SYSTEM MARGINALITY VALIDATION OF DDR3 | DDR4 MEMORY AND SERIAL I/O

BY AL CROUCH, ADAM LEY AND

ALAN SGUIGNA



By Adam Ley – Chief Technologist, Non-intrusive Board Test and JTAG



Adam serves customers by ensuring that ASSET's non-intrusive board test (NBT) methodologies comprise a best-in-class solution to meet the evolving need for improved coverage of board test in the face of ongoing erosion of physical access. Adam is an active participant in IEEE 1149.1, having previously served terms as working group vice chair

and as standard technical editor (2001 revision), as well as in nearly all related standards, including: 1149.4, 1149.5, 1149.6, 1149.7, 1149.8.1, 1500, 1532, 1581, P1149.1.1P1149.10, iNEMI boundary-scan adoption, PICMG MicroTCA, and SJTAG (system JTAG). Adam's prior experience spanned over a decade at TI where he had roles in application support for TI's boundary-scan logic products and for test and characterization of new logic families.

Alan Sguigna – Vice President of Sales & Customer Service

Alan has more than 20 years of experience in senior-level general management, marketing, engineering, sales, manufacturing, finance and customer service positions. Before joining ASSET, he worked in the telecom industry. He has had profit and loss responsibility for a \$150 million division of Spirent Communications, a supplier of test products and services. Prior to his tenure with Spirent, Mr. Sguigna also served in business development positions with Nortel Networks, overseeing the growth of its voice over Internet protocol (VoIP) products.

Al Crouch – Chief Technologist, Embedded Instrumentation Methodologies and IJTAG



Al investigates the use of embedded instruments for IC test and debug, board test and debug, and software debug. He is a Senior Member of the IEEE and serves as the vice chairman of the IEEE P1687 IJTAG working group that is developing this standard for embedded instruments. He has contributed significantly to its hardware architecture

definition. Al is also a member of the P1838 Working Group on 3D test and debug, and co-chair of the iNEMI BIST group, which is defining the use of embedded instruments for board test. Al's previous experience includes design-for-test and debug at various semiconductor companies, including TI, DEC and Motorola, as well as chief scientist at startup companies DAFCA and INOVYS.

Platform for Embedded Instruments

2

Table of Contents

Introduction	6
Definitions and Background	7
The Basics – 'SMV 101'	9
Oscilloscopes versus Embedded Instruments	10
Signal Integrity Validation (SIV) versus System Marginality Validation (SMV) - Part 1	13
SIV versus SMV – Part 2	15
SIV versus SMV – Part 3	17
Will BIST Kill T&M?	19
Memory Validation	19
Single-Bit Memory Errors	20
Signal Integrity Margins of Different DIMM Suppliers	21
DDR4, Signal Integrity and Power Integrity	23
Defects on High-Speed Memory – Part 1	24
Defects on High-Speed Memory – Part 2	29
Serial I/O Validation	30
Structural Defects on Intel® QuickPath Interconnect (QPI) – Part 1	30
Structural Defects on Intel® QuickPath Interconnect (QPI) – Part 2	32
Defects on High-Speed Serial I/O – Part 1	34
Defects on High-Speed Serial I/O – Part 2	36
Defects on High-Speed Serial I/O – Part 3	38
Testing SATA 3	39
Equalization	41
Adaptive Equalization – Part 1	41
Adaptive Equalization – Part 2	42
Adaptive Equalization and Power Consumption	43
Silicon Issues	44
The Intel Cougar Point SATA Bug	45
Margins (Eye Diagrams) Follow the Silicon – Part 1	46



Margins (Eye Diagrams) Follow the Silicon – Part 2	
Silicon Aging and Signal Integrity	49
Miscellaneous Technical Topics	
PRBS31 and Validation of High-Speed Serdes	
Signal Integrity Testing with CRCs versus Pattern Generation and Capture	52
The Statistical Basis of SMV	56
Conclusion	60
Learn More	60

Figures

Figure 1: Oscilloscope eye diagrams	8
Figure 2: I/O Built-In Self Test (BIST) eye diagrams	9
Figure 3: The World's Most Expensive Oscilloscope Unboxing!	
Figure 4: SMV plot of SATA HDD	
Figure 5: SMV plot of SATA SSD	
Figure 6: Derivation of the SMV Eye Mask	
Figure 7: Voltage margining plot of DIMM vendor 'X'	
Figure 8: Voltage margining plot of DIMM vendor 'Y'	
Figure 9: Simplified block diagram of DIMM pinout	
Figure 10: Stuck-at fault on DDR DQS pin	
Figure 11: Two DQ shorted together	
Figure 12: Diminishing test probe access	
Figure 13: PCB manufacturing defects and variances	
Figure 14: Memory DQ short	
Figure 15: Layout and 3-D X-ray of QPI land and via	
Figure 16: Optical inspection of QPI land and via	
Figure 17: Graphic representation of QPI land, via and socket	
Figure 18: Short on two QPI nets	
Figure 19: SMV plots for QPI lanes	
Figure 20: AC-coupled differential pair	
Figure 21: Missing capacitor on differential pair	



Figure 22:	Stuck-at fault on differential pair	36
Figure 23:	Short between Tx1- to Tx2+ on differential pair	37
Figure 24:	SATA margining eye diagram	40
Figure 25:	Arrowhead-shaped SATA margin	40
Figure 26:	Balanced fixed/adaptive equalization	42
Figure 27:	Over-equalized I/O	43
Figure 28:	Under-equalized I/O	43
Figure 29:	Adaptive equalization working too hard	44
Figure 30:	Silicon into packages and onto boards	46
Figure 31:	Silicon/package/board routing	47
Figure 32:	Wafer defects and variances	48
Figure 33:	Non-uniform routes within devices	48
Figure 34:	Electromigration induced defects	50
Figure 35:	PCI Express Transaction Layer Packet (TLP)	53
Figure 36:	SATA 3 margin plot	56
Figure 37:	Deriving the SMV eye mask	57
Figure 38:	The Normal Distribution	58
Figure 39:	A 5x1 (5 tests on 1 system) margin run	59

Tables

. 3	3	8	
		3	38

© 2014 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.



Introduction

For today's high-speed printed circuit board designs, designing sufficient operating margins into the system so that it will operate optimally under nominal conditions has become extremely challenging. A large contributor to the narrowing of system margins is the signal integrity associated with high-speed serial I/O and memory buses. And shrinking chip and board geometries, higher speeds and greater chip and board densities have made it almost impossible for engineers armed with legacy measurement technologies such as oscilloscopes to accurately measure the overall signal integrity on serial I/O and memory buses, and then correlate this with the behavior of the system.

A system that has insufficient margin to continue to operate optimally through the corner cases of process, voltage, temperature, frequency and other variables is said to be marginal – that is, it may or may not function at its peak performance when it is operating at the edge of its defined specifications. In certain cases, the system may seem to behave normally, but even then it will be more prone to substandard performance, intermittent failures, crashes and hangs. This results in warranty returns, diminished customer satisfaction and deterioration of the manufacturer's brand.

A new approach to design validation known as system marginality validation (SMV) has emerged and is compensating for the deficiencies of legacy methods such as signal integrity validation (SIV) with oscilloscopes. Unlike SIV, that only provides a limited (and very expensive) snapshot of the signaling on a few lanes of a few buses in the system, SMV determines the marginality of the entire system while taking into account silicon and circuit board process variances, voltage, temperature, humidity and other effects. SMV uses statistical modeling methods and measurements extracted by software tools from instrumentation embedded in silicon. Finally, SMV provides a truly holistic look at the margins of a system and not just a portion of the system in isolation from its nominal operating conditions.

Fortunately, a large body of knowledge on system marginality describes what it is, how it is measured and how it can be improved on a particular design. This eBook is a compendium of various blogs and articles published by ASSET InterTech that delve into SMV to explain how it can determine the level of confidence engineers will have in a design. Using the right tools and



methodologies such as SMV, engineers are better able to design and deliver robust systems that function within specifications over their lifespan.

Definitions and Background

mar∙gin [mahr-jin]

noun

- 1. the space around the printed or written matter on a page.
- 2. an amount allowed or available beyond what is actually necessary: to allow a margin for error.
- 3. a limit in condition, capacity, etc., beyond or below which something ceases to exist, be desirable, or be possible: the margin of endurance; the margin of sanity.
- 4. a border or edge.

The dictionary definition of margin above refers to the outer limits of some condition beyond which a system will either perform sub-optimally or even fail. In the more general technical sense, margin represents the gap between a system's viable state and its failure state. For example, a system may be warranted to operate up to a maximum operating temperature of 50°C. If it's currently running at 30°C, the system's operating margin for temperature is 20°C between where it is now and its failing point is.

Narrowing the definition somewhat and applying it to high-speed memory and serial I/O buses yields some interesting insights. Firstly, the margin of a bus is greatly determined by the size and shape of the "eye" plot (voltage versus time) of the signaling on the bus. In the traditional oscilloscope world, an eye diagram looks like this:



System Marginality Validation of DDR3|DDR4 Memory and Serial I/O









Note that the signal waveforms in the Good Eye photo above are in fact distinct and well defined. This allows easy recognition of the signal at the receiving silicon and fairly flawless transmission.

The Bad Eye photo, on the other hand, shows signaling that is essentially unrecognizable because the individual signals do not seem to conform to any expected and discernible pattern. The incoming bitstream, subject as it is to jitter, intersymbol interference, attenuation, and other factors, will likely suffer a high bit error rate at the receiver.

Taking such measurements on an oscilloscope is very costly, time-consuming and error-prone. Plus the added load resulting from physically probing a circuit board may alter signals on some sensitive high-speed traces. Fortunately, a newer approach, which measures the incoming bits at the silicon itself and then overlays time and voltage margins, gives a much better representation of the physics of the bus' operation. An equivalent graphical representation is depicted below:



System Marginality Validation of DDR3|DDR4 Memory and Serial I/O







Figure 2: I/O Built-In Self Test (BIST) eye diagrams

Oscilloscope measurements are part of a signal integrity validation (SIV) test suite. I/O BIST (Input/Output Built-In Self Test - that is, measurements derived from instruments embedded in silicon and processed by software tools) is part of a system marginality validation (SMV) methodology.

Although in some respects the two approaches are similar (for example, the larger the eye pattern, the better the signal), the technology supporting the two validation approaches is wildly different. SIV uses mathematical modeling to de-embed or extract the data signal from any other electrical interference caused by the fixture probe, cables, connectors and other sources. SMV measures and reports the signal as it is received at the silicon receiver internal to the chip.

The remainder of this eBook reviews much of the technology that is involved with SMV.

The Basics - 'SMV 101'

System marginality validation (SMV) is a new and innovative way of not only validating the signal integrity of a circuit board design, but also of assessing how much margin is in the design relative to chip/board characteristics and processes that vary over time, voltage, temperature, frequency, humidity, component aging and numerous other factors.



Given its flexibility and power, SMV has become the premier method of not only front-end design validation, but also a foundation for manufacturing test and failure analysis. With SMV, system margins can be measured at any point in a product's lifecycle.

This section covers the basics of SMV and how it differs from traditional signal integrity validation (SIV).

Oscilloscopes versus Embedded Instruments

What's cheaper, faster and more powerful than an oscilloscope when it comes to validating highspeed signal integrity? Of course, the answer is a software application using embedded instruments. How is this possible?

Now, software applications that use I/O built-in self test (I/O BIST)-based embedded instruments within silicon to perform various tasks, including serdes and memory signal integrity (SI) validation, are widely available. These sorts of software tools are able to observe and report directly on the signal at the receiver in the silicon, as opposed to viewing what seems to be a closed eye on the board interconnect and reconstructing it using higher mathematics, as oscilloscopes do. As is commonly said, nowadays "the math is in the chip" when it comes to the emphasis and equalization schemes (some of which are adaptive and change on the fly!), so the best place to observe the waveform is from within the silicon.

Let's take a look at three attributes of embedded instrumentation versus external oscilloscopes and you'll see why the software-based embedded instruments provide the most effective methodology.

1. Embedded Instruments are more powerful

Because of the extensive time required to gather data and their extremely high procurement costs, oscilloscopes typically are limited to capturing only the signal waveforms on a couple of lanes on a bus at a time. So, engineers usually choose to measure only the longest and the shortest lanes, hoping that these two will have the worst signal integrity and, therefore, the worst margins. But this is often a pipe dream (please excuse the pun) since it assumes that all other buses in the design will have better signal integrity than these. In fact, some other arbitrary lane



may have a defective capacitor or pass near a noisy power delivery pin field. So there's a lot of risk associated with such a small sampling.

As well, oscilloscopes measure signal integrity under artificially ideal conditions: that is, nominal process/voltage/temperature (PVT) characteristics and normal operating traffic on the link. If the board or its silicon is operating outside of the nominal and normal conditions, a scope will likely not detect it in the lab (Scopes don't effectively correlate all of the system environment settings while measuring the trace under investigation.), but it probably will be detected later by the user when performance degrades or the system fails in some way.

Embedded instrumentation tools, on the other hand, have none of these restrictions. All bus lanes can be saturated with "synthetic" traffic, using pseudo-random bit sequence (PRBS) worst-case patterns. Then, software tools are able to detect the effects due to crosstalk and inter-symbol interference (ISI) that are missed by scopes. And the silicon supplier can provide eye masks that take PVT into account. (Some cutting edge SoCs and microprocessors also include embedded instruments to measure internal temperature, voltage levels, operating clock frequencies, and other internal environmental variables.)

2. Embedded Instruments are faster

With oscilloscopes, the engineer must have direct physical access to the signal traces on the prototype or manufactured circuit board. Setting this up can be a very lengthy process, adding several weeks to a design cycle because engineers must select the number of boards and lanes to test, design in test access on the board, solder on the probe heads and then finally perform the design validation.

With embedded instruments, it's just plug-and-go. Access to the JTAG header on the board is typically all that is required. Shaving a few weeks from the design cycle can make a big difference when it comes to time-to-market.

3. Embedded Instruments are cheaper

What's the price of a good high-end scope these days, to test signal integrity on PCIe Gen 3, QPI, SATA 3, etc? Let's put it at \$200,000 - \$300,000 USD. And don't forget those expensive



amplifiers and probe heads: let's add another \$25,000. And don't forget the additional channels you want to test; the price goes up again.

So let's say hypothetically that you want to validate signal integrity (SI) on five (5) boards, because this will give you a high confidence level that variances in the chips and board manufacturing processes aren't going to cause problems for users when the design moves into volume manufacturing. A sample size at least this large is an excellent idea (an insightful white paper on this topic is "<u>Platform Validation Using Intel IBIST</u>" by Stephanie Akimoff. This paper demonstrates that signal integrity follows the silicon as well as the board). Let's further assume that a single test run on a single board takes approximately 100 hours; that you want to test each board a handful of times to eliminate any procedural variation in the test process; and that you want to be able to react to any silicon version changes throughout your prototype runs. To keep things simple, let's ignore the need to test for PVT effects. We'll just keep our fingers crossed that our nominal tests give us enough margin and a low enough bit error rate (BER) to accommodate any conditions our system will encounter in the field!

So if we do the math on this test process it would require a minimum of approximately 100 test days (5 X 5 X 100 hours = 2,500 test-hours or about 100 test days). It would be impossible to do this with a single scope and also perform all the testing needed between the prototype runs as well as the testing required leading up to volume manufacturing and deployment in the field. You're going to have to buy more scopes at \$200,000 each at a minimum. That gets pretty expensive, pretty fast.

Conversely, the cost of an embedded instrument is virtually free, because the cost of a transistor on an IC has decreased with Moore's Law and the embedded instrument has probably already been included in the silicon by the chip design team for cost-effective IC tests. What's needed is the ability to access the embedded instruments during board development, board test, board characterization and board debug, as well as tools to perform these test functions.

Of course, software-based tools that take advantage of embedded instrumentation are available at a tenth of the cost of the scopes you'll need. You're able to deploy a handful of software-based tools, save money, get more testing done and deliver your product to market faster.



The Wrap-Up

Of course, scopes will never go away completely. They are useful for functions like compliance testing and can provide a good spot-check against embedded instrumentation-based solutions in design validation applications. You just don't need to buy so many of them.

Signal Integrity Validation (SIV) versus System Marginality Validation (SMV) – Part 1

A few years ago, engineers used expensive high-end oscilloscopes to perform signal integrity validation (SIV) on their designs, and considered that an adequate indicator of the design's potential for success over its product lifecycle. But with today's products, process and parameter variations occur that require system marginality validation (SMV) to be done by less expensive software-based tools to determine if a design is ready for high volume production.

First, some definitions: SIV uses oscilloscope measurements of select voltage and timing parameters across certain process/voltage/temperature (PVT) corner conditions and interface configurations. The intent of SIV is to ascertain transmitter and link robustness via oscilloscope waveform captures.

SMV, on the other hand, uses embedded instrumentation within silicon to perform system level margining of I/O buffer control knobs such as VREF, buffer strength, slew rate and timing controls with PVT and interface configuration variability "baked into" an eye mask that defines the acceptable operating margin for the design. SMV validates the robustness, including transmitter, receiver and interconnect link, of the entire system.

The problem with SIV is that it is expensive, difficult and slow. If you want a chuckle, watch this <u>YouTube video</u> on the world's most expensive unboxing of a \$140,000 Agilent DSA91304A 13GHz oscilloscope. And if you're feeling really rich, you could invest EUR 228,062 (a whopping \$300,000 USD by today's exchange rate) into an Agilent DSOX93204A Infiniium 33GHz oscilloscope, which you'll need to do SIV on fast buses such as PCI Express Gen3.





\$140,000 Agilent 90000 13GHz Oscilloscope Unboxing

Figure 3: The World's Most Expensive Oscilloscope Unboxing!

And, of course, probe heads and amplifiers aren't cheap either. It takes weeks to set up access to and attach the heads – and you can only do that on early prototypes, because you don't want too many vias on production boards since they would mess up your signal integrity even further. Finally, an oscilloscope only has a limited number of channels, which means you would probably only measure the shortest and longest traces, because you assume these are the buses most likely to have poor signal integrity. But you really won't know for sure if, for example, on one of the other lanes that was not measured (like the one that, unbeknownst to you, is routed near some voltage regulator noise) actually has the worst margins. It could very well turn out later that the lane you neglected to measure will be responsible for intermittent system crashes when the product ends up in the hands of users.

So SMV came along to address these deficiencies. With software-based tools that see what the silicon sees, you can determine the operating margins for the entire system – Tx, Rx and interconnect. There's no access to worry about because embedded instruments provide data from the inside out, not from the outside in. Plus, there's no soldering to do. You can run SMV multiple times, across multiple systems, with multiple third-party add-in cards and DIMMs to get a high level of confidence in your designs across a wide spectrum of possible defects and variances that can raise their ugly heads once the system is in the field. This is important because margins can vary statistically and substantially from one production run to another due to all



sorts of variances. (Several good white papers on the topic of variances are: "<u>Margins (Eye</u> <u>Diagrams) Follow the Silicon</u>" and "<u>How to avoid poor serdes performance caused by circuit</u> <u>board manufacturing variances</u>".) And you can do all of this at a tiny fraction of the cost of just one oscilloscope.

Oscilloscopes are still needed during the early stages of product development for compliance testing and for limited SIV measurements. But SMV is a more optimal method to determine how much margin is in the design. Operating margins are simply a better predictor of the product's performance throughout its lifecycle than a limited snapshot of the system's SIV.

SIV versus SMV - Part 2

Aside from the price, what are the other advantages of embedded instrumentation-based system marginality validation tools?

Oscilloscopes are used to perform SIV while software-based tools like ScanWorks perform SMV. As high-speed buses have gotten faster and faster over time, the cost of oscilloscopes has risen exponentially, due in part to the hardware technology within them. A 33GHz Agilent DSOX93204A Infiniium oscilloscope costs \$300,000. And the faster scopes, such as the 63GHz Agilent DSAX96204Q, cost EUR 359,980 (\$470,000 by today's exchange rate). Once you add in the cost of the software options, probe heads, amplifiers and other paraphernalia, the price goes well over half a million dollars.

Apart from the capital cost avoidance and labor reduction associated with easy-to-use softwarebased embedded instrumentation tools vs. oscilloscopes, the software tools also provide functionality that simply cannot be replicated by the legacy heavy iron scopes. One example of this is automated shmooing of BIOS programmable equalization settings. For SATA 3, for example, the Discrete Time Linear Equalization (DTLE) of the receiver must be adjusted to take into account the different margins associated with different vendors' hard drives, solid state drives, different cables, whether they are hooked up directly to the motherboard versus through a 6-ft. cable, etc. A simple demonstration of this can be seen with a notebook design that has a default DTLE setting of 2 within the BIOS. This setting could possibly favor a hard disk drive on



the bus because more hard drives are configured in such systems than any other peripheral. The ScanWorks margin graph of the configuration with a hard drive is as follows:



Figure 4: SMV plot of SATA HDD

Switching the hard disk drive (HDD) to a solid state drive (SSD), unfortunately, yields a margin result that is nowhere near as favorable:



Figure 5: SMV plot of SATA SSD

The problem here is that the default DTLE setting favors the hard drive, which gives off quite a bit of EMI compared to the flash drive. The DTLE would need to be adjusted to tune the design



to accommodate both an HDD and a SSD and give acceptable margins for both. Otherwise, the solid state drive will not perform at its potential; it will run slower than its optimal performance and may even be subject to intermittent link failures (i.e. the operating system loses the drive).

Tuning the DTLE with an oscilloscope is an entirely manual process that could take weeks of laborious effort. With a software system that takes advantage of embedded instrumentation, DTLE settings can be automatically swept through a range while simultaneously taking voltage and timing margins to yield the optimal value for all sorts of SATA devices, such as drives from different vendors, drives with different operating speeds and other characteristics. This level of validation gives confidence across multiple possible customer configurations. Unfortunately, efficiently finding the right value is simply not possible with an oscilloscope.

SIV versus SMV - Part 3

The eye mask produced by SMV and that defines the design's acceptable operating margins takes into consideration variances in chip design and manufacturing, which can cause the devices on a circuit board to degrade over time. This can be shown visually:



Figure 6: Derivation of the SMV Eye Mask

High Volume Manufacturing (HVM) in the figure above refers to the fact that the performance of the same device can vary because of variances in the device's manufacturing processes. These



System Marginality Validation of DDR3|DDR4 Memory and Serial I/O

manufacturing process variances can cause variations in the internal signal integrity from one particular device to the next. The effect of these variations is typically a Gaussian distribution.

In Figure 6, Small Sample Size refers to the fact that SMV is based on measuring a finite number of circuit boards, which determines the size of the sample. For example, as part of a 5x5 methodology (five SMV tests done across five different boards), five different chips are used as part of the validation testing. To have a high degree of confidence and to reduce the effects of variances within these five chips on the calculations, the size of the sample must be statistically valid. In other words, a single measurement is insufficient. The larger the sample size, the higher the confidence in the resulting statistics. In recognition of this, Altera has created an extensive database of customer backplanes to calibrate its eye mask and refine its on-chip equalization schemes.

BER Adjustment in Figure 6 takes into account the fact that the bit error rate (BER) is measured at each margining point (defined as each voltage and timing step) to detect faults. Because the error threshold for many serial buses is a very big number, empirically gathering this data with a scope would require considerable time. For example, the acceptable bit error rate threshold for Intel QuickPath Interconnect is 1 bit in 10E14. An extremely long period of time (hours or days) would be needed at each margining point on a board design to detect errors. As a result, scopes apply a BER Adjustment factor that extrapolates the error rate based on a very short measurement, typically gathering data for less than a second.

The Other category in the figure represents adjustments to the eye mask made for silicon aging, humidity, pollution, process/voltage/temperature (PVT), and other factors.

SMV has emerged as a viable means of measuring system operating margins in real world environments. As a methodology, SMV has been adopted widely within the semiconductor community by companies such as Intel, PLX, Altera, Xilinx, Broadcom, TI and others. At these companies, SMV is the preferred means of determining how much margin is in a chip design.

For more information on the nature of chip and board defects and variances that affect system margins see the tutorial, "Detection and Diagnosis of Printed Circuit Board Defects and Variances using on-chip embedded instrumentation".



Will BIST Kill T&M?

To quote Ransom Stephens in the DesignCon Community Blog, "BIST (Built-In System Test), is an acronym that would keep executives at test and measurement companies awake at night, if they knew what it meant." What's he talking about?

In his blog, Ransom describes the use of BIST for accessing the "real eye diagram" (and thus the true margin of the bus) as seen within the chip. You want to see what the silicon sees, because an external probe introduces ISI (inter-symbol interference), multi-path interference and multiple internal reflections due to the frequency response of connectors, traces, cables, backplanes and all other introduced points of discontinuity. Sure, you can try to de-embed all that noise mathematically to obtain the true signal, but, in his words, "Suspicious? I hope so… In principle, *it's possible to de-embed all the way into the chip if the S-parameters are known, but in practice, you'd need superconducting test equipment.*"

Further, Ransom describes an experiment conducted by Eric Kvamme of LSI Corp., where BIST was used to display bathtub plots of a 25 Gbit/s part down to a bit error ratio of 1E-15. This experiment could not have been performed without BIST. In other words, forget about using an oscilloscope.

For many people in the industry, BIST is a more generic term to refer to embedded instruments (that is, instruments embedded in chips). These embedded instruments can be used for board test, as well as design validation and platform debug. A good example of an embedded instrument for the purpose of board test is IEEE 1149.1, also known as (aka) boundary scan. An embedded instrument for the purposes of design validation is Intel® IBIST. And an embedded instrument for platform debug is run-control (also referred to as debug-port control), which is used by such tools as the SourcePoint[™] and Arium probes.

The traditional T&M companies have good reasons for sleepless nights.

Memory Validation

Memories keep getting faster in data rates and clock frequencies, and lower in voltage. In fact, DDR4 will be the last of the wide, single-ended data buses; the laws of physics and signal



integrity have run their course for this technology. It gets difficult to put more than one DIMM on a channel at speeds in excess of 2,400 Mbps. New technologies such as Hybrid Memory Cube and High Bandwidth Memory will supplant it.

Already, SMV is showing that different memory suppliers have varying margins on their DIMMs and in their silicon. More sophisticated pattern generation techniques, such as row hammering, target silicon defects and variances where excessive ACTIVATE commands induce bit errors in adjacent memory rows. And with memory moving into the device stacks, as with embedded DRAM (eDRAM), more focus going forward will be on the role that chips play in overall system margins, because validating adequate margins on a design is an effective way to avoid product recalls and returns later. In the meantime, board designers need to take extra care to provide enough margining headroom on DDR.

Single-Bit Memory Errors

Ever wonder if a stray cosmic ray or alpha particle might double your bank account by causing an undetected RAM error?

In fact, lower-end systems without Error Correction Code (ECC) will crash (assuming you're lucky) when a single-bit memory error is encountered. These soft faults are often the results of high-energy neutrons from cosmic rays or alpha particles that can decay the isotopes within the silicon packaging or surrounding materials. Alternatively, poor signal integrity on a bus in terms of its susceptibility to crosstalk, for example, will affect the system's soft error rate.

Systems with ECC will transparently correct single-bit errors and log the results, resulting in a small performance hit. Putting ECC memory in your desktop or laptop is generally a good idea if you're involved in financial or scientific applications (as opposed to just doing Facebook). But conventional ECC memory can't handle double-bit or more catastrophic memory errors. For higher-end systems that demand high-reliability and availability, more sophisticated ECC that can detect and correct multi-bit errors within a single memory device are needed. Schemes like Chipkill, Extended ECC, Chipspare and SDDC scatter the bits of the ECC code across multiple chips.



System Marginality Validation of DDR3|DDR4 Memory and Serial I/O

The goal, of course, is to reduce the overall incidence of soft errors and to improve the performance and robustness of the system. Although we can't block cosmic rays and high-energy neutrons, the effects of alpha particle interaction can be mitigated by using purer materials. Of course, this comes with increased costs. In addition, signal integrity is the most important aspect of system performance to examine to determine whether the design has plenty of operating margin. We've seen how temperature, jitter, noise, voltage aberrations and manufacturing variances can all affect signal integrity. A good memory test program is essential for stable and reliable memory tuning.

Some subtle memory failures manifest themselves only with certain data patterns. Others might only flare up when certain addresses are accessed. These types of errors depend not only on the data being written or read, but also on the data in the surrounding bytes that are being transferred at the same time. A simple memory test -- one that writes a fixed pattern to all bytes -- will not discover this type of failure. Similarly, failures might only show up when accessing non-adjacent addresses, so memory test programs that perform accesses in a non-sequential pattern should be employed.

A good read on the general topic of memory testing can be found in this white paper: "<u>Testing</u> <u>High-Speed Memory with Embedded Instruments</u>".

Signal Integrity Margins of Different DIMM Suppliers

Not all memories are created equal. Some DIMM suppliers' cards have margins that are better than others. And, of course, the better the margins, the better the performance of the system and the fewer blue screen crashes.

We all know that signal integrity is critical for optimal system performance. High-speed signals on serial I/O such as PCI Express (PCIe) and on memory buses like DDR3 should operate with very low bit error rates. What this means is that under normal operating conditions, the margin (or shape and size of the eye diagram of the signaling on the bus) should be large enough so that there is a very low probability of a bit flip. Some single or multi-bit errors can be corrected, but the uncorrectable errors could degrade performance or cause a system crash or hang.



Notebook DIMMs are made by many companies, including Corsair, Kingston, Patriot, Samsung, Crucial and others. Measuring the design margins of different DIMM suppliers to determine which modules perform better and which are more reliable is fairly straightforward. The ScanWorks HSIO tool was connected to a garden-variety notebook board and ran a 1-D voltage margining test on two modules from different vendors. The results for Vendor X and Vendor Y are below:











System Marginality Validation of DDR3|DDR4 Memory and Serial I/O

Green indicates passing lanes and red indicates failing lanes. Vendor X has excellent margins; on the positive voltage side all lanes are passing at the maximum allowed range and on the negative voltage side the margins are comfortable at the -40 voltage ticks (well past the defined eye mask at which margins are unsatisfactory). Vendor Y, on the other hand, is right at the very edge of the guard band. On the positive voltage side, a couple of lanes are approaching 24 ticks, which is on the "hairy edge" of the acceptable range. There is enough margin on the negative voltage side, but maybe not enough to survive the variances that will occur across different lots of silicon and circuit boards.

So whose memory would you buy, Vendor X or Vendor Y?

These are, of course, just the results of a single test. As mentioned above, variances in the chips or boards may produce unacceptable margins once a design has gone into volume production. An excellent e-Book on the source of manufacturing variances is here: "<u>How to avoid poor serdes</u> performance caused by circuit board manufacturing variances".

DDR4, Signal Integrity and Power Integrity

One of the biggest design challenges today revolves around maintaining signal integrity in the presence of power and ground rail fluctuations due to simultaneously switching signals. This is particularly true for DDR4 memory.

DDR4 is a big step from DDR3; much bigger than DDR3 was over DDR2. The speed of DDR4 has increased up to 3200 Mb/s at the top end. The V_{dd} of DDR3 drops from 1.5V to 1.2V for DDR4. The Unit Interval (UI) shrinks from 469 ps to 313 ps. Channel interconnect skew and jitter on DDR4 easily consume 50% of the 2133 Mbps timing budget. These, combined with other factors, including the effects of DQS jitter, edge roll-off, impedance discontinuities, pin-to-pin capacitance variations, crosstalk and inter-symbol interference (ISI), make designs with DDR4 far more challenging to simulate and measure. One must also take into account variances in the manufacturing of printed circuit boards. These issues are described in this eBook: "How to avoid poor serdes performance caused by circuit board manufacturing variances". Manufacturing variances in silicon are described in this paper: "Platform Validation using Intel® Interconnect Built-In Self Test (Intel® IBIST)".



Most importantly, stability of the power distribution network (PDN) plays a key role in signal integrity and operating margins of the design. The maximum ripple of the PDN is specified as +/- 60mV for DDR4 as opposed to +/- 75mV for DDR3. Simultaneous switching noise (SSN) can have a major effect. For example, in the worst cases, all 64 bits of a data bus could transition simultaneously, causing large instantaneous changes in current across the PDNs. These fluctuations in voltage levels could impact the timing margins on the transitioning signals. Such simultaneously switching outputs (SSO) will have a decided effect on memory and other serial I/O data integrity issues on the board. On memory, this can be mitigated, to some extent, by Data Bus Inversion (DBI).

Nowadays, some board designers use power-aware signal integrity (SI) simulation tools to provide some level of assurance of proper operations. This involves modeling the copper shapes that comprise the power and ground planes, as well as the vias that run through them, along with their couplings to the signal traces. These vias essentially act as radial transmission lines that excite the parallel structures on the plate planes, disturb the power supplied to the chips and couple noise back onto the signals as well.

In addition, decoupling capacitors must also be incorporated into the model and simulation, as should the voltage regulator module (VRM).

Given the complexities of these simulations, design engineers should be cautious and not rely totally on measurements taken with oscilloscopes since their results also rely on simulations. More empirical measures are provided by on-chip embedded instrumentation that reports precisely what is being seen at the device transmit and receive buffers. Further, worst-case measurements are highly recommended. These can be obtained by generating worst-case bit patterns from embedded instruments to generate SSO as well as the maximum amount of jitter, crosstalk and ISI. These issues are described in an eBook: "<u>Bandwidth tests reveal shrinking eye diagrams and signal integrity problems</u>".

Defects on High-Speed Memory - Part 1

The following analysis examines the kinds of defects encountered with the DDR memory bus and the effects these defects could have on system performance and stability.



For the purposes of this discussion let's consider a standard 240-pin DDR3 DIMM. A simplified, high-level block diagram of the pinout is below:



Figure 9: Simplified block diagram of DIMM pinout

The DDR3 memory bus on high-end systems differs from serial I/O buses in many fundamental ways. First, the DDR3 bus is a parallel bus, as opposed to serial. Error detection and correction is via ECC memory at the physical layer, as opposed to high-speed serial buses, which relegate such tasks to the data link layer using cyclic redundancy checks (CRC) and upper portions of the protocol stack. In addition, serial buses usually employ embedded clocking, while DDR3 pairs separate differential, but not AC-coupled, strobes (the DQS signals), which are assigned per nibble or byte of data (the DQ signals). The DDR strobes are not continually running clocks as they are on serial I/O buses, but, rather, they turn on as needed and act as source-synchronous clocks.

To understand the behavior of a defective DIMM, it is important to understand the process whereby the memory bus is first initialized. (This is also referred to as training the bus or training a lane on the bus.) A BIOS or boot loader will run a minimal amount of code when a system is first booted in order to ensure that the memory is basically functional. So, in general, the BIOS will sync up DQ and DQS to optimize the system at the center timing and voltage point, and then it will do a basic test of the DQ at location 0 within each rank. So, if there are gross defects, the



BIOS will either (a) disable the channel, or (b) hang the boot process with a "memory failure" post code.

If the system quietly disables the channel, this may pose a problem to conventional functional memory testers because they may not be aware of the issue. And hanging the boot process is also an issue because, as we all know, when the screen (normal terminal output) goes dark, it takes a certain level of expertise to determine the cause of the failure.

But, of course, there are numerous defect scenarios that have effects far more nefarious. Let's look at some of them.

A Short-Circuit on DQS0+

Since DQS are differential pair, they are immune to common mode noise. The receivers operate by considering the difference in amplitude between the positive and negative nets of the pair. Usually, there may be enough residual signal even with one leg stuck at GND, for example, to ensure that the memory timing requirements are met. This will allow the BIOS to train the memory bus. However, a high number of bit errors will occur. Most of these errors will be automatically detected and corrected by the ECC, although their presence will impact the performance of the memory bus.



Figure 10: Stuck-at fault on DDR DQS pin



Two DQ Shorted Together

At first glance it might seem that this kind of defect would be easily detectable, but further examination reveals that the BIOS memory training process is not reliably pattern-sensitive. In a perfect world, the net signal of a shorted '1' and a '0' is 0.5, but the internal voltage biases of the receivers may, in fact, cause them to miss the defect and read back what was written. So, the memory bus may or may not train, but if it does train correctly, the system will subsequently fail under load.



Figure 11: Two DQ shorted together

Other Process and Random Variances

At lower DDR speeds, test vias were commonly added to memory nets in order to perform incircuit test (ICT) on DIMMs. This practice is now mostly defunct due to signal integrity issues caused by excess metal that resulted in reflections and signal attenuation.



System Marginality Validation of DDR3|DDR4 Memory and Serial I/O



Figure 12: Diminishing test probe access

However, these signal integrity issues can still be manifested through process and random variances in PCB manufacturing. These include, but are not limited to, the following variances:

- Stripline dimensions
- Trace surface finish
- Incompletely plated vias
- Flaws introduced during the imaging process (pin holes, nicks, cuts)
- Plating thickness
- Delamination
- Head-in-pillow



Figure 13: PCB manufacturing defects and variances

For more information on this fascinating subject, see our eBook, "<u>How to avoid poor serdes</u> <u>performance caused by circuit board manufacturing variances</u>".



Defects on High-Speed Memory – Part 2

When a short circuit between two DQ lanes occurs, it can escape detection by the BIOS because of a voltage bias on the memory controller that causes the value that was read out to be the same as that which was read in. For example, if there were a short circuit and DQ0 was written as a '1' while DQ1 was written as a '0', the resultant values stored in the memory cells would be indeterminate, because the short circuit may have yielded a level midway between high and low. Deterministically, it is possible that this might escape the simple testing that is part of the memory training algorithm within the boot loader.



Figure 14: Memory DQ short

In this instance, even if the memory training sequence does complete, the system will soon fail because data will be read in and out of main memory during the remainder of the board boot-up process. On some systems this will result in the infamous "blue screen" system crash, which yields very little diagnostic information. And, of course, many test routines can only run within main system RAM, which would be impossible with a system crash. It is best to try to catch the failure during the BIOS memory training phase.

An implementation of cache-based instrumented memory testing routines can be reviewed in our white paper, "<u>Cache-as-RAM to bring up non-booting boards</u>".



Serial I/O Validation

High-speed serial differential I/O is intended to be self-healing insofar as it is still able to operate in the presence of significant design flaws, defects and variances. Although this provides greater system availability, some circles would refer to it as the boot-at-all-costs mentality. The trade-off here is that system manufacturers who want to optimize performance need to pay special attention to overall serdes port margins because the survivability of differential high-speed I/O can be one of its weaknesses. Systems on the edge of booting, but that eventually do boot, will often exhibit marginal behaviors, such as intermittent crashes, hangs and drop-outs. Systems with slightly greater margins but still close to the edge will run continuously, albeit at a reduced throughput and performance.

In the following sub-sections we correlate the effects that structural defects (such as short circuits or open circuits) have on link margins.

Structural Defects on Intel® QuickPath Interconnect (QPI) - Part 1

A test engineer recently shared some empirical results from boundary-scan testing of Intel QuickPath Interconnect (QPI) nets on a new design. ICT testers offer no coverage on these nets and some short circuit and open circuit defects defy detection by conventional functional test. Here's what boundary-scan testing found.

Intel QPI runs at 9.6 GT/s per lane on Haswell Xeon systems, up from 8 GT/s on Sandy Bridge Xeon. This speed is only expected to increase in the future. At these speeds, signal integrity issues preclude the placement of ICT test pads on the circuit board's nets, rendering ICT unable to provide any test coverage. In addition, QPI employs differential signaling, so the I/O receivers may be able to reconstruct an incoming data stream even in the presence of board-level structural defects. As a result, lanes on a QPI bus with defects may initialize at the physical layer and train up, albeit at a degraded level of overall throughput. What happens next varies, depending on the overall margins of the board and chips. Typically, such systems will exhibit reduced performance, unexpected behaviors, lane drop-outs and even crashes/hangs, many of which may, unfortunately, occur at the user's premises.



One manufacturer recently fired up boundary-scan test for an Intel Xeon-based server platform and immediately saw a 2.9% failure rate. Somewhat perplexed because the boards were booting without problems and functional test had uncovered no failures, root cause analysis was initiated by first performing a 3-D X-ray of the CPU BGA sockets. This is what they saw:



Figure 15: Layout and 3-D X-ray of QPI land and via

The above pictures require some explanation. In the layout picture on the left, the yellow features represent the "dog bone" via and land for node QPI1_RX_4_DP (the land is at the bottom and the via is circled in red). The green feature circled in orange is a land for node GND. The 3-D X-ray picture on the right suggests that there may be a short between the GND land and the adjacent QPI1_RX_4_DP via.

When the processor's BGA socket was removed, a visual inspection of the PCB showed the following:



Figure 16: Optical inspection of QPI land and via



The photo above shows that via for QPI1_RX_4_DP (circled in red) is covered with solder. This makes it susceptible to being shorted against balls at either of the two adjacent lands, which are circled in green. The graphical cross-section (below) of the BGA socket ball, lands and vias for the surrounding area shows what's happening:



Figure 17: Graphic representation of QPI land, via and socket

As stated above, these defects on high-speed serial I/O will often escape detection by conventional functional test processes because the ports may actually train successfully and communicate data, appearing to operate normally. Depending on the bit error rate induced by the defect and the duration of the functional test (bit error count is a function of the bit error rate and the test's duration), even more sophisticated functional test algorithms, that report the contents of QPI error counter registers, may not indicate a potential failure.

For a more detailed treatment of the effects of defects on high-speed serial I/O and memory buses, and how to detect them, a free white paper is available here: "<u>Tutorial: Board Test of DDR3/DDR4 Memory and Serial I/O</u>".

Structural Defects on Intel® QuickPath Interconnect (QPI) - Part 2

What does a user experience when there's a short circuit on an Intel QPI net out of sight under the CPU socket?

As we've pointed out, boundary-scan tests can easily detect structural defects like shorts and opens on high-speed differential I/O buses such as Intel® QuickPath Interconnect (QPI). Of course, boundary scan is the preferred technology for this kind of testing because in-circuit test has no access to buses like QPI, PCIe Gen 3, SATA III and others since test pads on these nets



creates signal integrity issues. But the system manufacturer has to ask what will happen if boundary-scan tests are not being applied to these high-speed nets and, as a result, defects in systems shipped to users escape detection?

In one experiment, we shorted two QPI nets together on a server board:



Figure 18: Short on two QPI nets

Interestingly, the QPI port trained up normally and the system seemed to behave properly. As we know, this kind of defect is often invisible to conventional functional tests since differential I/O is inherently self-healing.

However, we do know that these defects will affect the overall operating margins for the system. To prove it, we ran tests based on embedded instrumentation on the board. The results are summarized below:



Figure 19: SMV plots for QPI lanes

The margins have collapsed the two lanes (lanes 14 and 17) that were shorted. The composite margin on this QPI port is poor, but there is sufficient margin for the system to initialize,



although it will subsequently perform at a less than optimal level. Looking more closely, this link has a very high number of correctable errors. So, several different things can happen, depending on the incidence of bit errors caused by the defect. (Note that because of the size of QPI's flits and its CRC, it has a BER threshold of 10E-14, much more stringent than PCIe Gen3, for example, whose BER threshold is 10E-12.) The various outcomes of the defect could be the following:

- High number of PHY layer re-initializations
- Many CRC errors, with accompanying data link layer re-transmissions
- Data lane failovers
- Intermittent kernel crashes with CATERR thrown

This system is defective. Although it may appear to be operational, it is compromised and will eventually fail in the field.

For more of the theory behind defects on high-speed serial I/O, check out this whitepaper: "Tutorial: Board Test of DDR3/DDR4 Memory and Serial I/O".

Defects on High-Speed Serial I/O – Part 1

Shorts and open circuits on high-speed serdes buses, such as PCIe, may have subtle and difficultto-diagnose effects on system performance. In other words, manufacturers might not know about them until users start complaining and the manufacturer starts receiving warranty returns. What kind of effects are these and how are they prevented?

Let's look at a typical AC-coupled differential bus, such as PCIe. A differential pair, transmitting from one chip to receiving on another, can be illustrated as follows:





Figure 20: AC-coupled differential pair

Now let's review a couple of failure scenarios and see what happens.

An Open-Circuit: Missing Capacitor

Suppose a capacitor was not soldered to a circuit board or it somehow was detached or disabled in the field. This open circuit on one net will not necessarily prevent all signals from getting through to the Rx1- net at the receiver, as shown below:



Figure 21: Missing capacitor on differential pair

Receivers reconstruct the differential signaling on the + and – legs of a pair. In some cases sufficient coupling may be present for a lane to train and operate, albeit at a reduced level of performance. This particular lane will be more susceptible to crosstalk, power distribution noise (PDN), jitter and inter-symbol interference (ISI), so it will likely operate with a higher bit error rate (BER). If the lane's performance crosses certain thresholds, PHY layer re-initializations, data link layer retransmissions and ultimately lane drop-outs (either soft (intermittent) or hard) will result.



A Short-Circuit: Tx1- TO GND

In this example, one of the transmit nets connected to ground has a short. Similar to the previous example, this will impair the propagation of the signal to its intended receiver.



Figure 22: Stuck-at fault on differential pair

But again, a receiver on a differential line operates by considering the *difference* in the signals received. As a result, the data stream may be reconstructed despite the short. Whether the link drops out or continues to operate depends on a large number of factors with the bit error rate ultimately determining whether the link operates or not.

There are many other kinds of failure scenarios, such as shorts between Tx1+ and Tx1-, Tx1+ and Tx2+, Tx1+ and Rx1+, two missing capacitors, etc. These are hard faults, resulting from missing components, excess solder and other common assembly defects. High-speed serial I/O buses are also sensitive to the quality of the interconnects, manifested by such defects as incompletely plated vias, high trace surface roughness or head-in-pillow faults. A good reference work describing many of these manufacturing assembly variances and their effects on high-speed serial I/O can be found here: "Tutorial: Board Test of DDR3/DDR4 Memory and Serial I/O".

Defects on High-Speed Serial I/O – Part 2

Short and open circuits can have either subtle or dramatic effects on bus performance. These can range from higher bit error rates and slower system performance due to link re-initializations and packet re-transmissions, to outright lane drop-outs (either intermittent or permanent).



These shorts and opens can be difficult to replicate in real-life experiments. Simulating an open by pulling a jumper is not the same thing as a missing ball, the presence of a solder void or a head-in-pillow defect. And simulating a short circuit by closing a jumper doesn't approximate excess solder bridging between two balls. Experiments will often result in catastrophic bus failures because high-speed buses aren't designed to survive these types of induced failures. But real-life opens may still permit a level of coupling on the nets, so the bus continues to operate. The same thing applies to shorts: the rejection of common mode noise may still allow the bus to run, albeit at a reduced level of performance.

Here's an example:

A Short Circuit: Tx1- to Tx2+



Figure 23: Short between Tx1- to Tx2+ on differential pair

The negative leg of a transmit lane (Tx1-) is shorted to the positive leg of an adjacent transmit lane (Tx2+). But, as in the previous example, the receiver continues to operate by rejecting common mode noise and energy from Tx2+ is already coupled to both Tx1+ and Tx1-, so some of the additional coupling energy will still be rejected. Again, the bus may continue to operate, but its performance will be impaired. How severe the impairment caused by an excessive bit error rate will determine the degree to which the bus' throughput will be impaired by packet re-



transmissions and PHY layer re-initializations or whether the performance degradation is intermittent or constant.

Defects on High-Speed Serial I/O – Part 3

A summary table of the handful of defects and their effects that have been reviewed is as follows:

Defect	Effect	Impact
Open (ope missing een)	Low	Some enhanced BER and crosstalk. Link
Open (one missing cap)	LOW	will run.
Short (Tyl to CND)	Modium	Performance impairment due to enhanced
Short (TXT- to GND)	Medium	crosstalk. Link may downgrade.
Short $(T_x 1 + T_x 2 \perp)$	Modium	Performance impairment due to enhanced
Short $(1x1 - to 1x2 +)$	Medium	crosstalk. Link may downgrade.
Short $(Tx1+ to Tx1-)$	High	No signal. Link will downgrade.
Open (two missing	Uich	Unfiltered DC voltage biases likely to cause
caps)	підіі	link to fail.

Table 1:	Serial I/O	defects and	their effects

Now let's consider three technologies that can detect some or all of these defects:

Boundary-Scan Test (BST) – A combination of the original IEEE 1149.1 boundary-scan standard (also known as JTAG) and its updated 1149.6 version will detect all of the defects listed above as long as the associated devices comply with these two IEEE specifications. It is important to have comprehensive DC and AC boundary scan coverage to detect the entire universe of short circuit and open circuit defects on both sides of the capacitors. Because of the complexity of a comprehensive 1149.6 implementation, many vendors' solutions fall short of 100% shorts and opens coverage.

Processor-Controlled Test (PCT) – The processor's debug port and run-control facilities can be used to examine the serial I/O status registers of devices and detect CRC errors, as well as link width and speed anomalies. In the table above, PCT will detect all defects listed as "High" and "Medium" under the effect column. A comprehensive library of devices supported by a PCT tool is critical because different devices from different manufacturers such as PLX, Broadcom, Mellanox, IDT and others all have different status register definitions. Researching these would



require man-months of manual effort. PCT also runs below the operating system and BIOS/boot loader, making it extremely effective for detecting defects that prevent a board from booting.

Intel IBIST – This is a bit error rate and margining tool that uses embedded instrumentation within Intel's silicon to detect the defects labeled "Low" under the effect column in the table above. Different types of defects will take I/O outside of its allowable range for voltage and/or timing. As a result, margining both and comparing these against a baseline to locate violations and/or skew is important.

All of these test technologies have trade-offs in terms of ease-of-test-implementation, elapsed test time and test coverage. Moreover, each technology will detect a class of defects that the others may miss. Consequently, all three technologies are needed to detect defects on high-speed serial I/O.

Testing SATA 3

Serial ATA 3 (SATA 3 or SATA III) is a differential bus running at 6 Gbps. It's commonly used on computer motherboards, such as notebooks, as a connection to mass storage devices. SATA 3 tests should analyze whether the performance of the mass storage device is impaired by poor performance on the bus.

Differential buses pose unique assembly test challenges. Because they present in common mode, some structural defects (such as shorts and opens) will be invisible to functional test. And some at-speed faults are invisible to traditional structural test technologies, such as ICT or boundary scan. A combination of test technologies is needed to capture the universe of potential structural, functional and performance-impacting failures.

As an illustration of this, we recently ran an Intel HSIO margining test on SATA 3 buses on two motherboards from two different vendors. The margin graph for the better performing bus looked like this:





Figure 24: SATA margining eye diagram

The poorly performing motherboard exhibited the following arrowhead-shaped eye diagram:



Figure 25: Arrowhead-shaped SATA margin

What does this mean? It's the classic question: "if a tree falls and no one is around to hear it, does it make a sound?" Many computer users won't notice the difference. Their computers will just run more slowly than they should, but the users won't notice because they have no basis for comparison. In the case above, the ODM was lucky enough to catch the failures illustrated in Figure 25 before a huge number of boards were built and they were inflicted upon unknowing users.



Here's a link to another interesting blog from Thierauf Design & Consulting on the pitfalls of taking shortcuts with regards to signal integrity: "<u>Walking through your layout database</u>".

Equalization

Equalization is a signal processing technique that enhances signal quality on the transmit or receive sides of a serdes connection. It is one of the key technologies that improves the margins of a signal and thus the size and shape of the eye diagram. This enables signal reconstruction despite the vagaries of crosstalk, inter-symbol interference, non-linear frequency response and other factors.

Adaptive equalization, in particular, reacts in real-time to the changing circumstances of a circuit board's operating environment and adjusts equalization parameters on the fly to optimize the strength of the signal. But, as with any technology, optimum and power-efficient results are typically only achieved when equalization is tuned to the specific system configuration since all configurations cannot be modeled.

Adaptive Equalization - Part 1

Modern high-speed I/O equalization schemes typically include both fixed (programmable) and adaptive components to ensure signal integrity even in adverse system conditions. What tools are available to ensure that these equalization techniques are working properly on a given system?

At higher speeds, inter-symbol interference (ISI) and crosstalk can distort signal integrity to the extent that the eye (margin) on a channel is essentially closed. Think of equalization as the math inside the chip (on both the transmit and receive sides) that reconstructs the signal and provides sufficient margin so that your board works properly. When done well and under good conditions, equalization can allow silicon receivers to distinguish bit levels at the bus' desired bit error rate (BER) even when the eye diagram measured on the board's interconnect appears to be compromised.

Although different serial bus technologies like PCIe, SATA, USB 3.0 and others use different approaches, equalization can be classified in two general categories: fixed and adaptive. Fixed techniques are locked in during board configuration. A fixed receiver scheme such as CTLE



System Marginality Validation of DDR3|DDR4 Memory and Serial I/O

(Continuous-Time Linear Equalization) has limited tuning ability and is quite susceptible to variances in both silicon and board manufacturing processes, but CTLE has the benefit of requiring a small footprint in logic circuitry on the chip. Adaptive equalization, such as Decision Feedback Equalizer (DFE) and Automatic Gain Control (AGC) continually self-adjust based upon the unique characteristics of the system, including board and silicon processes, temperature, voltage and other conditions. As you might expect, adaptive equalization technologies are more expensive than CTLE in terms of the silicon footprint required.

It is worthwhile noting that equalization can actually have an adverse effect on random noise, such as that induced by crosstalk. On the transmitter side, pre-emphasis and de-emphasis can make crosstalk worse. On the receiver side, CTLE and Feed-Forward Equalization (FFE) amplify crosstalk noise. DFE has no effect on crosstalk noise. A good discussion on this is in the Test & Measurement World article: "<u>Crosstalk problems are back</u>".

So, given the above, tools based on embedded instrumentation can be very helpful setting up and tuning equalization parameters, as well as measuring true margins at the silicon level (not on the board interconnect itself). A great technical paper on the effects of equalization on signal integrity is here: "<u>Margins (Eye Diagrams) Follow the Silicon</u>".

Adaptive Equalization – Part 2

Let's discuss why a chip's equalization parameters should be tuned.

A well-tuned system has a good balance between the equalization provided by the fixed and adaptive logic. This is illustrated below.



Figure 26: Balanced fixed/adaptive equalization

But there are cases where fixed and adaptive equalization are out of balance. In the example below, the fixed programmable settings are doing too much work, leaving adaptive equalization



very little wiggle room to improve the I/O margins and making collapsed eyes more likely in systems in the field.



Figure 27: Over-equalized I/O

And in another example below, the fixed settings are too conservative, making adaptive equalization do more work. And since the adaptive circuitry has a larger silicon footprint, it will draw more power on the chip.





It's important to measure and tune both the fixed and adaptive equalization settings prior to shipping systems into the field. Since each system shipped may possibly have different chips (e.g. different PCIe endpoints, DIMM suppliers, etc.), plus the PVT variability due to the chips themselves can be different and various PCB vendors' loss profiles will vary as well, etc., all these should be measured and tuned for as well.

Adaptive Equalization and Power Consumption

Poorly tuned adaptive equalization will cause systems to gulp power, instead of sipping it.

Adaptive equalization becomes very important at higher bus speeds, where signals are corrupted by reflections, coupling, and noise from chip and manufacturing board variances, which lead to higher attenuation. Such signals must be reconstructed to provide acceptable operating margins and for a circuit board to operate properly. Except for systems with very low requirements for power and size, adaptive equalization is routinely deployed for any system with bus speeds above 5Gbps. For example, DFE is provided on Intel® QuickPath Interconnect (QPI) and PCI



Express (PCIe) Gen3 buses. Adaptive equalization continually adjusts itself to the unique characteristics of each combination of chips, boards and variances on a particular board.

The problem with adaptive equalization is power consumption. DFE, for example, can add anywhere from 15% to 30% to the overall serdes power budget for an optimized PCIe design. At the 40 nm process node, a general industry rule is that DFE requires 5 mW per Gbps of data per each DFE TAP; so a standard 5-TAP implementation requires 250 mW/channel at 10 Gbps. At 28 nm, if we estimate total serdes power consumption to be 120 mW per full duplex channel at 8 Gbps, we see that DFE for a single x16 PCIe Gen3 port can consume about half a watt. On a given high-end design, which will have dozens and dozens of channels, adaptive equalization can easily consume 5-10W or more.

So, tuning adaptive equalization is extremely important to reduce overall power utilization. This involves striking a balance between fixed and adaptive equalization so that the adaptive logic is not doing too much work. For example, if a link is under-equalized, the fixed equalization settings (set in the BIOS or boot loader) are too low and the adaptive equalization carries more of the burden. Not only does this consume more power, but typically it strains the DFE at the extreme of its operating region, eroding margins on lanes that are on the edge of the guard band. This is visually represented below.



Figure 29: Adaptive equalization working too hard

This topic is discussed relative to a methodology for tuning the fixed and adaptive equalization on Intel Xeon platforms in this eBook: "Signal Integrity Validation for Intel Xeon Platforms".

Silicon Issues

The correlation between signal integrity and semiconductor process variances is not welldocumented, yet it directly affects circuit board system margins. Process variations within and between manufactured lots of devices, defects within chips, design flaws and silicon aging that



manifest themselves over time – all of these affect the performance of high-speed serial I/O and memory at the macroscopic circuit board scale. Most signal integrity engineers concern themselves solely with the signal integrity of the circuit board design without considering the contributions of devices on the board to its marginality.

The Intel Cougar Point SATA Bug

In early 2011, Intel discovered a design issue on its Cougar Point chipset and took an approximately \$700 million charge against earnings to repair and replace affected parts and systems. What could have been the root cause of this and how could it have been prevented?

Product recalls happen all of the time. Some of them make the news and some don't. Design issues and manufacturing variances can affect the operating margins of a product and over time erode the performance of a system to the point where owner/operator notice sluggish execution, excess power consumption, or hangs and crashes. These issues are associated with both the chips and boards that make up a system. The Intel Cougar Point problem is especially interesting because of its financial impact and its association with certain design errors.

In an interview with Intel's Steve Smith (found here), the root cause of the SATA problem on the Cougar Point chipset was traced back to a transistor in the 3 Gbps phase lock loop (PLL) clocking tree. This transistor was biased with too high of a voltage, which could result in a failure of the SATA ports 2 through 5 over time. In fact, the problem could be coaxed out of the chipset by running the part at elevated temperatures and voltage. Intel discovered this problem itself with thermal chamber testing. The differential AC-coupled SATA physical layer uses embedded strobes derived from the PLL to clock the 8b/10b encoding, so leakage and drift in the PLL logic ultimately leads to clocking marginalities and an increasing number of re-transmits over time, which degrade performance on the bus and ultimately (over months? years?) could cause the ports to fail.

Aside from testing their designs across a wide swing of temperature and voltage conditions, circuit board signal integrity engineers have other tools at their disposal to catch these kinds of defects. Certainly stressful patterns intended to exercise the I/O logic to its practical limits and generate crosstalk, ISI and clock recovery problems is an essential tool. The challenge is this: SI



design or manufacturing issues can be caused by both the chips as well as the boards where the chips are deployed, so critical marginalities in either or both of these can exacerbate and compound any system defects. User dissatisfaction with the product is the eventual outcome. Tools such as those described in the e-Book "<u>Bandwidth Tests Reveal Shrinking Eye Diagrams</u> and Signal Integrity Problems" can help catch these problems before they result in field issues or product recalls.

Margins (Eye Diagrams) Follow the Silicon - Part 1

We know from empirical evidence that a system's operating margins are as sensitive to the chips on the board as they are to the board's design and manufacturing processes. Why is this so?

The topic was discussed in this white paper: "<u>Margins (Eye Diagrams) follow the Silicon</u>", that was originally presented at a DesignCon as "*Platform Validation Using Intel® Interconnect Built-In Self Test (Intel® IBIST)*." The white paper shows that a system's overall signal integrity is quite sensitive to the chips deployed on the circuit board. So, if you test the same board design numerous times with different chips, you'll get swings in the size and shape of the eye diagrams for high-speed serdes I/O and memory. Obviously, the chip production process is subject to variances. This is apparent when silicon migrates from wafers to die, and especially in die shrinks (i.e. when the size of the die is reduced; for example, from 45nm to 32nm to 22nm and so on) with regards to process variation and leakage.

This is due to chip design being very analogous to board design, but on a nano-geometry level. It's easiest to see this analogy visually:



Figure 30: Silicon into packages and onto boards





Figure 31: Silicon/package/board routing

Given the complexity of I/O floor planning, logic synthesis, placement, clock insertion, routing, timing closure, manufacturing and so on, margins can vary considerably from chip to chip.

Add chip fabrication variances to the manufacturing variances from one circuit board to the next and it's easy to see that a typical signal integrity validation cycle might involve a minimum of five tests on five systems (sometimes referred to as a '5x5' procedure). For more information on a typical validation methodology using this approach, see our application brief "<u>Signal Integrity</u> <u>Validation for Intel Xeon Platforms</u>".

Margins (Eye Diagrams) Follow the Silicon - Part 2

Let's dive a little deeper into why signal integrity depends on the silicon. First, we'll look at wafer and die manufacturing variances.

The white paper "<u>Platform Validation Using Intel® Interconnect Built-In Self Test (Intel®</u> <u>IBIST)</u>" shows empirically that operating margins (the size and shape of the eye) are very dependent on the chips on a board. This is due to the design and manufacturing variances within the die and packages. Let's zero in on manufacturing variances in particular.

The impact of manufacturing variances has become much more pronounced below the 90 nm process node. Prior to 90 nm, wafer yield profiles were mostly impacted by imperfections in the



silicon, including those caused by dust or dirt in the environment, hence the refinement of clean room science. Defects tended to be in the range of $2/cm^2$ (two defects per square centimeter) and they tended to cluster. As a result, basic precautions produced satisfactory yields. In addition, small die sizes also statistically improved yields. This is what it looks like:



Figure 32: Wafer defects and variances

But below the 90 nm process node (Note that many complex chips are at 28 nm now. Intel is at 22 nm with its Ivy Bridge chip.) the yield profiles are almost entirely dependent upon manufacturing process. Wafer processing now includes phase-shift mask lithography, chemical-mechanical planarization and other complex steps, which are far more damaging to the routing layers and produce non-uniform routes. For high-speed I/O routes within the chip, this can analogously look like this:



Figure 33: Non-uniform routes within devices



The ultimate result is that process variation can have a significant impact across a wafer or a die based upon their sizes. And when the geometries are so infinitesimal, effects such as propagation delay, crosstalk and current leakage are much more pronounced. As a result, more rigorous PVT (Process/ Voltage/ Temperature) testing must be applied. A 5x5 testing methodology (five tests on five systems, each with different chips) is the minimum necessary to establish confidence in design margins. (See our eBook: "Signal Integrity Validation for Intel Xeon Platforms".)

Silicon Aging and Signal Integrity

Fruit rots. Tires wear out. And silicon ages. Let's look at the degradation process and its effects on signal integrity on-chip as well as chip performance.

As metal oxide semiconductor field effect transistors (MOSFETs) scale to ever-smaller geometries, speed and transistor density increase while active power per transition decreases. All of these are desirable in today's electronics industry, but the natural aging process for silicon is also accelerated as this scaling continues. Let's define aging within the context of degraded signal integrity (SI) on a chip, which in turn leads to more bit errors and reduced performance over time. Of course, aging will affect all attributes of a chip's performance, but SI is of particular interest due to the significant impact on system overhead at higher levels of the stack, such as errors at the serdes PHY layer.

Two sources of reliability degradation are:

- Charge trapping
- Electromigration

Some examples of charge trapping include random telegraph noise (RTN), bias temperature instability (BTI) and hot carrier injection (HCI). An excellent article <u>("Solve MOSFET</u> characteristic variations and reliability degradation issues") in EDN describes the effects of RTN and BTI. RTN occurs when a hole or an electron is captured in an oxide trap and the resulting charge is emitted from the trap. If this electron capture and resulting charge emission continues, the drain current (I_d) fluctuates, which causes the threshold voltage (V_{th}) to shift. RTN worsens at higher temperatures.



BTI is another example of charge-trapping that decreases I_d and shifts V_{th} . BTI in particular has a permanent component, from which the system almost never recovers.

HCI is a similar phenomenon where an electron or a hole gains sufficient kinetic energy to overcome a potential barrier and breaks an interface state. This can result in damage to the encasing dielectric material if the hot carrier disrupts its atomic structure. The presence of such mobile carriers in the oxides triggers numerous processes causing physical damage that can drastically change the device characteristics over prolonged periods.

Charge trapping degrades the chip performance over time until eventually the thresholds collapse.

Another source of reliability problems within chips is electromigration on the interconnects. Electromigration involves transporting material caused by the gradual movement of the ions in a conductor due to the momentum transfer between conducting electrons and diffusing metal atoms. Although electromigration damage ultimately results in failure of the affected IC, the first symptoms are intermittent glitches, which are almost impossible to diagnose. As described earlier, buses made up of differential pairs both on-chip and on a circuit board are, in fact, somewhat self-healing, insofar as open circuits may be overcome with sufficient coupling to allow successful data transmission, albeit at a higher error rate.



A view of interconnect breakdown from electromigration is below.

Figure 34: Electromigration induced defects



Certainly, the semiconductor industry intensively researches these reliability issues, and many mitigating technologies are in place to extend the life of chips. As usual, it is a race between shrinking geometries, form factors and process nodes, offset by new technology innovation.

As you might expect, different chips will have different levels of defects, variances and aging that affect the SI of any given design differently from other designs. In fact, we have conducted an empirical study of the effects of silicon variations on SI and found that poor SI is more closely related to the chips than to the circuit board where the chips are installed. This study can be reviewed here: "Margins (Eye Diagrams) Follow the Silicon".

Miscellaneous Technical Topics

Signal integrity is a fascinating topic. You could spend a lifetime becoming an expert in it. Some, such as Ransom Stephens, Eric Bogatin, Howard Johnson and many others have done so. These experts have contributed greatly to the engineering community by writing on the subject and inspiring a generation of engineers to explore this exciting field.

In the following section, the technology behind stressful or, what is sometimes referred to as "killer," pattern generation and error checking, as well as the statistical model for SMV, are explored.

PRBS31 and Validation of High-Speed Serdes

What's the right dwell time, or pattern length, to consider when checking for intersymbol interference, clock recovery and circuit drift?

A PRBS31 pattern (pseudo-random bit sequence of length $2^{31} - 1$, or 2,147,483,647 bits) is considered the gold standard when it comes to stressing high-speed I/O buses like PCIe, 40Gbps Ethernet, and OIF/CEI 11G-SR. PRBS31 provides a very stressful environment so that random jitter (RJ), sinusoidal jitter (SJ), intersymbol interference (ISI), crosstalk and other flaws can be detected. In other words, by using it, you'll be able to achieve a high level of confidence in the operating margins of a design.



The problem with PRBS31 is it costs too much. It usually takes about 20 repeats of the PRBS31 pattern to meet the requisite confidence levels. High-end 33GHz oscilloscopes can't store even one pattern of that size. And bit error rate tests (BERT) don't work either since RJ and SJ cause pattern elements to drift between repetitions.

Of course, there are different pattern generation and capture algorithms, most notably the OIF-CEI CID jitter tolerance pattern (mentioned in this EDN blog: "<u>PRBS31: Slower, costlier,</u> <u>worse</u>"), which are less costly in terms of test time. This particular implementation requires a pattern of about 21,000 bits, which, since it is arbitrarily defined, at least in part, rather than algorithmic, means that resources on the ICs are needed to store the pattern. Given this, a test tool like ScanWorks can be used to initiate the stress and retrieve the results through the chip's JTAG TAP.

To read more about alternative and better ways to validate signal integrity on high-speed I/O, check out this eBook: "<u>Bandwidth tests reveal shrinking eye diagrams and signal integrity</u> <u>problems</u>".

Signal Integrity Testing with CRCs versus Pattern Generation and Capture

The question came up recently about whether engineers could just check the CRC error counts coming from the operating system and deduce from this data the quality of the signal integrity and whether operating margins are being exceeded. After all, CRC checks for bit errors, right? Here's why this is not good enough:

This eBook has shown that design defects and/or silicon and board manufacturing variances all contribute to a reduction in the operating margins on today's circuit boards. In addition, increasing bus speeds only exacerbates this problem.

Many of the bus technologies use schemes such as encoding, scrambling, and adaptive equalization to ensure adequate margins or open eye patterns, but the effects of jitter, intersymbol interference (ISI), crosstalk and other impairments must be simulated in order to test for their effects.



Let's look at PCIe Gen3 for a moment. It's an AC-coupled, differential bus that uses embedded clocking to provide a robust and survivable data path. It uses 128b/130b encoding and data scrambling to avoid long strings of consecutive individual bits (CIDs) because CIDs are the bane of signal integrity. They dramatically impact the clock recovery (CR) circuits' ability to lock and hold. Encoding and scrambling also have the benefit of achieving DC balance in the bit stream, reducing data wander and improving error recovery.

PCIe encapsulates its data within a Transaction Layer Packet (TLP) with a Cyclical Redundancy Code (CRC) because this protects the entire packet with the exception of the framing start/end bytes. A TLP looks like this:



PCle Transaction Layer Packet (TLP)

Figure 35: PCI Express Transaction Layer Packet (TLP)

For PCIe 3, the Link CRC (LCRC) is 32 bits wide based on the large, variable-sized payload. The End-to-End CRC (ECRC) provides some level of data integrity for different link hops. For other buses like QPI, which use a smaller, fixed-sized payload, the LCRC is 8 bits.

Now that we've covered that background, let's look at the four reasons why CRC checking by itself is inadequate, especially when compared to the results returned by pattern-based checking.

1. It takes a long time to detect failures at nominal voltage and time

PCIe Gen3 runs at roughly 8 Gbps and is rated within the PCI-SIG specification for one bit error in 10¹². In this context, "rated" means that at nominal voltage and time, BER should be below this error rate. This is because the signaling schemes across all serial buses are never guaranteed to deliver the bits perfectly across the interconnects. The physical layer is always designed to minimize the probability of incorrect transmission and/or reception of a bit. The probability doesn't drop to zero, but rather to the rated BER in the specification. When the bus' actual BER is below the rated BER, routine transmission errors are allowed by the physical layer. Recovery



mechanisms are employed in the link layer so that higher level functions are not aware of or affected by these errors.

As a result of these factors, traffic must be generated on the bus for a considerable period of time to manifest errors above the rated BER threshold. The confidence level in the bus is given by the equation in our whitepaper "<u>Platform Validation using Intel Interconnect Built-In Self Test (Intel IBIST)</u>". For a bus like QPI, which has an error rate of 1 in 10¹⁴, achieving a high confidence level can take days or weeks of testing. Engineers don't have weeks to test signal integrity given today's aggressive design delivery schedules.

2. The design's true margins under real-world conditions are not determined

System-level OS-based testing with CRCs is usually performed at nominal time and voltage. In other words, the design is said to be "fresh," which means that the operating conditions are more or less perfect. But we've seen that process/voltage/temperature (PVT) effects can result in a wide swing in margins. That's why silicon vendors bin their devices based upon their performance at the fringe of the envelope. And drift in high-speed I/O circuits – aging of capacitors, variations in power supplies, the effects of current leakage on gates over time, etc. – will also negatively impact operating margins. Simply testing a board design under ideal conditions can be misleading with regards to the quality of signal integrity.

Testing the operating margins at the worst-case extremes with synthetic stress patterns can map margins to an eye mask that takes into account drift and PVT effects.

3. CRC is not perfect

CRCs use polynomial arithmetic to create a checksum against the data it is intended to protect. The design of the CRC polynomial depends on the maximum total length of the block to be protected (data + CRC bits), the desired error protection features and the type of resources for implementing the CRC, as well as the desired performance. Of course, this will involve tradeoffs. For example, a typical PCI Express 3.0 packet CRC polynomial is:

x32 + x26 + x23 + x22 + x16 + x11 + x10 + x8 + x7 + x5 + x4 + x2 + x + 1

Whereas for Ethernet frames, the CRC generator may use the following polynomial:



x32 + x26 + x23 + x22 + x16 + x12 + x11 + x10 + x8 + x7 + x5 + x4 + x2 + x + 1

The PCIe 3.0 CRC-32 for the TLP LCRC will detect 1-bit, 2-bit and 3-bit errors, but 4-bit errors may escape detection. Bit slips or adds have no guarantee of detection. Burst errors of 32 bits or less will likely be detected.

For QPI, the 8-bit CRC can detect the following within flits:

- All 1b, 2b, and 3b errors
- Any odd number of bit errors
- All bit errors of burst length 8 or less
 - Burst length refers to the number of contiguous bits in error in the payload being checked (i.e. '1xxxxx1').
- 99% of all errors with burst length 9
- 99.6% of all errors of burst length > 9

4. OS-based traffic CRC checking doesn't really stress much

Most CRC-based tests saturate the bus with heavy OS-based traffic such as streaming video. This can increase traffic to more than 90% of the bus' bandwidth. This normal functional traffic is subject to 128b/130b encoding and scrambling that reduces the occurrences of long strings of 31 consecutive identical bits (CIDs) to fewer than 10⁻¹², which is the BER threshold. But stressing clock recovery (CR) circuits requires checking with longer CIDs. Just running traffic and checking CRCs doesn't cut it.

What is known as synthetic or killer patterns are necessary to aggravate all intersymbol interference (ISI) that might reasonably be expected, challenge the ability of clock recovery circuits to lock and hold, and check receiver circuitry against drift. A PRBS31 pattern, as mentioned previously, fulfills these criteria. The intent is to generate the most stressful patterns as possible and then check the bits one-by-one. It doesn't get more precise than that.

So after all this, let's ask: why is this important? Well, poor signal integrity indicates that a design has a high probability for being plagued by uncorrectable errors or system crashes, resulting in costly field repairs or even product recalls. Also, let's not forget power consumption.



If SI is not optimized, any application of adaptive equalization can increase power requirements by 15% - 30%. So, if the signal integrity on a design is poor, it will only get worse over time in the field. All systems run slower and eventually start to hang or crash over time. You want your design to run clean when it is first shipped.

The Statistical Basis of SMV

To determine the operating margins on a system design the sample size from which data is collected must be large enough to yield a statistically valid result. Margins are a statistical projection based on measurements on a certain number of prototype circuit boards. An adequate sample size will account for variances in the silicon fabrication processes, changes in temperature and voltages, finite test time, and a number of other factors. What is the math behind this?

As previously explained, SMV uses on-chip embedded instrumentation to collect data that will project the design's margins. Examples of embedded instrumentation include Intel® Interconnect Built-In Self Test (IBIST) or Freescale[™] DDR Validation Tool. Embedded instruments are used in conjunction with PC-based validation and test systems, such as the ASSET ScanWorks® platform, which will perform pattern generation, capture, loopback, voltage and time margining, and error checking among many other functions. Margin data gathered from a SATA 3 bus and projected by a validation system might look like this:



Figure 36: SATA 3 margin plot



In the margin plot above, green indicates no errors were encountered on this lane during the dwell time at the margining point (typically one second or thereabouts). Yellow indicates a correctable error was detected. Red indicates an uncorrectable error was detected.

The results vary each time the data is gathered. The data is affected by the following:

- Variances in the silicon's high volume manufacturing (HVM) processes (between and within fabrication lots of devices)
- A finite number of test samples taken
- An adjustment for the bit error rate measured during the limited dwell time
- Chip and board effects such as temperature, voltage, humidity, silicon aging, etc.

This can be depicted visually:



Figure 37: Deriving the SMV eye mask

Silicon vendors typically gather numerous measurements on a large number of fabricated devices. From this sample, a normal distribution is constructed and an eye mask developed, usually based upon a 95% two-sigma (2σ) or higher confidence level. The standard deviation is, according to measures of dispersion, the square root of the variance of the data set. For measurements of grouped data, the arithmetic mean and standard deviation are defined as: (1) an average that the mean of the data set must exceed, and (2) the number of measurements needed to reach the 2σ confidence level.







To determine the size of the data sample needed to produce a statistically valid projection of SMV, both the number of measurements on one circuit board and the number of different circuit boards tested must be taken into account. We have shown that variances in HVM processes contribute a significant amount to the performance variations in a design. (See our white paper: "<u>Margins (Eye Diagrams) Follow the Silicon</u>".) To reach statistical validity, a so-called 5x5 methodology is often employed. That is: five prototype boards each configured with its own set of chips make up the number of boards in the sample. Each board is tested five times for a total dataset based on 25 tests. This sample size of data will produce a high degree of confidence in the distributions projected from the data and that variances in both the chips and the boards in the design have been taken into account. The following empirical data was taken from a SATA margin run performed five times on one board:





Figure 39: A 5x1 (5 tests on 1 system) margin run

The data above shows a large amount of variability from one margin run to the next. But, if the dataset is based on 5 tests on 5 boards (25 total tests) and the average exceeds the defined eye mask, the results will be within two sigma 2σ (two standard deviations) for a high confidence level. But, since the process is based on statistics, that only deals with probabilities, adequately passing the eye mask only means that the risk of system crashes or hangs will be lower than if the design had not passed the eye mask; passing a margin test does not guarantee that the system will work perfectly every time. Conversely, failing the eye mask does not necessarily mean that the system is unusable. Such tests project probabilities and the confidence level associated with each outcome.

For more information on this 5x5 validation methodology, see our white paper, "<u>Signal Integrity</u> <u>Validation for Intel Xeon Platforms</u>". (Note that registration and an NDA with ASSET InterTech are required to download this white paper.)



Conclusion

System marginality validation (SMV) has emerged as a new, robust methodology for predicting the performance of an electronic system over its lifespan. Using on-chip embedded instrumentation within silicon, in conjunction with software that accesses this on-board logic, SMV 'sees what the silicon sees' and measures the true operating margins at the devices' I/O buffers behind the pins. As such, SMV can remove the absolute dependency designers have had on legacy test and measurement technologies such as oscilloscopes, which can no longer cost-effectively validate today's very high-speed serial I/O and memory buses.

SMV uses a statistical margining approach to take into account the variances introduced during manufacturing, such as PCB fiber weave, silicon fabrication processes, temperature, humidity, voltage and other factors to provide a confidence-based approach to predict how a system will perform in the field. This systematic approach applies holistically to an electronic system as the sum of its parts: die, package, chip, printed circuit board and final assembly.

As the speed, density and complexity of today's modern designs continues to increase, SMV will play an even more important role in enhancing the reliability and availability of electronic systems.

Learn More

Want to learn more? The definitive reference on how circuit board and semiconductor defects and variances, in both the design and manufacturing processes, affect system quality and reliability. This is highly recommended reading for engineers and managers in Development and Production. Register today!



