The University of New Hampshire
InterOperability Laboratory

10GEC THE 10 GIGABIT ETHERNET CONSORTIUM

XAUI Electrical Test Suite

Version 1.1

Technical Document

Last Updated: February 4, 2003 3:20 AM

10 Gigabit Ethernet Consortium
InterOperability Laboratory
Research Computing Center
University of New Hampshire

121 Technology Drive, Suite 2
Durham, NH 03824
Phone: (603) 862-0205
Fax: (603) 862-0898
http://www.iol.unh.edu/consortiums/10gec

MODIFICATION RECORD

- December 4, 2002 Version 0.1 Released
- January 15, 2003 Version 1.0 Released
- February 4, 2003 Version 1.1 Released
  - Minor formatting changes
  - Fixed standards reference to read '802.3ae-2002', not '802.3-2002'

ACKNOWLEDGMENTS

The University of New Hampshire would like to acknowledge the efforts of the following individuals in the development of this test suite.

Andy Baldman UNH InterOperability Lab

INTRODUCTION

Overview

The University of New Hampshire's InterOperability Laboratory (IOL) is an institution designed to improve the interoperability of standards based products by providing an environment where a product can be tested against other implementations of a standard. This suite of tests has been developed to help implementers evaluate the functioning of their Clause 47 XAUI-based products. The tests do not determine if a product conforms to the IEEE 802.3 standard, nor are they purely interoperability tests. Successful completion of all tests contained in this suite does not guarantee that the tested device will operate with other 10Gb/s capable devices. However, combined with satisfactory operation in the IOL's interoperability test bed, these tests provide a reasonable level of confidence that the Device Under Test (DUT) will function well in many 10Gb/s environments.

Organization of Tests

The tests contained in this document are organized to simplify the identification of information related to a test and to facilitate in the actual testing process. Each test contains an identification section that describes the test and provides cross-reference information. The discussion section covers background information and specifies why the test is to be performed. Tests are grouped in order to reduce setup time in the lab environment. Each test contains the following information:

Test Number

The Test Number associated with each test follows a simple grouping structure. Listed first is the Test Group Number followed by the test's number within the group. This allows for the addition of future tests to the appropriate groups of the test suite without requiring the renumbering of the subsequent tests.

Purpose

The purpose is a brief statement outlining what the test attempts to achieve. The test is written at the functional level.

References

The references section lists cross-references to the IEEE 802.3 standards and other documentation that might be helpful in understanding and evaluating the test and results.

Resource Requirements

The requirements section specifies the hardware, and test equipment that will be needed to perform the test. The items contained in this section are special test devices or other facilities, which may not be available on all devices.

Last Modification

This specifies the date of the last modification to this test.

Discussion

The discussion covers the assumptions made in the design or implementation of the test as well as known limitations. Other items specific to the test are covered here.

Test Setup

The setup section describes the configuration of the test environment. Small changes in the configuration should be included in the test procedure.

Procedure

The procedure section of the test description contains the step-by-step instructions for carrying out the test. It provides a cookbook approach to testing, and may be interspersed with observable results.

Observable Results

The observable results section lists observables that can be examined by the tester to verify that the DUT is operating properly. When multiple values are possible for an observable, this section provides a short discussion on how to interpret them. The determination of a pass or fail for a certain test is often based on the successful (or unsuccessful) detection of a certain observable.

Comments

This section contains a description of known issues with the test procedure, which may affect test results in certain situations. It may also provide references to Test Suite Appendices that may provide more detailed information regarding any potential issues.

TABLE OF CONTENTS

GROUP 1: Transmitter Verification

Overview: The following group of tests pertains to the operation of the transmitter and the determination of various parametric values as defined in Std. IEEE 802.3ae-2002. Note, successfully passing these tests, or failing these tests does not necessarily indicate that the device under test will, or will not, be interoperable. Devices that pass these tests are more inclined to be interoperable with, not only existing products, but also all future standard compliant devices.

Test 47.1.1 - XAUI Baud

Purpose: To verify that the baud rate of the device under test (DUT) is within the conformance limits specified in 47.3.3 of IEEE 802.3.

References:

[1] IEEE Std. 802.3ae - 2002 Edition, Clause 47 (multiple subclauses).

Resource Requirements:

- Digital storage oscilloscope (DSO), 6GHz minimum bandwidth
- Two 50-ohm coaxial cables (24" or shorter)
- Post processing software

Last Modification: December 3, 2002 (Version 1.0)

Discussion:

The 10 Gigabit Attachment Unit Interface (XAUI) is defined in clause 47 of IEEE Std 802.3-2002. Section 47.3.3 defines the driver characteristics, the first of which is the XAUI Baud. The Baud is the transmitted symbol rate of a given device, and is specified in clause 47.3.3 to be 3.125Gbaud +/- 100ppm. This translates to 3.125Gbaud +/- 312.5KBaud, with a nominal Unit Interval (UI) of 320ps. This rate is to be verified for each of the four lanes of XAUI signaling for a particular DUT, which are identified as Lane 0, Lane 1, Lane 2, and Lane 3.

Test Setup:

Use the coaxial cables to connect the DUT's Lane 0 differential signal to channels 2 and 3 of the DSO. Channels 2 and 3 will be subtracted via post-processing to create the differential signal. Ensure that all unused DUT input and output ports are properly terminated.

Procedure:

1. Configure the DUT so that it is sourcing either the CJPAT or IDLE test pattern.
2. Configure the DSO to capture at least 100ms of waveform data.
3. Process the waveform data in blocks of 10669 samples each, recovering a nominal Baud rate for each block.
4. Compute the average Baud over all blocks.
5. Repeat Steps 1 through 4 for Lanes 1, 2, and 3.

Observable Results:

a. The average Baud rate over all blocks in a given Lane shall be 3.125Gbaud +/- 312.5KBaud.
b. The measured result for each Lane must fall within the conformance limits.

Comments:

A more detailed discussion of the post-processing algorithms used for this test can be found in Appendix 47.A of this Test Suite. The discussion also includes an explanation of the rationale behind the use of the 10669-sample block length.

Test 47.1.2 – Driver Output Amplitude

Purpose: To verify that the driver differential output amplitude of the device under test (DUT) is within the conformance limits specified in clause 47.3.3.2 of IEEE 802.3.

References:

[1] IEEE Std. 802.3ae - 2002 Edition, Clause 47 (multiple subclauses).

Resource Requirements:

- Digital storage oscilloscope (DSO), 6GHz minimum bandwidth
- Two 50-ohm coaxial cables (24" or shorter)
- Post processing software

Last Modification: December 3, 2002 (Version 1.0)

Discussion:

The 10 Gigabit Attachment Unit Interface (XAUI) is defined in clause 47 of IEEE Std 802.3-2002. Section 47.3.3 defines the driver characteristics, and subclause 47.3.3.2 specifies the Differential Output Amplitude and Swing. Because the XAUI receiver is intended to be AC-coupled, DC-referenced logic levels are not defined, but rather separate specifications are provided for both the absolute driver voltages in addition to the differential peak-to-peak amplitude. Figure 47-3 of Clause 47, reproduced below, illustrates the absolute driver voltage limits and the definition of the differential peak-to-peak amplitude.

This test verifies the differential peak-to-peak amplitude conformance, while Test 47.1.3 verifies the conformance for the individual single-ended signal halves of the differential signal. For this test, the differential peak-to-peak amplitude shall be less than or equal to 1600mV. This value is to be verified for each of the four lanes of XAUI signaling for a particular DUT, which are identified as Lane 0, Lane 1, Lane 2, and Lane 3.

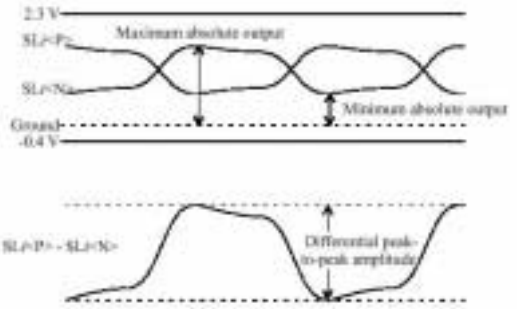

Figure 47–3—Driver output voltage limits and definitions. L*i*<P> and L*i*<N> are the positive and negative sides of the differential signal pair for Lane *i* (*i* = 0, 1, 2, 3).

Test Setup:

Use the coaxial cables to connect the DUT's Lane 0 differential signal to channels 2 and 3 of the DSO. Channels 2 and 3 will be subtracted via post-processing to create the differential signal. Ensure that all unused DUT input and output ports are properly terminated.

Procedure:

1. Configure the DUT so that it is sourcing either the CJPAT or IDLE test pattern.
2. Configure the DSO to capture at least 100ms of waveform data.
3. Subtract the two single-ended halves to create the differential signal.
4. Measure the maximum peak-to-peak voltage swing over the entire set of waveform data.
5. Repeat Steps 1 through 4 for Lanes 1, 2, and 3.

Observable Results:

a. The maximum peak-to-peak voltage swing for each Lane shall be less than or equal to 1600mV.
b. The measured result for each Lane must fall within the conformance limits.

Comments: None.

Test 47.1.3 – Driver Output Swing

Purpose: To verify that the single-ended output swing of the device under test (DUT) is within the conformance limits specified in clause 47.3.3.2 of IEEE 802.3.

References:

[1] IEEE Std. 802.3ae - 2002 Edition, Clause 47 (multiple subclauses).

Resource Requirements:

- Digital storage oscilloscope (DSO), 6GHz minimum bandwidth
- Two 50-ohm coaxial cables (24" or shorter)
- Post processing software

Last Modification: December 3, 2002 (Version 1.0)

Discussion:

The 10 Gigabit Attachment Unit Interface (XAUI) is defined in clause 47 of IEEE Std 802.3-2002. Section 47.3.3 defines the driver characteristics, and subclause 47.3.3.2 specifies the Differential Output Amplitude and Swing. Because the XAUI receiver is intended to be AC-coupled, DC-referenced logic levels are not defined, but rather separate specifications are provided for both the absolute driver voltages in addition to the differential peak-to-peak amplitude. Figure 47-3 of Clause 47, reproduced below, illustrates the absolute driver voltage limits and the definition of the differential peak-to-peak amplitude.

This test verifies the conformance for the individual single-ended signal halves of the differential signal. For this test, the maximum and minimum absolute amplitudes for both the Sli<P> and Sli<N> signals must be between -0.4V and 2.3V, with respect to signal ground. This value is to be verified for each of the four lanes of XAUI signaling for a particular DUT, which are identified as Lane 0, Lane 1, Lane 2, and Lane 3.

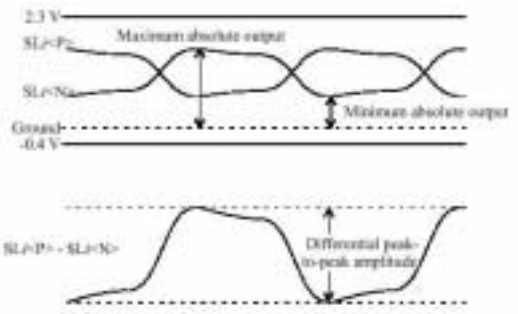

Figure 47–3—Driver output voltage limits and definitions. Li<P> and Li<N> are the positive and negative sides of the differential signal pair for Lane *i* (*i* = 0, 1, 2, 3).

Test Setup:

Use the coaxial cables to connect the DUT's Lane 0 differential signal to channels 2 and 3 of the DSO. Channels 2 and 3 will be subtracted via post-processing to create the differential signal. Ensure that all unused DUT input and output ports are properly terminated.

Procedure:

1. Configure the DUT so that it is sourcing either the CJPAT or IDLE test pattern.
2. Configure the DSO to capture at least 100ms of waveform data.
3. Use the post-processing software to determine the absolute maximum and minimum voltage levels for both the SLi<P> and SLi<N> signals over the entire set of waveform data (i = 0).
4. Repeat Steps 1 through 4 for Lanes 1, 2, and 3.

Observable Results:

a. The maximum and minimum values of Sli<P> and Sli<N> for all four Lanes shall be between -0.4V and 2.3V.

Comments: None.

Test 47.1.4 – Driver Output Impedance

Purpose: To verify that the output impedance (return loss) of the device under test (DUT) is within the conformance limits specified in clause 47.3.3.4 of IEEE 802.3.

References:

[1] IEEE Std. 802.3ae - 2002 Edition, Clause 47 (multiple subclauses).

Resource Requirements:

- Differential Vector Network Analyzer (VNA), 3.125GHz minimum bandwidth
- Two 50-ohm coaxial cables (24" or shorter)
- Post processing software

Last Modification: December 3, 2002 (Version 1.0)

Discussion:

The 10 Gigabit Attachment Unit Interface (XAUI) is defined in clause 47 of IEEE Std 802.3-2002. Section 47.3.3 defines the driver characteristics, and subclause 47.3.3.4 specifies the Differential Return Loss.

As stated in subclause 47.3.3.4, the differential return loss includes contributions from on-chip circuitry, chip packaging, and any off-chip components related to the driver. The output impedance applies to all valid output levels, and the reference impedance for differential return loss measurements is 100 ohms.

For the purpose of this test, return loss is defined as the magnitude of the reflection coefficient expressed in decibels. The reflection coefficient, Γ, is the ratio of the voltage in the reflected wave to the voltage in the incident wave. Note that this is also known as the S_{11} scattering parameter (s-parameter). For frequencies from 312.5 MHz to 3.125 GHz, the differential return loss of the driver shall exceed Equation (47-1):

$$\begin{aligned} S_{11} &= -10 \text{ dB} && (\text{for } 312.5\text{MHz} < f < 625\text{MHz}) \\ &= -10 + 10\log(f/625) \text{ dB} && (\text{for } 625\text{MHz} < f < 3.125\text{GHz}) \end{aligned} \qquad (47\text{-}1)$$

This value is to be verified for each of the four lanes of XAUI signaling for a particular DUT, which are identified as Lane 0, Lane 1, Lane 2, and Lane 3.

Test Setup:

Use the coaxial cables to connect the DUT's Lane 0 differential transmitter to the VNA. Configure the VNA to measure the S_{11} scattering parameter (reflection coefficient). Ensure that all unused DUT input and output ports are properly terminated.

Procedure:

1. Configure the DUT so that it is sourcing either the CJPAT or IDLE test pattern.
2. Calibrate the VNA to remove the effects of the coaxial cables.
3. Measure the reflection coefficient at the transmitter from 312.5 MHz to 3.125GHz
4. Compute the return loss from the reflection coefficient values.
5. Repeat Steps 1 through 3 for Lanes 1, 2, and 3.

Observable Results:

a. For all lanes, the differential return loss shall exceed the limits described by Equation 47-1.

Comments: None.

Test 47.1.6 – Driver Transmit Jitter

Purpose: To verify that the DUT conforms to the jitter requirements specified in clause 47.3.3.5 of IEEE 802.3

References:
[1] IEEE Std. 802.3ae - 2002 Edition, Clause 47 (multiple subclauses).

Resource Requirements:
- Digital storage oscilloscope (DSO), 6GHz minimum bandwidth
- Compliance Interconnect test channel (OPTIONAL)
- Two 50-ohm coaxial cables (24" or shorter)
- Post processing software

Last Modification: December 3, 2002 (Version 1.0)

Discussion:
The 10 Gigabit Attachment Unit Interface (XAUI) is defined in clause 47 of IEEE Std 802.3-2002. Section 47.3.3 defines the driver characteristics, and subclause 47.3.3.5 specifies the Driver Jitter requirements.

Subclause 47.3.3.5 states that the driver "shall satisfy either the near-end eye template and jitter requirements, or the far-end eye template and jitter requirements." The reason for this option is to allow for the use of signal pre-emphasis/de-emphasis in order to compensate for channel losses and improve performance. Two sets of specifications are given, and the designer is allowed to conform to either the near-end or far-end sets of requirements. The table below summarizes the near-end and far-end upper limits for both Total and Deterministic Jitter.

Table 47-5: Transmit Jitter Maximum Limits

	Total Jitter (UI)	Deterministic Jitter (UI)
Near End	+/- 0.175 pk from mean	+/- 0.085 pk from mean
Far End	+/- 0.275 pk from mean	+/- 0.185 pk from mean

Test Setup:
Each differential Lane of the DUT will be connected to channels 2 and 3 of the DSO using the coaxial cables. If a far-end measurement is being made, insert the Compliance Interconnect between the DUT and the DSO. Channels 2 and 3 will be subtracted via post-processing to create the differential signal. Ensure that all unused DUT input and output ports are properly terminated.

Procedure:
1. Configure the DUT so that it is sourcing the CJPAT test pattern.
2. Configure the DSO to capture 1Mpts of waveform data at 20GS/s.
3. Process the waveform data using the post processing software to determine the Total and Deterministic Jitter.
4. Repeat Steps 1 through 4 for Lanes 1, 2, and 3.

Observable Results:
a. For all lanes, the Total and Deterministic Jitter values shall not exceed the limits specified in Table 47-5.

Comments:
A detailed description of the post-processing operation can be found in Appendix 47.A of this Test Suite.

GROUP 2: Receiver Verification

Overview: The following group of tests pertains to the operation of the receiver and the determination of various parametric values as defined in Std. IEEE 802.3ae-2002. Note, successfully passing these tests, or failing these tests does not necessarily indicate that the device under test will, or will not, be interoperable. Devices that pass these tests are more inclined to be interoperable with, not only existing products, but also all future standard compliant devices.

Test 47.2.1 – Receiver Differential-Mode Input Impedance

Purpose: To verify that the receiver differential-mode input impedance (return loss) of the device under test (DUT) is within the limits specified in clause 47.3.4.5 of IEEE 802.3.

References:

[1] IEEE Std. 802.3ae - 2002 Edition, Clause 47 (multiple subclauses).

Resource Requirements:

- Differential Vector Network Analyzer (VNA), 3.125GHz minimum bandwidth
- Two 50-ohm coaxial cables (24" or shorter)
- Post processing software

Last Modification: December 3, 2002 (Version 1.0)

Discussion:

The 10 Gigabit Attachment Unit Interface (XAUI) is defined in clause 47 of IEEE Std 802.3-2002. Section 47.3.3 defines the driver characteristics, and subclause 47.3.4.5 specifies the Receiver Input Impedance.

As stated in subclause 47.3.4.5, the differential return loss includes contributions from on-chip circuitry, chip packaging, and any off-chip components related to the receiver. AC coupling components are also included in this requirement. The reference impedance for differential return loss measurements is specified to be 100 ohms. Subclause 47.3.4.5 specifies that the receiver input impedance shall result in a differential return loss better than 10 dB from 100 MHz to 2.5 GHz.

This value is to be verified for each of the four lanes of XAUI signaling for a particular DUT, which are identified as Lane 0, Lane 1, Lane 2, and Lane 3.

Test Setup:

Use the coaxial cables to connect the DUT's Lane 0 differential receiver to the VNA. Configure the VNA to measure the differential mode S_{11} scattering parameter (reflection coefficient). Ensure that all unused DUT input and output ports are properly terminated.

Procedure:

1. Calibrate the VNA to remove the effects of the coaxial cables.
2. Measure the reflection coefficient at the receiver from 100 MHz to 2.5GHz
3. Compute the return loss from the reflection coefficient values.
4. Repeat Steps 1 through 4 for Lanes 1, 2, and 3.

Observable Results:

a. For all 4 Lanes, the differential return loss shall exceed 10 dB from 100 MHz to 2.5 GHz.

Comments: None.

Test 47.2.2 – Receiver Common-Mode Input Impedance

Purpose: To verify that the receiver common-mode input impedance (return loss) of the device under test (DUT) is within the limits specified in clause 47.3.4.5 of IEEE 802.3.

References:

[1] IEEE Std. 802.3ae - 2002 Edition, Clause 47 (multiple subclauses).

Resource Requirements:

- Differential Vector Network Analyzer (VNA), 3.125GHz minimum bandwidth
- Two 50-ohm coaxial cables (24" or shorter)
- Post processing software

Last Modification: December 3, 2002 (Version 1.0)

Discussion:

The 10 Gigabit Attachment Unit Interface (XAUI) is defined in clause 47 of IEEE Std 802.3-2002. Section 47.3.3 defines the driver characteristics, and subclause 47.3.4.5 specifies the Receiver Input Impedance.

As stated in subclause 47.3.4.5, the common mode return loss includes contributions from on-chip circuitry, chip packaging, and any off-chip components related to the receiver. AC coupling components are also included in this requirement. The reference impedance for common mode return loss measurements is specified to be 25 ohms. Subclause 47.3.4.5 specifies that the receiver input impedance shall result in a common mode return loss better than 6 dB from 100 MHz to 2.5 GHz.

This value is to be verified for each of the four lanes of XAUI signaling for a particular DUT, which are identified as Lane 0, Lane 1, Lane 2, and Lane 3.

Test Setup:

Use the coaxial cables to connect the DUT's Lane 0 differential receiver to the VNA. Configure the VNA to measure the common mode S_{11} scattering parameter (reflection coefficient). Ensure that all unused DUT input and output ports are properly terminated.

Procedure:

1. Calibrate the VNA to remove the effects of the coaxial cables.
2. Measure the reflection coefficient at the receiver from 100 MHz to 2.5GHz
3. Compute the return loss from the reflection coefficient values.
4. Repeat Steps 1 through 4 for Lanes 1, 2, and 3.

Observable Results:

a. For all four Lanes, the common mode return loss shall exceed 6 dB from 100 MHz to 2.5 GHz.

Comments: None.

Appendix 47.A – Post-Processing Methodology

Purpose: To provide a detailed discussion of the post-processing methodologies used in this test suite.

References:

[1] IEEE Std. 802.3ae - 2002 Edition, Clause 47 (multiple subclauses).

Last Modification: January 12, 2003 (Version 1.0)

Discussion:

47.A.1 – Introduction

While each test in this Test Suite defines a specific procedure for the parameter being measured, in reality it is possible to perform several of the tests simultaneously during the post-processing of a single long-duration waveform capture. This is how the custom post-processing software created by the InterOperability Lab's 10Gigabit Ethernet Consortium (IOL 10GEC) operates. The purpose of this Appendix is to provide some of the background information about the software and algorithms used, in order to allow for feedback and discussion of the methods used for performing the tests.

The primary goal of the post processing was to be able to build an eye diagram from the captured data, and fit that eye to the appropriate near or far-end mask. This was to be accomplished in MATLAB using similar techniques to those used in the IOL's 100BASE-TX Fast Ethernet PMD Test Suite, whereby a reference clock is extracted from the data stream, and is then used to slice the waveform to build the eye. An additional result of this operation is that once the reference clock is determined, the jitter on the signal edges in relation to that reference can be determined and statistically analyzed. An FFT of these jitter values can show how the jitter is distributed in the frequency domain, and the nominal frequency of the recovered reference clock can be examined in reference to its conformance limits. These operations are all readily implemented in MATLAB, and can provide a general feel for the quality of the signal being observed. While the mask tests and nominal transmitter frequency values are directly specified in Clause 47, the jitter PSD is not specified in the Standard, and the jitter analysis, as it is currently implemented, is recognized to not be as thorough as the full specification outlined in Clause 47 and Annex 48B. The 10GEC intends to improve the capabilities of this software over time, and comments are welcome.

47.A.2 – Post-Processing Software

The post processing software consists of approximately a dozen core computational functions written in MATLAB, around which a Graphical User Interface (GUI) has been designed that creates an intuitive, easy-to-use system that is capable of performing most of the transmitter-based tests quickly and simply using only a high-speed Digital Storage Oscilloscope (DSO). With that, the bulk of this section will provide an outline of the post-processing operations from a code-level perspective. It is not necessary to know MATLAB in order to understand this section, although one might find it helpful.

In general, the DSO is configured to capture 1 million samples of data per capture on two simultaneous channels at 20 GS/s. Channel 2 is used for the positive half of the differential signal, and Channel 3 is used for the negative half. Each capture is downloaded to the PC running the analysis software, and processed in blocks of 10669 samples each (the reasoning behind this block length will be discussed later). The test duration is only limited by time, as the user can specify the total number of blocks to observe, and the software will continue to download and process additional waveform data until the indicated number of blocks have been analyzed.

For each block, the first step in the processing chain is to extract the reference clock. This is performed by a separate function that performs a linear best fit of the zero-crossing times for a given block of waveform samples. The slope of the best-fit line determines the nominal frequency of the recovered clock, and the y-intercept determines the phase. The reference clock created by the function actually represents the ideal sampling times located at the centers of the recovered UIs. (Note that this is equivalent to shifting the UI boundaries by 90 degrees.) Resampling the waveform at the times indicated by the returned clock would produce the bitstream (after slicing the analog

values.) The function also computes the timing error for each zero crossing in relation to the 'ideal' crossing time, and returns that array of error values along with an additional array that indicates the relative locations of the error values along the x-axis (time axis). These two values (the timing errors themselves, and the locations of the edges that produced those errors) can be used to construct the jitter waveform, which is essentially a plot of the timing error values vs. UI.

At this point, a brief comment must be made regarding block size. For several reasons, the 1MS capture is broken into smaller blocks and processed piecewise. However the results of the processing, specifically the timing/jitter results, will change as a function of this block size. Because the clock recovery operation is performed individually on each block, as the block size decreases, the clock recovery operation will more effectively 'track' low frequency drift that would otherwise show up as jitter if a longer block size is used. So then, the question arises as to what the proper block length for the purposes of these measurements should be. For the answer, the jitter specifications of clause 47.4.3 may be used as a guide. Clause 47.4.3 states that, "For the purpose of jitter measurement, the effect of a single-pole high pass filter with a 3 dB point at 1.875 MHz is applied to the jitter." Considering the jitter in the frequency domain, when looking at the jitter FFT, the *maximum* frequency that can be detected is determined by the Nyquist frequency of the "sample rate" of the jitter values, where this sample rate basically the minimum edge-to-edge time, i.e., the UI width. Thus a nominal UI of 320ps equates to a maximum jitter frequency of (1/2)*(1/320e-12) = 1.5625GHz.

Conversely, the *lowest* FFT frequency point is determined by the length of the waveform, i.e., the number of samples in the data set. The inter-frequency spacing in the frequency domain (df) is equal to 1/(N*dt), where *N* is the number of points and *dt* is the sampling interval. Thus, we can choose the block length so that the minimum detectable FFT frequency is 1.875MHz, as specified by Clause 47.4.3. It turns out that 1667 jitter samples at 320ps spacing results in a minimum frequency of 1/(1667*320e-12) = 1.8746MHz. Thus, 1667 UI's times 6.4 samples per UI (at 20GS/s) results in a block length of 10669 samples. This is the block length that will be used in order to effectively minimize the effects of jitter below 1.875MHz on the overall jitter measurement.

Returning to the post-processing code, a while loop breaks the long waveform into 10669-sample blocks and processes the blocks sequentially. For each block, a reference clock is computed. The recovered clock frequency for this block is stored in an array, and the computed jitter values for the edges are stored in a separate array. Both of these values will be used later. The waveform and clock data is then passed to the eye diagram generating function, which slices the waveform at the clock boundaries and accumulates the eye. The eye is a 2-dimensional histogram, where each UI is broken up into 320 horizontal and 512 vertical bins. The samples of the waveform block are processed UI per UI, and the samples are dealt to the appropriate histogram bins depending on their relation to the nearest edges of the recovered clock. Because the clock edges are 'analog' values and do not necessarily coincide with the sample intervals, the 1-ps-wide horizontal bin width (320 bins/UI) is sufficient to show adequate resolution of the nominal 6.4 samples/UI waveform.

The top-level function continues to download and process the waveform blocks until the specified number of blocks has been analyzed. The results are then displayed in a series of four plots:

The first plot shows the eye pattern itself displayed in color, along with a vertical colorbar showing the relative intensity of the 256-level colormap used. The indicated near-end or far-end mask is also shown on the plot, represented in blue. An additional function determines the number of mask violations (if any) by tallying up the number of waveform pixels that fall outside the specified mask. This value is recorded and displayed in the GUI.

The second plot is a 256-bin histogram of the observed timing error values that were measured for each of the blocks and accumulated into one large array. Depending on the nature of the jitter for a particular device, the shape of this histogram will tend to be either Gaussian if the jitter is purely random, or bi- or multi-modal if there is significant deterministic jitter. Currently, the peak-to-peak value of the total jitter histogram is recorded, and a crude attempt to locate the peaks of the deterministic components is made, which is used to determine the positive and negative limits of the deterministic jitter component (DJ).

The third plot is a plot of the recovered "TX Clock Frequency Deviation" for each processed block. This is a plot showing the recovered clock frequency value computed for each block, however because the specification is defined as a relatively large value (3.125 GHz), with a fairly small range (+/- 100ppm, or 312.5MHz), it is typically

simpler to consider the clock frequency in terms of the deviation in MHz from 3.125 GHz. This is what this plot shows. Note that this is essentially a representation of the low frequency wander present in the recovered transmit clock. The values per block are shown in blue, and the mean value across all blocks is indicated in red.

The final plot is a sample PSD of the jitter waveform for one waveform block. The jitter waveform for the last processed block is used (for no particular reason other than the fact that that block of data still exists after the while loop has exited.) Because the jitter waveform is a representation of the timing error values spaced appropriately in time, the magnitude of the FFT of these values shows the distribution of the jitter in the frequency domain. In general, the plot appears as a particular noise floor (the random jitter), with spikes at multiples of some fundamental frequency (typically 312 MHz or so), which represent the deterministic jitter. Although not visible in the reports, a zoom-in of the MATLAB generated plot shows that the lowest non-DC frequency component of the FFT is 1.875MHz (plus or minus a small deviation due to the actual recovered clock frequency value for that particular block.)

Summary/Conclusion:

While the capability of the post-processing code will continue to be expanded, it currently exists as a tool that can be used to observe and quantify the fundamental timing and amplitude characteristics of a 3.125 Gbps XAUI electrical signal. While there is still the opportunity for further enhancements, the current implementation allows for a quick measurement of the overall signal shape and jitter characteristics using only a high-speed real-time DSO and a handful of post-processing routines written in MATLAB. What began as an exploratory effort into the possibilities and practicalities of using a real-time DSO for multi-gigabit electrical signal characterization has now resulted in a GUI-based tool that performs a majority of the transmitter-related tests required by clause 47 of the XAUI specification. While the argument can easily be made in favor of a higher bandwidth measurement instrument (a common rule of thumb states that the bandwidth of the measurement instrument should be at least 10 times the bandwidth of interest, which would require upwards of 60GHz scope bandwidth based on a risetime of 60ps (see 47.3.3.3) and the Tr=.35/BW rule of thumb), an alternate perspective would be that measurements performed with a lower bandwidth instrument are simply providing stricter limits on the signals being measured. In other words, if a device generates a conformant eye using a lower-bandwidth instrument, one would expect the result to generally look *better* if a higher bandwidth instrument were used – for the most part. Note that this does not necessarily justify using the lower-bandwidth instrument in all cases, however it does allow the user to get a sense for the amount of leeway that could be considered reasonable when looking at an eye that might be borderline conformant.

In addition, the use of a real-time DSO and flexible post-processing allows for characteristics to be observed that might be more difficult to observe using a sampling scope, such as the low-frequency wander characteristics and jitter PSD. The flexible nature of MATLAB allows for easy algorithm modifications and enhancements.

Overall, it is believed that the original goal of the intended effort has been met, and a set of post-processing algorithms has been assembled that performs the basic transmitter test functions defined in clause 47. Input is welcome on any aspect of this Appendix, or any of the post-processing methodology itself. It is the intent of the 10GEC to work with any, and all, interested parties toward test solutions that are of value to the vendor community.

Comments regarding the contents of this Appendix may be directed to Andy Baldman (author) at aab@iol.unh.edu, Bob Noseworthy (ren@iol.unh.edu), or Eric Lynskey (elynskey@iol.unh.edu).

Appendix 47.B – XAUI Compliance Interconnect

Purpose: To provide clarification regarding the interpretation of the XAUI Compliance Interconnect definition specified in Clause 47.4.1 of IEEE 802.3, for the purposes of this test suite.

References:

[1] IEEE Std. 802.3ae - 2002 Edition, Clause 47 (multiple subclauses).

[2] "Characterization of Balanced Digital Components and Communication Paths", Agilent Technologies Application Note #EPSG084729, 2001.

Last Modification: January 6, 2003 (Version 1.2)

Discussion:

47.B.1 – Introduction

The 10 Gigabit Attachment Unit Interface (XAUI) specified in Clause 47 of IEEE 802.3-2002 is an electrical interface that is primarily intended as a chip-to-chip interconnect, implemented with traces on a printed circuit board. While the XGMII is electrically limited to distances of approximately 7cm, the optional XAUI interconnect allows the XGMII to be extended to reach distances of up to approximately 50cm [1].

The XAUI consists of eight 8B/10B encoded serial data streams, each operating at a nominal rate of 3.125 GBaud. Each stream (a.k.a., 'Lane') is a low-swing AC-coupled balanced differential signal running across a pair of PCB traces. Four of the pairs are used in the transmit direction, and four in the receive direction. Since the XAUI is essentially an embedded communications link, the performance of the system ultimately depends both on the integrity of the transmitted signals, as well as the characteristics of the channel.

To this end, Clause 47 includes electrical specifications for XAUI. Conformance limits for traditional signal characteristics such as amplitude, eye mask, and jitter are provided. However, the manner by which these specifications are defined is somewhat non-traditional, primarily due to the wide range of implementations the XAUI may embody. As such, the conformance specifications are potentially a point of confusion if not interpreted properly. This Appendix is intended to identify and clarify these areas and provide a specific interpretation for the purposes of this test suite.

47.B.2 – Eye Template and Jitter Requirements: Near-end vs. Far-end

Clause 47.3.3.5 ('Driver template and jitter') states that the driver "shall satisfy either the near end eye template and jitter measurements, or the far end eye template and jitter measurements." The reasoning behind such a specification was to allow for the two primary methods by which the XAUI transceiver may deal with the task of compensating for signal degradation due to channel losses (which may be particularly excessive for long-reach implementations). While some implementations may choose to employ equalization at the receiving end in order to improve system performance, others may elect to use the technique of signal pre-emphasis/de-emphasis at the transmitter end in order to pre-compensate for the anticipated channel losses, and improve the overall quality of the signal as it appears at the receiver. Since the pre-emphasis/de-emphasis method involves the re-shaping of the transmitted pulse at the near-end, the problem arises such that these distortions may in turn cause the pulse to violate a mask template that was designed for a normal, uncompensated transmitter. Conversely, a mask specification based on the far-end signal is equally problematic, in that a device choosing to implement receive equalization could, in theory, be perfectly able to tolerate lower quality signals that easily violate a far-end mask designed for pre-emphasis/de-emphasis based systems.

In order to accommodate both design types, the solution that was ultimately chosen was to provide two mask templates, and require conformance to only one of the two. A near-end template is defined which is intended for devices that do not use transmitter pre-emphasis/de-emphasis, and a far-end template is defined for those that do.

(Note that these are only guidelines suggested by the standard, not specific requirements. Either type of device may realistically be verified to either mask.)

47.B.3 – Compliance Interconnect Definition

One consequence of providing a far-end mask template specification is that in order for it to have any practical meaning, one must also say something about the channel. Clause 47.4.1 and Figure 47-6 specify the characteristics for a test channel that is referred to in the clause as a "Compliance Interconnect". Some ambiguity could potentially be seen here, as it is uncertain upon first glance if this specification is provided purely for informative purposes, or if operation across this channel is required for conformance. Because of the lack of any "shall" statements associated with Clause 47.4.1, one would tend to think the former, however the fact that the title of Clause 47.4 is "Electrical Measurement Requirements" would suggest that the interconnect specification is in fact normative.

The Compliance Interconnect is a key component of several tests in Clause 47. First, it is required for measuring the far-end eye (see 47.4.2), and the far-end jitter (see 47.4.3.1). (Note that technically, this could be considered an "optional requirement", since a device may alternately conform to the near-end mask and jitter specifications, and still be considered compliant.) However, the Compliance Interconnect is also used to calibrate the reference input signals used for performing the jitter tolerance test specified in 47.4.3.2, which is a test that pertains to *all* devices. Thus, it may be safely stated that a test channel that conforms to 47.4.1 and Figure 47-6 is to be considered a requirement for the purposes of this test suite.

This being said, there is also the potential for confusion regarding the interpretation of the limit line shown in Figure 47-6, reproduced below for convenience.

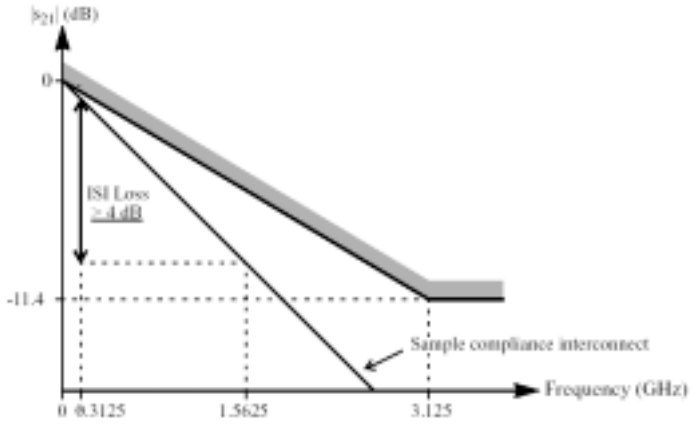

Figure 47–6—Compliance interconnect magnitude response and ISI loss.

Equation 47-2 in Clause 47.4.1 states that the transmission magnitude response, $|S_{21}|$, of the Compliance Interconnect in dB, should satisfy the following:

[EQ 47-2] $$|s_{21}| \le |s_{21}|_{limit} = -20\log(e) \times [a_1\sqrt{f} + a_2 f + a_3 f^2]$$

where *f* is the frequency on Hz, $a_1 = 6.5x10^{-6}$, $a_2 = 2.0x10^{-10}$, and $a_3 = 3.3x10^{-20}$. This limit applies from DC to 3.125 GHz, and the magnitude above 3.125GHz shall not exceed a flat level of -11.4dB.

Because of the use of a negative dB scale, in addition to the fact that a lossy channel is being described in terms of the S_{21} transmission response parameter (which is often thought of as a gain), it may be unclear to the casual reader whether or not the limit line specified by Equation 47-2 is intended to be treated as the upper bound on the channel attenuation, or the lower bound. Should a valid compliance interconnect be *more* lossy than the limit line, or less? In general, channel specifications typically provide a limit line for *attenuation*, which represents the maximum allowed attenuation value, i.e., the channel attenuation must be less than the limit line in order to be considered conformant. Losses are usually indicated in dB, and increase positively along the y-axis. However, we see that this is not the case for Figure 47-6. We also see that the response for a "Sample compliance interconnect" is provided on the plot, and that this response is below the limit line, however the y-axis is labeled in negative dB. What does this mean?

A closer analysis of Figure 47-6, combined with a review of s-parameters and the decibel, provides us with the answer. By definition, the S_{21} transmission response represents the forward *gain* of the system, Pout/Pin, expressed in decibels. Values greater than 0 dB represent a positive gain (i.e., Pout/Pin > 1), while values less than 0 dB represent a "negative gain" (or in other words, a positive loss). Thus, a "gain" of -11.4 dB is equal to a power multiplication factor of $10^{(-11.4/10)} = .0724$. Thus, a gain of -12.4 dB (i.e., below the line) would correspond to a power multiplication factor of $10^{(-12.4/10)} = .0575$. By stating that the magnitude response of the Compliance Interconnect must be *less* than the limit line from Figure 47-6, the standard implies that a valid test channel must have *at least* as much loss as the value shown by the limit line of Equation 47-2. Channels with greater loss are additionally considered valid for the purpose of conformance testing. Another way to think of this would be to flip the limit line in Figure 47-6 upside down (to represent *loss* instead of gain), re-label the y-axis from 0 to +11.4 dB, and state that the attenuation for a valid Compliance Interconnect must be above the limit line.

It is relevant to note here that the 'minimum' value for the limit line was determined empirically, by physically measuring the performance of a range of real, actual interconnects ranging in length from 46 to 56 cm, and taking the median of the observed data (see 47.3.5). It appears that the intent was to provide a baseline value for the expected "worst case" channel, which could serve as a reference point for designers.

47.B.4 – Verification of Compliance Interconnect

For the purposes of this test suite, it is thus required that a valid Compliance Interconnect be used, in order for performing both the far-end eye template measurements, as well as the calibration of the reference input signals used for testing the XAUI receiver. This test channel may be provided by the designer of the DUT, or a default test channel provided by the IOL may be used instead. In either case, the channel must be verified to provide *at a minimum*, the amount of attenuation required by 47.4.1/Figure 47-6. The channel may be even more severe, provided the signal at the far end can still meet the far-end driver requirements of 47.3.3.

In theory, various methods may be used to verify the response of the channel. For this test suite, the recommendation is that a fully differential Vector Network Analyzer (VNA) be used, such as the Agilent E5071B ENA Series Analyzer, or equivalent. A fully differential measurement device can allow for both common and differential mode s-parameter measurements on a given DUT. This type of tool provides the most accurate characterization of the response characteristics for differential systems. (For a more thorough discussion on the drawbacks of using TDR and multiport single-ended measurements for the purposes of characterizing differential systems, see Reference [2]). The channel that will be used for the tests covered in this suite must be verified prior to running the tests. For the purposes of this test suite, it will be considered standard practice to verify the test channel and provide a plot of the response every time a device is tested. This will allow for cases where the device vendor wishes to provide their own Compliance Interconnect for use with a particular DUT. If the device vendor does not wish to provide his or her own Compliance Interconnect, a default conformant channel will be provided. The 10GEC provides an assortment of sample test channels that have various attenuation and crosstalk characteristics, and is always open to accepting new test channels of any form; however only those channels that meet the requirements of 47.4.1/Fig. 47-6 shall be made available for selection as a Compliance Interconnect.

Summary/Conclusion:

This discussion was meant to formally state the 10GEC's interpretation of the Compliance Interconnect definition as it pertains to the XAUI conformance testing performed at the UNH InterOperability Lab. The minimum requirements for the Compliance Interconnect were discussed. In addition, statements were presented supporting the stance that the Compliance Interconnect is an essential and required component of XAUI testing. Comments regarding any of the contents of this appendix may be directed to Andy Baldman (aab@iol.unh.edu), Bob Noseworthy (ren@iol.unh.edu), or Eric Lynskey (elynskey@iol.unh.edu).