at 10^{-12} BER using the Dual-Dirac model. Amplitude measurements in both cases are voltage based as acquired by an oscilloscope.

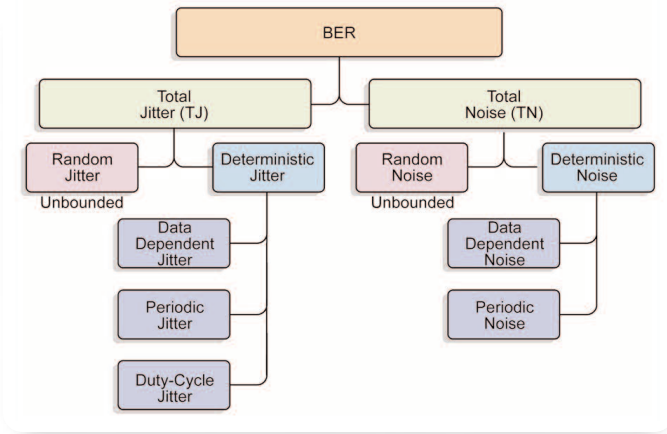
What about Noise and its contribution to Jitter?

The analysis techniques described thus far for jitter measurement are predominantly performed by real time oscilloscopes. The filtering techniques such as the 3500:250 method for the transmitter and the phase jitter filter for RefClk rely on edge information that is captured in real-time. Real time acquisition is the only way to insure that the frequency components of the jitter are



► **Figure 15.** PCI Express Waveform Mask Specifications.

known, and thus can be filtered in the frequency domain using post processing algorithms. However, vertical noise in a real-time oscilloscope does add error in both the jitter and the amplitude domain. This is due to the fact that real-time oscilloscopes acquire waveforms using an analog amplifier at the front end. Since the Noise x Bandwidth product for an analog amplifier is constant, the more bandwidth you have, the more noise you will have in the acquired signal. Also, real-time oscilloscope’s Analog to Digital (A/D) Converter resolution is limited to eight bits by design, resulting in unavoidable quantization errors. The combination of system noise and quantization combine to add error to the measurements. In many cases, these errors are not significant and can be ignored, however, as PCI Express developers design longer transmission channels, the resolution and signal-to-noise ratio of a real-time oscilloscope may become limiting. For these applications, equivalent time sampling oscilloscopes can be used for higher precision measurements. Additionally, new software tools have been developed to not only analyze the jitter of the high speed serial data signal, but the noise as well.

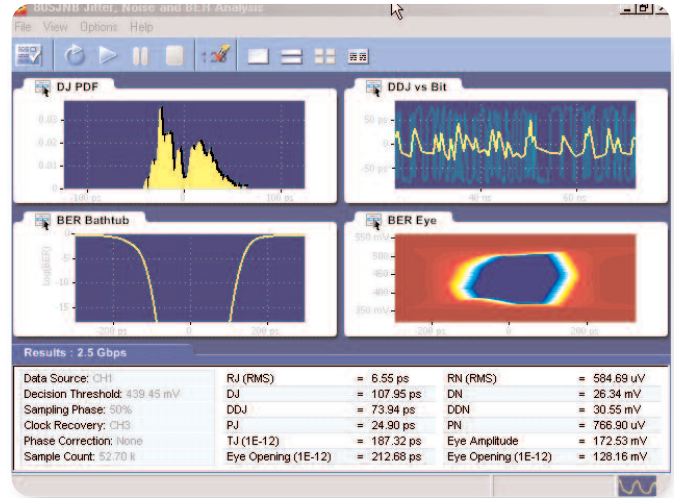


▶ **Figure 16a.** Jitter and Noise Decomposition.


In Figure 16a, the left side of the figure shows traditional decomposition of serial data jitter. The right side of the figure shows the addition of a new technique used in sampling oscilloscopes that decomposes the noise in the amplitude domain much the way jitter is decomposed in the time domain. The result is a BER Eye that represents a two dimensional Probability Density Function (PDF). Figure 16b shows the result of this analysis on PCI Express Add-In Card using TDS8JNB (Advanced, Jitter, Noise, and BER Analysis software) on a Tektronix sampling oscilloscope, providing further insight into precise causes of eye closure.

Frequency Domain Measurements

The interconnect link compliance testing increasingly require frequency domain measurements, such as insertion loss, return loss, and frequency domain crosstalk. Return loss measurement is also frequently required to characterize the input of a receiver. These frequency domain measurements, which are also referred to as S-parameters, are primarily done in differential mode, since the differential mode measurement directly relates to the bandwidth degradation and jitter in the digital system. Some standards begin to require common mode measurements or mode conversion measurements.

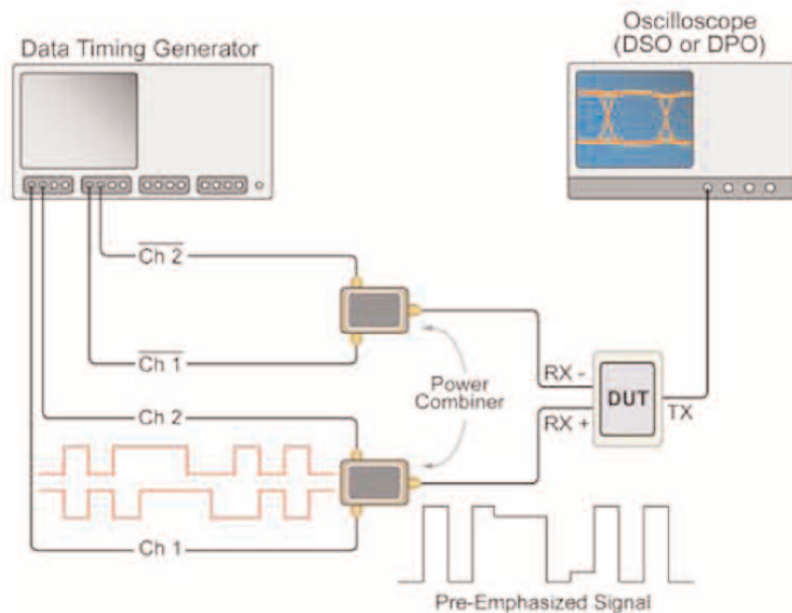


▶ **Figure 16b.** Jitter and Noise PDF from JNB Software.



▶ **Figure 17.** Frequency domain crosstalk measurement.

This measurement can be performed with a TDR-capable oscilloscope with true differential capability and of appropriate rise time, which directly relates to the measurement bandwidth. Figure 17 shows a frequency domain crosstalk measurement for a PCI Express interconnect as analyzed using 80SSPAR IConnect S-Parameter and Z-line software running on a Tektronix sampling oscilloscope.



► **Figure 18.** Creating a de-emphasized signal from two DTG channels.

Receiver Sensitivity and Jitter Tolerance Measurements

Gen1 receiver testing requires a stimulus source to drive the device under test (DUT) configured in a loop-back mode. The most economical tool for the job is a pattern generator (also known as a data timing generator) that can generate defined 2.5 Gb/s test patterns such as the training sequences (TS1 & TS2) and apply them to the device. When a training sequence is recognized by the receiver section of the DUT, the transmitter section sends out a similar sequence. This transmitted sequence can be observed and analyzed with an oscilloscope and/or logic analyzer. It is common practice to alter the training sequence such that the DUT's performance can be characterized under various stressful conditions. Specific stresses include amplitude level variations, eye-crossing level changes, differential skew variances and added noise and jitter.

Amplitude levels are not the only adjustable in a DTG; other adjustments include:

Timing tolerance – frequency can be varied to search for receiver limits.

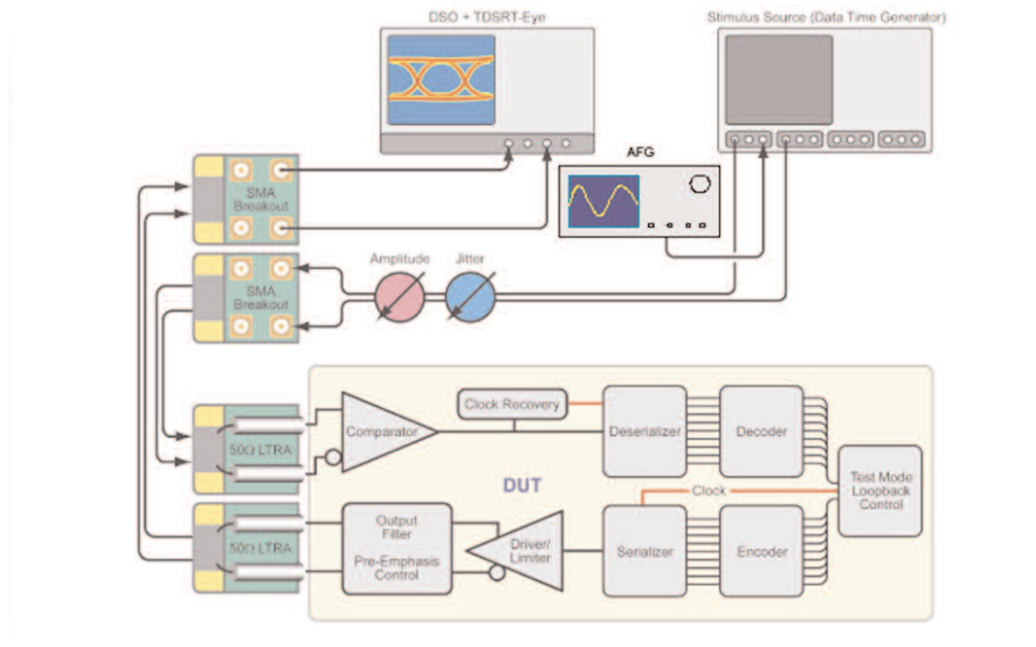
Crossing level tolerance – the crossing voltage between the D+ and D- legs of the differential signal can be varied to simulate common mode asymmetry.

Skew tolerance – the delay between the D+ and D- legs of the differential signal- can be varied to simulate differential skew.

Jitter tolerance – jitter frequency modulation and amplitude can be varied to exercise the device to specified requirements.

For Gen2 devices running at 5 Gb/s, a faster generator is required. An Arbitrary Waveform Generator like the AWG7102 allows for not only data generation but Direct Synthesis capability enables signal generation of equalization, De/Pre Emphasis or modulated jitter signal.

Jitter tolerance for receivers is defined as the ability to recover data successfully in the presence of jitter. Meeting the specification guarantees that the SerDes and PLL circuits can recover the clock even when a certain amount of jitter is present. The test for jitter tolerance is one of the more important tests used to guarantee interoperability. Rigorous jitter tests are especially critical in applications such as PCI Express. There is an absolute requirement for a signal source that can supply jitter with specific amplitude and frequency modulation characteristics.

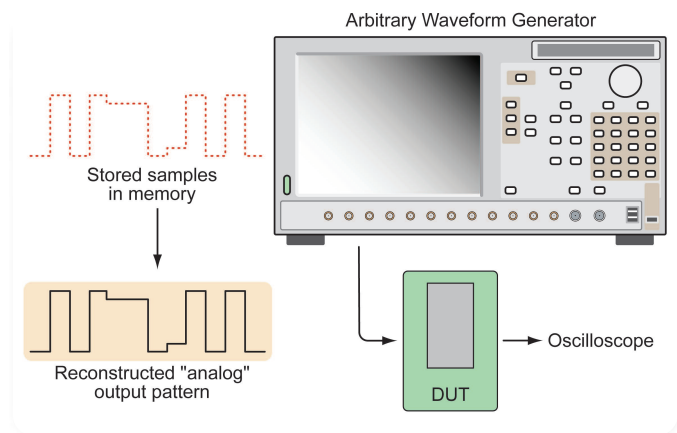


▶ **Figure 19.** Setup for Jitter Tolerance testing.

The PCI Express Base Specification specifies jitter as the variation of the eye diagram crossing points relative to a recovered Transmitter Unit Interval (UI). As shown in Figure 19, a DTG provides jitter control in terms of both amplitude and frequency content. With this combination, thorough jitter tolerance testing of PCI Express receivers can be accomplished.

Alternatively, with an AWG using the Direct synthesis technique to create signals that represent what data looks like after propagation, transmission line, and system effects have taken their toll, it is possible to eliminate the ancillary equipment formerly required for modifying and modulating DTG produced signals. With a suitable direct-synthesis AWG, jitter in any form can be merged into the test signal itself. The effects of both random and deterministic jitter can be modeled. The AWG's ability to vary imposed effects such as jitter in both quantitative and qualitative terms is also valuable, and easy to accomplish.

The same is true of noise and aberrations overshoot. Because of the fundamental architecture of the AWG platform allows essentially and signal shape to be defined, the small details of signal behavior are easy to model. And though small, these details are exactly what make up a "real world" signal. Using an AWG, it is possible to closely approximate the content of a signal that has passed through FR4 traces, connectors, and cables.



▶ **Figure 20.** The arbitrary waveform generator reconstructs sampled waveforms stored in its memory.

Insuring Signal Fidelity in Analog Measurements

Every PCI Express measurement, whether for compliance testing, validation, or debug purposes, consists of probing the device, acquiring the signal data, and analyzing the captured information. Selecting the right tools, and applying them correctly, is the way to ensure fast, accurate results. This section of the document discusses different techniques for connecting to the device under test and analog performance considerations when choosing an oscilloscope.

Connecting to the Serial Device Under Test

Measurements begin with a connection to the device under test. The question arises: what is the mechanical interface of the lane to be tested? This is a question usually answered in recommended test definitions or dictated by access points in the design. There are diverse configurations, each with its own unique characteristics.


The mechanical interface has a profound impact on probing for both electrical and logical testing. The probe may need to match tiny physical features, while delivering high-speed differential signals intact to the measurement instrument. Alternatively, the probe may need to adapt to prototype fixturing such as test points fitted with SMA connectors.

There are four approaches to the probing challenge:

- Pseudo-Differential Active Probing
- True Differential Active Probing
- Pseudo-Differential SMA Connection
- True Differential SMA Connection
- TriMode Probing

Remember, PCI Express signals are transmitted differentially, so all of the solutions explained here are meant to connect differentially.

Pseudo-Differential Active Probing - Two differential active probes, one on each side of the differential signal, can be used for pseudo-differential measurements. Figure 21 shows a Chip-Chip application example where two probes can be soldered to the D+ and D- legs of the differential pair with respect to ground. Two channels of the oscilloscope capture two channels of data. The probe can capture both AC and DC components of the signal which makes it the best tool for looking at common mode effects on a differential link. While the link's receiver sees only the




► Figure 21. Pseudo-Differential Active Probing.


differential swing from the transmitter, variations in the common-mode voltage can cause undesirable amplitude variations in alternating cycles. The signal measured in Figure 21, shows an example of asymmetrical signals captured by a pseudo-differential probing setup. The PCI Express Base specification calls out the maximum common-mode AC voltage (V_{CM}) and also stipulates an acceptable range for common mode DC voltage. The differential voltage waveform V_{DIFF} is achieved by subtracting the D+ and D- waveforms. Eye diagram, amplitude, and jitter and timing measurements are made on this math waveform. Because the two sides of the waveform enter two separate oscilloscope input channels in pseudo-differential probing, the inputs must be de-skewed before making any measurements.

True Differential Active Probing - Figure 22 shows a differential probe capturing the receive side of a connector-based Card-Card serial link. Unlike the pseudo-differential connection, this probe requires only one oscilloscope channel and makes the subsequent math steps unnecessary. This offers, among other advantages, the ability to use the multiple channels of the oscilloscope to capture multiple lanes. It is also useful for debugging of multiple high-speed test points. The differential waveform in Figure 22 requires only one channel of the oscilloscope and does not need deskewing. In this case, only differential mode measurements are available. It is important to note that any probe will impose some loading to the Device Under Test (DUT). Every probe has its own circuit model whose impedance can change with increasing frequency. This can affect the behavior of the observed circuit and influence the measurement, factors that must be considered when evaluating results. When designing “Chip-Chip” interfaces it is important to design test pads into the design of the board, and to consider the mechanical requirements of the intended probes. Probing pads should be placed as close as possible to the receiver termination resistors to avoid signal distortions due to reflections.

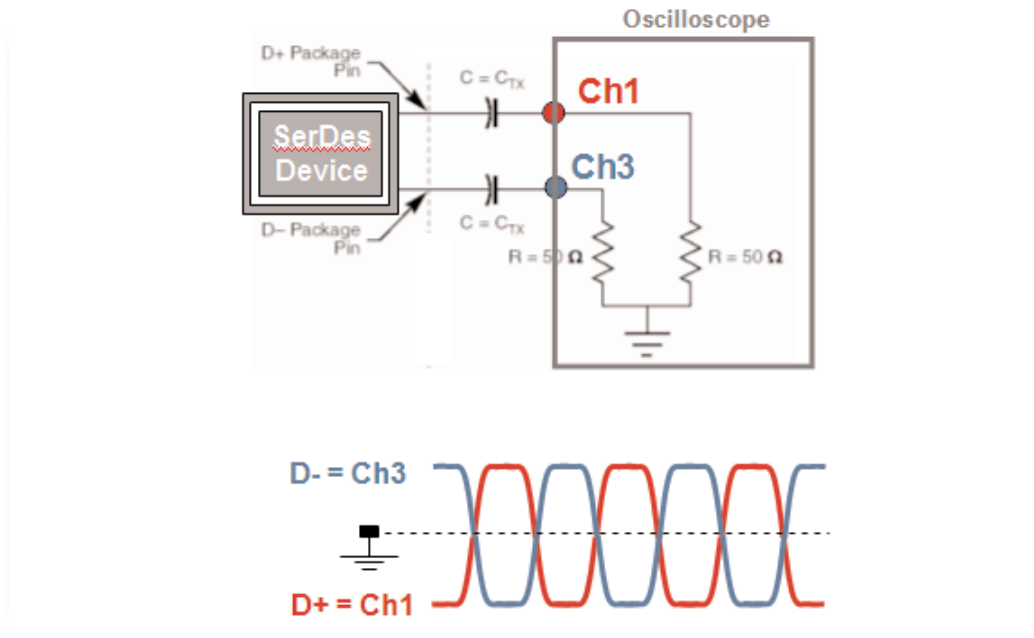
TriMode Probing - With the use of a new probing technique from Tektronix called TriMode probing. Each leg of the differential signal D+ and D- can be acquired separately by choosing (A) or (B); the differential signal Vdiff can be acquired using (A-B); or the common mode signal can be captured using the $[(A+B)/2]$ selection on the probe body. This technique combines the probing techniques in the previous two sections into a single probe. Figure 23 shows the probe tip and how the connections are made.



▶ Figure 22. True Differential Active Probing.



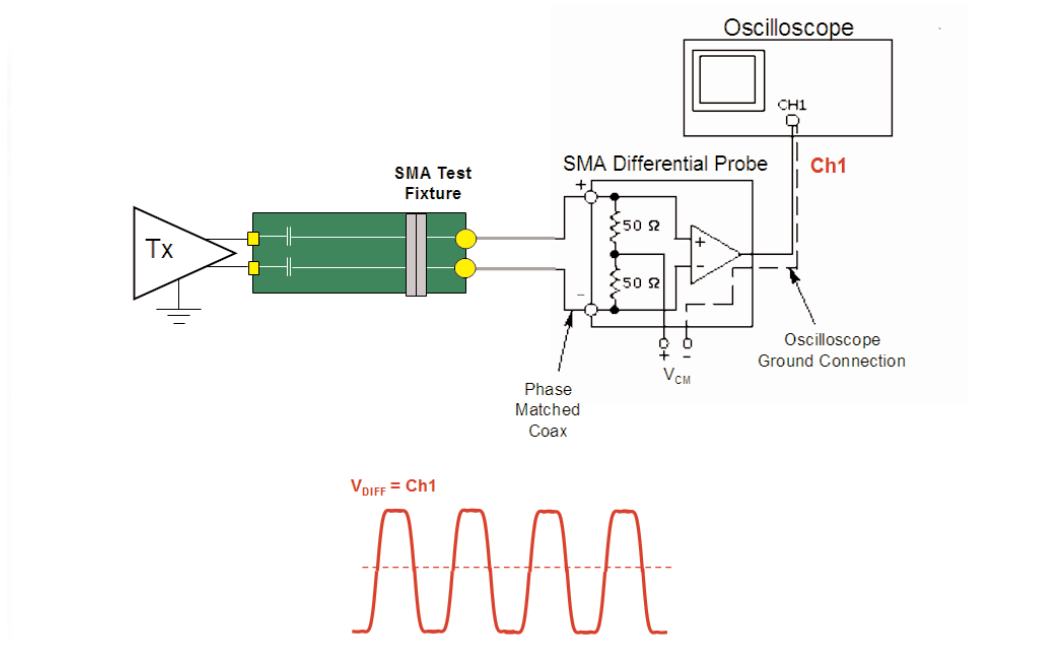
▶ Figure 23. TriMode Probing.



► **Figure 24.** Pseudo-differential SMA Connection.

Pseudo-Differential SMA Connection - Many compliance test fixtures and prototype circuits expect characteristic load at their outputs. If the fixture/prototype is fitted with SMA high-frequency connectors, the SMA pseudo-differential approach provides access to the signals. An example of this is a transmitter compliance test shown in Figure 24. Here the output of the DUT connects directly to two inputs of the oscilloscope, each of which has an input impedance of 50 Ω. SMA inputs are common on high performance oscilloscopes and do not require the purchase of a differential probe to make

measurements. As shown in Figure 24, common mode AC measurements can be made on both legs of the differential pair. The signal must be AC coupled because the oscilloscope is terminated to ground. As in the active probing case, the differential mode waveform is created by subtracting one common mode signal from the other. When using this technique, care must be taken that that channels are deskewed and that any loss associated with the cable from the DUT to the input of the oscilloscope is taken into account in the measurements.



▶ **Figure 25.** True Differential SMA Probing.

True Differential SMA Connection – Figure 25 shows a new class of oscilloscope probe designed to address transmitter path compliance measurements. The SMA input differential probe is a good fit for compliance tests where interoperability points are defined at the “Card-Card” or “Card-Cable” interface. A 100 Ω matched termination network properly terminates both legs of the differential signal to any common mode voltage. The common mode voltage can be provided by the user or in some cases directly from the oscilloscope. Figure 25 illustrates a true differential

SMA probe that converts the single ended signal is to differential using a differential amplifier. This eliminates the need to deskew the channels. In the case of the Tektronix P7380SMA (Differential Signal Acquisition System with SMA Inputs), the loss associated with the matched SMA cables is calibrated out of the probe when it is manufactured, providing more accurate measurement results than the psuedo-differential SMA approach where cable loss needs to be either calibrated out or considered as margin.



► **Figure 26.** Pulse response of an ideal low pass filter.

Analog Acquisition Considerations


System Performance Drives Measurement Accuracy

The pairing of oscilloscopes and their probes must be considered as a system when choosing solutions for critical high-speed measurements. The instrument's performance becomes part of the measurement. The analog acquisition characteristics of the system as well as its digitizing specifications all play a role in delivering meaningful measurement results. Understanding the performance specifications also plays a key role in choosing the right instrument for the task at hand.

Evaluating The Acquisition System's Performance

Bandwidth – The term “bandwidth” denotes the frequency range that the oscilloscope's input circuit (vertical amplifier) will admit, within certain tolerances. The instrument's sample rate is usually optimized to

handle this same range. It is important to insure that the bandwidth of the system is adequate for performing the measurements. How much bandwidth is enough? The answer to this question depends on the frequency content of the signal under test. Imagine the input of the oscilloscope as an ideal “brick wall” filter, and your 5 Gb/s PCI Express signal (transmitting 101010 repeating pattern) as a perfect 2.5 GHz square wave, the signal displayed on the oscilloscope screen would be as shown in Figure 26. The three plots show what happens as the ideal bandwidth filter is varied in BW. Note that an oscilloscope that only passes the 3rd harmonic will result in much higher amplitude and transition time error than that of an oscilloscope that passes the 5th harmonic. For best signal fidelity, it is clear that at least the 5th harmonic is required to minimize measurement error.



► **Figure 27.** 5 Gb/s captured with 15GHz oscilloscope.

In practice, most oscilloscopes are not “brick wall” filters and the roll-off characteristics of the oscilloscope can have tremendous implications on signal fidelity. For example, a 15 GHz oscilloscope is calibrated to have flat frequency response (0dB of loss) at 12.5 GHz; it’s 3 dB down at 15 GHz; and has some energy at 17.5 GHz (the 7th harmonic). This results in a very accurate eye diagram response as shown in Figure 25 as it measures the 5 Gb/s output of a high speed BERT generator. The mask in the figure represents the Gen2 PCI Express transmitter transition bit mask. Not much margin if a lower bandwidth instrument is used.

Risetime – As shown in Figure 27, bandwidth and rise time are related. The PCI Express Base specification calls out the minimum transition time of a transmitter device. For Gen1, the minimum transition time is 0.125 UI (50 ps) and for Gen2 it is 0.15 UI (30 ps). It’s important not to confuse the traditional rise time specification of the oscilloscope (which is measured at 10-90% levels)

with the transition time measurement in the specification (which is measured at 20-80% levels). For this reason, most oscilloscope manufacturers will put both 10-90% and 20-80% risetime specifications on their datasheets. If the system risetime is much less than that of the specified measurement, then the accuracy will be higher. For example, in the case of the 15 GHz real time oscilloscope, the 20-80% risetime specification is 19 ps, which results in around a 5% transition time measurement error on the fastest Gen2 transition of 30 ps.

5th harmonic accuracy is a minimum requirement for transmitter compliance testing (at the pins of the SerDes). Further precision for characterization purposes can be achieved with a sampling oscilloscopes with bandwidths up to 70GHz and risetimes down to 5 ps (10-90%).

For downstream measurements at the CEM and other interconnect test points, it is often sufficient to use a lower bandwidth (slower rise time) oscilloscope. The reason for this is once the signal leaves the pins of the transmitter and proceeds down the copper transmission channel, the signal is further filtered and the 5th harmonic content of the signal is reduced. If the transition times of the signal are slowed to a point where they are a high percentage of the unit interval, the error in the rise time and amplitude are greatly reduced when using a lower bandwidth instrument. However, the only way to truly know if you’re getting an accurate measurement with a lower bandwidth oscilloscope is to make the same measurement with an instrument with 5th harmonic performance and compare the results.

Sample Rate – In real time oscilloscopes, the sample rate required to accurately reconstruct an analog signal is defined by the Nyquist sampling theorem. This states that of a sample rate greater than two times the highest frequency content of the input signal is required to capture enough data to accurately reconstruct the waveform. Therefore, an instrument with 15 GHz input bandwidth is supported with a 40 GS/s sample rate. In sampling oscilloscopes sample rate is orders of magnitude lower, but is not a primary performance measure.

Interpolation – For signals with fast transition times (like the 30-50 ps transitions in PCI Express), $\sin(x)/x$ interpolation of the sampled data increases the accuracy of amplitude-based measurements and eye diagrams. The interpolation factor should be set to insure that there are at least three sample points that occur on the edge being measured. For jitter measurements in which a reference level is taken at the switching threshold or “zero level” of the differential signal, interpolation typically has negligible affect on the results. Some real time oscilloscopes provide $\sin(x)/x$ interpolation as part of the acquisition path; it is the default means of interpolation. This provides accurate results without slowing down the waveform and measurement throughput.


Noise Floor – In real-time oscilloscopes, jitter noise floor (JNF) can be defined by the random component of the Time Interval Error (TIE) on a signal with all data dependent jitter removed. It can be measured directly by providing an ideal clock signal (101010) repeating pattern to the input of the scope and measuring the standard deviation (RMS) of the TIE. Highest performance real time oscilloscopes provide a random jitter noise floor of typically <400 fs RMS. A similar setup with a sampling oscilloscope (although not real time) measures typically <200 fs RMS due to the lower noise floor and higher digitizer resolution. Real time oscilloscopes

have 8 bit digitizer resolution DC, where sampling oscilloscopes provide 14 bits. As discussed earlier, the lower resolution of the real time oscilloscope combined with the noise contribution from its input amplifier contribute to the 2X higher relative JNF.


Record Length – Record length in a real time oscilloscope essentially defines the frequency content of the signal that can be captured. This becomes especially important when making jitter measurements. Low frequency spectral components of the TIE (Jitter Spectrum) can only be measured if they are captured. Long record length oscilloscopes allow you to capture these frequencies at full sample rate. For example, with a 100M record length, you can capture 5 msec of waveform data at 50 ps/pt. This represents 12.5 Million UIs of Gen1 PCI Express traffic or 150 cycles of the spread spectrum clock.

In PCI Express, low frequency jitter concerns have led to the Rev1.1 specification to require 1Million UIs of statistical certainty. In order to capture 1Million UIs in a single acquisition, an oscilloscope must have at least 8M record length. In order to filter out high frequency components and resolve the 30 KHz SSC profile from the data signal, more than 10 cycles (333usec); or 13M of data must be captured.

Triggering – In a real-time oscilloscope, trigger performance can be paramount in debug applications. In order to capture signal anomalies that can cause bit errors, the instrument must be able to trigger at the same bandwidth of the signal. For example, anomalies in a 5 Gb/s signal could be glitches on the order of 100 ps in width. At a system level, triggering on specific symbols can also be important. The next section describes a few use cases where Tektronix unique trigger system can be used to debug PCI Express design issues.



▶ Figure 28a. Bit Error Detected.



▶ Figure 28b. Trigger on the Bit Error.

Debug the Link Using Pinpoint® Triggering and Data Decoding Tools


Locate and Trigger on a Bit Error

Often it is the case that in the middle of making signal integrity measurements, unexpected behavior occurs that suggests that there may be a design issue. In these instances, debug that become critical to isolating and resolving the issue quickly. Often analysis tools can be used in conjunction with the instrument's trigger system to pinpoint these issues. Figure 28a shows a case where the RT-Eye software detects an infrequent bit error. At first glance it appears to be a 200 ps pulse occurring randomly and infrequently within the 2.5 Gb/s data. By turning on the RT-Eye software's Bit Error

Locator, the eye diagram (upper left) and the acquired waveform (upper right) are displayed simultaneously. When a mask violation occurs, every sample point that violates the mask is highlighted with a red dot on the waveform view. By zooming in on the waveform plot detail (bottom), the true nature of the anomaly is brought to light. The failed bit is actually a 600 ps wide bit! Now that the failed bit has been discovered, what can be done to debug it? Since we know that the pulse is 600 ps wide, we can setup a Pulse Width trigger to trigger the oscilloscope only when 600 ps pulses occur. Figure 28b shows the trigger setup that pinpoints the 600ps wide pulse in the data stream enabling further investigation into the root cause.

An Introduction to PCI Express™ Measurements

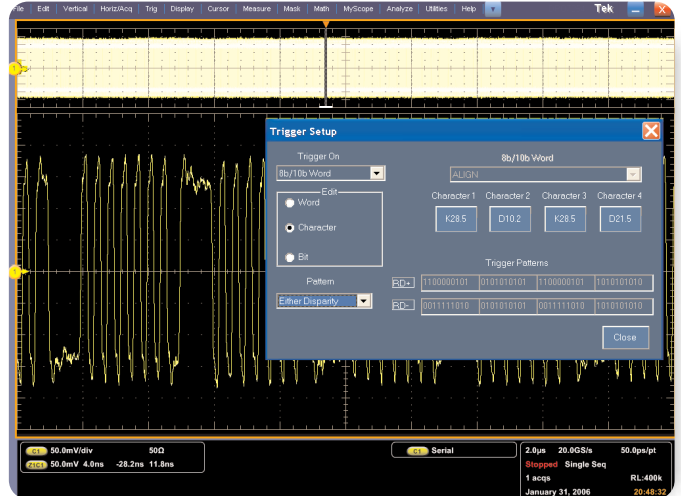
► Primer



► **Figure 29a.** Decoding the data from the waveform.

Validate and Trigger on 8b/10b Symbols

In addition to bit errors occurring in the analog portion of the PHY, it is important that the proper digital data is being transmitted over the link. While the trained eye can look at an analog waveform and convert it to its symbols (K28.5, D10.2, etc.) doing so can be tedious and error prone. Using a real time oscilloscope with PTD (Protocol Trigger and Decode) software (Figure 29a),



► **Figure 29b.** Trigger on any combination of four symbols.

the data is decoded automatically from the acquired waveform (bottom) and a listing of the symbols is also provided (middle) enabling the validation of the digital data. Figure 29b shows the PTD software's trigger setup menu that allows you to trigger on any four symbols (forty bits) of the data. The SerDes based trigger also allows you to trigger on disparity and character errors in real time.

Digital Validation and Debug

Physical layer compliance tests are often the beginning of a long and arduous product development process. Beyond analog complinace, the designer must validate the correctness of the protocol and ensure devices are compatible during system integration. After all, the device has to pass in a multitude of system configurations. Debugging and design validation extend beyond compliance testing and the electrical sub-block of the physical (PHY) layer up to the logical sub-block of the PHY and the Data Link and Transaction layers.

Within the protocol stack, observing the logical sub-block of the PHY layer calls for a specialized analyzer to acquire and interpret the data as 8b/10b symbols. Historically logic analyzers have been the preferred tool for the job. Protocol analyzers are designed to capture the highest layers of the protocol, and are commonly limited to just one type of serial bus. If the design task requires a simultaneous view of diverse buses or general purpose signals in both serial and parallel formats, the logic analyzer is the tool of choice.

Protocol analyzers have historically worked well when the bus is operating correctly, whereas a logic analyzer displays physical layer aspects of the bus, including detailed insight into link training and power management states.


In a logic analyzer, captured data is displayed in the waveform window for cross bus analysis or analog correlation and the state listing window, which provides protocol decode and error detection. With appropriate analysis tools, this recorded data can be automatically interpreted to help engineers confirm that higher-level programmatic instructions are being carried out correctly, and discover the source of errors that have caused a failure on the PCIe link.

Until recently, acquiring serial data with a logic analyzer meant using a complex external pre-processor to interface the instrument to the device under test. But with serial bus architectures becoming almost universal in digital systems, a more efficient solution is needed... and has emerged. Serial acquisition is now integrated into the logic analyzer itself in the form of serial modules such as the Tektronix TLA7Sxx Series Serial Analyzer Modules. In effect these acquisition instruments are interchangeable in a mainframe with their parallel counterparts. This new capability dramatically simplifies serial acquisition and enables designers to mix serial and parallel functions as needed.

Table 1 compares some of the strengths of the logic analyzer versus protocol analyzer platforms.

Serial Acquisition Feature	Logic Analyzer	Protocol Analyzer
Packet Triggering	√	√
PHY Layer Analysis	√	
General Purpose Debug	√	
Correlated serial and parallel acquisition data	√	
Correlate Digital with Analog in a Common Data Display	√	

▶ **Table 1.** A comparison of logic analyzer and protocol analyzer features.



► **Figure 30.** 16 channel, full width, mid-bus footprint profile and pinout.

Probing Means Making the Right Connections

All measurements start with probing, and preserving signal fidelity is just as much a concern here as it was with the analog acquisitions described in earlier sections. In those sections we discussed the concept of equipping prototype boards with SMA connectors to deliver maximum signal fidelity to an oscilloscope. Equivalent preparations are important when acquiring high-speed serial data with a logic analyzer. A probing scheme that “slows down” edges and delays pulse can contribute false errors to the data stream.

At PCI Express data rates of 2.5 Gb/s or 5.0 Gb/s, it is not enough to simply clip on a common LA probe. Flying leads and micro-grabbers will introduce ground loops and other unpredictable analog effects. One common practice is to route signals to “headers” made up of pins that mate to mass-terminated cables connected to the LA. But this approach too will impact low-voltage, high frequency serial signals. Fast-changing data edges

simply will not tolerate the electrical stub and the losses contributed by a connector. The ideal probe for PCI Express acquisition will offer high analog bandwidth, high impedance, low capacitance, and minimal stub effects.

For PCI Express and other high-speed serial protocols, two probing methods are prevalent: the “interposer” probe that plugs into an existing PCI Express slot; and the “mid-bus” probe that sits atop a footprint made up of pads that are a non-intrusive part of the signal traces. The mid-bus method requires advance planning and a small amount of board real estate but the landing pads have almost no impact on the signal when the probe is not attached. In addition the mid-bus probe is ideal for chip to chip inter-connectors that may require full visibility during system validation. Figure 30 shows the pinout of a single-direction 16 channel mid-bus probe. The footprint of the 8 channel probe is similar though it reflects the reduced pin count.



▶ **Figure 31.** The mid-bus serial probe connects to dedicated test pads on the circuit board's surface, consuming minimal board real estate.

Both mid-bus and interposer probes are available with the new generation of Serial Analyzers. Importantly, both utilize an active-buffered architecture that passively observes the signal, taps off the signal, and buffers it to the serial analyzer modules.

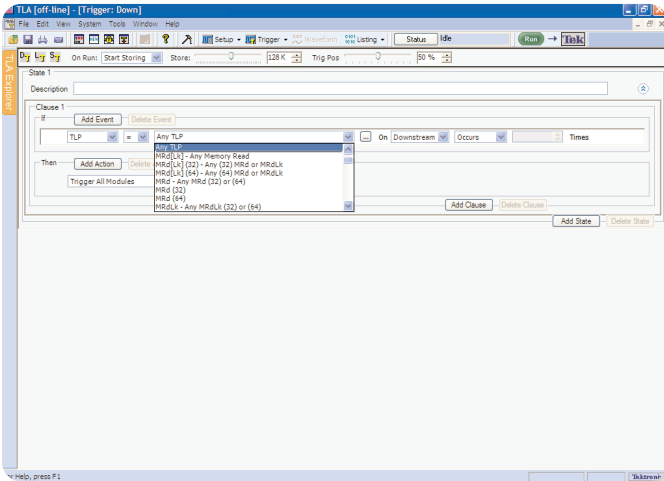
A repeater design creates a new copy of the signal, potentially with different edge characteristics. In effect you are not acquiring the actual signal from the physical

layer, as a result some types of errors may be artificially corrected. An active-buffered architecture is the only solution that offers visibility into the true signal at full speed.

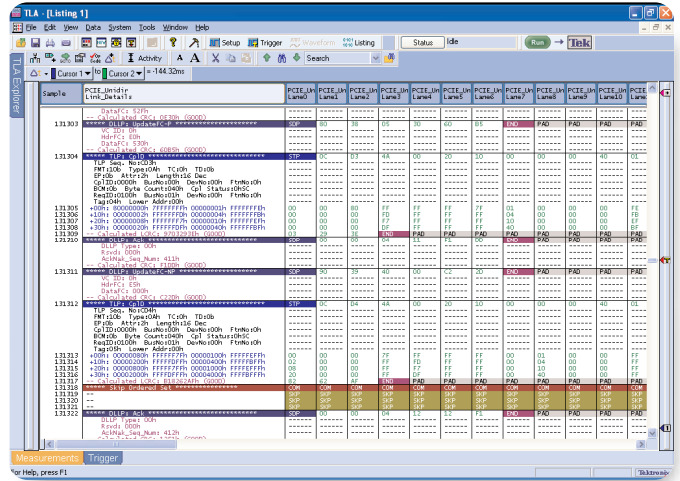
Figure 31 depicts a mid-bus probe connected to a PCI Express serial analyzer module housed in a Tektronix TLA7012 logic analyzer mainframe.

An Introduction to PCI Express™ Measurements

► Primer



► **Figure 32.** Protocol-based trigger criteria simplify the setup step.



► **Figure 33.** The serial listing view provides deep details about PCI Express packets and symbols.

Triggering Tools Boost Troubleshooting Efficiency

One of the defining features of any logic analyzer is the flexibility of its triggering system. Debug work proceeds quickly when the logic analyzer permits real-time triggering on packet elements such as headers and payload or raw symbols.

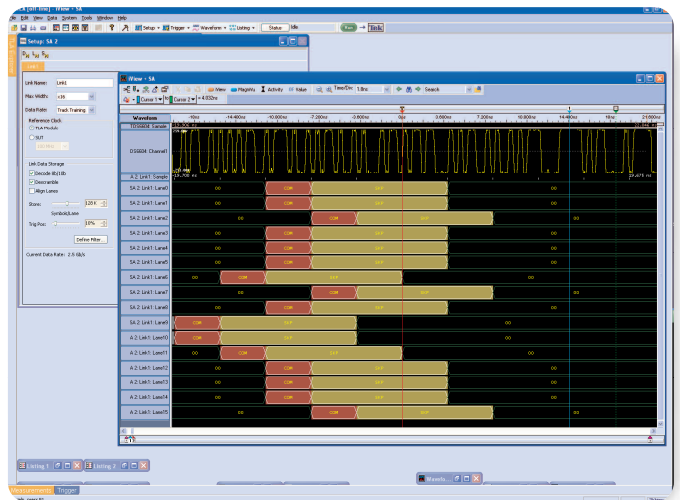
The TLA7Sxx Series simplifies serial triggering with a protocol specific trigger user interface. The serial analyzer's advanced triggering capabilities can quickly isolate and capture specific transactions, packets, ordered sets, link events, or link errors.

There are various triggering templates such as the one shown in Figure 32. These allow the user to select from menu items and fill in a “form” to specify events of interest. The fields are optimized for the unique requirements of the PCI Express serial protocol.

Analyzing the Results

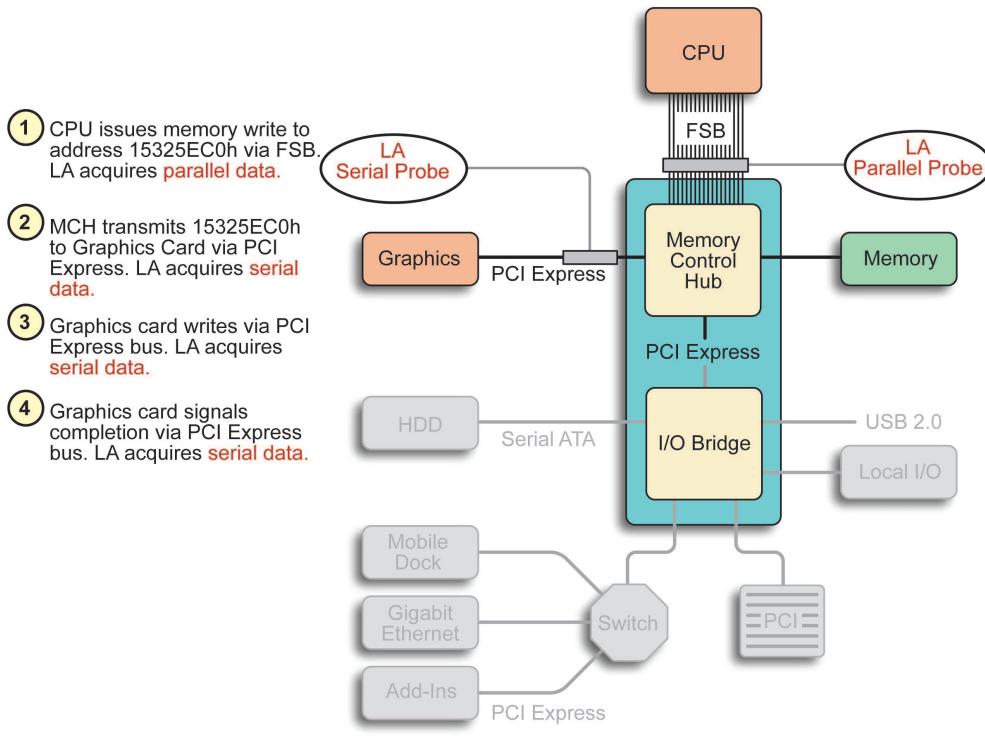
With the acquisition completed, the stored data must be decoded into meaningful results and displayed in data windows for analysis. The serial analyzer PCI Express package includes sophisticated software tools to disassemble, decode, and display the captured data in a packet-style view using the listing window. A x16 PCIe acquisition is shown in the listing window in Figure 33.

The display includes three elements: a packet summary, detailed decoding of the fields of a packet/control symbol, and the raw data. Color-coding differentiates the text in each of these three elements and distinguishes control symbols from packets in the packet/control symbol summary.



► **Figure 34.** Time-aligned packet data shown as digital waveforms , correlated with an analog waveform. This display is equivalent to the conventional logic analyzer timing view, but it represents decoded serial data.

Figure 34 is a waveform view of serial data. The serial analyzer delivers advanced disassembly features, such as deep capture of serial buses time correlated to all other system buses using the logic analyzer's common system timestamp. The disassembler provides control symbol decoding and display of individual symbols for all layers of the protocol. Simultaneous decoding of both transmit and receive data ports is available along with any other system busses being probed by the logic analyzer. Additionally, using iView, analog and digital waveforms can be time corrected.



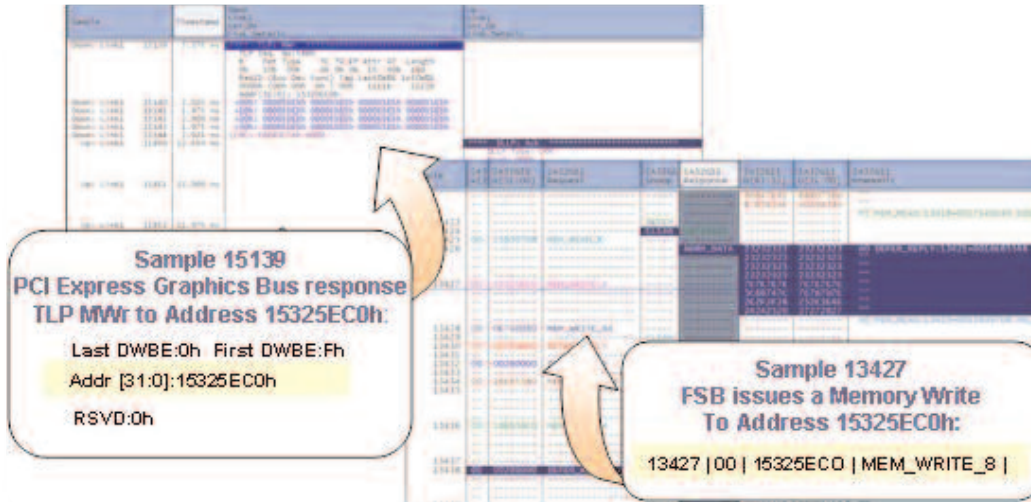
▶ **Figure 35.** Typical motherboard block diagram showing the parallel and serial test points.

Looking at the “Big Picture” of a Multibus System

As previously explained, a logic analyzer equipped with a serial module is in a unique position to capture uncompromised time-correlated data from any and all buses within a system, whether serial or parallel. While it is rarely necessary to capture every single bus at once, the need to acquire a full set of PCI Express lanes along with the concurrent Front Side Bus transaction or the parallel output of a critical register, for example, is common. The following example illustrates this point.

Figure 35 depicts a typical (though simplified) motherboard architecture. Buses that are not relevant to this discussion are shown in gray, although they too may require time-correlated acquisition and analysis at some point.

Suppose the CPU issues a Memory Write instruction to the graphics card. Let’s assume the destination address is 12325EC0h, a 32-bit value expressed in hexadecimal form.



► **Figure 36.** A time-correlated acquisition from serial and parallel buses on the motherboard.

The instruction exits the CPU and travels across the Front Side Bus (FSB) to the Memory Control Hub (MCH). A parallel logic analyzer probe acquires the FSB. This data is recorded by acquisition modules such as the TLA7AA4. Figure 36 shows the result of this acquisition. The window in the foreground is the FSB trace.

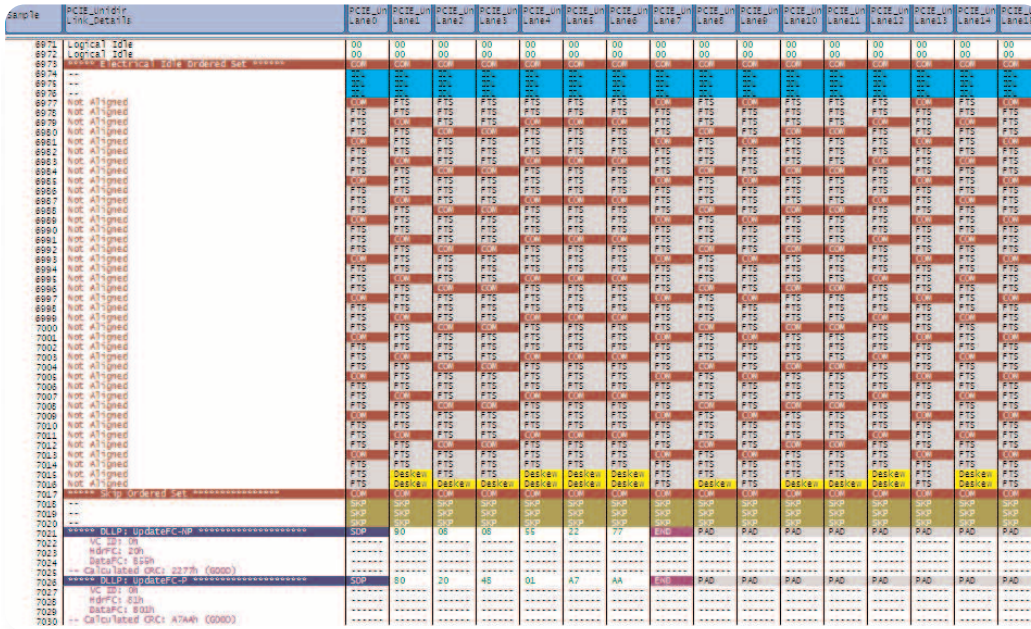
The MCH, receiving the instruction, tells the Graphics card to Write to memory address 15325EC0h. The instruction is relayed through the serial PCI Express bus connecting the MCH to the Graphics card. The logic analyzer’s mid-bus “connectorless” serial probe sits on a group of pads on the bus, capturing the serial traffic. The result acquired by the TLA7S16 PCI Express module is shown in the background window in Figure 36. In this example, the transaction has executed correctly—the Graphics card has written to the specified address.

Note the samples in which the two respective events

occurred. The FSB event was recorded 1712 samples earlier than the response from the Graphics card. The two events are time-correlated. The 1712-sample separation is equivalent to a time interval that remains consistent across all transactions of this type.

Capturing Power Management Transitions

The importance of Active State Power Management (ASPM, or simply power management) has grown with each successive generation of mobile computing devices. And with the increasing emphasis on “green” products and energy conservation, power management plays a role in desktop and server systems as well. Consequently the validation of power management features in all types of computing platforms has taken on a new level of urgency.



▶ **Figure 37.** The logic analyzer syncs up within the timespan of ten FTS packets, maximizing the likelihood of capturing errors during the LOs-to-LO transition.

The logic analyzer is the tool of choice for analyzing events that occur during the transitions to and from power management states. Inevitably some of the signals of interest appear on one or more PCI Express buses within a system, which necessitates a logic analyzer that can carry out uncompromised acquisition of multiple PCI Express links.

The most critical requirement for observing the change of power management states is the ability to synchronize quickly to the link as power states are exited. The transition passes quickly and a tool that’s slow to respond might easily miss the very cycles in which a problem occurs!

Two PCI Express states and the transition between them are the subject of the acquisition. The normal operating state of the PCI Express bus—routine

transmission and reception of packets—is known as “L0.” One power-saving state is known as L0 Standby, or L0s. A Fast Training Sequence (FTS) Ordered Set is used to achieve bit and symbol lock when transitioning from the L0s state to the L0 state. While there is no way to avoid losing some of these FTS packets in the earliest cycles of the transition, the logic analyzer must minimize these losses.

Figure 37 is a logic analyzer display of a transition sequence from L0s to L0. The information was captured with a TLA7S16 PCI Express module, which has industry-leading performance in this key area. As the screen clearly shows, synchronization has occurred within the time span of ten FTS packets. Now the instrument will acquire all remaining FTS packets and continue to capture L0 cycles.

Overview of PCI Express Measurement Solutions

At this point, it should be clear that serial compliance measurements require the resources of high-performance measurement instruments. Five classes of instrumentation are most often used for PCI Express validation and compliance testing as well as debug tasks.

Oscilloscopes

The tool of choice for measuring the electrical portion of the active PHY is a high-performance oscilloscope. State-of-the-art digital oscilloscopes can capture waveforms in the multigigabit range with very good signal integrity.

Once a clock is recovered from a serial bit stream, from the waveform the oscilloscope can create an eye diagram. Eye diagrams created by oscilloscopes provide a good view of signal characteristics. An extensive library of built-in measurements is available for immediate quantitative evaluation of both the eye diagrams and of the acquired waveform.

Further processing with dedicated tools will yield the Time Interval Error (TIE). The TIE can be analyzed to separate random and deterministic jitter, an analysis that can then also provide an estimate for total jitter at the 10^{-12} bit error rate (BER).

Aside from combining the functions of several instruments into one, the oscilloscope has other advantages: probing flexibility, rich display, triggering capabilities, and more.

Real Time (RT) Oscilloscopes

Real-time oscilloscopes capture continuous, contiguous data records. In case of serial link analysis the oscilloscope needs to capture at least several samples from each bit. Today's state-of-the-art digital storage oscilloscopes (DSO), digital phosphor oscilloscopes (DPO), and digital serial analyzers (DSA) can sample up to 20 ps sample-to-sample (sampling rate of 50 GS/s); with up to 20 GHz bandwidth they can characterize well signals up to 8 Gb/s. The continuous, real-time record feature of real-time oscilloscopes gives them these special advantages:

- Digital Signal Processing (DSP) software algorithm can recover the embedded clock from the digitized serial data bit stream. This (sw-based) method of clock recovery is the most flexible one; moreover it avoids the need for a clock recovery hardware and its unavoidable jitter.
- Since the signal path in the best real-time oscilloscopes performs up-to or nearly up-to the BW of the instrument, the extensive triggering capabilities of the hw triggers in these real-time oscilloscopes can be used to trigger on data or events of interest – often capturing events so rare as to be invisible through other means.
- Real-time capture of the data allows analysis with the least number of constraints; the whole bandwidth of the instrument is captured, so e.g. the jitter spectral information is complete and unaliased for all jitter types; and even completely random data can be captured and stored/decoded.
- Finally, the completeness of the data capture is invaluable in debug of unexpected, unpredictable behaviors.

Equivalent Time (ET) Sampling Oscilloscopes

The oscilloscopes for analyzing signal integrity at very high-speed the equivalent-time Sampling oscilloscopes, also known as “sampling oscilloscopes” and as “communication (signal) analyzers”. With a maximum bandwidth of 70+ GHz, these instruments are capable of analyzing optical and electrical signals from below 1 Gb/s to 40 Gb/s and beyond. Because sampling oscilloscopes do not attempt to capture at real-time speed they have the opportunity to perform more precise signal capture: so the sampling oscilloscopes offer higher digitizer resolution, and a superior noise performance – both of these features are advantageous for acquisition of low voltage signals common in serial links today. Similarly, since the samples are captured in only “equivalent-time” sequence they can be spaced within femtoseconds of each other, removing any concerns about sample-to-sample spacing or interpolation. On the other hand the jitter analysis is more complex and sometimes constraints the signal measured, e.g. to a repetitive pattern.

Because of their high bandwidth, sampling oscilloscopes also offer TDR and S-parameter measurement capability, so eliminating the need for a separate VNA (Vector Network Analyzer) for performing the S-parameter measurements on the serial data devices.

Sampling oscilloscopes can also be equipped with clock recovery; in this case the function is provided by hardware.

Some real-time oscilloscopes do offer both Real Time (RT) and Equivalent Time (ET) Eye reading techniques using hardware clock recovery, each of which has their advantages.

Signal Generators

Good high-speed engineering practices include exercising designs under “real-world” conditions. The right tool for mimicking these conditions as closely as possible is a programmable signal generators. Generating test signals at today's data rates requires high-speed Data Timing Generators (DTG) and Arbitrary Waveform Generators (AWG). Without these instruments there would be no way to test and validate new physical layer designs. Many signal sources can replay signals that have been captured with an oscilloscope. The signals can act as a reference signal or may be modified to stress the device under test.

Data Timing Generators are especially useful for generating multiple streams of channels of parallel data—up to 96 channels in today's most advanced instruments. These tools deliver 3.3 Gb/s data rates. At the same time these advanced instruments provide a host of signal manipulation features including independent level, rise/fall and jitter controls.

High-speed digital signals inevitably have analog attributes. Arbitrary Waveform Generators can provide stimulus signals with analog content (usually deliberate impairments) on a bus channel. Capable of delivering any type or shape of waveform, AWGs are universally applied in design and manufacturing. Current AWGs have sample rates up to 20 GS/s with 5.8 GHz data rate bandwidth.

Logic Analyzers

The preferred tool for measuring the logical sub-block of the PHY and the Data Link and Transaction layers of PCI Express is the logic analyzer. Unlike both RT and ET oscilloscopes, logic analyzers provide protocol disassembly of all layers of the link with packet level, symbol level, and link event triggering across all lanes of the link. The purpose of the logic analyzer is to simplify acquisition and analysis of the purely digital aspects of both serial and parallel transmission. To carry out the serial bus debug mission, the LA must deliver features consistent with the needs of high-speed buses: physical layer acquisition, deep memory, flexible triggering and synchronization with other system buses, and more. And like the oscilloscope, it must offer low impact probing tools that provide direct insight into physical layer signals.

Summary

PCI Express Technology is here to stay, and is destined to grow in importance as the computer market continues to demand ever-accelerating data rates. Design and validation engineers have a new and perhaps unfamiliar discipline— serial compliance measurement—to learn even as they confront aggressive development schedules and fast-changing standards.


Fortunately, industry working groups such as the PCI-SIG exist to stabilize the technology and disseminating information about design, architecture, and compliance requirements.

At the same time, measurement tools, ranging from real time and sampling oscilloscopes to logic analyzers and signal sources, help engineers deal with PCI Express measurement challenges. These solutions deliver the performance to capture, display, and analyze the most complex serial signals. Thanks to these innovative, automated tools, engineers can perform compliance and validation tests quickly and easily which enables faster time to market.




DSA70000 Series

The DSA70000 Series are the new generation of real-time DPO and are the industry's best solution to the challenging signal integrity issues.



TDS8200 Series

The TDS8200 Series are designed for research, design evaluation, and manufacturing test for applications requiring bandwidths into tens of GHz.




TLA7000 Series

The TLA7000 Series provides breakthrough digital systems analysis tools to capture and analyze the source of elusive problems. The TLA7000 Series provides the speed you need to capture the source of those elusive problems, plus the visibility you want with large displays and fast system data throughput.



DTG5000 Series

The DTG5000 Series combines the power of a data generator with the capabilities of a pulse generator in a versatile, bench-top form factor. Its modular platform allows you to easily configure the performance of the instrument to your existing and emerging needs to minimize equipment costs.



AWG7000 Series

The AWG7000 Series Arbitrary Waveform Generator delivers a unique combination of superior signal stimulus, unrivaled sample rate, bandwidth and signal fidelity and uncompromised usability. This family offers the industry's best solution to the challenging signal stimulus issues.

Contact Tektronix:

ASEAN / Australasia (65) 6356 3900

Austria +41 52 675 3777

Balkan, Israel, South Africa and other ISE Countries +41 52 675 3777

Belgium 07 81 60166

Brazil & South America (11) 40669400

Canada 1 (800) 661-5625

Central East Europe, Ukraine and the Baltics +41 52 675 3777

Central Europe & Greece +41 52 675 3777

Denmark +45 80 88 1401

Finland +41 52 675 3777

France +33 (0) 1 69 86 81 81

Germany +49 (221) 94 77 400

Hong Kong (852) 2585-6688

India (91) 80-22275577

Italy +39 (02) 25086 1

Japan 81 (3) 6714-3010

Luxembourg +44 (0) 1344 392400

Mexico, Central America & Caribbean 52 (55) 5424700

Middle East, Asia and North Africa +41 52 675 3777

The Netherlands 090 02 021797

Norway 800 16098

People's Republic of China 86 (10) 6235 1230

Poland +41 52 675 3777

Portugal 80 08 12370

Republic of Korea 82 (2) 6917-5000

Russia & CIS +7 (495) 7484900

South Africa +27 11 206 8360

Spain (+34) 901 988 054

Sweden 020 08 80371

Switzerland +41 52 675 3777

Taiwan 886 (2) 2722-9622

United Kingdom & Eire +44 (0) 1344 392400

USA 1 (800) 426-2200

For other areas contact Tektronix, Inc. at: 1 (503) 627-7111

Updated 1 June 2007

Our most up-to-date product information is available at: www.tektronix.com



Copyright © 2007, Tektronix. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

09/07 FLG/WOW

4HW-19375-1

Tektronix
Enabling Innovation

