BANDWIDTH TESTS REVEAL SHRINKING EYE DIAGRAMS AND SIGNAL INTEGRITY

Avoid Performance degradations on high-speed buses like DDR3, PCI Express, Intel® QPI, USB and others

PROBLEMS



A SIGNAL INTEGRITY WHITEPAPER



A Signal Integrity Whitepaper

By Tim Caffee and Eric Johnson



Tim Caffee, Vice President Design Validation & Test – obtained dual degrees, B.S. degrees in Mathematics and a B.S. degree in Computer Science, at Old Dominion University. He is one of the founders at ASSET InterTech, Inc. He as spend the last 21 years working at ASSET in different roles of the company to help customers achieve success in new technology deployment and business development. Over the last ten years, he has been collaborating with Intel to enable new markets in validation and test based on embedded IP within Intel silicon. Prior to ASSET, he spent several years at Texas Instruments and held numerous positions in the Defense Group. Outside the office, he enjoys traveling, mountain biking, and has a passion for great cuisine.

Eric Johnson, Technical Product Manager, High-Speed I/O Tools – has been involved with circuit board design and test for 15 years. Following graduation from Iowa State University with a Bachelor's Degree in Electrical Engineering, Eric was a hardware designer for Rockwell Collins. He designed a variety of circuit boards, including several that were integrated into test equipment. At Rockwell Collins, he became an expert in boundary-scan test development. He joined ASSET 12 years ago as an applications engineer and soon advanced to the position of Senior Technologist. He currently serves as the Technical Product Manager for ASSET's ScanWorks High-Speed I/O products and solutions.





Table of Contents

Executive Summary	4
Less room for error	5
What to do about it	12
Signal integrity validation from the inside out	13
Conclusions	15
Learn More	15

Table of Figures

Figure 1: Example ideal eye diagram of three unit intervals
Figure 2: An eye mask (in grey) indicates the limits of the operating margins for the signaling on a bus. A signal that strays onto the masked area would have poor signal integrity and likely not meet the requirements of its specification
Figure 3: Eye diagram of the first generation of a high-speed serdes channel with a transfer rate of 6 GT/s
Figure 4: Eye diagram of the second generation of a high-speed serdes bus with a transfer rate of 8 GT/s
Figure 5: Eye diagram of the third generation of a high-speed serdes bus with a transfer rate of 10 GT/s
Figure 6: Superimposing the eye diagrams for three generations of a high-speed bus illustrates the dramatic reduction in the size of the eye diagram and the resulting erosion of operating margins on the bus
Figure 7: An eye diagram generated by ASSET InterTech's ScanWorks® from data gathered by instruments embedded in chips

 $\ensuremath{\mathbb{C}}$ 2012 ASSET InterTech, Inc.

ASSET and ScanWorks are registered trademarks while the ScanWorks logo is a trademark of ASSET InterTech, Inc. All other trade and service marks are the properties of their respective owners.



Executive Summary

Factors like jitter, inter-symbol interference (ISI), crosstalk and others can create havoc on the signal integrity of high-speed serdes and memory channels, making maximum bus speeds difficult to achieve in practice. Compounding this predicament is the fact that channel speeds keep increasing from one generation bus technology to the next. With each step upward to a higher speed and higher signaling frequencies, a serdes or memory bus becomes more susceptible to distortions and anomalies which can effectively disrupt bus traffic and stall system throughput. For serdes buses like PCI Express[™], Serial ATA, USB, Intel's QuickPath Interconnect (QPI) and memory buses like DDR, the higher the frequency of the signaling, the more susceptible the interconnect becomes to errors, re-transmissions and other anomalies.

To avoid potential problems with high-frequency bus traffic the signal integrity on the bus must be validated during each of the major phases of a system's life cycle, including design/development, manufacturing and as an installed system in the field. If the signal integrity on a serdes channel is not what it should be, steps should be taken to correct the problem and improve system performance.

Unfortunately, effectively and economically validating the signal integrity on a high-speed bus has become more difficult as the limitations of legacy probe-based test equipment have become more obvious in recent years. Now though, non-intrusive software-driven test methods based on embedded instrumentation are providing alternative validation solutions that are more cost-effective and deliver observed signal integrity data. These methods provide soft access to hard data. In addition, industry specifications like the IEEE P1687 Internal JTAG (IJTAG) standard for embedded instrumentation are emerging to simplify and streamline the adoption of signal integrity validation techniques based on embedded instruments.



Less room for error

There's an old story around the electronics industry that compares advances in chip technology with progress in the automotive industry. The story goes like this: *If automotive technology had kept pace with Silicon Valley and the development of semiconductor technology, motorists today could buy a vehicle with a V-32 engine and it would be capable of speeds of 10,000 mph; or a 30-pound car that gets 1,000 miles to the gallon — either one at a sticker price of less than \$50.*

The high-speed serdes and memory buses in today's systems are a lot like that V-32, 10,000 mph car screaming down a crowded highway with thousands of other cars doing the same thing. There isn't much margin for error at 10,000 mph! Drawing a parallel with serdes signals, what if one car (bit) is going too fast? Or another is too slow? Or the one in the next lane is swerving into the wrong lane? They will start crashing into each other and causing problems just like jitter, inter-symbol-interference (ISI) and crosstalk affect real-world serdes and memory buses. And as the speed of a bus increases, the margin for error decreases. Disruptions in traffic are more likely to occur.

Each new generation of a high-speed bus brings with it a faster speed and bigger design challenges. As another example, consider a car that can travel at a speed of 60 mph one day, 80 mph the next and 100 mph the day after that. While the driver may arrive at work sooner on day 3, there is far less margin for error at 100 mph than 60 mph. At 100 mph a momentary distraction, mechanical problem with the car or an object in the road could cause a serious accident. The margin for error on today's high-speed buses is similarly slight and it shrinks with each incremental step up in signaling frequency. In other words, as the transfer rates on these buses increase (Table 1) with each new generation of bus technology, operating margins become smaller and smaller. If the quality of the signaling on a bus is not optimal, bus performance can suffer, causing an adverse effect on system performance. Even the fastest processors cannot overcome degraded transfer rates on clogged I/O and memory buses.



Bus Speeds	
DDR	
DDR1	400 MT/s
DDR2	1.066 GT/s
DDR3	2.133 GT/s
DDR4	4.266 GT/s
PCIe	
PCIe 1	2.5 GT/s
PCIe 2	5 GT/s
PCIe 3	8 GT/s
PCIe 4 (2014/2015)	16 GT/s
USB	
USB 1-Low	1.5 Mb/s
USB 1-Full	12 Mb/s
USB 2	480 Mb/s
USB 3	5 Gb/s
Serial ATA (SATA)	
SATA 1	1.5 Gb/s
SATA 2	3 Gb/s
SATA 3	6 Gb/s
Direct Media Interface (DMI)	
DMI 1 (2004)	10 Gb/s
DMI 2 (2011)	20 Gb/s

Table 1: Escalating speeds on several prevalent high-speed buses

Most engineers think of signal integrity in terms of an eye diagram, which is an overlay of multiple data bits over one or two bit cycles or unit intervals (UI's). The eye diagram plots time on the X axis and voltage along the Y axis. (Figure 1). For some bus types, the center portion of the plot resembles an open human eye, hence the name eye diagram.



A Simplified Eye Diagram





An eye mask is placed over portions of the eye diagram (Figure 2) to aid in determining the integrity of the signal. If the signal strays into the masked portion of the diagram, the bus is not meeting the requirements of its specification and, as a result, many errors could be encountered during its operation. A good signal – one with high signal integrity – would never touch the masked areas and ideally have a certain additional margin for error between signal and mask.





Figure 2: An eye mask (in grey) indicates the limits of the operating margins for the signaling on a bus. A signal that strays onto the masked area would have poor signal integrity and likely not meet the requirements of its specification.

Typically, the signals on each new generation of a high-speed serdes or memory bus operate at a higher frequency to increase data speeds. As the frequency of a signal increases, the size of the eye will decrease, bringing the edges of the signal closer to the masked area in the center. While masks may change from one bus generation to the next, one can assume that each new generation will have less margin for error. In essence, this reduces the margin on the bus to the point where the slightest anomaly will adversely affect signal integrity. A discussion of three generations of a hypothetical serdes bus will illustrate this point.

When the first generation of this high-speed serdes was introduced it had a transfer rate of six gigatransfers per second (GT/s). The eye diagram at this speed is decidedly open (Figure 3). In fact, the width of the eye is 166 picoseconds (ps).





Figure 3: Eye diagram of the first generation of a high-speed serdes channel with a transfer rate of 6 GT/s.

As often happens, a second generation of this particular high-speed serdes with a higher theoretical speed was introduced several years later. This higher speed version of the bus featured higher signal frequencies, increasing the transfer rate to eight GT/s (Figure 4). The higher frequencies caused the eye to shrink by 25 percent, reducing the width of the signal from 166 ps to 125 ps. Such a reduction in the eye leaves considerably less operating margin. Consequently, chances will be greater that a slight deviation from the acceptable signal path would result in the signal straying into the unacceptable mask area of the eye diagram.







As soon as the faster second generation of this hypothetical serdes was introduced, designers quickly began developing a faster third generation with a higher base signal frequency and a 10 GT/s transfer rate. Plotting the eye diagram for the third generation (Figure 5) shows that the eye has once again shrunk significantly. From generation 2 to generation 3, the width of the eye was reduced from 125 to 100 ps or by 20 per cent. Operating margins have shrunk yet again.





Figure 5: Eye diagram of the third generation of a high-speed serdes bus with a transfer rate of 10 GT/s.

A critical point to note is that the significant reductions in the eye diagram could have taken place over the course of a few short years. Assuming the elapsed time between the replacement of generation 1 with generation 2 and the introduction of generation 3 was only two years, the operating margins on this bus could have been reduced by 40 percent – or nearly half – in just two years. That is, the width of the eye went from 166 ps in generation 1 to 100 ps with the introduction of generation 3 two years later. Superimposing the eye diagrams for all three generations (Figure 6) highlights the dramatic reduction in eye size and the consequent erosion in operating margins over what could have been a very short period of time.





Figure 6: Superimposing the eye diagrams for three generations of a high-speed bus illustrates the dramatic reduction in the size of the eye diagram and the resulting erosion of operating margins on the bus.

The discussion above illustrates how faster bus speeds reduce the ideal eye width or UI. A shrinking UI is challenging enough, but there are also many other influences that affect signal integrity and the eye diagram. Among these, jitter, voltage attenuation and PVT (process, voltage, temperature) variations can significantly affect signal integrity during every stage of the product lifecycle.

Jitter is an anomaly in the signal that causes signal transitions to occur before or after they should. Jitter is made up of many component parts, but it is often broadly categorized as deterministic jitter (Dj) and random jitter (Rj). Depending upon its causes, deterministic jitter can usually be minimized through proper board design and signal routing. But random jitter is just that – random – and normally it cannot be controlled because it occurs unexpectedly or in a nondeterministic way.

No matter how well designed the chips, packaging and circuit boards are, some amount of jitter will be present on every bus. On an actual eye diagram, this is shown by the fact that the plot of the signal is not a single compact line. Instead, the plot of the signal is thick, indicating the range of the variability in the signal caused by jitter. This can be seen by comparing the ideal plot in Figure 1 with the real-world eye diagram in Figure 2.



Additionally, the amount of jitter on a serdes or memory bus does not typically shrink as signaling frequencies are increased to achieve higher data transfer rates. Instead, as the eye closes at higher frequencies, jitter becomes a more significant portion of the signal. In other words, at higher frequencies it is more likely that jitter will cause signal integrity problems because the operating margins on the bus have been reduced significantly, while the amount of jitter has remained relatively the same.

The earlier example of the hypothetical high-speed serdes can also illustrate the effects of jitter on signal integrity. Assuming that jitter will contribute 10 ps of variability to the timing on this serdes bus and that jitter does not decrease as speed increases, then jitter on generation 1 of the bus, which has a transfer rate of six GT/s and an eye width of 166 ps, will take up just six percent of the signal eye width (10 / 166 = 6%). On generation 2 with a smaller eye width of only 125 ps, jitter will comprise eight percent of the eye (10 / 125 = 8%). And lastly, on generation 3, with a transfer rate of 10 GT/s and an eye width of only 100 ps, jitter becomes an even more significant portion of the eye width at 10 percent (10 / 100 = 10%).

This demonstrates that as the speed of a serdes or memory bus increases, jitter takes on greater and greater importance relative to the signal integrity on the bus. All the more reason to validate signal integrity throughout a system's life cycle.

Voltage attenuation – another factor that has detrimental effects on signal integrity – is also exacerbated by increasing bus speeds. Attenuation, which measures the drop in a signal's voltage between the transmitter and the receiver, occurs as a result of characteristics of the circuit board where the bus has been implemented. Three major component parts of attenuation are insertion loss, skin effects and via stubs.

Insertion loss is a measure of how much of the signal is lost per inch of trace on the board. At higher frequencies, insertion loss increases. Skin effect refers to the effects of electrical resistance, which also increases at higher signal frequencies. A via stub is the unused and unterminated portion of a via (a physical connection between layers of a circuit board). A via stub acts as an antenna for the signals traveling on the trace where the stub is located. As an antenna,



it will dissipate a portion of the signal before it reaches the receiver and can also receive noise and disruptions from other sources. This antenna effect caused by a via stub increases as the frequency of the signal increases. The cumulative effect of these and others factors on voltage attenuation is to decrease the height of the signal as it travels across a circuit board between the transmitting and receiving devices. The signal near the transmitter would have a much taller eye than the same signal near the receiver. As the height of the signal decreases, the eye "closes," making it harder and harder for the receiver to decipher the original signal.

Another major factor affecting the signal integrity of high-speed serdes and memory buses is known as PVT, referring to the process, voltage, and temperature variances that arise during the manufacturing of chips, circuit boards and systems. All of these factors have an effect on the operating margins on high-speed buses. For example, different instances of the same device can be faster or slower as a result of manufacturing variations. Or the power supply voltage can be higher on some boards than others and this will affect how the chips on these boards perform. In addition, many parameters will vary with temperature, causing further variations in signal integrity across a serdes bus. All of these characteristics can and will change over time, disrupting signal integrity and necessitating signal integrity validation.

What to do about it

Because signal integrity can be disrupted in so many ways, validating the quality of signal integrity during every phase of a system's life cycle has taken on greater importance. During design and development, the transfer rates and margins on serdes and memory buses on prototype circuit boards must be validated before the board can be released to manufacturing. Failing to identify and correct signal integrity problems before volume manufacturing begins will cause huge problems and incur much greater costs to fix later in the board's life cycle. During manufacturing, process variances in the silicon, packaging and board materials, as well as process variations in manufacturing the board itself can significantly alter signal integrity. Validating signal integrity at this phase will reduce user dissatisfaction and excessive product returns. Once installed in the field, troubleshooting a poorly performing system should include



validating signal integrity on the system's high-speed buses to determine whether they are within valid margins. This can identify some of the root causes of unacceptable performance.

Historically, a variety of intrusive, probe-based test and measurement (T&M) devices and various kinds of testers have provided engineers the eye diagrams and other measurements they have needed to assess signal integrity. Unfortunately, this legacy test equipment is limited because of its reliance on placing a physical probe on a device or a test pad on a circuit board to extract test data. As signal frequencies increase on high-speed buses, placing a probe on the bus to monitor signal integrity becomes decreasingly feasible. Many designs prohibit the placement of test pads on high-speed buses because of the capacitive coupling effects that pads and probes have on the sensitive signaling. Pads and probes will introduce anomalies into the signaling at a time when operating margins on these buses are shrinking dramatically. As a result, today's best practices in board design will typically prohibit test pads. Without access to the bus, intrusive probe-based testers like oscilloscopes or logic analyzers cannot monitor signal integrity.

Fortunately, non-intrusive software-based tools based on instrumentation embedded in chips – embedded instruments – rather than physical probes can be used in place of intrusive testers. Instead of extracting signaling data from the outside which will include the anomalies introduced by the probe and test pad, non-intrusive embedded instruments deliver only the signal data that is seen by the receivers. In other words, soft access is provided to the hard data validation engineers need by way of these embedded instruments and the tools that support them.

Signal integrity validation from the inside out

For several years many of the semiconductor providers have embedded instruments into their high-performance or complex chips for the purposes of chip-level characterization, verification and test. Recently, companies such as ASSET InterTech have re-applied these embedded instruments beyond their original rather limited application in chip test. Embedded instrumentation is able to provide a more cost-effective and straightforward way to gather empirical signal integrity data by monitoring and reporting data as it is received at the receiver. Now, various validation, test and debug instruments are being embedded into chips and re-used throughout the life cycle. These embedded instruments are typically accessed by way of standard



technologies, such as the IEEE 1149.1 boundary-scan (JTAG) Test Access Port (TAP) which is present on practically all processors and many other types of digital devices. JTAG provides access to an external software-based platform that can manage the embedded instruments in the system as well as compile and analyze the test and measurement data they gather. Figure 7 shows an eye diagram generated by ASSET's ScanWorks[®] platform for embedded instruments.



Figure 7: An eye diagram generated by ASSET InterTech's ScanWorks® from data gathered by instruments embedded in chips.

Quite often the instrumentation embedded by a chip vendor is that vendor's proprietary intellectual property (IP), such as Intel® Corporation's Interconnect Built-In Self-Test (IBIST). Intel is embedding IBIST into its advanced microprocessors and chipsets, and in some cases licensing the technology to other chip companies. To date, ScanWorks is the only third-party tools platform that can make use of IBIST to validate signal integrity on high-speed serdes and memory buses such as Intel QuickPath Interconnect, DDR3, Direct Media Interface (DMI), PCI Express and PCI Express Graphics (PEG).



Conclusions

The validation of signal integrity on high-speed buses has never been as critical as it is today. Without the high transfer rates of contemporary serial buses, system throughput could quickly degrade to unacceptable levels. The situation is exacerbated by the fact that higher frequency signaling on each successive generation of high-speed bus technology makes the bus much more sensitive and susceptible to distortion by reducing the operating margins for the signaling. The margin for error on today's high-speed buses is far less than it was on previous generations of the same bus and the challenge continues to increase with each new generation of technology.

At the same time, the speed, complexity and sensitivity of these high-speed buses is rendering legacy test methods practically obsolete because they are based on intrusive probe-based measurements which distort signals instead of providing an accurate snapshot of signal integrity. Fortunately, instruments embedded in chips are able to gather validation data at the receiver so that engineers can have an unadulterated and empirical view of signal integrity.

Learn More

Learn more about Intel® IBIST. Register for our technical paper, "Platform Validation Using Intel® IBIST," and find out how you can test multiple links of a high-speed serial bus in a single system. The paper demonstrates the benefits of a validation method based on embedded instruments over traditional methods using oscilloscopes.



