Proposed Working Draft

X3T10/1143

Revision 04
Jan 23, 1996

Information Technology - SCSI Enhanced Parallel Interface - EPI

This is an internal working document of X3T10 of the Accredited Standards Committee X3. As such, this is not a completed report and has not been approved by Task Group X3T10. The contents may be modified by the X3T10 Task Group. This document is made available for review and comment only.

Permission is granted to members of X3, its technical committees, and their associated task groups to reproduce this document for the purposes of X3 standardization activities without further permission, provided this notice is included. All other rights are reserved. Any duplication for commercial or for-profit use is prohibited.

[Technical Editor's Note: This Rev contains a proposed outline and content based on technical information and editing time available to the editor as of Jan 23, 1996.

ASC X3T10 EPI Technical Editor:

Bill Ham
Digital Equipment
334 South Street
Shrewsbury MA 01545
(508) 841-2629 Ph
(508) 841-5266 Fx
HAM@subsys.enet.dec.com

Other Points of Contact:

X3T10 Chair
John Lohmeyer
Symbios Logic Inc.
1635 Aeroplaza Drive
Colorado Springs, CO 80916
Voice: (719) 573-3362
Fax: (719) 597-8225
Email: John.Lohmeyer@symbios.com

X3T10 Vice-Chair
Lawrence Lamers
Adaptec
691 Milpitas Blvd
Milpitas, CA 95035
Voice: (408) 957-7817
Fax: (408) 957-7193
ljlamers@aol.com

X3 Secretariat
Lynn Barra
Administrator Standards Processing
X3 Secretariat
1250 Eye Street, NW Suite 200
Washington, DC 20005
Telephone: 202-626-5738
Facsimile: 202-638-4922

SCSI Reflector
Internet address for subscription to the SCSI reflector: scsiadm@wichitaks.ncr.com
Internet address for distribution via SCSI reflector: scsi@wichitaks.ncr.com

SCSI Bulletin Board 719-574-0424

Document Distribution
Global Engineering
or
15 Inverness Way East
Englewood, CO 80112-5704
Telephone: 303-792-2181
800-854-7179
Facsimile: 303-792-2192

ABSTRACT

This document describes SCSI configurations that may be achieved within the context of the specifications in the SCSI-2, SCSI-3 SPI and SCSI-3 FAST 20 standards. These configurations have one or more configuration parameters expanded beyond that formally specified in the standard documents and may require special components or an interpretation of the underlying technical reasons for parameters specified in the standards. This technical report describes the considerations that can lead to effective implementations of special components but does not describe the detailed design of any component. The information in this document does supersede any requirements in the referenced standards for formal compliance with standards.

PATENT STATEMENT

CAUTION: The developers of this technical report have requested that holder's of patents that may be required for the implementation of the configurations described in this document, disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this document.

As of the date of publication of this document and following calls for the identification of patents that may be required for the implementation of the configurations described in this document, no other such claims have been made. No further patent search is conducted by the developer or the publisher in respect to any standard or document it processes. No representation is made or implied that licenses are not required to avoid infringement in the use of this document.

Proposed document organization

1. Scope

2. References

3. Definitions, symbols, and abbreviations

4. General

This document describes configurations and extensions that are possible within the context of the existing SCSI standards but that may not be obvious or formally allowed by the standards. It is the intent of this document to provide technical information and guidance to enable these expanded capabilities in an effective way.

The common theme on all of these enhancements is expansion of applications that can be addressed without changing the features of existing SCSI device implementations that comply with the existing standards. Of special interest are interconnecting devices capable of different widths, extending the physical length of the SCSI domain, increasing the dynamic reconfigurability of domains, and extending the number of addressable devices in a single domain.

5. Bus segment concept

Existing SCSI standards define parameters for SCSI busses based on the assumption that there is a single electrically conducting path between bus terminators for each signal and that a SCSI domain contains all the devices between these two terminators. This electrical path is assumed to pass signals in both directions without delay other than that caused by the propagation delay of the transmission line associated with the path. It is assumed that there are no intervening active components in the path between the bus terminators.

A more general concept recognizes that it is possible to build SCSI domains that use more complex physical implementations where there may be active electrical components between SCSI devices. A building block for these more complex implementations is the bus segment which is defined as two bus terminators and the associated single electrically conducting path between these terminators (for each signal) that satisfies the assumptions in the first paragraph of this section. Multiple bus segments may be functionally connected together by special coupling circuits described in section 7.

Each bus segment has its own TERMPOWER sources and TERMPOWER distribution parameters.

Bus segments must use the same transmission type (differential, LVD SCSI, or single ended) within the segment. A domain may contain segments that use different transmission types.

Bus segments largely follow the same rules individually that are described in the existing standards (with important exceptions). Using multiple bus segments with coupling circuits in the same domain allows much more of the full properties supported by the SCSI bus protocol to be realized than when using single segment domains.

Some of the salient properties impacted are device count limits, physical length limits, ground voltage shifts, dynamic removal and replacement of portions of domains, and mixing of device types (single ended and differential).

5.1. Segment length parameters

The existing SCSI-2 and SCSI-3 SPI standards mention physical lengths as recommended or required maximums for certain conditions. The SCSI-3 Fast 20 standard incorporates information that allows implementers significant leeway in the maximum physical lengths through the use of informative material provided in annexes. The reasons behind the length numbers contained in the existing standards are not always stated in the standards and some of the technology that existed some years ago has improved significantly since these standards were stabilized. Both of these points invite re-examination of length related parameters based on the latest knowledge and technology.

The concept of a bus segment is introduced for the first time in this document. Using multiple bus segments in the same SCSI domain directly challenges the notion of using physical length alone as a domain configuration limit. In order to understand and give guidance concerning appropriate physical length of bus segments and domains one must look at the properties that affect this length.

Bus segments and SCSI domains have physical length limits that are determined by (1) the time required for signal propagation or (2) by loss of signal quality due to resistance, dispersion, delay skew, transmission line reflections, jitter or other factors. As long as the physical plant delivers the signals to SCSI device connectors and bus terminators with the proper timing and signal quality there is no technical reason to further restrict the physical length. These criteria must be applied to the lines that first fall outside of acceptable bounds as the length is increased. In some cases the TERMPWR lines limit the length. See section 12.

For the domain length limits one must consider the propagation time consumed by the elements that connect the bus segments as well as that consumed in the individual segments. The length limits for individual segments may be independently considered as long as the domain length (protocol timing) limits are not exceeded. A single domain may contain many segments since some of the most attractive versions of SCSI have very short segment length limits (caused by signal quality) that consume only a small part of the domain timing budget.

Section 7 considers the propagation delay issues for the elements that connect bus segments. The remainder of this section considers the factors that affect the physical length of single segment domains and segments within multisegment domains.

5.1.1. Protocol timing limits

The SCSI arbitration and phase boundary protocols define a time of 400 ns for the bus to "settle" after certain conditions described in the SCSI-3 SIP document are possible. The intent of this "bus settle delay" is to allow a full round trip time for electrical disturbances to decay before proceeding to the next steps in the protocol. This means that the overall time required for a signal to travel from one end of the bus to the other and to "reflect back to the beginning"

is a maximum of 400 ns. If there is no reflection a full 400 ns is available for a one way transit time. In most cases, however reflections will be present and a one way transit time of 200 ns must be used.

If one has a uniform media and a known signal propagation velocity, Vp, this establishes the maximum physical length as Lmax = 200ns/Vp. With the presently allowed slowest Vp of approximately 5ns/m this gives a maximum length of 40 meters. One cannot use this number as a general limit because the transmission media is not uniform and the signal are delayed by the non-uniformities. The next section considers the most important additional effects cause by device loading.

5.1.1.1. Device loading parameters

When real devices are present on a segment their stubs and capacitive loads add to the bus delay for the signals. The amount of this delay depends on the details of the loading. In the normal worst case with a 0.1 meter stub and a 25 pF capacitive load this delay has been reported to be approximately 0.5 ns per device.

In general the allowed transit time must be reduced by the delay caused by the loads. The device delay is denoted Tdi for the "i"th device and the total delay is the sum of the delays from all devices on the domain (including any segment connecting elements) and is denoted Tdd.

This gives a general formula for the maximum domain length as:

$$Lmax = (200ns - Tdd) / Vp$$

With 14 devices having a total of 0.5 ns per device there is a 7 ns reduction. One should also allow some delay for connectors and other disturbances. A reasonable number for this may be 2 ns for a domain having several connectors. Tdd therefore becomes approximately 9 ns for a fully loaded wide SCSI segment

5.1.1.2. Caution when using differential interfaces

It is important to note that the times relating to the signal propagation are defined as measured from the device connectors. This is especially important when using differential interfaces with separate transceivers since the propagation time through the transceivers can be several tens of nanoseconds. The SCSI standard does not allow any special budget for these extra transceiver delays so designers of differential interfaces must ensure that the timings of the transceivers are considered.

5.1.2. Propagation - time - limited domain length

Using the numbers in section 5.1.1.1 and assuming a Vp or 5 ns/m we find a maximum single segment length (with reflections) to be 191/5 = 38.2 meters. The number in the present standards use a much more conservative Vp of 6.6 ns/m (28.93m) and adds extra margin to report a maximum length of 25 meters. Using the real numbers without the extra margin we gain an extra 13.2 meters in total segment length. For media specially built for high propagation velocity one can achieve at least 3.76ns/m for a loaded length of 50.8m or double

that in the present standards without changing any protocol timings (if the signal quality remains adequate).

Lengths above 6 meters are suggested as mainly applicable to differential transmissions in the standards. We will examine this assumption in more detail in section 5.1.3.

It is not necessary to abandon the length based scheme to achieve the benefits of multisegment domains but it is necessary to adopt some relationship between length and time to allow for the coupling circuits between domains. This relationship can be provided by the equation in section 5.1.1.1. One may think of the length based scheme as a significantly conservative specification method.

5.1.3. Signal quality limits

This section explores the effects of signal quality on segment lengths.

5.1.4. Wire/Printed circuit board parameters

5.2. Other segment properties

6. Mixed width operation

SCSI is specified to operate with any of three data path widths: 8 bit, 16 bit, and 32 bit. This document does not consider the 32 bit operation. Conditions exist where it is desirable to implement both 8 and 16 bit SCSI devices within the same bus segment. When this happens the bus segment is said to be using mixed width operation.

Since the timing requirements become more strict at higher data rates the risk of using mixed width operation increases at higher data rates. The details of these risks are explored in this section.

6.1. Architectural Options

This section describes four different ways to configure mixed width bus segments. Each requires separate consideration and each must adhere to the wiring tables shown in 6.2.2.

6.1.1. 16 bit main path

The simplest form of mixed width segment has a 16 bit main path to which either 8 bit devices, 16 bit devices, or both are attached. Only the 16 bit main path has bus termination. No 8 bit devices may provide any bus termination.

Figure 1 shows the architecture of this implementation. Every 8 bit connection only contacts the SCSI lines that are used for 8 bit devices. The upper 9 bits (data and parity) are not contacted by the 8 bit device. This condition produces more electrical load on the lower 9 bits (data and parity) than on the upper 9 bits.

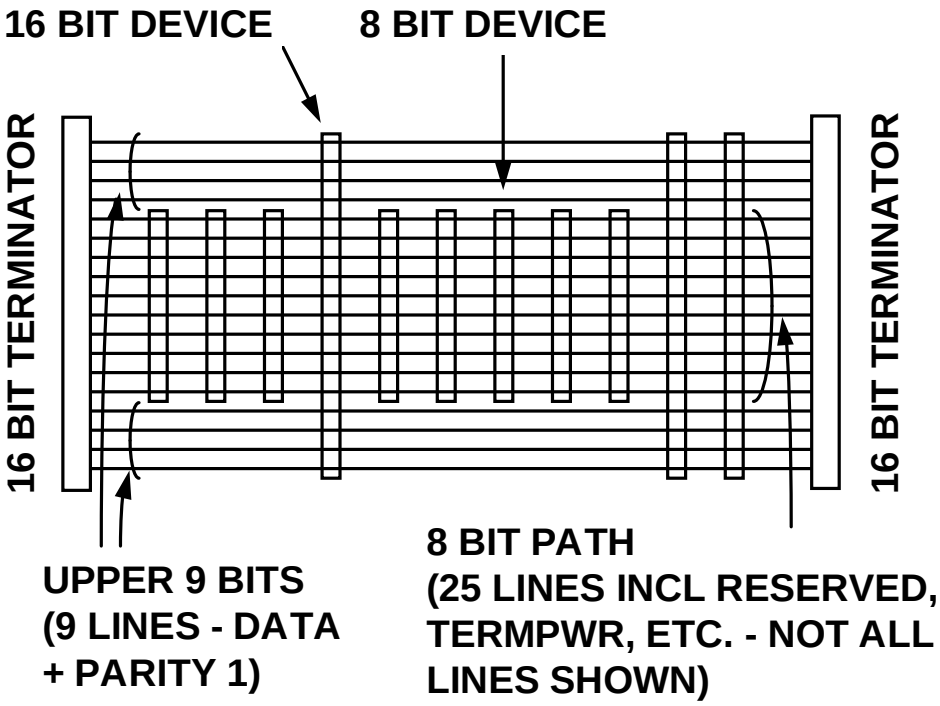

Figure 1 - Mixed width connections using 16 bit main path

This non-uniform loading produces timing skew in direct proportion to the number of 8 bit devices used and the electrical load they present. The worst case is when two 16 bit devices are near the opposite ends of the bus and there are 8 8-bit devices in between. One could have 14 8 bit devices in between but this case could not work since there are only 8 address bits available on 8 bit devices and only 8 devices can be addressed. In this case there is 8 times the individual device loading skew on the data lines. Measurements have shown that a typical delay induced by a single device is 0.5 ns. The one way timing skew in this worst case situation is therefore approximately 4 ns. This is larger than the timing skew produced by the cable media itself and could be responsible for data errors in the upper 9 bits if not treated.

One way to lessen the impact of the non-uniform loading is to add electrical load to the upper 9 bits by the connection scheme used between the 8 bit devices and the 16 bit path.

If the 16 bit path uses the common 0.025" centerline flat ribbon cable there will be a connector conversion necessary to attach the 8 bit device. The connection to the flat ribbon cable will require a high density 68 pin connector while the 8 bit device will require a 50 pin low density connector. It is convenient to add discrete capacitors within the connector converter device between each of the upper 9 bits and a ground line. The value of the capacitors should be approximately the same as the pin capacitance of the 8 bit device. Of course not all 8 bit devices have the same pin capacitance but a reasonable approximation is from 10 to 20 pF.

This compensation scheme can significantly reduce the skew to levels that are generally negligible.

If one is using round cable (either shielded or unshielded) for the 16 bit path it is possible to use the 50 pin low density connector directly to the 16 bit path. In this case one should be careful not to break the conductors while attaching this connector in order to maintain equal path lengths for the lines. There is no readily convenient access to the upper 9 bits in this case so the non-uniform loading will still be present.

A better way is to use the 68 pin high density connectors for the round cable and use the scheme described above within the connector converter device.

One also needs to consider the TERMPOWER distribution for this case. The 16 bit path has 4 conductors available for the TERMPOWER. These 4 lines should be connected together by the 8/16 bit connector converter and passed to the TERMPOWER lines on the 8 bit side. There is an issue however on the 8 bit side because the single ended and differential TERMPOWER requirements are different. The single ended uses only one TERMPOWER line while the differential uses two lines.

The risk is that some older single ended devices may ground one of the pins used for differential TERMPOWER. If one of these devices is attached it will connect the TERMPOWER line to ground and disable the entire bus. The simple solution is to connect only the single ended TERMPOWER line from the 8 bit side to the 4 lines on the 16 bit side. This connector converter will therefore work for both single ended and differential applications.

For the differential case any TERMPOWER sources on the 8 bit devices will only use one TERMPOWER line through the connector converter. Since the 8 bit devices are not allowed to have bus termination in this case one does not need to consider this condition.

6.1.2. 8 bit main path

6.1.3. Single 16 bit path and single 8 bit path

6.1.4. Multiple 8 or 16 bit paths

6.2. References and tables

6.2.1. SFF-8017 reference

6.2.2. Wiring tables

7. Bus expanders

Bus expanders are elements used for connecting segments together. They do not occupy a SCSI ID and are intended to be "invisible" to the protocols.

The following features are considered to be desirable properties of bus expanders:

- Some propagation delay budget is consumed
- No SCSI ID's, no arbitrations or messages can be sent
- Retransmitted signal timing skew (both delay and hi/lo) are no worse than from a valid SCSI device
- Does not interfere with the REQ/ACK offset count
- Min/Max pulse widths are maintained
- A reset filter is required
- Placement of the initiator and targets with respect to the isolator in arbitrary
- TERMPOWER is not connected between the segments being coupled

7.1. Homogeneous type

7.2. Heterogeneous types

8. Logical bus length parameters

8.1. Effects of Wired-Or glitches

8.2. Expander propagation delay effects

9. Topologies and configuration rules with simple expanders (static)

9.1. Series combinations

9.2. Stubbing combinations

9.3. Mixed

10. Dynamic reconfigurations

10.1. Addition and removal of devices

10.1.1. Framework of entire process

10.1.2. Electrical considerations

10.1.3. Logic input states during transitions

10.1.4. Protocol implications

10.2. Addition and removal of bus segments

11. Topologies and configuration rules for dynamically reconfigurable systems

11.1. Initiator placement

11.2. Redundant paths

11.3. Initialization

11.4. Power considerations

12. Terminator power distribution

12.1. Wire effects

12.2. Mixed power configurations

12.2.1. High voltage / Low voltage

12.2.2. Power on / Power off segments