X3T9.2/93-174 Rev 0

To: John Lohmeyer
Chairman, ANSI X3T9.2 (SCSI)

From: Doug Hagerman
Bill Ham
Digital Equipment Corporation
SHR 3-2/W28
334 South Street
Shrewsbury, MA 01545

Phone: 508-841-2145
FAX: 508-841-6100
Mail: hagerman@starch.enet.dec.com

Date: 5 November 1993

Subject: Ultra-SCSI Proposal Presentation Slides

These are the slides from the X3T9.2 presentation in which we propose a faster timing option for SCSI.

Ultra-SCSI Proposal Overview

Doug Hagerman
Bill Ham
Digital Equipment Corporation
9 November 1993

What Is Ultra-SCSI?

- a proposal for an informative annex to SPI-1
- fully backwards compatible with current parallel SCSI
- uses exactly the same physical environment:
    - . cables
    - . connectors
    - . terminators
- the same mindset as "fast" was over "synchronous"
- useful in two situations:
    - . differential cable busses
    - . single ended backplane busses
- faster timings give 40 & 80 MBytes/sec on wide bus

Ultra-SCSI enables the management of the technical and financial risk of further incremental performance enhancements to storage subsystems.

Typical Application - External Bus

External subsystem interconnect cable using parallel, wide, differential physical layer:

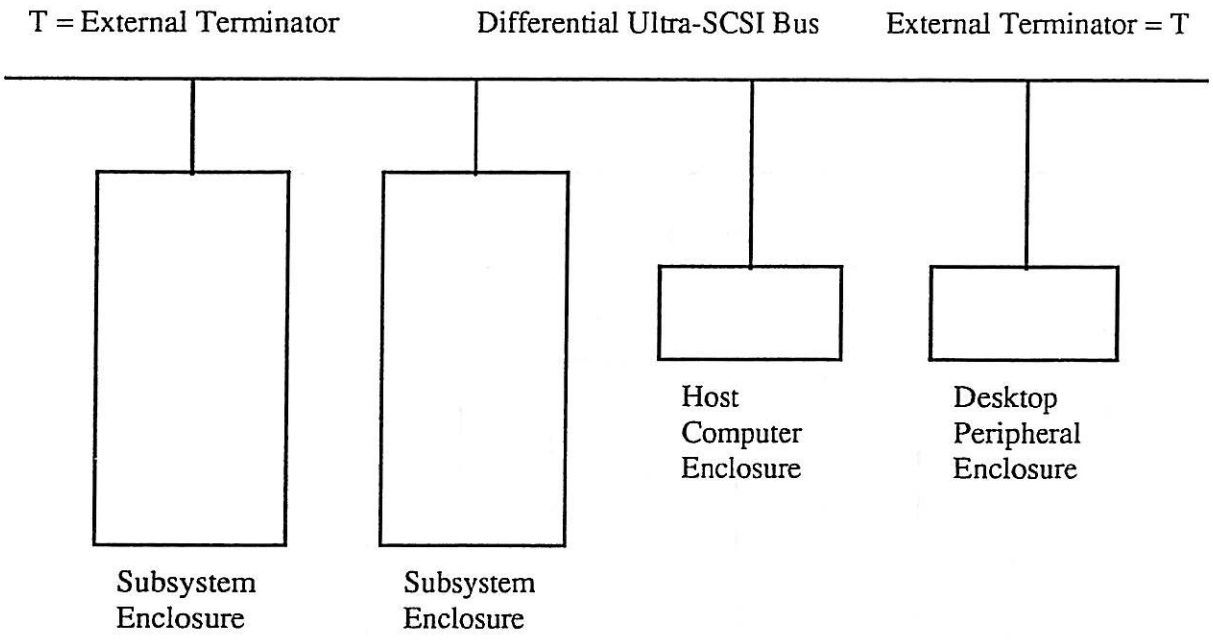

This sketch shows an environment with a limited number of enclosures, a reasonable bus length, terminators external to every enclosure, and a fairly static configuration.
A configuration like this could support fairly large data servers.

Typical Application - Internal Bus

Internal backplane interconnect using parallel wide single-ended physical layer.

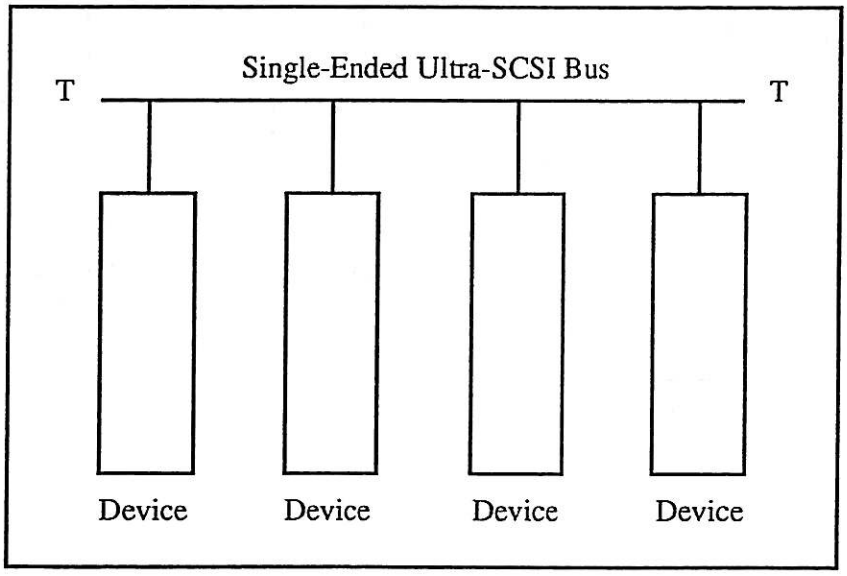

This sketch shows an environment with a fixed number of device slots, a very short bus, terminators external to every device, and a well-controlled electrical configuration. This would be useful in storage subsystem enclosures.

Features

- works with other new stuff to give a more capable SCSI interconnect:
    - . smaller connectors
    - . incremental SCSI tuning efforts
    - . warm swap

- offers minimal market perturbation because it's incremental:
    - . compatible with existing testers, code, devices, chips
    - . recognizeable as "SCSI"

Is This Really Useful?

- real SCSI implementations don't use full bandwidth
    - . message and command timings are asynch (slow)
    - . chip interrupt service time
    - . think time in microcode
    - . wasted time on bus (e.g. could disconnect sooner)
    - . less than optimal req-ack timing

Example of Real Bus Overheads

- experimental bus phase times on "fast" bus (10 MByte/second) are as follows:
  - . this is real lab data
  - . numbers selected from several devices
  - . one sector read (512 Bytes)

Bus Phase	Time (microseconds)
Arbitration	2.60
Selection	2.60
Message Out	6.69
Command	13.03
Message In	14.37
Bus Free	
2nd Arbitration	2.35
Reselection	2.19
Message In	23.90
Data In ->Req0	0.60
Req0 -> Ack0	4.13
Ack0 -> Ack511	51.16 (512B Data transfer time)
Ack511->end Data In	4.87
Status	2.20
Message In	5.70
Total Bus Busy time	136.39 (Total)
SCSI overhead	85.23 (Total less transfer time)
Resulting effective speed	3.75 MBytes/second

Extrapolating to faster data phase gives:

Read Transfer Profile	Projected Resulting Speed (MBytes/second)
512B transfers at 10 MB/sec	3.75
4kB transfers at 10 MB/s	8.28
4kB transfers at 20 MB/s	14.12
4kB transfers at 40 MB/s	21.83
4kB transfers at 80 MB/s	30.02

Timing Proposal

Synchronous Data Transfer Request (SDTR) Message

Faster timings must be signalled to the chip in some fashion. The current method for Fast SCSI uses the Transfer Period specified in the SDTR message to control the timing selection. Ultra-SCSI simply expands this method as follows.

Transfer type	async	sync	fast	ultra-1	ultra-2
Transfer period factor	0	>50	>25	>12	>6
Transfer period (nsec)	async	>200	>100	>50	>25
Megatransfers/second	0-5	0-5	5-10	10-20	20-40
MB/sec on wide bus	0-10	0-10	10-20	20-40	40-80

Timing Parameters *** INITIAL PROPOSAL ***

The following timing parameters must be changed to support the two levels of Ultra-SCSI. This version of the proposal simply has the current fast values divided by two or four for ultra-1 and ultra-2 timing, but this may need to be changed based on chip implementation issues.

Timing Parameter	sync	fast	ultra-1	ultra-2
Transfer period (nsec)	200	100	50	25
Transmit assertion period	80	30	15	7
Transmit negation period	80	30	15	7
Transmit hold time	53	33	16	8
Transmit setup time	23	23	11	5
Receive assertion period	70	22	11	5
Receive negation period	70	22	11	5
Receive hold time	25	25	12	6
Receive setup time	15	15	7	3

SOME PROTOCOL ISSUES

- ALL ARBITRATIONS, SELECTIONS, AND INITIAL NEGOTIATIONS FOR SPEED, WIDTH, OFFSET, ETC ARE EXACTLY AS WITH TODAY'S PARALLEL SCSI
- AS SOON AS TWO DEVICES AGREE TO TALK "ULTRA SCSI" ALL TRANSMISSION AFTER ARBITRATION AND SELECTION PROCEED ACCORDING TO THE ULTRA SCSI PROTOCOL
- SYNCHRONOUS BURST SPEEDS OF AT LEAST 20 MEGATRANSFERS/SEC AND POSSIBLY AS MUCH AS 40 MEGATRANSFERS/SEC ARE ALLOWED
  - THIS ALLOWS 16 BIT BANDWIDTHS OF 40 AND 80 MEGABYTES/SEC (ACTUALLY 9BIT "PBYTES")
  - 80 MEGAPBYTES/SEC IS 720 MEGABITS/SEC

PARALLEL BUS DATA INTEGRITY TESTER

- A LABORATORY SET OF EQUIPMENT THAT CAN DIRECTLY MEASURE THE NATURE AND FREQUENCY OF OCCURRENCE OF ERRORS ON A BUS
- CAPABLE OF OPERATING WITH PARAMETERS EXCEEDING THAT NORMALLY USED IN SERVICE:
  - SPEED
  - DATA PATTERNS
  - TIMING MARGINS
  - VOLTAGE OFFSETS
  - DATA PATH WIDTH
- PROGRAMMABLE, FLEXIBLE, BUS MEDIA INDEPENDENT
- MARGIN TO FAILURE QUANTITATIVELY MEASUREABLE
- INDEPENDENT FROM COMPLEX BUS PROTOCOLS -- DATA PHASE FOCUS

PARALLEL BUS DATA INTEGRITY TESTER HARDWARE

- PROGRAMMABLE EXCITER / COMPARATOR / COUNTER (UP TO 20 MEGATRANSFERS/SEC-- 40 SOON))
- LOAD / TRANSCEIVER BOARDS (SPECIFIC TO BUS MEDIA)
  - DIFFERENTIAL SCSI (UP TO 27 BITS WIDE)
  - SINGLE ENDED SCSI (UP TO 27 BITS WIDE)
- COMPUTERS FOR SOFTWARE DEVELOPMENT, DOWNLOADING, AND DATA MANAGEMENT
- TIMING MODIFICATION EQUIPMENT (FOR MEASURING MARGIN)
- TESTING SOFTWARE
  - ARCHITECTURE
  - DATA PATTERN CREATION
  - ERROR COMPARATORS / DIAGNOSIS
  - AUTOMATIC TIMING MARGIN TESTING
- DATA OUTPUTS
  - WAVEFORMS (MICROVAX BASED)
  - ERROR RATE PLOTS (PC AND MICROVAX)
- BUS MEDIA

PARALLEL BUS DATA INTEGRITY TESTER

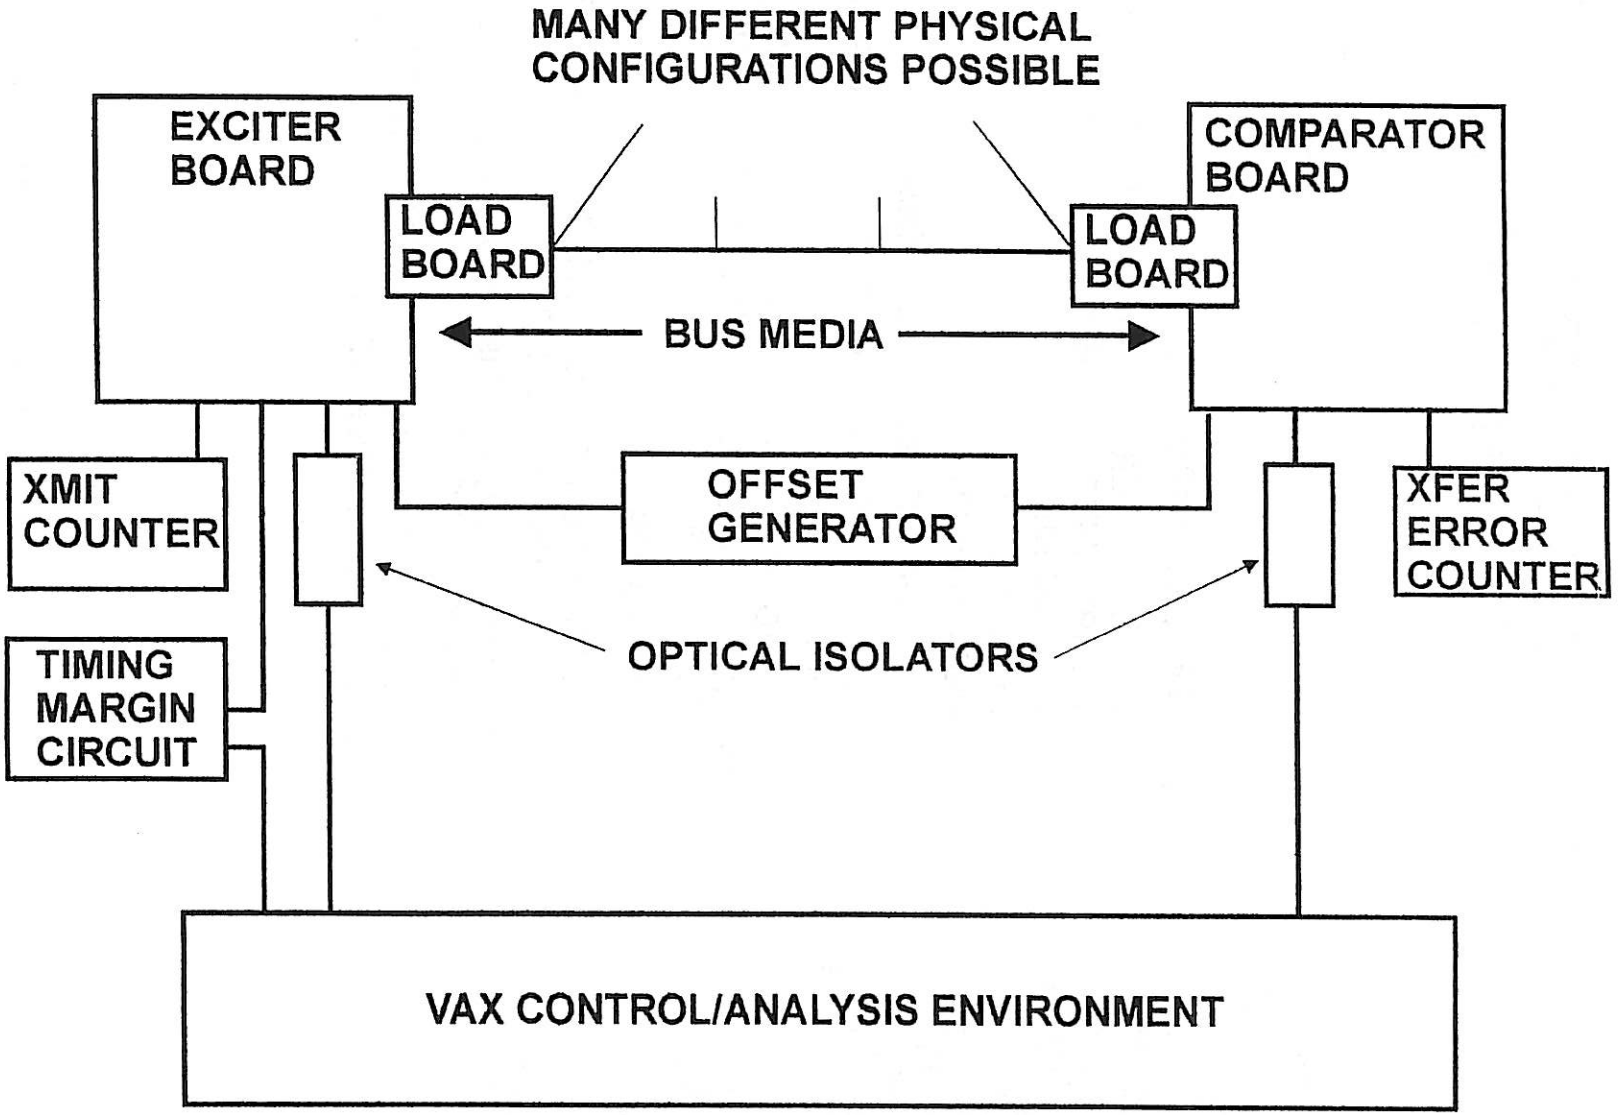

PARALLEL BUS DATA INTEGRITY TESTER HARDWARE

- **TIMING MARGIN MEASUREMENTS**
  - DATA IS TRANSFERRED FOR A SPECIFIED PERIOD WITH A SPECIFIED PATTERN AND PHYSICAL CONFIGURATION
  - CLOCK IS SET AT A KNOWN POSITION WITH RESPECT TO DATA EDGES
  - DATA ERRORS (IF ANY) ARE MEASURED
  - PROCESS IS REPEATED FOR DIFFERENT CLOCK POSITIONS
  - CLOCK POSITIONS ARE VARIED UNTIL 100% GOOD AND 100% BAD DATA ARE DETECTED IN THE SAMPLED TIME
- **ERROR ANALYSIS**
  - DETECTED ERRORS ARE COMPARED TO THE KNOWN TRANSMITTED DATA
  - USUALLY A SINGLE BIT IS RESPONSIBLE FOR ALL THE INITIAL ERRORS
  - WAVEFORMS MAY BE CAPTURED UNDER THE VARIOUS ERROR CONDITIONS

TIMING MARGIN TESTING

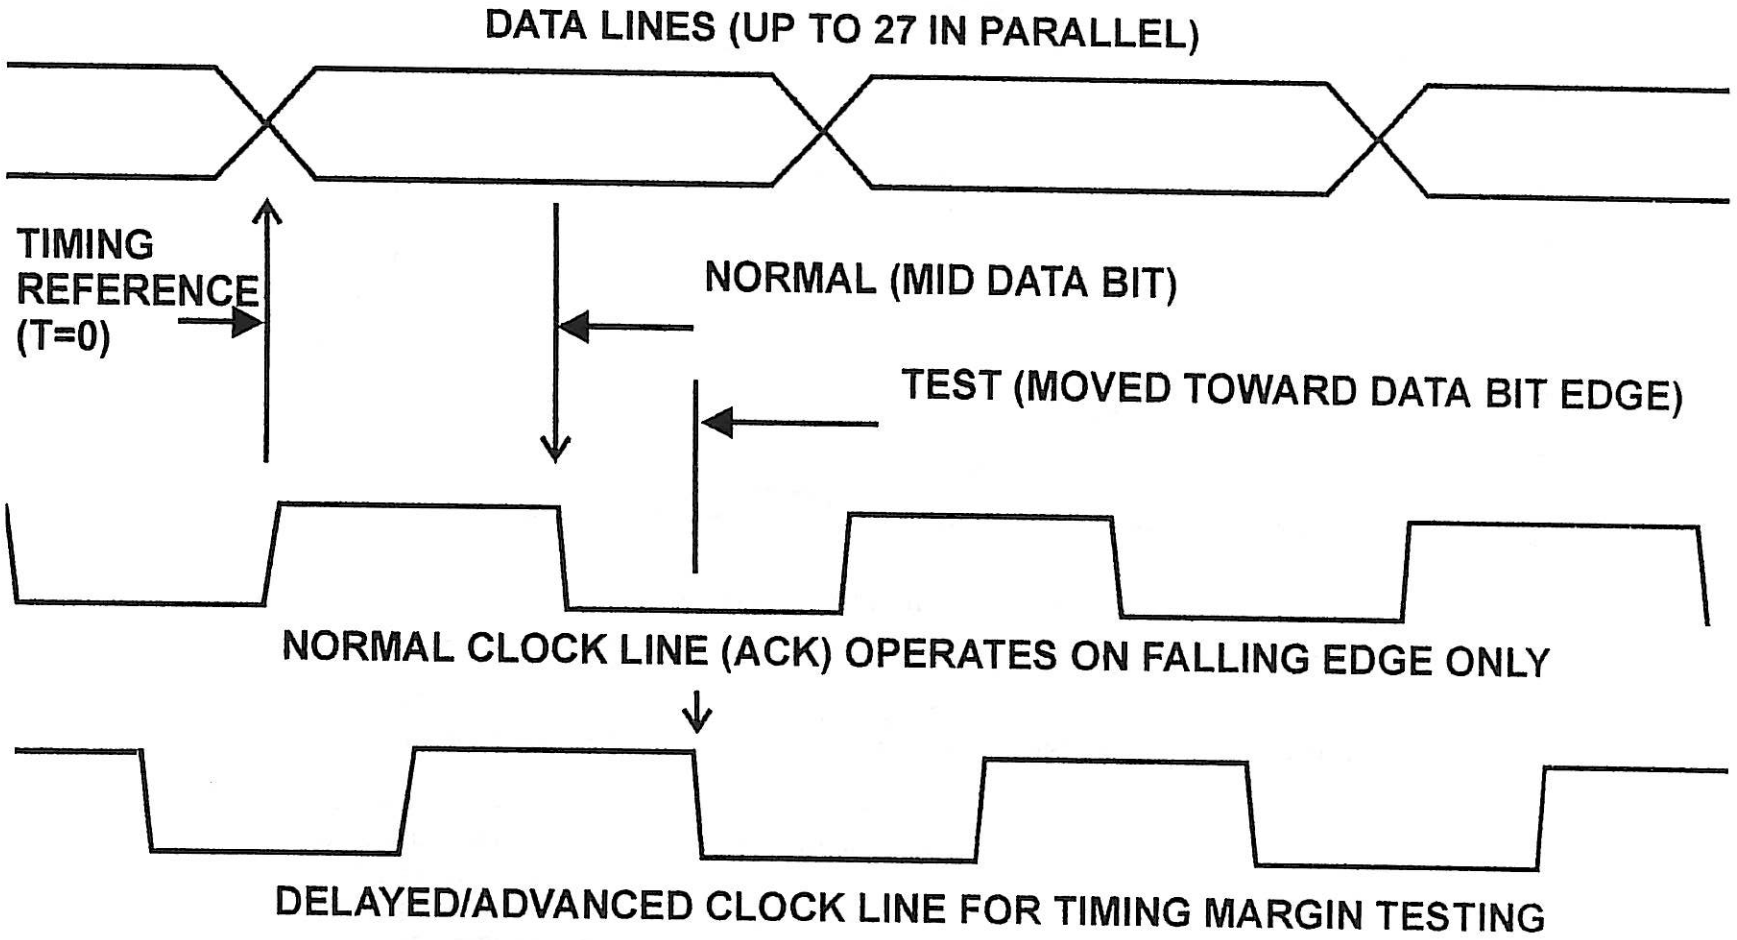

SYSTEM SETUP

Time per test cycle	300 Secs
Data Rate	20 Megatransfers/sec
Total Transfers	6 x10^9
Data Pattern	16K Random (repeating)
Transceivers	DS36954V
Cable Configuration	Exciter=T======== 20M ========T=Comparator
Bus Type	P-Cable Differential

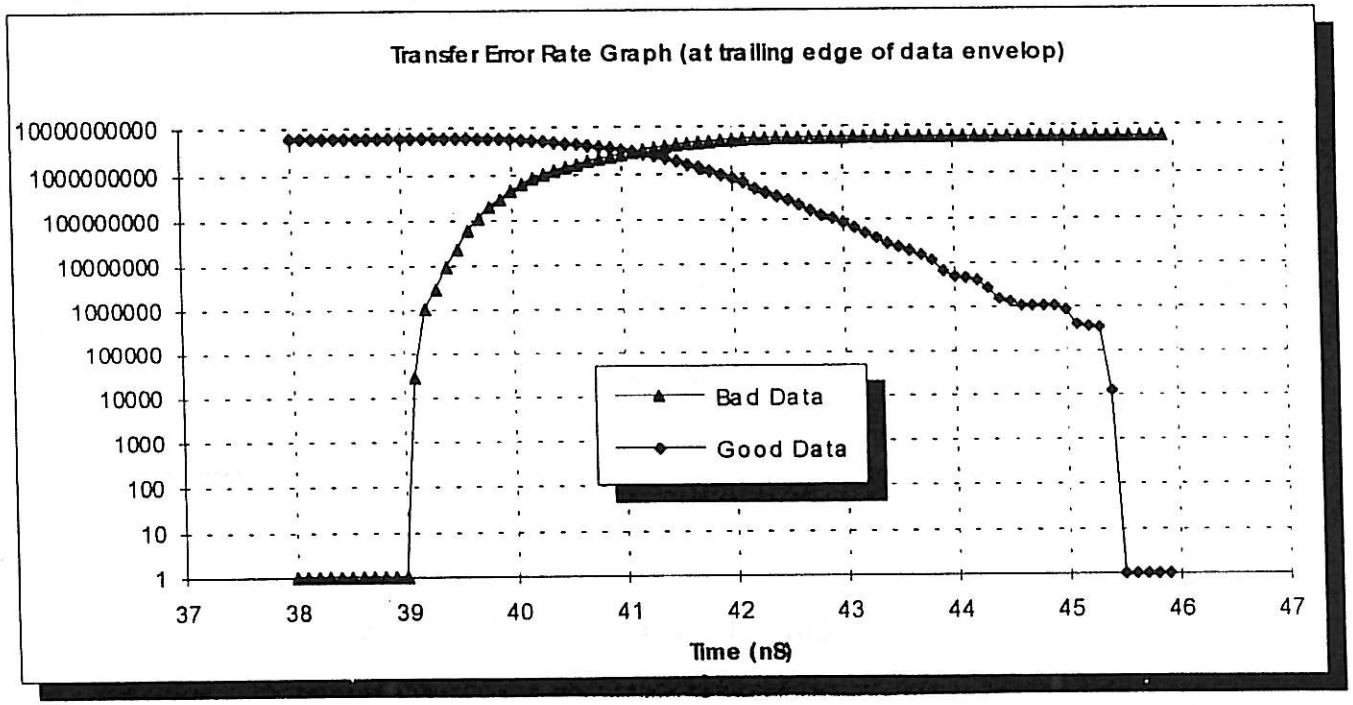

F.W.Lee 19-Oct-1998

172

SYSTEM SETUP

Time per test cycle	300 Secs
Data Rate	20 Megatransfers/sec
Total Transfers	6×10^9
Data Pattern	16K Random (repeating)
Transceivers	DS36954V
Cable Configuration	Exciter=T======== 20M ========T=Comparator
Bus Type	P-Cable Differential

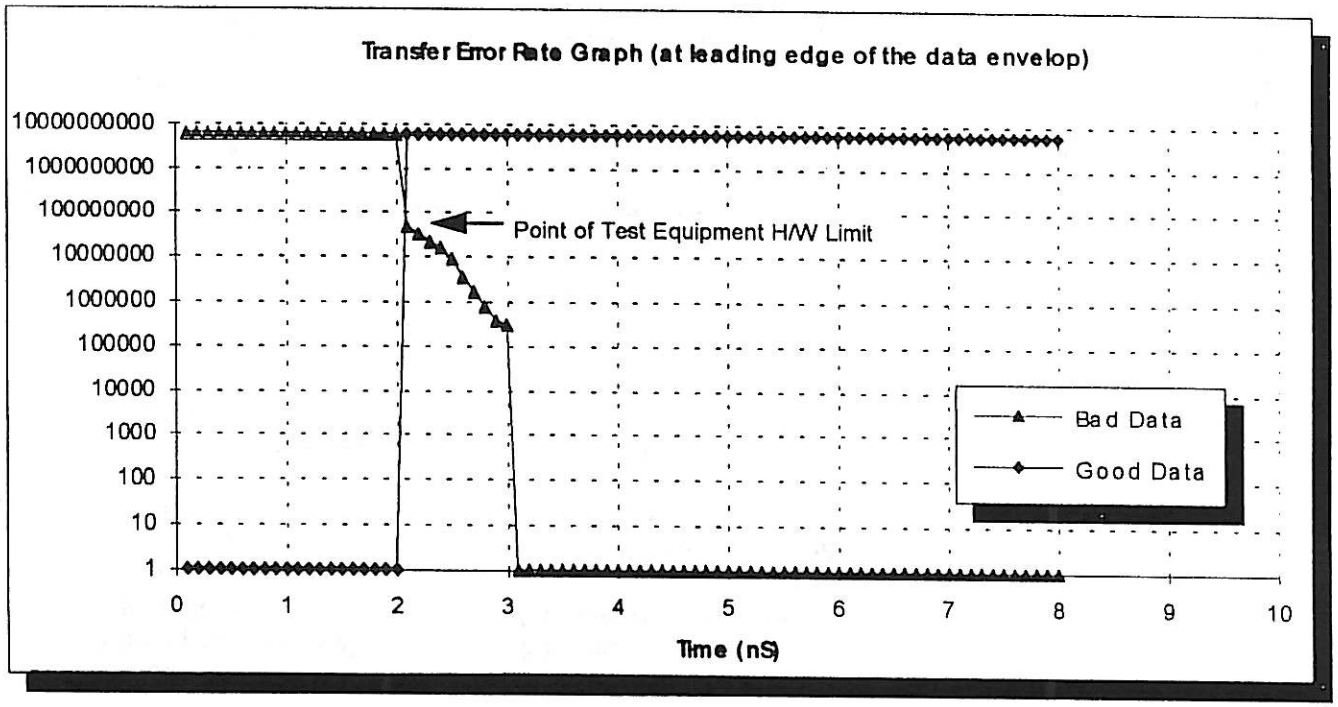

F.W.Lee 19-Oct-1993

SYSTEM SETUP

Time per test cycle	300 Secs
Data Rate	20 Megatransfers/sec
Total Transfers	6 x10^9
Data Pattern	16K Random (repeating)
Transceivers	DS36954V
Cable Configuration	Exciter=T======== 20M ========T=Comparator
Bus Type	P-Cable Differential

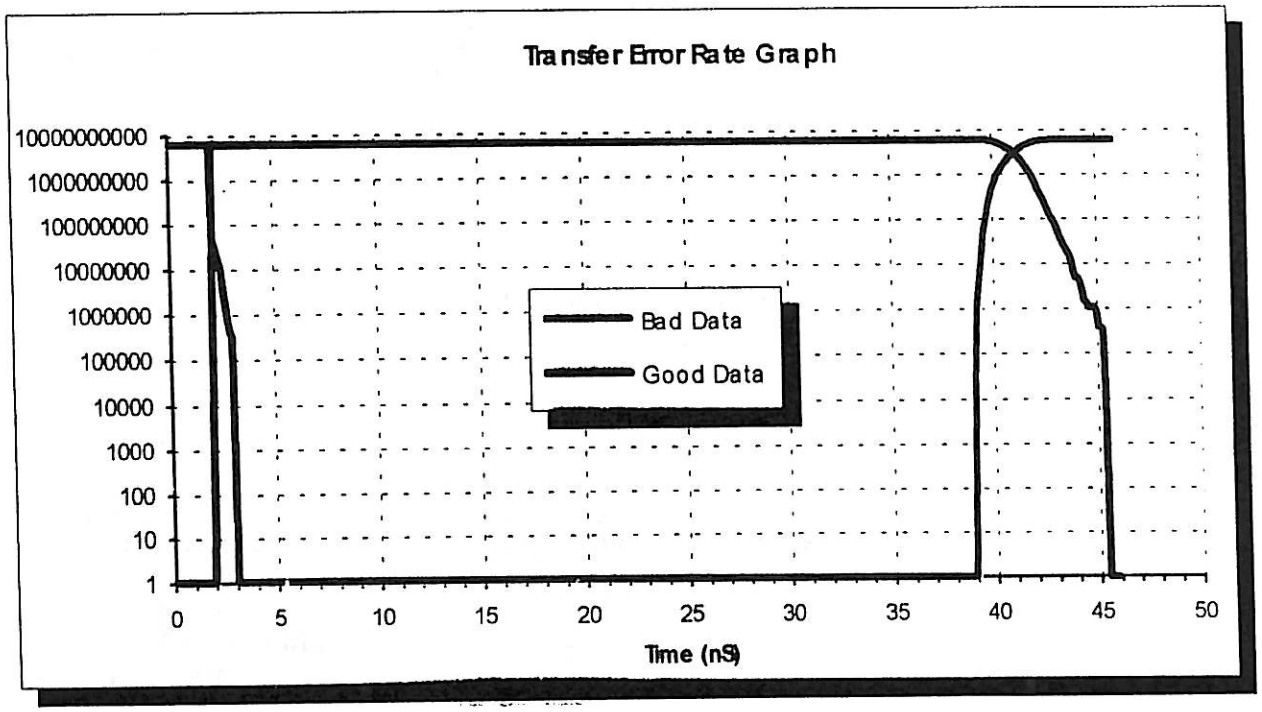

F.W.Lee 19-Oct-1993

174

20 MEGATRANSFERS/SEC
2X NORMAL OPERATION

DRIVER END

DIFFERENTIAL SIGNAL

INDIVIDUAL PLUS AND MINUS SIGNALS

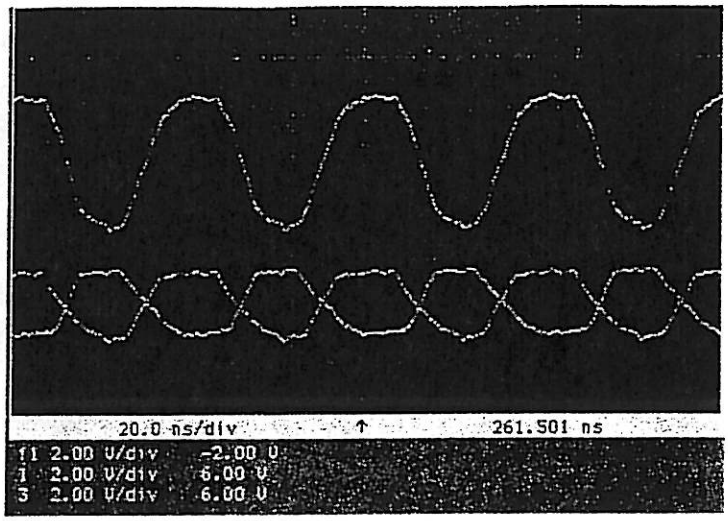

RECEIVER END
(20 METERS FROM TRANSMITTER)

DIFFERENTIAL SIGNAL

INDIVIDUAL PLUS AND MINUS SIGNALS

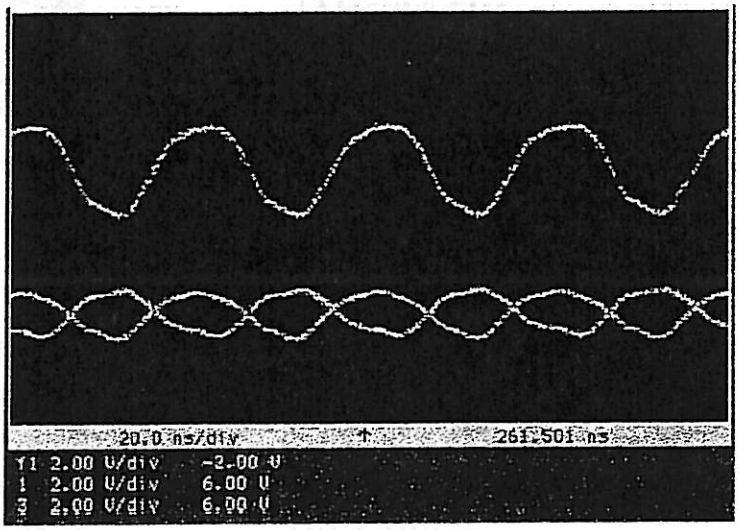

STANDARD 30 GAUGE SCSI P CABLE

30 MEGATRANSFERS/SEC
(3X NORMAL OPERATION)

DRIVER END

DIFFERENTIAL SIGNAL

INDIVIDUAL PLUS AND MINUS SIGNALS

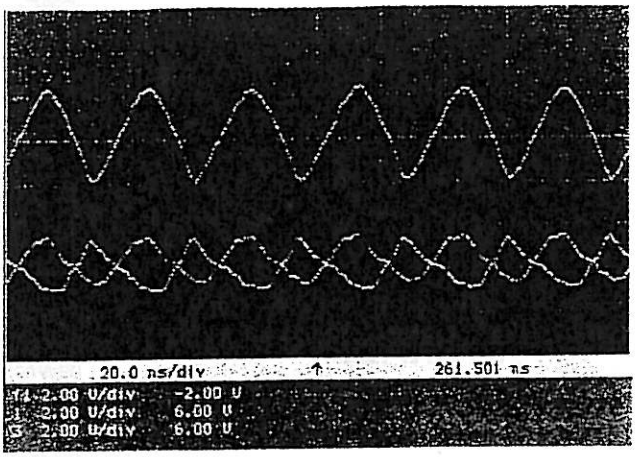

RECEIVER END (20 METERS)

DIFFERENTIAL SIGNAL

INDIVIDUAL PLUS AND MINUS SIGNALS

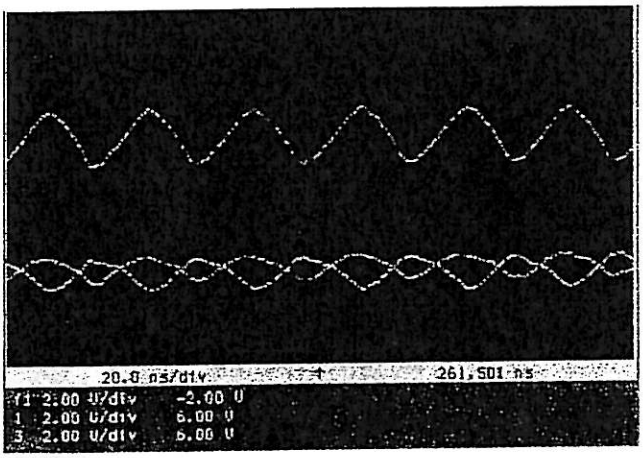

RECEIVER END (40 METERS)

DIFFERENTIAL SIGNAL

INDIVIDUAL PLUS AND MINUS SIGNALS

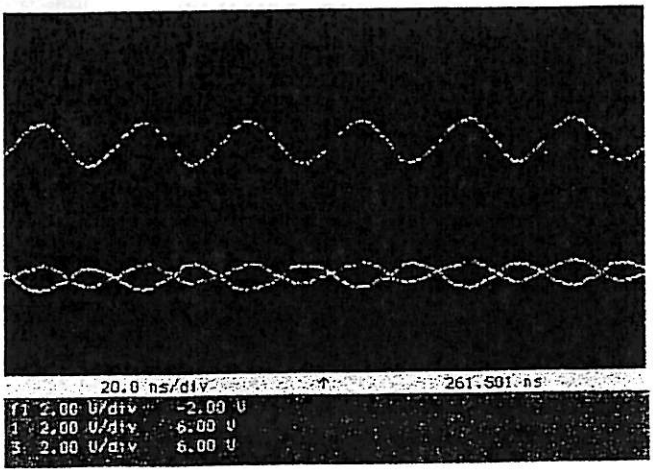

SILICON ISSUES

- THE MAJOR ISSUE WITH ULTRA SCSI IS CHIP AVAILABILITY
- THE PRESENT SCSI SPECS WERE DEVELOPED WITH SILICON TECHNOLOGY AT LEAST TWO GENERATIONS BACK
- SEVERAL CHIP VENDORS HAVE EXPRESSED SIGNIFICANT INTEREST AND APPARENTLY HAVE THE TECHNOLOGY TO OFFER MUCH HIGHER SPEED OPERATION WITH THE LATEST CHIP TECHNOLOGY
- SERIAL CHIPS ARE BEING PRODUCED THAT OPERATE AT CLOCK RATES SEVERAL TIMES FASTER THAN ULTRA SCSI REQUIRES
- TIMING MARGINS NEED CAREFUL EXAMINATION FOR ULTRA SCSI