Date: 2/4/05

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli)

Subject: SAS 1.1 Signal Performance Measurements Annex

Annex A
(normative)

Signal Performance Measurements

A.1 Introduction

This annex specifies the configuration requirements for making electrical performance measurements. These measurements consist of signal output, signal tolerance, and return loss. Standard loads are used in all cases so that independent specification of connection components and transportability of the measurement results are possible.

NOTE 1 - The methodology for making return loss measurements are specified in this annex although this standard does not define any return loss values. Statements in this annex that imply return loss values are specified should be ignored.

A.2 A simple connection

A.2.1 Overview

In the basic structure considered the physical link consists of three component parts:

a) the transmitter device;
b) the interconnect; and
c) the receiver device.

Each of those components is connected by a separable connector. On a duplex connection signals travel in opposite directions down the same nominal path.

Figure A.1 shows a duplex link and the location of the connectors.

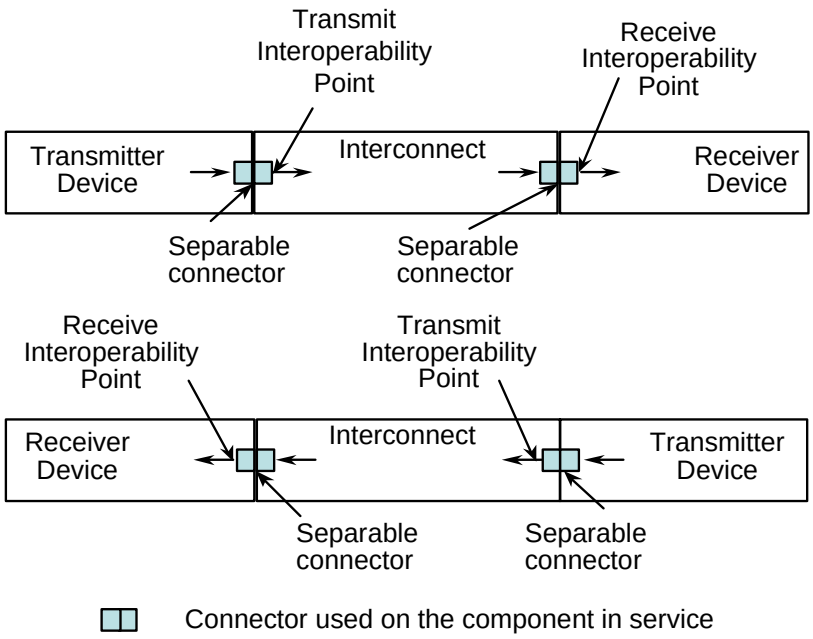

Figure A.1 — A simple duplex link physical connection

Since connectors are always used in the mated condition the only access to the signals is before the signal enters the mated connector (i.e. upstream) or after the signal exits the connector (i.e. downstream). Even if signals were able to be practically accessed at the point of mating within the connector such access would disturb the connector to the point that the measurement of the signal would be compromised (e.g., attempting to access the unmated connector with probes does not provide valid results because the connector is not the same when unmated as when mated and the probe contact points are not at the same location as the connector contact points). Using probes the contacts are not deflected and the shields are not connected in unmated electrical connectors.

In this annex the signal outputs are always measured downstream of the mated connector (see figure A.1) so that the contribution of the connector to the signal properties is included in the measurement. This approach assigns a portion of the connector losses to the upstream component but it also makes the signal measurement conservative. If the connectors on both ends of the interconnect are the same the additional loss at the downstream connector is balanced by the reduced loss at the upstream connector. For transmitter devices a slightly stronger transmitter is required to pass the signal through the downstream half of the connector that does not belong to the transmitter device. The signal coming into receiver devices is specified after the signal has passed through the connector.

Examination of the details of the measurement methods described in this annex shows that the mated connector issue may not be as severe as it appears.

The TxRx connection has an assumed reference impedance of 100 ohms

A.2.2 Assumptions for the structure of the transmitter device and the receiver device

Figure A.2 shows the details of a transmitter device.

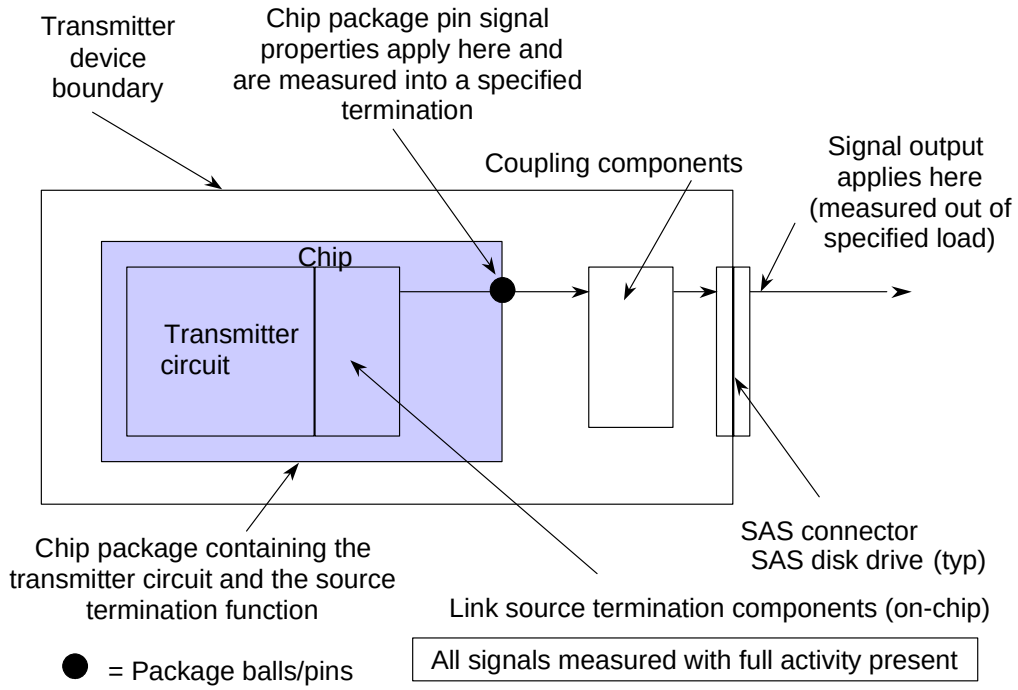

Figure A.2 — Transmitter device details for a disk drive type interoperability point

As figure A.2 shows any of the following internal parts of this transmitter device may be labeled as the transmitter:

a) the transmitter circuit in the chip;
b) the chip itself; and
c) the chip and its associated chip package

The term transmitter is therefore not well defined and is not used in the terminology without a modifier.

The transmitter device contains:

a) a connector (i.e., half a mated pair);
b) optional coupling circuits;
c) the source termination;
d) the transmitter circuit;
e) PCB traces and vias;
f) the chip package; and
g) possibly ESD devices.

It is assumed that the source termination is contained within the chip package.

Interoperability points may be defined at the chip package pins in some network standards (e.g., Ethernet XAUI). This standard does not define requirements at chip package pins.

Figure A.3 shows the details of a receiver device. It is similar to the transmitter device.

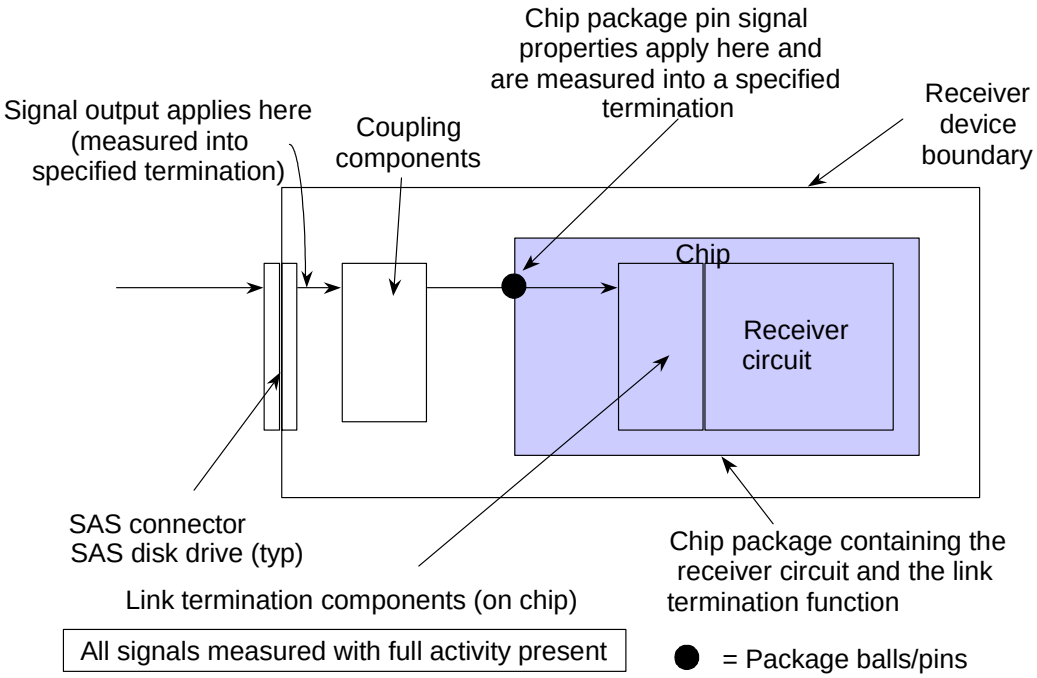

Figure A.3 — Receiver device details for a disk drive type interoperability point

As figure A.3 shows any of the following internal parts of this receiver device may be labeled as the receiver:

a) the receiver circuit in the chip
b) the chip itself
c) the chip and its associated chip package

The term receiver is therefore not well defined and is not used in the terminology without a modifier.

The receiver device contains:

a) a connector (i.e., half a mated pair);
b) optional coupling circuits;
c) the link termination;
d) the receiver circuit;
e) PCB traces and vias;
f) the chip package; and
g) possibly ESD devices.

It is assumed that the link termination is contained within the chip package.

A.2.3 Definition of receiver sensitivity and receiver device sensitivity

The term receiver sensitivity is problematic in common usage. This term is not used for interoperability in standards. A related term applicable to the receiver device input signal is the receiver device sensitivity. While these two terms are related they are significantly different because of the noise environment assumed. The description in this subclause is used to uniquely define these terms with the understanding that this standard discourages usage of either term.

Receiver sensitivity:

a) The receiver in the receiver sensitivity refers to signal properties at the chip package pin for the chip package that contains the receiver circuit;
b) Receiver sensitivity is defined as the minimum vertical inner eye opening at which the receiver chip delivers the required BER. The horizontal eye opening shall be minimum (i.e., maximum jitter present) and all activity is quiesced except for the receiver itself; and

c) Receiver sensitivity is not defined in the SAS context because there are no chip pin specifications.

Receiver device sensitivity:

The term receiver device sensitivity is defined as the minimum vertical inner eye opening measured at the signal output point for the input to the receiver device at which the receiver chip (i.e., the chip in the chip package on the board containing the receiver device interoperability point) delivers the required BER with all activity expected in the application for the receiver circuit present (i.e., not quiesced as for the receiver sensitivity definition), with the CJTPAT (see annex A), and the minimum horizontal eye opening in the signal at the receive device interoperability point.

Special test conditions are required to measure these sensitivities (see A.8). The terminology used in this standard is signal tolerance instead of receiver device sensitivity.

A.3 Measurement architecture requirements

A.3.1 General

Signal specifications are only meaningful if the signals are able to be measured with practical instrumentation. Another requirement is that different laboratories making measurements on the same signal get the same results within acceptable measurement error (i.e., the measurements must be accessible, verifiable, and transportable). As of the publishing of this standard there are no accepted standard for creating signals with traceable properties and with all the properties required for an effective signal specification architecture for high speed serial applications.

Some of the elements required for practical, verifiable, and transportable signal measurements are included in this standard.

Having signal specifications at interoperability points that do not depend on the actual properties of the other link components not under test requires that specified known loads be used for the signal measurements. In service, the load presented to the interoperability point shall be that of the actual component and environment existing in service.

Interfacing with practical instruments also requires that specified impedance environments be used. This forces a signal measurement architecture where the impedance environment is 50 single ended (i.e., 100 ohms differential). It also forces the requirement for instrumentation quality loads of the correct value.

Instrumentation quality loads are readily available for simple transmission line termination. However, none are available for more complex loads that include specified propagation time, insertion loss properties, crosstalk properties, and jitter creation properties.

For signal tolerance measurements the signal shall be calibrated before applying it to the interoperability point under test. This calibration is done by adjusting the properties of the signal measured out of a known load (i.e., just like the signal output case) and then removing the known load and applying the signal unchanged to the interoperability point under test. It is assumed that any changes to the signal from the calibration state to the measurement state are caused by the interoperability point under test and is therefore part of the performance sought by the measurement.

A.3.2 Relationship between signal compliance measurements at interoperability points and operation in systems

The signal measurements in this standard and return loss measurements in this annex apply under specified test conditions that simulate some parts of the conditions existing in service (e.g., this simulation includes duplex traffic on all ports and under all applicable environmental conditions). Other features existing in service (e.g., non ideal return loss in parts of the link that are not present when measuring signals in the specified test conditions) are included in the specifications themselves. This methodology is required to give each side of the interoperability point signal performance requirements that do not depend on knowing the properties of the other side.

Measuring signals in an actual functioning system at an interoperability point does not verify compliance for the components on either side of the interoperability point although it does verify that the specific combination of components in the system at the time of the measurement produces compliant signals. Interaction

between components on either side of the interoperability point may allow the signal measured to be compliant but this compliance may have resulted because one component is out of specification while the other is better than required.

It is recommended that additional margin be allowed when performing signal compliance measurements to account for conditions existing in service that may not have been accounted for in the specified measurements and specifications.

A.4 De-embedding connectors in test fixtures

Connectors are part of the test fixtures required for obtaining access to the interoperability points. This is intrinsic for most practical measurements because the connectors used on the service components are different from those used on the instrumentation.

The effects of the portions of the connector that is used on the test fixture need to be accounted for in order to not penalize the point under test by the performance of the test fixture connector. This accounting process is termed de-embedding in this subclause.

Figure A.4 shows two cases that apply.

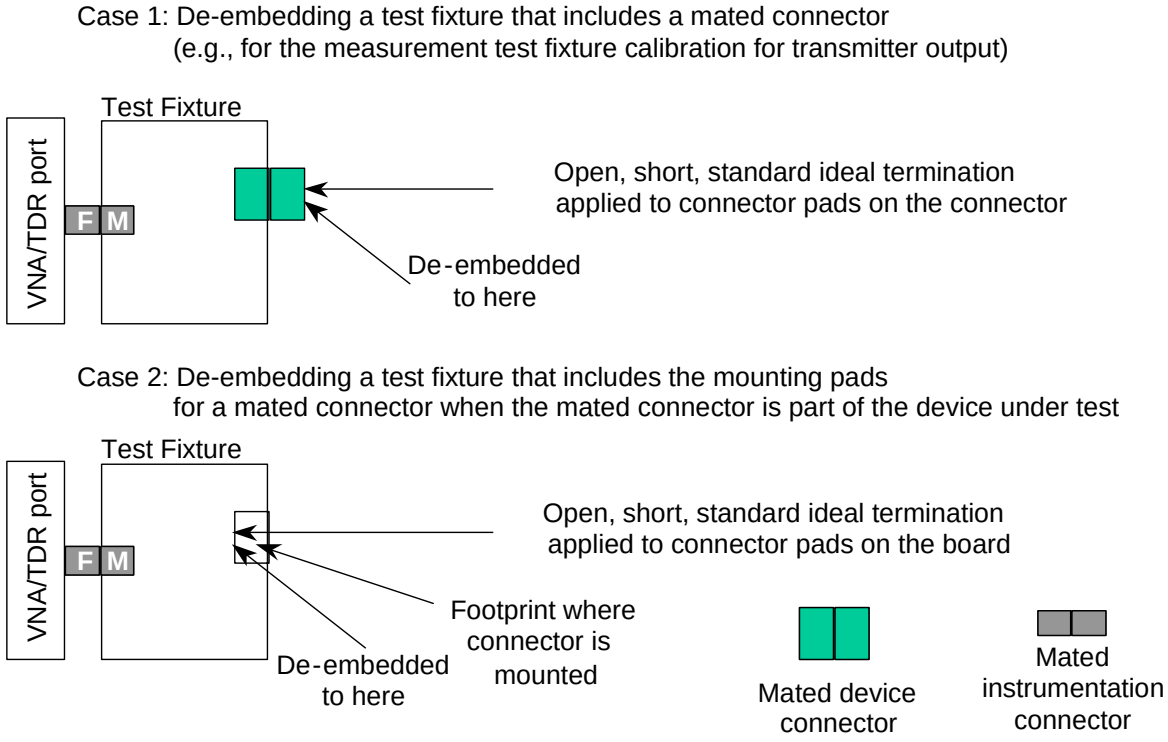

Figure A.4 — De-embedding of connectors in test fixtures

The de-embedding process assumes that the test fixture is linear and that S parameter methodologies are used. Fundamentally an S parameter model for the test fixture with or without the connector in place is the result. Knowing this model for the test fixture, with or without the connector in place, allows simulation of the impact of the test fixture on the signal measurement.

A.5 Measurement conditions for device signal output at the transmitter device

The measurement conditions required for a differential transmitter device signal output are shown in figure A.5. The two required cases described in figure A.5 are as follows:

a) the transmitter device is directly attached to the receiver device; and

b) the transmitter device is attached to the receiver device through an interconnect assembly (e.g., cable assembly or PCB).

To simulate some of the properties of the interconnect assembly an instrumentation quality compliance interconnect is used. This compliance interconnect is assumed embedded in the compliance interconnect test fixture as shown in more detail in figure A.6.

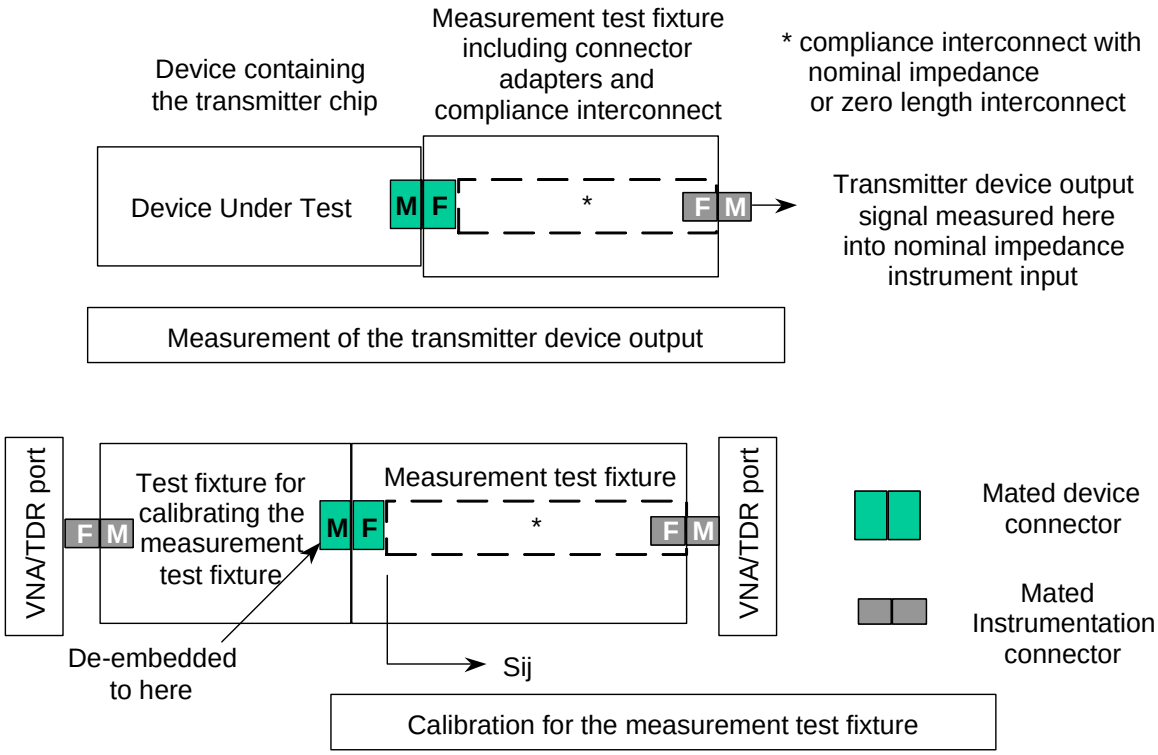

Figure A.5 — Measurement conditions for transmitter device signal output specifications

In the lower portion of figure A.5 a cable assembly connecting the measurement test fixture to the instrumentation port is shown. This cable assembly is considered part of the instrumentation and is not specifically shown in the top portion of figure A.5, figure A.6, figure A.7, figure A.8, figure A.9, figure A.10, nor figure A.11. The gender of the connector that connects the instrument or the instrument plus the connecting cable to the set up may need to be changed in specific instrumentation connections.

A measurement test fixture may be constructed from a
standard compliance interconnect with instrumentation connectors
and a connector adapter as shown in this figure.
This gives both zero length and compliance interconnect parts.

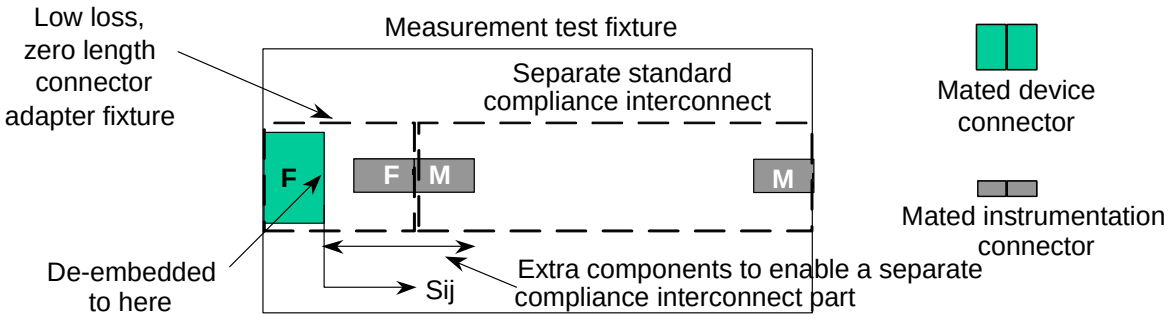

Extra components add loss and delay compared to
the standard compliance interconnect. By using high quality
instrumentation connectors and short, low loss designs the extra
components may be tolerable. Extra components make it more difficult
for the transmitter to meet the required output specs.

Figure A.6 — Transmitter device output signal measurement test fixture details

A.6 Measurement conditions for signal tolerance at the transmitter device

The measurement conditions required for the signal tolerance at the differential transmitter device interoperability point are shown in figure A.7. The two required cases described in figure A.7 are as follows:

a) the transmitter device is directly attached to the receiver device; and
b) the transmitter device is attached to the receiver device through an interconnect assembly (e.g., cable assembly or PCB).

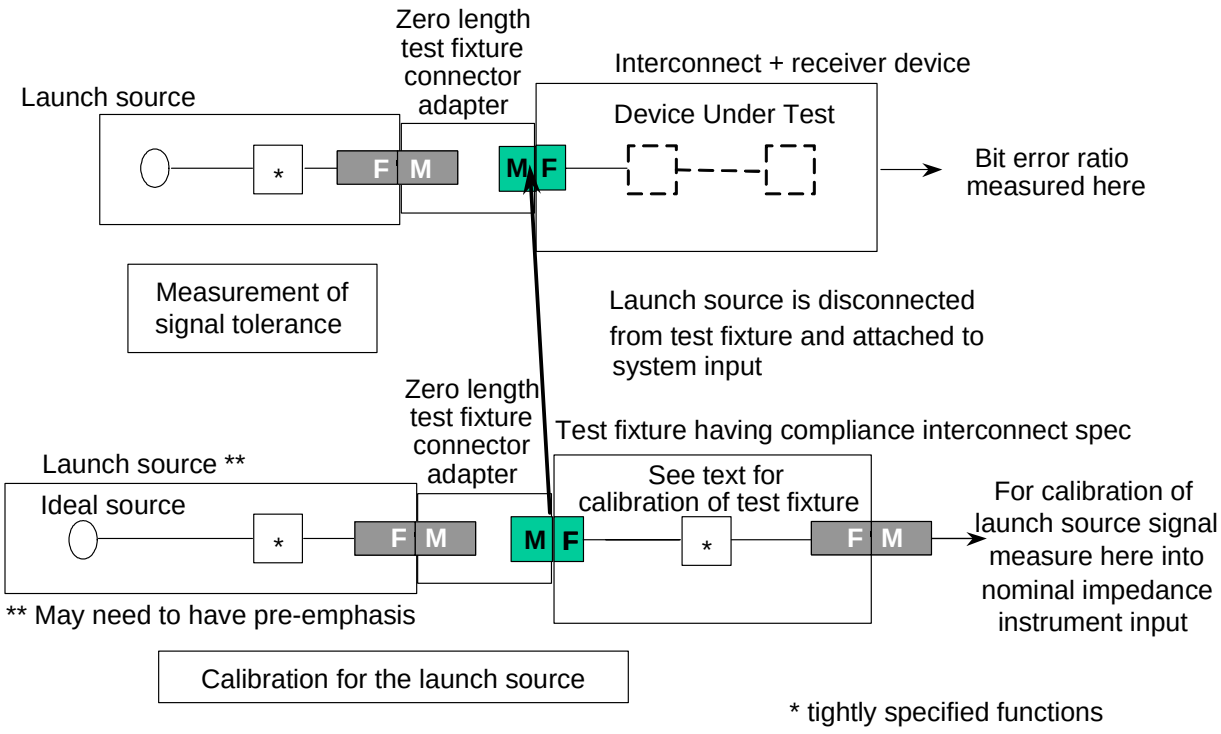

Figure A.7 — Measurement conditions for system (interconnect + receiver device) signal tolerance

The test fixture for this measurement is shown in figure A.8.

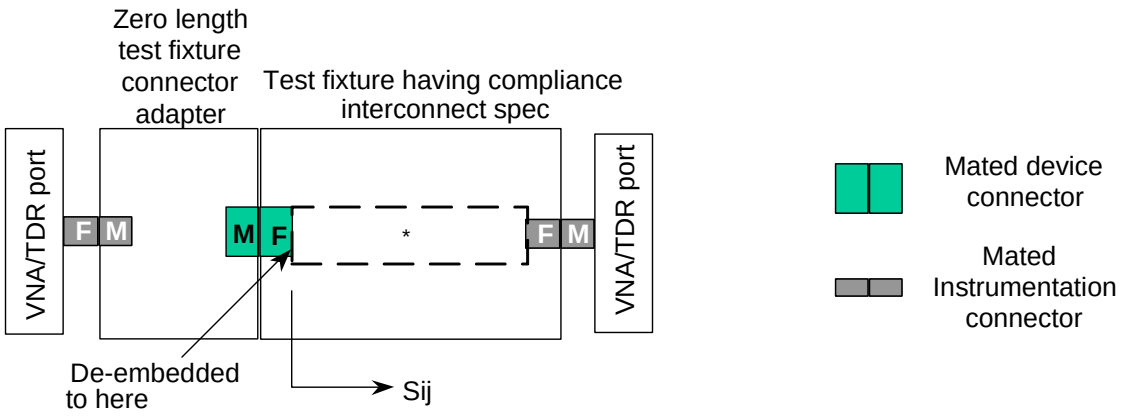

Note: This test fixture is identical to the measurement test fixture used for the transmitter device output signal (i.e., the connector genders and pins used in the SAS connector are the same).

* tightly specified function

Figure A.8 — Calibration of test fixture having compliance interconnect specification for transmitter device (interconnect + receiver device) signal tolerance

A.7 Measurement conditions for signal output at the receiver device

The measurement conditions for the signal output at the receiver device are shown in figure A.9.

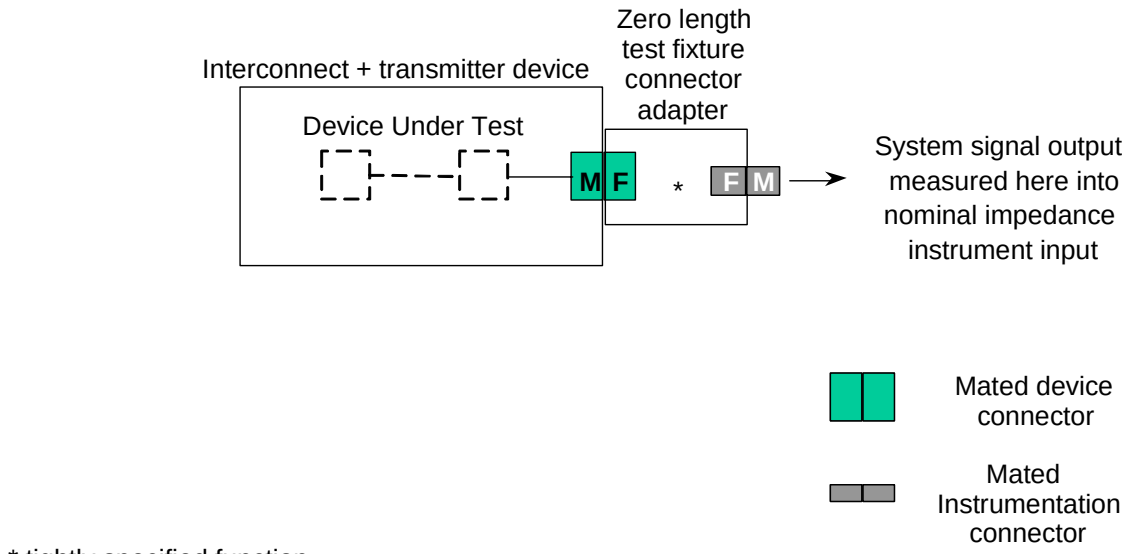

* tightly specified function

Figure A.9 — Measurement conditions for signal output at the receiver device

The interconnect may be the zero length interconnect where the transmitter device is connected directly to the receiver device.

A.8 Measurement conditions for signal tolerance at the receiver device

The measurement conditions required for the signal tolerance at the differential receiver device interoperability point are shown in figure A.10. The two required cases described in figure A.10 are as follows:

a) the transmitter device is directly attached to the receiver device; and
b) the transmitter device is attached to the receiver device through an interconnect assembly (e.g., cable assembly or PCB).

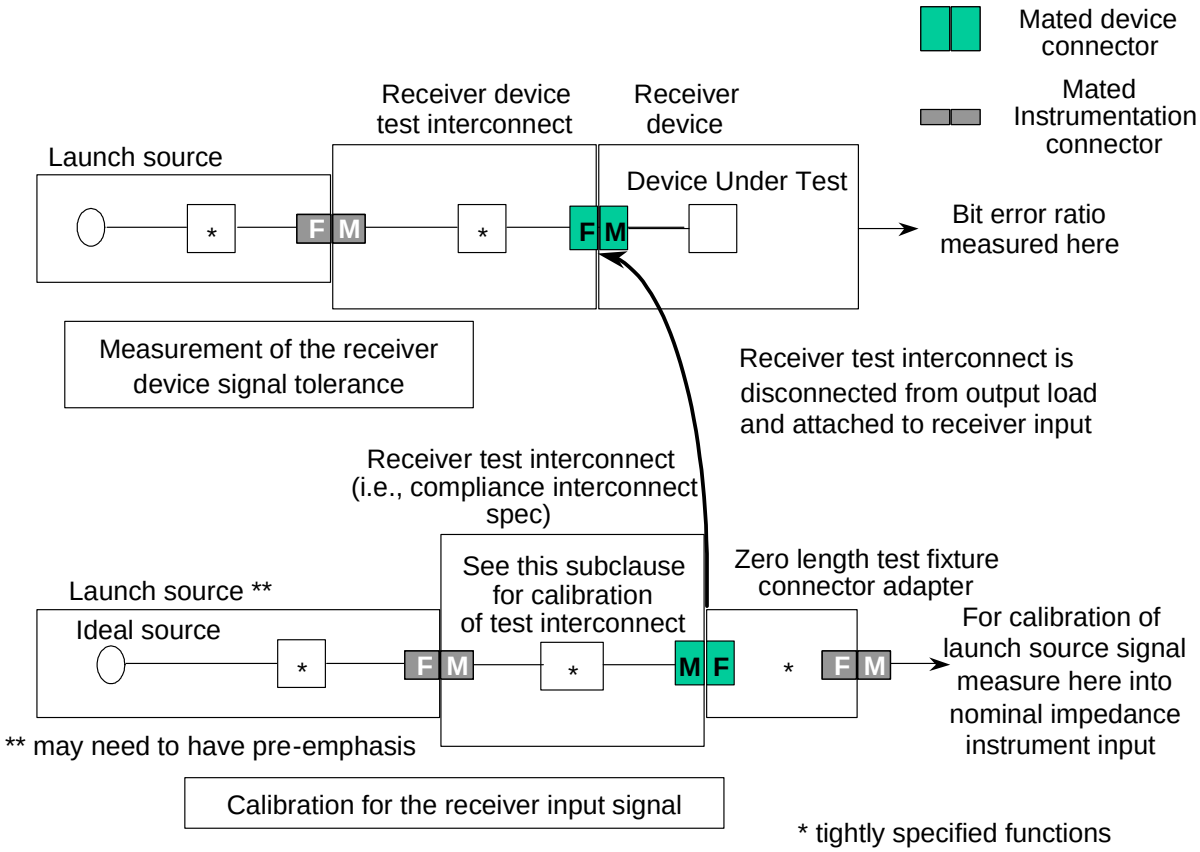

Figure A.10 — Measurement conditions for signal tolerance at the receiver device

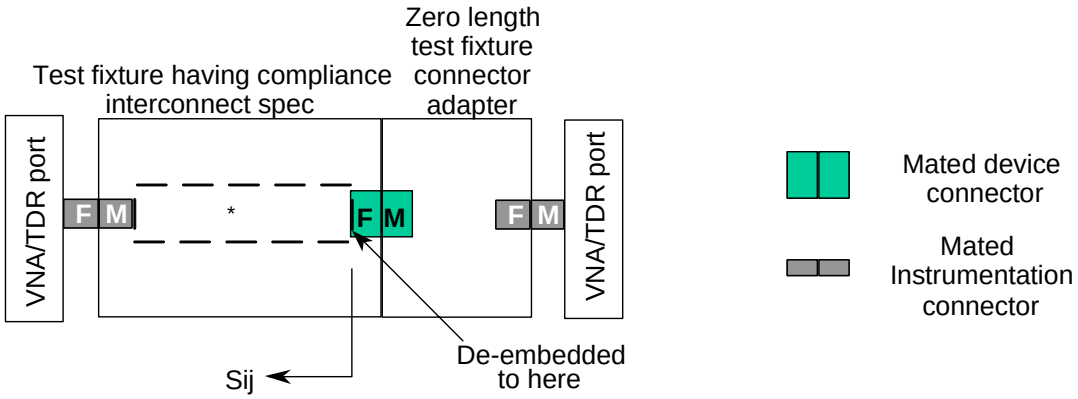

Note:

1 - This is not identical to the measurement test fixture used for the transmitter output signal even though the connector genders are the same. The pins used in the SAS connector are for the Rx (i.e., not the TX) and the signals flow the other way. The S22 measurement here is the same as the S11 measurement for the transmitter output signal but on different pins.

2 - The S21 and S12 are used to create the desired jitter in this application and are not as critical.

* tightly specified function

Figure A.11 — Calibration of receiver test interconnect test fixture for signal tolerance at the receiver device

A.9 S-parameter measurements

A.9.1 Introduction

Properties of link elements that are linear may be represented by S-parameter spectra. There are two problematic areas when applying S-parameters to differential electrical links:

a) Naming conventions; and
b) Use of single ended vector network methods on differential and common mode systems.

This subclause explores both of these areas.

Measurement architecture for the most common conditions are described in some detail.

A.9.2 Naming conventions in high speed serial links

Significant confusion exists concerning the naming of Sij. There is general agreement that Sij is the ratio of the signal coming out of the ith port to the signal coming into the jth port (e.g., for a 2 port linear element having ports i and j with the signals being differential or common mode for SAS electrical links). The confusion may happen when numbers are assigned to i and j in specific cases. Another confusion factor may come from naming the type of measurement to be performed.

There are two types of measurement:

a) return loss from the same port of the element; and
b) transfer function or insertion loss across the element.

A return loss measurement may be referred to as an S11 measurement and the transfer function or insertion loss may be referred to as S21.

When a return loss measurement is performed on port 2 of the element the measurement is reporting the S22 property of the element even though it is exactly the same kind of measurement that is done for the S11 of the element on port 1.

A port number convention is used where the downstream port is always port 2 and the upstream port is always port 1. The stream direction is determined by the direction of the primary signal launched from the transmitter device to the receiver device.

Measurement types should not be referred to as S11 or S21 but rather by the return loss, insertion loss, or transfer function.

Figure A.12 shows the port naming conventions for link elements, loads, and where those elements exit.

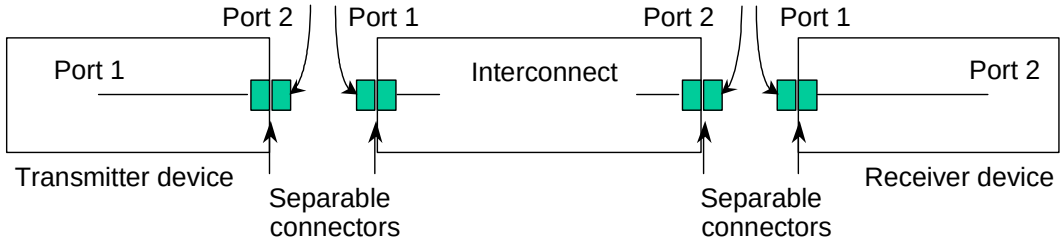

Port definitions for loads used for signal output testing and S-parameter measurements in multiline configurations

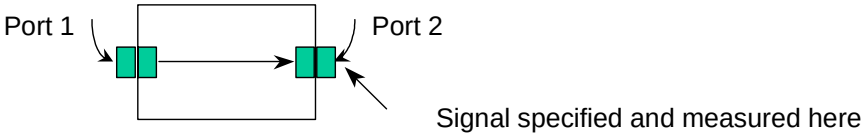

This load has ideal differential and common mode properties

Note: The transmitter device port 1 and receiver device port 2 are internal and are not defined. They are required to be an ideal source and an idea sink respectively.

Figure A.12 — Sij nomenclature conventions

A.9.3 Use of single ended instrumentation in differential applications

Figure A.13 shows the connections that are made to a four port vector network analyzer (VNA) or a time domain network analyzer (TDNA) for measuring S parameters on a four single ended port black box device. These analyzers recognize incident signals denoted by the ‘A’ subscript and reflected signals from the same port denoted by the ‘B’ subscript.

All the measurements specified in this standard relate to differential signal pairs. It requires all four analyzer ports to measure the properties of two differential ports.

VNA ports are all single ended and the differential and common mode properties for differential ports are calculated internal to the VNA. If using a TDNA consult the details for the specific instrument.

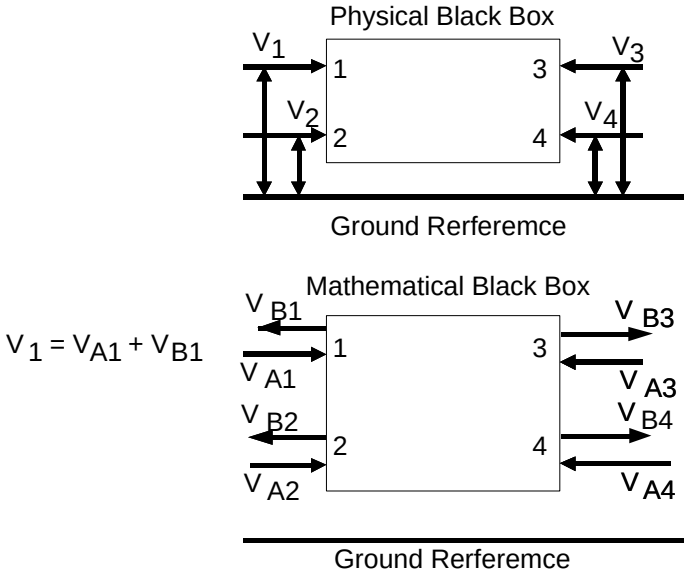

Figure A.13 — Four single ended port or two differential port element

A.9.4 Measurement configurations for link elements

A.9.4.1 Overview

Special test fixtures are required to make S-parameter measurements partly because the connectors used on real link elements are different from those used on instrumentation . The goal is for these test fixtures to be as invisible as possible.

See A.9.4 for the description of the measurement configurations. All of these measurements are return loss in this annex. A more complete set of S-parameters is used as part of the calibration process for test fixtures.

A.9.4.2 Transmitter device return loss

Figure A.14 shows the configuration to be used for the transmitter device return loss measurement.

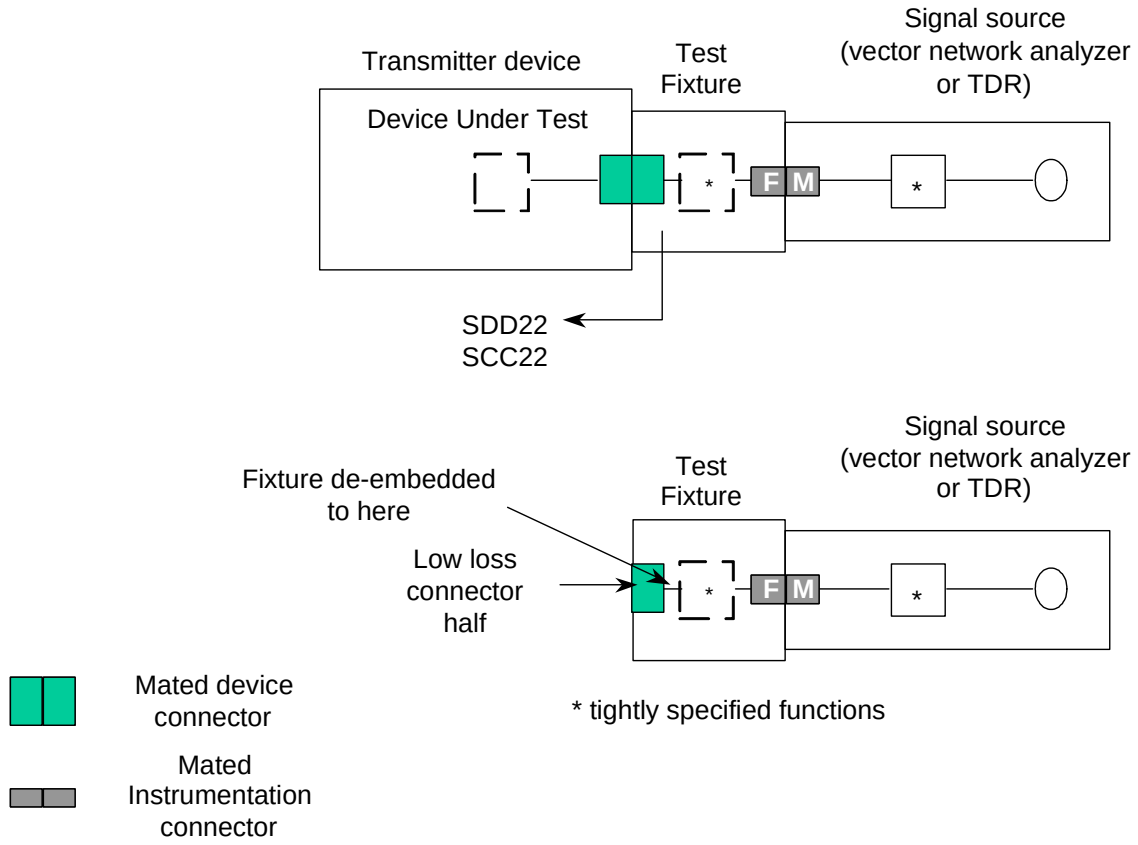

Figure A.14 — Measurement conditions for upstream return loss at the transmitter device connector

The test fixture in figure A.14 uses low loss connectors to avoid penalizing the transmitter device under test for the test fixture half of the connector. If the test fixture half of the device connector is poor then the transmitter device has to compensate for those losses.

The test fixture losses up to the mounting points for the device connector are de-embedded using the methods described in figure A.4.

A.9.4.3 Receiver device return loss

Figure A.15 shows the configuration to be used for the receiver device return loss measurement.

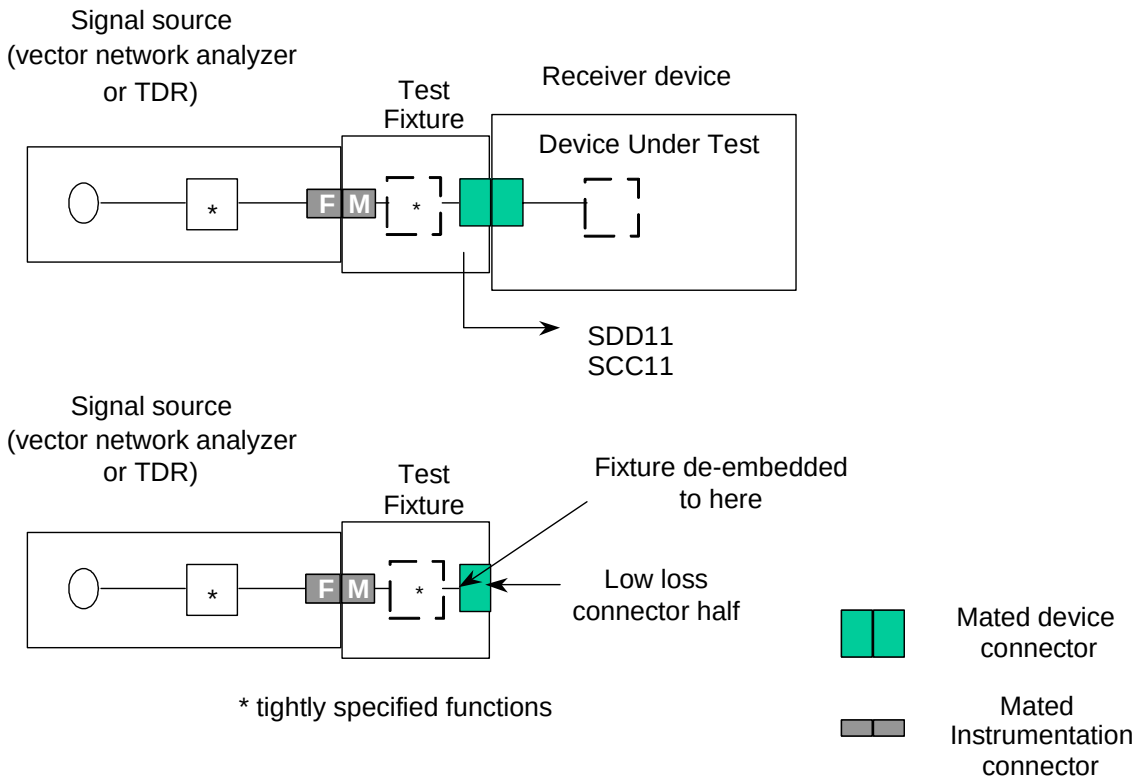

Figure A.15 — Measurement conditions for downstream return loss at the receiver device connector

The test fixture in figure A.15 uses low loss connectors to avoid penalizing the receiver device under test for the test fixture half of the connector. If the test fixture half of the device connector is poor then the receiver device has to compensate for those losses.

The test fixture losses up to the mounting points for the device connector are de-embedded using the methods described in figure A.4.

A.9.4.4 Return loss at the transmitter device connector (interconnect input)

Figure A.16 shows the conditions for making the return loss measurement into the interconnect attached to the transmitter device.

This measurement, like the signal tolerance measurement at the transmitter device connector, requires both the interconnect and the receiver device to be in place and the combination is measured. If the receiver device is replaced by an ideal load then the return loss does not represent in service conditions. If the interconnect is very lossy then the effects of the load on the far end (i.e., where the receiver device is located), are not significant and an ideal load may be used. However, if the interconnect is not very lossy as in the zero length case, then the measured return loss may be dominated by the properties of the receiver device and not the properties of the interconnect.

For short links this return loss performance may be the limiting factor for the entire link due to severe unattenuated reflections that create large deterministic jitter.

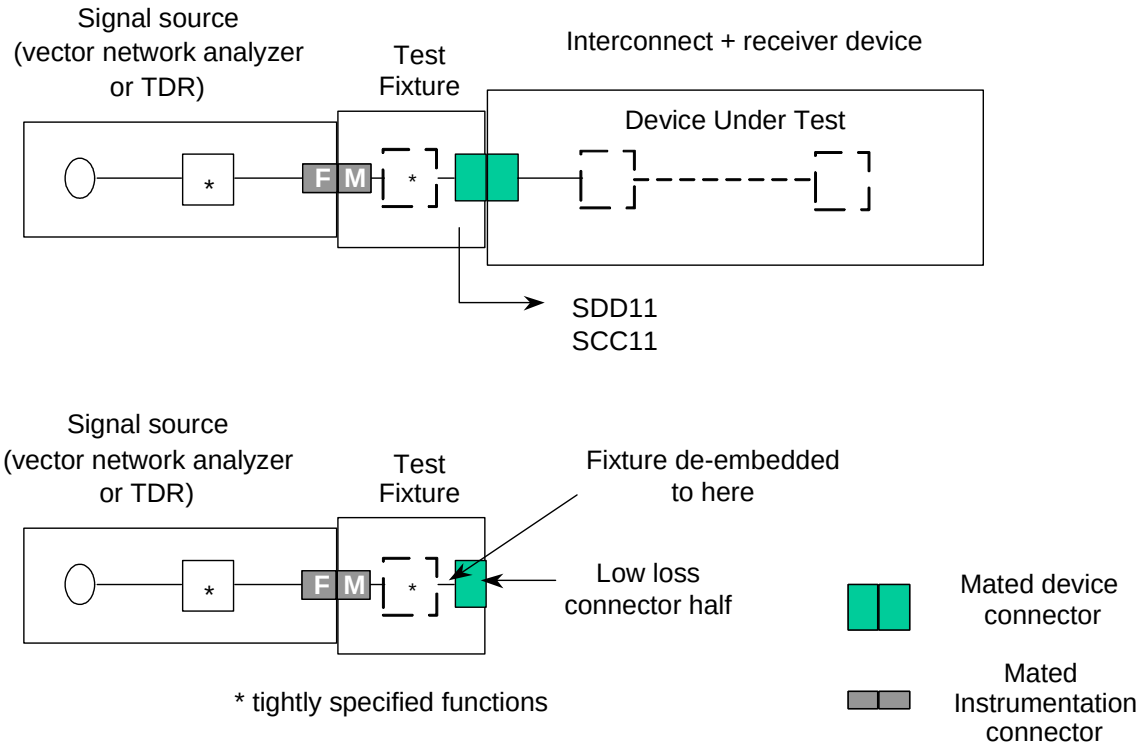

Figure A.16 — Measurement conditions for downstream return loss at the transmitter device connector

A.9.4.5 S22 at the receiver device connector (interconnect output)

Figure A.17 shows the conditions for making the return loss measurement out of the interconnect attached to the receiver device.

This measurement is unique in that it requires both the interconnect and the transmitter device to be in place and the combination is measured. This may be considered a reverse direction signal tolerance measurement. If the transmitter device is replaced by an ideal load then the return loss does not represent in service conditions. If the interconnect is very lossy then the effects of the load on the far end (i.e., where the transmitter device is located) are not significant and an ideal load may be used. However, if the interconnect is not very lossy as in the zero length case, then the measured return loss may be dominated by the properties of the transmitter device and not the properties of the interconnect.

For short links this return loss performance may be the limiting factor for the entire link due to severe unattenuated reflections that create large deterministic jitter.

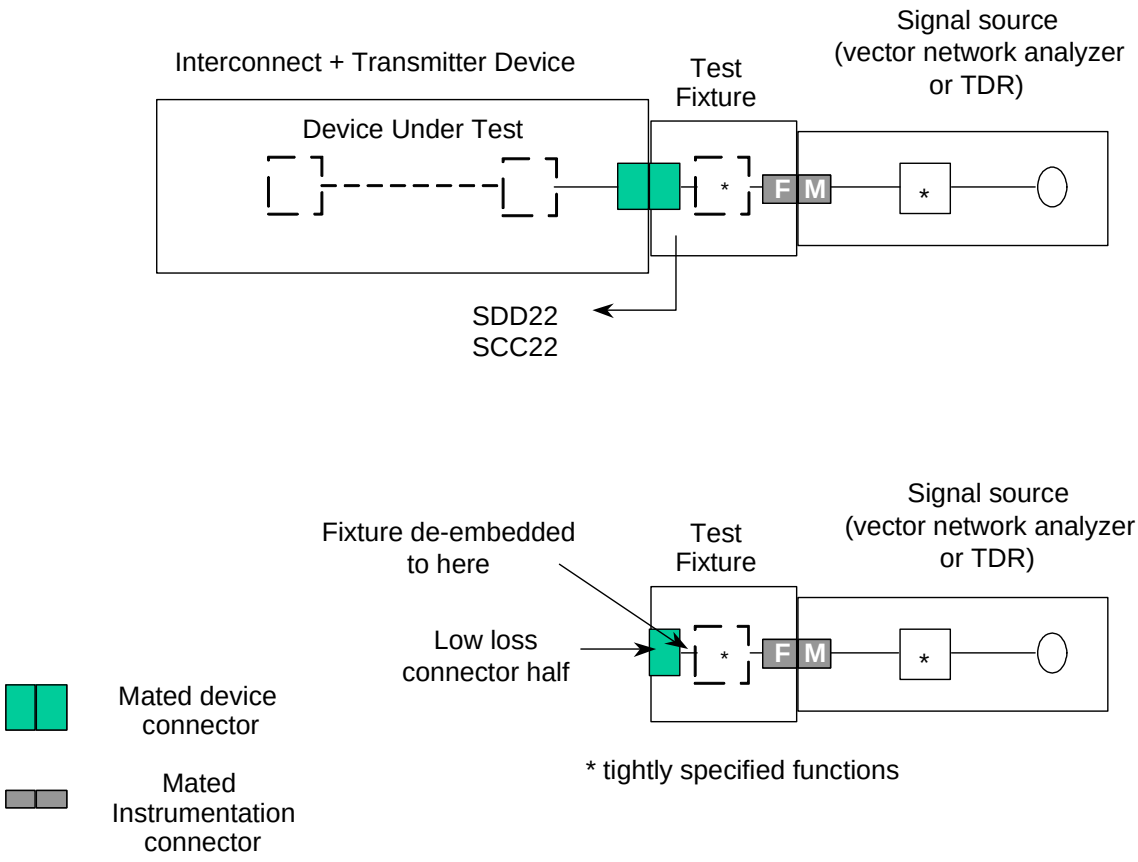

Figure A.17 — Measurement conditions for upstream return loss at the receiver device connector

A.9.5 Summary for S-parameter measurements

S-parameters are the preferred method of capturing the linear properties of link elements. Complex, but tractable, methods are required to use single ended instruments for differential and common mode applications. Careful attention to test configuration details is required.

A frequency domain spectrum output is required for all S-parameters and specifying pass fail limits to such a spectrum may over constrain the system because some peaks and properties are benign to the application.