University of New Hampshire InterOperability Laboratory
Gigabit Ethernet Consortium

Superseded

As of September 19th, 2003 the Gigabit Ethernet Consortium Clause 40 Physical Medium Attachment Conformance Test Suite Version 1.1 has been superseded by the release of the Clause 40 Physical Medium Attachment Conformance Test Suite Version 1.2. 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Please refer to the following site for both current and superseded test suites: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/ge/

GIGABIT ETHERNET

Clause 40
Physical Medium Attachment Test Suite
Version 1.1

Technical Document

MODIFICATION RECORD

- June 18, 2003 Version 1.1
  Jon Beckwith: changed format, updated references, and corrected Appendix A. Made other minor formatting changes.
- October 8, 1999 Initial Version

ACKNOWLEDGMENTS

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

INTRODUCTION

Overview

The University of New Hampshire's InterOperability Laboratory (IOL) is an institution designed to improve the interoperability of standards based products by providing an environment where a product can be tested against other implementations of a standard. This suite of tests has been developed to help implementers evaluate the functioning of their Physical Medium Attachment on their 1000BASE-T products.

These tests are designed to determine if a product conforms to the IEEE 802.3 standard. The tests in this suite use specifications defined in Clause 40 of the 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 well in most environments.

Organization of Tests

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

Test Number

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

Purpose

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

References

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

Resource Requirements

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

Last Modification

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

Discussion

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

Test Setup

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

Procedure

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

Observable Results

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

Possible Problems

This section contains a description of known issues with the test procedure, which may effect test results in certain situations.

TABLE OF CONTENTS

Group 1: PMA Electrical Specification

Scope: The following tests are those currently implemented by the University of New Hampshire InterOperability Lab.

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

Test 40.6.1.1 – Peak Differential Output Voltage and Level Accuracy

Purpose: To verify correct transmitter output levels.

References:

      [1] IEEE standard 802.3-2002, subclause 40.6.1.1.2, figure 40.19 – Test modes
      [2] Ibid., subclause 40.6.1.1.3, figure 40-22 – Test fixtures
      [3] Ibid., subclause 40.6.1.2.1 – Peak differential output voltage and level accuracy
      [4] Test suite appendix 40.6.A – 1000BASE-T transmitter test fixtures

Resource Requirements:

      • Refer to test suite appendix 40.6.A

Last Modification: March 6, 2003 (version 1.1)

Discussion:

Subclause 40.6.1.1.2 of IEEE standard 802.3-2002 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 figure 40-19 of the standard.

Test mode 1 is used to verify the 1000BASE-T transmitter waveform in terms of differential output voltage and droop. In this test, the peak differential output voltage is measured at points A, B, C, and D as indicated in figure 40-19. The magnitude of the voltage at points A and B shall be between 670 and 820 mV. 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%. The magnitudes of the voltages at points C and D shall not differ from 0.5 times the average of the voltage magnitudes at points A and B by more than 2%.

Test Setup: Refer to test suite appendix 40.6.A.

Test 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. Measure the peak voltage of the waveform at points A, B, C, 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:

    • The magnitude of the voltage at points A and B shall be between 670 and 820 mV.
    • The magnitudes of the voltage at point B and shall not differ from the magnitude of the voltage at point A by more than 1%.
    • 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%.
    • 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 40.6.1.2 – Maximum Output Droop

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

References:

> [1] IEEE standard 802.3-2002, subclause 40.6.1.1.2, figure 40.19 – Test modes
> [2] Ibid., subclause 40.6.1.1.3, figure 40-23 – Test fixtures
> [3] Ibid., subclause 40.6.1.2.2 – Maximum output droop
> [4] Test suite appendix 40.6.A – 1000BASE-T transmitter test fixtures

Resource Requirements:

- Refer to test suite appendix 40.6.A.

Last Modification: March 6, 2003 (version 1.1)

Discussion:

Subclause 40.6.1.1.2 of IEEE standard 802.3-2002 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 figure 40-19 of the standard.

Test mode 1 is used to verify the 1000BASE-T transmitter waveform in terms of differential output voltage and droop. In this test, the differential output voltage is measured at points F, G, H, and J as indicated in figure 40-19. Points G and J are exactly 500 ns from points F and H respectively. The magnitude of the voltage at point G shall be greater than 73.1% of the magnitude of the voltage at point F. The magnitude of the voltage at point J shall be greater than 73.1% of the magnitude of the voltage at point H.

Test Setup: Refer to test suite appendix 40.6.A.

Test 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. Measure differential output voltage at points F, G, H, 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:

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

Possible Problems: None.

Test 40.6.1.3 – Differential Output Templates

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

References:

[1] IEEE standard 802.3-2002, subclause 40.6.1.1.2, figure 40.19 – Test modes
[2] Ibid., subclause 40.6.1.1.3, figure 40-22 – Test fixtures
[3] Ibid., subclause 40.6.1.2.3, figure 40-26, tables 40-10 and 40-11 – Differential output templates
[4] Test suite appendix 40.6.A – Transmitter test fixtures

Resource Requirements:

- Refer to test suite appendix 40.6.A.

Last Modification: March 6, 2003 (version 1.1)

Discussion:

Subclause 40.6.1.1.2 of IEEE standard 802.3-2002 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 figure 40-19 of the standard.

Test mode 1 is used to verify the 1000BASE-T transmitter waveform in terms of differential output voltage and droop. In this test, the waveforms around points A, B, C, D, F, and H, as shown in figure 40-19, are compared to the differential output templates shown in figure 40-26. The waveforms around points A, B, C, and D are compared to normalized time domain transmit template 1 after the following normalization factors are applied:

- The waveform around point A is normalized to the peak voltage at point A.
- The waveform around point B is normalized to the negative of the peak voltage at point A.
- The waveform around point C is normalized to 0.5 times the peak voltage at point A.
- The waveform around point D is normalized to the negative of 0.5 times the peak voltage at point A.

The waveforms around points F and H are compared to time domain transmit template 2 after the following normalization factors are applied:

- The waveform around point F is normalized to the peak voltage at point F.
- The waveform around point H is normalized to the peak voltage at point H.

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

Test Setup: Refer to test suite appendix 40.6.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:

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

Possible Problems: None.

Test 40.8.3.1 – MDI Return Loss

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

References:

　　[1] IEEE Std 802.3-2002: Section 40.8.3.1

Resource Requirements:

- Network analyzer
- Post-processing PC
- Connecting Hardware

Last Modification: March 6, 2003 (version 1.1)

Discussion:

A 1000BASE-T device compliant with the IEEE 802.3-2002 standard 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 results in reflections of the transmitted signals. Because the impedances can never be exactly 100Ω, and because the termination impedance varies with frequency, some limited amount of reflections must be allowed. Return loss is a measure of the signal power that is reflected due to the impedance mismatch. 802.3-2002 specifies that the reflected power at the MDI must be at least 16 dB less than the incident power over the range of 1.0 to 40MHz. The return loss must be at least $10 - 20*\log_{10}(f/80)$ dB from 40 to 100MHz. This 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: Setup the devices as shown in Figure 40.8.3.1 using test jig #2 as shown in figure 40.8.3.2.


Figure 40.8.3.1: Return loss test setup

Figure 40.8.3.2: Test Jig 2

In order to perform tests for all four pairs, four short cables were constructed. Each cable had a standard termination on one end, and the pair to be tested terminated on pins 1,2 on the other end. The effect of the cable is removed during calibration of the Network Analyzer.

Procedure:

1. The DUT is put into test mode 4 (as defined in section 40.6.1.1.2) and the BI_DA pair is connected to the reflection port of a network analyzer. This ensures that the 1000BASE-T physical layer is active and transmitting data.
2. The network analyzer measures the reflections at the MDI referenced to a 50Ω characteristic impedance.
3. Data is processed on the PC to calculate the reflections for characteristic impedances of 85, and 115Ω.
4. Repeat above steps for the BI_DB, BI_DC, and BI_DD pairs.

Observable Results:

- The return loss measured at each MDI shall be at least 16 dB from 1.0 to 40 MHz, and at least $10 - 20 * \log_{10} (f /80)$ dB from 40 to 100MHz when referenced to a characteristic impedance of $100Ω \pm 15\%$.

Possible Problems: None.

Appendix 40.6.A – 1000BASE-T Transmitter Test Fixtures

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

References:

[1] IEEE standard 802.3-2002

Resource Requirements:

- Disturbing signal generator, Tektronix AWG2021 or equivalent
- Digital storage oscilloscope, Tektronix TDS7104 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: June 13, 2003 (version 1.1)

Discussion:

40.6.A.1 – Introduction

Subclause 40.6.1.1.3 of IEEE standard 802.3-2002 defines 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 100 Ω 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 only differ in the specification of the disturbing signal and the inclusion of a highpass test filter. The test fixture characteristics are given in table 40.6.A-1.

Table 40.6.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 Figure 40-22 of the IEEE 802.3-2002 Standard, 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 100 Ω 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 40.6.A-1, and the modular break out board used is shown in figure 40.6.A-2. Each test fixture is realized through the settings of the disturbing voltage generator and configuration of the post-processing block.


Figure 40.6.A-1: Test setup block diagram

Figure 40.6.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 2 V peak-to-peak DUT output using test fixture 3, the vertical range of the oscilloscope must be set to at least 7.4 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 3. 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 40.6.A-1 uses power splitters. Like 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 $2.8V_{p-p}$. 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 $1.4V_{p-p}$. 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 $1.4V_{p-p}$.

Figure 40.6.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 2 V peak-to-peak than 7.4 V peak-to-peak, the quantization error on the DUT output will be smaller.

The test setup block diagram in figure 40.6.A-1 may be implemented with the equipment listed in table 40.6.A-2. The remainder of this appendix discusses the test setup in the context of this implementation.


Figure 40.6.A-3: Power splitter operation

Table 40.6.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, 8 GS/s sample rate, 16 million sample memory
Power splitter	Mini-Circuits ZSC-2-1W	2-way 0^o, 1 to 650 MHz

40.6.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 50 Ω so that a differential 100 Ω 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 40.6.A-4. The measurement was performed using the HP 8753C network analyzer with the HP 85047A S-parameter test set. 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.

From figure 40.6.A-4, it can be seen that the power splitters are well matched to about 700 MHz. Also, the insertion loss is about 3.2 dB from 1 to 150 MHz. Note that the power splitters are AC-coupled devices. The highpass cut-off frequency of the power splitters must be known so that their impact on droop measurements can be removed. Since the network analyzer is also AC-coupled, the test setup shown in figure 40.6.A-5 was used to properly measure the low-frequency response.

The test setup shown in figure 40.6.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° out of phase with the input to power splitter A. The LeCroy LC584AXL 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 40.6.A-6. The differential input amplitude was 200 mV.


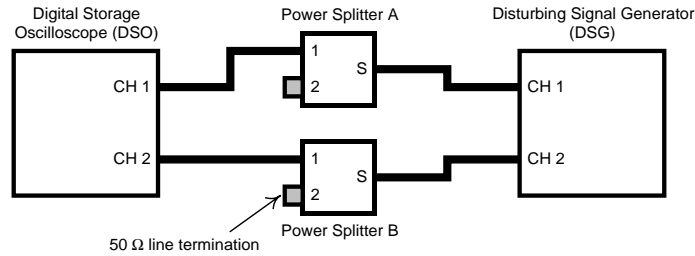Figure 40.6.A-4: Power splitter insertion loss

Figure 40.6.A-5: Test setup for highpass cut-off frequency measurement

The highpass cut-off frequency 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.


Figure 40.6.A-6: Low-frequency response of power splitter pair

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 50 Ω 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 40.6.A-7 and 40.6.A-8 respectively. These figures show that all harmonics are at least 40 dB below the fundamental.


Figure 40.6.A-7: Spectrum of 31.25 MHz synthesized sine wave from the Tektronix AWG2021

Figure 40.6.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 40.6.A-3.

Table 40.6.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	1.25 V peak-to-peak	2.5 V peak-to-peak
Filter	50 MHz	50 MHz
Offset	0	0

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

A 2041 may be used in replacement of the 2021 to transmit the disturbing signal for test fixture 1 & 2. Test Fixture 3 requires a peak-to-peak amplitude greater than 2 V, and therefore the 2041 cannot be used.

Table 40.6.A-4: Tektronix AWG2041 channel 1 settings

Setting	Test Fixtures 1 and 2
Sample Rate	1 GS/s
Samples Per Cycle	32
Amplitude	1.1 V peak-to-peak
Filter	50 MHz
Offset	0

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

The linearity of the Tektronix AWG2021 was tested using the setup shown in figure 40.6.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 40.6.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 4 V peak-to-peak.


Figure 40.6.A-9: Disturbing signal generator linearity measurement

Figure 40.6.A-10: Spectrum at port 3 of the resistive splitter

40.6.A.4 – Digital Storage Oscilloscope

A digital storage oscilloscope (DSO) with at least 3 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 50 Ω 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.

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

Another common 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 2 GS/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 2 GS/s sample rate, this would require a sample memory of 32 K-samples. Deeper sample memories are useful for jitter measurements, but that is beyond the scope of this appendix.

40.6.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 highpass test filter when required.

Figure 40.6.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 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.

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 40.6.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 highpass 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 40.6.A-13.

Figure 40.6.A-11: Observed transmitter test mode 1 waveform before post-processing


Figure 40.6.A-12: Input waveform and scaled output waveform with best-fit sine wave removed

Figure 40.6.A-13: Output waveform with droop compensation

Figure 40.6.A-14: Output of transmitter test filter

Note 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 highpass 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 40.6.A-14.

40.6.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 break-out 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 50 Ω.

The reflection coefficient of each circuit with respect to a 50 Ω 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 40.6.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 40.6.A-1.

$$\text{Gain Error} = \frac{Z_A}{50 + Z_A} + \frac{Z_B}{50 + Z_B} \qquad \text{(Equation 40.6.A-1)}$$

Equation 40.6.A-1 assumes that the differential source impedance is a precisely balanced 100 Ω resistance. The impedance of each circuit was derived from the reflection coefficient and the gain error is plotted in figure 40.6.A-16.

In section 40.6.A.2, the frequency response of the power splitters was measured for each differential component and again as a pair. Comparing figures 40.6.A-4 and 40.6.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 40.6.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 40.6.A-17.

Figure 40.6.A-15: Differential return loss at the input to the test setup


Figure 40.6.A-16: Differential gain error due to impedance imbalance in the test setup

Figure 40.6.A-17: Test setup common-mode rejection

40.6.A.7 – Conclusion

This appendix presented a reference implementation for test fixtures 1 through 4 from IEEE standard 802.3-2002. 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 40.6.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 highpass 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 40.6.A-5: Realization of 1000BASE-T Transmitter Test Fixtures

Setting	Test Fixture 1	Test Fixture 2	Test Fixture 3	Test Fixture 4
Tektronix AWG2021				
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)	1.25 V	1.25 V	2.5 V	—
Offset	0	0	0	—
Tektronix AWG2041				
Sample Rate	1 GS/s	1 GS/s	—	—
Samples Per Cycle	32	32	—	—
Filter	50 MHz	50 MHz	—	—
Amplitude (peak-to-peak)	1.11 V	1.11 V	—	—
Offset	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: The settings for channels 1 and 2 of the AWG2021 and AWG2041 are identical except for a 180^o phase-shift.

Note 2: If the peak-to-peak amplitude of the disturbing signal is not correct at port S, the amplitude of the AWG may be changed accordingly.