

DesignCon 2009

Platform Validation Using Intel® Interconnect Built-In Self Test (Intel® IBIST)

Stephanie Akimoff, ASSET InterTech, Inc.



Abstract

This paper describes platform validation using Intel® Interconnect Built In Self Test (IBIST). IBIST utilizes the embedded instrumentation technology present on high-end chipsets. This paper concentrates on testing multiple links of a high-speed serial bus in a single system. It will show validation differences between traditional methods using an oscilloscope and a third-party test methodology. These differences include test coverage, test time, and eye capture results.

Author Biography

Stephanie graduated from Pacific Lutheran University in 1999; double majoring in electrical engineering and computer engineering. She worked for Intel Corporation for seven years before taking a position with ASSET InterTech, Inc. as an application engineer. She is currently working with ASSET InterTech's ScanWorks® IBIST toolkit for Next Generation Intel® Microarchitecture (Nehalem)-based platforms development team on validation and customer issues.

Introduction

Platform validation using Intel® Interconnect Built-In Self Test (Intel® IBIST) reveals important device and platform relationships that were present before but were not fully understood. This paper discusses:

- The relevancy of these relationships based on a thorough understanding of what is being tested
- The statistical accuracy of the test described in confidence levels
- The stressfulness of the test
- The effects of testing all the Intel® QuickPath Interconnect (Intel® QPI) interconnects at the same time
- The amount of validation test time. The cost of validation cycles is very important. Man hours often make up a significant amount of the validation budget.

The effects of the devices used in the unit under test (UUT) have always been present, but now some can be characterized by the signal integrity platform validation engineer. This paper will show some of the characteristics seen and some of the effects they have on a particular system using ASSET InterTech's ScanWorks® IBIST toolkit for Next Generation Intel® Microarchitecture (Nehalem)-based platforms which utilizes JTAG to run Intel IBIST actions.

What is Intel® IBIST?

Intel® IBIST is a design validation and test architecture. Embedded into many Intel processors and chipsets, Intel® IBIST enables chip-to-chip interconnect testing and design validation of high-speed buses on a printed circuit board. "IBIST leverages the boundary-scan IEEE 1149.1 specification as the hardware and software communication methodology for accessing and controlling its embedded on-chip capabilities."¹

The Intel® QPI uses Intel IBIST for validation. To use IBIST in this case, each link needs to be put into loopback (see Figure 1). Loopback consists of a master IBIST device transmitting patterns from an internal buffer to a remote device. The remote device receives the patterns at the physical protocol layer and retransmits the data back to the master device without processing the data. For this to work, both devices have to be enabled and be able to send and receive data on the interconnect.




Figure 1. Loopback Mode Example

In Figure 1, both devices support IBIST. In general, the loopback remote device does not need to be an IBIST device. On this platform, for the Intel QPI interconnect testing, all remote devices contained an IBIST engine. Note that the master is both the transmitter and the receiver for the test data. The black arrows represent the traces on the platform under test. These traces are the portion of the test that is of concern to the signal integrity platform validation engineer.

Types of IBIST Tests

This paper will explore three types of IBIST tests:

- Pattern Generation and Checking (PG&C)
- Bit Error Rate Testing (BERT)
- Margining

Pattern Generation and Checking (PG&C)

The PG&C test makes certain that patterns can be run across the interconnect. PG&C uses pattern buffers in the IBIST master device to send a pattern across the link to the remote device. The remote device retransmits the same data it received at the physical protocol layer back to the master's receivers. The master device compares the patterns that were sent with what it received. If any mismatches are detected, the lane(s) and number of bits in error within a group of patterns are reported to the user. For the purposes of this paper, we are going to define data mismatches as errors.

The PG&C test takes a small sample. The complete test time can range from less than a second to a minute.

Bit Error Rate Testing (BERT)

BERT builds upon the PG&C test. An Intel® IBIST BERT takes the PG&C test specifications and runs it thousands of times; counting the number of times an error is reported. BERT testing involves a statistical analysis to determine a confidence level for a link. This methodology allows the signal integrity validation engineer to trade test time with measurement accuracy. "The statistical confidence level is defined as the probability, based on a set of measurements, that the actual probability of an event is better than some specified level"² (see Figure 2).

For example, using the formula in Figure 2 and a confidence level of 99 percent, with a probability for the link set at 10^{-14} and a speed of 6.4GT/s, the test should complete in about 20 hours. If 1 error is seen, then for the same confidence level, the test time is now extended to 29 hours. For 5 errors seen, the test time now extends to 57 hours.

$$CL = 1 - \sum_{k=0}^N \frac{(np)^k}{k!} e^{-np}$$

N= Number of errors allowed
n = Total number of trials (i.e., total bits transmitted)
k = Number of events occurring in n trials (i.e., bit errors)
p = Probability that an event occurs in one trial (i.e., probability of bit error)
CL= Confidence level

Figure 2. Equation to Determine Confidence Level

Margin Testing

The third type of IBIST test is margin testing. The data transfer mechanism is similar to PG&C and BERT except in this test the master varies the voltage threshold (VOC – Voltage Offset Control) and the sampling point relative to time (PI - Phase Interpolation). Margin testing does not test to the accuracy of a BERT but it tests at multiple margin points. This gives the engineer better insights to the amount of margin on a link especially when different stresses are added.

The individual test point results are plotted as pass or fail. The plots for each link resemble the eye diagram seen on a traditional oscilloscope. The test time per point is far less than for BERT. The typical dwell time at a particular test point is two seconds. This is equivalent to a confidence level of much less than one percent.

There are three types of margin tests; Cross, Full, and Limited Range.

- Cross margin tests only check the values along the horizontal and vertical axis. The number of IBIST steps on each axis is programmable. An IBIST step is equivalent to an IBIST unit. The margin point is reached when any lane on the link finds at least one data miscompare. This test gives the signal integrity validation engineer a rough look into the amount of margin on the link. Since minimal margin points are tested, the test takes about one tenth of the time of a full margin test.
- Full margin tests each voltage value for each PI value. If the PG&C test passes the next voltage increment is tested. If the test fails, the voltage and time values are noted, and the test moves to either the negative values to be tested or to the next PI to be tested. When the test is complete, the failing points can be charted resulting in a plot that is similar to a traditional oscilloscope eye.
- Limited range margin tests provide margining information in a specified range. The boundaries for the test are programmable. This test was not used for this paper.

An IBIST unit is:

- Undefined in terms of picoseconds (ps) and millivolts (mV).
- A register setting inside the IBIST engine.
- In the current family of devices being used to test, each IBIST unit is not of the same value in terms of picoseconds and millivolts. Therefore there is not a direct correlation in terms of real time and voltage units and IBIST units.
- Discussed in terms of voltage (VOC) and timing (PI) in this paper.

IBIST is used more qualitatively than quantitatively. This means that the shape, position and opening of the eye are important, not the raw values.

Test Platform Configuration

The platform for the tests in this paper has three devices- CPU0, CPU1 and IOH. They are connected by a point-to-point QPI interconnect. Both CPUs were socketed, which allows for interchangeability. The IOH was soldered on the board. Figure 3 shows the IBIST device configuration used.

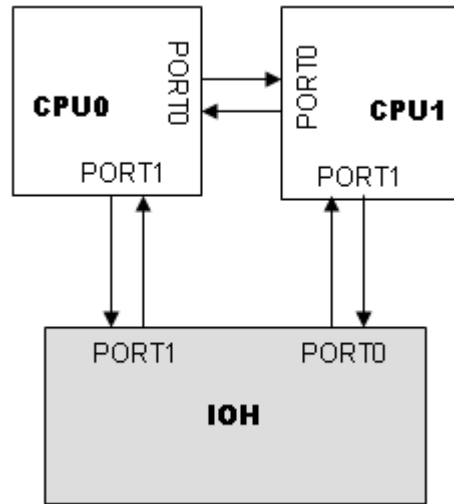


Figure 3. Platform Block Diagram

For this paper, we will use the following definitions:

- **Device**
A physical component that contains the IBIST engine. In Figure 3, there are three devices: CPU0, CPU1 and the IOH.
- **Lane**
A lane is a single differential pair. Each lane connects the transmitter of the IBIST device to the receiver of a second IBIST device. This connection is point to point. Daisy chaining devices is not allowed. There is granularity to this level in testing, but testing is done at the link level.

- Link
A link is a group of twenty lanes connecting two IBIST devices. Each link contains one master and one remote in loopback mode. There are three devices on the platform. Each can be either a master and/or a remote device. Therefore, there are six separate links on the platform. Testing is done at the link level.
- Port
A port contains two links, one in each direction. Each device can contain multiple ports. Each port is connected by a link to an equivalent port on the other device (see Figure 3).

BERT Testing Results

Multiple tests were run to a confidence level of 99 percent. No errors were found on any of the lanes during these tests. One reason for the lack of errors is that while a BERT test runs to a high confidence level, the point being tested has no variance in voltage or timing. This is the most optimal place to sample the signal. The link would not be reliable if an error were detected at this point.

Margin Test Results

For the cross test margining and full margin testing, a value of one IBIST step for PI and two IBIST steps for VOC were used. This allowed for good granularity while cutting the test time nearly in half compared to using a VOC step size of one IBIST unit. See the example in Figure 4.

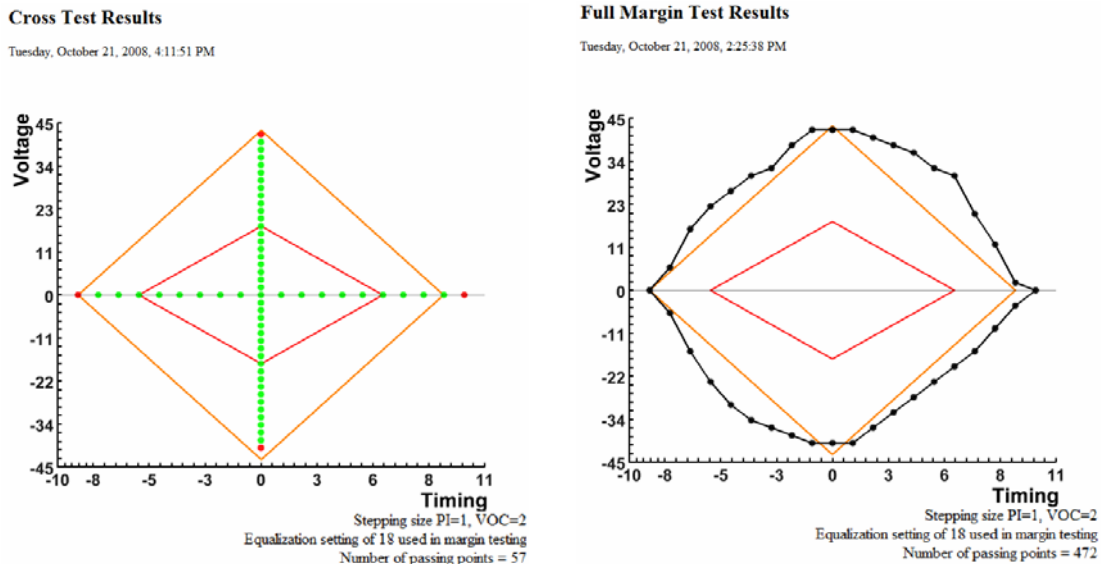


Figure 4. Cross and Full Margin Examples

For the cross test, the green dots represent a passed margin point. The red dots on the outer axis of the margin are the points at which the test failed. There are two eye masks. The inner mask is such that the test must pass this margin point. The outer mask is where

the margin should pass. For the full margin tests, the black dots represent the same margin criteria as the red dots for the cross test. The masks are the same between the two tests.

Variables that Impact Test Results

Odd vs. Even Samplers

Because of the speed of Intel® QPI, two alternating samplers (defined as even and odd) are required. These samplers are not defined as deterministic. For this testing, both the even and odd tests gave repeatable results. In all cases, the even samplers gave less margin than the odd samplers. Therefore all testing for this paper was done with even samplers.

Figure 5 is an example of the results taken with the same device, same link, and on the same platform.

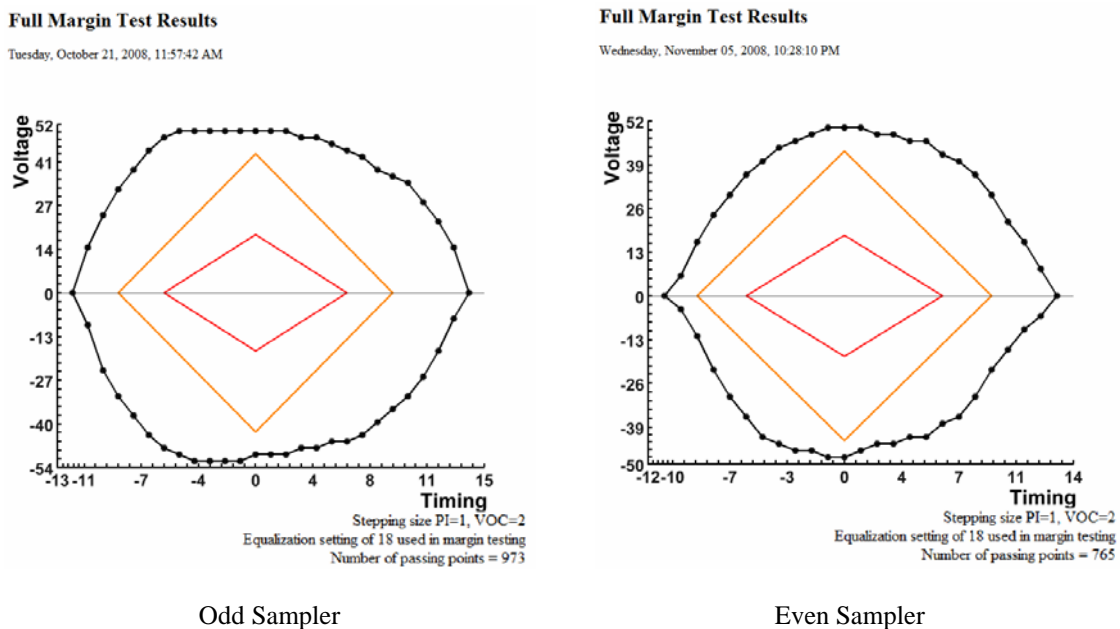


Figure 5. Full Margin Example with Odd and Even Samplers

Equalization Testing

Equalization is a function that is applied to the remote device for margin testing. The purpose is to change the strength of the signal driving to the original master's receivers. All lanes in the link have this value applied equally. One of the main reasons for changing this strength is to overcome the effects of ISI (Inter Symbol Interference).

Multiple equalization tests were run on several devices. Cross tests were performed incrementing the equalization settings from 3 to 27 out of a possible range of 0 to 31.

These were run on the same system, same device, and same link). This was repeated many times to assure data accuracy. In every case, the equalization settings affected the position and the amount of margin on the eye. The results can be grouped by their equalization values (see Figure 6).

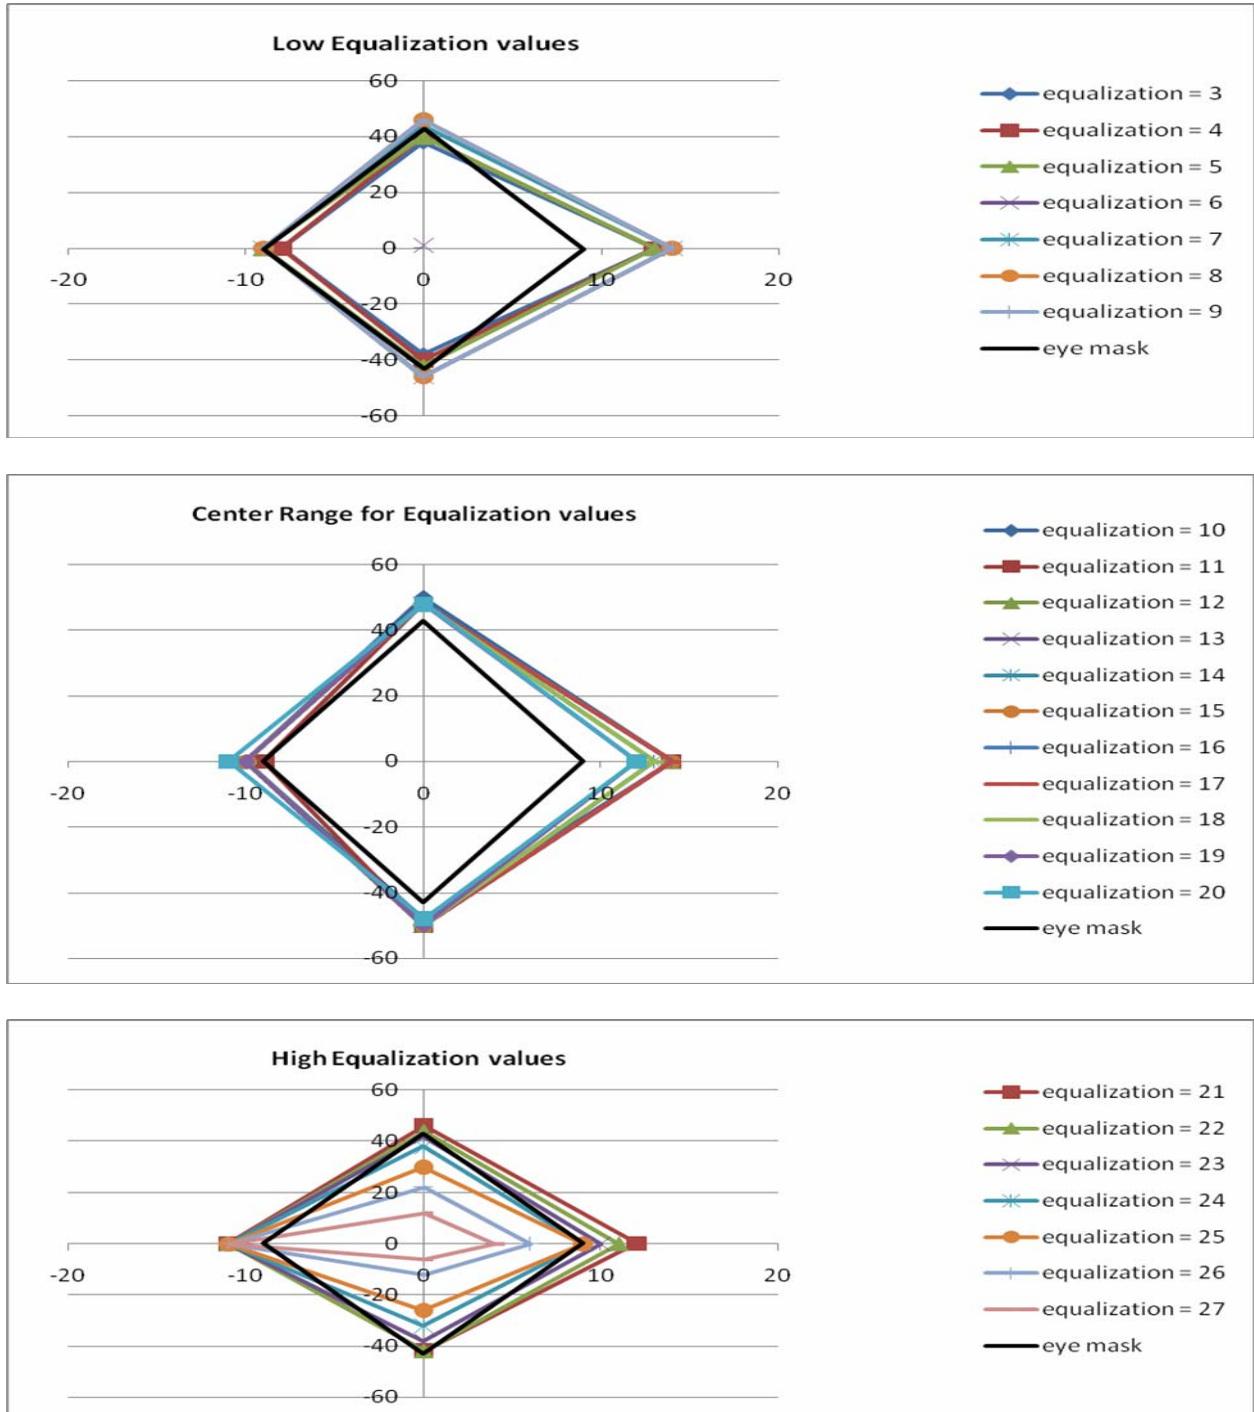


Figure 6. Equalization Results

The IBIST eye mask is only met in the center range of the equalization values. For the low equalization values, while the timing values are met, the eye looks to be shifted in the positive PI direction and the VOC does not always meet the mask criteria. For the high equalization values, not only does VOC decrease, but as the equalization values get higher, the eye starts eroding the amount of positive PI margin while the negative PI margin remains the same value. This data was consistent on multiple tests of the equalization values and multiple devices.

Device Being Tested

To look at the effects of the device on the platform testing, four CPU devices were tested on two identical systems using the test configuration shown in Figure 3. The matrix below describes which device was in which socket, and which socket was master for each test.

Master	Device	Slave	Device	System Used
CPU0	D1	CPU1	D2	original
CPU1	D2	CPU0	D1	original
CPU0	D2	CPU1	D1	original
CPU1	D1	CPU0	D2	original
CPU0	D3	CPU1	D4	original
CPU1	D4	CPU0	D3	original
CPU0	D4	CPU1	D3	original
CPU1	D3	CPU0	D4	original
CPU0	D1	CPU1	D2	second
CPU1	D2	CPU0	D1	second

Multiple cross margin tests were taken with each combination. Through this series of tests, we found:

- For each device (D1, D2, D3 and D4) the training values, failing point, and lane that failed followed the device and not its position on the platform.
- The training values were noted for VOC and PI for both the odd and even samplers on each lane. For each device, the values did not change more than ± 2 in any of the tests per lane.
- When that device was tested on a second system, the training values, failing point, and lane that failed were consistent with the results from the original system.

For example, at an equalization value of 18, the cross test results were:

Master	Device	Slave	Device	Cross Test Results			
				+VOC	-VOC	+PI	-PI
CPU1	D1	CPU0	D2	48	-50	13	-10
CPU0	D2	CPU1	D1	44	-42	10	-8

The devices were then swapped on the same system (device D2 was on CPU1 and device D1 was on CPU0). The cross test results were nearly identical - see the following table:

Master	Device	Slave	Device	Cross Test Results			
				+VOC	-VOC	+PI	-PI
CPU1	D2	CPU0	D1	44	-42	10	-9
CPU0	D1	CPU1	D2	48	-50	13	-10

- The training values for each lane were compared to the training values for the same device in the previous socket. These values followed the device.
- The failing lanes that caused the margin points were also compared. These too were identical to the failing lanes for the same device in the previous socket.

The lane failures and the training values are unique values that change with the device. This is significant as they did not follow the platform under the same test conditions (the only variable being the exchange of the devices).

The cross margin graphs in Figure 7 show the results of moving the D3 and D4 devices between sockets on the same system. The graphs show the amount of margin follows the device being used for the test more than the socket that it is installed in.

The testing shows that most of the IBIST test margin is device driven. The goal is to determine the amount of margin on the motherboard. What does this mean? Thinking will have to be turned around *from* how much of the effects are from the motherboard *to* how much of the effects of testing are not impeded by the platform. This consideration is equally as important as having all devices function with the platform. It is not feasible to match components for certain boards. The makers of the devices have given the amount of margin that is needed. If the board impedes this, then that margin will not be met.

Speed of the Interconnect

The speed of the interconnect is important. In testing, it was noted that there were far greater margins with a speed of 4.8GT/s than at 6.4GT/s. All tests were performed at 6.4GT/s.

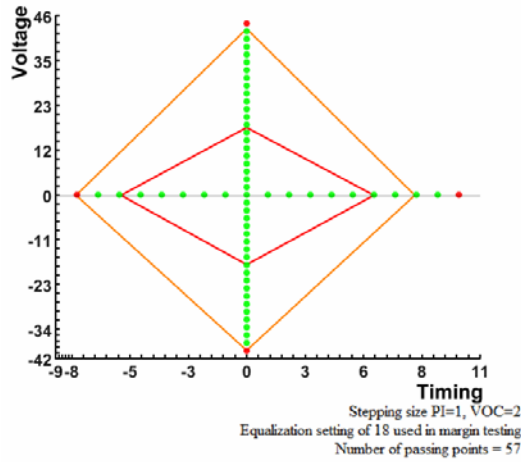
Stress on Links

Testing tries to maximize the amount of stress on the link's signals. This enhances possible signal degradation of the link. To see if the stress is maximized:

1. Baseline a link between particular devices and verify that the test results are repeatable.
2. Change the stress on the link and retest. The changes in the stress will cause changes in the margin.

Cross Test Results


Saturday, October 11, 2008, 6:52:23 PM



CPU0 (D3) master to CPU1 (D4)

Cross Test Results


Friday, October 10, 2008, 12:25:00 PM



CPU1 (D4) master to CPU0 (D3)

Cross Test Results

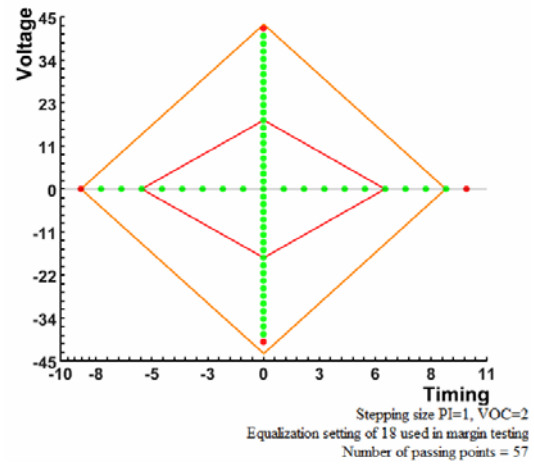
Tuesday, October 14, 2008, 10:42:57 AM



CPU0 (D4) master to CPU1 (D3)

Cross Test Results

Tuesday, October 21, 2008, 4:11:51 PM



CPU1 (D3) master to CPU0 (D4)

Figure 7. Cross Test Results for D3 and D4 Devices

3. Look at the margin to see if the stress increased or decreased.
4. Retest the margin several times with the same stress to verify the results are repeatable.

One can change the stress by changing the patterns applied, the dwell time, by running tests concurrently, or by combining some of these factors.

Stress from Patterns

All tests of randomly changing the pattern buffers from the default patterns showed improved margins. This shows that patterns have some effect on the amount of margin seen for the platform. The four default test patterns consist of four patterns that target different sensitivities.

The IBIST engine includes a Linear Feedback Shift Register (LFSR) that can be combined with the internal patterns in the data buffers. This combination produces a set of pseudo-random test patterns. Tests between using static patterns and pseudo-random patterns with the LFSR enabled were compared. The pseudo-random data patterns showed less margin, proving that they were more stressful than static patterns.

Stress from Dwell Time

Small variances in the IBIST margin point are normal and expected since the dwell time at each margin point was two seconds. The test was repeated multiple times with the two second dwell time. Most tests showed the exact same results in all four margin end points. A few tests showed a small variance in these margins. Using the equation in Figure 2, this margin point has a 0.0085 percent confidence level.

Multiple cross margin tests were taken while expanding the dwell time at a margin point to 60 seconds and 15 minutes. The following changes were seen with lengthening the dwell time (see Figure 8):

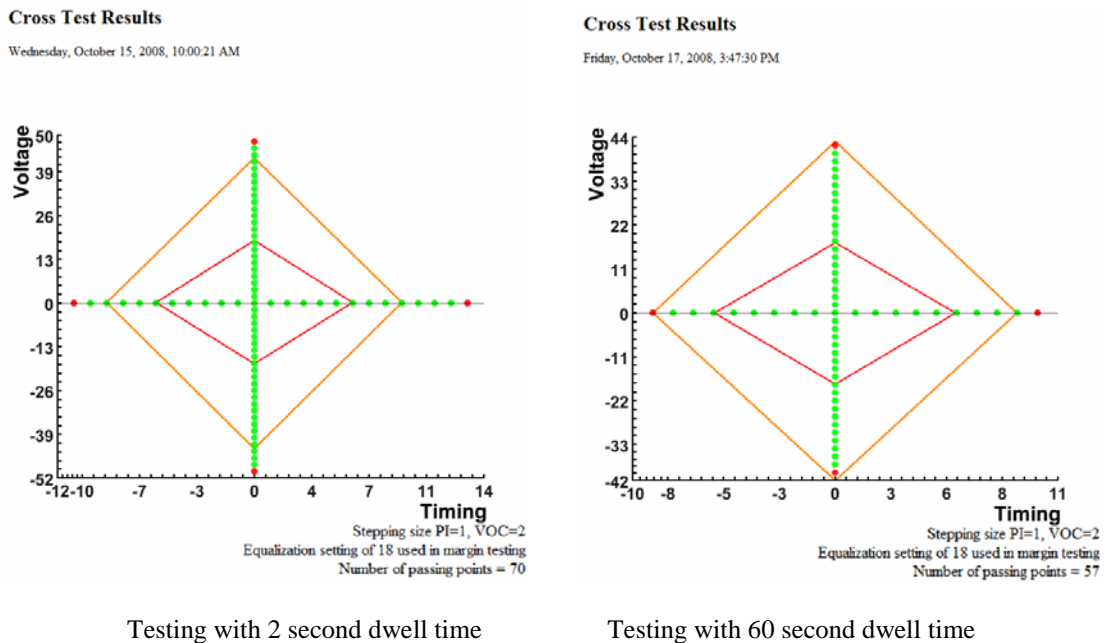


Figure 8. Cross Test Results with Default and Increased Dwell Times

- For the VOC margin end points, the 60 second dwell time yielded a 12 percent decrease in margin in both positive and negative margins. The results for a 15 minute dwell time per IBIIST margin point was the same as for 60 seconds.
- The increased confidence level for a 60 seconds dwell time is 0.38 percent. The confidence level for a 15 minute dwell time increases to 5.6 percent.

The overhead of the tool was evaluated for each dwell time. It was seen to be constant and there was no increase with increased dwell times.

The number of margin points directly impacts how long the test will take. While there is no guarantee that increasing the dwell time will decrease your margins, the preliminary test results show (for the dwell times tested) that the two second dwell time was the best case margin for that link, not the worst case.

Increasing the dwell time to reach a 99 percent confidence level is not feasible. If a full margin test is performed and the VOC and PI are set at one IBIIST margin step, then there would be about 800 margin test points. For a 99 percent confidence level, each of the 800 margin points would take 20+ hours. This means that it would take over 16,000 hours to validate a link. Tradeoffs need to be made between the confidence level wanted and the time available for testing.

Stress from Running Tests Concurrently

Tests can be run concurrently. This means that one device can be a master on more than one link, or a second test can be run on a separate link. Several links can be tested concurrently. Two different types of tests were used. One test, used as control, stressed a link that was unrelated to the master. A second concurrent test stressed another port on the master. The control experiment showed no significant change in margin between the stressed and non-stressed conditions.

The concurrent test with stress on more than one port of a single master showed loss of margin in most tests. This indicates that the additional IO stress caused by testing both ports on the device at the same time may reduce the amount of margin available. Testing with both ports active simulates the real world functionality of the device better than testing an isolated port.

Stress from a Combination of Sources

Combining the stresses of concurrency and a 60-second dwell time showed some degradation of margin and shifted the center of the eye (see Figure 9). The test with this combination of stresses took 1 1/2 hours compared with a 12 minute test time for the 2 second dwell time by itself.

The test results for the concurrent testing with increased dwell time (Figure 9) show that the decrease in margin was in both the negative VOC and the negative PI compared to the default dwell time of two seconds. Up until now all the tests were CPU to CPU. Not all

ports of a device have the same sensitivity to stress. For all the devices tested, CPU to CPU passed the eye mask under all conditions. For the CPU to IOH link, this was not always the case. Most of the issues with this testing were due to a known issue with the chipset revision used.




Figure 9. Concurrent Cross Test Results with Default and Increased Dwell Time

Non-Optimal Equalization

Testing was done to show the effects of non-optimal equalization values. For this test, two different equalization settings were chosen. The BIOS for the platform under test had a CPU to CPU equalization set at 18. This was chosen as a test control setting. In the previous Equalization Testing discussion we saw that 18 is in the center range for the equalization settings. Ten was chosen as the second test equalization setting.

Note that the center point for all equalization settings is 16 which is 0x10h. It is quite conceivable that someone could accidentally confuse decimal with hex. The equalization setting for the IOH in BIOS is 16 decimal.

Testing showed that some devices are more sensitive to non-optimal equalization than others. Using the CPU as the master, changing the equalization had little effect on the CPU to CPU link. Yet, with the CPU as master to the IOH, changing the equalization had a large effect on margin. Figure 10 shows concurrent testing results with the same device on the same system. The only variable was the equalization setting on this link. The results show that for the optimal equalization setting, the link passes the margin mask. For the non-optimal setting, the margins are compromised.

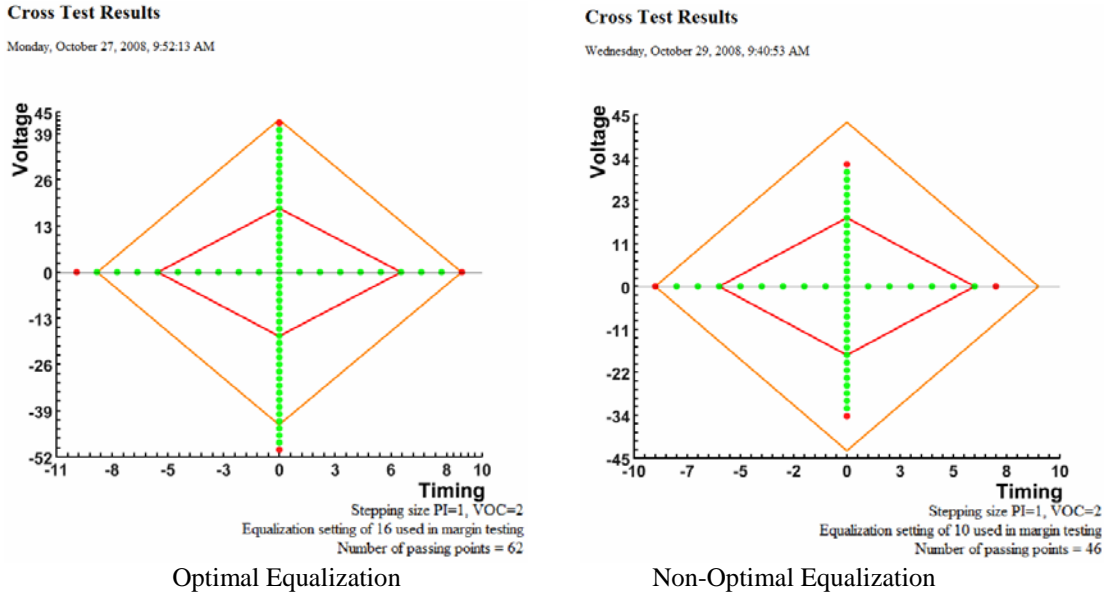


Figure 10. Cross Test Results Versus Equalization

Oscilloscope Test While Running an IBIST BERT

An oscilloscope was used to look at the interconnect on one lane while running the IBIST BERT tests as performed on the link as described in the sections above. The same system was used as characterized in prior tests. The only variant from previous tests was connecting the oscilloscope. Probes were attached to vias beneath the CPU socket on the back of the motherboard.

No BERT errors were seen with the ScanWorks tool. The oscilloscope showed a clean eye for this one lane (see Figure 11). No special equipment was used to capture this eye. The oscilloscope capture was taken over the course of the 20+ hour test and shows an open eye. As discussed previously, the BERT test takes a statistical confidence of the best case point on an eye. The multiple hues in the picture graphically capture the signals placement in relation to the test. For a very few captures, the eye was degraded by up to 60 mV. This can be related to the IBIST BERT test in that no errors should be seen on this interconnect under this type of stress.

Figure 11 can be compared to the IBIST margin test. The oscilloscope capture shows the statistical history of what happened on the lane during testing. The different colors represent the amount of times the signal crossed that point. The redder the color (hotter), the more times that this was the eye result. The bluer the color (colder), the fewer times this was the eye result. A small statistical sample shows that the margin was uniformly decreased and showed the limit. There seems to be a definite divided region which is hotter. There is a definite cold region also in which very few samples are contained. This is like an inverse of the confidence level that was used for IBIST testing.

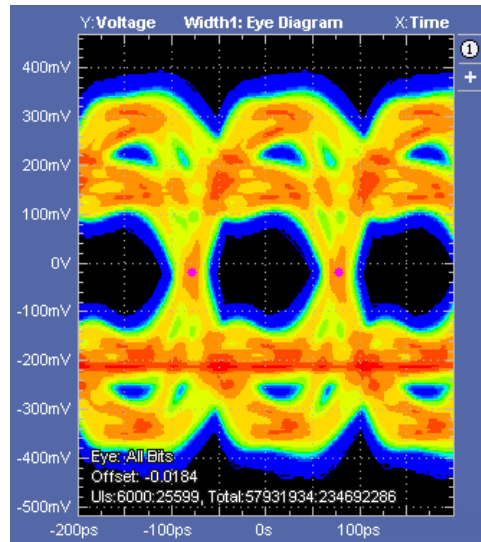


Figure 11. Oscilloscope capture while running BERT

For an IBIST confidence level, there is a certainty that an error will not be seen at this VOC/PI point. For an oscilloscope confidence level, there is certainty that a signal will meet this mV/ps level. The hotter the color, the higher the confidence that the signal will reach this point. The colder the color, the lower the confidence that the signal will reach this point. The eye in Figure 11 was taken during 20+ hour test so there is a high level of assurance in the eye results.

The IBIST full margin tests use a statistical confidence at each margin point. As seen in the Stress on Links section above, for the full margin tests at a dwell time of two seconds, other stresses could cause a decrease in the amount of margin. Increasing the dwell time also increases the confidence level of the test, which yields a decrease in the amount of margin. This can be compared to the hotter region of the oscilloscope capture. There seems to be a saturation point or leveling off in which an increase in dwell time has no effect on the test results. This can be compared to the colder region of the oscilloscope capture.

In comparing the full margin test results with the oscilloscope results, there are some dissimilarities:

- The full margin test looks at all the lanes and the oscilloscope only looks at one lane.
- The full margin test is in loopback mode so essentially it is looking at two sets of differential pairs, the oscilloscope is only looking at one set of differential pairs.
- The full margin test actually looks at what the receiver sees, and the oscilloscope is far from the receiver. While both show adequate margin on this link, the shapes are different between the oscilloscope and the full margin test.

Taking all the dissimilarities into account there are also multiple similarities.

- Both the oscilloscope capture and the full margin test are statistical in nature.
- Both show a saturation point, with the approach to the saturation point having a higher confidence level than an more open eye.
- Both show lots of margin under the same conditions.

Note there are limitations to using the oscilloscope:

- It is impossible for an oscilloscope to take measurements of all 20 lanes at once. Viewing the stress on all lanes at the same time shows the actual stress of each lane while under test. The IBIIST data is more complete.
- The oscilloscope capture in Figure 11 was taken on the vias on the back side of the motherboard. The signal passes through multiple layers of the board, to a socket, to the balls of the processor, through the processor package, to the receiver on the die of the processor. This is several layers and disconnects away from the receiver. The IBIIST tests are testing the signal the receiver actually sees.
- There are mechanical limitations to applying the oscilloscope probes. For the system used in this experiment, there was a solid metal plate over the vias that needed to be removed before the oscilloscope probes could be added. Modifications needed to be done to reattach the processor heat sink for testing. When using JTAG during development, a header is already on the board. No modifications need to be made to the board.

It was noted that when using the oscilloscope, there was sensitivity to the noise from an unshielded system ten feet away. The unshielded system caused noise on the oscilloscope even when the system under test was turned off. While the noisy lab system was on, accurate test measurement could not be taken with an oscilloscope and the BERT test could not be run. The probes were acting as an antenna and injecting this noise to the vias. Once the noisy lab system was turned off, the IBIIST tests could be run and accurate measurements could be taken.

Summary

The major variable of an IBIIST margin test is the characteristics of the devices attached to the platform whether socketed or soldered onto the board. Through testing it was shown that the amount of margin, the training values, and the margin lane failures all follow the device.

To a platform validation engineer this is important since the only control they have is the platform under test. A basic assumption needs to be made that the devices that are used in the platform test will meet the margin requirements given by the device vendors. It is up to the platform engineer to make certain that the platform does not infringe upon this margin.

While IBIIST testing is statistical in nature, the testing has shown to be repeatable within reasonable limits given the same conditions for the test. There were however several factors that could affect the results of the test margins. The factors explored were

odd/even samplers, equalization settings, devices, interconnect speed, patterns, dwell time, concurrency, or a combination of stresses. Each of these factors can affect the test results and needs to be understood by the platform validation engineer.

References

¹ <http://edageek.com/2007/10/23/intel-ibist-intertech/>

² Maxim Integrated Products, Application Note 1095:HFTA-05.0: "Statistical Confidence Levels for Estimating BER Probability", October 26, 2000