University of New Hampshire InterOperability Laboratory
Ethernet Consortium

Superseded

As of November 22nd, 2004 the Gigabit Ethernet Consortium Clause 40 Physical Medium Attachment Conformance Test Suite version 2.0 has been superseded by the release of the Clause 40 Physical Medium Attachment Conformance Test Suite version 2.1. This document along with earlier versions, are available on the Ethernet Consortium test suite archive page.

Please refer to the following site for both current and superseded test suites:
http://www.iol.unh.edu/testsuites/ethernet/archive.html

GIGABIT ETHERNET CONSORTIUM

Clause 40
Physical Medium Attachment (PMA) Test Suite
Version 2.0

Technical Document

Last Updated: 31 March 2004 2:39 PM

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

121 Technology Drive, Suite 2
Durham, NH 03824
Phone: (603) 862-0166
Fax: (603) 862-4181
http://www.iol.unh.edu/consortiums/ge

TABLE OF CONTENTS

The University of New Hampshire
InterOperability Laboratory

MODIFICATION RECORD

- March 31, 2004 (Version 2.0)
  Jon Beckwith: Added Tests 40.1.5, 40.2.1 and Appendices B-E.

- Sep 19, 2003 (Version 1.2)
  Mostly formatting changes, plus one technical typo fix
  Andy Baldman: Updated cover page to include consortium name, full test suite name, and new IOL logo
  Reorganized document to put Table of Contents first
  Revised and reorganized Introduction section
  Changed referencing style to distinguish between internal/external references
  Modified test numbers by removing subclause indicator (e.g. "40.6.1.1" became "40.1.1")
  All references to disturber voltage levels in Appendix 40.A now show correct values

- Jun 18, 2003 (Version 1.1)
  Jon Beckwith: General formatting changes
  Updated references to reflect latest standards
  Added schematics for return loss jig and 8-pin modular breakout board

- Oct 08, 1999 (Version 1.0)
  Initial release

The University of New Hampshire
InterOperability Laboratory

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	University of New Hampshire
Jon Beckwith	University of New Hampshire
Adam Healey	University of New Hampshire
Eric Lynskey	University of New Hampshire
Bob Noseworthy	University of New Hampshire
Matthew Plante	University of New Hampshire
Gary Pressler	University of New Hampshire

INTRODUCTION

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 particular suite of tests has been developed to help implementers evaluate the functionality of the Physical Medium Attachment (PMA) sublayer of their 1000BASE-T products.

These tests are designed to determine if a product conforms to specifications defined in the IEEE 802.3 standard. Successful completion of all tests contained in this suite does not guarantee that the tested device will operate with other 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 properly in many 1000BASE-T environments.

The tests contained in this document are organized in such a manner as to simplify the identification of information related to a test, and to facilitate in the actual testing process. Tests are organized into groups, primarily in order to reduce setup time in the lab environment, however the different groups typically also tend to focus on specific aspects of device functionality. A three-part numbering system is used to organize the tests, where the first number indicates the clause of the IEEE 802.3 standard on which the test suite is based. The second and third numbers indicate the test's group number and test number within that group, respectively. This format allows for the addition of future tests to the appropriate groups without requiring the renumbering of the subsequent tests.

The test definitions themselves are intended to provide a high-level description of the motivation, resources, procedures, and methodologies pertinent to each test. Specifically, each test description consists of the following sections:

Purpose

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

References

This section specifies source material *external* to the test suite, including specific subclauses pertinent to the test definition, or any other references that might be helpful in understanding the test methodology and/or test results. External sources are always referenced by number when mentioned in the test description. Any other references not specified by number are stated with respect to the test suite document itself.

Resource Requirements

The requirements section specifies the test hardware and/or software needed to perform the test. This is generally expressed in terms of minimum requirements, however in some cases specific equipment manufacturer/model information may be provided.

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 initial configuration of the test environment. Small changes in the configuration should not be included here, and are generally covered in the test procedure section, below.

Procedure

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

Observable Results

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

Possible Problems

This section contains a description of known issues with the test procedure, which may affect test results in certain situations. It may also refer the reader to test suite appendices and/or whitepapers that may provide more detail regarding these issues.

GROUP 1: PMA ELECTRICAL SPECIFICATIONS

Overview:

This group of tests verifies several of the electrical specifications of the 1000BASE-T Physical Medium Attachment sublayer outlined in Clause 40 of the IEEE 802.3-2002 standard.

Scope:

All of the tests described in this section have been implemented and are currently active at the University of New Hampshire InterOperability Lab.

Test 40.1.1 - Peak Differential Output Voltage and Level Accuracy

Purpose: To verify correct transmitter output levels.

References:

[1] IEEE Std 802.3-2002, subclause 40.6.1.1.2 - Test modes
[2] Ibid., Figure 40-19 - Example of transmitter test mode 1 waveform
[3] Ibid., subclause 40.6.1.1.3 - Test fixtures
[4] Ibid., subclause 40.6.1.2.1 - Peak differential output voltage and level accuracy

Resource Requirements: Refer to appendix 40.A

Last Modification: September 14, 2003 (version 1.2)

Discussion:

Reference [1] states that all 1000BASE-T devices must implement four transmitter test modes. This test requires the Device Under Test (DUT) to operate in transmitter test mode 1. While in test mode 1, the DUT shall generate the pattern shown in [2] on all four transmit pairs, denoted BI_DA, BI_DB, BI_DC, and BI_DD, respectively.

In this test, the peak differential output voltage is measured at points A, B, C, and D as indicated in [2] while the DUT is connected to test fixture 1 defined in [3]. The conformance requirements for the peak differential output voltage and level accuracy are specified in [4].

Test Setup: Refer to appendix 40.A

Procedure:

1. Configure the DUT so that it is sourcing the transmitter test mode 1 waveform.
2. Connect pair BI_DA from the MDI to test fixture 1.
3. Measure the peak voltage of the waveform at points A, B, C, and D.
4. For enhanced accuracy, repeat step 3 multiple times and average the voltages measured at each point.
5. Repeat steps 2 through 4 for pairs BI_DB, BI_DC, and BI_DD.

Observable Results:

a. The magnitude of the voltages at points A and B shall be between 670 and 820 mV.
b. The magnitude of the voltage at point B and shall not differ from the magnitude of the voltage at point A by more than 1%.
c. The magnitude of the voltage at point C shall not differ from 0.5 times the average of the voltage magnitudes at points A and B by more than 2%.
d. The magnitude of the voltage at point D shall not differ from 0.5 times the average of the voltage magnitudes at points A and B by more than 2%.

Possible Problems: None.

Test 40.1.2 - Maximum Output Droop

Purpose: To verify that the transmitter output level does not decay faster than the maximum specified rate.

References:

[1] IEEE Std 802.3-2002, subclause 40.6.1.1.2 - Test modes
[2] Ibid., Figure 40-19 - Example of transmitter test mode 1 waveform
[3] Ibid., subclause 40.6.1.1.3 - Test fixtures
[4] Ibid., subclause 40.6.1.2.2 - Maximum output droop

Resource Requirements: Refer to appendix 40.A

Last Modification: September 14, 2003 (version 1.2)

Discussion:

Reference [1] states that all 1000BASE-T devices must implement four transmitter test modes. This test requires the Device Under Test (DUT) to operate in transmitter test mode 1. While in test mode 1, the DUT shall generate the pattern shown in [2] on all four transmit pairs, denoted BI_DA, BI_DB, BI_DC, and BI_DD, respectively.

In this test, the differential output voltage is measured at points F, G, H, and J as indicated in [2] while the DUT is connected to test fixture 2 defined in [3]. The conformance requirements for the maximum output droop are specified in [4].

Test Setup: Refer to test suite appendix 40.A

Procedure:

1. Configure the DUT so that it is operating in transmitter test mode 1.
2. Connect pair BI_DA from the MDI to test fixture 2.
3. Measure differential output voltage at points F, G, H, and J.
4. For enhanced accuracy, repeat step 3 multiple times and average the voltages measured at each point.
5. Repeat steps 2 through 4 for pairs BI_DB, BI_DC, and BI_DD.

Observable Results:

a. The voltage magnitude at point G shall be greater than 73.1% of the voltage magnitude at point F.
b. The voltage magnitude at point J shall be greater than 73.1% of the voltage magnitude at point H.

Possible Problems: None.

Test 40.1.3 - Differential Output Templates

Purpose: To verify that the transmitter output fits the time-domain transmit templates.

References:

[1] IEEE Std 802.3-2002, subclause 40.6.1.1.2 - Test modes
[2] Ibid., Figure 40-19 - Example of transmitter test mode 1 waveform
[3] Ibid., subclause 40.6.1.1.3 - Test fixtures
[4] Ibid., Figure 40-6 - Normalized transmit templates as measured at MDI using transmit test fixture 1
[5] Ibid., subclause 40.6.1.2.3 - Differential output templates

Resource Requirements: Refer to appendix 40.A

Last Modification: September 14, 2003 (version 1.2)

Discussion:

Reference [1] states that all 1000BASE-T devices must implement four transmitter test modes. This test requires the Device Under Test (DUT) to operate in transmitter test mode 1. While in test mode 1, the DUT shall generate the pattern shown in [2] on all four transmit pairs, denoted BI_DA, BI_DB, BI_DC, and BI_DD, respectively.

In this test, the differential output waveforms are measured at points A, B, C, D, F, and H as indicated in [2] while the DUT is connected to test fixture 1 defined in [3]. The various waveforms will be compared to the normalized time domain transmit templates specified in [4]. The waveforms around points A and B are compared to normalized time domain transmit template 1 after they are normalized to the peak voltage at point A. The waveforms around points C and D are compared to normalized time domain transmit template 1 after they are normalized to 0.5 times the peak voltage at point A. The waveforms around points F and H are compared to normalized time domain transmit template 2 after they are normalized to the peak voltages at points F and H, respectively.

The waveforms may be shifted in time to achieve the best fit. After normalization and shifting, the waveforms around points A, B, C, D, F, and H shall fit within their corresponding templates, as specified in [5].

Test Setup: Refer to appendix 40.A

Procedure:

1. Configure the DUT so that it is operating in transmitter test mode 1.
2. Connect pair BI_DA from the MDI to test fixture 1.
3. Capture the waveforms around points A, B, C, D, F, and H.
4. For more thorough testing, repeat step 3 multiple times and accumulate a 2-dimensional histogram (voltage and time) of each waveform. This is often referred to as a *persistence waveform*.
5. Normalize the waveforms around points A, B, C, and D and compare them with normalized time domain transmit template 1. The waveforms may be shifted in time to achieve the best fit.
6. Normalize the waveforms around points F and H and compare them with normalized time domain transmit template 2. The waveforms may be shifted in time to achieve the best fit.
7. Repeat steps 2 through 6 for pairs BI_DB, BI_DC, and BI_DD.

Observable Results:

a. After normalization, the waveforms around points A, B, C, and D shall fit within normalized time domain transmit template 1.
b. After normalization, the waveforms around points F and H shall fit within normalized time domain transmit template 2.

Possible Problems: None.

Test 40.1.4 - MDI Return Loss

Purpose: To measure the return loss at the MDI for all four channels

References:

[1] IEEE Std 802.3-2002, subclause 40.8.3.1 - MDI return loss
[2] Ibid., subclause 40.6.1.1.2 - Test modes

Resource Requirements:

- RF Vector Network Analyzer (VNA)
- Return loss test jig
- Post-processing PC

Last Modification: September 14, 2003 (version 1.2)

Discussion:

A compliant 1000BASE-T device shall ideally have a differential impedance of 100Ω. This is necessary to match the characteristic impedance of the Category 5 cabling. Any difference between these impedances will result in a partial reflection of the transmitted signals. Because the impedances can never be exactly 100Ω, and because the termination impedance varies with frequency, some limited amount of reflection must be allowed. Return loss is a measure of the signal power that is reflected due to the impedance mismatch. Reference [1] specifies the conformance limits for the reflected power measured at the MDI. The specification states that the return loss must be maintained when connected to cabling with a characteristic impedance of 100Ω ± 15%, and while transmitting data or control symbols.

Test Setup:

Connect the devices as shown in Figure 40.1.4-1 using the test jig shown in Figure 40.1.4-2.

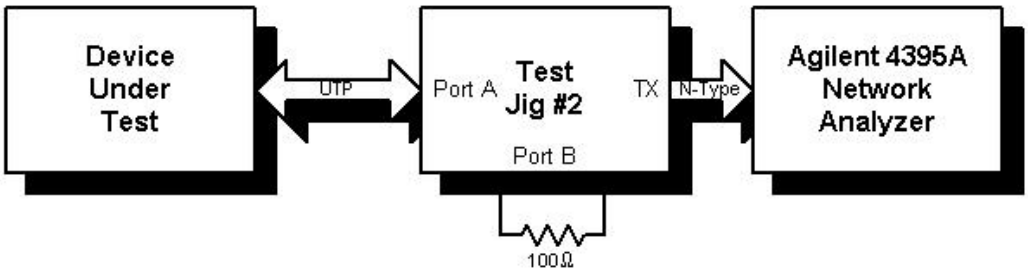

Figure 40.1.4-1: Return loss test setup

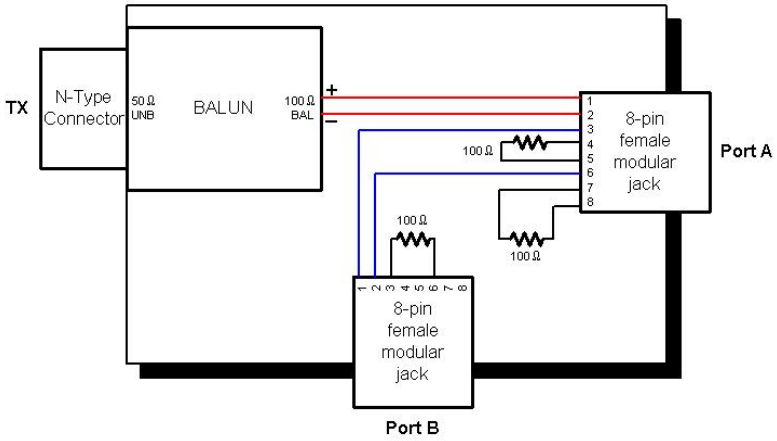

Figure 40.1.4-2: Test Jig #2

Note that Test Jig #2 is a standard jig used by the IOL for various return loss tests. In 100Base-Tx PMD testing, Port B is utilized to send IDLE to the DUT. Here, we do not need to send IDLE to the DUT, and thus, Port B is not used. Also, because the network analyzer is connected to pins 1 and 2 of the 8-pin modular jack, four short UTP cables (approximately 4" long) are needed in order to map the BI_DA, BI_DB, BI_DC, and BI_DD signals from the DUT to the 1-2 pair of the test jig Port A. The effect of each of these cables is removed during calibration of the Network Analyzer.

The specification states that the return loss must be maintained while transmitting data or control symbols. Therefore, it is necessary to configure the DUT so that it is transmitting a signal meeting these requirements. The test mode 4 signal specified in [2] is used in this case to approximate a valid 1000BASE-T symbol stream.

Procedure:

1. Configure the DUT so that it is operating in transmitter test mode 4.
2. Connect the BI_DA pair of the DUT to the reflection port of the network analyzer.
3. Calibrate the network analyzer to remove the effects of the test jig and connecting cable.
4. Measure the reflections at the MDI referenced to a 50Ω characteristic impedance.
5. Post-process the data to calculate the reflections for characteristic impedances of 85 and 115Ω.
6. Repeat steps 2 to 5 for the BI_DB, BI_DC, and BI_DD pairs.

Observable Results:

a. The return loss measured at each MDI pair shall be at least 16 dB from 1 to 40 MHz, and at least $10-20\log_{10}(f/80)$ dB from 40 to 100MHz when referenced to a characteristic impedance of 100Ω ± 15%.

Possible Problems: None.

Test 40.1.5 – Transmitter Timing Jitter, FULL TEST (EXPOSED TX_TCLK)

Purpose: To verify that the DUT meets the jitter specifications defined in Clause 40.6.1.2.5 of IEEE 802.3.

References:

[1] IEEE standard 802.3-2002, subclause 40.6.1.1.1 – Test channel
[2] Ibid., subclause 40.6.1.1.2, figure 40-20 – Test modes
[3] Ibid., subclause 40.6.1.1.3, figure 40-25 – Test fixtures
[4] Ibid., subclause 40.6.1.2.5 – Transmitter Timing Jitter
[5] Test suite appendix 40.6.A – 1000BASE-T transmitter test fixtures

Resource Requirements:

- A DUT with an exposed TX_TCLK clock signal
- A Link Partner device which also provides an exposed TX_TCLK
- Digital storage oscilloscope, Tektronix TDS7104 or equivalent
- (Optional) High-impedance differential probe, Tektronix P6248 or equivalent (2)
- Jitter Test Channel as defined in [1]
- 8-pin modular plug break-out board
- 50 Ω coaxial cables, matched length
- 50 Ω line terminations (6)

Last Modification: March 22, 2002 (Version 1.1)

Discussion:

The jitter specifications outlined in Clause 40.6.1.2.5 define a set of measurements and procedures that may be used to characterize the jitter of a 1000BASE-T device. The clause defines multiple test configurations that serve to isolate and measure different aspects of the jitter in the overall system. While the spec makes distinctions between MASTER mode jitter and SLAVE mode jitter, additional distinctions are made between filtered and unfiltered jitter. Also, there are different timing references by which the jitter is determined depending on the configuration.

For the purpose of this test suite, a step-by-step procedure is outlined that will determine all MASTER and SLAVE mode jitter parameters for a particular DUT. The entire test is separated into three distinct sections in order to minimize test setup complexity and facilitate understanding of the measurement methodology.

The purpose of the first section will be to measure J_{txout} , which is defined as the peak-to-peak jitter on the MDI output signal relative to the TX_TCLK while the DUT is operating in either Test Mode 2 (MASTER timing mode), or Test Mode 3 (SLAVE timing mode). This value is measured for each of the four MDI pairs, BI_DA, BI_DB, BI_DC, and BI_DD for when the DUT is configured as MASTER, and when the DUT is configured as SLAVE. This produces eight J_{txout} values for a particular DUT.

The purpose of the second section will be to measure both the unfiltered and filtered peak-to-peak jitter on the TX_TCLK itself, relative to an “unjittered reference”, while the DUT is configured as the MASTER and is operating under normal conditions (i.e., linked to the Link Partner using a short piece of UTP). While the standard does not provide any further definition for what exactly an “unjittered reference” is or how it is to be derived, for the purposes of this test suite it is to be defined as the straight-line best fit of the zero crossings for any specific capture of the signal under test. Thus, the jitter for any particular edge is defined as the time difference between the actual observed zero crossing time and the corresponding “ideal” crossing time. The setup for this section is relatively straightforward, and is much less complicated than the setup required for the third and final section.

The third and most involved part of the test will measure both the unfiltered and filtered TX_TCLK jitter for the case where the DUT is operating in SLAVE mode. Note that while the MASTER TX_TCLK jitter of the previous section was defined with respect to an “unjittered reference”, the SLAVE TX_TCLK jitter of this section is

instead defined *with respect to the MASTER TX_TCLK*. Thus in order to perform this test, both the DUT and the Link Partner TX_TCLK's must be *simultaneously* monitored with the DSO. In addition, the standard also requires that the DUT and Link Partner be connected by means of the Jitter Test Channel defined in [1], instead of the short piece of UTP used in the previous section.

Note that in order to perform these tests as specified in the standard, it is a requirement that the DUT provide access to the TX_TCLK clock signal (which is not always the case). In addition, the test setup requires a functioning Link Partner device that also provides access to the TX_TCLK. While access to the TX_TCLK signal is relatively straightforward and easy to provide on evaluation boards and prototype systems, it can become quite impractical in more complicated systems. In the case where no exposed TX_TCLK signal is available, it may be possible to perform a simplified version of the full jitter test procedure, which could provide some useful information about the jitter in the system, and possibly verify some subset of the full set of specifications to some degree. Please refer to Appendix 40.6.B for more on this issue.

The full jitter test procedure, in three parts, is presented in the following three sections.

PART I – MASTER/SLAVE Jtxout MEASUREMENTS

Test Set-Up:

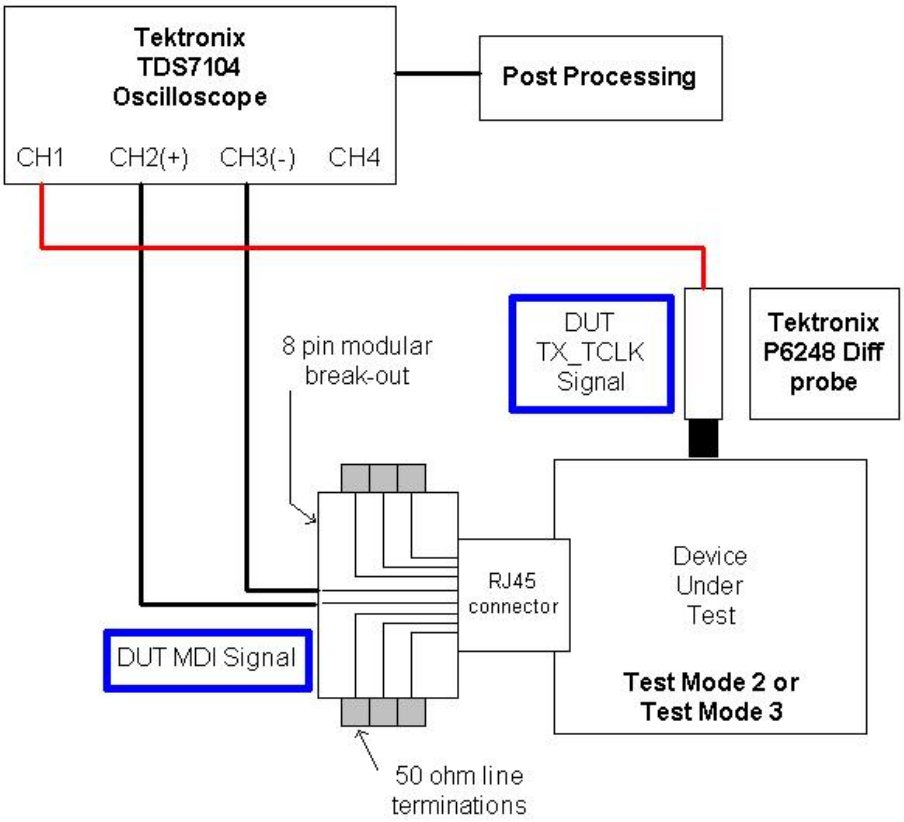

Procedure:

1. Configure the DUT for transmitter Test Mode 2 operation (MASTER timing mode).
2. Connect the TX_TCLK and BI_DA signals to the DSO.
3. Capture 100ms to 1000ms worth of edge data for both the TX_TCLK and BI_DA signals.
4. Compute and record the peak-to-peak jitter on the BI_DA output signal relative to the TX_TCLK.
5. Repeat steps 2, 3, and 4 for pairs BI_DB, BI_DC, and BI_DD.
6. Configure the DUT for Test Mode 3 (SLAVE timing mode), and repeat steps 2 through 5.

Observable Results:

The results of this section will be combined with the results of Parts II and III in order to produce the final pass/fail jitter values. While the 8 values determined here do ultimately affect the final results, no specific pass/fail criteria are assigned to the J_{txout} values themselves.

Possible Problems: None.

PART II – UNFILTERED AND FILTERED TX_TCLK JITTER (MASTER MODE)

Test Set-Up:

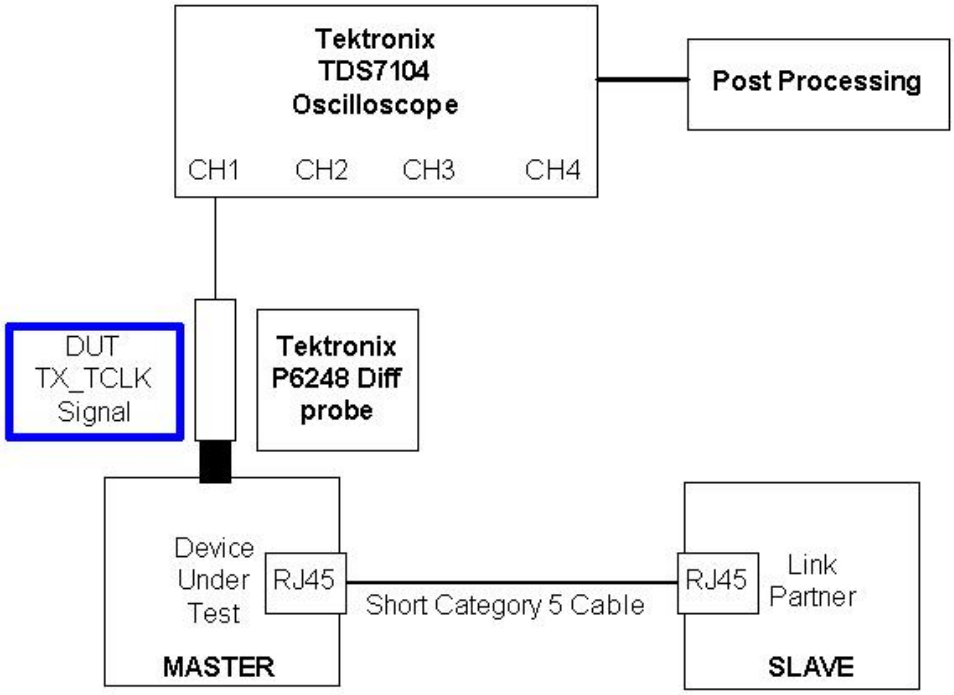

Procedure:

1. Configure the DUT for normal operation in the MASTER timing mode.
2. Configure the Link Partner for normal operation in the SLAVE timing mode.
3. Connect the DUT to the Link Partner using a standard UTP patch cable, and verify that a valid link exists between the two devices.
4. Connect the DUT TX_TCLK signal to the DSO.
5. Capture 100ms to 1000ms worth of TX_TCLK edge data.
6. Compute and record the peak-to-peak jitter on the TX_TCLK relative to an unjittered reference.
7. Pass the sequence of jitter values from Step 6 through a 5KHz high-pass filter, and record the peak-to-peak value of the result. Add to this value the worst pair MASTER J_{txout} value measured in Part I. Record the result.

Observable Results:

The result of Step 6 should be less than 1.4 ns.
The result of Step 7 should be less than 0.3 ns.

Possible Problems:

Clause 40.6.1.2.5 states that, "for all unfiltered jitter measurements, the peak-to-peak value shall be measured over an interval of not less than 100ms and not more than 1 second." In general, it is well beyond the ability of most current DSO's to perform single-shot captures of this length at the sample rates required for this test (1GS/s recommended minimum). To compensate for this, it will generally be necessary to perform multiple captures such that that the total number of observed clock edges satisfies the required limits. In this case, a new "unjittered reference clock" must be computed for each capture in order to measure the jitter. One should note that as the single-shot capture length decreases, the reference clock extraction function (PLL) will be less effective in its

ability to "track" any low frequency modulation in the transmit clock. If a longer duration single-shot capture is possible, these slow variations will show up as jitter. For this test, it is recommended that the DSO be set to utilize the maximum possible single-shot memory depth in order to minimize the impact of this effect.

Note that this issue only pertains to the unfiltered jitter measurements, since the standard requires that all filtered jitter measurements be performed over an unbiased sample of, "at least 10^5 clock edges", which is easily within the single-shot memory depth of most current DSO's.

PART III – UNFILTERED AND FILTERED TX_TCLK JITTER (SLAVE MODE)

Test Set-Up:

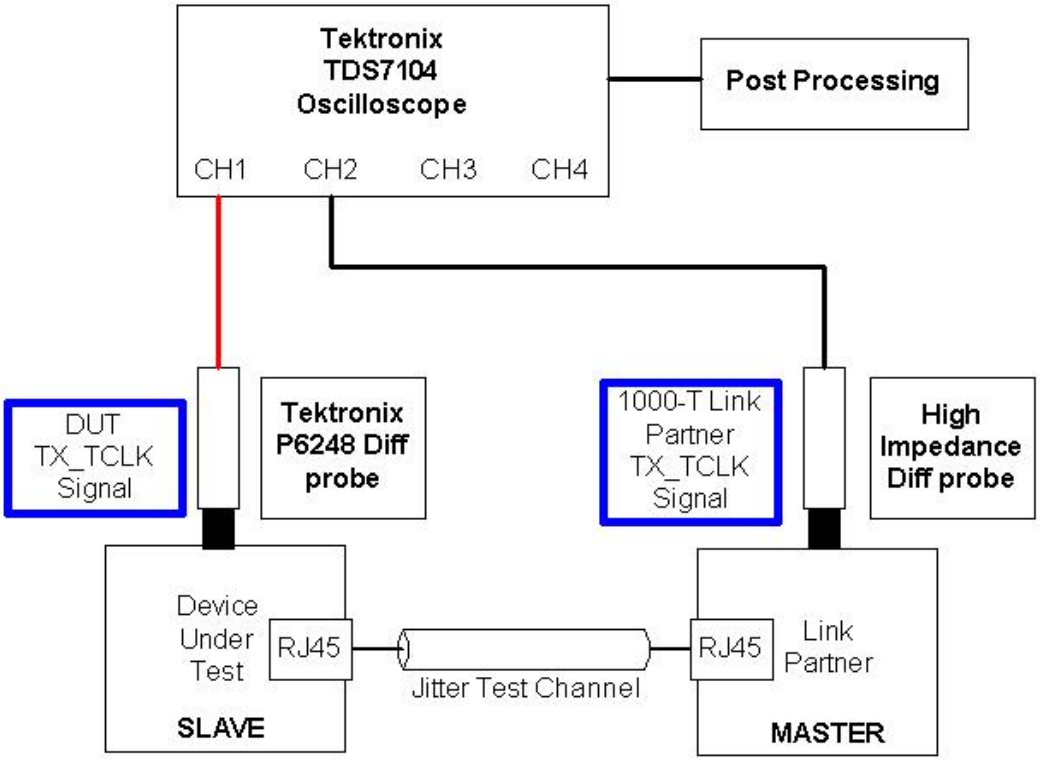

Procedure:

1. Configure the DUT for normal operation in the SLAVE timing mode.
2. Configure the Link Partner for normal operation in the MASTER timing mode.
3. Insert the Jitter Test Channel between the DUT and the Link Partner, oriented such that Port A of the Test Channel is connected to the DUT.
4. Connect both the DUT and the Link Partner TX_TCLK signals to the DSO.
5. Ensure that the DUT is receiving valid data by verifying that the DUT GMII Management Register bit 10.13 is set to 1.
6. Capture 100ms to 1000ms worth of TX_TCLK edge data for both the DUT and Link Partner.
7. Compute the jitter waveform on the Link Partner TX_TCLK, relative to an unjittered reference. Filter this waveform with a 5KHz HPF. Store the peak-to-peak value of the result.
8. Compute the jitter waveform on the DUT TX_TCLK, relative to the Link Partner TX_TCLK. Record the peak-to-peak value.
9. Pass the jitter waveform from Step 8 through a 32KHz HPF, and record the peak-to-peak value of the result. Add to this the worst pair SLAVE mode J_{txout} value from Part I. Subtract the result obtained in Step 7 above. Record the result.

Observable Results:

The result from Step 8 should be less than 1.4 ns.
The result from Step 9 should be less than 0.4 ns.

Possible Problems: (See possible problems discussion from Part II.)

GROUP 2: PMA RECEIVE TESTS

Overview:

This section verifies the integrity of the 1000BASE-T PMA Receiver through frame reception tests.

Scope:

All of the tests described in this section have been implemented and are currently active at the University of New Hampshire InterOperability Lab.

Test 40.2.1 – Bit Error Rate Verification

Purpose: To verify that the device under test (DUT) can maintain low bit error rate in the presence of the worst-case input signal-to-noise ratio.

References:

[1] IEEE Std. 802.3-2002, clause 40
[2] Ibid, Clause 40.4.2.3, PMA Receive Function
[3] Ibid, Clause 40.7, Link Segment Characteristics
[4] Ibid, Clause 40.6, PMA Electrical Specifications

Resource Requirements:

- Transmit station capable of producing a worst case signal
- Category 5e cable plants
- Monitor

Last Modification: January 9, 2004 (Version 1.0)

Discussion:

The operation of the 1000BASE-T PMA sublayer is defined in [1], to operate with a bit error rate of 10^{-10}, as specified in [2], over a worst case channel, as defined in [3]. This test shall verify a 10^{-11} bit error rate, as is done in 100Base-Tx PMD. The results from the 10^{-11} bit error rate test are informative. If the DUT is unable to meet this BER, 10^{-10} is performed to verify compliance.

Based on the analysis given in appendix 40.F, if more than 7 errors are observed in $3x10^{11}$ bits (about 24,700,000 1,518-byte packets), it can be concluded that the error rate is greater than 10^{-11} with less than a 5% chance of error. Note that if no errors are observed, it can be concluded that the BER is no more than 10^{-11} with less than a 5% chance of error.

The transmit station is configured to transmit the worst case rise time and output amplitude, while still meeting the requirements set in [4]. Two worst-case scenarios are utilized. A slow rise time of 5.12ns creates worst-case quantization error; a fast rise time of 4.61ns maximizes the signal bandwidth. Both of the transmit settings utilize the lowest transmit amplitude possible. The electrical specifications for these transmit conditions are provided in Appendix 40.C. Rise time estimation is determined using the techniques described in Appendix 40.D.

Note that in the cases where specific equipment models are specified, any piece of equipment with similar capabilities may be substituted. For multiple port devices, note that the length of the unshielded twisted pair (UTP) cable used to connect to the monitor station should be kept as short as possible (less than a foot). If longer lengths are necessary, the impact of the cable on the measurement must be evaluated and steps taken to remove its effect.

Test Setup:

Connect the transmit station to the DUT across minimum and maximum attenuation cable plants as shown in figure 40.2.1-1.

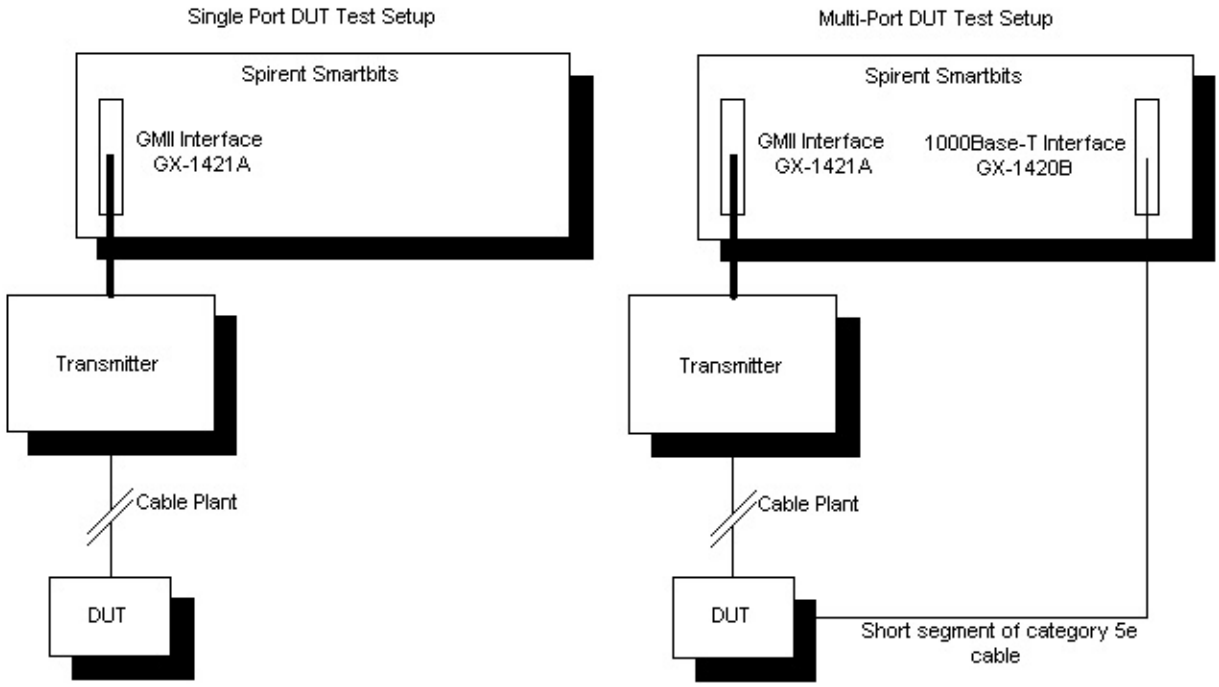

Figure 40.2.1-1: Receiver Test Setups

Procedure:

1. Configure the transmit station such that it generates the slowest worst-case rise time and output amplitude, while maintaining the minimum electrical requirements discussed in [4].
2. The test station shall send 24,700,000 1,518-byte packets (for a 10^{-11} BER) and the monitor will count the number of packet errors.
3. Repeat steps 1 through 2 for the fastest worst-case rise time.

Observable Results: There shall be no more than 7 errors for any iteration.

Possible Problems:

If a device fails to meet the 10^{-11} BER criteria, the BER can be reduced to 10^{-10} to verify conformance to the IEEE 802.3 Standard. Thus, reducing the number of frames sent to 2,470,000. In other cases, the rate at which the device under test can process incoming packets may make the test duration prohibitive. In such cases, fewer packets may be sent resulting in a lower confidence that a bit error rate of 10^{-10} is being met.

TEST SUITE APPENDICES

Overview:

The appendices contained in this section are intended to provide additional low-level technical details pertinent to specific tests defined in this test suite. Test suite appendices often cover topics that are beyond the scope of the standard, but are specific to the methodologies used for performing the measurements covered in this test suite. This may also include details regarding a specific interpretation of the standard (for the purposes of this test suite), in cases where a specification may appear unclear or otherwise open to multiple interpretations.

Scope:

Test suite appendices are considered informative, and pertain only to tests contained in this test suite.

Appendix 40.A - 1000BASE-T Transmitter Test Fixtures

Purpose: To provide a reference implementation of test fixtures 1 through 4

References:

[1] IEEE Std 802.3-2002, subclause 40.6.1.1.3 - Test fixtures
[2] Ibid., Figure 40-22 - Transmitter test fixture 1 for template measurement
[3] Ibid., Figure 40-23 - Transmitter test fixture 2 for droop measurement
[4] Ibid., Figure 40-24 - Transmitter test fixture 3 for distortion measurement
[5] Ibid., Figure 40-25 - Transmitter test fixture 4 for jitter measurement

Resource Requirements:

- Disturbing signal generator, Tektronix AWG2021 or equivalent
- Digital storage oscilloscope, Tektronix TDS7104 or equivalent
- Vector Network Analyzer, HP 8753C or equivalent
- Spectrum analyzer, HP 8593E or equivalent
- Vector Network Analyzer, HP 8712B or equivalent
- Power splitters, Mini-Circuits ZSF-2-1W or equivalent (2)
- 8-pin modular plug break-out board
- 50 Ω coaxial cables, matched length (3 pairs)
- 50 Ω line terminations (6)

Last Modification: September 14, 2003 (version 1.2)

Discussion:

40.A.1 - Introduction

References [1] through [5] define four test fixtures to be used in the verification of 1000BASE-T transmitter specifications. The purpose of this appendix is to present a reference implementation of these test fixtures.

In test fixtures 1 through 3, the Device Under Test (DUT) is directly connected to a 100Ω differential voltage generator. The voltage generator transmits a sine wave of specific frequency and amplitude, which is referred to as the disturbing signal, V_d. An oscilloscope monitors the output of the DUT through a high impedance differential probe. The three test fixtures differ only in the specification of the disturbing signal and the inclusion of a high pass test filter. The test fixture characteristics are given in Table 40.A-1.

Table 40.A-1: Characteristics of test fixtures 1 through 3

Test Fixture	V_d Amplitude	V_d Frequency	Test Filter
1	2.8 V peak-to-peak	31.25 MHz	Yes
2	2.8 V peak-to-peak	31.25 MHz	No
3	5.4 V peak-to-peak	20.83 MHz	Yes

The purpose of V_d is to simulate the presence of a remote transmitter (1000BASE-T employs bi-directional transmission on each twisted pair). If the DUT is not sufficiently linear, the disturbing signal will cause significant distortion products to appear in the DUT output. Note that while the oscilloscope sees the sum of the V_d and the DUT output, only the DUT output is of interest. Therefore, a post-processing block is required to remove the disturbing signal from the measurement.

Upon looking at the diagrams shown in [2], [3], and [4], it is important to note that V_d is defined as the voltage *before* the 50Ω resistors. Thus, the amount of voltage seen at the transmitter under test is 50% of the original amplitude of V_d.

In test fixture 4, the DUT is directly connected to a 100Ω resistive load. Once again, the oscilloscope monitors the DUT output through a high impedance differential probe.

This appendix describes a single test setup that can be used as test fixtures 1 through 4. A block diagram of this test setup is shown in Figure 40.A-1, and the modular break out board used is shown in Figure 40.A-2. Each test fixture is realized through the settings of the disturbing voltage generator and configuration of the post-processing block.

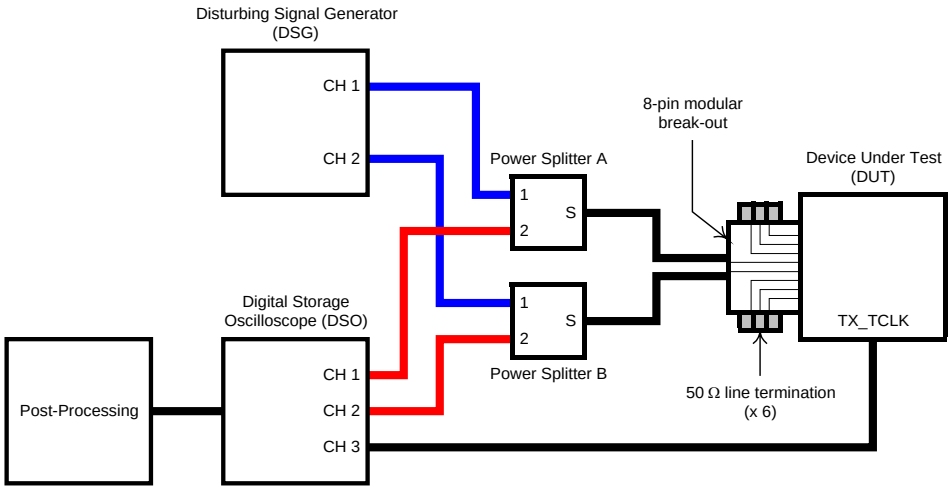

Figure 40.A-1: Test setup block diagram

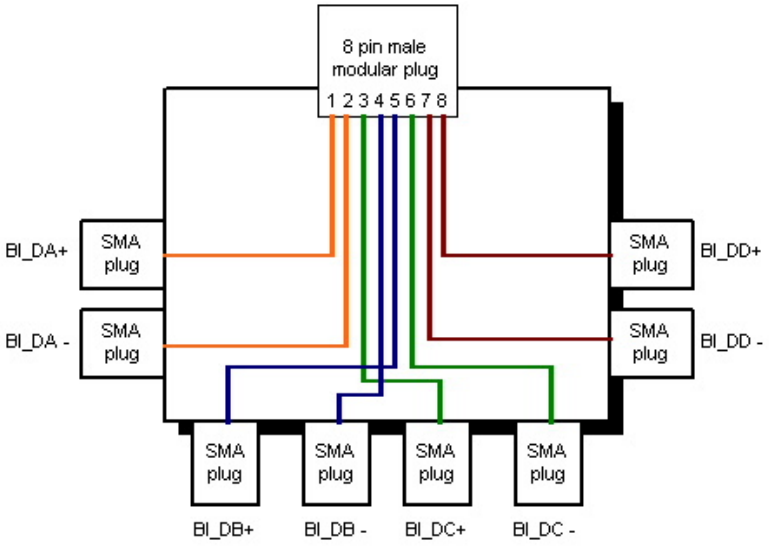

Figure 40.A-2: 8-pin modular breakout board

Note that this test setup does not employ high impedance differential probes. In order to use high impedance differential probes, the vertical range of the oscilloscope must be set to accommodate the sum of V_d and the DUT output. For example, in order to analyze the 2V peak-to-peak DUT output using test fixture 3, the vertical range of the oscilloscope must be set to at least 4.7 V peak-to-peak. If a digital storage oscilloscope (DSO) is used, this increases the quantization error on the DUT output by more than a factor of two. Since a DSO must be used to make post-processing possible, it is beneficial to use the smallest vertical range possible.

To this end, the test setup in Figure 40.A-1 uses power splitters. As its name implies, the power splitter divides a power input to port S evenly between ports 1 and 2. Conversely, inputs to ports 1 and 2 are averaged to produce the output at port S. The key feature of the power splitter is that ports 1 and 2 are isolated. The test setup uses this feature to apply the disturbing signal to the DUT while having a minimum amount of it reach the DSO. In effect, the test setup replicates the hybrid function present in 1000BASE-T devices.

Due to the nature of the setup, V_d is not set to 2.8V peak-to-peak. The magnitude of V_d as seen at port S should be equal to half that defined in the standard. For test fixtures 1 and 2, this is 1.4V peak-to-peak. This means that the actual output voltage of the Disturbing Signal Generator should be approximately 1.4V+3dB. Prior to each test performed, the voltage at port S is verified to be 1.4V peak-to-peak.

Figure 40.A-3 shows the signal flow through the power splitter. Note that the isolation between ports 1 and 2 is no more than 6 dB better than the return loss of the termination at port S. For example, an input to port 1 loses 3 dB on its way to port S. The termination at port S reflects some amount of the power back into the splitter, which is then split evenly between ports 1 and 2 (another 3 dB loss). For conformant 1000BASE-T devices, the return loss at the MDI is greater than 16 dB from 1 to 40 MHz. Therefore, the isolation between ports 1 and 2 is expected to be better than 22 dB when port S is connected to a conformant 1000BASE-T device. In this configuration, the vertical range of the DSO must be set to accommodate the sum of the residual V_d and the DUT output. Since this is much closer to 2V peak-to-peak than 7.4V peak-to-peak, the quantization error on the DUT output will be smaller.

The test setup block diagram in Figure 40.A-1 may be implemented with the equipment listed in Table 40.A-2. The remainder of this appendix discusses the test setup in the context of this implementation.

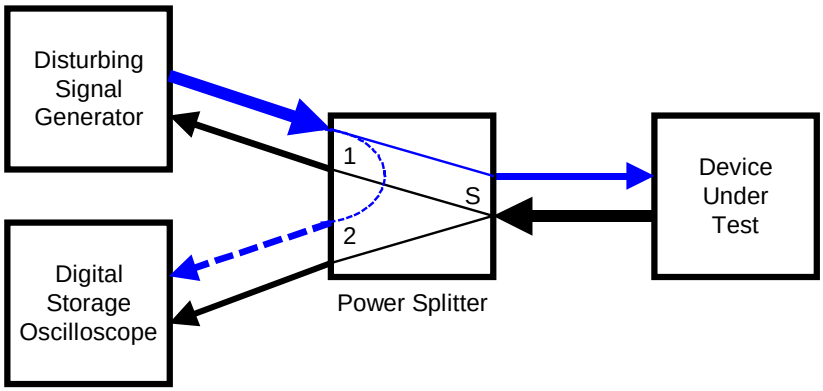

Figure 40.A-3: Power splitter operation

Table 40.A-2: Equipment list

Functional Block	Equipment	Key Features
Disturbing signal generator	Tektronix AWG2021	2 channels, 5 V peak-to-peak output per channel, 250 MS/s sample rate
Digital storage oscilloscope	Tektronix TDS7104	4 channels, 1 GHz bandwidth, 8GS/s sample rate, 16 million sample memory
Power splitter	Mini-Circuits ZSC-2-1W	2-way 0°, 1 to 650 MHz

40.A.2 - Power splitters

Since the power splitters are single-ended devices, two of them are required to make differential measurements. This imposes two constraints. First, the port impedance of the power splitter must be 50Ω so that a differential 100Ω load is presented to the DUT. Second, the power splitters must be matched devices. Differences in the insertion loss, delay, and port impedance of the power splitters will degrade the common-mode rejection of the test setup.

The insertion loss of power splitters A and B are plotted on the same axis in Figure 40.A-4. The measurement was performed using the HP 8753C network analyzer with the HP 85047A S-parameter test set. From this figure, it can be seen that the power splitters are well matched to about 700 MHz. In addition, the insertion loss is about 3.2 dB from 1 to 150 MHz. Note that a 3 dB insertion loss is intrinsic to the operation of a power splitter. The performance of a power splitter is gauged by how much the insertion loss exceeds 3 dB.

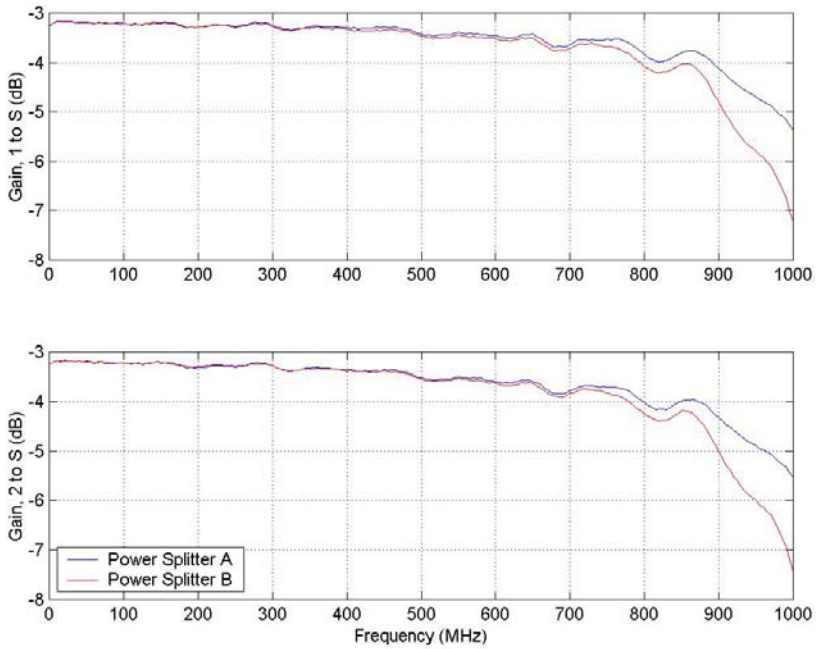

Figure 40.A-4: Power splitter high-frequency insertion loss

Note that the power splitters are AC-coupled devices. The low frequency -3dB cut-off point of the power splitters must also be known so that their impact on droop measurements can be removed. Since the network analyzer is an AC-coupled instrument with a minimum frequency of 300 kHz, the test setup shown in figure 40.A-5 was used to properly measure the low-frequency response.

The test setup shown in Figure 40.A-5 uses the Tektronix AWG2021 to inject low-frequency sine waves into port S of the power splitters. The power splitters are driven differentially. In other words, the input to power splitter B is 180^{o} out of phase with the input to power splitter A. The DSO captures the resultant sine waves at port 1 of the splitters and takes the difference to get a differential signal. The ratio of the differential output amplitude to the differential input amplitude is recorded for a range of frequencies and the results are presented in Figure 40.A-6. The differential input amplitude was 200 mV.

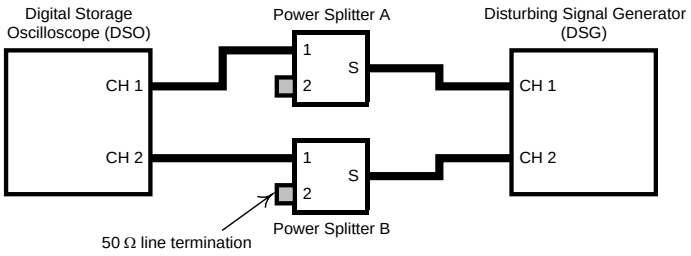

Figure 40.A-5: Test setup for low-frequency cut-off measurement

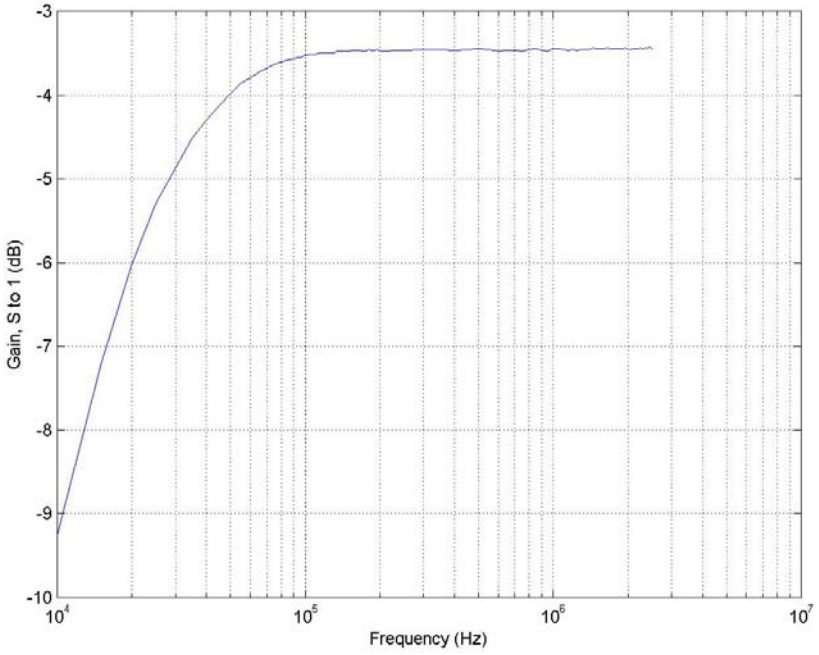

Figure 40.A-6: Low-frequency response of power splitter pair

The low-frequency -3dB cut-off point of the power splitter pair was determined to be 18.3 kHz. This number will be used in the post-processing block to compensate for the low-frequency response of the power splitters and improve the accuracy of droop measurements.

40.6.A.3 – Disturbing signal generator

The disturbing signal generator (DSG) must be able to output a sine wave with the amplitude and frequency required by the test fixture. Furthermore, the DSG must meet spectral purity and linearity constraints and it must have a port impedance of 50Ω to match the power splitters.

The spectral purity and linearity constraints stem from the typical method used to remove the disturbing signal during post-processing. This method uses standard curve fitting routines to find the best-fit sine wave at the disturbing signal frequency. The best-fit sine wave is subtracted from the waveform leaving any harmonics and distortion products behind. Significant harmonics and distortion products can lead to measurement errors. Therefore, the standard requires that all harmonics be at least 40 dB down from the fundamental. Furthermore, the standard states that the DSG must be sufficiently linear so that it does not introduce any "appreciable" distortion products when connected to a 1000BASE-T transmitter.

Note that the use of power splitters makes these constraints easier to satisfy. First, thanks to the isolation between ports 1 and 2, the disturbing signal and the accompanying harmonics and distortion products are greatly attenuated when they reach the DSO. Second, due to the nature of the power splitter, only half of the power output by the 1000BASE-T transmitter reaches the DSG. This reduces the amplitude of any distortion products generated by the DSG. However, since only half of the power output by the DSG reaches the DUT, the DSG is forced to output twice the power in order to get the amplitude required by a given test fixture.

Synthesized 31.25 and 20.83 MHz sine waves from the Tektronix AWG2021 were measured directly with an HP 8593E spectrum analyzer. The results are presented in Figures 40.A-7 and 40.A-8 respectively. These figures show that all harmonics are at least 40 dB below the fundamental.

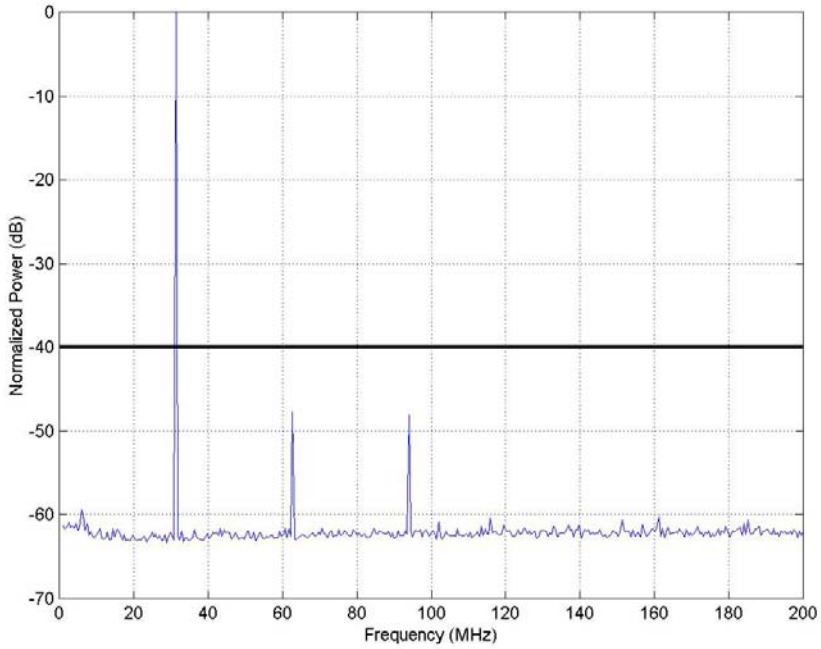

Figure 40.A-7: Spectrum of 31.25 MHz synthesized sine wave from the Tektronix AWG2021

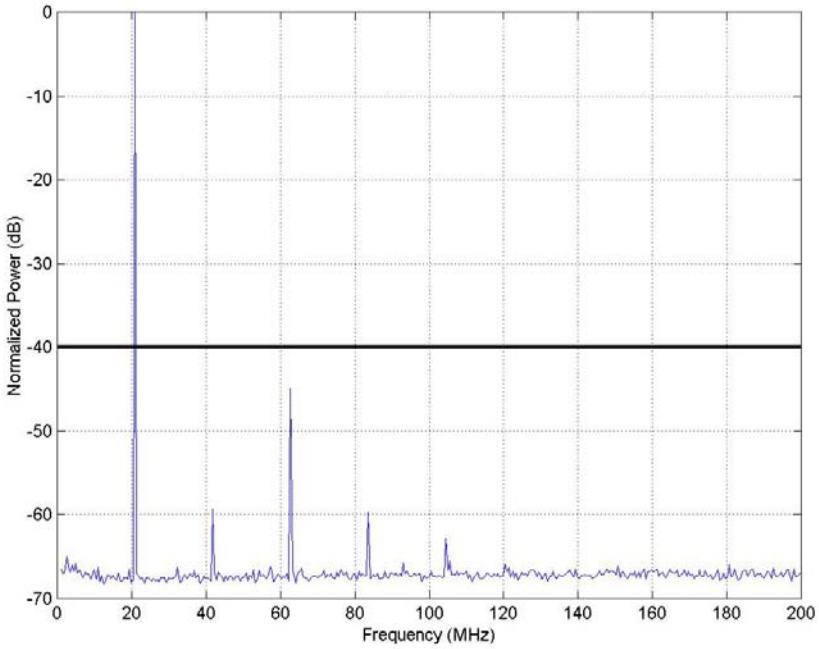

Figure 40.A-8: Spectrum of 20.83 MHz synthesized sine wave from the Tektronix AWG2021

The Tektronix AWG2021 includes built-in filters, which were used to achieve greater harmonic suppression. In order to provide the correct disturbing signal amplitude at the DUT, the output of the Tektronix AWG2021 was set to a level that would compensate for the combined insertion loss of the filter and the power splitter. A complete list of the settings is included in Table 40.A-3.

Table 40.A-3: Tektronix AWG2021 channel 1 settings

Setting	Test Fixtures 1 and 2	Test Fixture 3
Sample Rate	250 MS/s	250 MS/s
Samples Per Cycle	8	12
Amplitude	1.26V peak-to-peak	2.12V peak-to-peak
Filter	50 MHz	50 MHz
Offset	0	0

Note: The settings for channel 2 are identical except that the amplitude of the sine wave is inverted.

The linearity of the Tektronix AWG2021 was tested using the setup shown in Figure 40.A-9. The resistive splitter shown in the test setup has an insertion loss of 6 dB between any two ports. The spectrum measured at the output of port 3 is shown in Figure 40.A-10. This figure shows that all harmonics and distortion products are at least 40 dB below the fundamental. Note that the outputs from channels 1 and 2 are both 4V peak-to-peak.

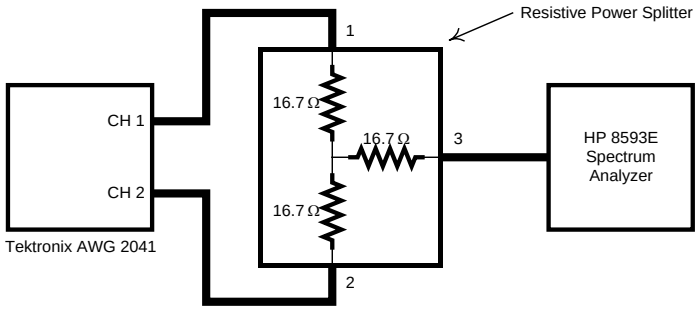

Figure 40.A-9: Test setup for disturbing signal generator linearity measurement

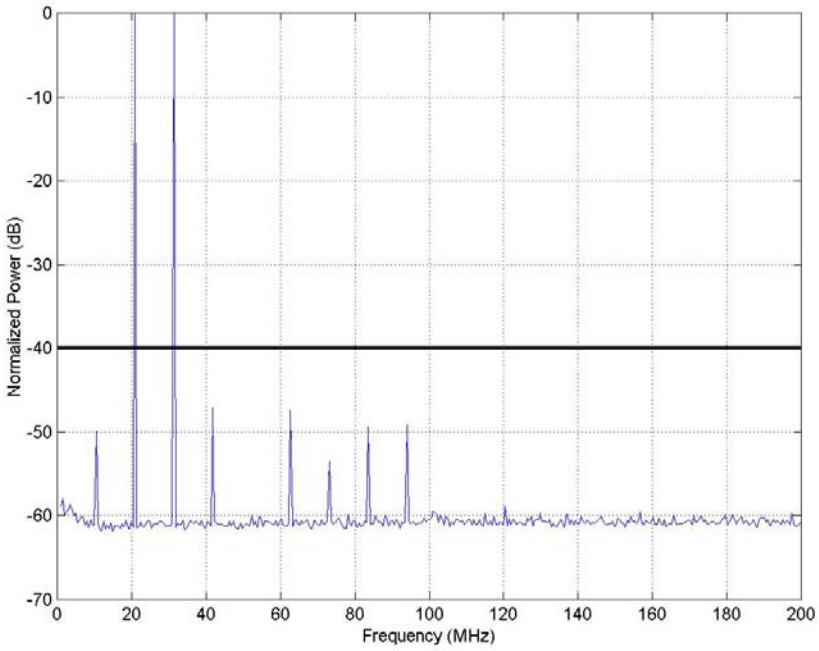

Figure 40.A-10: Spectrum measured at port 3 of the resistive splitter

40.A.4 – Digital Storage Oscilloscope

A digital storage oscilloscope (DSO) with at least three channels is required. Two channels are required to measure the differential signal present at port 2 of the power splitters. These channels must be DC-coupled and they must present a 50Ω characteristic impedance. The third channel is used in test fixtures 3 and 4 to monitor TX_TCLK. The requirements for this channel depend on how TX_TCLK is presented.

Ideally, the frequency response of the oscilloscope would be flat across the bandwidth of interest. Given a 3 ns rise time, the fastest rise time expected for a 1000BASE-T signal, the bandwidth of interest would be roughly 117 MHz, using the *bandwidth=0.35/risetime* rule of thumb.

Another rule of thumb states that the bandwidth of the instrument should be 10 times the bandwidth of interest. If the instrument is assumed to be a first-order low pass filter, the gain only drops 0.5% at one-tenth of the cut-off frequency. Therefore, if the bandwidth of the instrument were on the order of 1 GHz, the frequency response would be reasonably flat out to 117 MHz.

A third rule of thumb is that the sample rate must be at least 10 times the bandwidth of interest for linear interpolation to be used. A minimum sample rate of 2GS/s is recommended for 1000BASE-T signals.

Finally, the DSO should have sufficient sample memory to store the 1000BASE-T transmitter test waveforms. These waveforms are on the order of 16 μs in length. At a 2GS/s sample rate, this would require a sample memory of 32K samples. Deeper sample memories are useful for jitter measurements, but that is beyond the scope of this appendix.

40.A.5 – Post-Processing Block

The post-processing block removes the disturbing signal from the measurement, compensates for the insertion loss and low-frequency response of the power splitters, and applies the high pass test filter when required.

Figure 40.A-11 shows the waveform seen by the oscilloscope when the test setup is functioning as test fixture 1. This waveform is the sum of the transmitter test mode 1 waveform and some residual disturbing signal.

The residual disturbing signal can be removed by subtracting the best-fit sine wave at the disturbing signal frequency. Note that only amplitude and delay (phase) must be fit, since the exact frequency can be measured *a priori*. If multiple waveforms were captured for the purpose of measurement averaging, the amplitude would only need to be fit for the first iteration, leaving phase as the only uncertainty. These shortcuts can be employed to reduce the execution time of the curve-fitting routines.

For the example in Figure 40.A-11, the curve-fitting routine determined that the best-fit amplitude was 48 mV and the best-fit phase was 3.1 μs. The best-fit sine wave was subtracted from the waveform and a scale factor 1.44 ($10^{3.2/20}$) was applied to compensate for the insertion loss of the power splitters. Figure 40.A-12 shows the processed waveform and the DUT output, also referred to as the test setup input, plotted on the same axis. This figure demonstrates the impact that the power splitter's low-frequency response has on the waveform.

The low-frequency response of the power splitter is modeled as first-order high pass filter with a cut-off frequency of 18.3 kHz. Applying the inverse function of this filter to scaled output waveform yields the waveform shown in Figure 40.A-13.

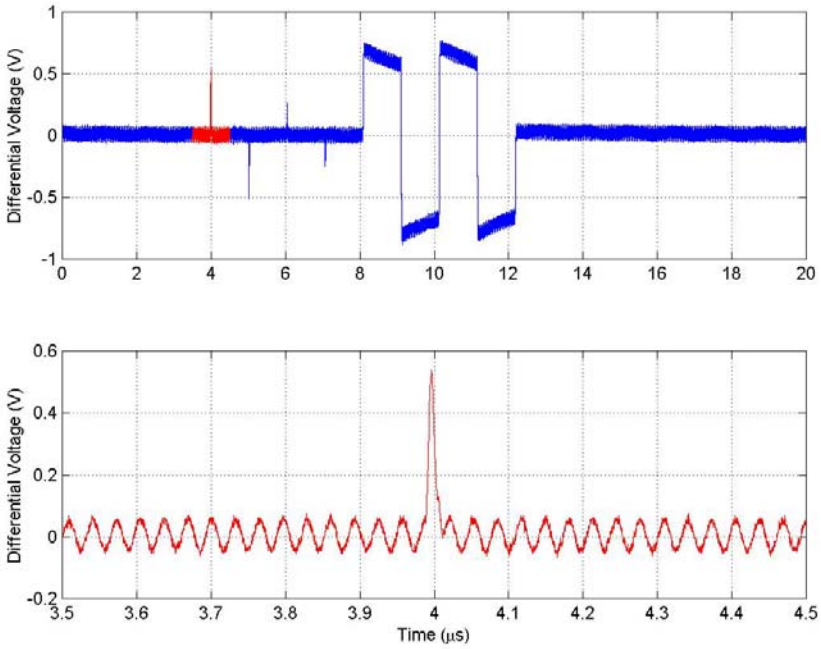

Figure 40.A-11: Observed transmitter test mode 1 waveform before post-processing

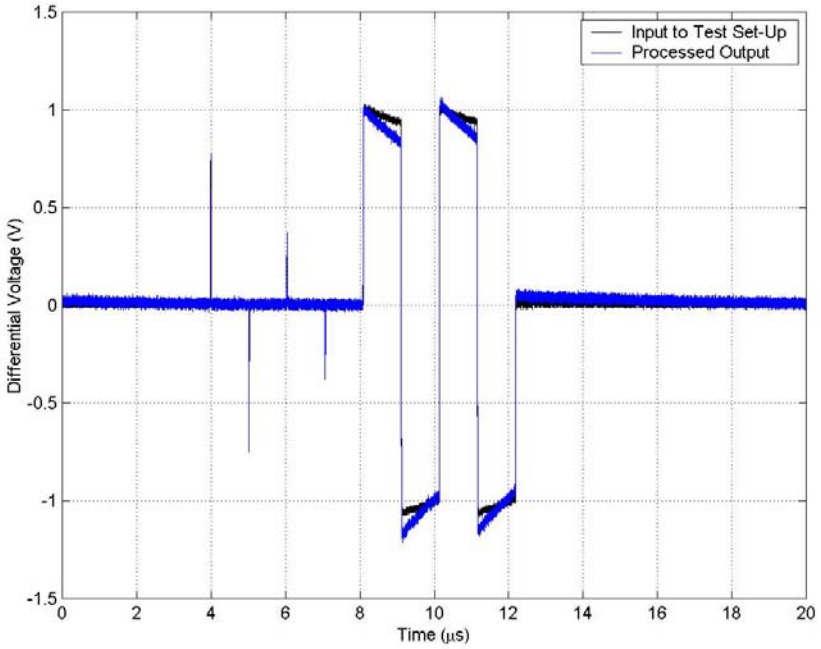

Figure 40.A-12: Input waveform and scaled output waveform with best-fit sine wave removed

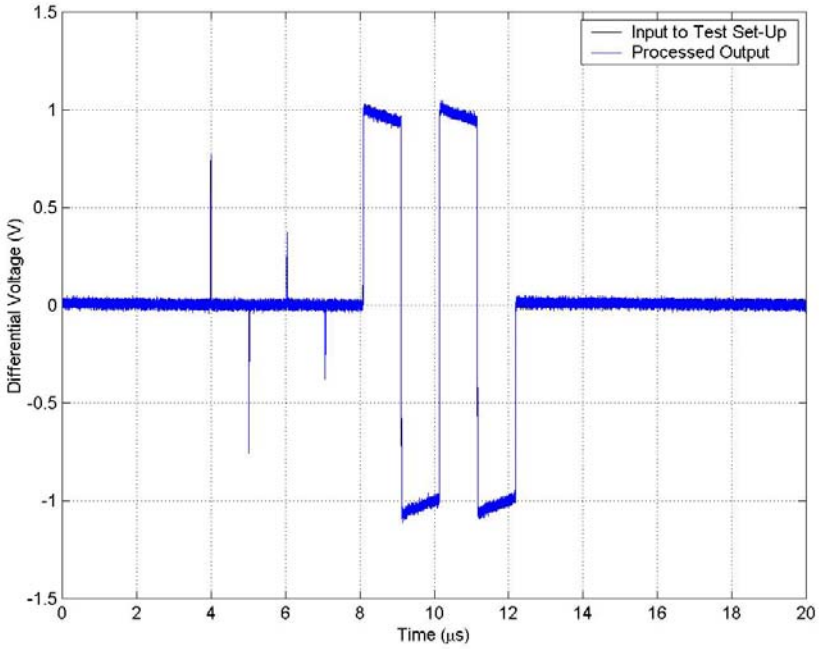

Figure 40.A-13: Output waveform with droop compensation

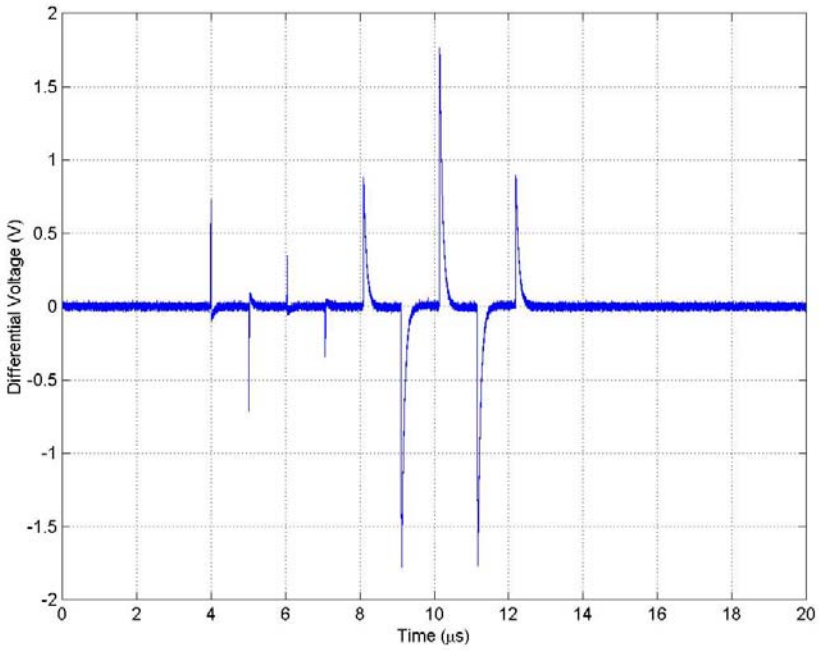

Figure 40.A-14: Output of transmitter test filter

Note from Figure 40.A-13 that the processed waveform is now indistinguishable from the DUT output. This implies that the post-processing successfully removed the distortion of the test setup and that the DUT was linear. If the DUT was not sufficiently linear, then the output would have been distorted due to the presence of the disturbing signal.

Test fixtures 1 and 3 require the presence of a high pass test filter whose cut-off frequency is 2 MHz. While the test filter may be a discrete component, the test setup described in this appendix implements the filter in the post-processing block. An example of the output from this test filter is provided in Figure 40.A-14.

40.A.6 – Complete test setup

The complete test setup must be evaluated in terms of the differential impedance presented to the DUT and the common-mode rejection ratio. Since the test setup is composed of two single-ended circuits, each circuit was measured independently and their differential equivalent was computed. This requires the 8-pin modular plug breakout board to be removed from the measurement. If care is taken with the construction of the board, it will have a minimal impact on the performance of the test setup. This means that the traces from the 8-pin modular plug to the RF connectors must be as short as possible and the trace length must be matched on a pair-for-pair basis. If for some reason the traces must be long (more than 2”), steps must be taken to ensure that the trace impedance is 50Ω.

The reflection coefficient of each circuit with respect to a 50Ω resistive source was measured using an HP 8712B network analyzer. It can be shown that the differential reflection coefficient is the average of the single-ended reflection coefficients. The return loss, which is the magnitude of the reflection coefficient expressed in decibels, is given in Figure 40.A-15.

Note that any differences in the impedance of the two circuits will result in an error in the differential gain of the test setup. If the input impedance of circuit A is Z_A and the input impedance to circuit B is Z_B, the gain error is given in Equation 40.A-1.

$$\text{Gain Error} = \frac{Z_A}{50 + Z_A} + \frac{Z_B}{50 + Z_B} \qquad \text{(Equation 40.A-1)}$$

Equation 40.A-1 assumes that the differential source impedance is a precisely balanced 100Ω resistance. The impedance of each circuit was derived from the reflection coefficient and the gain error is plotted in Figure 40.A-16.

In section 40.A.2, the frequency response of the power splitters was measured for each differential component and again as a pair. Comparing Figures 40.A-4 and 40.A-6, the pass-band gain of each individual power splitter is greater than the gain of the differential pair. This difference is due to the impedance imbalance, and the magnitude of the difference agrees with the data in Figure 40.A-16.

Impedance unbalance also causes common-mode noise to appear as a differential signal. The performance of a differential probe is measured in terms of how well it rejects common-mode noise. This is referred to as the common-mode rejection ratio (CMRR). The CMRR can be computed that difference between the transfer function of the individual circuits. An HP 8712B network analyzer was used to measure the transfer function of each individual circuit and the difference is plotted in Figure 40.A-17.

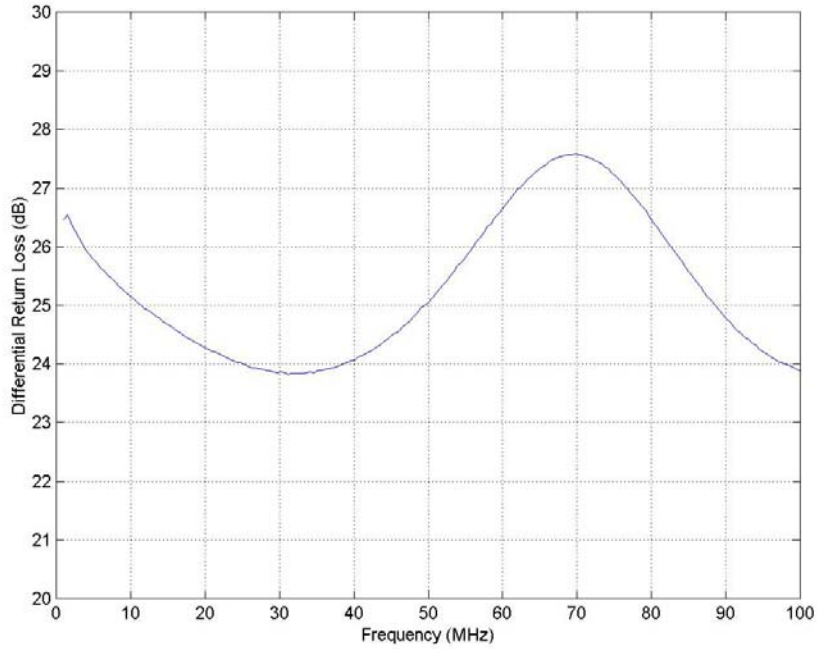

Figure 40.A-15: Differential return loss at the input to the test setup

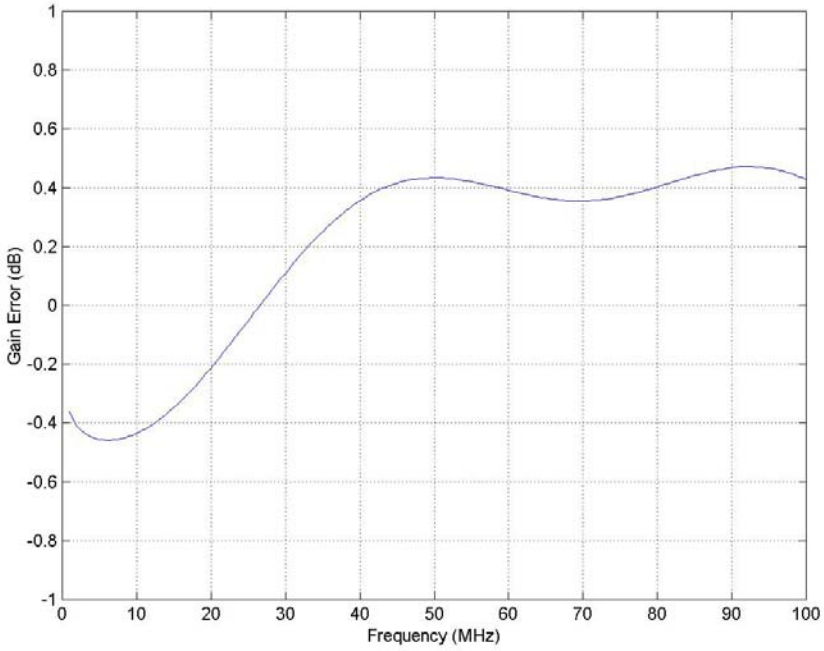

Figure 40.A-16: Differential gain error due to impedance imbalance in the test setup

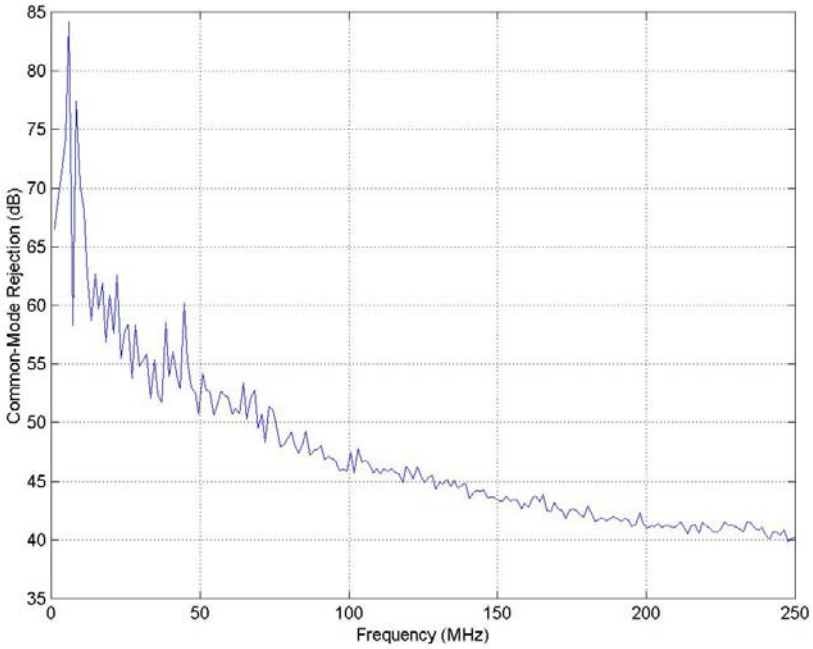

Figure 40.A-17: Test setup common-mode rejection

40.A.7 - Conclusion

This appendix has presented a reference implementation for test fixtures 1 through 4. A single physical test setup was used and each individual test fixture was realized through the configuration of the disturbing signal generator and the post-processing block. Table 40.A-5 summarizes the configuration required to realize each test fixture.

The test setup utilizes a hybrid function to minimize the level of the disturbing signal that reaches the oscilloscope. This allows a smaller vertical range to be used, which in turn reduces the quantization noise on the measurement. Furthermore, it relaxes the constraints placed on the disturbing signal generator in terms of spectral purity. However, the hybrid function also requires additional steps in the post-processing block to deal with insertion loss and the high pass nature of the hybrid.

The test setup was shown to present a reasonable line termination to the device under test. Despite the fact that the test setup uses two single-ended circuits to perform the differential measurement, the matching was sufficient to provide good impedance balance and common-mode rejection.

Table 40.A-5: Realization of 1000BASE-T Transmitter Test Fixtures

Setting	Test Fixture 1	Test Fixture 2	Test Fixture 3	Test Fixture 4
AWG2021 Channel 1				
Sample Rate	250 MS/s	250 MS/s	250 MS/s	—
Samples Per Cycle	8	8	12	—
Filter	50 MHz	50 MHz	50 MHz	—
Amplitude (peak-to-peak)	1.26 V	1.26 V	2.12 V	—
Offset	0	0	0	—
Post-Processing				
V_d Removal	Yes	Yes	Yes	No
Waveform Scaling	Yes	Yes	Yes	Yes
Droop Compensation	Yes	Yes	Yes	Yes
Test Filter	Yes	No	Yes	No
Miscellaneous				
Monitor TX_TCLK	No	No	Yes	Yes

Note 1: The settings for channels 1 and 2 of the AWG2021 are identical except for a 180^o phase-shift.

Appendix 40.B – Transmitter Timing Jitter, No TX_TCLK Access

Purpose: To provide an analysis of the Transmitter Timing Jitter test method defined in Clause 40.6.1.2.5 of IEEE 802.3, and to propose an alternative method that may be used in cases where a device does not provide access to the TX_TCLK signal.

References:

[1] IEEE standard 802.3-2002, subclause 40.6.1.1.1 – Test channel
[2] Ibid., subclause 40.6.1.1.2, figure 40-20 – Test modes
[3] Ibid., subclause 40.6.1.1.3, figure 40-25 – Test fixtures
[4] Ibid., subclause 40.6.1.2.5 – Transmitter Timing Jitter
[5] Test suite appendix 40.6.A – 1000BASE-T transmitter test fixtures

Resource Requirements:

- A DUT without an exposed TX_TCLK clock signal
- Digital storage oscilloscope, Tektronix TDS7104 or equivalent
- 8-pin modular plug break-out board
- 50 Ω coaxial cables, matched length
- 50 Ω line terminations (6)

Last Modification: March 25, 2002 (Version 1.1)

Discussion:

40.B.1 – Introduction

In addition to supporting the standard transmitter Test Modes, the jitter specifications found in Clause 40.6.1.2.5 require a device to provide access to the internal TX_TCLK signal in order to perform the Transmitter Timing Jitter tests. While access to the TX_TCLK signal is relatively straightforward and easy to provide on evaluation boards and prototype systems, it can become impractical in more formal implementations. In the case where no exposed TX_TCLK signal is available, it may be possible to perform a simplified version of the full jitter test procedure which could provide some useful information about the quality and stability of a device's transmit clock. This Appendix will discuss the present test method, and will propose an alternate test procedure that may be used to perform a simplified jitter test for devices that support both transmitter Test Mode 2 (TM2) and Test Mode 3 (TM3), but do not provide access to the TX_TCLK signal. Because this procedure deviates from the specifications outlined in Clause 40.6.1.2.5, it is not intended to serve as a legitimate substitute for that clause, but rather as an informal test that may provide some useful insight regarding the overall purity and stability of a device's transmit clock.

40.B.2 – MASTER timing mode tests

The formal MASTER timing mode jitter procedure of Clause 40.6.1.2.5 can basically be summarized by the following steps (with the DUT configured as MASTER):

- Measure the pk-pk jitter from the TX_TCLK to the MDI (i.e., "J_{txout}").
- Measure the pk-pk jitter on the TX_TCLK, relative to an unjittered reference.
  - This must be less than 1.4ns.
- HPF (5KHz) the TX_TCLK jitter, take the peak-to-peak value, and add J_{txout}
  - This result must be less than 0.3 ns.

We see that there are essentially specifications on the following two parameters:

1) Unfiltered jitter on the TX_TCLK.
2) Sum of the filtered TX_TCLK jitter plus the unfiltered J_{txout}.

In actual systems, it should be fairly reasonable to assume that J_{txout} will be relatively small compared to the filtered TX_TCLK jitter. If J_{txout} were zero, access to the internal TX_TCLK wouldn't be necessary, because the TM2 jitter at the MDI would be identical to the jitter on the internal TX_TCLK. In effect, you would essentially be able to "see" the TX_TCLK jitter through the MDI.

It is this idea that allows us to design a hypothetical test procedure for the case when a device does not provide access to TX_TCLK. Suppose the following procedure is performed:

- Measure the unfiltered peak-to-peak jitter on the TM2 output at the MDI, relative to an unjittered reference.
- Filter the MDI output jitter with the 5KHz HPF to determine the filtered peak-to-peak jitter.

Note that the TM2 jitter measured at the MDI is actually the sum of the TX_TCLK jitter plus J_{txout}. Given this fact, one could argue that if the TM2 jitter, relative to an unjittered reference, is less than 1.4ns, then the TX_TCLK jitter component alone **must** be less than 1.4ns as well. (In other words, if the results are conformant when J_{txout} is included, the results would be even better if J_{txout} could be separately measured and subtracted.) Thus, the device could be given a legitimate passing result for the unfiltered MASTER TX_TCLK jitter if the measured TM2 jitter relative to an unjittered reference is less than 1.4ns.

A similar argument can be made for the filtered TX_TCLK jitter case. In the formal jitter test procedure, J_{txout} is **not** filtered before it is added to the filtered TX_TCLK jitter. For our hypothetical test, the jitter at the MDI (after filtering) is effectively the sum of the filtered TX_TCLK jitter plus the filtered J_{txout}. Thus, we can conclude that if the filtered TM2 jitter is greater than 0.3ns, it would only fail in a **worse** manner if J_{txout} were not filtered prior to being added to the filtered TX_TCLK jitter.

Note that this test is *inconclusive* if the peak-to-peak value of the filtered MDI jitter is less than 0.3ns. This is because it can't be known for sure exactly how the filtered jitter is distributed between J_{txout} and actual TX_TCLK jitter. For example, suppose that in our hypothetical test, the result for the filtered jitter was just under 0.3ns, and the device was given a passing result for the filtered TX_TCLK jitter test. If the filtered jitter was 100% due to J_{txout} (i.e., TX_TCLK jitter was zero), then the device would actually fail the formal test, where J_{txout} is measured *sans filter* before being added to the filtered TX_TCLK jitter. Thus, the original passing result of our hypothetical test would have been incorrect.

By the same logic, the results are also inconclusive for the unfiltered jitter case when the peak-to-peak result is greater than 1.4ns. Again, this is because it is not possible to know how much of this value is due to J_{txout}. Thus, assigning a failing result to a device whose unfiltered TM2 jitter was just above 1.4ns could be incorrect if it

was otherwise determined that a large part of the jitter was due to J_{txout}, which would not have been included in the unfiltered TX_TCLK jitter value had the formal jitter test procedure been performed.

The table below summarizes the possible outcomes of the hypothetical test, and lists the pass/fail result that may be assigned for the given outcome.

Table 40.B-1: Hypothetical test outcomes and results

Paramter	Conformance Limit	Result < Limit	Result > Limit
Unfiltered TM2 jitter	1.4ns	PASS	Inconclusive
Filtered TM2 jitter	0.3ns	Inconclusive	FAIL

40.B.3 – SLAVE timing mode tests

The question remains as to the possibility of designing a similar hypothetical test for the SLAVE timing mode case based on the Test Mode 3 (TM3) signal observable at the MDI. Unfortunately, this is not as straightforward as was the case for the MASTER timing mode. This is due to the fact that the formal procedure of Clause 40.6.1.2.5 relies heavily on access to both the MASTER and SLAVE TX_TCLK signals for SLAVE jitter measurements, in addition to the fact that the SLAVE measurements are to be made *with both devices operating normally, connected to each other via their MDI ports*, which precludes the use of the MDI for the purpose of gaining access to the internal TX_TCLK.

Furthermore, the meaning of the Test Mode 3 mode itself is somewhat confusing as it is described in Clause 40.6.1.1.2:

"When test mode 3 is enabled, the PHY shall transmit the data symbol sequence {+2, –2} repeatedly on all channels. The transmitter shall time the transmitted symbols from a 125.00 MHz +/-0.01% clock in the SLAVE timing mode. A typical transmitter output for transmitter test modes 2 and 3 is shown in Figure 40–20."

A SLAVE physical layer device is defined in Clause 1.4.255 as, *"the PHY that recovers its clock from the received signal and uses it to determine the timing of transmitter operations."* If it is truly intended that a device be operating in the SLAVE timing mode while in Test Mode 3, it would need to be provided with a signal at the MDI from which to determine the recovered clock. This, however, would preclude the measurement of the SLAVE J_{txout} values due to the fact that one cannot simultaneously provide a reference clock and monitor the TM3 waveform on the same bi-directional MDI wire pair. The most reasonable interpretation of intended TM3 operation (on the part of the author, anyway,) would be that a DUT would use it's own MASTER clock as the "received signal", and provide it internally to the SLAVE clock recovery mechanism, which would then generate the clock used for transmitting the {+2, -2} symbol sequence for TM3. The problem with this method from a conformance perspective is that it is impossible to verify that a device is truly operating in this manner when it is in TM3. (Perhaps a better implementation of TM3 would be to simply send another device's TM2 signal into the DUT's MDI while the DUT's transmitter remains silent. Then, the jitter on the DUT (SLAVE) TX_TCLK could be measured with respect to the incoming TM2 signal.) Regardless, it is still difficult to design an abbreviated test for SLAVE mode jitter that strictly adheres to the specifications of Clause 40.6.1.2.5, and does not require access to the TX_TCLK.

It may be possible however, to design a test that attempts to emulate the intentions of the formal procedure, while deviating from it as little as possible. To begin, note that the formal method for measuring the SLAVE-related jitter parameters can be summarized by the following steps:

- Configure DUT (SLAVE) for TM3. Measure the jitter from the TX_TCLK to the MDI (i.e., "J_{txout}").
- Connect the DUT to the Link Partner (MASTER) through the Jitter Test Channel.
- Measure the jitter on the MASTER TX_TCLK, relative to an unjittered reference. Filter this jitter waveform with a 5KHz HPF. Record the peak-to-peak value of the result. (This value will be subtracted later from the measured SLAVE jitter value.)
- Measure the jitter on the DUT TX_TCLK, relative to the MASTER TX_TCLK.
    - This must be less than 1.4ns peak-to-peak.
- Filter the DUT TX_TCLK jitter waveform with a 32KHz HPF, take the peak-to-peak value, add J_{txout}, and subtract the recorded peak-to-peak filtered MASTER jitter value.
    - This result must be less than 0.4 ns.

The key concepts of this method are basically:

1) Measure the filtered jitter on the "source clock".
2) Pass the clock through a worst-case echo environment.
3) Measure the unfiltered jitter on the "recovered clock", with respect to the source clock.
4) Filter this jitter, subtract J_{txout}, and subtract the filtered jitter from the source clock.

If a device is intended to use its own MASTER clock as the input from which the SLAVE clock is derived, a hypothetical approximation for this procedure for the case where one only has access to the MDI signaling might be:

1) Measure the DUT's TM2 jitter relative to an unjittered reference, filter with a 5KHz HPF, and record both the filtered and unfiltered peak-to-peak values.
2) Measure the DUT's TM3 jitter relative to an unjittered reference. Subtract the unfiltered TM2 peak-to-peak jitter value.
    - This result must be less than 1.4ns.
3) Filter the TM3 jitter with a 32KHz HPF, subtract the filtered TM2 pk-pk jitter value.
    - This result must be less than 0.4ns.

This procedure approximates the formal procedure, with two exceptions. The first is that it is obviously not possible to insert the jitter test channel between the source clock and the recovered clock. The second difference is that in addition to the jitter test channel, the MASTER's J_{txout} is also present between the source and recovered clocks in the formal procedure, but is not present in the hypothetical test procedure (although it should be zero if the DUT's internal MASTER TX_TCLK is being used directly as the input to the PLL.)

Given that these two differences actually make the clock recovery operation easier for the DUT, it is technically inappropriate to apply the same SLAVE mode conformance limits specified in Clause 40.6.1.2.5. (If somehow the alternate test conditions were more difficult, the same argument from the hypothetical MASTER test could be used, i.e., if the device can still pass under tougher conditions, we can be fairly certain that it would pass under the formal test conditions.) One solution to this problem would be to revise the conformance limits to stricter values, however this would require research into what these values should be, and these values would need to be verified an accepted by the general community. Not having this, a possible alternative would be to perform the tests and report the numerical results for purely informational purposes without judging them on a pass/fail basis, with the only exception being the results of the MASTER mode (TM2) tests when the results are within the pass/fail regions shown in Table 40.6.B-1.

40.B.4 – Conclusion

This appendix was intended as an analysis of the jitter test procedure of Clause 40.6.1.2.5, for the case where a device does not provide access to the TX_TCLK signal. An attempt was made to basically "do the best with what you've got", and determine what subset (if any) of the jitter specifications can be verified if the TX_TCLK signal is not available. The analysis provides a method that is solely based on the Test Mode 2 and Test Mode 3 signals as observed at the MDI. The method for the MASTER mode jitter parameters can, under some circumstances, yield legitimate pass/fail results for a particular DUT however, depending on the measured values, will produce inconclusive results. In these cases, while it may not be possible to assign a pass/fail judgment, the determined jitter values may still be useful from a design perspective and could be reported for informational purposes only.

It was concluded that it is not possible to strictly verify any of the SLAVE mode jitter parameters without access to the TX_TCLK, however an alternate method was presented which approximates the intentions of the formal procedure. Because the method is a simplified version of the formal procedure, it is not possible to apply the same conformance limits specified in the standard, thus reducing it to a purely informal test. Depending on the validity of the analysis and the ultimate need for such a test, it might be possible to develop this method into a valid alternative, although new conformance limits would need to be determined and the method would need to be accepted by the standards body.

Appendix 40.C – Transmitter Specifications

Purpose: To present an example transmitter electrical specification to implement the 1000BASE-T PMA Receiver test suite.

Last Modification: January 9, 2004 (Version 1.0)

Discussion:

40.C.1 – Introduction

This appendix describes of the transmitter electrical specifications for the BER verification test suite used by the University of New Hampshire.

40.C-2 – Transmitter Specifications

Table 40.C-1: Summary of results from 1000BASE-T PMA testing performed on the 4.61ns Rise Time Transmitter

Test	Parameter	BI_DA	BI_DB	BI_DC	BI_DD	Units
40.1.1	**Peak Differential Output Voltage and Level Accuracy**					
	Magnitude of the voltage at point A	681	685	681	682	mV
	Magnitude of the voltage at point B	681	684	678	681	mV
	Difference between the magnitudes of the voltages at points A and B	0.000	0.146	0.441	0.147	%
	Difference between the magnitude of the voltage at point C and 0.5 times the average of the voltage magnitudes at points A and B	0.587	0.292	0.588	0.880	%
	Difference between the magnitude of the voltage at point D and 0.5 times the average of the voltage magnitudes at points A and B	0.587	0.292	0.000	0.293	%
40.1.2	**Maximum Output Droop**					
	Ratio of the voltage at point G to the voltage at point F	95.8	95.0	95.6	95.7	%
	Ratio of the voltage at point J to the voltage at point H	96.12	95.72	95.75	95.62	%
40.1.3	**Differential Output Templates**					
	Waveform around point A	Pass	Pass	Pass	Pass	
	Waveform around point B	Pass	Pass	Pass	Pass	
	Waveform around point C	Pass	Pass	Pass	Pass	
	Waveform around point D	Pass	Pass	Pass	Pass	
	Waveform around point F	Pass	Pass	Pass	Pass	
	Waveform around point H	Pass	Pass	Pass	Pass	

Table 40.C-2: Summary of results from 1000BASE-T PMA testing performed on the 5.12ns Rise Time Transmitter

Test	Parameter	BI_DA	BI_DB	BI_DC	BI_DD	Units
40.1.1	**Peak Differential Output Voltage and Level Accuracy**					
	Magnitude of the voltage at point A	693	696	692	695	mV
	Magnitude of the voltage at point B	693	696	688	692	mV
	Difference between the magnitudes of the voltages at points A and B	0.000	0.000	0.578	0.432	%
	Difference between the magnitude of the voltage at point C and 0.5 times the average of the voltage magnitudes at points A and B	0.578	0.862	0.290	0.576	%
	Difference between the magnitude of the voltage at point D and 0.5 times the average of the voltage magnitudes at points A and B	0.578	0.000	0.290	0.288	%
40.1.2	**Maximum Output Droop**					
	Ratio of the voltage at point G to the voltage at point F	95.8	95.3	95.4	95.6	%
	Ratio of the voltage at point J to the voltage at point H	96.14	95.69	95.99	95.66	%
40.1.3	**Differential Output Templates**					
	Waveform around point A	Pass	Pass	Pass	Pass	
	Waveform around point B	Pass	Pass	Pass	Pass	
	Waveform around point C	Pass	Pass	Pass	Pass	
	Waveform around point D	Pass	Pass	Pass	Pass	
	Waveform around point F	Pass	Pass	Pass	Pass	
	Waveform around point H	Pass	Pass	Pass	Pass	

Appendix 40.D – Rise Time Calculation

Purpose: To present the methodology used to find the rise time of a 1000Base-T transmitter.

Last Modification: January 9, 2004 (Version 1.0)

Discussion:

40.D.1 – Introduction

This appendix describes of the methodology used by the University of New Hampshire to determine the rise time of the transmitter configuration used in the 1000Base-T PMA Receiver Test Suite. This description is intended to be an example for those that wish to implement the test suite in their own lab.

40.D-2 – Rise Time Estimation

Signal rise is defined as a transition from the baseline voltage to $+V_{out}$. The signal rise time is defined to be the time difference between the points where the signal transition crosses 10% and 90% of V_{out}.

The standard does not define a rise time requirement for 1000Base-T, nor does it describe a method in which to measure the rise time. This test suite utilizes the "A" reference pulse in the Test Mode 1 waveform to calculate the transmitter rise time. The rise time of this pulse is measured from the 10% to 90% marks of the rising edge of the pulse, as shown below in Figure 40.D-1

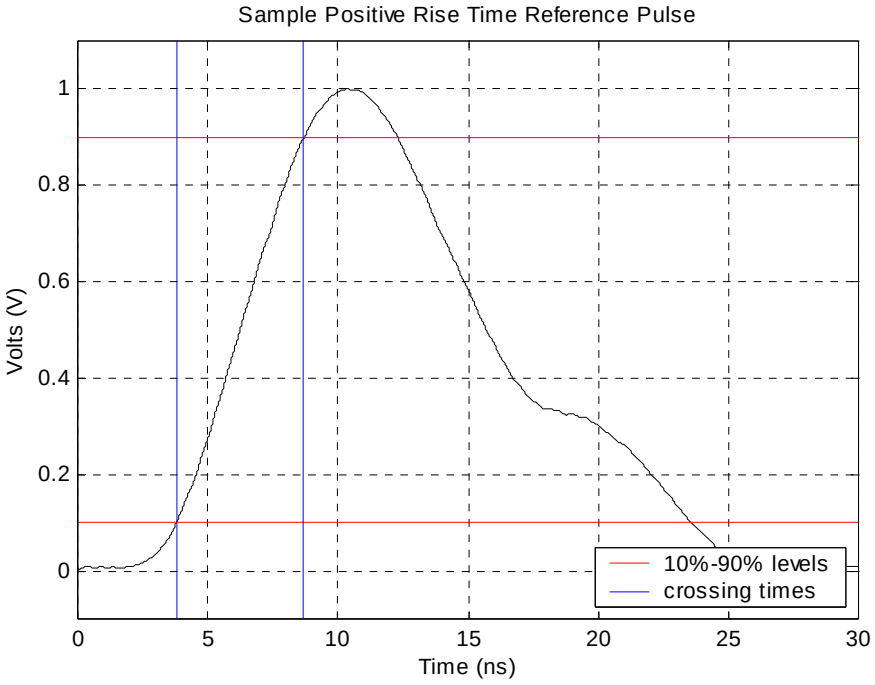

Figure 40.D-1: Sample Positive Rise Time Measurement

Appendix 40.E – Category 5e Cable Test Environment

Purpose: To examine the specifications of a category 5e cable test environment.

Last Modification: January 9, 2004 (Version 1.0)

Discussion:

Since equalizers often tend to be optimized for particular cable conditions the test procedure uses both high attenuation and a low attenuation environment. The high attenuation testing is done over a Category 5e compliant channel attenuated to simulate a worst-case environment equivalent of 60 degrees (Refer to Table 40.E-1). The low attenuation testing is done over a Category 5e compliant channel specified in Table 40.E-1. Each of these channels must be tested to ensure that they meet the expected characteristics as defined by their associated standards.

Table 40.E-1: UTP Channel Definitions

		Insertion Loss – Low (+/- 1 dB)[a]			**Insertion Loss – High (+/- 1 dB)[a]**		
Technology	**Media Type**	**16 MHz**	**32 Mhz**	**100 Mhz**	**16 MHz**	**32 MHz**	**100 MHz**
1000BASE-T	Category-5 UTP	9.9	14.2	25.7	0.6	1.0	1.2

[a]Insertion loss is the sum of channel attenuation and connector losses.

Appendix 40.F – Bit Error Rate Measurement

Purpose: To develop a procedure for bit error rate measurement through the application of statistical methods.

References:

[1] Miller, Irwin and John E. Freund, Probability and Statistics for Engineers (Second Edition), Prentice-Hall, 1977, pp. 194-210, 240-245.

Last Modification: January 9, 2004 (Version 1.0)

Discussion:

40.F.1 – Introduction

One key performance parameter for all digital communication systems is the bit error rate (BER). The bit error rate is the probability that a given bit will be received in error. The BER may also be interpreted as the average number of errors that would occur in a sequence of n bits.

While the bit error rate concept is quite simple, the measurement of this parameter poses some significant challenges. The first challenge is deciding the number of bits, n, that must be sent in order to make a reliable measurement. For example, if 10 bits were sent and no errors were observed, it would be foolish to conclude that the bit error rate is zero. However, common sense tells us that the more bits that are sent without error, the more reasonable this conclusion becomes. In the interest of keeping the test duration as short as possible, we want to send the smallest number of bits that provides us with an acceptable margin of error.

This brings us to the second challenge of BER measurement. Given that we send n bits, what reasonable statements can be made about the bit error rate based on the number of errors observed? Returning to the previous example, if 10 bits are sent and no errors are observed, it is unreasonable to say that the BER is zero. However, it may be more reasonable to say that the BER is 10^{-1} or better. Furthermore, you are absolutely certain that the bit error rate is not 1.

In this appendix, two statistical methods, hypothesis testing and confidence intervals, are applied to help us answer the questions of how many bits we should be sent and what conclusions can be made from the test results.

40.F.2 – Statistical Model

A statistical model for the number of errors that will be observed in a sequence of n bits must be developed before we apply the aforementioned statistical methods. For this model, we will assume that every bit received is an independent Bernoulli trial. A Bernoulli trial is a test for which there are only two possible outcomes (i.e. a coin toss). Let us say that p is the probability that a bit error will occur. This implies that the probability that a bit error will not occur is (1-p).

The property of independence implies that the outcome of one Bernoulli trial has no effect on the outcomes of the other Bernoulli trials. While this assumption is not necessarily true for all digital communications systems, it is still used to simplify the analysis.

The number of successful outcomes, k, in n independent Bernoulli trials is taken from a binomial distribution. The binomial distribution is defined in equation 40.F-1.

$$b(k; n, p) = C_{n,k}\, p^k (1-p)^{n-k} \qquad \text{(Equation 40.F-1)}$$

Note that in this case, a successful outcome is a bit error. The coefficient $C_{n,k}$ is referred to as the binomial coefficient or "n-choose-k". It is the number of combinations of k successes in n trials. Returning to coin toss

analogy, there are 3 ways to get 2 heads from 3 coin tosses: (tails, heads, heads), (heads, tails, heads), and (heads, heads, tails). Therefore, $C_{3,2}$ would be 3. A more precise mathematical definition is given in equation 40.F-2.

$$C_{n,k} = \frac{n!}{k!(n-k)!} \quad \text{(Equation 40.F-2)}$$

This model reflects the fact that for a given probability, p, a test in which n bits are sent could yield many possible outcomes. However, some outcomes are more likely than others and this likelihood principle allows us to make conclusions about the BER for a given test result.

40.F.3 – Hypothesis Test

The statistical method of hypothesis testing will allow us to establish a value of n, the number of bits to be sent, for the BER measurement. Naturally, the test begins with a hypothesis. In this case, we will hypothesize that the probability of a bit error, p, for the system is less than some target BER, P_0. This hypothesis is stated formally in equation 40.F-3.

$$H_0 : p \le P_0 \quad \text{(Equation 40.F-3)}$$

We now construct a test for this hypothesis. In this case, we will take the obvious approach of sending n bits and counting the number errors, k. We will interpret the test results as shown in table 40.F-1.

Table 40.F-1: Acceptance and rejections regions for H_0

Test Result	**Conclusion**
k = 0	H_0 is true
k > 0	H_0 is false

We now acknowledge the possibility that our conclusion is in error. Statisticians define two different categories of error. A type I error is made when the hypothesis is rejected even though it is true. A type II error is made when the hypothesis is accepted even though it is false. The probability of a type I and a type II error are denoted as α and β respectively. Table 40.F-2 defines type I and type II errors in the context of this test.

Table 40.F-2: Definitions of type I and type II errors

Type I Error	k > 0 even though p ≤ BER
Type II Error	k = 0 even though p > BER

A type II error is arguably more serious and we will define n so that the probability of a type II error, β, is acceptable. The probability of a type II error is given in equation 40.F-4.

$$\beta = (1-p)^n < (1-P_0)^n \quad \text{(Equation 40.F-4)}$$

Equation 40.F-4 illustrates that the upper bound on the probability of a type II error is a function of the target bit error rate and n. By solving this equation for n, we can determine the minimum number of bits that need to sent in order to verify that p is less than a given P_0 for a given probability of type II error.

$$n > \frac{\ln(\beta)}{\ln(1-P_0)} \quad \text{(Equation 40.F-5)}$$

Let us now examine the probability of a type I error. The definition of α is given in equation 40.F-6.

$$\alpha = 1-(1-p)^n \le 1-(1-P_0)^n \quad \text{(Equation 40.F-6)}$$

Equation 40.F-6 shows that while we increase n to make β small, we simultaneously raise the upper bound on α. This makes sense since the likelihood of observing a bit error increases with the number of bits that you send, no matter how small bit error rate is. Therefore, while the hypothesis test is very useful in determining a reasonable value for n, we must be very careful in interpreting the results. Specifically, if we send n bits and observe no errors, we are confident that p is less than our target bit error rate (our level of confidence depends on how small we made β). However, if we do observe bit errors, we cannot be quick to assume that the system did not meet the BER target since the probability of a type I error is so large. In the case of k > 0, a confidence interval can be used to help us interpret k.

40.F.4 – Confidence Interval

The statistical method of confidence intervals will be used to establish a lower bound on the bit error rate given that k > 0. A confidence interval is a range of values that is likely to contain the actual value of some parameter of interest. The interval is derived from the measured value of the parameter, referred to as the point estimate, and the confidence level, (1-α), the probability that the parameter's actual value lies within the interval.

A confidence interval requires a statistical model of the parameter to be bounded. In this case, we use the statistical model for k given in equation 40.F-1. If we were to compute the area under the binomial curve for some interval, we would be computing the probability that k lies within that interval. This concept is shown in figure 40.F-1.

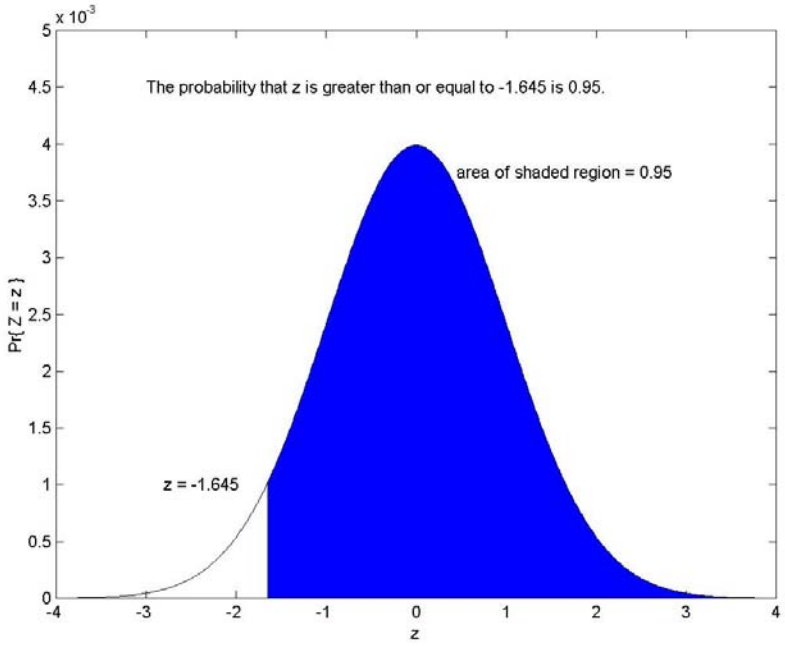

Figure 40.F-1: Computing the probability that z ≥ -1.645 (standard normal distribution).

To compute the area under the binomial curve, we need a value for the parameter p. To compute a confidence interval for k, you assume that k/n, the point estimate for p, is the actual value of p.

Note that figure 40.F-1 illustrates the computation of the lower tolerance bound for k, a special case where the confidence interval is $[k_l, +\infty]$. A lower tolerance bound implies that in a percentage of future tests, the value of k will be greater than k_l. In other words, actual value of k is greater than k_l with probability equal to the confidence

level. Therefore, if k_l/n is greater than P_0, we can say that the system does not meet the target bit error rate with probability (1-α). By reducing α, we reduce the probability of making a type I error.

To determine the value of k_l, it is useful to assume that the binomial distribution can be approximated by a normal (Gaussian) distribution when n is large. The mean and variance of this equivalent distribution are the mean and variance of the corresponding binomial distribution (given in equations 40.F-7 and 40.F-8).

$$\mu_K = np \quad \text{(Equation 40.F-7)}$$

$$\sigma_K^2 = np(1-p) \quad \text{(Equation 40.F-8)}$$

Now, let α be the probability that $Z \leq z_\alpha$ where Z is a standard normal random variable. A standard random variable is one whose mean is zero and whose variance is one. The random variable K can be standardized as shown in equation 40.F-9.

$$Z = \frac{K - \mu_K}{\sigma_K} \quad \text{(Equation 40.F-9)}$$

Note that Z is greater than z_α with probability (1-α), the confidence level. We apply this inequality to equation 40.F-9 and solve for K to get equation 40.F-10.

$$K > \mu_K + z_\alpha \sigma_K$$
$$K > np + z_\alpha \sqrt{np(1-p)} \quad \text{(Equation 40.F-10)}$$

As mentioned before, we assume that p is k/n. We can now generate an expression for k_l, the value that K will exceed with probability (1-α). This expression is given in equation 40.F-11.

$$k_l = k + z_\alpha n \sqrt{\frac{(k/n)(1-k/n)}{n}} \quad \text{(Equation 40.F-11)}$$

Finally, we argue that if K exceeds k_l, then the actual value of p must exceed k_l/n. Therefore, we can generate an expression for p_l, the value that p will exceed with probability (1-α), and compare it to the target bit error rate. By applying this comparison (given in equation 40.F-12) the probability of a type I error can be greatly reduced. For example, by setting z_α to –1.645, the probability of a type I error is reduced to 5%.

$$P_0 \geq p_l = \frac{k_l}{n} = \frac{k}{n} + z_\alpha \sqrt{\frac{(k/n)(1-k/n)}{n}} \quad \text{(Equation 40.F-12)}$$

40.F.5 – Sample Test Construction

We now compress the theory presented in sections 40.F-2 through 40.F-4 into two inequalities that may be used to construct a bit error rate test. First, we take equation 40.F-5 and assume that $\ln(1-P_0)$ is $-P_0$ (valid for P_0 much less than one). The result is equation 40.F-13.

$$n > \frac{-\ln(\beta)}{P_0} \quad \text{(Equation 40.F-13)}$$

Second, we examine equation 40.F-12. Assuming that (1-k/n) is very close to 1 and substituting $-\ln(\beta)/P_0$ for n, we get equation 40.F-14.

$$-\ln(\beta) \geq k + z_\alpha \sqrt{k}$$ (Equation 40.F-14)

The largest value of k that satisfies equation 40.F-14 is k_l. The benefit of these two equations is that a bit error rate test is uniquely defined by β and α and that the test scales with P_0. Table 40.F-3 defines n and k_l in terms of β and α.

Table 40.F-3: n and k_l as a function of β and α.

β	–ln(β)	n	α	z_α	k_l
0.10	2.30	$2.30/P_0$	0.10	–1.29	5
0.10	2.30	$2.30/P_0$	0.05	–1.65	6
0.05	3.00	$3.00/P_0$	0.05	–1.65	7
0.05	3.00	$3.00/P_0$	0.01	–2.33	10
0.01	4.60	$4.60/P_0$	0.05	–1.65	9
0.01	4.60	$4.60/P_0$	0.01	–2.33	13

As an example, let us construct a test to determine if a given system is operating at a bit error rate of 10^{-12} or better. Given that a 5% chance of a type I error is acceptable, the test would take the form of sending $3x10^{12}$ bits and counting the number of errors. If no errors are counted, we are confident that the BER was 10^{-12} or better.

Given that a 5% chance of a type II error is acceptable, we find that k_l is 7. If more than 7 errors are counted, we are confident that the bit error rate is greater than 10^{-12}. However, what if between 1 and 7 errors are counted? These cases may be handled several different ways. One option is to make a statement about the bit error rate (whether it is less than or greater than 10^{-12}) at a lower level of confidence. Another option would be to state that the test result is success since we cannot establish with an acceptable probability of error that the BER is greater than 10^{-12}. Such a statement implies that we failed to meet the burden of proof for the conjecture that the BER exceed 10^{-12}. Of course, the burden of proof could be shifted to the device under test which would imply that any outcome other than k = 0 would correspond to failure (the device under test failed to prove to us that the BER was no more than 10^{-12}). If neither of these solutions are acceptable, it is always an option to perform a more vigorous bit error rate test in order to clarify the result.

40.F.6 – Packet Error Rate Measurement

It is often easier to measure packet errors than it is to measure bit errors. In these cases, it is helpful to have some linkage between the packet error rate and the bit error rate. To make this linkage, we assume that the bit error rate is low enough and the packet size is small enough so that each packet error contains exactly one bit error.

To complete the linkage, some care must be taken regarding how many packets to send. A bit error is only detectable in the region of the packet that is covered by the cyclic redundancy check (CRC). In the context of Ethernet, this region is the first bit of the destination address to the last bit of the CRC. There is no guarantee that errors in the preamble, start-of-frame delimiter, and inter-packet gap will be detected. Therefore, we must translate n from the number of bits are sent to the number of "observable" bits that are sent. This will increase the test duration since a portion of the time will be spent sending unobservable bits.

For packets of length x bits, at least n/x packets must be sent to perform the equivalent bit error rate test. If no packet errors are observed, the conclusion is that the bit error rate is less than P_0. If more than k_l packet errors are observed, the conclusion is that the bit error rate is greater than P_0.

Note that to reinforce the assumption that there is only one bit error per packet error, a test should be run with the shortest possible packets. However, if extremely low bit error rates are to be verified, it may be favorable to use long packets to increase the percentage of observable bits and reduce the test duration.