

Passive interconnect performance working group (PIP) 00-379r1
December 12, 2000
Nashua, NH

Subject: Approved minutes for the SCSI passive interconnect performance working group, PIP, in Santa Cruz CA on October 10, 2000

In the absence of both Zane Daggett of Hitachi, chair and Dave Chapman of Amphenol, vice chair, Bill Ham, secy led the meeting. Bill Ham of Compaq, secretary, also took these minutes. There was a good attendance from a broad spectrum of the industry. Umesh Chandra of Seagate hosted the meeting.

Previous approved minutes: 00-301r1

1. Introduction.....	2
2. Attendance.....	2
3. Agenda development.....	2
4. Approval of previous minutes.....	3
5. Review of action items.....	3
6. Administrative structure:.....	3
7. Review of industry activities.....	3
8. Presentations on new topics.....	3
8.1 ATN vector network analyzer presentation, Bill Ham (for ATN).....	4
8.2 Additional measurement results from an extended set of cable media, Bruce Manildi, Seagate.....	4
9. Effects of non-uniformities and/or periodic structures.....	4
9.1 Effects of periodic structures continued, Larry Barnes, LSI Logic.....	4
10. Rationalization of existing attenuation data, all.....	4
11. Exploratory cable media round robin (round robin 1).....	6
11.1 Progress report, Umesh Chandra, Seagate.....	6
11.2 Testing methodology and parameters.....	9
11.3 Participating companies, all.....	10
11.4 Present status.....	10
12. Proposed new round robins.....	10
12.1 Expanded parameter set round robin 2, Chuck Grant, Madison Cable and Greg Vaupotic, Amphenol Spectra Strip.....	11
12.2 Backplane round robin 1, Umesh Chandra, Seagate.....	12
12.3 Cable Assembly round robin 1, Martin O., Molex.....	12
13. Frequency dependence of dielectric constant test methodology - Barnes, LSI Logic.....	12
14. Worst case vs spec numbers strategy.....	13
15. SCSI passive interconnect as an N-port construction - all.....	13
15.1 Local neighborhood concepts.....	14
15.2 Length measurements.....	15
15.3 Interoperability points.....	15
15.4 Approach to concatenated constructions.....	17
16. Project proposal for PIP - all.....	17
17. PIP documentation - Daggett.....	18
17.1 Topics for consideration for the passive interconnect test document.....	18
17.1.1 Components of passive interconnect.....	18
17.1.2 Construction.....	18
17.1.3 Specific technical concentration areas.....	19

17.1.4 Test types	20
17.1.5 Instrumentation / measurement methods:	22
17.1.6 Acceptable performance values	22
18. Starting point for specifying application of tests.....	22
18.1 Definitions for the document - Barnes	23
19. Architectural definitions.....	23
20. Next meetings.....	24
21. Action Items:.....	24
21.1 Old action items from previous meetings	24
21.2 New actions from this meeting	25

1. Introduction

In the absence of both Zane Daggett and Dave Chapman, Bill Ham, secy, opened the meeting, conducted the introductions, and reviewed the meeting purpose.

2. Attendance

The following folks were present:

Name	Company	E-Mail
Paul Aloisi	TI	Paul_Aloisi@TI.com
Larry Barnes	LSI	larry.barnes@lsil.com
Joe Basista	C&M	jbasista@cm-corp.com
Umesh Chandra	Seagate	umesh_chandra@notes.seagate.com
Jason Chou	Foxconn	JasonC@foxconn.com
Robert Christopher	IBM	bobchris@us.ibm.com
Steve Gasparovic	EDA (IO2GO)	steve@io2go.com
Don Getty	TI	donald_getty@ti.com
Chuck Grant	Madison Cable	cgrant@madisoncable.com
Bill Ham	Compaq	bill_ham@ix.netcom.com
Lee Hearn	Adaptec	lee_hearn@corp.adaptec.com
Thom Kreusel	HP	thom_kreusel@hp.com
David MacQuown	Adaptec	david_macquown@corp.adaptec.com
Bruce Manildi	Seagate	bruce_manildi@notes.seagate.com
Richard McMillan	Adaptec	Richard_mcmillan@corp.adaptec.com
Martin Ogbuokiri	Molex	mogbuokiri@molex.com
Bill Troop	IBM	troop@us.ibm.com
Richard Uber	Quantum	richard.uber@quantum.com
Greg Vaupotic	Amphenol Spectra-Strip	greg.vaupotic@snet.net

3. Agenda development

The agenda shown was that used (moved by Vaupotic / seconded by Aloisi)

4. Approval of previous minutes

Bill Ham moved and Greg Vaupotic seconded that the draft minutes from the previous meeting be approved as modified. Motion passed unanimously.

The methodology for minutes uses the draft/approved minutes scheme with posting to the t10 web site of the minutes as the vehicle for publication. Postings are announced to the SCSI reflector after the posting is verified to be on the web site.

Minutes will be in .pdf format.

5. Review of action items

Action items were reviewed and the status is listed below in the action items section.

6. Administrative structure:

The present administrative structure is:

Chair: Zane Daggett, Hitachi
Vice Chair: Dave Chapman, Amphenol
Secretary: Bill Ham, Compaq

Document editors: Zane Daggett, editor in chief, Bill Ham, assistant editor, Greg Vaupotic, assistant editor, others welcome (but work is expected)

7. Review of industry activities

Bill Ham briefly reviewed the T10, T11, and SFF activities relating to testing and modeling. He noted that the T11 modeling activity is now being headed by Dean Wallace and that the first meeting is scheduled for the December T11 week in Austin, TX.

8. Presentations on new topics

8.1 ATN vector network analyzer presentation, Bill Ham (for ATN)

Bill Ham went through most of the slides recently presented to the SFF HPEI SSWG. With help from Greg Vaupotic most of the main points were communicated. This vector network analyzer methodology seems to have significant applicability to the SCSI measurement challenge.

It is expected that this methodology will be used for some measurements.

8.2 Additional measurement results from an extended set of cable media, Bruce Manildi, Seagate

A document exists but no number has been assigned. This was similar to the round robin 1 except that different baluns were used and several different cable media types were measured. The expanded sample set included twist and flat and several different kinds of round shielded types. Suppliers included Madison, Hitachi, and Spectra Strip. Lengths ranged from 38.5 to 250 ft. A large set of swept frequency plots were shown. The details are too complex to attempt to describe in the minutes.

This presentation will be uploaded.

Bruce Manildi is actioned to put the presentation from the extended set of cable media on the T10 web site.

9. Effects of non-uniformities and/or periodic structures

9.1 Effects of periodic structures continued, Larry Barnes, LSI Logic

Larry presented a spread sheet that reports the loaded impedance and velocity of propagation based on frequency, unloaded impedance, separation of loads, capacitance of loads assuming no stub length. Larry agreed to make this spreadsheet available on the web.

Larry Barnes is actioned to place his spreadsheet on the T10 web site.

10. Rationalization of existing attenuation data, all

[This table was not updated but a key factor in the disagreements between the Seagate results and the other measurements was noted.]

The table generated in the last meeting was re-examined. A summary of the data available at meeting time is given in the table below.

The data in the table is for cable assemblies or media only, media impedance varies in the 128-132 ohm range, data presented in units of dB/m, fail limit set at 12 dB @ 25m @ 200 MHz in SPI-3. Note that there

are significant differences in the measurement methodologies and that direct comparisons may not be valid. Nevertheless, the data presented shows the magnitude of difference that has been reported. This framework will be used to compare the round robin data.

	Seagate HP 8110 pulse generator with clock-like pattern with slow edges, spectrum analyzer used for measuring outputs and inputs	Hitachi network analyzer with 100 ohm baluns	Madison 4-port network analyzer	Amphenol twisted pair ribbon (no flats) network analyzer with TP101 with 121Ω matching pads
Sample descripti on	cable assys	cable media	cable media	cable media
80 MHz				
30 strand		0.32/m (8/25m)	0.31/m (7.75/25m)	
30 solid	0.16/m?? (4/25m) media type not known for sure at time of writing	0.26/m (6.5/25m)		TPE 0.26dB/m (6.5/25m)
28 strand	0.12/m (3/25m)	0.19/m (4.75/25m)		
28 solid			0.25/m (6.25/25m)	
200 MHz				
30 strand		0.59/m (14.75/25m)	0.52/m (13/25m)	
30 solid	0.38/m?? media type not known for sure at time of writing	0.44/m (11/25m)		TPE 0.46dB/m (11.5/25M)
28 strand	0.28/m (7 dB/m)	0.36/m (9/25m)	0.40/m (10/25m)	
28 solid			0.40/m (10/25m)	

Note that there are few direct comparisons and some minor inconsistencies in the data in the above table. There is considerably more than the wire gauge that influences the attenuation. Thus the wisdom of doing a performance rather than a physical description specification.

Nevertheless, the present specification of 12 dB @ 25m @ 200 MHz seems to be reasonable given the 11 dB/25m data for 30 AWG solid.

It was noted that the Seagate data has backed out the d.c. contribution which is a couple of dB. This helps the agreement considerably.

11. Exploratory cable media round robin (round robin 1)


11.1 Progress report, Umesh Chandra, Seagate

[much of this material is retained for continuity - Madison data and a summary table was added]

The following samples have been collected:

- 30AWG solid round from Hitachi 57m (190 ft) [sample 1]
- 28AWG stranded round from Hitachi 57m (190 ft) [sample 2]
- 30 AWG solid twist and flat [sample 3]

A test procedure has been documented:



Test Procedure using HP4396A/45046A

- Apply 'shorts' (thru) (2) in place of DUT
- Set 'SOURCE' level to +20 dB, 'START' to 0.5 MHz 'STOP' to 300 MHz
- 'CAL' -> 'Cal Menu' -> 'Response' -> 'Thru' -> 'Done'
- 'SAVE' -> 'State' -> 'Done'
- Replace 'shorts' with the DUT (Cable under Test)
- Collect Data
- 'SAVE' -> 'ASCII Save' -> 'Data Only (ASCII)'
- Import file into XL and create graph

Test results:

- 30 AWG Solid Hitachi #49557-068 SCSI3
- dc R = 20.5 OHMS
- CHANNEL: 1
- MEASURE TYPE: S21
- FORMAT TYPE: LOG MAG
- NUMBER of POINTS: 201
- SWEEP TIME: 70 ms
- SWEEP TYPE: LIN FREQ
- SOURCE POWER: 20 dBm
- IF BANDWIDTH: 40 kHz

Data summary from Seagate:

Attenuation at 200 MHz for sample 1 9.21 dB/25m [30AWG solid] average of some 2 pairs

Attenuation at 200 MHz for sample 2 10.82 dB/25m tan/white single pair [28AWG stranded]

Attenuation at 200 MHz for sample 3 the number is somewhat meaningless (partly due to making the measurement in a "bunched up" mode)

Hitachi data:

Test set up: same as SPI-3 except used 100 ohm baluns (350 MHz) and no matching pads

Attenuation at 200 MHz for sample 1 9.25 dB/25m [30AWG solid] average of all 34 pairs

Attenuation at 200 MHz for sample 2 10.39 dB/25m [28AWG stranded] average of all 34 pairs 10.60 dB/25m for the same pair as used by Seagate

Attenuation at 200 MHz for sample 3 not measured

Madison data:

Attenuation at 200 MHz for sample 1 [30 AWG solid] 9.22 dB/25m average for all 34 pairs ranging from 8.94 to 9.67 4 port VNA used for both samples

Attenuation at 200 MHz for sample 2 10.39 dB/25m average for all 34 pairs ranging from 9.83 to 11.03 [28 AWG stranded] tan/white was 10.77
 Attenuation at 200 MHz for sample 3 not measured

Data summary:

company	sample 1 (dB/25m)	sample 2 (dB/25m)	sample 3 (dB/25m)
Seagate	9.21	10.82 tan/white	na
Hitachi	9.25	10.39 avg	na
Madison	9.22	10.39 avg 10.77 pair tan/white	na

Conclusions so far:

- One conclusion is that measuring either way gives very similar results.
- Another conclusion is that the pulse amplitude method used previously by Seagate does not match well the results from the network analyzer methods (although the difference is less when d.c. attenuation is accounted for).

The round robin will continue as planned. To those who have not yet reported their results: no fair adjusting the results because the answers are now available.

Umesh agreed to document the round robin activities and results. Umesh to create the draft document for the next meeting.

11.2 Testing methodology and parameters

Test methods to be modified as follows:

[Following copied from SPI-3 and modified appropriately]

E.7.5 Measurement test fixture and measurement equipment

An instrument capable of supplying a sinusoidal signal is used as the signal source and an instrument capable of detecting the amplitude of a sinusoidal signal is used as the signal sink. Two measurement test fixtures are required: one for the source end and one for the sink end. Since most source and sink instruments capable of using variable frequency sinusoidal signals are single ended, a balun or a hybrid (M/A - Com) may be used between the instruments and the test fixtures. Impedance matching networks may be required. (Lpad). If a source or sink is used differential signals then no balun is required for the differential source or sink.

Equipment Required: Network Analyzer (HP 87xx Series) or equivalent.

The test fixture having 61 ohm single ended paths for each signal line shown in SPI-3 will NOT be used for the tests. Samples shall be directly soldered to the Lpads. The opens, shorts, through calibration method shall be used.

The tested lines shall be terminated by using 122 ohm nominal resistors between the + signal and the - signal line. [some observations have indicated different results when using the terminations instead of keeping the untested line open as presently implied in SPI-3]. The question of whether to require this termination may be answered in subsequent testing.

11.3 Participating companies, all

Companies agreeing to participate in this initial exploratory measurement effort: Seagate, Hitachi, Madison, Amphenol Spectra-Strip, Adaptec,

Contacts:

Hitachi: Clint Heiser (603 669-4347 ext 362)
Seagate: Bruce Manildi (831 439-7229) [Coordinator]
Amphenol Spectra-Strip: Greg Vaupotic (203 287-7425)
Adaptec: Lee Hearn
Madison: Chuck Grant (508 752-2881 ext 306)

Samples requested: all round shielded raw media 25 meters long, one sample of each from Hitachi and Madison of 28AWG solid, 28 AWG stranded, 30 AWG stranded, 30 AWG solid (total of 8 samples)

11.4 Present status

The following order was used:

Hitachi: Clint Heiser (603 669-4347 ext 362)
Seagate: Bruce Manildi (831 439-7229) [Coordinator]
Madison: Chuck Grant (508 752-2881 ext 306)

Samples just delivered to:

Amphenol Spectra-Strip: Greg Vaupotic (203 287-7425)
Adaptec: Lee Hearn

12. Proposed new round robins

12.1 Expanded parameter set round robin 2, Greg Vaupotic, Amphenol Spectra Strip and Chuck Grant, Madison Cable

Round robin 2 is based on a significantly more precise specification of the measurement details.

Specifically the following is proposed for round robin 2:

- Attenuation from 1 MHz to 1 GHz with specific values recorded at 80 and 200 MHz in addition to the plot. Use scheme 2 described below.
- Differential impedance (swept frequency 30 Hz IF bandwidth, 801 points across the range of 10 MHz to 1 GHz) Measurement to be done using both open and short for far end termination and calculating the reported result: $Z_{diff} = (Z_{oc} \times Z_{sc})^{1/2}$ Scheme 1 used for test fixture
- Skew within the pair (use TDT per SPI-3)
- Samples are: one round (26 meters) and two flat with different constructions (both also 26 meters)
- Number of pairs to be measured for all measurements: round: 2 specific identified pairs in the inner layer; 2 pairs in the middle and 2 in the outer; flat: pair 1, 2, 5, 17 (pair 1 contains a unique color code, other pairs are counted from pair 1)

Samples to be measured should include both uniform and non uniform constructions. Other details will be proposed by Greg.

The impedance, skew within the pair and the sample length has direct impact on attenuation data. The fixture is important as well. When we report out test results, we should include all of the above so that we do not mislead people.

Greg and Chuck are actioned to prepare a detailed proposal for samples, measurements to be performed, and logistics.

Umesh showed two approaches to test fixtures and calibration: (1) based on separate hybrid junctions (also known as magic T) (50 to 100 ohm) and the (2) based on a balun (50 SE to 50 DIFF) with L pads for matching.

Greg Vaupotic offered the additional information:

The 100 ohm scheme (1) can be calibrated by using precision 50 ohm discrete opens shorts and loads which calibrates the system at the point where the test fixture is attached. Then the test fixture is attached and the reference plane is moved to the point where the IUT will be attached. The reference plane position is set by adjusting until the empty test fixture shows a zero degree phase response (that is expected for an open circuit). Linearity can be validated by using different known resistor loads. The 100 ohm scheme is intended for use with swept frequency impedance.

The 122 ohm scheme (2) (see document 00-339r0), is intended for use with attenuation and swept frequency measurements. Calibration cannot be used with traceable standards due to lack of available values and lack

of suitable attachment points. Calibration is therefore approached by recording the insertion loss of the pair of devices in a through calibration scheme. A two port calibration can also be done by using non-traceable standards and a smoother measurement result is expected. The through calibration is recommended because it is considered to be more reproducible and more robust.

12.2 Backplane round robin 1, Umesh Chandra, Seagate

This effort is needed to add backplane testing methodology to the test suite. The basic idea is to acquire some small selection of SCSI backplanes and send these around for testing at different companies.

Umesh has a non-product backplane design that could be used. Bill Ham thinks that he may have some older backplanes that are not products that could be used.

Umesh is actioned to select 2 boards for use in the round robin.

Ham is actioned to select 2 boards for use in the round robin.

Umesh is actioned to propose a set of tests to be used including things like test fixtures and specific slots to be measured.

12.3 Cable Assembly round robin 1, Martin O., Molex

Martin O. to propose a cable assembly round robin: sample set, tests, participants, schedule, etc.

Martin is actioned to create this proposal.

13. Frequency dependence of dielectric constant test methodology - Barnes, LSI Logic

Deferred (again) to December due to Larry Barnes' personal time constraints. This work is intended to use the HP polished probe method according to the following table.

Greg Vaupotic noted that data from the swept frequency slab method on partially processed material could be made available.

This discussion produced collaboration containing the following:

Polyethylene samples to be provided by Spectra Strip	slab method	coax probe method
--	-------------	-------------------

raw material processed into a slab	Spectra strip	LSI Logic
raw material extruded over single conductor 26 AWG with approximately 12 to 15 mils dielectric thickness	NA	LSI logic -- strip material off conductor and measure by inserting the probe through slit in the conductor to the opposite side (C-)
material removed from conductor processed into slab	Spectra strip	LSI Logic

The range of frequency will be at least 1 MHz to 1 GHz linear sweep.

The samples will be targeted to Larry by mid November 2000.

Greg Vaupotic and Larry Barnes are actioned to complete the work outlined in the above table.

14. Worst case vs spec numbers strategy

[This item was not discussed at this meeting but is still unresolved.]

This item is intended to address the basic issues created by using a collection of worst case component specifications instead of recognizing the typical integrated impact of a collection of real parts.

Worst case specifications frequently result in excessive demands to improve component performance. On the other hand, given a large population of components, there is some chance that end-to-end failures will be experienced unless the worst case component performance requirements are met.

Therefore, conditions for SCSI are approaching the point where some recognition of the overspecification that results from worst case methodologies is reflected in the component performance specifications. When the worst case numbers are added up for SPI-4 the SCSI segment will not work. Therefore, either the specification methodology or some component specifications need to change.

This is one reason why there is presently a focus on the interconnect specifications.

In any case there were no breakthroughs in this area at this meeting.

15. SCSI passive interconnect as an N-port construction - all

[This section was discussed in earlier meetings but will be retained in the minutes until transferred into the PIP document.]

It was previously agreed that the SCSI cable assembly will be considered as an N'-port element where every connector constitutes the approximate location of the ports. Since SCSI is a parallel bus every connector contains a multiplicity of somewhat independent ports (one for every differential signal).

For purposes of the SPIP work a lower case "N", n, refers to the number of the specific signal in a connector. An upper case "N'", N' refers to the number of a specific connector. Thus a SCSI passive interconnect is characterized by N' connectors and n signals. N' is used instead of N so that when referring to connectors or ports verbally there will be distinction. Typically n ranges from 1 to 27 for SCSI applications. N' is determined by the structure of the interconnect and ranges from 2 to 18 (16 devices + 2 terminators) in most cases.

Therefore, a SCSI passive interconnect many contain up to $18 \times 27 = 486$ ports. Each N'th port can be represented by a matrix of n ports. The structure of the matrix will be based on the names of the signals.

Each port is characterized by (1) the signal launched into the port and the signal reflected back from the launched signal (2) the signal transferred to the port from other ports in the cable assembly.

The signals delivered out of every port when the most degraded allowed signal is launched from every other port (one at a time), when the most aggressive noise sources are present on all other ports that can couple into the port under test, and when the resonant conditions are within acceptable bounds shall meet at least the minimum requirements for a received signal.

15.1 Local neighborhood concepts

[This section was discussed in earlier meetings but will be retained in the minutes until transferred into the PIP document.]

For signals, the basic idea is to not test for interactions that are insignificant to the port under test. For example in a flat cable signals removed from the signal under test by at least 5 signal pairs do not significantly couple into the signal under test and do not need to be considered. The level of interaction deemed to be significant is left to be defined.

For physical constructions the dimensional precision within which the construction shall be considered identical is 1/10 of the rise length of the fastest signal to be used in the interconnect. This is approximately 1 inch for 1ns rise time signals in shielded twisted pair media. In other words, two connectors placed 0.5 inch apart may be treated identically regardless of which is actually tested. Similarly, the placement of connectors on nominally identical flat ribbon cables

shall be considered identical if they are within 1 inch of being at the same position.

15.2 Length measurements

[This section was discussed in earlier meetings but will be retained in the minutes until transferred into the PIP document.]

The length of the interconnect needs to be understood. For example there are a few percent difference between different "length" parameters. For example it may require 103 feet of wire to produce a cable assembly with 100 feet overall connector to connector path length. (2-103 foot wires to produce a single 100 foot twisted pair). Cabling applied to a bundle of pairs (cable lay) also affects the total path length.

The electrical length is also important where the propagation time is part of the interest for the specification.

The following was proposed as the way that PIP will consider length issues:

That the length parameters be separated into two pieces:

- (1) the physical length along the geometrical center line (e.g. center line of the jacket for round cables to the center line of the unmated connector) of the completed cable assembly (not necessarily the actual wire length for any specific conductor)
- (2) the propagation time between electrical access points (typically connectors) in the cable assembly

Other lengths such as those internal to the media will NOT be used as descriptors in PIP. These internal lengths may be important for creating accurate models but are not essential to specify how to do proper measurements on cable assemblies and therefore do not belong in the PIP effort.

15.3 Interoperability points

[This section was discussed in earlier meetings but will be retained in the minutes until transferred into the PIP document. Bill Ham is actioned to extract this section for presentation at the next SCSI plenary meeting and to produce a version suitable for use in the SSM document.]

Interoperability points are physical points in the system where separable connectors exist and where it is required that the components on either side of the connector may be supplied from different compliant vendors. Following is a list where interoperability might be expected

in a SCSI segment. A "Y" following the position designation means that this will be considered an interoperability point for PIP purposes. Similarly, a "N" following the position designation means that the point will NOT be considered an interoperability point for PIP purposes.

Disk drive connector mounted directly on the disk drive (Y)

HBA connector external connector (Y)

HBA internal SCSI connector to internal cables (Y)

HBA internal SCSI connector to the mother board (N)

Motherboard SCSI connector where the mother board contains the HBA (in an ASIC) on board (Y)

Backplane connectors:

Any connector that directly accepts a disk drive or other SCSI device (Y)

Any connector that directly connects to an external cable assembly through an expander on the backplane (Y)

Any connector on an external cable assembly that connects to an external connector of an HBA (Y)

Any connector on an external cable assembly that connects to an external connector of a disk drive array containing an expander immediately behind the external connector (Y)

Any connector on an internal cable that directly connects to a disk drive or other SCSI device. (Y)

Question: should the external connector to a disk drive array that does not contain an expander be considered an interoperability point? The group agreed that this should NOT be an interoperability point until proven otherwise in the SSM group. (N)

The external connector to a box that has external cable assembly attached and an internal cable assembly attached internally to the same connector. (N)

Note: this means that one may NOT have a cable to cable connection at the bulkhead if interoperability is required.

Question: should separable connectors that belong to terminators be considered interoperability points? The group agreed that these connectors should be included in the interoperability suite.(Y)

15.4 Approach to concatenated constructions

[This section was discussed in earlier meetings but will be retained in the minutes until transferred into the PIP document.]

The group identified two basic situations:

In the first situation the SCSI passive interconnect performance is considered under the conditions where the bus segment interconnect consists of a single media type and construction. For example, in this situation, two dissimilar (e.g. round to flat) cable assemblies connected together in series would not have an interoperability point at the point of common connection. Similarly, a backplane connected directly to a round shielded cable would NOT have an interoperability point at the backplane connector.

While the first situation is relatively easy to construct performance requirements around, it leaves several important constructions without clear definition.

Thus the second situation:

Four examples are described where interoperability is probably expected in common constructions:

- where a short cable assembly is used between the HBA and the bulkhead in a PC-like packaging
- where a short cable assembly is used between the disk drive and the backplane
- where an HBA is used between the external bulkhead and both internal and external cables
- where an external cable is attached directly to a backplane

Each of these cases has the property that the performance at the connector is significantly affected by the details of the passive interconnect on BOTH sides of the connector. This complicates specifying unique performance requirements the connector because of interactions on both sides.

Because these are important practical applications, some approach is needed in the PIP work. This subject is left for further consideration by the modeling group and in future PIP efforts. Pending definition of a viable strategy for these situation 2 cases work in PIP will be focused on the situation 1 cases.

16. Project proposal for PIP - all

Bill Ham noted that the project proposal for PIP was approved at the T10 plenary in May. The project proposal is document number 00-238r0. The project number 1439-D. This numbering implies that PIP is a standard and not a technical report. Bill Ham is actioned to check this out.

Motion Ham/Vaupotic that the PIP ad hoc recommend to the T10 plenary that the PIP project be to develop a standard (if that is not already the case).

6 in favor, 0 opposed, 3 abstain.

17. PIP documentation - Daggett

This working group will proceed to develop an internal committee document Titled: "SCSI Passive Interconnect Performance Requirements" whose schedule is independent from SPI-x standards schedule. Zane Daggett is editor, Bill Ham and Greg Vaupotic are assistant editors. The document will follow the same general format as 99-219rx.

The document number is 00-160rx.

17.1 Topics for consideration for the passive interconnect test document

[This section was discussed in earlier meetings and modified slightly in this meeting. It will be retained in the minutes until transferred into the PIP document.]

The material in this section was reviewed from the last meeting as possible candidates for consideration for the PIP document 00-160rx.

17.1.1 Components of passive interconnect

The following constitute the basic building blocks of passive interconnect:

media (wire and backplane)
connectors
transition regions (connector termination / comb out / lacing regions / vias)

17.1.2 Construction

The following physical constructions are part of PIP:

Point to point:

- two connector shielded
- two connector unshielded

Multidrop:

- multi connector shielded (e.g. external daisy chain)
- multi connector unshielded
- multi connector backplanes

Stubs:

- backplane stubs (length of conductor extending beyond a terminator)
- device circuit board stubs
- unshielded cable stubs
- shielded cable stubs

Overall length and specific placement and properties of stubs are essential parts of the description of the construction. Note that the length and position may not be measured in inches but rather in nanoseconds.

In general the passive interconnects for SCSI are complex multiport circuits whose performance must be considered from every connector in the interconnect.

17.1.3 Specific technical concentration areas

- Cable assembly design
 - Non uniform media issues (e.g. twisted flat)
 - Transition regions for cable assemblies
 - Connector performance specifications
 - Connector variations
 - Assembly construction variations
- Backplane design
 - Distance between connectors
 - Trace impedance
 - Overall backplane size constraints
 - Power and ground distribution
 - Routing restrictions (e.g. no traces over breaks in power/ground planes)
 - Holes for cooling air
 - Mechanical rigidity
 - EMC containment
 - Effect of vias
 - Effects of connector attachment scheme (vias, pads)
 - Proximity effects of extraneous materials near the signal paths
 - Et cetera
- EMC - reference SFF-8410 for CMPT and EMR for emissions - applies to shielded versions only

- Susceptibility issue for backplanes?

17.1.4 Test types

[This list was not reviewed at this meeting but will be placed high on the agenda for the next meeting.]

The following tests are presently viewed as the candidate list

Local impedance
Extended distance impedance
Capacitance (SE, DF)
Frequency dependence of dielectric constant
Propagation time - differential signal for each signal pair
Propagation time skew - difference between pairs
+ signal to - signal balance - within the pair (balance degradation)
Attenuation (voltage transfer function) - within the pair
Attenuation skew (difference in voltage transfer function between pairs)
Eye diagrams (signal degradation)
Rise time degradation
Common mode (treat each pair as a single conductor) impedance
Common mode capacitance
Common mode noise
Near end crosstalk
Far end crosstalk
Attenuation to cross talk ratio (ACR)
EMC (CMPT, EMR) shielded versions only

The Level 1 and Level 2 approach described in SFF-8410 will be used. Level 1 is required for performance and has specific acceptable limits defined. Level 2 is diagnostic and has no specific limits defined.

There is some support for including additional swept frequency tests (possibly as level 1 tests) but these have not been defined. The issue of how to construct valid time domain signals from frequency domain measurements is part of this discussion. Also the of the choice of interoperability points significantly interacts with the test results and needs to be considered.

The initial cut at the frequency/time domain testing distribution is listed below:

- Local impedance (time domain)
- Extended distance impedance (frequency domain)
- Capacitance (SE, DF) (frequency domain)
- Frequency dependence of dielectric constant (frequency domain)
- Propagation time - differential signal for each signal pair (time domain and frequency domain - needs more work)

- Propagation time skew - difference between pairs (time domain and frequency domain - need more work)
- + signal to - signal balance - within the pair (balance degradation)
- Attenuation (voltage transfer function) - within the pair (frequency domain)
- Attenuation skew (difference in voltage transfer function between pairs) (frequency domain)
- Eye diagrams (signal degradation)
- Rise time degradation
- Common mode (treat each pair as a single conductor) impedance
- Common mode capacitance
- Common mode noise
- Near end crosstalk (time domain and frequency domain - needs more work)
- Far end crosstalk (time domain and frequency domain - needs more work)
- Attenuation to cross talk ratio (ACR) (frequency domain)
- EMC (CMPT, EMR) shielded versions only (frequency domain)

The following represents the present thinking on the division of tests between the level 1 and level 2 types.

Level 1

- Local impedance
 - Extended distance impedance
 - Propagation time - within the pair
 - A.C. signal degradation - all pairs to clock (Full signal characterization - e.g. Eye diagrams)
 - D.C. leakage to ground [impacts receiver bias / d.c. offset]
 - A.C. balance degradation within the pair (+ signal to - signal balance /common mode)
 - End to end d.c. resistance difference within the pair
 - Capacitance (SE, DF)
 - Near end crosstalk (for noise induced on wired-or signals)
 - Far end crosstalk (for noise induced on wired-or signals)
- EMC (CMPT, EMR) shielded versions only

Level 2

- Signal degradation within the pair (Full signal characterization - e.g. Eye diagrams [note that the full set of signals with respect to the clock is the level 1 requirement - a single signal performance is not adequate (example being different jitter and proptime on different signals)])
- Rise time degradation
- Frequency dependence of dielectric constant

- Attenuation to cross talk ratio (ACR)
- Attenuation (voltage transfer function) - within the pair
- Attenuation skew (difference between voltage transfer function pairs)
- pair to pair
- Propagation time skew - pair to pair

Problem areas needing future attention - not classified yet:

- Common mode impedance
- Common mode capacitance
- Common mode noise
- Resonance effects

17.1.5 Instrumentation / measurement methods:

- Baluns
- Eye diagram / signal degradation testing (including cross talk noise)
- Filtering schemes for eye pattern generation

These topics are in addition to other issues already identified for media.


17.1.6 Acceptable performance values

All level 1 tests will have specific acceptable values proposed.


18. Starting point for specifying application of tests

The following configurations were agreed to be used for the next steps in developing specifications for the application of the tests.

SIMPLEST POINT TO POINT



SIMPLEST MULTIDROP



The point to point construction allows inclusion of connectors and transition regions, multiple line cross talk, resonance issues and other issues not previously considered in uniform cable media.

The simplest multidrop construction adds a single connector and greatly increases the number of possible interactions between lines as well as adding a known non-uniformity in addition to the connectors and transitions regions in the point to point construction.

Clearly these simple constructions do not have the desired complexity of some interesting applications. These more complex constructions, such as backplanes will be considered after we complete these initial constructions. [Note: the complexity for these simple constructions is significantly greater than that previously considered.]

18.1 Definitions for the document - Barnes

Not addressed at this meeting.

19. Architectural definitions

This refers to issues like defining the test points, nomenclature, and the like. It was decided to use the same conventions commonly used for modeling and transmission lines if possible. Larry B to propose a specific syntax for the next meeting.

All measurements will be through a mated connector. This means that the test fixturing specification will be critical since part of the tested interconnect will remain with the test environment and part will be removable with the IUT.

Zane is creating a summary table for all tests defined above and to start the document.

It was agreed that a special filtering function is needed for some tests to account for the filtering that may occur in the receivers. See 00-149r0 for more detail.

20. Next meetings

Approved schedule:

December 12-13, 9AM to 5 PM 12/12; 9AM to 12PM 12/13, Manchester, NH
(Hitachi Cable)

Requested schedule:

February 20-21, 9AM to 5 PM 02/20; 9AM to 12PM 02/21, Southern CA
(Foxconn)

April 3-4, 2000 9AM to 5 PM 04/03; 9AM to 12PM 04/04 Worcester, MA
(Madison)

21. Action Items:

21.1 Old action items from previous meetings

Larry Barnes to acquire data from the polished coax probe method for dielectric constant frequency variations.
Status: equipment now in hand, test results now expected before December meeting due to time availability

Zane to provide data from the HP slab method for dielectric constant frequency variations
Status: carried over

Bill Ham to post the minutes to the T10 web site
Status: done 00-301r0

Zane to create a summary table for all tests defined above and to start the document.

Status: carried over

Larry Barnes to look at list of proposed tests and suggest revisions / additions to incorporate possible frequency domain tests and design validation tests and production tests.

Status: carried over

Bill Ham to create a minutes framework that includes all the necessary boilerplate. The T11 format will be used as a starting point.

Status: done

Greg Vaupotic to supply details on how to build the interface circuit for attenuation testing.

Status: done 00-339r0

Bruce M. to initiate collection of samples for testing.

Status: done

Larry B to propose a specific syntax for defining the test points, nomenclature, and the like for the next meeting.

Status: carried over

Greg/Ham to get the balun description posted

Status: done 00-339r0

Umesh to create the draft document for round robin 1 for next meeting.

Status: done to be presented at this meeting

Umesh to forward the samples from round robin 1 to Jie Fan

Status: done

Greg and Jie to prepare a detailed proposal for samples, measurements to be performed, and logistics for round robin 2.

Status: carried over pending group decisions on measurement methodologies to be used

Umesh to select 2 boards for use in the backplane round robin.

Status: done

Ham to select 2 boards for use in the backplane round robin.

Status: partly done new

Umesh to propose a set of tests to be used including things like test fixtures and specific slots to be measured.

Status: carried over

21.2 New actions from this meeting

Bill Ham to post the minutes to the T10 web site

Status: new

Greg V to provide Larry with polyethylene samples per section 13.
Status: new

Larry Barnes to execute the tests defined in section 13.
Status: new

Bill Ham to create a document outlining the interoperability points
issues for SCSI
Status: new

Larry Barnes to place his spreadsheet for backplane on the T10 web site.
Status: new

Bruce Manildi to put the presentation from the extended set of cable
media on the T10 web site.
Status: new

Umesh to create a draft document describing the details of and results
from the cable media round robin 1 for the next meeting
Status: new

Martin to create a proposal for cable assembly round robin 1
Status: new

Greg Vaupotic and Larry Barnes to complete the work outlined in the
table relating to the frequency dependence of dielectric constant.
Status: new

Bill Ham to extract this section on interoperability points for
presentation at the next SCSI plenary meeting and to produce a version
suitable for use in the SSM document.
Status: new