SFF Committee documentation may be purchased in hard copy or electronic form.
SFF specifications are available at ftp://ftp.seagate.com/sff

SFF Committee

SFF-8429 Specification for

Signal Specification Architecture for HSS Links

Rev 1.1 June 27, 2005

Secretariat: SFF Committee

Abstract: This specification defines the architectural requirements for specifying the performance of the signals and signal tolerance in links where the components are intended to be interoperable and interchangeable. This subject is important to most enclosure to enclosure and intra-building wiring applications. Specifically the following transport technologies may be affected: Fibre Channel, Ethernet, Infiniband, SAS, and others. Other applications such as power and telecom for these general-purpose considerations are also possible.

This specification provides a common specification for systems manufacturers, system integrators, and suppliers of interconnect components. This is an internal working specification of the SFF Committee, an industry ad hoc group.

This specification is made available for public review, and written comments are solicited from readers. Comments received by the members will be considered for inclusion in future revisions of this specification.

The description of a methodology in this specification does not assure that the methodology is actually used in published or draft standards. Methodologies must comply with this specification to achieve interoperability and interchangeablity between suppliers of link components.

Support: This specification is supported by the identified member companies of the SFF Committee.

POINTS OF CONTACT:

Bill Ham
Hewlett Packard
165 Dascomb Road
Andover, MA 01810
978 828-9102
Bill.Ham@HP.com

I. Dal Allan
Chairman SFF Committee
14426 Black Walnut Court
Saratoga CA 95070
408-867-6630
endlcom@acm.org

EXPRESSION OF SUPPORT BY MANUFACTURERS

The following member companies of the SFF Committee voted in favor of this industry specification.

EMC
FCI
Foxconn
Hewlett Packard
Sun Microsystems

The following member companies of the SFF Committee voted to abstain on this industry specification.

Adaptec
Amphenol
Dell
Emulex
Hitachi GST
Infineon
Intel
LSI Logic
Molex
Seagate
Toshiba America
Tyco AMP
Unisys
Vitesse Semi

If you are not a member of the SFF Committee, but you are interested in participating, the following principles have been reprinted here for your information.

PRINCIPLES OF THE SFF COMMITTEE

The SFF Committee is an ad hoc group formed to address storage industry needs in a prompt manner. When formed in 1990, the original goals were limited to defining de facto mechanical envelopes within which disk drives can be developed to fit compact computer and other small products.

Adopting a common industry size simplifies the integration of small drives (2 1/2" or less) into such systems. Board-board connectors carrying power and signals, and their position relative to the envelope are critical parameters in a product that has no cables to provide packaging leeway for the integrator.

In November 1992, the SFF Committee objectives were broadened to encompass other areas which needed similar attention, such as pinouts for interface applications, and form factor issues on larger disk drives. SFF is a forum for resolving industry issues that are either not addressed by the standards process or need an immediate solution.

Specifications created by the SFF Committee are expected to be submitted to bodies such as EIA (Electronic Industries Association) or an ASC (Accredited Standards Committee). They may be accepted for separate standards, or incorporated into other standards activities.

The principles of operation for the SFF Committee are not unlike those of an accredited standards committee. There are 3 levels of participation:

- Attending the meetings is open to all, but taking part in discussions is limited to member companies, or those invited by member companies
- The minutes and copies of material which are discussed during meetings are distributed only to those who sign up to receive documentation.
- The individuals who represent member companies of the SFF Committee receive documentation and vote on issues that arise. Votes are not taken during meetings, only guidance on directions. All voting is by letter ballot, which ensures all members an equal opportunity to be heard.

Material presented at SFF Committee meetings becomes public domain. There are no restrictions on the open mailing of material presented at committee meetings. In order to reduce disagreements and misunderstandings, copies must be provided for all agenda items that are discussed. Copies of the material presented, or revisions if completed in time, are included in the documentation mailings.

The sites for SFF Committee meetings rotate based on which member companies volunteer to host the meetings. Meetings have typically been held during the ASC T10 weeks.

The funds received from the annual membership fees are placed in escrow, and are used to reimburse ENDL for the services to manage the SFF Committee.

If you are not receiving the documentation of SFF Committee activities or are interested in becoming a member, the following signup information is reprinted here for your information.

Membership includes voting privileges on SFF Specs under development.

CD_Access Electronic documentation contains:
- Minutes for the year-to-date plus all of last year
- Email traffic for the year-to-date plus all of last year
- The current revision of all the SFF Specifications, as well as any previous revisions distributed during the current year.

Meeting documentation contains:
- Minutes for the current meeting cycle.
- Copies of Specifications revised during the current meeting cycle.

Each electronic mailing obsoletes the previous mailing of that year e.g. July replaces May. To build a complete set of archives of all SFF documentation, retain the last SFF CD_Access mailing of each year.

Name: _____ Title: _____

Company: _____

Address: _____

_____

Phone: _____ Fax: _____

Email: _____

Please register me with the SFF Committee for one year.

___ Voting Membership w/Electronic documentation $ 2,160

___ Voting Membership w/Meeting documentation $ 1,800

___ Non-voting Observer w/Electronic documentation $ 660 U.S.
$ 760 Overseas

___ Non-voting Observer w/Meeting documentation $ 300 U.S.
$ 400 Overseas

Check Payable to SFF Committee for $_____ is Enclosed

Please invoice me for $_____ on PO #: _____

MC/Visa/AmX_____ Expires_____

SFF Committee
14426 Black Walnut Ct
Saratoga CA 95070

408-867-6630
408-867-2115Fx
endlcom@acm.org

Foreword

When 2 1/2" diameter disk drives were introduced, there was no commonality on external dimensions e.g. physical size, mounting locations, connector type,

connector location, between vendors.

The first use of these disk drives was in specific applications such as laptop portable computers in which space was at a premium and time to market with the latest machine was an important factor. System integrators worked individually with vendors to develop the packaging. The result was wide diversity, and with space being such a major consideration in packaging, it was not possible to replace one vendor's drive with a competitive product.

The desire to reduce disk drive sizes to even smaller dimensions such as 1.8" and 1.3" made it likely that devices would become even more constrained in dimensions because of a possibility that such small devices could be inserted into a socket, not unlike the method of retaining semiconductor devices.

The problems faced by integrators, device suppliers, and component suppliers led to the formation of an industry ad hoc group to address the marketing and engineering considerations of the emerging new technology in disk drives. After two informal gatherings on the subject in the summer of 1990, the SFF Committee held its first meeting in August.

During the development of the form factor definitions, other activities were suggested because participants in the SFF Committee faced problems other than the physical form factors of disk drives. In November 1992, the members approved an expansion in charter to address any issues of general interest and concern to the storage industry. The SFF Committee became a forum for resolving industry issues that are either not addressed by the standards process or need an immediate solution.

At the same time, the principle was adopted of restricting the scope of an SFF project to a narrow area, so that the majority of specifications would be small and the projects could be completed in a rapid timeframe. If proposals are made by a number of contributors, the participating members select the best concepts and uses them to develop specifications which address specific issues in emerging storage markets.

Those companies which have agreed to support a specification are identified in the first pages of each SFF Specification. Industry consensus is not an essential requirement to publish an SFF Specification because it is recognized that in an emerging product area, there is room for more than one approach. By making the documentation on competing proposals available, an integrator can examine the alternatives available and select the product that is felt to be most suitable.

Suggestions for improvement of this specification will be welcome. They should be sent to the SFF Committee, 14426 Black Walnut Ct, Saratoga, CA 95070.

The development work on this specification was done by the SFF Committee, an industry group. The membership of the committee since its formation in 1990 has included a mix of companies which are leaders across the industry.

SFF Committee --

Signal Specification Architecture for HSS Links

1. Scope

Connections between end points in high speed serial data and communication systems, sometimes called links, are often made from several components like connectors, bulk cable, optical fiber, cable assemblies, transmitters, receivers. This document explores the issues and architectures that apply when it is desirable to form the connection from components that are specified independently of other components in the same connection. This independence allows construction of a physical connection by assembling the components without needing to consider the actual properties of the other components (as long as the other components also comply with their independent specifications).

Independent specification of connection components allows suppliers of the components to design products that may be used in a variety of applications and may be ‘field assembled’ to other compliant components to form the complete connection. These properties not only facilitate higher volumes and competition for the components themselves but also allow significant reduction in the qualification and verification time required for the final connections.

The properties of components that produce the desired interoperability and interchangeability are several including mechanical, overall size, electrical properties, optical properties, environmental stability, safety, power consumption, reliability, color, cost, thermal properties, installability, serviceability, robustness and others. Many of the ideas brought forward in this document apply to several of these properties. However, the electrical properties and the optical properties are the main focus for this document.

This simple sounding goal of independent specification is far more complicated than it may first appear. Interactions between the components and the direction of signal flow provide first order impact to the independence desired.

One of the first things to address in realizing the stated goals is identifying the points in the connection that define the boundaries of the components. These points should map directly to the points where the link components may be assembled and disassembled. For purposes of this document attention is confined to points where separable connectors (electrical or optical) exist within the connection. In this context the requirements for specification of signal performance at these connectors are considered.

Much of the material in this document is derived from a presentation developed by the Fibre Channel physical group and is available on the T11 web site (www.T11.org) as document 04-024v6 and from the recently completed Methodologies for Jitter and Signal Quality Specification (MJSQ) that is also available on the T11 web site. Bill Ham, HP, is the editor for both of these documents.

The SFF Committee was formed in August, 1990 to broaden the applications for storage devices, and is an ad hoc industry group of companies representing system integrators, peripheral suppliers, and component suppliers.

2. SFF Specifications

There are several projects active within the SFF Committee. At the date of printing specification numbers had been assigned to the following projects. The status of Specifications is dependent on committee activities.

F = Forwarded	The specification has been approved by the members for forwarding to a formal standards body.
P = Published	The specification has been balloted by members and is available as a published SFF Specification.
A = Approved	The specification has been approved by ballot of the members and is in preparation as an SFF Specification.
C = Canceled	The project was canceled, and no Specification was Published.
D = Development	The specification is under development at SFF.
E = Expired	The specification has been published as an SFF Specification, and the members voted against re-publishing it when it came up for annual review.
e = electronic	Used as a suffix to indicate an SFF Specification which has Expired but is still available in electronic form from SFF e.g. a specification has been incorporated into a draft or published standard which is only available in hard copy.
i = Information	The specification has no SFF project activity in progress, but it defines features in developing industry standards. The specification was provided by a company, editor of an accredited standard in development, or an individual. It is provided for broad review (comments to the author are encouraged).
s = submitted	The document is a proposal to the members for consideration to become an SFF Specification.

Spec #	Rev	List of Specifications as of June 28, 2005
SFF-8000		SFF Committee Information
INF-8001i	E	44-pin ATA (AT Attachment) Pinouts for SFF Drives
INF-8002i	E	68-pin ATA (AT Attachment) for SFF Drives
SFF-8003	E	SCSI Pinouts for SFF Drives
SFF-8004	E	Small Form Factor 2.5" Drives
SFF-8005	E	Small Form Factor 1.8" Drives
SFF-8006	E	Small Form Factor 1.3" Drives
SFF-8007	E	2mm Connector Alternatives
SFF-8008	E	68-pin Embedded Interface for SFF Drives
SFF-8009	4.1	Unitized Connector for Cabled Drives
SFF-8010	E	Small Form Factor 15mm 1.8" Drives
INF-8011i	E	ATA Timing Extensions for Local Bus
SFF-8012	3.0	4-Pin Power Connector Dimensions
SFF-8013	E	ATA Download Microcode Command
SFF-8014	C	Unitized Connector for Rack Mounted Drives
SFF-8015	E	SCA Connector for Rack Mounted SFF SCSI Drives
SFF-8016	C	Small Form Factor 10mm 2.5" Drives
SFF-8017	E	SCSI Wiring Rules for Mixed Cable Plants
SFF-8018	E	ATA Low Power Modes
SFF-8019	E	Identify Drive Data for ATA Disks up to 8 GB
INF-8020i	E	ATA Packet Interface for CD-ROMs
SFF-8025	0.7	SFF Committee Specification Categories
INF-8028i	E	- Errata to SFF-8020 Rev 2.5
SFF-8029	E	- Errata to SFF-8020 Rev 1.2

SFF-8030	2.0	SFF Committee Charter
SFF-8031		Named Representatives of SFF Committee Members
SFF-8032	1.6	SFF Committee Principles of Operation
INF-8033i	E	Improved ATA Timing Extensions to 16.6 MBs
INF-8034i	E	High Speed Local Bus ATA Line Termination Issues
INF-8035i	E	Self-Monitoring, Analysis & Reporting Technology
INF-8036i	E	ATA Signal Integrity Issues
INF-8037i	E	Intel Small PCI SIG
INF-8038i	E	Intel Bus Master IDE ATA Specification
INF-8039i	E	Phoenix EDD (Enhanced Disk Drive) Specification
SFF-8040	1.2	25-pin Asynchronous SCSI Pinout
SFF-8041	C	SCA-2 Connector Backend Configurations
SFF-8042	C	VHDCI Connector Backend Configurations
SFF-8043	E	40-pin MicroSCSI Pinout
SFF-8045	4.7	40-pin SCA-2 Connector w/Parallel Selection
SFF-8046	E	80-pin SCA-2 Connector for SCSI Disk Drives
SFF-8047	C	40-pin SCA-2 Connector w/Serial Selection
SFF-8048	C	80-pin SCA-2 Connector w/Parallel ESI
SFF-8049	E	80-conductor ATA Cable Assembly
INF-8050i	1.0	Bootable CD-ROM
INF-8051i	E	Small Form Factor 3" Drives
INF-8052i	E	ATA Interface for 3" Removable Devices
SFF-8053	5.5	GBIC (Gigabit Interface Converter)
SFF-8054	0.2	Automation Drive Interface Connector
INF-8055i	E	SMART Application Guide for ATA Interface
SFF-8056	C	50-pin 2mm Connector
SFF-8057	E	Unitized ATA 2-plus Connector
SFF-8058	E	Unitized ATA 3-in-1 Connector
SFF-8059	E	40-pin ATA Connector
SFF-8060	1.1	SFF Committee Patent Policy
SFF-8061	E	Emailing drawings over the SFF Reflector
SFF-8062		Rolling Calendar of SSWGs and Plenaries
SFF-8064		Unshielded HD Cable/Board Connector System
SFF-8065	C	40-pin SCA-2 Connector w/High Voltage
SFF-8066	C	80-pin SCA-2 Connector w/High Voltage
SFF-8067	3.4	40-pin SCA-2 Connector w/Bidirectional ESI
INF-8068i	E	Guidelines to Import Drawings into SFF Specs
SFF-8069	E	Fax-Access Instructions
INF-8070i	1.3	ATAPI for Rewritable Removable Media
SFF-8072	1.2	80-pin SCA-2 for Fibre Channel Tape Applications
SFF-8073	C	20-pin SCA-2 for GBIC Applications
INF-8074i	1.0	SFP (Small Formfactor Pluggable) Transceiver
SFF-8075	1.0	PCI Card Version of SFP Cage
SFF-8076	-	SFP Additional IDs
INF-8077i	3.1	XFP (10 Gbs Small Form Factor Pluggable Module)
SFF-8078	C	XFP-E
SFF-8079	1.7	SFP Rate and Application Selection
SFF-8080	E	ATAPI for CD-Recordable Media
SFF-8082	5.1	Labeling of Ports and Cable Assemblies
SFF-8084	0.2	0.8mm SFP Card Edge Connector Dimensioning
SFF-8085	0.9	100 Mbs Small Formfactor Transceivers
SFF-8086	1.2	Compact Multilane Series: Common Elements
SFF-8087	1.3	Compact Multilane Series: Unshielded
SFF-8088	1.2	Compact Multilane Series: Shielded
SFF-8089	1.3	SFP Rate and Application Codes
INF-8090i	6.09	ATAPI for Multimedia Devices (Mt Fuji5)
SFF-8101	C	3 Gbs and 4 Gbs Signal Characteristics
SFF-8110	C	5V Parallel 1.8" drive form factor

SFF-8111	1.3	1.8" drive form factor (60x70mm)
SFF-8122		1.8" (60x70mm) w/SCA-2 Connector
SFF-8120	2.6	1.8" drive form factor (78x54mm)
SFF-8123	C	1.8" (60x70mm) w/Serial Attachment Connector
SFF-8124	0.2	Memory Form Factor Disk Drive Connections
SFF-8131		40mmx50mm Form Factor
SFF-8200e	1.1	2 1/2" drive form factors (all of 82xx family)
SFF-8201	2.4	2 1/2" drive form factor dimensions
SFF-8212e	1.2	2 1/2" drive w/SFF-8001 44-pin ATA Connector
SFF-8221	C	Pre-Aligned 2.5" Drive >10mm Form Factor
SFF-8222	2.1	2.5" Drive w/SCA-2 Connector
SFF-8223	2.4	2.5" Drive w/Serial Attachment Connector
SFF-8225	C	2.5" Single Voltage Drive
SFF-8300	1.2	3 1/2" drive form factors (all of 83xx family)
SFF-8301	1.4	3 1/2" drive form factor dimensions
SFF-8302e	1.1	3 1/2" Cabled Connector locations
SFF-8323	1.4	3 1/2" drive w/Serial Attachment Connector
SFF-8332e	E	3 1/2" drive w/80-pin SFF-8015 SCA Connector
SFF-8337e	E	3 1/2" drive w/SCA-2 Connector
SFF-8342e	1.3	3 1/2" drive w/Serial Unitized Connector
INF-8350i	E	3 1/2" Packaged Drives
SFF-8400	C	VHDCI (Very High Density Cable Interconnect)
SFF-8401		Optical Transceiver for Short-Reach Appcns
SFF-8410	16.1	High Speed Serial Testing for Copper Links
INF-8411	1.0	High Speed Serial Testing for Backplanes
SFF-8412	12.2	HSOI (High Speed Optical Interconnect) Testing
SFF-8415	4.1	HPEI (High Performance Electrical Interconnect)
SFF-8416	15.0	HPEI Bulk Cable Measurement/Performance Reqmnts
SFF-8420	11.1	HSSDC-1 Shielded Connections
SFF-8421	2.4	HSSDC-2 Shielded Connections
SFF-8422	C	FCI Shielded Connections
SFF-8423	C	Molex Shielded Connections
SFF-8424	0.5	Dual Row HSSDC-2 Shielded Connections
SFF-8425	1.4	Single Voltage 12V Drives
SFF-8426		HSSDC Double Width
SFF-8429	1.1	Signal Specification Architecture for HSS Links
SFF-8430	4.1	MT-RJ Duplex Optical Connections
SFF-8431		SFP+
SFF-8441	14.1	VHDCI Shielded Configurations
SFF-8448	0.4	SAS Sideband Utilization
SFF-8451	10.1	SCA-2 Unshielded Connections
SFF-8452	3.1	Glitch Free Mating Connections for Multidrop Aps
SFF-8453		Shielded High Speed Serial connectors
SFF-8454		SCA-2 Enhanced HSS
SFF-8460	1.2	HSS Backplane Design Guidelines
SFF-8464	C	Improved MM HSS Optical Link Performance
SFF-8470	2.9	Multilane Copper Connector
SFF-8471	C	ZFP Multilane Copper Connector
SFF-8472	9.5	Diagnostic Monitoring Interface for Optical Xcvrs
INF-8475i	2.2	XPAK Small Formfactor Pluggable Receiver
SFF-8480	2.1	HSS (High Speed Serial) DB9 Connections
SFF-8482	1.8	Unshielded Dual Port Serial Attachment Connector
SFF-8483	C	External Serial Attachment Connector
SFF-8484	1.6	Multilane Unshielded Serial Attachment Connector
SFF-8485	0.5	Serial GPIO (General Purpose Input/Output) Bus

SFF-8500e	1.1	5 1/4" drive form factors (all of 85xx family)
SFF-8501e	1.1	5 1/4" drive form factor dimensions
SFF-8508e	1.1	5 1/4" ATAPI CD-ROM w/audio connectors
SFF-8523	1.4	5 1/4" drive w/Serial Attachment Connector
SFF-8551	3.2	5 1/4" CD Drives form factor
SFF-8552	1.1	5 1/4" 9.5mm/12.7mm Optical Drive Form Factor
SFF-8572	C	5 1/4" Tape form factor
SFF-8610	C	SDX (Storage Device Architecture)
SFF-8617		SAS Transition cables

3. Sources

Copies of ANSI standards or proposed ANSI standards may be purchased from Global Engineering.

15 Inverness Way East 800-854-7179 or 303-792-2181
Englewood 303-792-2192Fx
CO 80112-5704

Copies of SFF Specifications are available by joining the SFF Committee as an Observer or Member or by download at ftp://ftp.seagate.com/sff

14426 Black Walnut Ct 408-867-6630x303
Saratoga 408-867-2115Fx
CA 95070

TABLE OF CONTENTS

1. Introduction

Connections between end points in high speed serial data and communication systems, sometimes called links, are often made from several components like connectors, bulk cable, optical fiber, cable assemblies, transmitters, receivers. This document explores the issues and architectures that apply when it is desirable to form the connection from components that are specified independently of other components in the same connection. This independence allows construction of a physical connection by assembling the components without needing to consider the actual properties of the other components (as long as the other components also comply with their independent specifications).

Independent specification of connection components allows suppliers of the components to design products that may be used in a variety of applications and may be ‘field assembled’ to other compliant components to form the complete connection. These properties not only facilitate higher volumes and competition for the components themselves but also allow significant reduction in the qualification and verification time required for the final connections.

The properties of components that produce the desired interoperability and interchangeability are several including mechanical, overall size, electrical properties, optical properties, environmental stability, safety, power consumption, reliability, color, cost, thermal properties, installability, serviceability, robustness and others. Many of the ideas brought forward in this document apply to several of these properties. However, the electrical properties and the optical properties are the main focus for this document.

This simple sounding goal of independent specification is far more complicated than it may first appear. Interactions between the components and the direction of signal flow provide first order impact to the independence desired.

One of the first things to address in realizing the stated goals is identifying the points in the connection that define the boundaries of the components. These points should map directly to the points where components may be assembled and disassembled. For purposes of this document attention is confined to points where separable connectors (electrical or optical) exist within the connection. In this context the requirements for specification of signal performance at these connectors are considered.

Much of the material in this document is derived from a presentation developed by the Fibre Channel physical group and is available on the T11 web site (www.T11.org) as document 04-024v6 and from the recently completed Methodologies for Jitter and Signal Quality Specification (MJSQ) that is available from http://www.techstreet.com/cgi-bin/new_results search for ‘MJSQ’. Bill Ham, HP, is the editor for both of these documents. Both FC-PI-2 and SAS 1.1 have adopted key portions of the methodology defined in this document.

2. A simple connection

2.1 Overview

In a simple example, that becomes less simple soon enough, the connection, or physical link, consists of three component parts: the transmitter device, the interconnect, and the receiver device each joined together by a separable connector. Each component part needs its own performance specification that is independent of the properties of any other specific component in the connection. Signals travel in

opposite directions down the same nominal path if a duplex connection is used. For some applications like Gigabit Ethernet the same wire is used simultaneously for both directions (the received signal is detected by subtracting the known launched signal from the total). This document is optimized for the more common case where each direction of signal travel has its own dedicated wires. The general approach is useful in both cases.

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

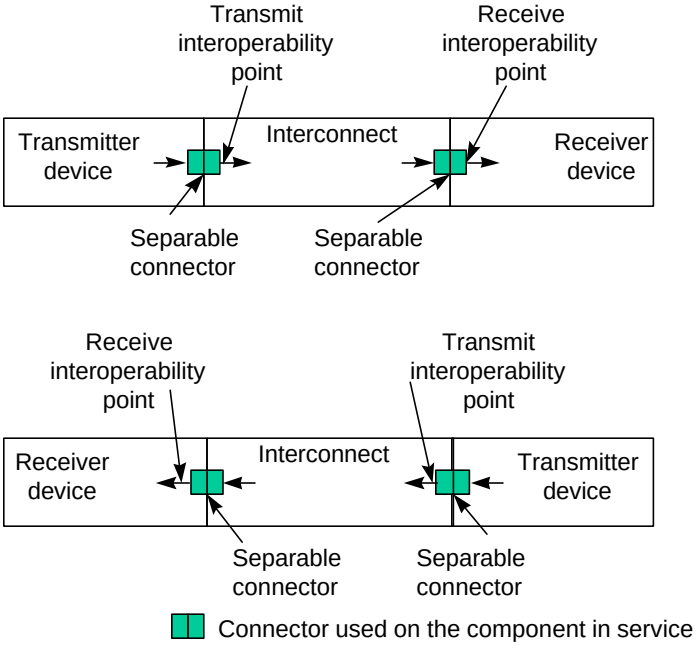

Figure 1 -A simple duplex link physical connection

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 could 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. Attempting to access the unmated connector with probes, for example, is also not acceptable because the connector is not the same when unmated as when mated and the probe contact points will not be at the same location as the connector contact points. When using probes the contacts are not deflected and the shields are not connected in unmated electrical connectors. For optical connectors the only practical access method to an unmated connector is by adding the mating half and that makes it a mated connector.

In this document the signal outputs are always measured downstream of the mated connector (as shown in Figure 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 later shows that the mated connector issue is not a severe as it may appear.

Each part intrinsically has some properties that may be expressed via impedance and/or S parameter methodologies

The TxRx connection has an assumed ‘reference impedance’ e.g. 100 ohms.

2.2 Under the hood for the transmitter device and the receiver device

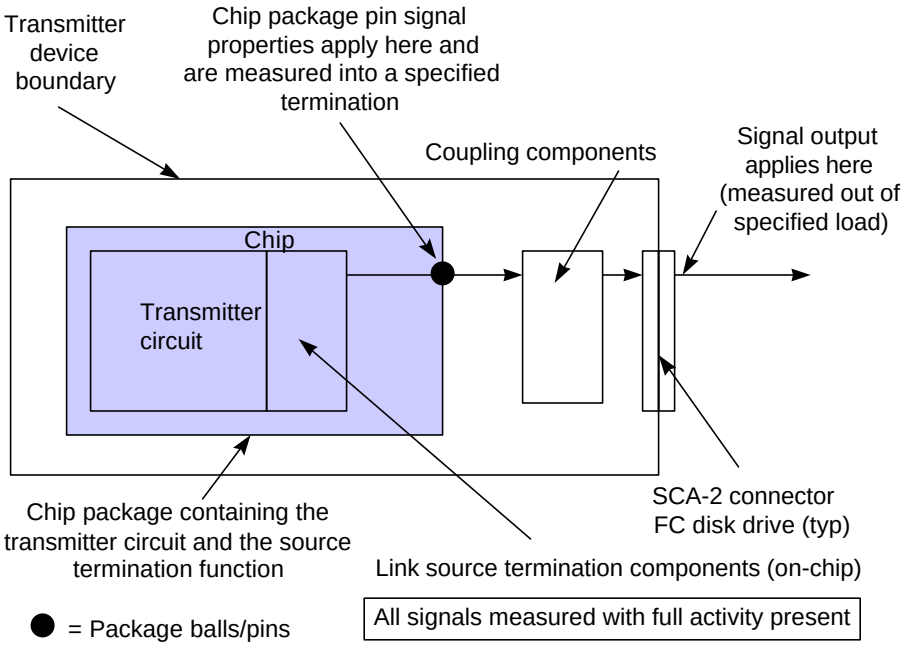

Figure 2 - Transmitter device details for an HDD type interoperability point

Figure 2 shows the details of a transmitter device. Notice that there are at least three internal parts of this transmitter device that could be called a ‘transmitter’:

1. the transmitter circuit in the chip
2. the chip itself
3. 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 connector (half a mated pair), optional coupling circuits, the source termination, the transmitter circuit, PCB traces and vias, the chip package, and possibly ESD devices. It is assumed that the source termination is contained within the chip package.

Interoperability specifications are defined at the connector of the transmitter device for storage standards (FC, SAS, SSA etc.)

Interoperability points might be defined at the chip package pins in some network standards (e. g., Ethernet XAUI).

The practice of defining interoperability points at the chip pins optimizes the specification from a chip point of view but may leave the rest of the interoperability problem in the link unspecified. This document assumes that the more important interoperability points are at the separable connectors where the user interfaces exist rather than at the chip interfaces that are usually permanently embedded inside some device. Other problems with specifying signals at chip pins include practical measurement difficulties and the impact of getting the signal to and through the transmitter device connector.

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

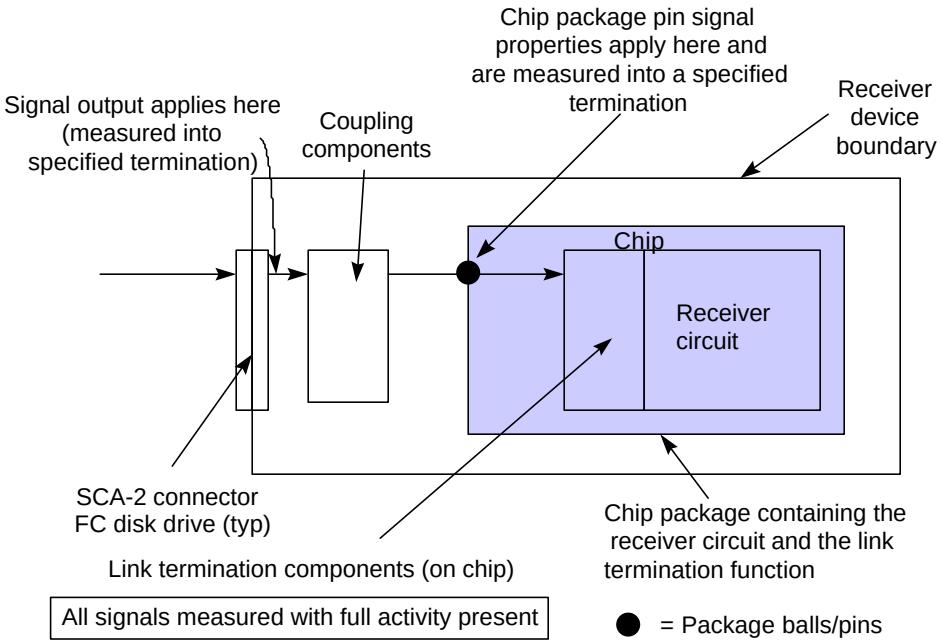

Figure 3 - Receiver device details for an HDD type interoperability point

Notice that there are at least three internal parts of this receiver device that could be called a ‘receiver’:

1. the receiver circuit in the chip
2. the chip itself
3. 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 connector (half a mated pair), optional coupling circuits, the link termination, the receiver circuit, PCB traces and vias, the chip package, and possibly ESD devices. It is assumed that the link termination is contained within the chip package.

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 specifications but it has proven impossible to purge the term. 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 following description is used to uniquely define these terms with the understanding that this document discourages usage of either term.

Receiver sensitivity:

- The receiver in the receiver sensitivity refers to signal properties at the chip package pin for the chip package that contains the receiver circuit.
- 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 (maximum jitter present) and all activity is **quiesced** except for the receiver itself.
- Receiver sensitivity is not defined where there are no chip pin specifications

Receiver device sensitivity:

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

Special test conditions are required to measure these sensitivities as described later. The terminology used is ‘signal tolerance’ instead of ‘receiver sensitivity’.

3. Basic signal performance properties

3.1 Introduction

At any given interoperability point a set of signal or link performance requirements are used to define the interoperability requirements. If these requirements are all met then, within the limitations of specific system implementation noise properties and damage that may have occurred during shipping and installation, the complete link should deliver the require BER without further testing and qualification.

For both optical and electrical links the following signal properties require specification at every interoperability point:

- Signal output (eye type measurement of the signal output at the interoperability point)
- Downstream signal tolerance (BER measurement using minimum quality signal as input at the interoperability point)
- Upstream return loss (S parameter 22 looking upstream at the interoperability point)
- Downstream return loss (S parameter 11 looking downstream at the interoperability point)

The details of these measurements depend strongly on the details of the electrical or optical connection.

In this document only the differential electrical type is described since that is the most complex. Single ended electrical and optical types essentially use the differential signal measurement in a single ended mode.

3.2 Differential electrical links

At a differential electrical interoperability point both the differential and common mode performance requirements are specified. The complete list of the eight required properties is in Table 1.

Table 1 - List of specifications required at differential electrical interoperability points

Name	Symbol	Description
Differential signal output	DSO	Jitter eye measurement (Note 1) with defined data pattern out of a laboratory grade electrical load - signal measured through the mated connector used in service Measures the performance of all upstream portions of the link
Differential signal tolerance	DST (Note 2) (Note 3)	BER measurement with a defined data pattern using a signal generated from laboratory grade calibration and launch conditions - signal calibrated through a laboratory grade mated connector Measures the performance of all downstream portions of the link
Differential upstream return loss (Note 4)	SDD22	Return loss measurement looking upstream using differential S parameter methodologies assuming an ideal differential and common mode reference impedance - effects of the service connector are included in this measurement Measures the performance of all upstream portions of the link
Differential downstream return loss (Note 4)	SDD11	Return loss measurement looking downstream using differential S parameter methodologies assuming an ideal differential and common mode reference impedance - effects of the service connector are included in this measurement Measures the performance of all downstream portions of the link
Common mode signal output	CSO	Jitter eye measurement (Note 1) with defined data pattern out of a laboratory grade electrical load - signal measured through the mated connector used in service Measures the performance of all upstream portions of the link
Common mode signal tolerance	CST (Note 2 (Note 3))	BER measurement with a defined data pattern using a signal generated from laboratory grade calibration and launch conditions - signal calibrated through a laboratory grade mated connector Measures the performance of all downstream portions of the link
Common mode upstream return loss (Note 4)	SCC22	Return loss measurement looking upstream using common mode S parameter methodologies assuming an ideal differential and common mode reference impedance - effects of the service connector are included in this measurement Measures the performance of all upstream portions of the link
Common mode downstream return loss (Note 4)	SCC11	Return loss measurement looking downstream using common mode S parameter methodologies assuming an ideal differential and common mode reference impedance - effects of the service connector are included in this measurement Measures the performance of all downstream portions of the link

Note 1 - see MJSQ for details on jitter eye measurements

Note 2 - Signal tolerance measurements for differential and common mode properties cannot be separated because the result of a tolerance measurement is the BER for the link. The signal used for signal tolerance measurements contains the worst case combination of differential, common mode, and data pattern properties.

The differential and common mode content in the signal used for signal tolerance is the specified quantity in this table.

Note 3 - Signal tolerance methods are described in more detail in MJSQ where four kinds of jitter content are described in these signals.

Note 4 - Use of S parameters assumes that the relevant portions of the links behave linearly - such behavior may not always exist in active devices or where magnetic coupling elements are used - the active devices are set to their nominal bias conditions during this measurement to minimize the impact of non-linear properties.

4. Measurement architecture requirements

4.1 Introduction

Signal specifications are only meaningful if the signals can 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. In other words the measurements must be accessible, verifiable, and transportable. As of this writing there are no accepted standards 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 document.

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 will 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 or possibly 75 ohms single ended (100 or 150 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. For more complex loads the industry is still working on how to make these available. The properties of more complex loads include specified propagation time, insertion loss properties, crosstalk properties, and jitter creation properties. More discussion on the complex loads is given in clause 6.

For signal tolerance measurements one must calibrate the signal 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 (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.

4.2 Relationship between signal compliance measurements at interoperability points and operation in systems

The signal and return loss measurements in this document apply under specified test conditions that simulate some parts of the conditions existing in service. This simulation includes, for example, duplex traffic on all Ports and under all applicable environmental conditions. Other features existing in service, such as non ideal return loss in parts of the link that are not present when measuring signals in the specified test conditions, are accounted for 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. Such a measurement may 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 specification limits.

5. The interoperability penalty and related issues

There is an interoperability penalty in the signal specifications caused by differences in the signal output and signal tolerance at the same interoperability point. These differences arise because components on both sides of the interoperability point each are allowed independent tolerances for properties.

For example, in a simple d.c. system with two compliant components, if the downstream component has a lower impedance than nominal and the upstream component has a higher impedance than nominal the signal at the interface is reduced by both the upstream and downstream components. If the signal output specification is measured with nominal impedance for the signal load the effect of the high impedance on the upstream component is included in the measurement. However, the signal will still not meet the minimum amplitude requirement when a compliant, but not nominal, low impedance downstream component is used in service.

The output signal specification is increased to accommodate the worst case possible from compliant components. This increase is effectively an interoperability penalty that would not exist if nominal components were used.

This document does not specify the methodology for determining the interoperability penalty values as this is a complex topic that has not yet been documented in detail. Extensive use of sophisticated simulations are required for determining the interoperability penalty in specific cases.

Another somewhat related issue is the applicability of the measurements to applications where small differences exist between the measurement conditions and the service conditions. For example, the actual link length between the transmitter device and the receiver device may be different from the cable assembly length assumed or used during measurement because of board traces between the cable assembly and the link termination in the system.

One approach used to address this condition for electrical interconnect assemblies in SFF-8415 (see www.sffcommittee.com) requires that eye measurements on interconnect assemblies be valid at data rates 10% above and 10% below the nominal data rate as well as at the nominal data rate. This ensures that there are no important suckouts near the data rate.

6. Compliance interconnect and standard receiver device methods

6.1 Compensation

Compensation is the attempt to mitigate the deleterious effects occurring during signal transmission by adding or subtracting features from the signal. Compensation may be executed in the transmitter device, in the interconnect, and in the receiver

device. Compensation may be applied to properties of the signal that depend on the specific data pattern and to properties that are predictable upon subsequent use such as line to line propagation time skew and DCD. The basic assumption for compensation is that the degradation intensity and type remains stable over periods of at least several bit times.

Compensation schemes that may adjust the parameters of the compensating mechanism are termed “adjustable”. Compensation schemes that use active circuitry are termed “active”. Compensation schemes that pass the signal through a transfer function are termed “filtered” or “equalized”, the latter being derived from the common practice of matching the transmission losses across part of the frequency spectrum. Adjustable schemes that change the parameters of the compensating mechanism in response to specific received signal measurements are termed “adaptive”. It is usually assumed that adaptive schemes will use some sort of automatic means to do the adapting.

The mechanisms of degrading signals fall broadly into two modes: (1) primary losses along the transmission path such as attenuation, reflections, and resonances and (2) secondary losses due to causes that are external to the transmission path such as crosstalk noise. In some cases the secondary losses may exceed the primary losses. Losses may be amplitude or timing precision or both.

Details concerning the important differences between the location of the compensation are explored in the following sub clauses.

6.2 Transmitter device compensation

Any compensation scheme that is implemented in the transmitter device either (a) makes assumptions concerning the nature and intensity of the degradation that occurs during the transmission or (b) has some sort of feedback from the receiver that allows adjustment of the parameters of the compensation scheme. Such a feedback mechanism relies on higher level protocol.

When transmitter device compensation is used, the results are visible in the signal launched from the transmitter device. However, it may be necessary to use special methods to determine the quality of the signal from the transmitter device since the compensation process may significantly alter the signal.

The general method for verifying transmitter device compensation is to pass the transmitter device signal through a transfer function that emulates the loss mechanism, in both magnitude and phase, for a standard interconnect (called a compliance interconnect) attached to the transmitter device before examining the signal. One measurement set up for such a methodology is shown in Figure 4 where models are used instead of hardware where possible.

The mathematical description of a compliance interconnect is not a specification for the interconnect used in the system itself (although the specifications for the compliance interconnect are derived from assumptions about the system interconnect); it is a specification of the load that is placed on the transmitter device for purposes of enabling measurement of transmitter device signals. See T11/04-604v0 for a more detailed discussion of compliance interconnect methodologies.

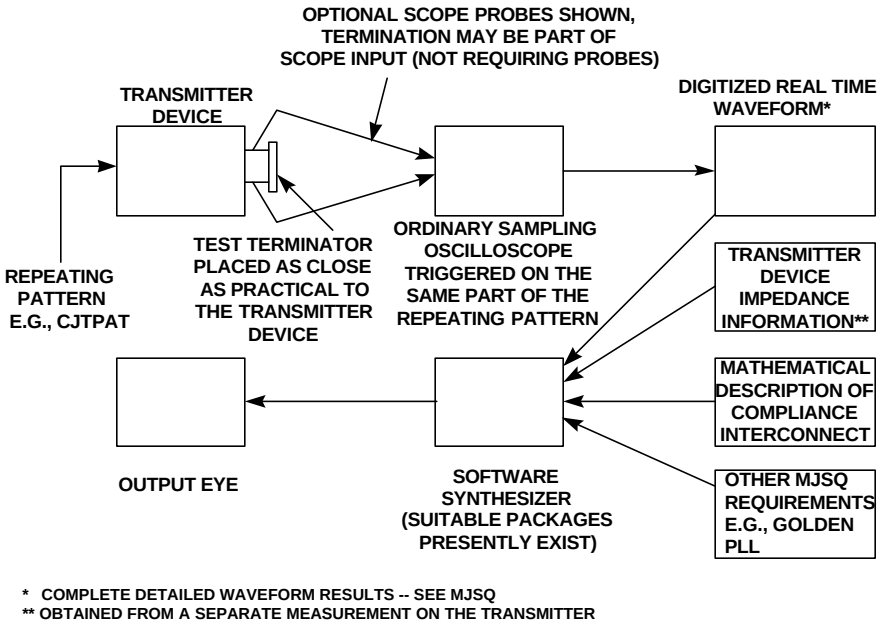

Figure 4 - Measurement set up for evaluating transmitters

It is also possible to evaluate the transmitter device by using a golden physical sample of compliance interconnect attached to the transmitter. In this case the mathematical description of the compliance interconnect is built into the golden hardware and direct observation of the output eye is possible without the software synthesizer. Even in this case the specification for a physical compliance interconnect is not the same as the specification for the link interconnect where a much more complete set of requirements are specified including items such as crosstalk and EMI performance.

6.3 Interconnect compensation

Interconnect is specified by its ability to transport a minimum quality launched signal to an adequate quality signal coming out the far end. If compensation is incorporated into the interconnect itself then the interconnect includes the compensation as part of its performance and no special treatment is required.

6.4 Receiver device compensation

Figure 5 shows the basic scheme used to specify input signals for devices whose internal circuity (receiver chip itself or other circuitry) incorporates compensation. If this compensation is included as part of the link budgeting process the effects of this compensation must be visible at the connector where the interoperability specifications apply. In an extreme example the signal may exhibit no jitter eye opening and/or no waveform eye opening at all at the connector unless the effects of the compensation are included in the signal measurement process. The receiver device itself is evaluated based on its ability to produce small error ratios with a minimum quality incoming signal. This evaluation demands the ability to specify and calibrate a signal at the interoperability point for use with signal tolerance measurements.

Since the receiver device processes the signal internally, the observed signal at

the connector needs to be processed to emulate the internal signal in the device at point “A” during the signal measurement. This is done by passing the signal at the connector through a standard compensation transfer function (possibly, but not necessarily, the same as that used internally in the device) before evaluating the signal. The standard compensation transfer function (or standard receiver device) is that assumed when creating the link budgets. Real receiver devices need not implement exactly the same function or method but the properties of the signal used for signal tolerance measurements are known and the device may be designed to work with this signal. The standard compensation function (standard receiver device) allows

- Visibility of the signal properties at the receiver device interoperability point
- Use of the same signal quality specification methodology as for non-compensated signals
- Link budgeting to be done without compromising the ability to do proprietary receiver designs.

For signal tolerance testing the signal observed downstream of the standard compensation function is then adjusted to have the properties of a minimum quality signal. This calibrated signal then applied to the device and the device bit error ratio is measured. Bit errors are detected at point “B” after passing through the internal receiver.

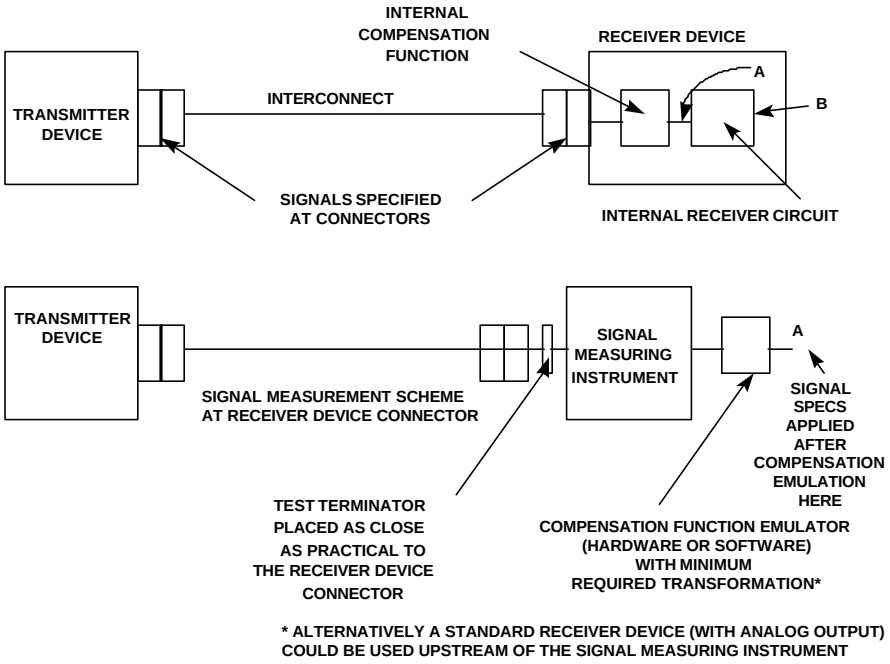

Figure 5 - General measurement configuration for the signal going into the receiver device that uses compensation

Since the measurement system shown in Figure 5 is linear, one could place the compensation function before the signal measurement instrument by using a hardware version of a golden receiver signal processor. In all cases specification of a standard compensation function is required.

6.5 Multiple compensation functions in the same link

Interactions between compensation implemented at different points in the link (transmitter device, interconnect, receiver device) is expected. Work is underway to quantify the impacts of this interaction. This interaction is another part of the interoperability penalty discussed in section The interoperability penalty and related issues.

7. De-embedding connectors in test fixtures

Connectors are necessarily 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 section.

Figure 6 shows two cases that apply.

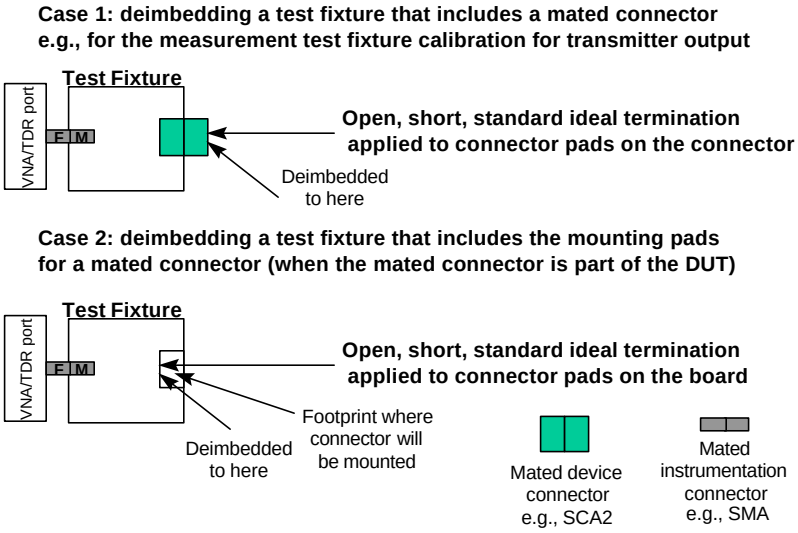

Figure 6 - De-embedding of connectors in test fixtures

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 on the signal measurement.

8. Measurement conditions for signal output (DSO) at the transmitter device

The measurement conditions required for a differential transmitter device signal output (DSO) are shown in Figure 7. Two required cases are described in this figure: one where the transmitter device is directly attached to the receiver device and the other where the transmitter device is attached to the receiver device through an interconnect assembly (cable assembly or PCB).

To simulate some of the properties of the interconnect assembly an instrumentation quality compliance interconnect is used as described in clause 6. This compliance interconnect is assumed embedded in the compliance interconnect test fixture as shown in more detail in Figure 8.

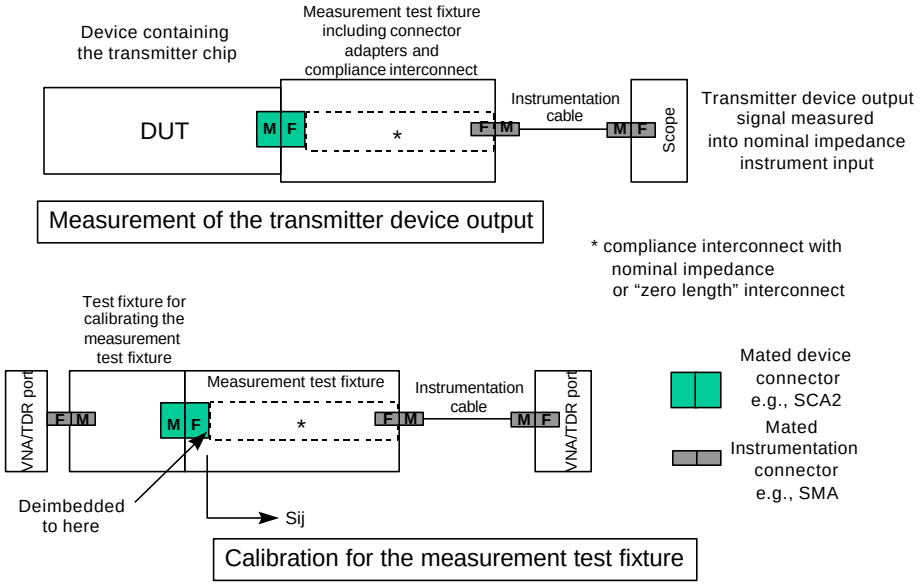

Figure 7 - Measurement conditions for transmitter device signal output specifications

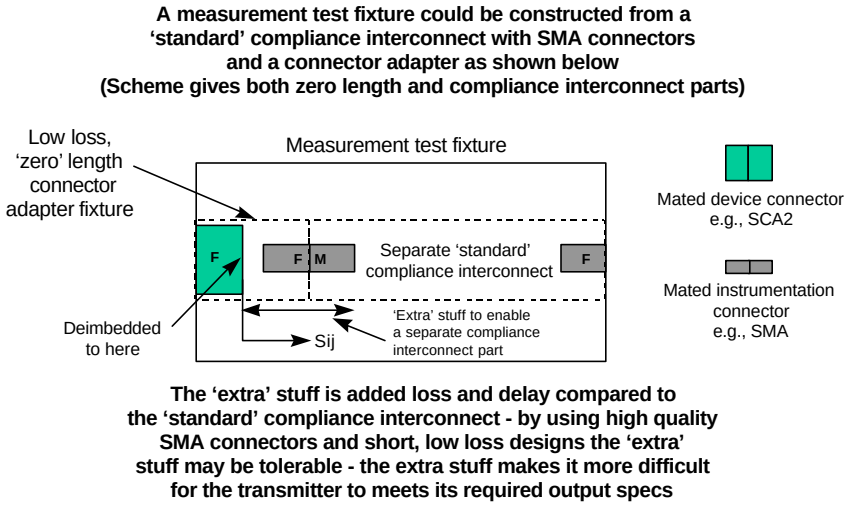

Figure 8 - Transmitter device output signal measurement test fixture details

9. Measurement conditions for signal tolerance (DST) at the transmitter device

The measurement conditions required for the signal tolerance (DST) at the differential transmitter device interoperability point are shown in Figure 9. Two required cases are described in this figure: one where the transmitter device is directly attached to the receiver device and the other where the transmitter device is attached to the receiver device through an interconnect assembly (cable assembly or PCB).

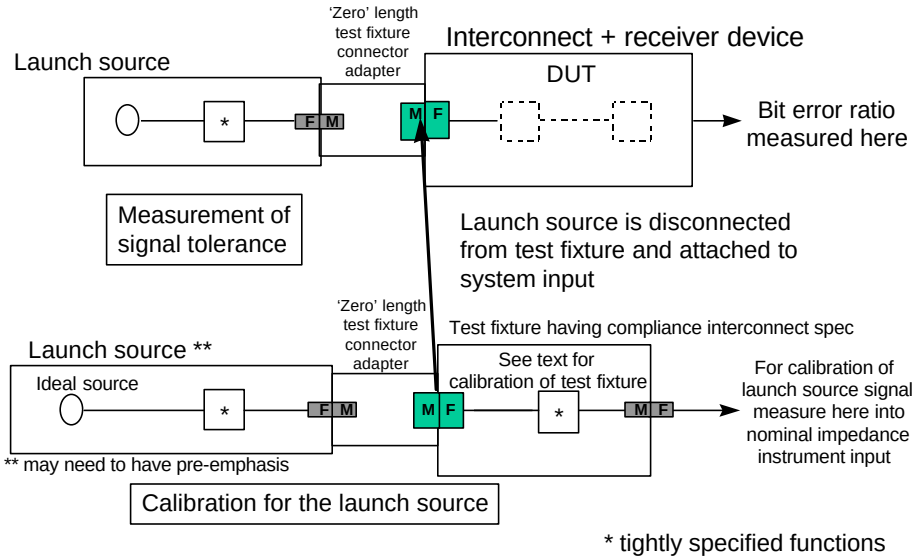

Figure 9 - Measurement conditions for system (interconnect + receiver device) signal tolerance

The test fixture for this measurement is shown in Figure 10.

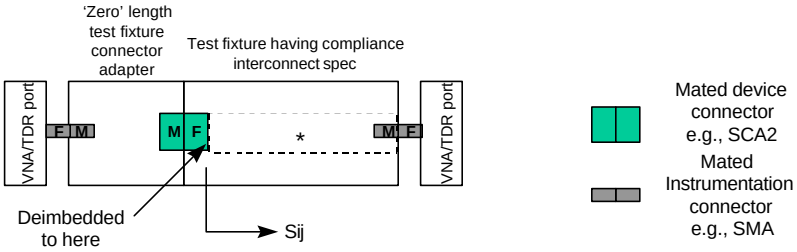

Note that this test fixture is identical to the 'measurement' test fixture used for the transmitter device output signal (even the connector genders and pins used in the SCA-2 connector are the same)

* tightly specified function

Figure 10 - Calibration of test fixture having compliance interconnect specification for transmitter device (interconnect + receiver device) signal tolerance

10. Measurement conditions for signal output (DSO) at the receiver device

The measurement conditions for the signal output at the receiver device are shown in Figure 11.

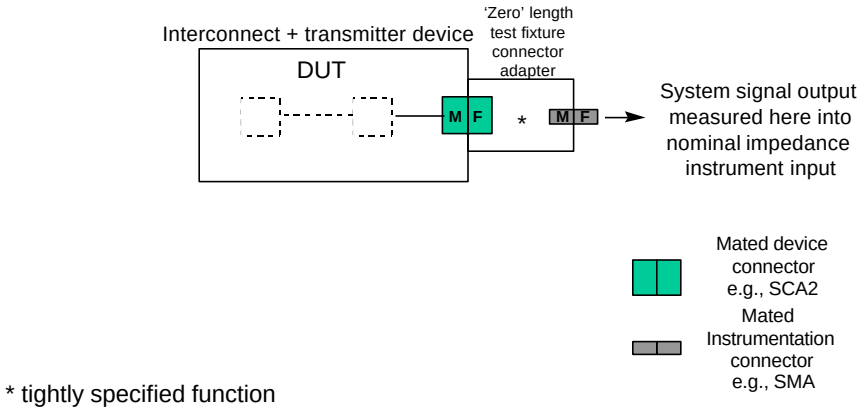

Figure 11 - Measurement conditions for signal output at the receiver device

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

11. Measurement conditions for signal tolerance (DST) at the receiver device

The measurement conditions required for the signal tolerance (DST) at the differential receiver device interoperability point are shown in Figure 12. Two required cases are described in this figure: one where the transmitter device is directly attached to the receiver device and the other where the transmitter device is attached to the receiver device through an interconnect assembly (cable assembly or PCB).

Figure 13 shows the calibration of receiver test interconnect test fixture for the receiver device signal tolerance measurement.

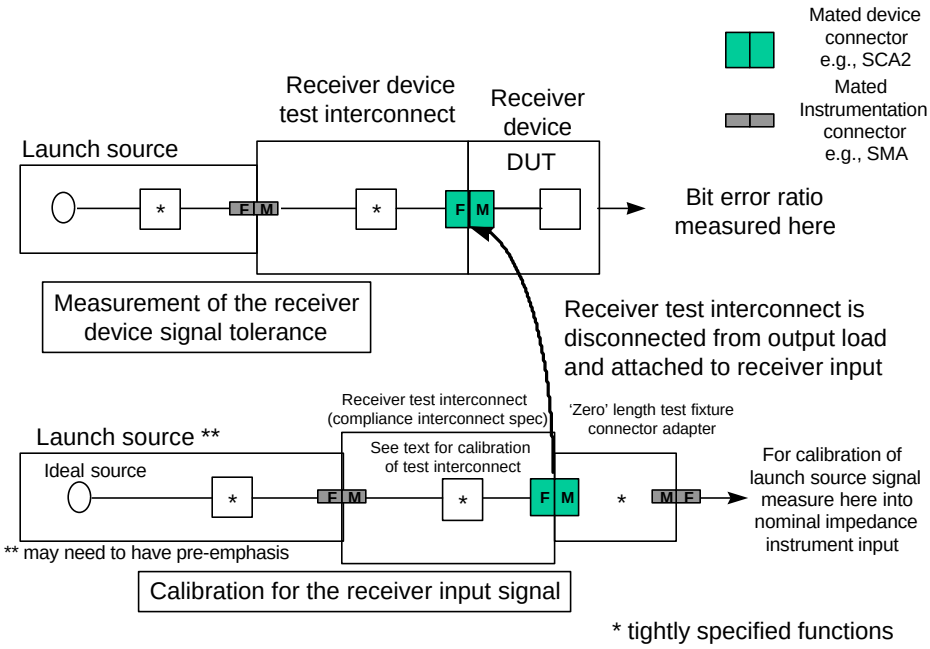

Figure 12 - Measurement conditions for receiver device signal tolerance

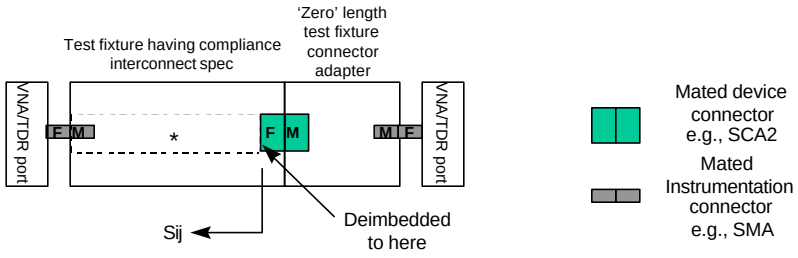

Note that this is <u>not</u> identical to the 'measurement' test fixture used for the transmitter output signal even though the connector genders are the same. The pins used in the SCA-2 connector are for the Rx (not the Tx) and the signals flow the other way. The S22 measurement here is the same as the S11 measurement for the transmitter output signal but on different pins

Also note that the S21 and S12 are used mainly to create the desired jitter in this application and are not as critical

* tightly specified function

Figure 13 - Calibration of receiver test interconnect test fixture for receiver device signal tolerance

12. S-parameter measurements

12.1 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:

- Naming conventions
- Use of single ended vector network methods on differential and common mode systems.

This clause explores both of these areas.

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

12.2 Naming conventions in high speed serial links

Significant confusion has existed concerning the naming of Sij in Fibre Channel and other high speed serial electrical links. There seems to be general agreement that Sij is the ratio of the signal coming out of the ith port to the signal coming into the jth port (for a 2 port linear element having ports i and j - the signals are differential or common mode for FC 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 basically two types of measurement: (1) return loss from the same port of the element and (2) transfer function or insertion loss across the element. In common parlance 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 14 shows the port naming conventions for link elements and loads.

- The following figure shows specifically where the element ports exist and how they are named
- Note that transmitter device port 1 and receiver device port 2 are internal and are not defined - they would be an ideal source and an ideal sink respectively

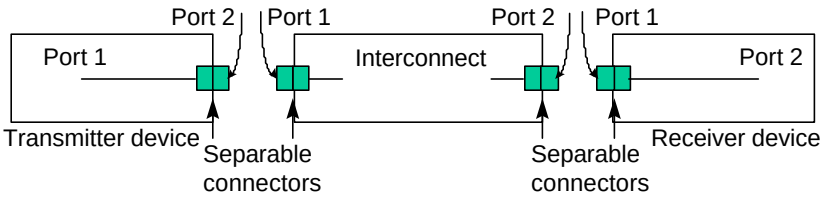

Port definitions for loads used for signal output testing and S-parameter measurements in multiline configurations

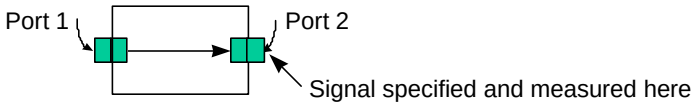

This load has ideal or ‘Golden’ differential and common mode properties

Figure 14 - Sij nomenclature conventions

12.3 Use of single ended instrumentation in differential applications

Figure 15 shows the connections that would be made to a four port vector network analyzer (VNA) for measuring S parameters on a four single ended port ‘black box’ device. VNA’s 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 document relate to differential signal pairs. It requires all four VNA 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.

Figure 15, Figure 16, and Figure 17 illustrate the conventions and requirements for differential measurements. Common mode measurements use a similar scheme but are not specifically documented here.

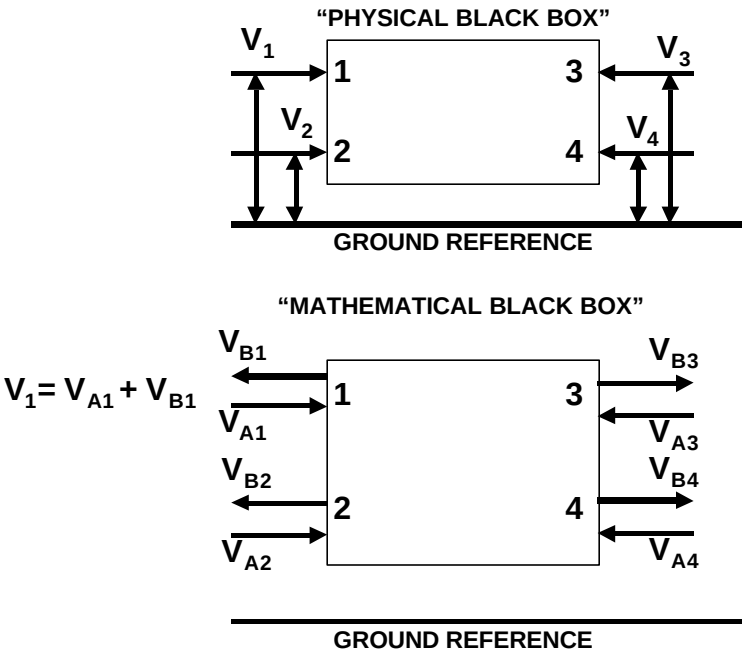

Figure 15 - Four single ended port element

Figure 16 shows the definition of the differential ports and the differential S parameters that may be acquired from a two pair element. Since the VNA has only four single ended ports only two differential ports may be measured at one time. Physical reconfiguration is required to access all the differential S parameters listed in Figure 17.

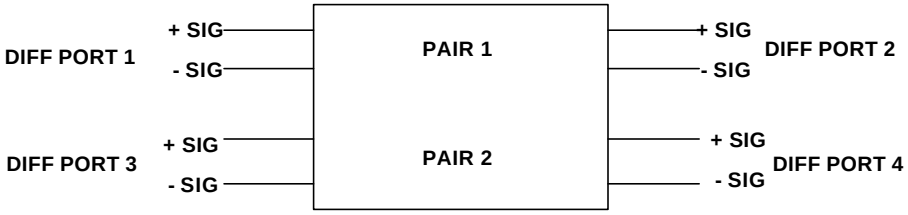

SDD11 = DIFFERENTIAL RETURN LOSS FROM DIFF PORT 1
SDD22 = DIFFERENTIAL RETURN LOSS FROM DIFF PORT 2
SDD33 = DIFFERENTIAL RETURN LOSS FROM DIFF PORT 3
SDD44 = DIFFERENTIAL RETURN LOSS FROM DIFF PORT 4
SDD21 = DIFFERENTIAL INSERTION LOSS AT DIFF PORT 2 FROM DIFF PORT 1
SDD31 = DIFFERENTIAL NEAR END CROSS TALK AT DIFF PORT 3 FROM DIFF PORT 1
SDD41 = DIFFERENTIAL FAR END CROSS TALK AT DIFF PORT 4 FROM DIFF PORT 1
SDD42 = DIFFERENTIAL NEAR END CROSS TALK AT DIFF PORT 4 FROM DIFF PORT 2
SDD43 = DIFFERENTIAL INSERTION LOSS AT DIFF PORT 4 FROM DIFF PORT 3
SDD32 = DIFFERENTIAL FAR END CROSS TALK AT DIFF PORT 3 FROM DIFF PORT 2
SDD31 = DIFFERENTIAL NEAR END CROSS TALK AT DIFF PORT 3 FROM DIFF PORT 1
SDD12 = DIFFERENTIAL INSERTION LOSS AT DIFF PORT 1 FROM DIFF PORT 2
SDD13 = DIFFERENTIAL NEAR END CROSS TALK AT DIFF PORT 1 FROM DIFF PORT 3
SDD14 = DIFFERENTIAL FAR END CROSS TALK AT DIFF PORT 1 FROM DIFF PORT 4
SDD23 = DIFFERENTIAL FAR END CROSS TALK AT DIFF PORT 2 FROM DIFF PORT 3
SDD24 = DIFFERENTIAL NEAR END CROSS TALK AT DIFF PORT 2 FROM DIFF PORT 4

Figure 16 - Definition of possible differential S-parameters for a 2-pair differential element

Figure 17 shows all the possible configurations for a two pair link element. Note that termination is required on all differential ports that are not connected to the VNA.

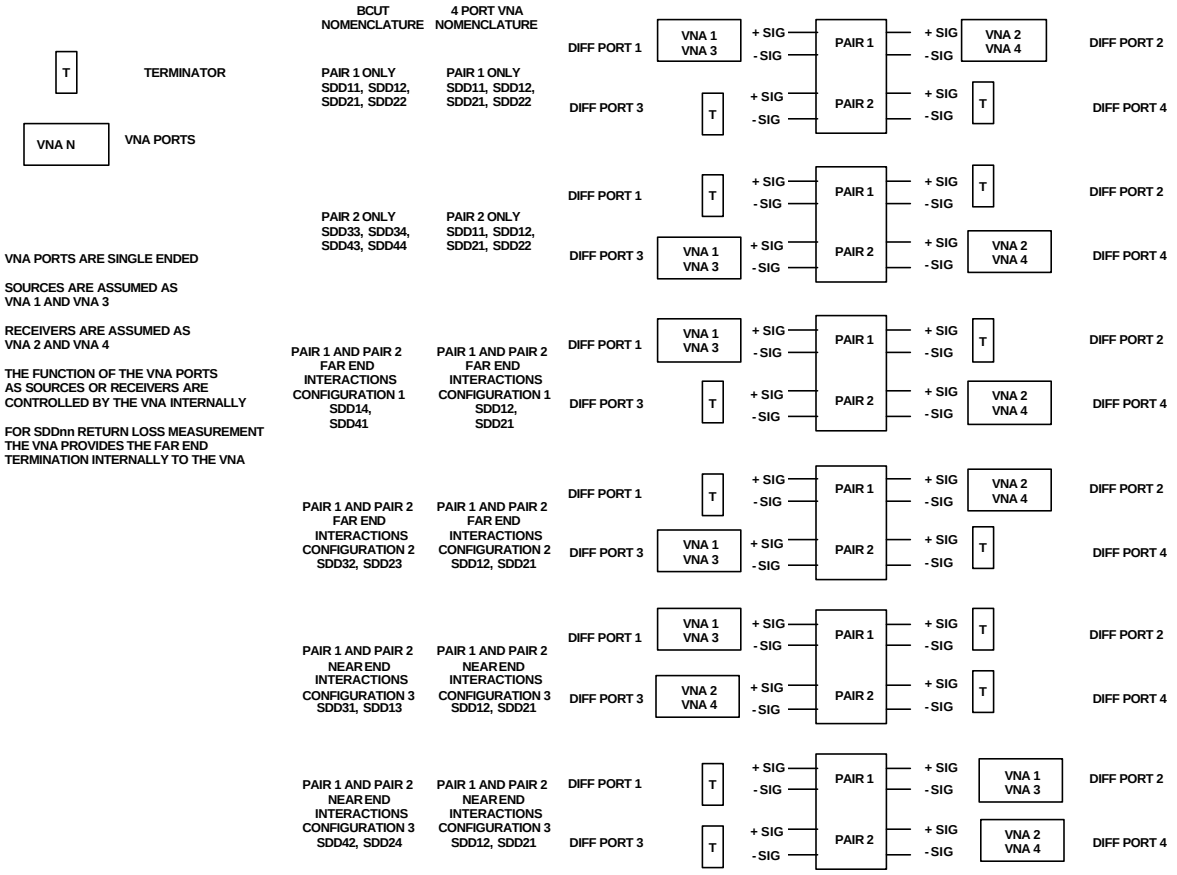

Figure 17 - Measurement configurations for a 2-pair differential element using single ended instrumentation

Although the above figures show the connections for a VNA the same information may be obtained by using a TDR/TDT measurement to acquire a time domain waveform. The waveforms may then be converted to S parameters using a suitable software package.

12.4 Measurement configurations for link elements

12.4.1 Overview

Special test fixtures are required to make S-parameter measurements in part 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.

Sub clause 12.4 describes the measurement configurations used for the four common conditions:

- Transmitter device return loss
- Receiver device return loss

- S11 at the transmitter device connector (interconnect input)
- S22 at the receiver device connector (interconnect output)

12.4.2 Transmitter device return loss

Figure 18 shows the configuration to be used for the transmitter device return loss measurement.

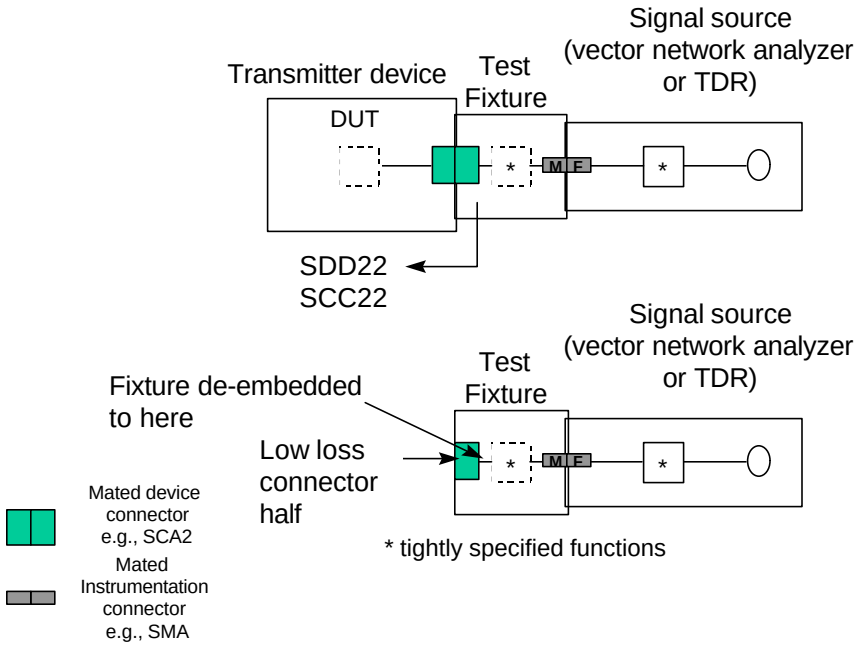

Figure 18 - Measurement conditions for transmitter device S22 requirements

Notice that the test fixture 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 be that much better to accommodate.

The test fixture losses up to the mounting points for the device connector are de-embedded using the methods described in Figure 6.

12.4.3 Receiver device return loss

Figure 19 shows the configuration to be used for the receiver device return loss measurement.

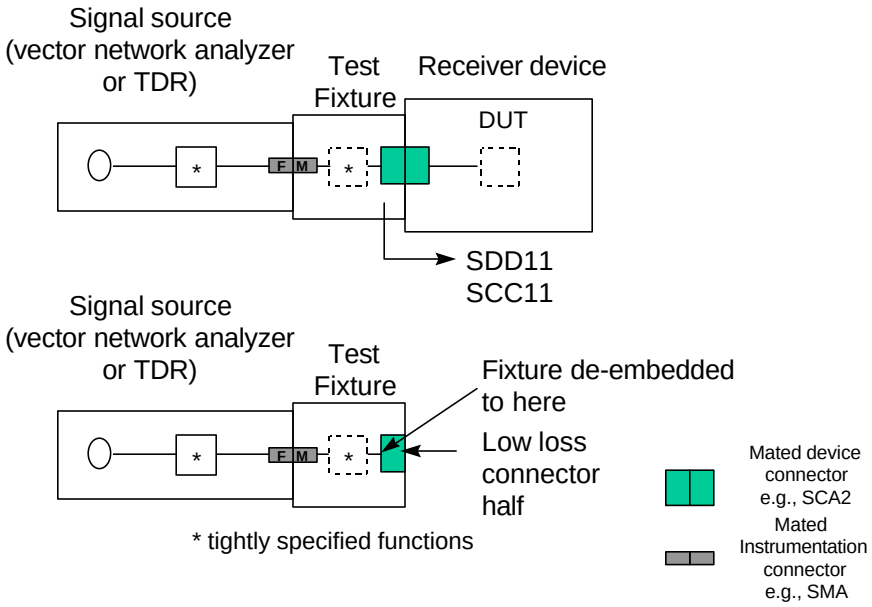

Figure 19 - Measurement conditions for receiver device S11 requirements

Notice that the test fixture 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 be that much better to accommodate.

The test fixture losses up to the mounting points for the device connector are de-embedded using the methods described in Figure 6.

12.4.4 S11 at the transmitter device connector (interconnect input)

Figure 20 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 will not represent in service conditions. If the interconnect is very lossy then the effects of the load on the far end (where the receiver device would be) 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.

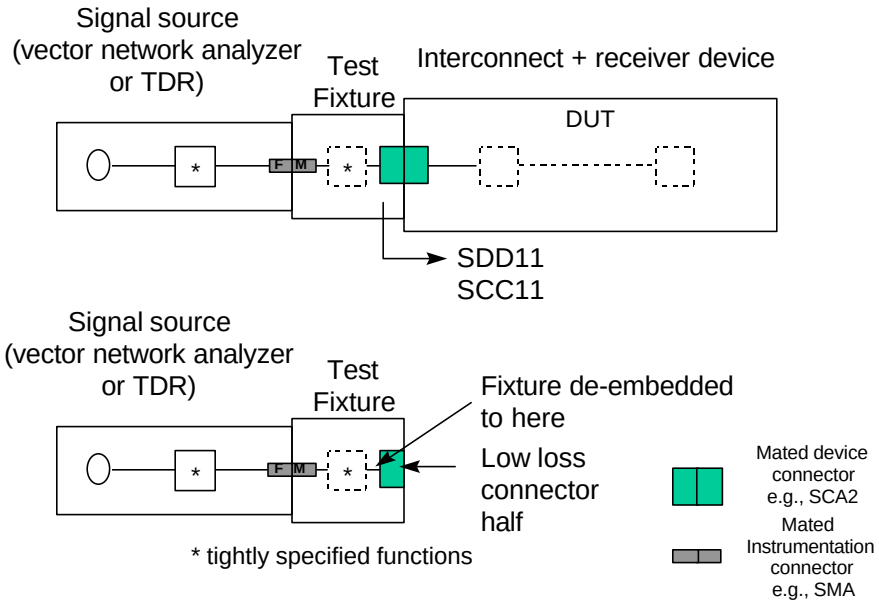

Figure 20 - Measurement conditions for S11 at the transmitter device connector

12.4.5 S22 at the receiver device connector (interconnect output)

Figure 21 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 is sort of like a reverse direction signal tolerance measurement. If the transmitter device is replaced by an ideal load then the return loss will not represent in service conditions. If the interconnect is very lossy then the effects of the load on the far end (where the transmitter device would be) 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.

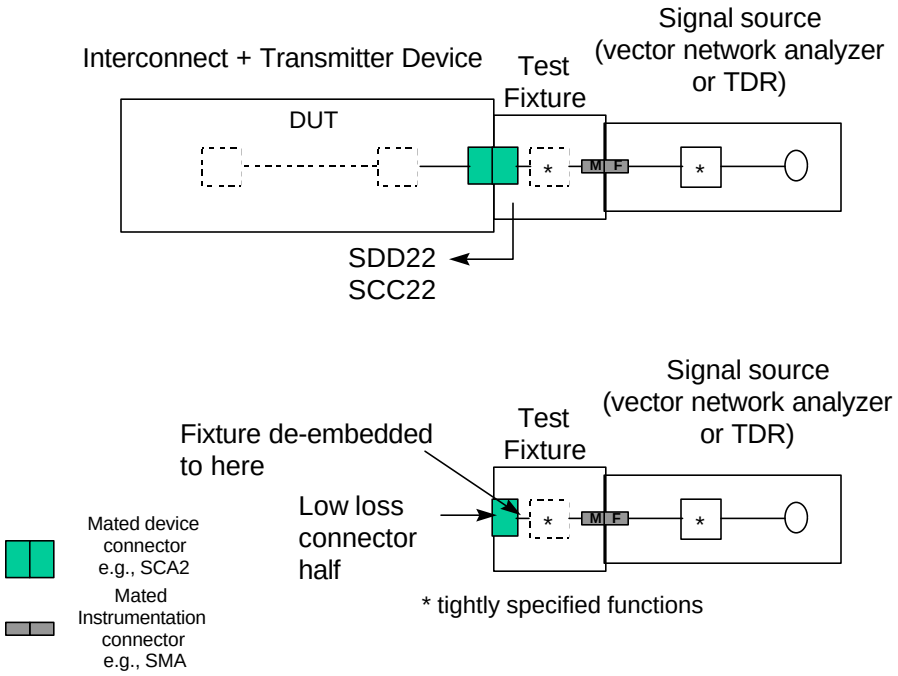

Figure 21 - Measurement conditions for S22 at the receiver device connector

12.5 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 essential.

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

The main power of the S-parameter methodology is the model for the element that results if a complete set of S-parameters is measured (including common mode and crosstalk from other sources in the environment. This model may be used with time domain waveforms and other known application conditions to simulate the resulting signal at the interoperability point of interest.

Unfortunately this simple sounding statement understates the complexity of the problem of developing effective numerical pass fail limits.

13. Development of performance specification values

This revision of this document has not incorporated the methodologies required to generate the performance specification values. However, some key points are listed somewhat randomly below as a framework.

Since receivers detect the state of the signal incoming to the receiver circuit at a particular signal level and within a particular window of time (not of frequency) the specifications for signal performance are done in a time vs. signal level manner. This is typically an eye specification of some sort. Attempting to place boundaries on frequency domain spectra properties is not a good idea for signal specifications.

The interoperability penalty is a huge consideration that appears to be only approachable via Monte Carlo simulation methods where the important properties of the link elements are varied over the entire tolerance range that is allowed. Among the important properties to vary are the length, capacitance and inductance values, features of the S-parameter properties at different frequencies, crosstalk levels, and common mode properties.

Generation of S-parameter properties is likewise best approached via Monte Carlo variation of the properties known to affect the losses.

The eight different properties are required at every interoperability point (See Table 1) must be incorporated into the methods for generating performance value limits. Formerly it was common to have only one or two values (e.g. meet the differential eye and impedance level for example).

The limited utility of making measurements in operating systems for determining compliance of components is only recently becoming understood.

The use of compensation in transmitter devices, interconnect, and receiver devices may cause interactions between the compensation methods that require restrictions on the type, intensity, and physical placement of the compensation. For example, compensation methods that assume properties of interconnect place a dynamic tension between the implementation of the compensation and the freedom to design interconnect.

The basic foundation has been laid for how to do the measurements and how to attain independence of one link component from another. Methods for extracting the benefits of the compensation done in transmitter devices and receiver devices are now known.

The challenge now emerging across several new applications is how to develop the numbers for the specifications that allow true element interchangeability and still allow the link to have the speed, distance, and robustness required.

In some sense this problem of developing the performance specifications that deliver

true interoperability is not new at all. It is just a set of design trade-offs. However, there are some new features.

One thing new is we have been working for a long time with large hidden margins in the link receivers (like factors of 4) that allowed a multitude of suboptimal conditions to exist and still have the link deliver a usable BER. These large hidden margins are disappearing as the higher data rates force signals to shrink and the basic signal integrity to disappear into jitter and noise. We are finally at the point where we must address this basic specification architecture since we are running out of this hidden margin.

Another new thing is the realization that the interoperability penalty may be quite high.

Yet another is that return loss is a key in determining whether short links can even work at all.

No scheme comes without a cost. Multilevel encoding, for example, comes with a high power cost even if the highest frequency in the systems can be limited by the encoding.

14. Summary

If performance specification values exist in a standard, this document can help to get the data that determines whether compliance to the standard exists.

The much harder question to answer is whether the specifications in these standards actually deliver the performance required for the interoperability expected. Unless the methodology described in this document is used for both the creation of the values in the standards and for the measurement of the performance at the interoperability points for the components of the link it is likely that things like interoperability penalty, test fixture deimbedding, and compensation methodologies were not adequately accounted for. Present systems actually work largely because the components used significantly exceed the requirements (and therefore carry an intrinsically higher cost than necessary).

There is a fundamental difference between having a set of hardware and software come together and work in a plugfest and truly effective set of specifications that make it very likely that ANY compliant collection of link elements will deliver the required error rate performance when connected together. Even if the specifications are quite poor it is very likely that a plugfest methodology will not show any problems because very few if any of the components used are actually worst case (or even close to worst case).

One should not worry when plugfest results show a lot of problems. Under those conditions work can be done to identify and eliminate the gross problems that prevented the successful operation. The real worry should be when the plugfest shows no problems and no one has any idea about how far from actual failure the collections of components used for the plugfest really are. It is very likely that the problems will show up later after significant production quantities have been made and when it is much more expensive to fix. It is also commonly very difficult to determine which component is at fault in this case.