

PRELIMINARY ANSI WORKING DOCUMENT

DOCUMENT NO. X3T9.3/87-005

REVISION 1.3 5/18/87

ENHANCED SMALL DEVICE INTERFACE

This document is a draft proposed American National Standard. As such, it is subject to change before approval and publication as an American National Standard. This document is not intended to be a device design or procurement specification.

For current information on the status of this document, contact the individuals shown below:

Gary S. Robinson
Chairman X3T9.3
MS MLO 12-3/E51
Digital Equipment Corp.
146 Main Street
Maynard MA 01754

617-493-4094

All rights reserved

I. Dal Allan
Vice Chairman X3T9.3
ENDL Consulting
14426 Black Walnut Court
Saratoga CA 95070

408-867-6630

X3T9.2/87-105
Rev 1.3

103

Expository Remarks for Enhanced Small Device Interface

The initial products introduced as small form factor 5 1/4" magnetic disk drives were relatively low in capacity (5-10 MB), had a limited data rate (5 Mbs), and slow access times (85-100 msec). The market response to these products generated significant demand for their use in personal computers and single user workstations.

The smaller size and low power requirements of small form factor drives led to the development of higher capacity (100 MB+), faster data rate (10 Mbs+) and fast access time (sub-20 msec+). To support these kinds of products it was felt that a new interface which was electrically compatible with that used on the lower end products would provide a simpler upgrade path than any other alternative.

An ad hoc industry group of controller and device manufacturers agreed to pursue this direction as a common effort and the first working document was introduced in May, 1983. This initial effort defined the Enhanced Small Disk Interface but after the definition of the Enhanced Small Tape Interface it was decided to merge the two into a single document in October, 1983 as the Enhanced Small Device Interface. In January, 1985 it was proposed that a version suitable for optical disk be incorporated and the first working draft was issued in March, 1985. NOTE: In May, 1987 It was agreed that tape did not have broad industry acceptance and would not be incorporated in the standard.

In January, 1986 a number of changes were incorporated in the document that were the result of manufacturers having implemented to the specification. At this time it was agreed ESDI was beginning an era of industry wide usage and that it was suitable for proposal as an American National Standard.

In early 1986 a number of manufacturers interested in ensuring that ESDI would become an American National Standard as well as promoting the use of ESDI within industry joined together as the ESDI Steering Committee to support both these efforts. The Accredited Standards Committee (ASC) X3T9.3 ad hoc ESDI working group first met in July, 1986.

The primary objectives in developing this specification were to:

- (1) provide a low cost, high performance interface definition suitable for the smaller, high performance secondary memory devices being introduced,
- (2) develop a standard which could support higher data transfer rates as well as provide for additional performance features that are desirable on higher performance systems.
- (3) provide a single interface definition which could effectively support magnetic and optical disk drives on the same controller.

In order to accommodate the latter, the standard defines that which is common and specifies any differences by device type. An ESDI document referred to as Revision F.3 represents the documented version of an industry de facto specification which covers the early product designs using ESDI. This standard is intended to be downwards compatible with the industry Revision F.3 ESDI specification.

1.	SCOPE AND PURPOSE	1-1
1.1	SCOPE	1-1
1.2	PURPOSE	1-1
1.3	EDITORIAL CONVENTIONS	1-1
1.4	DESCRIPTION OF SECTIONS	1-1
2.	REFERENCED AND RELATED AMERICAN NATIONAL STANDARDS	2-1
2.1	REFERENCED AMERICAN NATIONAL STANDARDS	2-1
2.2	RELATED AMERICAN NATIONAL STANDARDS	2-1
3.	GENERAL DESCRIPTION	3-1
3.1	APPLICATION ENVIRONMENT	3-1
3.2	COMMAND STRUCTURE	3-1
3.3	MAGNETIC DISK SUPPORT	3-1
3.4	OPTICAL DISK SUPPORT	3-1
4.	GLOSSARY	4-1
5.	PHYSICAL INTERFACE	5-1
5.1	ELECTRICAL	5-1
5.2	CONTROL SIGNALS DRIVERS AND RECEIVERS	5-1
5.3	DATA LINE DRIVERS AND RECEIVERS	5-3
5.3.1	Line Driver	5-3
5.3.2	Line Receiver	5-3
5.4	SPECIAL TERMINATION	5-3
5.4.1	WRITE GATE Termination	5-3
5.4.2	Control Signals Termination (Optional)	5-5
5.5	RECOMMENDED CONFIGURATIONS	5-6
5.6	MECHANICAL	5-7
5.6.1	J1/P1 Connector	5-7
5.6.1.1	Edge Connector	5-7
5.6.1.2	Header Connector	5-7
5.6.2	J2/P2 Connector	5-7
5.6.2.1	Edge Connector	5-7
5.6.2.2	Header Connector	5-7
5.6.3	J3/P3 Connector	5-8
5.6.4	J4/P4 Frame Ground Connector	5-8
6.	SIGNAL LINES	6-1
6.1	PIN ASSIGNMENTS	6-1
6.2	CONTROL OUT LINES (D-O)	6-2
6.2.1	DRIVE SELECT 2(0), 2(1), 2(2) (D-O)	6-3
6.2.2	HEAD SELECT 2(0), 2(1), 2(2), and 2(3) (D-O)	6-3
6.2.2.1	Disk	6-3
6.2.2.2	Optical	6-3
6.2.3	WRITE GATE (D-O)	6-3
6.2.3.1	Disk	6-3
6.2.3.2	Optical	6-3

6.2.4	READ GATE (D-O)	6-3
6.2.5	COMMAND DATA (D-O)	6-5
6.2.6	TRANSFER REQ (D-O)	6-5
6.2.7	ADDRESS MARK ENABLE (D-x)	6-5
6.2.7.1	Disk Soft Sector (ADDRESS MARK ENABLE) (Optional)	6-5
6.2.7.2	Disk Fixed Sector (ADDRESS MARK ENABLE) (Optional)	6-5
6.3	CONTROL IN LINES	6-6
6.3.1	DRIVE SELECTED (D-O)	6-6
6.3.2	READY (D-O)	6-6
6.3.2.1	Disk	6-6
6.3.2.2	Optical	6-6
6.3.3	CONFIG/STATUS (D-O)	6-6
6.3.4	TRANSFER ACK (D-O)	6-6
6.3.5	ATTENTION (D-O)	6-7
6.3.6	INDEX (D-O)	6-7
6.3.7	SECTOR/ADDRESS MARK FOUND (D-O)	6-7
6.3.7.1	Disk and Optical Hard Sector (SECTOR)	6-7
6.3.7.2	Disk Soft Sector (ADDRESS MARK FOUND)	6-7
6.3.8	COMMAND COMPLETE (D-O)	6-7
6.3.8.1	Disk	6-8
6.3.8.2	Optical	6-8
6.4	DATA TRANSFER LINES (D-O)	6-8
6.4.1	NRZ WRITE DATA (D-O)	6-8
6.4.2	NRZ READ DATA (D-O)	6-8
6.4.3	READ/REFERENCE CLOCK (D-O)	6-9
6.4.4	WRITE CLOCK (D-O)	6-9
7.	COMMANDS AND RESPONSES	7-1
7.1	COMMAND STRUCTURE	7-1
7.2	COMMANDS	7-1
7.2.1	Command Summary	7-2
7.2.2	Magnetic Disk Commands	7-3
7.2.3	Optical Disk Commands	7-4
7.3	RESPONSES AND STATUS	7-5
7.4	SEEK (0000) (D-O)	7-5
7.4.1	Disk (SEEK)	7-5
7.4.2	Optical (SEEK)	7-5
7.4.2.1	Seek Absolute	7-5
7.4.2.2	Seek Distance/Direction	7-5
7.5	RECALIBRATE (0001)	7-5
7.5.1	Disk (RECALIBRATE)	7-5
7.5.2	Optical (RECALIBRATE)	7-6
7.6	REQUEST STATUS (0010) (D-O)	7-6
7.6.1	Disk Status Response	7-8
7.6.1.1	Disk Standard Status Response	7-8
7.6.1.2	Disk Extended Status Response (Optional)	7-9
7.6.2	Optical Status Response	7-10
7.6.2.1	Optical Standard Status	7-10
7.6.2.2	Optical Device Status	7-11

7.7	REQUEST CONFIGURATION (0011) (D-O)	7-13
7.7.1	Disk	7-13
7.7.1.1	General Configuration Response	7-13
7.7.1.2	Transfer Rate	7-14
7.7.1.3	Head Group Skew	7-14
7.7.1.4	Write Splice Information	7-14
7.7.1.4.1	READ GATE Window	7-14
7.7.1.4.2	Size of Write Splice	7-14
7.7.1.5	Drive Delays	7-14
7.7.1.5.1	Read Data Propagation Delay	7-16
7.7.1.5.2	Write Data Delay	7-16
7.7.1.5.3	Mark Detection Skew	7-16
7.7.1.6	Specific Configuration Response	7-17
7.7.1.7	Seek Overhead Skew	7-17
7.7.2	Optical	7-17
7.7.2.1	General Configuration Response Bits	7-17
7.7.2.2	Specific Configuration Response	7-20
7.7.4	Vendor Identification Response (Optional)	7-21
7.7.4.1	Standard Vendor Information	7-21
7.7.4.2	Extended Vendor Information	7-21
7.8	SELECT HEAD GROUP or SELECT TRACK (0100) (D-x)	7-22
7.8.1	Disk (SELECT HEAD GROUP)	7-22
7.9	CONTROL (0101) (D-O)	7-22
7.9.1	Disk	7-22
7.9.2	Optical	7-23
7.10	DATA STROBE OFFSET or DATA RECOVERY OFFSET (0110) (D-O)	7-24
7.10.1	Disk (DATA STROBE OFFSET)	7-24
7.10.2	Optical (DATA RECOVERY OFFSET)	7-24
7.11	TRACK OFFSET (0111) (D-O)	7-25
7.12	INITIATE DIAGNOSTICS (1000) (D-O)	7-25
7.13	SET UNFORMATTED BYTES/SECTOR (1001) (D-x) (Optional)	7-26
7.14	HIGH ORDER SEEK ADDRESS (1010) (x-O)	7-26
7.15	FORMAT (1011) (x-O)	7-27
7.15.1	Optical (FORMAT) (Optional)	7-27
7.17	Reserved (1101)	7-27
7.18	SET CONFIGURATION (1110) (D-O)	7-27
7.19	Reserved (1111)	7-28
8.	READ, WRITE AND FORMAT PARAMETERS	8-1
8.1	CRITICAL READ FUNCTION TIMING PARAMETERS (D-x)	8-1
8.1.1	Read Initialization Time.	8-1
8.1.2	READ GATE Timing	8-1
8.1.3	Read Propagation Delay	8-1
8.1.4	READ CLOCK Timing	8-1
8.1.5	REFERENCE CLOCK Valid Time	8-1

8.1.6	READ CLOCK Valid Time	8-1
8.2	CRITICAL WRITE FUNCTION TIMING PARAMETERS (D-x)	8-1
8.2.1	Read-to-Write Recovery Time	8-1
8.2.2	WRITE CLOCK-to-WRITE GATE Timing	8-1
8.2.3	Write Driver Plus Data-Encoder Turn-On From WRITE GATE	8-1
8.2.4	Write-Driver Turn-Off from WRITE GATE	8-1
8.2.5	Write-to-Read Recovery Time	8-1
8.2.6	Head Switching Time	8-1
8.2.7	Write Propagation Delay	8-1
8.3	HARD SECTORED DRIVE (D-x)	8-3
8.3.1	Fixed Sector Format Rules	8-3
8.3.2	Intersector Gap (ISG)	8-4
8.3.3	Address Area	8-5
8.3.3.1	PLO Sync Field	8-5
8.3.3.2	Byte Sync Pattern	8-5
8.3.3.3	ID	8-5
8.3.3.4	ADR Check Bytes - (ID Check Codes)	8-5
8.3.3.5	ADR Pad - (ID Pad)	8-5
8.3.4	Data Area	8-5
8.3.4.1	Write Splice	8-5
8.3.4.2	PLO Sync bytes	8-5
8.3.4.3	Byte Sync Pattern	8-5
8.3.4.4	Data Field	8-5
8.3.4.5	Data Check Bytes	8-5
8.3.4.6	Data Field Pad	8-5
8.3.5	Format Speed Tolerance Gap	8-5
8.3.6	Fixed Sector PLO Sync Format Timing	8-7
8.4	SOFT SECTORED DRIVE (D-x)	8-7
8.4.1	Soft Sector Format Rules	8-7
8.4.2	Address Mark Field	8-9
8.4.3	Address Mark Pad	8-9
8.4.4	Intersector Gap (ISG) Speed Tolerance Gap	8-9
8.4.5	Soft Sector Format Timing	8-9
8.5	CRITICAL READ FUNCTION TIMING PARAMETERS (x-O)	8-9
8.5.1	Read Initialization Time.	8-9
8.5.2	READ GATE Timing	8-9
8.5.3	Read Propagation Delay	8-9
8.5.4	READ CLOCK Timing	8-9
8.5.5	REFERENCE CLOCK Valid Time	8-9
8.5.6	READ CLOCK Valid Time	8-9
8.6	CRITICAL WRITE FUNCTION TIMING PARAMETERS (x-O)	8-9
8.6.1	Read-to-Write Recovery Time	8-9
8.6.2	WRITE CLOCK-to-WRITE GATE Timing	8-9
8.6.3	Write-Driver Plus Data-Encoder Turn-On From WRITE GATE	8-9
8.6.4	Write-Driver Turn-Off from WRITE GATE	8-9
8.6.5	Write-to-Read Recovery Time	8-11
8.6.6	Head Switching Time	8-11

8.6.7	Write Propagation Delay	8-11
8.7	OPTICAL FORMAT (x-O)	8-11
8.7.1	Format Rules	8-11
8.7.2	Intersector Gap (ISG)	8-11
8.7.3	Address Area	8-11
8.7.3.1	PLO Sync Field	8-11
8.7.3.2	Sync Pattern Byte(s)	8-11
8.7.3.3	Address Field	8-11
8.7.3.4	Address Field Check Code	8-13
8.7.3.5	Sector Write Status	8-13
8.7.3.6	Address Field Pad	8-13
8.7.4	Data Area	8-13
8.7.4.1	Write Splice	8-13
8.7.4.2	PLO Sync bytes	8-13
8.7.4.3	Recording Pattern Byte(s)	8-13
8.7.4.4	Data Field	8-13
8.7.4.5	Data Field Check Codes	8-13
8.7.4.6	Data Field Pad	8-13
8.7.5	Format Speed Tolerance Gap	8-13
8.7.6	Post Field Area	8-13
8.7.6.1	Write Splice	8-13
8.7.6.2	PLO Sync bytes	8-13
8.7.6.3	Recording Pattern Byte(s)	8-13
8.7.6.4	Post Field	8-15
8.7.6.5	Post Field Check Codes	8-15
8.7.6.6	Post Field Pad	8-15
9.	RECOMMENDATIONS	9-1
9.1	NORMAL OPERATION COMMAND TRANSFER	9-1
9.2	HANDLING COMMAND TRANSFER ERRORS	9-2
9.2.1.1	Interface Fault	9-2
9.2.1.2	Parity Error	9-2
9.2.1.3	Invalid or Unimplemented Command Fault	9-2
9.2.2	Detection and Resynchronization	9-2
9.2.3	Recovery Protocol	9-4
9.3	CONFIGURATION DATA	9-4
9.3.1	Interpreting Magnetic Disk General Config Data	9-4
9.3.1.1	Bit 15: Not a Magnetic Disk Drive	9-4
9.3.1.2	Bit 14: Format Speed Tolerance Gap Required	9-6
9.3.1.3	Bit 13: Track Offset Option Available	9-6
9.3.1.4	Bit 12: Data Strobe Offset Option Available	9-6
9.3.1.5	Bit 11: Rotational Speed Tolerance > 0.5% and < 1.0%	9-6
9.3.1.6	Bit 10: Transfer Rate > 10 MHz	9-6
9.3.1.7	Bit 9: Transfer Rate > 5 MHz, < 10 MHz	9-6
9.3.1.8	Bit 8: Transfer rate < 5 MHz	9-6
9.3.1.9	Bit 7: Removable Cartridge Drive	9-6
9.3.1.10	Bit 6: Fixed Drive	9-6
9.3.1.11	Bit 5: Spindle Motor Control Option Implemented	9-6
9.3.1.12	Bit 4: Head Switch Time > 15 Usec	9-8
9.3.1.13	Bit 3: Not MFM	9-8
9.3.1.14	Bit 2: Controller Soft Sector (Address Mark)	9-8

9.3.1.15	Bit 1: Drive Hard Sector (Sector Pulses)	9-8
9.3.1.16	Bit 0: Reserved	9-8
9.3.2	Magnetic Disk Commands Determined Invalid by Config Data	9-9
9.4.	Recommended Controller Responses to Status Conditions	9-10
9.4.1	General Error Recovery Philosophy	9-10
9.4.2	I/O Status for Error Recovery	9-10
9.5	STANDARD STATUS	9-12
9.5.1	Interpreting Magnetic Disk Standard Status	9-12
9.5.1.1	Bit 15: Reserved	9-12
9.5.1.2	Bit 14: Removable Media Not Present	9-12
9.5.1.3	Bit 13: Write Protected Removable Media	9-12
9.5.1.4	Bit 12: Write Protected, Fixed Media	9-12
9.5.1.5	Bit 11: Reserved	9-12
9.5.1.6	Bit 10: Reserved	9-12
9.5.1.7	Bit 9: Spindle Motor Stopped	9-12
9.5.1.8	Bit 8: Power On Reset Conditions Exist	9-12
9.5.1.9	Bit 7: Command Data Parity Fault	9-12
9.5.1.10	Bit 6: Interface Fault	9-12
9.5.1.11	Bit 5: Invalid or Unimplemented Command	9-12
9.5.1.12	Bit 4: Seek Fault	9-14
9.5.1.13	Bit 3: Write Gate With Track Offset	9-14
9.5.1.14	Bit 2: Vendor Unique Status Available	9-14
9.5.1.15	Bit 1: Write Fault	9-14
9.5.1.16	Bit 0: Removable Media Changed	9-14
9.5.2	Error Recovery Commands for Magnetic Disk	9-14
9.6	MISCELLANEOUS	9-18
10.	TIMING	10-1
10.1	DESKEWING	10-1
10.2	SYMBOLS	10-1
10.3	TERMS	10-1
APPENDIX	A: MAGNETIC DISK DEFECT LIST	App-1
APPENDIX	B: DISK DEFECTS	App-5
APPENDIX	C: DISK ERROR CORRECTION	App-7
APPENDIX	D: OPTICAL DISK FORMATS	App-12
APPENDIX	E: TRANSMISSION LINES AND THEIR EFFECTS ON LINE DRIVERS	App-16
APPENDIX	F: HIGH SPEED DATA PORT CONSIDERATIONS USING ECL	App-17
APPENDIX	G: RS-422 DIFFERENTIAL COMPONENTS	App-18
APPENDIX	H: DISK SERVOS	App-22
APPENDIX	I: DETERMINING READ GATE ASSERTION/NEGATION TIMES	App-25
APPENDIX	J: MAGNETIC DISK FORMATTING	App-28

FIGURES

FIGURE 5-1	CONTROL SIGNALS DRIVER/RECEIVER COMBINATION	5-2
FIGURE 5-2	DATA LINE DRIVER/RECEIVER COMBINATION	5-3
FIGURE 5-3	WRITE GATE TERMINATION	5-3
FIGURE 5-4	OPTIONAL DATA CABLE CONTROL SIGNALS TERMINATION	5-5
FIGURE 5-5	UNIDIRECTIONAL SINGLE ENDED LINES FROM DEVICE	5-6
FIGURE 5-6	UNIDIRECTIONAL SINGLE ENDED LINES FROM CONTROLLER	5-6
FIGURE 7-1	COMMAND DATA WORD STRUCTURE	7-1
FIGURE 7-2	WORD STRUCTURE OF CONFIGURATION/STATUS DATA	7-5
FIGURE 8-3	FIXED SECTOR OPTICAL FORMAT	8-11
FIGURE 10-1	ONE BIT TRANSFER TIMING (TO DRIVE)	10-2
FIGURE 10-2	ONE BIT TRANSFER TIMING-FROM DRIVE	10-3
FIGURE 10-3	WRITE ADDRESS MARK TIMING	10-4
FIGURE 10-5	TYPICAL OPERATION(S)	10-5
FIGURE 10-8	READ ADDRESS MARK TIMING	10-7
FIGURE 10-9	LOW SPEED PORT (0-15kHz) DISK NRZ READ DATA TIMINGS	10-8
FIGURE 10-10	HIGH SPEED PORT (10-21kHz) DISK NRZ READ DATA TIMINGS	10-9
FIGURE 10-11	LOW SPEED PORT (0-15kHz) DISK NRZ WRITE DATA TIMINGS	10-10
FIGURE 10-12	HIGH SPEED PORT (10-21kHz) DISK NRZ WRITE DATA TIMINGS	10-11
FIGURE 10-13	HIGH SPEED PORT AND OPTICAL NRZ TIMING PERIODS	10-12
FIGURE 10-14	FIXED SECTOR FORMAT TIMING - ALTERNATIVE 1	10-13
FIGURE 10-15	FIXED SECTOR FORMAT TIMING - ALTERNATIVE 2	10-14
FIGURE 10-16	SOFT SECTOR FORMAT TIMING	10-15
FIGURE 10-17	NORMAL OPERATION - NO ERRORS OR FAULTS	10-15
FIGURE 10-18	ABNORMAL OPERATION - FAULT DETECTED BY DRIVE WHILE	10-16
FIGURE 10-19	ABNORMAL OPERATION - POWER TURN ON OR FAULT DETECTED BY	10-16

FIGURE 10-20	COMMUNICATIONS TRANSFER FAULT, ATTENTION PREVIOUSLY NEGATED	10-17
FIGURE 10-21	COMMUNICATIONS TRANSFER FAULT, ATTENTION PREVIOUSLY ASSERTED	10-17
FIGURE 10-22	INTERFACE FAULT (<17 BITS TRANSFERRED, TRANSFER REQ=1)	10-18
FIGURE 10-23	INTERFACE FAULT WITH ATTENTION PREVIOUSLY ASSERTED	10-18
FIGURE 10-24	PARITY ERROR OR INVALID COMMAND FAULT WITH STATUS DATA RESPONSE	10-19
FIGURE 10-25	PARITY ERROR OR INVALID COMMAND FAULT WITH STATUS DATA RESPONSE	10-20
FIGURE 10-26	COMMUNICATIONS FAULTS RECOVERY PROCEDURE	10-21
FIGURE 10-27	TIMING DIAGRAM FOR A SUCCESSFUL SPINDLE START SEQUENCE	10-22
FIGURE A-1	DEFECT LIST FORMAT	App-4
FIGURE D-1	RECOMMENDED FIXED SECTOR OPTICAL FORMAT	App-13
FIGURE D-2	FIXED SECTOR OPTICAL DISK TIMING	App-13
FIGURE D-3	OPTIMEM ESDI FORMAT READ TIMING (PRELIMINARY)	App-15
FIGURE D-4	OPTIMEM ESDI FORMAT WRITE TIMING (PRELIMINARY)	App-15
FIGURE E-1	APPLICATION ENVIRONMENT	App-17

TABLES

TABLE 5-3	J3/P3 CONNECTOR PIN ASSIGNMENTS	5-13
TABLE 6-1	CONTROL CABLE (J1/P1) PIN ASSIGNMENTS	6-1
TABLE 6-2	DATA CABLE (J2/P2) PIN ASSIGNMENTS	6-1
TABLE 7-1	SUMMARY OF DEFINED COMMANDS FOR ALL DEVICES	7-2
TABLE 7-2	MAGNETIC DISK COMMAND (CMD) DATA DEFINITION	7-3
TABLE 7-3	OPTICAL DISK COMMAND (CMD) DATA DEFINITION	7-4
TABLE 7-5	REQUEST STATUS MODIFIER BITS	7-6
TABLE 7-6	MAGNETIC DISK STANDARD STATUS RESPONSE BITS	7-8
TABLE 7-6A	MAGNETIC DISK EXTENDED STATUS RESPONSE BITS	7-9

TABLE 7-7	OPTICAL DISK STANDARD STATUS RESPONSE BITS	7-10
TABLE 7-8	OPTICAL DISK DEVICE STATUS RESPONSE BITS	7-11
TABLE 7-9	OPTICAL DISK CURRENT TRACK STATUS RESPONSE	7-12
TABLE 7-11A	MAGNETIC DISK GENERAL CONFIGURATION RESPONSE BITS	7-13
TABLE 7-11B	MAGNETIC DISK GENERAL CONFIGURATION RESPONSE BITS	7-14
TABLE 7-12	MAGNETIC DISK SPECIFIC CONFIGURATION RESPONSE BITS	7-17
TABLE 7-13	OPTICAL DISK GENERAL CONFIGURATION RESPONSE BITS	7-19
TABLE 7-14	OPTICAL DISK SPECIFIC CONFIGURATION RESPONSE BITS	7-20
TABLE 7-17	VENDOR IDENTIFICATION CODES RESPONSE	7-21
TABLE 7-18	MAGNETIC DISK CONTROL COMMAND MODIFIER BITS	7-22
TABLE 7-19	OPTICAL DISK CONTROL COMMAND MODIFIER BITS	7-23
TABLE 7-21	DATA STROBE OFFSET COMMAND MODIFIERS	7-24
TABLE 7-23	TRACK OFFSET COMMAND MODIFIER BITS	7-25
TABLE 7-24	DIAGNOSTIC PARAMETER BITS	7-26
TABLE 7-25A	HIGH ORDER SEEK ADDRESS MODIFIERS	7-27
TABLE 7-29	SET CONFIGURATION SOFT SWITCH PARAMETER BITS	7-28
TABLE 9-2	I/O CONTROL SIGNAL STATES FOR ERROR RECOVERY	9-10
TABLE 9-3A	COMMANDS USED TO RECOVER FROM REPORTED STATUS	9-16
TABLE 9-3B	COMMANDS USED TO RECOVER FROM REPORTED STATUS	9-17
TABLE 9-4	STANDARD STATUS BIT PRIORITY	9-18
TABLE G-2.	APPROPRIATE RS-422 DRIVERS/RECEIVERS/TRANSCIEVERS	App-20

DRAWINGS

DRAWING 5-1	J1/P1 EDGE CONNECTOR DIMENSIONS	5-9
DRAWING 5-2	J1/P1 HEADER CONNECTOR DIMENSIONS	5-10
DRAWING 5-3	J2/P2 EDGE CONNECTOR DIMENSIONS	5-11
DRAWING 5-4	J2/P2 HEADER CONNECTOR DIMENSIONS	5-12
DRAWING 5-6	TYPICAL PCB CONNECTOR ORIENTATION	5-13
DRAWING 5-7	J4/P4 FRAME GROUND CONNECTOR	5-14

1. SCOPE AND PURPOSE

1.1 SCOPE

This document defines the Enhanced Small Device Interface (ESDI) specification. The ESDI supports magnetic and optical disk with a like physical interface, the same command/response protocol and similar commands. Any differences result from the close physical dependency of device characteristics which are different between the device classes.

The interface is capable of handling data rates from 0 to 24 Megabits per second, depending on the driver/receiver class used.

1.2 PURPOSE

The purpose of this standard is to define the interconnection of small form factor devices by using an interface that provides programmed functionality. Such functionality simplifies the intermix of devices either within a class, or between classes, on controllers.

1.3 EDITORIAL CONVENTIONS

Certain terms are used in this standard. The proper names of signals are printed wholly in uppercase to avoid possible confusion with other uses of the same word e.g. READY. Any lowercase use of these words have the normal American-English meaning.

Commands, responses and other conditions or events which have specific meaning are printed with the first letter of each word in uppercase and the remainder in lowercase e.g. Header field, Select Head Group. Any lowercase uses of these words have the normal American-English meaning.

Where D and O are used they refer to magnetic disk and optical disk respectively. To assist the reader associate information with the proper device, the expression (D-O) is used, and an x is placed where the device has no application e.g. (D-x) indicates that the content applies to magnetic disk but not optical disk.

Hexadecimal values are shown with "x" preceding a pair of quotes enclosing the value e.g. x'0A'. If the character "x" is shown within quotes it means that the nibble is undefined and can be any value e.g. x'0x' means the first nibble shall be zero and the second can be any value 0-F.

1.4 DESCRIPTION OF SECTIONS

Section 1 contains the Scope and Purpose.

Section 2 contains Referenced and Related American National Standards.

Section 3 contains the General Description.

Section 4 contains the Glossary.

Section 5 contains the electrical and mechanical characteristics; covering drivers, receivers, line termination, connectors, and cable parameters.

Section 6 contains the signal descriptions.

Section 7 contains the command descriptions and their associated responses.

Section 8 contains the assumed formatting characteristics.

Section 9 contains recommendations for miscellaneous conditions including modes of recovery from command communication faults, recommended controller interpretation of configuration data and recommended error recovery procedures based on device status.

Section 10 contains the timing diagrams.

2. REFERENCED AND RELATED AMERICAN NATIONAL STANDARDS

2.1 REFERENCED AMERICAN NATIONAL STANDARDS

None.

2.2 RELATED AMERICAN NATIONAL STANDARDS

None.

THIS

PAGE

INTENTIONALLY

LEFT

BLANK

3. GENERAL DESCRIPTION

3.1 APPLICATION ENVIRONMENT

The controller provides control of one to seven devices with a device class specific set of commands, responses and data transfers. The classes of device defined are magnetic disk and optical disk.

The physical and logical interface characteristics of magnetic disk and optical disk drives have been included in the standard.

The interface consists of a control cable and one or more data cables.

The control cable is a daisy chained cable connecting up to seven disk devices. Only the last device has line termination for the control cable.

The data cables are radially connected to the controller. The data cables utilize differential driver and receiver pairs for NRZ data and clocks. The differential signal lines are terminated at the receiving end.

3.2 COMMAND STRUCTURE

The standard defines the method of transferring commands and responses serially, generally referred to as the Serial mode of operation, which utilizes NRZ data transfer along with serial commands and serial configuration and status reporting across the command control cable.

3.3 MAGNETIC DISK SUPPORT

The interface provides for two implementations:

- o Fixed Sector drives which utilize sector pulses to define sector boundaries,
- o Soft Sector drives which use a special Address Mark that is recorded on the media to define sector boundaries.

This standard does not require that both of these be available on the drive; rather it is the choice of the manufacturer on which to implement. If desired, both may be offered with a selection method provided by the manufacturer. It is strongly recommended that controller implementations be designed to support both Fixed and Soft Sector drives.

3.4 OPTICAL DISK SUPPORT

The interface provides for the attachment of optical disk removable cartridge drives. These have unique characteristics in that data may be Read Only (ROM), Write Once/Read Many times (WORM), or Erasable.

The high track density of optical technology typically requires using media which is pre-formatted by the media supplier, or by the drive itself. Under such circumstances the controller has no need to execute what is generally referred to as a format write.

Optical WORM disks differ from magnetic disks in that although they can be

read many times, they can be recorded only once. Recording is most like an update write on magnetic disk since it consists of reading the header prior to writing the data.

Optical disks use removable media, and there is no concept of a cylinder as on magnetic disk. Typically there is only one head, and media may be single or double sided, being reversed by the user as necessary. Drives which may offer two heads so that both sides of media can be read or written do not actually support the cylinder concept because at the high track densities involved, it is necessary for the drive to re-servo when it has to switch to the second surface.

Where the term cylinder is used for magnetic disk, the term track is used with optical disk. An optical disk drive may offer fast access within a number of tracks, and if this feature is offered it is referred to as a "band", and may be viewed as a cylinder on a horizontal plane.

The raw error rate of recording at such high densities is high, and it is necessary to use extended and sophisticated error correction techniques so that the final error rate achieved is comparable to magnetic media. Provision has been made for implementation of an erasure pointer correction scheme.

Wherever possible, the established techniques used by controllers in the support of magnetic disk drives have been incorporated. It should be possible for a magnetic disk controller to support optical disks without significant re-design of electronics. Some level of change will be needed, and the extent of the changes required are to a large degree dependent on the controller design itself.

4. GLOSSARY

The definitions used in this document are listed alphabetically:

Address Mark

This term refers to a field recorded by a disk drive capable of managing soft sectors. The Address Mark is used to recognize the beginning of a sector.

Bidirectional

This term refers to the use of signal lines which may be asserted by either the master or a slave, but not concurrently.

Byte

Eight bits.

Cylinder

The set of tracks which may be accessed on a disk without performing a Seek operation. This is typically all tracks which may be accessed without physical movement of the heads.

Defect Map

A list of the defects on the device which was recorded by the manufacturer.

Head

A single read/write recording element. Each head on a disk may read or write a serial data stream. Magnetic disks typically have more than one head, with each reading or writing a different surface. Optical disks typically have only one head, requiring that the media side to be read or written be inserted in the correct plane.

Index

The nominal start of each track of the disk. The first physical sector of each track follows index.

ISG (Inter Sector Gap)

The area between the end of the data area in one sector and the start of the address area of the next sector. It contains no recognizable data and provides spacing between sectors.

MO (Magneto Optic)

A method of recording data which requires the presence of a magnetic field and a high heat source (typically supplied by a laser).

Optional

This term describes features which are not required by the standard. However, if any feature defined by the standard is implemented, it shall be done in the same way as defined by the standard.

PLD (Phase Lock Oscillator)

A circuit located in the drive which synchronizes the read data decoder to data from the disk.

Response

The information transferred from the drive to the controller following a command request. The information returned is specific to the response type.

Sector

This term refers to the address and data areas identified by the disk drive, relative to index, with zero as the first.

Servo

"Servo Data" is information written on the drive which is used by the drive to keep the read/write heads on the same track as the disk rotates. Servo data may be kept on a dedicated drive surface, embedded between data fields, both, or in some other drive specific manner.

Skew

This term refers to the practice of moving the position of the first logical sector on each track to mask the time required to switch heads or cylinders. A skew factor is added when formatting a disk to improve disk performance by reducing wasted disk revolutions.

Track

That area of the disk accessible by a single head within a single cylinder. A track may be read or written during one rotation of the disk.

Unidirectional

This term refers to the use of signal lines which are not asserted by both the controller and the device (either concurrently or successively).

Vendor Unique

This term defines those features which can be defined by a vendor in a specific implementation. Caution should be exercised in defining and using such features since they may or may not be standard between vendors.

Volume

This term refers to removable media i.e. each removable entity such as an optical disk cartridge is called a volume.

Write Protect - an attribute of removeable media, usually requiring some physical sensing by the device, indicating whether the device is allowed to write data on the medium. When a volume is write protected the device is prevented from writing on the medium.

Write Splice

An area of the disk produced when writing begins. Data in the write splice area on the disk is undefined and must not be read from the disk.

THIS

PAGE

INTENTIONALLY

LEFT

BLANK

114

5. PHYSICAL INTERFACE

5.1 ELECTRICAL

The Enhanced Small Device Interface can be divided into the physically separated categories of:

- a) Control signals.
- b) Data signals
- c) DC power.

All control lines are digital in nature (open collector TTL) and either provide signals to the drive (input) or signals to the controller (output).

The data transfer signals are differential in nature and provide data either to (write) or from (read) the drive.

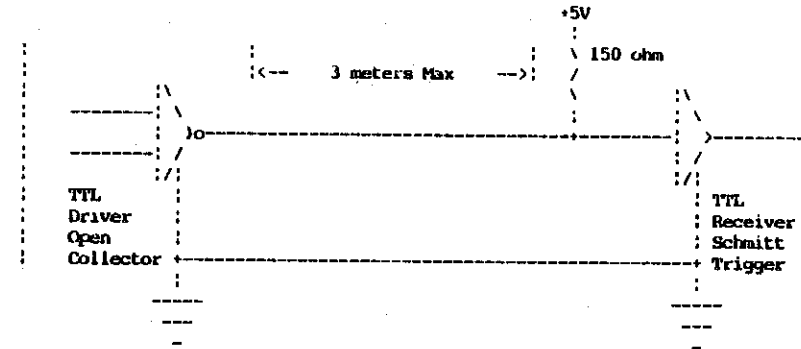
5.2 CONTROL SIGNALS DRIVERS AND RECEIVERS

The drivers have the following electrical specifications. See Figure 5-1 for the recommended configuration.

OPEN COLLECTOR:

Output:	Asserted:	0.0 VDC to 0.40 VDC @ $I = -48\text{mA}$
	Negated:	2.5 VDC to 5.25 VDC @ $I = +250\text{ uA}$
Input:	Asserted:	0.0 VDC to 0.80 VDC
	Negated:	2.0 VDC to 5.25 VDC
Maximum Total Input Load:		-0.4 mA @ 0.4 VDC
Maximum Input Hysteresis:		0.2 VDC

The signals are driven with an open collector output stage capable of sinking at least 48mA when asserted with maximum voltage of 0.4V measured at the driver. When the line driver is negated the driver transistor is off and collector leakage current is a maximum of 250uA.



NOTE: Termination resistors for lines originating at the controller are located in the last device. Lines originating at the device are terminated at the controller.

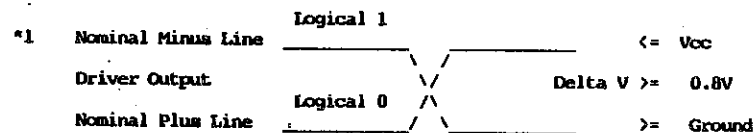
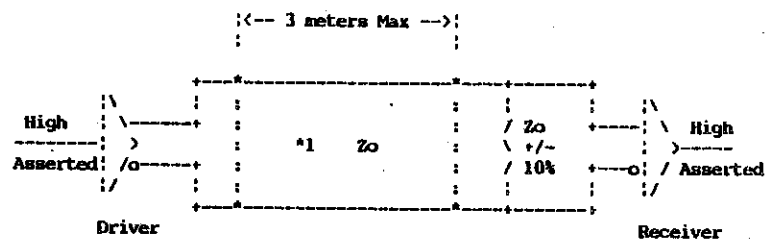
FIGURE 5-1 CONTROL SIGNALS DRIVER/RECEIVER COMBINATION

5.3 DATA LINE DRIVERS AND RECEIVERS

The data drivers and receivers are differential in nature, and may use either TTL or ECL logic. It is possible to mix implementations provided the common mode voltage between TTL driver and ECL receiver is sufficiently small. NOTE: Implementations need to consider the potential lack of common mode noise rejection in a mixed implementation.

In high speed applications, driver and receiver skew specifications are critical for timing calculations.

The recommended circuit is shown in Figure 5-2.



NOTE: Z_o = Flat Ribbon cable impedance rated at 105 ohms

FIGURE 5-2 DATA LINE DRIVER/RECEIVER COMBINATION

5.3.1 Line Driver

The driver circuit shall result in a balanced voltage source that will produce a differential voltage to the interconnecting cable. The difference between the Low output level and the High output level shall be greater than or equal to 0.8V with a 100 ohm load. With no load the difference between the output voltages shall be less than 6V and neither output voltage may be greater than 6V.

NOTE: A driver which complies with EIA Standard RS-422 meets these requirements.

5.3.2 Line Receiver

The receiver shall have an input sensitivity of $\pm 200\text{mV}$ minimum over its entire VCM (Common Mode Voltage) range (at least -1V to +7V referenced to the receiver ground). The VCM is defined as the algebraic mean of the two voltages appearing at the receiver input terminals with respect to the:

receiver circuit ground.

The receiver input impedance shall be greater than or equal to 4,000 ohms, and the receiver shall maintain correct operation for a differential input signal ranging between 200mV and 6V in magnitude.

The maximum voltage between either receiver input terminal and receiver circuit ground shall not exceed $\pm 10\text{V}$ (3V signal + 7V VCM), and the circuit shall be able to tolerate a maximum differential signal of 12V applied across its input terminals without being damaged.

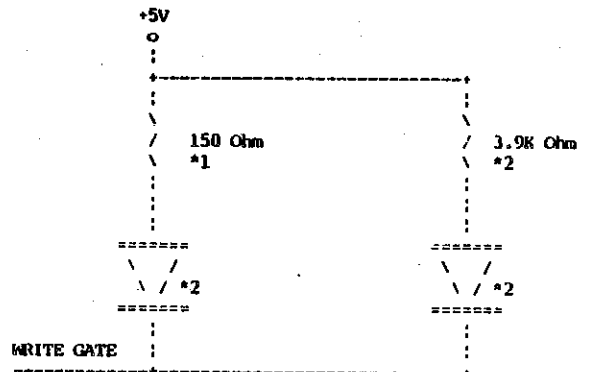
NOTE: A receiver which complies with EIA Standard RS-422 meets these requirements.

5.4 SPECIAL TERMINATION

For differential applications, the termination shall consist of a resistor with impedance equal to $Z_o \pm 10\%$ connected between the (+) data line and the (-) data line at the receiving end of the cable.

5.4.1 WRITE GATE Termination

This line shall be protected from terminator power loss by the circuit shown in Figure 5-3. The circuit is intended to prevent the WRITE GATE signal from being unintentionally asserted upon loss of terminator power. The 3.9K Ohm resistor permanently located in each drive assures line negation regardless of terminator power status.



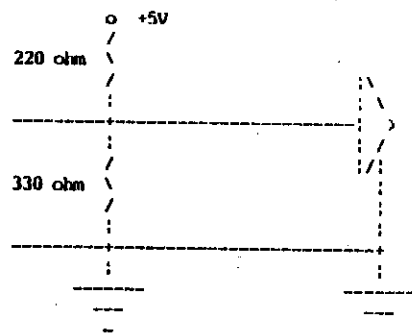
*1 Part of terminator pack in last drive of the daisy chain.

*2 Permanently located in the drive. Voltage drop shall be $< 0.75\text{V}$.

FIGURE 5-3 WRITE GATE TERMINATION

5.4.2 Control Signals Termination (Optional)

The optional termination for control signals only is applicable to the data cable.

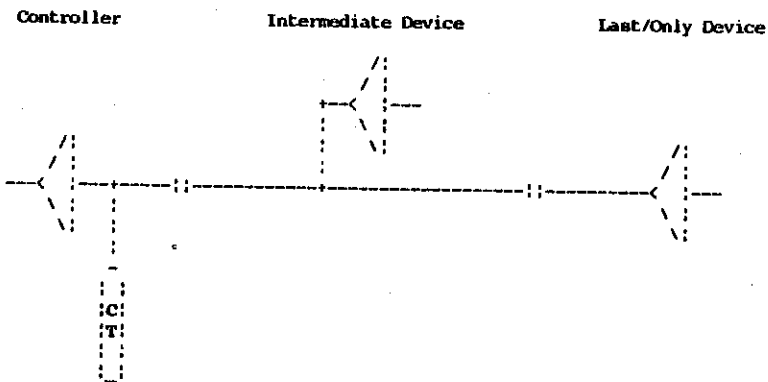


	Signal
Drive Termination	Address Mark Enable
Controller Termination	Drive Selected
	Sector/Address Mark Found
	Command Complete
	Index

FIGURE 5-4 OPTIONAL DATA CABLE CONTROL SIGNALS TERMINATION

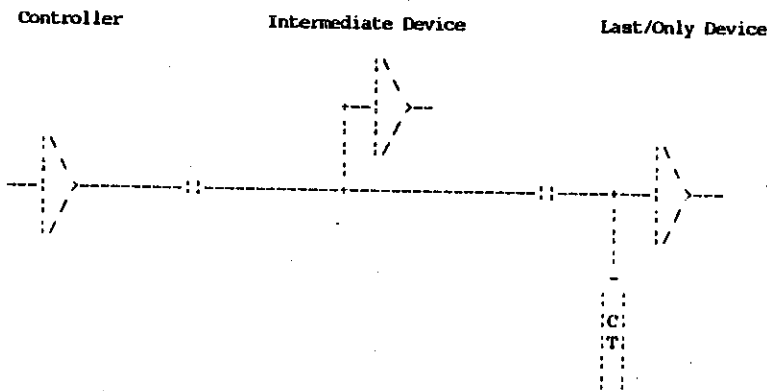
5.5 RECOMMENDED CONFIGURATIONS

Figure 5-4, Figure 5-5, and Figure 5-6 illustrate cabling for recommended configurations with the controller at one end of the cable, an ESDI device at the other, and with intermediate ESDI devices. Other electrically valid configurations may be used, even if they are not shown below.



Note: Terminator CT, shall be at extreme end of the cable

FIGURE 5-5 UNIDIRECTIONAL SINGLE ENDED LINES FROM DEVICE



Note: Terminator CT, shall be at extreme end of the cable

FIGURE 5-6 UNIDIRECTIONAL SINGLE ENDED LINES FROM CONTROLLER

5.6 MECHANICAL

The connectors and cables that interconnect a device and the controller are shown in Table 5-1.

Dev/Ctrl Receptacle	Cable/Connector Description	Cable Plug
J1	Control signals (multiplexed)	P1
J2	Data signals (radial)	P2
J3	DC power input	P3
J4	Frame ground	P4

TABLE 5-2 CABLES AND CONNECTORS

5.6.1 J1/P1 Connector

5.6.1.1 Edge Connector

The dimensions for the J1 34 pin PCB edge connector and its P1 mating cable connector are shown in Drawing 5-1. The pins are numbered 1 through 34 with the even pins located on one side of the PCB. A key slot is provided between pins 4 and 6.

See Drawing 5-6 for connector orientation.

5.6.1.2 Header Connector

The dimensions for the J1 34 pin header and its P1 mating cable connector are shown in Drawing 5-1.

5.6.2 J2/P2 Connector

5.6.2.1 Edge Connector

The dimensions for the J2 20 pin PCB edge connector and its P2 mating cable connector are shown in Drawing 5-3. The pins are numbered 1 through 20 with the even pins located on one side of the PCB. A key slot is provided between pins 4 and 6.

See Drawing 5-6 for orientation.

5.6.2.2 Header Connector

The dimensions for the J2 20 pin header and its P2 mating cable connector are shown in Drawing 5-4.

5.6.3 J3/P3 Connector

The dimensions for the J3 DC Power connector and its P3 mating cable connector are shown in Drawing 5-5.

The pin assignments are as shown in Table 5-2.

See Drawing 5-6 for connector orientation.

5.6.4 J4/P4 Frame Ground Connector

The J4/P4 frame ground connection and its mating P4 cable connector are shown in Drawing 5-7.

If wire is used, the hole in J4 shall accommodate a wire size of 18AWG maximum.

TO BE SUPPLIED BY AMP

DRAWING 5-1 J1/P1 EDGE CONNECTOR DIMENSIONS

TO BE SUPPLIED BY AMP

HEADER
CABLE
304

DRAWING 5-2 J1/P1 HEADER CONNECTOR DIMENSIONS

119

TO BE SUPPLIED BY AMP

DRAWING 5-3 J2/P2 EDGE CONNECTOR DIMENSIONS

TO BE SUPPLIED BY AMP

DRAWING 5-4 J2/P2 HEADER CONNECTOR DIMENSIONS

TO BE SUPPLIED BY AMP

DRAWING 5-5 J3/P3 CONNECTOR

J3 Connector Pin	Voltage
1	+12V DC +5%
2	12V RETURN
3	5V RETURN
4	+ 5V DC +5%

TABLE 5-3 J3/P3 CONNECTOR PIN ASSIGNMENTS

TO BE SUPPLIED BY AMP

DRAWING 5-6 TYPICAL PCB CONNECTOR ORIENTATION

TO BE SUPPLIED BY AMP

DRAWING 5-7 J4/P4 FRAME GROUND CONNECTOR

THIS

PAGE

INTENTIONALLY

LEFT

BLANK

6. SIGNAL LINES

6.1 PIN ASSIGNMENTS

Pin assignments for the control cable are shown in Table 6-1 and Table 6-2. The direction --> is Output to the drive and <-- is Input to the controller.

Disk Signals	Optical Signals	Signal Pin	Ground Pin
Cable is Flat Ribbon (3 meters maximum)			
--> HEAD SELECT 2(3)	Reserved	2	1
--> HEAD SELECT 2(2)	Reserved	4	3
--> WRITE GATE	=	6	5
<-- CONFIG/STATUS DATA	=	8	7
<-- TRANSFER ACK	=	10	9
<-- ATTENTION	=	12	11
--> HEAD SELECT 2(0)	SIDE SELECT	14	13
<-- SECTOR/AM FOUND	SECTOR	16	15
--> HEAD SELECT 2(1)	Reserved	18	17
<-- INDEX	=	20	19
<-- READY	=	22	21
--> TRANSFER REQ	=	24	23
--> DRIVE SELECT 2(0)	=	26	25
--> DRIVE SELECT 2(1)	=	28	27
--> DRIVE SELECT 2(2)	=	30	29
--> READ GATE	=	32	31
--> COMMAND DATA	=	34	33

TABLE 6-1 CONTROL CABLE (J1/P1) PIN ASSIGNMENTS

Disk Signals	Optical Signals	Signal Pin	Ground Pin
Cable is Flat Ribbon (3 meters maximum)			
<-- DRIVE SELECTED	=	1	
<-- SECTOR/AM FOUND	SECTOR	2	
<-- COMMAND COMPLETE	=	3	
--> ADDRESS MARK ENABLE	Reserved	4	
GROUND	=	5	
--> +/-WRITE CLOCK	=	7/8	6
GROUND	=	9	
<-- +/-READ REFERENCE CLOCK	=	10/11	12
--> +/-NRZ WRITE DATA	=	13/14	15/16
<-- +/-NRZ READ DATA	=	17/18	19
<-- INDEX	=	20	

TABLE 6-2 DATA CABLE (J2/P2) PIN ASSIGNMENTS

6.2 CONTROL OUT LINES (D-0)

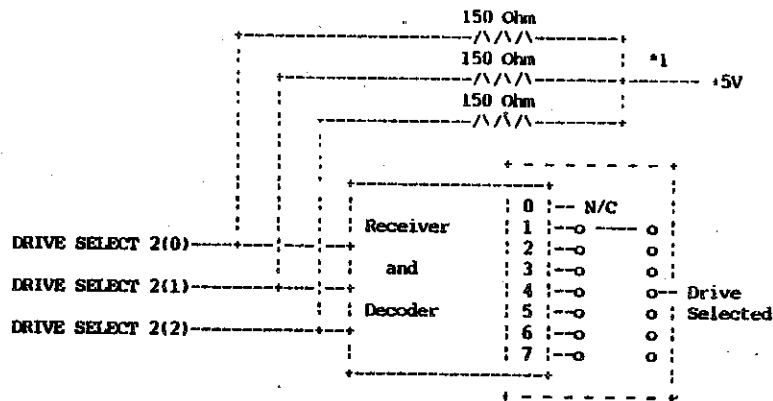
The control out signals are of two types: those to be multiplexed in a multiple drive system and those intended to do the multiplexing. The control signals to be multiplexed are WRITE GATE, READ GATE, HEAD SELECT 2(0), HEAD SELECT 2(1), HEAD SELECT 2(2), HEAD SELECT 2(3), TRANSFER REQ and COMMAND DATA. The signals to do the multiplexing are DRIVE SELECT 2(0), DRIVE SELECT 2(1), and DRIVE SELECT 2(2).

ADDRESS MARK ENABLE (D-x) is a control output in the radial cable. It is not multiplexed.

Any lines not used should be terminated.

6.2.1 DRIVE SELECT 2(0), 2(1), 2(2) (D-O)

The three DRIVE SELECT lines are to be decoded for drive select. Decode: 000 shall be a no select. See Figure 6-1 and Table 6-3. Drives should be deselected prior to power down as insurance against destructive writing.



NOTE: Illustrative implementation only. If a minimum delay time greater than zero is necessary for any reason a delay may be added between the selection circuit and the assertion of Drive Selected.

*1 Termination resistors are located in last drive only.

FIGURE 6-1 DRIVE SELECT TERMINATION

Drive Selected	None	1	2	3	4	5	6	7
Drive Select 2(2)	0	0	0	0	1	1	1	1
Drive Select 2(1)	0	0	1	1	0	0	1	1
Drive Select 2(0)	0	1	0	1	0	1	0	1

TABLE 6-3 DRIVE SELECTION MATRIX

6.2.2 HEAD SELECT 2(0), 2(1), 2(2), and 2(3) (D-O)

6.2.2.1 Disk

These four lines allow selection of each individual read/write head in a binary coded sequence. HEAD SELECT 2(0) is the least significant line. Heads are numbered 0 through 15. When all HEAD SELECT lines are negated, head 0 shall be selected. Addressing more than 16 heads is allowed by use of the SELECT HEAD GROUP command.

Head addressing is continuous from 0 through both removable and fixed drives. If a removable drive is present, head 0 shall be on the removable

media drive.

Addressing more heads than contained in the drive shall result in a write fault when attempting to perform a write operation.

A 150 OHM resistor pack allows for line termination.

6.2.2.2 Optical

One line, SIDE SELECT, is used to allow selection of either side of the cartridge platter. Heads are numbered as 0 and 1. When SIDE SELECT is negated, head 0 shall be selected.

If only one side is supported by a drive then this line may be ignored.

6.2.3 WRITE GATE (D-O)

6.2.3.1 Disk

This signal allows data to be recorded on the disk.

The assertion of this signal creates a write splice and initiates the writing of the header or the data PLO Sync field by the drive. The write splice is defined as the point at which WRITE GATE turns on or off relative to Index or Sector pulse. See Figure 10-14 through Figure 10-16.

When formatting, WRITE GATE should be negated for 2 bit times minimum between the address area and the data area to identify to the drive the beginning of the data PLO sync field.

This line shall be protected from terminator power loss by implementation of the circuit shown in Figure 5-3.

6.2.3.2 Optical

This signal allows data to be recorded on the disk.

The assertion of this signal initiates the writing of the data PLO Sync field by the drive. See Figure 10-15.

This line shall be protected from terminator power loss by implementation of the circuit shown in Figure 5-3.

6.2.4 READ GATE (D-O)

This signal allows data to be read from the disk.

READ GATE should only be asserted during a PLO sync field and at least the number of bytes defined by the drive prior to the ID or Data Sync Bytes. The PLO sync field length is determined by the response to the Request PLO Sync Field Length command. READ GATE shall be negated when passing over a write splice area.

A 150 OHM resistor pack allows for line termination.

6.2.5 COMMAND DATA (D-O)

When presenting a command, 16 information bits of serial data plus parity, will be presented on this line. This data is to be controlled by the handshake protocol with signals TRANSFER REQ and TRANSFER ACK. Upon receipt of this serial data, the drive shall perform the required function as specified by the bit configuration. Data is transmitted MSB first.

See Section 7 for the meaning of the various bit combinations.

See Figure 10-1 for timing.

The parity utilized in all commands shall be odd. The parity bit shall be a "1" when the number of "1's" in a 16 bit command is an even number. The number of "1" bits in the command, plus parity, shall be an odd number.

No communications should be attempted unless the COMMAND COMPLETE line is asserted just prior to the start of the transfer. Note: This line shall be negated when not in use.

A 150 OHM resistor pack allows for line termination.

6.2.6 TRANSFER REQ (D-O)

The Transfer Request line functions as a handshake signal in conjunction with TRANSFER ACK during command and configuration/status transfers. See Figure 10-1 and Figure 10-2 for timing.

6.2.7 ADDRESS MARK ENABLE (D-x)

This line shall be permanently terminated in the drive. See Figure 5-1.

6.2.7.1 Disk Soft Sector (ADDRESS MARK ENABLE) (Optional)

For soft sector drives, this signal, when WRITE GATE is asserted, causes an Address Mark to be written. ADDRESS MARK ENABLE shall be asserted for 24 +/-1 bit times. See Figure 10-3 for timing. The Address Mark written shall be left to the drive manufacturer's discretion.

ADDRESS MARK ENABLE, when asserted without WRITE GATE or READ GATE, causes a search for Address Marks. If WRITE GATE is asserted, the negation of this signal causes the drive to begin writing the ID PLO Sync field. See Figure 10-16.

6.2.7.2 Disk Fixed Sector (ADDRESS MARK ENABLE) (Optional)

In fixed sector drives, ADDRESS MARK ENABLE does not cause an Address Mark to be written on the media. The trailing edge of ADDRESS MARK ENABLE with WRITE GATE asserted initiates the writing of the header PLO sync field.

The beginning of a Header PLO Sync Field shall be defined at format time by using ADDRESS MARK ENABLE or by the leading edge of WRITE GATE assertion. See Figure 10-14 and Figure 10-15.

6.3 CONTROL IN LINES

All control cable input lines are enabled by their respective DRIVE SELECT decodes.

All data cable input lines are always enabled.

Figure 5-1 shows the recommended circuit.

6.3.1 DRIVE SELECTED (D-O)

A status line provided at the data cable connector to inform the controller of the selection status of the drive. The DRIVE SELECTED line is driven by a driver as shown in Figure 5-1. This signal shall be asserted only when the drive is selected as defined in section 6.2.1. The DRIVE SELECT output lines on the control cable are asserted by the controller. See Figure 10-4.

6.3.2 READY (D-O)

6.3.2.1 Disk

This signal indicates only that the spindle is up to speed.

When this interface signal is asserted it indicates that the drive spindle is up to speed. When this signal is negated the drive spindle is not up to speed.

6.3.2.2 Optical

This signal indicates that the spindle is up to speed. This interface signal when asserted, together with COMMAND COMPLETE indicates that the drive is ready to read, write or, seek. When the line is negated, all reading, writing and seeking is inhibited.

6.3.3 CONFIG/STATUS (D-O)

The drive presents serial data on the Configuration/Status Data line upon request from the controller. See Figure 10-5 for typical operation. This config-status serial data shall be presented to the interface and transferred using the handshake protocol with signals TRANSFER REQ and TRANSFER ACK. See Figure 10-2. Once initiated, 16 bits plus parity shall be transmitted MSB first. The parity utilized shall be odd.

6.3.4 TRANSFER ACK (D-O)

The Transfer Acknowledge signal functions as a handshake signal along with TRANSFER REQ during COMMAND and CONFIGURATION-STATUS transfers. See Figure 10-1 and Figure 10-2.

6.3.5 ATTENTION (D-O)

ATTENTION is asserted when the drive wants the controller to request its standard status. Generally, this is a result of a fault condition or a change of status.

If a selected device encounters an unexpected condition which causes it to become busy and unable to respond to the controller it shall assert ATTENTION in conjunction with the negation of COMMAND COMPLETE.

Writing is inhibited when ATTENTION is asserted.

ATTENTION is negated by the Control Command with the Reset Interface Attention modifier set only if the condition which caused it to occur no longer exists.

6.3.6 INDEX (D-O)

This pulse is provided by the drive once each revolution to indicate the beginning of a track. This signal is asserted to indicate INDEX. Only the transition at the leading edge of the asserted pulse is accurately controlled. See Figure 10-6. This signal is available on the control cable (gated) and on the radial data cable (ungated).

6.3.7 SECTOR/ADDRESS MARK FOUND (D-O)

These signals are available on the control cable (gated) and on the radial data cable (ungated). Both of these signals shall be implemented by the drive manufacturer.

6.3.7.1 Disk and Optical Hard Sector (SECTOR)

This interface signal, which is mutually exclusive with ADDRESS MARK FOUND, indicates the start of a sector. The leading edge of the asserted sector pulses is the only edge that is accurately controlled. The index pulse indicates sector zero. See Figure 10-7.

No short sector lengths are allowed.

6.3.7.2 Disk Soft Sector (ADDRESS MARK FOUND)

This interface signal, which is mutually exclusive with SECTOR, indicates the detection of the end of an address mark. See Figure 10-8 for timing.

6.3.8 COMMAND COMPLETE (D-O)

A status line provided at the radial data cable connector. This ungated input to the controller allows the drive's COMMAND COMPLETE status to be monitored during overlapped commands without selecting the drive.

This signal shall be negated in the following cases:

- o During a power up sequence, this line shall stay negated until the power up sequence is complete.

- o Upon receipt of the first COMMAND DATA bit, COMMAND COMPLETE shall stay negated during the entire command sequence.
- o Whenever the drive is unable to respond to the interface e.g. during recovery from internally detected error conditions. If this should occur during the time that the device is selected, then ATTENTION shall be asserted to advise the controller that an unexpected event has occurred and that the device is busy and unable to respond to the interface.

This signal is driven by an open collector driver as shown in Figure 5-1.

6.3.8.1 Disk

The power up sequence includes eventual spin up and recalibration.

This signal line is also negated within 15 usec from a head select change if head selection time is >15usec as indicated in Configuration Data.

6.3.8.2 Optical

The power up sequence includes a recalibration sequence initiated by drive logic if the R/W heads are not over track zero.

6.4 DATA TRANSFER LINES (D-O)

All lines associated with the transfer of data between the drive and the controller are differential in nature and may not be multiplexed. These lines are provided at the radial data cables of each drive.

Four pairs of balanced signals are used with magnetic and optical disk for the transfer of data and clock: NRZ WRITE DATA, NRZ READ DATA, WRITE CLOCK, and READ/REFERENCE CLOCK.

Figure 5-2 illustrates the recommended driver/receiver circuit.

6.4.1 NRZ WRITE DATA (D-O)

This is a differential pair that defines the data to be written on the track. This data shall be clocked by the WRITE CLOCK signal.

See Figure 10-11 and Figure 10-12 for timing.

6.4.2 NRZ READ DATA (D-O)

The data recovered by reading previously written information is transmitted to the controller via the differential pair of NRZ READ DATA lines. This data is clocked by the READ CLOCK signal. See Figure 10-9 and Figure 10-10 for timing. NRZ READ DATA shall be held at a zero level until PLO sync has been obtained and data is valid.

NOTE: NRZ READ DATA carries erasure pointer data during an erasure read in an optical disk drive.

6.4.3 READ/REFERENCE CLOCK (D-O)

The timing diagrams as shown in Figure 10-9 through Figure 10-12 depict the necessary sequence of events (with associated timing restrictions for proper read/write operation of the drive). The REFERENCE CLOCK signal from the drive shall determine the data transfer rate.

REFERENCE CLOCK is present and stable when READY is asserted (the drive is spinning).

READ CLOCK is valid when READ GATE is active and PIO Synchronization has been established.

REFERENCE CLOCK is valid when READ GATE is inactive.

All transitions between REFERENCE CLOCK and READ CLOCK shall be performed without glitches. Two missing clock cycles are permissible.

NOTE: On magnetic disk drives above 10 MHz transfer rate and on optical disks it is possible that the negated period of the clock may vary i.e. not be symmetrical in width with the asserted period. This occurs on a repetitive basis but not necessarily on every clock cycle. The leading edge of the signal shall be used to clock data.

6.4.4 WRITE CLOCK (D-O)

WRITE CLOCK is provided by the controller and shall be at the bit data rate. This clock frequency shall be dictated by the READ/REFERENCE CLOCK during the write operation. See Figure 10-11 and Figure 10-12 for timing.

WRITE CLOCK need not be continuously supplied to the drive. WRITE CLOCK should be supplied before beginning a write operation and should last for the duration of the write operation.

Optical disks may not produce a symmetrical wave form due to device specific implementations. Both standard and extended period wave forms are permitted, as defined in Figure 10-13.

THIS

PAGE

INTENTIONALLY

LEFT

BLANK

127

THIS

PAGE

INTENTIONALLY

LEFT

BLANK

7. COMMANDS AND RESPONSES

7.1 COMMAND STRUCTURE

Each command consists of 17 bits (16 command data + 1 parity) transferred serially. The structure of the command word is defined in Figure 7-1.

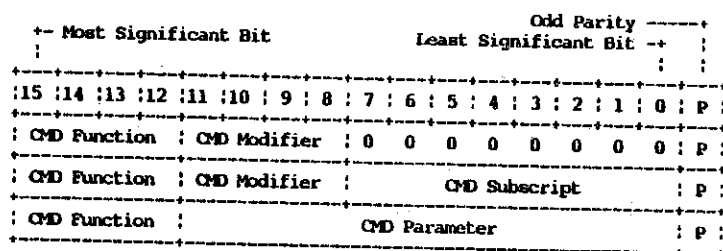


FIGURE 7-1 COMMAND DATA WORD STRUCTURE

7.2 COMMANDS

Within the commands, all unused or Not Applicable bits shall be set to zero and any "Reserved" Command Function received shall be treated as invalid.

Within this document, references to commands and modifiers are by bit e.g. 0011. References to subscripts are made as a value in the range of 0-255 since they are typically encoded e.g. 3.

7.2.1 Command Summary

Table 7-1 summarizes the commands defined for disk and optical implementations.

Magnetic Disk	Optical Disk
Seek	Seek
Recalibrate	Recalibrate
Request Status	Request Status
Request Configuration	Request Configuration
Select Head Group *	Reserved
Control	Control
Data Strobe Offset *	Data Recovery Offset *
Track Offset *	Track Offset *
Initiate Diagnostics *	Initiate Diagnostics *
Set Bytes per Sector *	Reserved
Reserved	High Order Seek Addr
Reserved	Format *
Reserved	Reserved
Reserved	Reserved
Set Configuration *	Set Configuration *
Reserved	Reserved
* Optional Commands	

TABLE 7-1 SUMMARY OF DEFINED COMMANDS FOR ALL DEVICES

7.2.2 Magnetic Disk Commands

Table 7-2 summarizes the magnetic disk commands.

CMD Fctn Bit 15-12	CMD Function Definition	CMD Modifier Applicable Bits 11-8	CMD Subscript Applicable Bits 7-0	CMD Parameter Applicable Bits 11-0	Status/ Config Data to Ctrl
0000	Seek	No	No	Yes	No
0001	Recalibrate	No	No	No	No
0010	Request Status	Yes	Yes	No	Yes
0011	Request Configuration	Yes	Yes	No	Yes
0100	Select Head Group *	No	No	Yes	No
0101	Control	Yes	No	No	No
0110	Data Strobe Offset *	Yes	Yes	No	No
0111	Track Offset *	Yes	Yes	No	No
1000	Initiate Diagnostics *	No	No	Yes	No
1001	Set Bytes per Sector *	No	No	Yes	No
1010	Reserved	-	-	-	-
1011	Reserved	-	-	-	-
1100	Reserved	-	-	-	-
1101	Reserved	-	-	-	-
1110	Set Configuration *	Yes	Yes	No	No
1111	Reserved	-	-	-	-
* Optional Commands					

NOTE: Simultaneous Data Strobe and Track Offsets are allowed by multiple commands

TABLE 7-2 MAGNETIC DISK COMMAND (CMD) DATA DEFINITION

7.2.3 Optical Disk Commands

Table 7-3 summarizes the optical disk commands.

CMD Fctn Bit 15-12	CMD Function Definition	CMD Modifier Applicable Bits 11-8	CMD Subscript Applicable Bits 7-0	CMD Parameter Applicable Bits 11-0	Status/ Config Data to Ctrl
0000	Seek	No	No	Yes	No
0001	Recalibrate	No	No	No	No
0010	Request Status	Yes	Yes	No	Yes
0011	Request Configuration	Yes	Yes	No	Yes
0100	Reserved	-	-	-	-
0101	Control	Yes	No	No	No
0110	Data Recovery Offset *	Yes	Yes	No	No
0111	Track Offset *	Yes	Yes	No	No
1000	Initiate Diagnostics *	No	No	Yes	No
1001	Reserved	-	-	-	-
1010	Hi Order Seek Address	Yes	No	Yes	No
1011	Format *	No	No	No	No
1100	Reserved	-	-	-	-
1101	Reserved	-	-	-	-
1110	Set Configuration *	Yes	Yes	-	-
1111	Reserved	-	-	-	-
* Optional Commands					

TABLE 7-3 OPTICAL DISK COMMAND (CMD) DATA DEFINITION

7.3 RESPONSES AND STATUS

When response or status information is requested by the controller via the proper commands 17 bits (16 data + 1 parity) of information are returned to the controller.

The specific information requested is specified by bits 11-8 of the command. The format of the response or status information returned is defined in Figure 7-2.

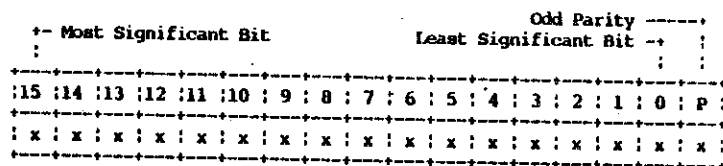


FIGURE 7-2 WORD STRUCTURE OF CONFIGURATION/STATUS DATA

7.4 SEEK (0000) (D-O)

7.4.1 Disk (SEEK)

The Seek command causes the drive to seek to the cylinder value contained in bits 11-0. A Seek command shall restore data strobe and track offsets to zero.

7.4.2 Optical (SEEK)

There are two alternative methods to initiate a Seek operation in the drive.

7.4.2.1 Seek Absolute

This Seek command causes the drive to begin the seek to the track value contained in bits 11-0 plus the value set by the high order 4 bits of the High Order Seek Address command (1010).

7.4.2.2 Seek Distance/Direction

This Seek command causes the drive to begin a seek of the distance specified in bits 11-0; plus the high order 4 bits value and the direction set by CMD modifier bits 8-9 in the High Order Seek Address command (1010). See Table 7-25A.

7.5 RECALIBRATE (0001)

7.5.1 Disk (RECALIBRATE)

The Recalibrate command shall cause the actuator to return to cylinder zero and restore Data Strobe Offset and Track Offset to zero.

7.5.2 Optical (RECALIBRATE)

On drives which implement Seek Absolute the Recalibrate command shall cause the actuator to return to track zero and restore Data Recovery Offset to zero.

On drives which implement Seek Distance/Direction the Recalibrate command causes no physical movement and restores Data Recovery Offset to zero.

7.6 REQUEST STATUS (0010) (D-O)

This command causes the drive to send 16 bits of standard or vendor unique status information to the controller as determined by the command modifier bits. The parity utilized in all status responses shall be odd.

Bits 15-12 shall not cause ATTENTION to be asserted. Bits 11-0 are fault or change of status bits that may cause ATTENTION to be asserted each time one is set (see Table 7-6 through Table 7-10).

Command Modifier Bits 11-8	Sub-command Bits 7-0	Function
0 0 0 0	0	Request Standard Status
0 x x x	1	Request Extended Standard Status
1 0 0 0		Request Vendor Unique Status
to		(D-x) Reserved
1 1 1 1		
1 0 0 0		(x-O) Reserved
1 0 0 1		(x-O) Request Optical Device Status
1 0 1 0		(x-O) Request Current Track Position
1 0 1 1		
to		(x-O) Reserved
1 1 1 1		

NOTE: The controller should respond with a Request Status command for any command transfer protocol timeout. Some drives have been built which implement Additional Vendor Unique Status in 8-15. The controller can manage this by requesting as many Vendor Unique status words as reported in its Configuration response.

TABLE 7-5 REQUEST STATUS MODIFIER BITS

o Request Standard Status: When the command modifier bits 11-8 of the Request Status command is 0000, the drive shall respond with 16 bits of Standard Status. Settings in this status may or may not be fault or change of status bits that cause ATTENTION to be asserted each time one is set.

o Request Vendor Unique Status: When the command modifier bits 11-8 of the Request Status command is 0001-0111, the drive responds with Vendor Unique Status (undefined in this specification). The number of words available is specified by configuration data. Each word of Vendor Unique Status is requested using a different command modifier configuration. The command modifier for the first word is 0001 and subsequent words are requested by incrementing the command modifier.

o Request Device Type or Additional Vendor Unique Status or Reserved: When the command modifier bits 11-8 of the Request Status command is 1000-1111, the drive responds with all zeros if it is Reserved. The number of words available is specified by configuration data. If Device Type Status is defined in this specification it shall be reported, or additional vendor unique status may be provided. Each word of Additional Vendor Unique Status is requested using a different command modifier configuration. The command modifier for the first word is 1000 and subsequent words are requested by incrementing the command modifier by one.

7.6.1 Disk Status Response

7.6.1.1 Disk Standard Status Response

There are 16 bits of status information returned to the controller in response to the Request Standard Status command with a subscript of 0.

Table 7-6 lists the conditions under which the Status Response bits shall be set and if ATTENTION is asserted.

Bit	Status Response	ATTN
15	Reserved	0
14	1 = Removable Media Not Present	0
	0 = If not Removable	0
13	1 = Write Protected - Removable Media	0
	0 = If not Removable	0
12	1 = Write Protected - Fixed Media	0
11	Reserved	0
10	Reserved	0
9	1 = Spindle Motor Stopped by Stop command	0
	1 = Spindle Motor Stopped for Other (e.g. Power On, Reset)	1
8	1 = Power On Condition *1	1
7	1 = Command Data Parity Fault	1
6	1 = Interface Fault	1
5	1 = Invalid or Unimplemented Command Fault	1
4	1 = Seek Fault	1
3	1 = Write Gate with Track Offset Fault	1
2	1 = Vendor Unique Status Not Associated with Error Condition	0
	1 = Vendor Unique Status Available for Error Condition	1
1	1 = Write Fault *2	1
0	1 = Removable Media Changed (has been changed since last status request)	1

*1 The drive is set to the initial condition that follows the Power On sequence of the drive electronics. This bit is set at initial Power On or as the result of an unexpected fault condition detected by the drive which resulted in an internal reset which returns the electronics to the same condition as after initial Power On. The controller will typically have to configure the drive when this condition occurs and may have to issue a Start Spindle Motor command.

*2 Conditions that can cause Write Fault are:

- Write current in a head without WRITE GATE asserted or no write current with WRITE GATE asserted and the drive selected.
- Multiple heads selected, no head selected, or improperly selected with WRITE GATE asserted.
- WRITE GATE asserted to a write protected drive.
- Simultaneous assertion of READ GATE and WRITE GATE.
- DC voltages grossly out of tolerance.
- Other drive specific error conditions.

TABLE 7-6 MAGNETIC DISK STANDARD STATUS RESPONSE BITS

7.6.1.2 Disk Extended Status Response (Optional)

There are 16 bits of status information returned to the controller in response to the Request Standard Status command with a subscript of 1.

Table 7-6A lists the conditions under which the Status Response bits shall be set and if ATTENTION is asserted.

Bit	Extended Status Response	ATT
15	1 = Spindle is Synchronized *1	1
14	Reserved	0
1		

*1 In a synchronized spindle system one drive (the master) or an external source generates a signal which drives use to synchronize their spindles. A separate cable outside the ESDI definition supplies this signal to each drive. There are three conditions:

- o Spindle is not synchronized to an external signal.
- o Spindle is synchronized.
- o Spindle had reported it was synchronized, but is no longer.

ATTENTION shall be asserted whenever there is a change from the synchronized condition (bit 15=0).

TABLE 7-6A MAGNETIC DISK EXTENDED STATUS RESPONSE BITS

7.6.2 Optical Status Response

7.6.2.1 Optical Standard Status

There are 16 bits of status information returned to the controller in response to the Request Standard Status command.

Table 7-7 lists the conditions under which the Status Response bits shall be set and if ATTENTION is asserted.

Table 7-7 is returned when the Command Modifier bits are 0000 and lists the conditions under which the Status Response bits shall be set and if ATTENTION is asserted.

Bit	Status Response	ATT
15	Reserved	0
14	Removable Media Not Present	0
13	Write Protected, Removable Media	0
12	Reserved	0
11	Reserved	0
10	Media Type has Changed (RO, WORM or Eraseable) *2	1
9	1 = Spindle Motor Stopped by Stop command	0
8	1 = Spindle Motor Stopped for Other (e.g. Power On, Reset)	1
7	1 = Power On Condition *1	1
6	Command Data Parity Fault	1
5	Interface Fault	1
4	Invalid or Unimplemented Command Fault	1
3	Seek Fault	1
2	Reserved	0
1	1 = Vendor Unique Status Not Associated with Error Condition	0
0	1 = Vendor Unique Status Available for Error Condition	1
	Write Fault *3	1
	Removable Media Changed (has been changed since last status request)	1

*1 The drive is set to the initial condition that follows the Power On sequence of the drive electronics. See also Magnetic Disk Status.

*2 The type of media loaded in the drive has changed, as well as the cartridge itself e.g. RO to WORM or differently formatted WORM.

*3 Conditions that can cause Write Fault are:

- a) Excessive write current.
- b) Multiple heads selected, no head selected, or improperly selected with WRITE GATE asserted.
- c) DC voltages grossly out of tolerance.
- d) WRITE GATE asserted to a previously recorded or flagged defective sector
- e) Simultaneous assertion of READ GATE and WRITE GATE if DRDW not implemented.

TABLE 7-7 OPTICAL DISK STANDARD STATUS RESPONSE BITS

7.6.2.2 Optical Device Status

There are two words of Device Status information returned to the controller in response to the Request Device Status command.

Table 7-8 is returned when the Command Modifier bits are 1001 and lists the conditions under which the Status Response bits shall be set and if ATTENTION is asserted.

Bit	Status Response	ATT
15	Drive Initialization Failure	1
14	Sensor Failure	1
13	Cartridge Load/Unload Failure	1
12	Spindle Not at Speed Failure	1
11	Focus Failure	1
10	Phase Lock/Tracking Failure	1
9	Tracking Failure	1
8	PLO Failure due to Laser Dying	1
7	Not Track Following	1
6	Not on Correct Track	1
5	Coarse Seek Failure	1
4	Write was Terminated	1
3	Eject Request	1
2	MO (MagnetoOptic) - Erase Encoded:10=Erase 11=Transition	x
1	~ Write 01=Write 00=Off	x
0	Reserved	0

TABLE 7-8 OPTICAL DISK DEVICE STATUS RESPONSE BITS

Table 7-9 is returned when the Command Modifier bits are 1010. Bits 15-8 define the current Track High order byte and bits 7-0 define the current Track Low order byte. NOTE: A drive which implements the Distance/Direction type of Seek only shall respond to this status request with an Invalid Command fault.

Bit	Status Response	ATT
15	Current Track High Order Byte - MSB	0
14		0
13		0
12		0
11		0
10		0
9		0
8	Current Track High Order Byte - LSB	0
7	Current Track Low Order Byte - MSB	0
6		0
5		0
4		0
3		0
2		0
1		0
0	Current Track Low Order Byte - LSB	0

TABLE 7-9 OPTICAL DISK CURRENT TRACK STATUS RESPONSE

7.7 REQUEST CONFIGURATION (0011) (D-0)

This command causes the drive to send 16 bits of configuration data to the controller. The parity utilized in all configuration responses shall be odd.

General Configuration response flags shall be used to verify the validity of the responses.

7.7.1 Disk

The specific configuration requested is specified by command modifier bits 11-8 and possibly by a subscript in bits 7-0.

7.7.1.1 General Configuration Response

If the command modifier bits 11-8 are set to 0000 then the general configuration status information shown in Table 7-11 is returned for the subscript shown.

Command Modifier Bits 11-8	Subscript Bits 7-0	Configuration Response
0 0 0 0	0	Bit Significant Configuration of Drive and Format
		Bit
		0 = Magnetic Disk Drive 15
		1 = Format Speed Tolerance Gap Required 14
		1 = Track Offset Option Available 13
		1 = Data Strobe Offset Option Available 12
		1 = Rotational Speed Tolerance is $>0.5\% < 1.0\%$ 11
		1 = Transfer Rate > 10 Mhz 10
		1 = Transfer Rate > 5 Mhz ≤ 10 Mhz 9
		1 = Transfer Rate ≤ 5 Mhz 8
		1 = Removable Cartridge Drive 7
		1 = Fixed Drive 6
		1 = Spindle Motor Control Option Implemented 5
		1 = Head Switch Time > 15 usec *1 4
		1 = Not MEM *2 3
		1 = Controller Soft Sectoring (Address Mark) 2
		1 = Hard Sectoring (Sector Pulses) 1
		1 = Subscripting Supported 0

*1 Command Complete shall be negated within 15 usec of a head change if this bit is set to 1.

*2 The controller may use the setting of this bit to select an appropriate disk data error correction method. See also 5.4.3.6

TABLE 7-11A MAGNETIC DISK GENERAL CONFIGURATION RESPONSE BITS

Command Modifier Bits 11-8	Subscript Bits 7-0	Configuration Response
0 0 0 0	1	Bit Significant Configuration of Drive and Format
		Bit
		1 = Synchronized Spindles supported 15
		1 = High Speed Data Port (See 7.7.1.3) 14
		1 = Notched Drive with Equal Zones exc first/last 13
		1 = Notched drive with Unequal Zone Sizes 12
		Reserved 11
		to 1
0 0 0 0	x	Values for Configuration of Drive and Format
		8 Transfer Rate in Kilohertz (rounded)
		9 Rotational Speed
		10 Head Group Skew
		Bits 15-8 Reserved
		Bits 7-0 Head Group Skew
		11 Notched Drive Zones Supported (See 7.18)
		Bits 15-0 Reserved
		Bits 7-0 Number of Zones
		12 Notched Drive Characteristics (See 7.18)
		Bits 15-8 Tracks per Zone
		Bits 7-0 Sectors per Track
		13 Notched Drive Characteristics (See 7.18)
		Bits 15-0 Address of First Track in Zone
		14 Write Splice Information
		Bits 15-8 READ GATE Window Size in bits (maximum)
		Bits 7-0 Size of Write Splice in bits (maximum)
		15 Drive Delays
		Bits 15-8 Read Data Propagation Delay in bits
		Bits 7-0 Write Data Delay in bits
		16 Drive Delays
		Bits 15-8 Reserved
		Bits 7-0 Mark Detection Skew

TABLE 7-11B MAGNETIC DISK GENERAL CONFIGURATION RESPONSE BITS

7.7.1.2 Transfer Rate

If the command modifier bits 11-8 are set to 0000, and the subscript is 8 then the data rate information shall be returned. Bits 15-0 shall contain a binary value that specifies the nominal clock frequency of the drive in kilohertz. Non-integral transfer rates shall be rounded up to the next integer value.

This response is required if the drive has reported in Bit Significant General Configuration that it supports a high speed data port. The range of nominal clock frequencies for a low speed data port is from zero to 15 Mhz and for a high speed data port it is from 10 to 24 Mhz. If the drive does not support this command it has a low speed data port, and the data rate is

determined from the General Configuration response settings.

7.7.1.3 Head Group Skew

This value is reported in modulo 256, and reflects the time it takes for a drive to switch between head groups. It need not be implemented if a drive does not support head groups. The value reported by the drive is in the same form and has the same restrictions as the cylinder skew described in Section 7.7.1.5.

7.7.1.4 Write Splice Information

This subcript provides values to assist the controller to manage avoiding the write splice.

7.7.1.4.1 READ GATE Window

This value specifies the maximum time window, in bits, during which the controller may assert READ GATE following a write splice i.e. the maximum time, in bits, from where the write function was previously initiated until READ GATE has to be asserted.

The time is measured from the end of the write splice to the end of the last bit where the controller may assert READ GATE and successfully achieve PLO Sync lock on the drive to assure a successful read operation. If no window is required by the drive, it shall set this value to zero, indicating to the controller that there is no window requirement.

If this value is not returned by the drive, then the controller shall assert READ GATE no later than 16 bit times from where the write function was previously initiated. This is defined by either the assertion of WRITE GATE, or the negation of ADDRESS MARK ENABLE. See also 8.1.2.

IMPLEMENTORS NOTE: This value is intended to support drives which use PLO Sync fields that are split into a 'high frequency' section followed by a 'low frequency' section to speed PLO lock. Drives which do not require this feature shall set this value to zero to allow the controller to increase the PLO Sync size to meet its own needs. The only restriction which then applies is that READ GATE shall be asserted by the controller in sufficient time to guarantee the minimum number of PLO Sync bytes.

7.7.1.4.2 Size of Write Splice

This value is the time, in bits, measured from the assertion of WRITE GATE to the time that it would be safe to assert READ GATE during a subsequent read operation. The minimum time is assumed to be zero bits.

If this value is not returned by the drive, then the controller shall assume that the maximum write splice size is 7 bits. See also 8.2.3.

IMPLEMENTORS NOTE: This value measures the time from WRITE GATE assertion to when the drive write channel is able to write valid data. It should not include any delays associated with getting WRITE DATA bits to the media.

7.7.1.5 Drive Delays

7.7.1.5.1 Read Data Propagation Delay

This value measures the time, in bits, from when the data is first detected at the media to when it appears on the interface (assertion of READ CLOCK).

If this value is not returned by the drive, the controller shall assume that the Read Data Delay is 9 bits. See also 8.1.3.

7.7.1.5.2 Write Data Delay

This value measures the time, in bits, from when the drive latches a WRITE DATA bit from the interface (assertion of WRITE CLOCK) to when that data bit has been recorded on the media.

If this value is not returned by the drive, the controller shall assume that the Write Data Delay is 8 bits.

7.7.1.5.3 Mark Detection Skew

This value measures the time, in bits, that INDEX, SECTOR, or ADDRESS MARK FOUND may vary relative to data recorded on the media i.e. the maximum distance in bits between the earliest possible and the latest possible, detection of INDEX, SECTOR, or ADDRESS MARK FOUND.

This value is a measure of the distance that a mark may move relative to a write splice between the write function and a subsequent read.

If this value is not returned by the drive, the controller shall assume that the Mark Detection Skew is 0 bits.

7.7.1.6 Specific Configuration Response

If command modifier bits 0001-1111 are used, the specific configuration information shown in Table 7-12 is returned for each Request Configuration command with those modifiers.

Command Modifier Bits 11-8	Configuration Response
0 0 0 1	Number of Cylinders - Fixed (FIXCYL)
0 0 1 0	Number of Cylinders - Removable Media (0 if not) (REMCYL)
0 0 1 1	Number of Heads
	Bits 15-8 Removable Drive Heads (REMIDS)
	Bits 7-0 Fixed Heads (FIXHDS)
0 1 0 0	Minimum Unformatted Bytes per Track
0 1 0 1	Minimum Unformatted Bytes per Sector (Hard Sector only)
0 1 1 0	Number of Sectors per Track
	Bits 15-8 Reserved
	Bits 7-0 Sectors per Track
0 1 1 1	Minimum Bytes in ISG Field (not inc InterSector Speed Tol)
	Bits 15-8 ISG Bytes after Index/Sector Pulse to WSplice
	Bits 7-0 Bytes per ISG
1 0 0 0	Minimum Bytes per PLO Sync Field
	Bits 15-8 Reserved
	Bits 7-0 Bytes per PLO Sync Field required when READ GATE is asserted
1 0 0 1	Number of Words of Vendor Unique Status Available
	Bits 15-8 Reserved
	Bits 7-0 Number of Vendor Unique Status Words
1 0 1 0 to 1 1 0 1	Reserved
1 1 1 0	Seek Overhead Skew
	Bits 15-8 Cylinder Switch Skew
	Bits 7-0 Head Switch Skew
1 1 1 1	Vendor Identification (Optional - See Table 7-17)
	Vendor Id - Extended Information (See 7.7.4.2)

Note: Removable heads map from 00 to (REMIDS-1)
 Fixed heads map from REMIDS to (REMIDS+FIXHDS-1)
 Removable cylinders map from 0000 to (REMCYL-1)
 Fixed cylinders map from 0000 to (FIXCYL-1)

TABLE 7-12 MAGNETIC DISK SPECIFIC CONFIGURATION RESPONSE BITS

7.7.1.7 Seek Overhead Skew

This information is supplied by the drive to assist the controller in determining the amount of skew necessary to allow contiguous data transfers. To provide granularity and consistency across all implementations the skew factor is returned as $n/256$ of the rotation time i.e. modulo 256.

Example: A drive with 32 sectors/track spinning at 3,600 rpm with a 5 msec cylinder switch seek time, no overhead on head switch, and a skew value of

77 reported in bits 15-8 ($60/3600 \times 256 = 65.1$ uses per unit of skew and $5/65.1 = 76.8$). The controller uses this to allow an actual skew of 10 ($77 \times 32/256 = 9.6$) sectors between cylinders excluding any controller overhead.

Drives which deassert COMMAND COMPLETE on a head switch need to supply values that include both head switch and cylinder switch skew values, while others need only supply the cylinder switch value. The values reported by the drive are chosen by the manufacturer to suit its anticipated application and make no provision for controller overhead.

7.7.2 Optical

The specific configuration requested is specified by command modifier bits 11-8 and possibly the subscript.

7.7.2.1 General Configuration Response Bits

If the command modifier bits 11-8 are set to 0000 the general configuration status information shown in Table 7-13 is returned as shown for the subscript used.

Command Modifier Bits 11-8	Sub- script 7-0	Configuration Response	Bit
0 0 0 0	0	General Configuration of Drive and Format	Bit
		1 = Not Magnetic Disk	15
		Encoded 2 '010' = Optical WORM Disk	14
		Encoded 1 '011' = Optical Erasable Disk	13
		Encoded 0	12
		Reserved	11
		Reserved	10
		Reserved	9
		Reserved	8
		Removable Cartridge Drive	7
		Reserved	6
		Spindle Motor Control Option Implemented	5
		Head Switch Time	4
		Not MFM	3
		Reserved	2
		Hard Sectoring	1
		1 = Subscripting Supported	0
0 0 0 0	1	General Configuration of Drive and Format	Bit
		Reserved	15
		1 = High Speed Data Port	14
		1 = Notched Drive with Equal Zones exc first/last	13
		1 = Notched drive with Unequal Zone Sizes	12
		Reserved	11
		Seek Alternative Encode 1 '00'=Absolute Seek	10
		Seek Alternative Encode 0 '01'=Distance/Direction	9
		'10'=Absolute and D/D	
		'11'=Reserved	
		Erasable Pointers Supported	8
		Reserved	7
		Defective Sector Identification Supported	6
		Reserved	5
		World Wide ID Implemented	4
		Reserved	3
		Post Field Option Available	2
		Reserved	1
		DRDW (Direct Read During Write)	0
0 0 0 0	x	Values for Configuration of Drive and Format	
	8	Data Rate	
		Bits 15-8 Reserved	
		Bits 7-0 Data Rate in Megahertz (rounded)	
	9	Rotational Speed	

TABLE 7-13 OPTICAL DISK GENERAL CONFIGURATION RESPONSE BITS

7.7.2.2 Specific Configuration Response

If command modifier bits 0001-1111 are used, the specific configuration information shown in Table 7-14 is returned for each Request Configuration command with those modifiers.

Command Modifier Bits 11-8	Configuration Response
0 0 0 1	Number of Tracks on Fixed Media (Erasable only)
0 0 1 0	Number of Tracks per Surface
0 0 1 1	Number of Heads
	Bits 15-8 Removable Drive Heads
	Bits 7-0 Fixed Heads (Erasable only)
0 1 0 0	Number of Recording Bytes per track
0 1 0 1	Number of Recording Bytes per Sector
0 1 1 0	Number of Sectors per Track
	Bits 15-8 Reserved
	Bits 7-0 Sectors per Track
0 1 1 1	Minimum Bytes in ISG Field
	Bits 15-8 Reserved
	Bits 7-0 Bytes per ISG
1 0 0 0	Minimum Bytes per PLO Sync Field
	Bits 15-8 Reserved
	Bits 7-0 Bytes per PLO Sync Field required when READ GATE is asserted
1 0 0 1	Number of Words of Vendor Unique Status Available
	Bits 15-8 Reserved
	Bits 7-0 Number of Vendor Unique Status Words
1 0 1 0	Header Address Information
	Bits 15-8 Format Type (x'01'=ESDI Synthesized Hdr)
	(x'02'=ESDI Raw Header)
	Bits 7-0 Size of Header Address Field
1 0 1 1	Pre-Recorded Data Pointer
	Bits 15-8 MSB of Track Address *1
	Bits 7-0 LSB of Track Address *1
1 1 0 0	Sync Pattern
	Bits 15-8 Number of Bytes in Pattern
	Bits 7-0 Least Significant or only Byte of Pattern
1 1 0 1	Recording Pattern
	Bits 15-8 Number of Bytes in Pattern
	Bits 7-0 Least Significant or only Byte of Pattern
1 1 1 0	Number of Sectors per Track in Concentric Operation *2
1 1 1 1	Vendor Identification (See Table 7-17)
	Vendor Id - Extended Information (See 7.7.4.2)

*1 If no pointer, both bytes shall be set to x'FFFF'

*2 If not equal to value in 0110 then media has spiral tracks

TABLE 7-14 OPTICAL DISK SPECIFIC CONFIGURATION RESPONSE BITS

7.7.4 Vendor Identification Response (Optional)

7.7.4.1 Standard Vendor Information

If a subscript of 0 was issued in the command, bits 15-8 return a binary value which identifies the vendor and bits 7-0 identify the vendor model, revision and any other information the vendor wishes to provide.

The following table lists the identification code assigned to vendors of ESDI drives:

Vendor Identification Code		Vendor Identification Code	
Company Name		Company Name	
00	Vendor Not Identified	10	PRIAM
01	CONTROL DATA/MPi	11	SIEMENS COMMUNICATIONS
02	CYBERNEX-AST	12	TOSHIBA
03	ELECTRONIC PROCESSORS	13	IBM
04	FUJITSU	14	MINISCRIBE
05	HITACHI	15	HEWLETT PACKARD
06	INFORMATION STORAGE	16	RICOH
07	LASERDRIVE	17	MICROSCIENCE INTERNATIONAL
08	MAXTOR	18	VERBATIM
09	MICROPOLIS	19	QUANTUM
0A	MITSUBISHI ELECTRIC	1A	CALIPER
0B	NEC	1B	NAKAMICHI
0C	NEWBURY DATA RECORDING	1C	not assigned
0D	NISSEI SANGYO	1D	not assigned
0E	OPTIMEM	1E	not assigned
0F	PERTEC PERIPHERALS	1F	not assigned

NOTE: Any manufacturer of ESDI products which are not listed above should contact the editor to have an Identification Code assigned.

TABLE 7-17 VENDOR IDENTIFICATION CODES RESPONSE

7.7.4.2 Extended Vendor Information

If the command contains a subscript other than 0 extended vendor unique information can be provided. The subscripts to provide such data shall be contiguous from 1.

7.8 SELECT HEAD GROUP or SELECT TRACK (0100) (D-x)

7.8.1 Disk (SELECT HEAD GROUP)

This optional command causes the drive to select a group of 16 heads. Heads 0-15 are considered as group 0. Bits 7-4 are an encoded binary value used to indicate the group to be selected i.e.

0000 - Head Group 0
 0001 - Head Group 1

 1110 - Head Group 14
 1111 - Head Group 15

The head group selected remains active until reset by issuing this command again to select a different head group. The individual head selected within each group is controlled by the four head select lines.

7.9 CONTROL (0101) (D-0)

7.9.1 Disk

This command causes the control operations specified by bits 11-8 to be performed as shown in Table 7-18.

Command Modifier Bits 11-8	Function
0 0 0 0	Reset Interface attention and Standard Status (Bits 0-11)
0 0 0 1	Reserved
0 0 1 0	Stop Spindle Motor (Optional)
0 0 1 1	Start Spindle Motor (Optional)
0 1 0 0	Reserved
to	
1 1 1 1	

TABLE 7-18 MAGNETIC DISK CONTROL COMMAND MODIFIER BITS

7.9.2 Optical

This command causes the control operations specified by bits 11-8 to be performed as shown in Table 7-19.

Command Modifier Bits 11-8: 7-0	Subs- cript:	Function
0 0 0 0	0	Reset Interface attention, Standard Status and Standard Optical Status (Bits 0-11)
0 0 0 1	0	Reserved
0 0 1 0	0	Stop Spindle Motor
0 0 1 1	0	Start Spindle Motor
0 1 0 0	0	Lock Cartridge (Optional)
0 1 0 1	0	Unlock Cartridge (Optional)
0 1 1 0	0	Eject Cartridge
0 1 1 1	0	Erase Read - Modulation
	1	- Amplitude
	2	- Synchronization
1 0 0 0	0	Magnet Control (MO) - Read
1 0 0 1	0	- Write
1 0 1 0	0	- Erase
1 0 1 1	0	Reserved
1 1 0 0	0	Spiral Operation Off
1 1 0 1	0	Spiral Operation On
1 1 1 x	0	Reserved

NOTE: A change in operation (Spiral On/Off or Off/On) is time dependent based on position relative to end of track at time it is received.

TABLE 7-19 OPTICAL DISK CONTROL COMMAND MODIFIER BITS

7.10 DATA STROBE OFFSET or DATA RECOVERY OFFSET (0110) (D-0)

7.10.1 Disk (DATA STROBE OFFSET)

This optional command causes the drive to offset the data strobe in the direction and amount specified by bits 11-8 as shown in Table 7-21:

Command Modifier Bits 11-8	Subs- cript: 7-0	Function
0 0 0 x	0	Restore Offset to Zero
0 0 1 0	0-255	Early Offset One
0 0 1 1	0-255	Late Offset One
0 1 0 0	0	Early Offset Two
0 1 0 1	0	Late Offset Two
0 1 1 0	0	Early Offset Three
0 1 1 1	0	Late Offset Three
1 x x x	0	Reserved for Disk
1 0 0 0	0-255	Laser Power Adjust Negative on Optical
1 0 0 1	0-255	Laser Power Adjust Positive on Optical
1 0 1 0	0	Reserved for Optical
1 1 1 1	0	Reserved for Optical

TABLE 7-21 DATA STROBE OFFSET COMMAND MODIFIERS

Seek or recalibrate commands restore offsets to zero. Simultaneous Data Strobe and Track offsets are allowed by use of multiple commands. Support of controller definable offset values is optional. If the drive does support subscripting, a value may be used to specify up to 255 instead of being limited to 3 by the separate commands. NOTE: The values of zero and one both have the same effect.

The Laser Power Adjust (Positive or Negative) causes the laser reading power to be changed from nominal power, either in the positive or negative direction, by some incremental amount as specified by the value in the subscript.

7.10.2 Optical (DATA RECOVERY OFFSET)

This optional command causes the drive to offset the data strobe in the direction and amount specified by the command modifier. See Table 7-21. Seek or Recalibrate commands restore offsets to zero.

Disks which can modify gain and use other techniques to recover data shall interpret the modifiers in a drive specific manner. Disks which cannot provide such functions shall ignore this command.

7.11 TRACK OFFSET (0111) (D-O)

This optional command causes the drive to perform a track offset in the direction and amount specified by bits 11-8 as shown below in Table 7-23:

Command Modifier Bits 11-8	Sub- cript: 7-0	Function
0 0 0 x	0	Restore Offset to Zero
0 0 1 0	0-255	Positive Offset One
0 0 1 1	0-255	Negative Offset One
0 1 0 0	0	Positive Offset Two
0 1 0 1	0	Negative Offset Two
0 1 1 0	0	Positive Offset Three
0 1 1 1	0	Negative Offset Three
1 x x x		Reserved

TABLE 7-23 TRACK OFFSET COMMAND MODIFIER BITS

Seek and recalibrate commands restore offsets to zero. Simultaneous Track and Data Strobe offsets are allowed by use of multiple commands.

Drives that implement only one value of offset, Data Strobe or Track, shall respond to unimplemented offset commands as a legal offset function. Support of controller definable offset values is optional. If the drive supports subscribing, a value may specify up to 255 instead of being limited to 3 by the separate commands.

7.12 INITIATE DIAGNOSTICS (1000) (D-O)

This optional command causes the drive to perform internal diagnostics. COMMAND COMPLETE indicates the completion of the diagnostics. ATTENTION with COMMAND COMPLETE indicates that a fault was encountered and status should be requested to determine the proper course of action.

The command modifier bits 11-8 shall be zero to perform the standard diagnostics. Alternatively, these bits may be used by the device to invoke alternate vendor diagnostics. See Table 7-24.

The alternate diagnostic routines shall be numbered in order beginning with x'01', and command reject issued when any unimplemented routines are requested. If the alternate diagnostics are not supported by the device then the command modifier bits may be ignored.

The diagnostic parameter bits may be used to modify the routine per vendor specifications, but all routines shall execute when a default value of zero is present.

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	P
CMD Function				Diagnostic No				Diagnostic Parameter								x

TABLE 7-24 DIAGNOSTIC PARAMETER BITS

7.13 SET UNFORMATTED BYTES/SECTOR (1001) (D-x) (Optional)

This optional command causes the drive to set the number of unformatted bytes per sector indicated in bits 11-0 (if implemented). This command is valid only if the drive is configured to be in the drive hard sector mode. This command is used only if the drive uses a settable counter for the number of bytes per sector and that counter is controllable from the interface.

The unformatted bytes per sector may be adjusted by the drive to meet the drive's requirements. If adjusted, the drive should attempt to retain the number of sectors per track desired by the controller (unformatted bytes per track divided by controller-requested unformatted bytes per sector) but may set a different number if format restrictions on the drive require this.

After setting a new value for Unformatted Bytes per Sector the controller shall re-request the Unformatted Bytes per Sector and Sectors per Track Configuration Responses to verify that the drive is set to the expected values.

7.14 HIGH ORDER SEEK ADDRESS (1010) (x-O)

If Seek Absolute Address is implemented this command is used to define the value of the upper 4 bits of the Seek Address in bits 3-0. This command shall be issued prior to the Seek command (0000) if the high order value in the address has changed since the previous seek was executed. COMMAND COMPLETE is asserted when the high order seek address is valid, has been accepted, and the drive is ready to execute the next command. If the address is invalid e.g. out of range, ATTENTION shall be asserted in conjunction with COMMAND COMPLETE.

If Seek Distance/Direction is implemented this command is used to define the value of the upper four bits of the seek distance in bits 3-0. This command shall be issued prior to the seek command (0000) if the high order distance value has changed, or if the direction has changed, since the previous seek was executed. COMMAND COMPLETE is asserted when the high order seek distance is valid, has been accepted, and the drive is ready to execute the next command. If the seek distance is invalid, i.e. out of range, ATTENTION shall be asserted in conjunction with COMMAND COMPLETE.

Command	Function
Modifier	
Bits 11-8:	
0 0 0 0	Absolute Seek Address
0 0 0 1	Direction towards Track Minimum
0 0 1 0	Direction towards Track Maximum
0 0 1 1	Reserved
to	
1 1 1 1	

TABLE 7-25A HIGH ORDER SEEK ADDRESS MODIFIERS

7.15 FORMAT (1011) (x-0) (Optional)

This command is used to format one track on the disk. It is used only when unformatted media is supported by the drive. Header fields and clocking bits shall be laid down by the drive. The controller is responsible to check the validity of the recorded fields.

This command is optional and manufacturer specific.

7.17 Reserved (1101)

7.18 SET CONFIGURATION (1110) (D-0)

This optional command should be rejected if it is not supported.

The Soft Switch Number provides up to 16 identification values and the Soft Switch Parameter is a modifier. See Table 7-29.

Soft Switch modifiers are used to configure disk drives capable of handling variable frequency recording, sometimes referred to as MCAV (Modified Constant Angular Velocity) or as Notched drives. The disk may be divided into zones of recording frequency that support a different number of sectors per track. If a disk can adapt itself to any frequency it shall do so by reporting x'FF' in response to Request Configuration with Subscript 11.

The controller may use Set Configuration with synchronized drives to set this (the selected) drive to act as master (7-0 = x'01') or as slave to another drive (7-0 = x'00').

The controller may use Set Configuration to identify the zone to be worked with (the first zone is numbered as 1 and begins at cylinder 0).

If the drive is configurable to operate with any valid zoning configuration the controller chooses, the drive shall accept the commands to set the number of Tracks per Zone and Sectors per Track to apply in that zone.

When a drive has been set to operate with a zone, then all information reported is relative to that zone. To find the configuration of a notched drive the controller shall repeat the same procedure for each zone (as it would for a drive which does not support notches) until the command is rejected because there are no more zones.

When set to Zone 0 the drive shall not respond as a notched drive but as a regular drive with only one recording frequency, that of the inner radius.

Soft Switch modifiers are available for the vendor to use as a method of defining configuration information. The implementation of this feature provides users the advantage of reduced installation effort and vendors the advantage of being able to set up automatic testing procedures for different drive configurations.

As an example of the way in which this command may be used by a vendor, the dip switches could be numbered and parameter used to identify how the switches are to be set (1-On, 0-Off). This command would then override any physical position to which the switches are set.

It is recommended that the device be capable of retaining the switch configuration information between power cycles.

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	P				
CMD Function				Switch No				Switch Parameter												
1	1	1	0	0	0	0	0	Vendor Unique												
																to				
				0	x	x	x	Reserved												
				1	0	0	0													
								to				Set Drive as Synchronized Master								
				1	0	1	1													
				1	1	0	0	Notched Drive Zone Number												
				1	1	0	1													
				1	1	1	0	Notched Drive Tracks/Zone												
				1	1	1	1													
				1	1	1	1	Notched Drive Sectors/Track												

Example: If a magnetic disk drive is capable of supporting soft or hard sector operation according to the setting of Dipswitch 1 then the configuration would be described as follows:

1	1	1	0	0	0	0	1	xxxx xx01 = Set Soft Sector								
								xxxx xx10 = Set Hard Sector								

TABLE 7-29 SET CONFIGURATION SOFT SWITCH PARAMETER BITS

7.19 Reserved (1111)

8. READ, WRITE AND FORMAT PARAMETERS

8.1 CRITICAL READ FUNCTION TIMING PARAMETERS (D-x)

Controller variations of the read timing are allowed if the following drive-dependent parameters are met:

8.1.1 Read Initialization Time.

A read operation may not be initiated until (0.7*ISG) following head switch. Drives not able to meet 15 usec head switching time shall drop COMMAND COMPLETE upon a head switch.

8.1.2 READ GATE Timing

READ GATE may not be asserted during a Write Splice area (READ GATE shall be negated one bit time minimum before a Write Splice area and may be asserted one bit time minimum after a Write Splice area).

READ GATE shall be enabled within 16 bit times from where the write function was previously initiated or according to the subscripted values provided in Request General Configuration subscripts. The write function is defined by either:

- o leading edge of WRITE GATE or
- o trailing edge of ADDRESS MARK ENABLE during hard sector format (as defined in Figure 10-14) or soft sector format.

8.1.3 Read Propagation Delay

Data (read) at the interface is delayed by up to 9 bit times from the data recorded on the disk media or according to the subscripted value provided in Request General Configuration subscripts. See manufacturer's drive product specification for exact value if subscripting is not supported.

8.1.4 READ CLOCK Timing

READ CLOCK and READ DATA are valid within the number of PLO sync field bytes specified by the drive configuration after Read Enable and a PLO sync field is encountered. The READ/REFERENCE CLOCK line may contain no transitions for up to two REFERENCE CLOCK periods for transitions between reference and read clocks. The transition period shall also be one-half of a REFERENCE CLOCK period minimum with no shortened pulse widths.

8.1.5 REFERENCE CLOCK Valid Time

The READ/REFERENCE CLOCK lines shall contain valid reference clocks within two REFERENCE CLOCK periods after the negation of READ GATE. Pulse widths shall not be shortened during the transition time but clock transitions may not occur for up to two REFERENCE CLOCK periods.

8.1.6 READ CLOCK Valid Time

The READ/REFERENCE CLOCK line shall contain valid Read clocks within two

Clock periods after PLO synchronization is established. Pulse widths shall not be shortened during the REFERENCE CLOCK to READ CLOCK transition time, but missing clocks may occur for up to two clock periods.

8.2 CRITICAL WRITE FUNCTION TIMING PARAMETERS (D-x)

Controller timing variations in the record-update function are allowed if the following drive-dependent write (and inter-related read) timing parameters are met:

8.2.1 Read-to-Write Recovery Time

Assuming head selection is stabilized, the time lapse from negating READ GATE to asserting WRITE GATE shall be five REFERENCE CLOCK periods minimum.

8.2.2 WRITE CLOCK-to-WRITE GATE Timing

WRITE CLOCKS shall precede WRITE GATE by a minimum of two and a half REFERENCE CLOCK periods.

8.2.3 Write Driver Plus Data-Encoder Turn-On From WRITE GATE

The write driver plus data-encoder turn-on time (write splice width) is between 3 and 7 REFERENCE CLOCK periods or defined according to the subscripted value provided in Request General Configuration subscripts.

8.2.4 Write-Driver Turn-Off from WRITE GATE

To account for data-encoding delays, WRITE GATE shall be held on for at least two byte times after the last bit of the information to be recorded.

8.2.5 Write-to-Read Recovery Time

The time lapse before READ GATE or ADDRESS MARK ENABLE can be asserted after negating the WRITE GATE is defined by the "ISG Bytes after Index/Sector" in Configuration Data Response.

8.2.6 Head Switching Time

WRITE GATE shall be negated at least 1 usec before a head change.

WRITE GATE shall not be asserted until 15 usec after a head change or COMMAND COMPLETE is asserted.

8.2.7 Write Propagation Delay

Write data received at the I/O connector shall be delayed by the write data encoder by up to 8 bit times maximum prior to being recorded on the media.

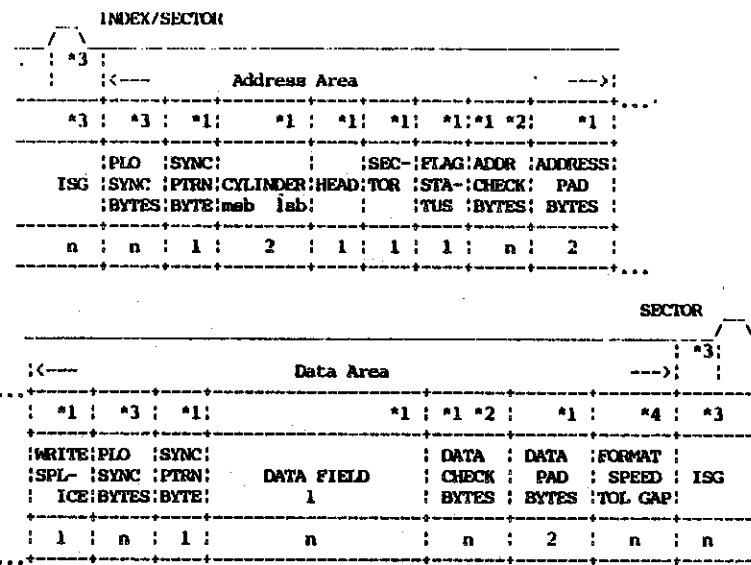
8.3 HARD SECTORED DRIVE (D-K)

The following section is included as an example of a fixed sector format implementation.

8.3.1 Fixed Sector Format Rules

The record format on the disk is under control of the controller. The INDEX and SECTOR pulses are available for use by the controller to indicate the beginning of a track and allow the controller to define the beginning of a sector. A suggested format for fixed data records is shown in Figure 8-1.

The format presented in Figure 8-1 consists of four functional areas; Intersector Gap, Address, Data, and Format Speed Tolerance Gap. The Data area is used to record the system's data files. The Address area is used to locate and verify the track and sector location on the disk where the Data areas are to be recorded. INDEX and SECTOR pulses are provided by the drive.



- *1 These areas are examples only, and may be structured to suit individual customer requirements.
- *2 The number of check bytes is controller-defined.
- *3 Established by device and reported in response to Request Configuration Commands.
- *4 Format Speed tolerance gap is required if REFERENCE CLOCK is not tied to rotational speed. The applicability of this gap is defined in the Configuration data.

NOTE: All byte numbers indicated are minimums.

FIGURE 8-1 FIXED SECTOR FORMAT

8.3.2 Intersector Gap (ISG)

The minimum Intersector Gap size is determined from the configuration data. The Intersector Gap provides a separation between each sector. The gap size is chosen to provide for the following:

- Drive required write-to-read recovery time (the minimum time between negation of WRITE GATE and assertion of READ GATE or ADDRESS MARK ENABLE which is specified in the "ISG Bytes after Index/Sector" in Configuration Data Response.
- Other drive required ISG times.
- Variations in detecting INDEX and SECTOR.

d) Controller decision time between sectors.

All but the last are used by the manufacturer as the basis of determining the "Bytes per ISG" specified in the Configuration Data Response.

8.3.3 Address Area

The address area (Figure 8-1) provides a positive indication of the track and sector locations. The address area is normally read by the controller and the address bytes verified prior to a data area read or write. The address area is normally only written by the controller during a format function and thereafter only read to provide a positive indication of the sector location and establish the boundaries of the data area. The address area consists of the fields described below.

8.3.3.1 PIO Sync Field

These bytes are required by the drive to allow the drive's read-data phase-locked oscillator to become phase and frequency synchronized with the data bits recorded on the media. The controller shall send zeros during this time.

8.3.3.2 Byte Sync Pattern

This field (one byte minimum) establishes byte synchronization (i.e. the ability to partition this ensuing serial bit stream into meaningful information groupings, such as bytes) and indicates to the controller the beginning of the ID information. The Byte Sync Pattern shall contain more than a single one bit to improve the confidence level of detection.

8.3.3.3 ID

These bytes are user-defined and interpreted by the controller. A suggested format consists of five bytes, which allows two bytes to define the cylinder address, one byte to define the head address, one byte to define the sector address, and one byte to define flag status.

8.3.3.4 ADR Check Bytes - (ID Check Codes)

An appropriate error-detection mechanism is generated by the controller and applied to the address for data integrity purposes. These codes are written on the media during formatting. Data integrity is maintained by the controller recalculating and verifying the address field check codes when the ID is read. ADR check bytes are user defined.

8.3.3.5 ADR Pad - (ID Pad)

The ID Pad (two bytes minimum) shall be written by the controller and are required by the drive to ensure proper recording and recovery of the last bits of the address-field check codes. These pad bytes shall be zeros.

8.3.4 Data Area

The data area (Figure 8-1) is used to record data fields. The contents of the data fields within the data area are specified by the host system. The remaining parts of the data area are specified and interpreted by the disk

controller to recover the data fields and ensure their integrity. The data area consists of the fields described below.

8.3.4.1 Write Splice

This area is required by the drive to allow time for the write drivers to turn on and reach recording amplitude sufficient to ensure data recovery. This area shall be allowed for in the format and the controller shall send zeros during this time. The resulting splice from a minimum 2 bit pulse shall be at most one byte long.

Since a write splice is caused by a transition of WRITE GATE, there is no write splice created when ADDRESS MARK ENABLE transitions while WRITE GATE is active (see Figure 10-14 and Figure 10-16).

8.3.4.2 PIO Sync bytes

These bytes are required when reading to allow the drive's phase-locked oscillator to become phase and frequency synchronized with the data bits recorded in the media. The controller shall send zeros during these byte times.

8.3.4.3 Byte Sync Pattern

The Sync Pattern (one byte minimum) establishes byte synchronization and indicates, to the controller, the beginning of the data field. The Byte Sync Pattern shall contain more than a single one bit to improve the confidence level of detection.

8.3.4.4 Data Field

The Data Field contains the information recorded by the controller.

8.3.4.5 Data Check Bytes

The Data Check bytes, commonly referred to as ECC (Error Check Code or Error Correction Code), are generated by the controller and written on the media at the end of the Data Field. Data integrity is maintained by the controller recalculating and verifying the Data Check Bytes or applying error correction algorithms if applicable after the Data Field is read. The Data Check Field is user defined.

8.3.4.6 Data Field Pad

The Data Field Pad bytes shall be a two byte minimum issued by the controller. The field is required by the drive to ensure proper recording and recovery of the last bits of the data field check codes. The controller shall send zeros during these byte times.

8.3.5 Format Speed Tolerance Gap

This gap is required if the Format Speed Tolerance Gap is Required bit is set in the General Configuration Response. If this gap is required, the size is determined by the controller and is:

Unformatted sector length * .01 (if configuration flag bit 11 is 0)

or

Unformatted sector length * .02 (if configuration flag bit 11 is 1)

This gap, if required, shall be between each sector. The byte pattern in this gap shall be zeros.

8.3.6 Fixed Sector PLO Sync Format Timing

The beginning of each PLO Sync field shall be specified by the controller.

For compatibility with Controller Soft Sector mode of operation, the beginning of the Header PLO Sync Field may be specified by the trailing edge of the ADDRESS MARK ENABLE signal when WRITE GATE is asserted. See Figure 10-14 and Figure 10-16.

An alternative method that may be used to define the beginning of the Header PLO sync field is to negate and then assert WRITE GATE immediately prior to the start of the Header PLO sync field. See Figure 10-15.

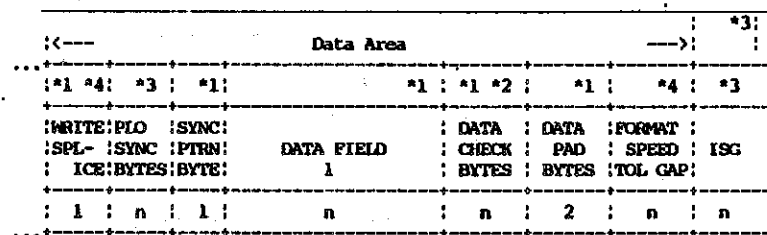
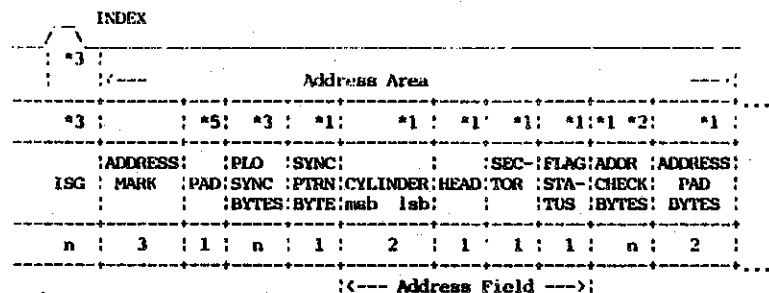
8.4 SOFT SECTORED DRIVE (D-x)

The following section is included as an example of an Address Mark implementation of a soft sector format.

The purpose of the format is to organize a data track into smaller sequentially numbered blocks of data called sectors.

8.4.1 Soft Sector Format Rules

The format shown below in Figure 8-2 is similar to the format commonly used for hard sectored disk drives and indicates minimum requirements.



- *1 These areas are examples only, and may be structured to suit individual customer requirements.
- *2 The number of check bytes is user-defined.
- *3 Established by device and reported in response to Request Configuration Commands.
- *4 Format Speed tolerance gap is required if REFERENCE CLOCK is not tied to rotational speed. The applicability of this gap is defined in the Configuration data.
- *5 This is part of the PLO Sync Field to allow for READ GATE assertion delays. The controller should treat this as an additional byte in the PLO SYNC field.

NOTE: All byte numbers indicated are minimums.

FIGURE 8-2 SOFT SECTORED FORMAT

The soft sectored format has the beginning of each sector defined by an ID Address Mark followed by a prewritten identification (ID) field which contains the sector, cylinder and head address plus flag information. The ID field is then followed by a user supplied data field.

The definitions of the functional areas shown in the soft sectored format are identical to those described for the hard sectored format. There are some additional fields in this format and they are the Address Mark field, Address Mark Pad, and ISG speed tolerance gap.

8.4.2 Address Mark Field

The address mark field is a field 24 +/-1 bits long and is found before the PLO sync field in the address area. The content of this field is drive dependent and is written by the drive when so commanded by WRITE GATE and ADDRESS MARK ENABLE asserted simultaneously.

Detection of Address Mark indicates the location of the beginning of a sector.

8.4.3 Address Mark Pad

The Address Mark Pad byte follows the Address Mark field and is to be considered as an additional byte in the PLO Sync field. It's purpose is to allow for READ GATE assertion delays after detecting the ADDRESS MARK FOUND signal.

8.4.4 Intersector Gap (ISG) Speed Tolerance Gap

The ISG is included in the format to allow for all those items discussed in section 8.3.2. In addition it shall also account for intersector speed tolerance if required.

There may also be an extended ISG at the end of the track if required. This gap is written during a format and is used as a filler and to allow for speed tolerances. This gap shall extend from the end of the last sector on the track up to the leading edge of index. The number of bytes in this gap are not included in the Configuration Data Response returned by the drive.

8.4.5 Soft Sector Format Timing

This timing is mainly to support drives that utilize unique encoding for PLO sync fields. The beginning of each PLO sync field shall be specified by the controller. The beginning of the header PLO sync field shall be specified by the trailing edge of the ADDRESS MARK ENABLE signal when WRITE GATE is asserted. See Figure 10-16 for timing.

8.5 CRITICAL READ FUNCTION TIMING PARAMETERS (x-0)

Controller variations of the read timing are allowed if the following drive-dependent parameters are met:

8.5.1 Read Initialization Time.

A read operation may not be initiated until 15 us following a head change to switch to the other surface.

8.5.2 READ GATE Timing

READ GATE shall be negated during a Write Splice area on an Optical disk.

8.5.3 Read Propagation Delay

Data (read) at the interface is delayed by up to 12 bit times from the data recorded on the disk media. See specific drive product specification for

exact value.

8.5.4 READ CLOCK Timing

READ CLOCK and READ DATA are valid within the number of PLO Sync field bytes specified by the drive configuration after Read Enable and a PLO Sync field is encountered. The READ/REFERENCE CLOCK line may contain no transitions for up to two REFERENCE CLOCK periods for transitions between REFERENCE CLOCK and READ CLOCK. The transition period shall also be one-half of a REFERENCE CLOCK period minimum with no shortened pulse widths.

8.5.5 REFERENCE CLOCK Valid Time

The READ/REFERENCE CLOCK signal shall contain valid REFERENCE CLOCKS within two REFERENCE CLOCK periods after the negation of READ GATE. Pulse widths shall not be shortened during the transition time but clock transitions may not occur for up to two REFERENCE CLOCK periods.

The clock pulses may not always be symmetrical in wave form, but the leading edge correlates to the bit recording rate of the drive. Asserted pulse widths shall not be shortened, but negated pulse widths may be extended.

8.5.6 READ CLOCK Valid Time

The READ/REFERENCE CLOCK line shall contain valid READ CLOCKS within two Clock periods after PLO Synchronization is established. Pulse widths shall not be shortened during the REFERENCE CLOCK to READ CLOCK transition time.

8.6 CRITICAL WRITE FUNCTION TIMING PARAMETERS (x-0)

Controller timing variations in the recording function (read header/write data) are allowed if the following drive-dependent write (and interrelated read) timing parameters are met:

8.6.1 Read-to-Write Recovery Time

Assuming head selection is stabilized, the time lapse from negating Read Gate to asserting WRITE GATE shall be five REFERENCE CLOCK periods minimum.

8.6.2 WRITE CLOCK-to-WRITE GATE Timing

Write Clocks shall precede WRITE GATE by a minimum of two and a half REFERENCE CLOCK periods.

8.6.3 Write-Driver Plus Data-Encoder Turn-On From WRITE GATE

The write driver plus data-encoder turn-on time (write splice width) is between 3 and 7 REFERENCE CLOCK periods.

8.6.4 Write-Driver Turn-Off from WRITE GATE

To account for data-encoding delays, WRITE GATE shall be held on for at least two byte times after the last bit of the information to be recorded.

8.6.5 Write-to-Read Recovery Time

Write to Read recovery time is $< \text{ or } = 10 \text{ } \mu\text{secs}$.

8.6.6 Head Switching Time

WRITE GATE shall be negated at least 1 μs before a head change to switch to the other surface. WRITE GATE may not be asserted until 15 μs after a head change or COMMAND COMPLETE is asserted.

8.6.7 Write Propagation Delay

Write data received at the I/O connector shall be delayed by the write data encoder by up to 12 bit times maximum prior to being recorded on the media.

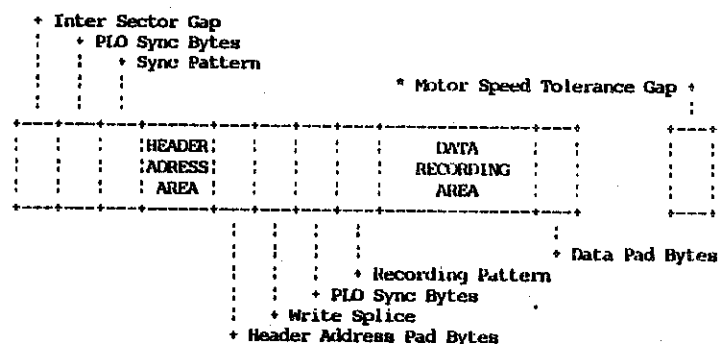
8.7 OPTICAL FORMAT (x-O)

This section uses SECTOR pulses available from the drive and describes a fixed sector implementation.

8.7.1 Format Rules

The record format on the disk is not necessarily under control of the controller, but is either pre-formatted by the media supplier or under control of the drive. There is more than one method for manufacturers to pre-format disks. The INDEX and SECTOR pulses are available for use by the controller to indicate the beginning of a track and sector zero, and the beginning of sectors other than zero respectively.

The controller is responsible for the contents of the available recording area as represented in Figure 8-3.



* Optional - may or may not be present

FIGURE 8-3 FIXED SECTOR OPTICAL FORMAT

A variety of physical formats are supported by optical disk manufacturers. The correlation between format and interface is shown in the appendices.

Fields commonly used by such formats are described below.

8.7.2 Intersector Gap (ISG)

The minimum Intersector Gap size is determined from the configuration data. The Intersector Gap provides a separation between each sector. The gap size is chosen to provide for:

- o Drive required write-to-read recovery time (minimum time between negation of WRITE GATE and assertion of READ GATE).
- o Drive required head switching time (between read-write and write-read).
- o Controller decision making time between sectors.
- o Other drive required ISG times.
- o Variations in detecting INDEX and SECTOR.

8.7.3 Address Area

The address area provides a positive indication of the track and sector locations. The address area is normally read by the controller and the address bytes verified prior to a data area read or write. The address area may be pre-recorded on the media by the manufacturer. It is read by the controller to provide a positive indication of the sector location and establish the boundaries of the data area. The address area typically consists of the following bytes.

8.7.3.1 PLO Sync Field

These bytes may be required by the drive to allow the drive's read-data phase-locked oscillator to become phase and frequency Synchronized with the data bits recorded on the media. The controller shall send zeros during this time.

8.7.3.2 Sync Pattern Byte(s)

This field establishes byte Synchronization (i.e. the ability to partition this ensuing serial bit stream into meaningful information groupings, such as bytes) and indicates to the controller the beginning of the address field information. It is recommended that the Sync Pattern Byte(s) contain more than a single one bit for a greater confidence level of detection.

The drive may synthesize a Sync Pattern, and if it fails to be presented on the first attempt to read a sector, the drive is responsible to recover on the next revolution. The sector pulse preceding a Sync Pattern may or may not have been presented under the circumstances when a drive fails to present Sync Pattern.

8.7.3.3 Address Field

These bytes are user-defined and interpreted by the user's controller.

8.7.3.4 Address Field Check Code

An appropriate error-detection mechanism may be generated and applied to the address field for identification purposes. These codes are typically recorded on the media during formatting. If data integrity is maintained by the drive these fields are not presented over the interface. If they are presented, integrity is maintained by the controller recalculating and verifying the Address Field Check Code when the address field is read.

8.7.3.5 Sector Write Status

If the media has been certified by the manufacturer this field contains zero to indicate the sector has been certified as usable by the manufacturer. If this byte contains any ones, then the sector cannot be recorded because it has been flagged as defective by the manufacturer.

The controller may choose to ignore this field, and consider it part of the gap but it is possible that the drive will abort any attempted write operations if the byte contains any ones.

8.7.3.6 Address Field Pad

The Address Field Pad byte(s) are written by the controller if they are required by the drive to ensure proper recording and recovery of the last bits of the address-field check codes. The pad byte(s) shall be zeros.

8.7.4 Data Area

The Data Area is used to record data fields. The contents of the data fields within the Data Area are specified by the host system. The remaining parts of the Data Area are specified and interpreted by the disk controller to recover the data fields and ensure their integrity. The Data Area typically consists of:

8.7.4.1 Write Splice

The controller shall send zeros during this area.

8.7.4.2 PLO Sync bytes

These bytes may or may not be required when reading to allow the drive's phase-locked oscillator to become phase and frequency Synchronized with the data bits recorded in the media. The controller shall send zeros during these byte times (if any).

8.7.4.3 Recording Pattern Byte(s)

If this field is used in conjunction with PLO Sync bytes it is used to establish byte Synchronization and indicates the beginning of the data field to the controller.

Disks which retain synchronization without the use of PLO Sync require that the Recording Pattern bytes be used to indicate the beginning of the data field to the controller.

Optical disks with pre-recorded data require that the controller recognize

the manufacturer's Recording Pattern Byte(s).

Drives shall provide for the possibility that a Recording Pattern may not be detected, and provision be made for synthesizing the presentation of one.

8.7.4.4 Data Field

The data field contains the host system's data files. A data field may be constructed from segments which contain data (in multiples of 128 bytes) and its associated Error Correction Code (in multiples of 8 bytes). NOTE: The multiples need not be the same value. By using segments to construct large data fields, a simple and straightforward encoder/syndrome calculator can be used instead of having to use the type of more complex ECC needed to handle large fields of contiguous data.

8.7.4.5 Data Field Check Codes

The Data Field Check Codes are generated by the controller and written at the end of each data segment. Data integrity is maintained by the controller recalculating and verifying the Data Field Check Codes, and applying error correction algorithms if applicable on each Data Field segment read.

8.7.4.6 Data Field Pad

The Data Field Pad bytes may not need to be issued by the controller if the ISG is an acceptable period for the drive to ensure proper recording and recovery of the last bits of the data field check codes. If any bytes are required, the controller shall send zeros during these byte times.

8.7.5 Format Speed Tolerance Gap

If this gap is required after the last sector of a track, the byte pattern in the gap shall be zeros.

8.7.6 Post Field Area

A Post Field Area may be used to point to an updated record when the current sector is to be deleted. The Post Field Area consists of:

8.7.6.1 Write Splice

The controller shall send zeros during this area.

8.7.6.2 PLO Sync bytes

These bytes may or may not be required when reading to allow the drive's phase-locked oscillator to become phase and frequency Synchronized with the data bits recorded in the media. The controller shall send zeros during these byte times (if any).

8.7.6.3 Recording Pattern Byte(s)

If this field is used in conjunction with PLO Sync bytes it is used to establish byte Synchronization and indicates the beginning of the data field to the controller.

Disks which retain synchronization without the use of PLO Sync require that the Recording Pattern byte(s) be used to indicate the beginning of the data field to the controller.

Optical disks with pre-recorded Post Fields require that the controller recognize the manufacturer's Recording Pattern Byte(s).

8.7.6.4 Post Field

This field may contain addressing information which can point to a replacement sector.

8.7.6.5 Post Field Check Codes

The Post Field Check Code bytes are generated by the controller and written on the media at the end of the Post Field. Data integrity is maintained by the controller recalculating and verifying the Post Field Check codes or applying Error Correction algorithms, if applicable when the Post Field is read. The Post Check field is user defined.

8.7.6.6 Post Field Pad

The Post Field Pad bytes shall be issued by the controller and are required by the drive to ensure proper recording and recovery of the last bits of the Post Field check codes. The controller shall send zeros during these byte times.

THIS

PAGE

INTENTIONALLY

LEFT

BLANK

9. RECOMMENDATIONS

The control of a device is based on the Command Data transferred serially to the drive and the Configuration and Status data transferred serially from the drive to the controller. This section describes recommended initialization and error recording procedures.

The normal operation is shown in Figure 10-17. Abnormal conditions of the drive are summarized in Figure 10-18 through Figure 10-25.

The following sections are organized as follows:

Section 9.1 defines normal operation.

Section 9.2 defines communication error (fault) handling.

Section 9.3 defines the recommended controller interpretation of configuration data.

Section 9.4 defines error recovery techniques based on the received status data from the drive.

Section 9.5 defines the recommended controller interpretation of standard status.

Section 9.6 defines some miscellaneous integrity points.

9.1 NORMAL OPERATION COMMAND TRANSFER

A command is sent to the drive via the COMMAND DATA line (the Interlocked REQ-ACK handshake protocol is not shown). The drive negates the COMMAND COMPLETE line to signify reception of the first bit of Command Data and asserts COMMAND COMPLETE when the command has been executed in its entirety. COMMAND COMPLETE being asserted with ATTENTION negated signifies that no errors were encountered by the drive during reception or execution of the command.

See Figure 10-17.

For normal command operation, none of the status or configuration data responses need be transferred. This simple interface operation should be easily implemented by both drive and controller manufacturers. The recommended "Fault/Abnormal" condition handling procedure is provided to increase the compatibility between drive and controller manufacturers. None of the recommended procedures violate or take exception to this standard.

9.2 HANDLING COMMAND TRANSFER ERRORS

The intent of this section is to define orderly detection, termination and resynchronization techniques for errors (or faults) which occur as a result of communication of Command and Configuration/Status data. Both the drive actions and the controller actions are defined.

The orderly detection, termination and resynchronization techniques for these faults covered in Section 9.2.2.

Section 9.2.3 defines the next higher level of controller and drive recovery techniques for communications faults.

Communication faults can be grouped into the following three categories:

9.2.1.1 Interface Fault

An interface fault occurs when either the drive or the controller fails to respond to the appropriate edge of the TRANSFER REQ/TRANSFER ACK interlocked handshake protocol within the specified 10 msec limit. With this fault, less than 17 bits of command or status data may be transferred prior to terminating the communications protocol. This type of fault may occur with or without ATTENTION previously asserted.

9.2.1.2 Parity Error

This occurs when either the drive or the controller detects a parity error on appropriate Command or Configuration/Status data. This fault is detected after a complete 17 bits of Command or Configuration/Status data has been transferred. This type of fault may occur with or without ATTENTION previously asserted.

9.2.1.3 Invalid or Unimplemented Command Fault

These types of faults (from a communications viewpoint) could occur at the drive due to an even number of bit errors (from 2 to 16) which were not detected by the "odd parity bit" logic. The drive logic may then decode the resultant command as an Invalid or Unimplemented command.

Note: The drive cannot determine whether an Invalid or Unimplemented command decode is due to an even number of bit errors or an actual Invalid or Unimplemented command sent by the controller. These types of faults may occur with or without ATTENTION previously asserted.

9.2.2 Detection and Resynchronization

In normal operation, the ATTENTION line is negated when the controller first attempts to communicate with the drive using the command protocol (see Figure 10-20 Communications Fault).

For this case (ATTENTION initially negated), the termination and resynchronization technique is that the drive asserts ATTENTION followed by activation of COMMAND COMPLETE. The activation of ATTENTION informs the controller that a fault has been detected by the drive and also requires that the controller terminate the transfer.

The drive should assert COMMAND COMPLETE after ATTENTION is asserted and after it has been determined that the controller has terminated any transfer in process. After COMMAND COMPLETE is asserted the drive and the controller will both be resynchronized, with the drive expecting the first bit of a command transfer (usually a Status Request command) over the interface.

- It is also possible for the ATTENTION line to be asserted (Figure 10-21: Communications Fault) when the controller is attempting to communicate with the drive (i.e. when the controller attempts to send the Read Status Command to the drive to determine the cause of another fault detected in the drive such as a Write Fault, Seek Fault, etc).

If a communications fault occurs when ATTENTION is already asserted, it is not possible for the drive to assert ATTENTION to inform the controller that a communications fault has been detected and to terminate any transfer in process.

For this case, (ATTENTION already asserted); the drive shall create an Interface Fault on the next request for a bit transfer, if applicable, to request that the controller terminate the transfer. Thus, the controller must allow for the possibility that a controller detected Interface Fault could be due to the drive informing the controller that a communications error has occurred.

The controller should respond to a detected Interface Fault by terminating this transfer and waiting for the drive to assert COMMAND COMPLETE. When the drive asserts COMMAND COMPLETE, the drive and the controller shall be resynchronized with the drive expecting the first bit of a command transfer over the interface.

These concepts are further defined in Figure 10-22 through Figure 10-25.

Figure 10-22 and Figure 10-23 define the interface protocol for a drive detected Interface Fault which occurred when TRANSFER REQ was asserted with ATTENTION previously negated and ATTENTION previously asserted respectively.

- In Figure 10-22, the drive signifies detection of the fault by activating ATTENTION. After ATTENTION is asserted the controller should negate TRANSFER REQ and not send another TRANSFER REQ until after COMMAND COMPLETE is asserted.

- In Figure 10-23, the drive does not respond to another TRANSFER REQ after the interface fault is detected and prior to the activation of COMMAND COMPLETE. If the controller sent another TRANSFER REQ after the fault was detected by the drive, the "no response by the drive" would create an Interface Fault at the controller, notifying the controller that the drive had detected a previous Interface Fault from the controller.

- Figure 10-22 and Figure 10-23 illustrate an Interface Fault which resulted from a 10 msec time out measured from the leading edge of TRANSFER ACK. An Interface Fault would also occur if the controller fails to send (request) the next bit of data within 10 msec (measured from the trailing edge of TRANSFER ACK). For this Interface Fault condition the notes and timing of Figure 10-22 and Figure 10-23 are also applicable except that both TRANSFER REQ and TRANSFER ACK are already negated when the Interface Fault condition is detected.

- Figure 10-24 and Figure 10-25 define the interface protocol for either a Parity Error, Invalid Command or Unimplemented Command fault with ATTENTION previously negated and ATTENTION previously asserted. In Figure 10-24 and Figure 10-25, the error occurred on a command sent by the controller which would result in status or configuration data being sent back by the drive. Thus, the controller is expecting to assert TRANSFER REQ for the first bit of Configuration/Status data after sending the parity bit of the command data.

- In Figure 10-24, the drive asserts ATTENTION to signify that the fault has been detected. However, the drive cannot guarantee whether TRANSFER REQ is asserted or negated when ATTENTION is asserted. In addition, for this type of fault, ATTENTION will probably be asserted after the TRANSFER ACK, for the parity bit of the command data has been negated. Thus it is the controllers responsibility to monitor the ATTENTION and COMMAND COMPLETE lines to insure that correct resynchronization occurs per Figure 10-24.

- In Figure 10-25, ATTENTION was previously asserted and cannot be used to inhibit the TRANSFER REQ for the first bit of Configuration/Status data (or terminate this TRANSFER REQ). For the case shown in Figure 10-25, the key assumption made by the drive is that the controller will assert TRANSFER REQ for the first bit of Configuration/Status data within 10 msec after the trailing edge of the TRANSFER ACK for the parity bit of the Command.

- If the Controller has not requested the first bit of Configuration/Status data within this 10 msec time out, the drive shall assert COMMAND COMPLETE. If the controller asserts TRANSFER REQ for the first bit of Configuration/Status data within this 10 msec, the drive shall create an Interface Fault to the Controller which will then signify that resynchronization is to occur. (The drive should also not respond to any subsequent TRANSFER REQs which occur within 10 msec of the last TRANSFER REQ for this sequence due to the 10 msec timeout denoted by note *4 in Figure 10-25).

- If a communication fault, as defined by Figure 10-24 or 10-25, occurred on a command which did not expect status or configuration data to be transferred back to the controller, the controller should not issue a TRANSFER REQ for the first bit of Configuration/Status data. For these Communication fault conditions, the notes and timing of Figure 10-24 and Figure 10-25 are applicable except that TRANSFER REQ shall not be asserted by the controller for the first bit of Configuration/Status data as defined at *3 in Figure 10-24 and 10-25.

9.2.3 Recovery Protocol

This section defines the controller and drive recovery actions necessary to recover from communications faults. This section assumes that a communications fault was the only type of fault that has occurred. (Multiple fault recovery is covered in Section 9.4).

- The recovery protocol for handling communications faults is summarized in Figure 10-26 (Communications Faults - Recovery Protocol). Figure 10-26 explanation is as follows:

- a) Assume a communications fault occurred on the original command transfer

(and was detected and resynchronized per Section 9.2.2), and ignore possibilities *1, *3 or *6. For this case, the controller should request status, obtain the status from the drive, observe that only a single communications fault bit was set, reset attention and then resend the original command.

b) Assume a communications fault occurred on the original command transfer but also allow possibility *1 to occur. For this case, a communications fault is also going to be detected on the Request Status Command. This shall result in a communications fault bit being set in the current drive status response for the Request Status Command as well (the communications fault bit set for the Request Status Command may or may not be the same bit that was set for the original command). If possibility *1 occurs, the controller should recover by reissuing the Request Status command and then proceeding as in case a. Note - the controller should remember that a retry was required (see Section 9.4); however, if only Communications fault bits are set, the recovery mode does not change.

c) Assume a communication fault occurred on the original command transfer but also allow possibility *3 to occur. For this case, a communications fault is also detected on the original status read from the drive (and therefore uninterpretable by the controller). An additional communications fault bit should be set in the drive's status response if an Interface Fault was detected by the drive. (No status bit should be set if the Controller detected a parity error fault). If possibility *3 occurs, the controller shall reissue the Request Status Command in order to obtain the status. Subsequent recovery is the same as case b.

d) Assume a communications fault occurred on the original command transfer but also allow possibility *6 to occur. For this case, the controller received the original status response from the drive (which was assumed to have only a single communications fault bit set) but the drive detected a communications error on the Reset Attention Command and was unable to reset the ATTENTION line and Standard Status bits. The controller should first attempt to resend the Reset Attention command at least 3 times before proceeding as shown in Figure 10-26.

If the ATTENTION line is negated with the reissued Reset Attention Command, the controller may then reissue the original command.

e) Assume all combinations of Figure 10-26.

9.3 CONFIGURATION DATA

Configuration data should be read by the controller whenever a drive is first powered on. After a drive has completed a power turn on sequence, the drive shall assert the ATTENTION line and also set the "Power On Condition Exists" bit in the Standard Status response.

9.3.1 Interpreting Magnetic Disk General Config Data

9.3.1.1 Bit 15: Not a Magnetic Disk Drive

If this bit is a zero, the device is a magnetic disk drive acting in accordance with the standard, otherwise this bit is a one.

9.3.1.2 Bit 14: Format Speed Tolerance Gap Required

The interpretation of this bit is dependent upon the sectoring mechanism which is defined by bits 1 or 2 of the General Configuration Response.

If Bit 14 is a zero, the required format speed tolerance gap size is zero bytes between sectors. If Bit 14 is a one, then a format speed tolerance gap is required between each sector with a minimum number of bytes as defined in Section 8.3.5 for hard sectored drives and Section 8.4.5 for soft sectored drives.

Note: On soft sectored drives the controller should format the last sector on the track and then send format data of zeros until the Index Pulse is received. At physical end of track, zero bytes (in addition to the gap of the last sector) shall result if the disk is rotating at the upper limit (disk rotating fast) but if the disk is rotating more slowly then the gap in the last sector shall be extended to include more bytes.

9.3.1.3 Bit 13: Track Offset Option Available

If Bit 13 is a zero, the controller should not attempt to utilize the "Track Offset" command as an extended retry method of recovering data from the media which cannot be read without a track offset. If Bit 13 is a zero and a "Track Offset" command is received, the drive should respond with an "Invalid or Unimplemented Command Fault." If Bit 13 is a one, the controller should utilize the "Track Offset" command as an extended retry method of recovering data from the media which cannot be read without a track offset.

9.3.1.4 Bit 12: Data Strobe Offset Option Available

Same definition as Bit 13 with the words "Track Offset" replaced by "Data Strobe Offset."

9.3.1.5 Bit 11: Rotational Speed Tolerance is $> 0.5\%$ and $< 1.0\%$

Typically, this bit is only interpreted by the controller if the "Format Speed Tolerance Gap Required" (Bit 14) bit is a 1. See Bit 14 definition for the controller's interpretation of Speed Tolerance.

If Bit 11=1, the rotational speed plus data clock rate tolerance is $< 0.5\%$;

If Bit 11=0, the rotational speed plus data clock rate tolerance is $> 0.5\%$ and $< 1.0\%$.

9.3.1.6 Bit 10: Transfer Rate > 10 MHz

If Bit 10 is a 1, the disk drive NRZ Data Transfer Rate is greater than 10 MHz. Per the current specification, the rate is > 10 MHz and < 15 MHz. Otherwise, Bit 10 is a zero. If this bit is a 1, the controller should be capable of operating with up to at least a 15 MHz Data Rate.

9.3.1.7 Bit 9: Transfer Rate > 5 MHz, < 10 MHz

If Bit 9 is a 1, the disk drive NRZ Data Transfer Rate is greater than 5 MHz but less than or equal to 10 MHz. Otherwise, Bit 9 is a zero. If this bit is a 1, the controller should be capable of operating with up to at least a 10

MHz Data Transfer Rate.

9.3.1.8 Bit 8: Transfer rate <5 MHz

If Bit 8 is a 1, the disk drive NRZ Data Transfer Rate is equal to or less than 5 MHz. Otherwise, Bit 8 is a zero. If this bit is a one, the controller should be capable of operating with up to at least a 5 MHz data transfer rate.

9.3.1.9 Bit 7: Removable Cartridge Drive

If Bit 7 is a 1, the disk drive is capable of containing removable media. If Bit 7 is a 0, the disk drive is not capable of containing removable media. If Bit 7 is a 1, all status bits and configuration bits for removable media are applicable and should be interpreted by the controller. If Bit 7 is a 0, none of the status bits or configuration bits for removable media are valid and should be logic zeros if read from the drive.

When checking for removable media capability, the controller should also check if fixed media is present in order to determine the appropriate head address(es) encoding to reference the removable and the fixed disk(s) in the drive. Head addressing is continuous from 0 through both removable and fixed disks. If removable media is present, head 0 should be on the removable media drive. (The number of heads, both fixed and removable, is contained in the "Number of Heads" General Configuration response.

9.3.1.10 Bit 6: Fixed Drive

If Bit 6 is a 1, the disk drive contains fixed media (whether the drive also contains removable media is defined by Bit 7). If Bit 6 is a 1, all status bits and configuration bits for fixed media are applicable and shall be interpreted by the controller. If Bit 6 is a 0, none of the status or configuration bits for fixed media are valid and should be zeros if read from the drive. Also, if Bit 6 is a zero, then either Bit 15 (not a magnetic disk drive) or Bit 7 (removable cartridge drive) should be set to indicate the device type.

9.3.1.11 Bit 5: Spindle Motor Control Option Implemented

If Bit 5 is a 0, the drive should automatically turn the spindle motor on upon power up. The Spindle Motor Stopped status bit (Bit 9) indicates the current state of the spindle motor when the Standard Status bits are read from the drive. The Control commands with command modifiers 0010 (Stop Spindle Motor) and 0011 (Start Spindle Motor) are invalid. The Invalid or Unimplemented Command Fault status bit (Bit 5) should be set if the motor control commands are sent to the drive.

If Bit 5 is a 1, the controller is responsible for turning on the spindle motor with the Start Spindle Motor command. The Spindle Motor Stopped status bit (Bit 9) indicates the current state of the spindle motor when the standard status bits are read from the drive. The Stop Spindle Motor Command is an optional command if the Spindle Motor Control Option Implemented bit (Bit 5) is a 1. The Invalid or Unimplemented Command Fault status bit (Bit 5) should be set if the Stop Spindle Motor command is received but is not implemented in the drive.

9.3.1.12 Bit 4: Head Switch Time >15 Usec

This bit shall be set to zero if the head switch time is <15 usec. COMMAND COMPLETE shall not be affected by changes in the Head Select lines (HEAD SELECT 2(0), 2(1), 2(2) and 2(3)) if Bit 4 is a zero.

If Bit 4 is a 1, the time required to switch heads is greater than 15 usec. The drive should negate COMMAND COMPLETE within 15 usec. from a change of state of the HEAD SELECT lines and should assert COMMAND COMPLETE when the head change has been performed.

Note: If a head change is not desired when a unit is selected or deselected, then the HEAD SELECT lines should contain the desired head address as previously selected at least one microsecond prior to the reselection of the unit and one microsecond after the UNIT SELECTED signal is negated for a unit deselect sequence.

9.3.1.13 Bit 3: Not MFM

If Bit 3 is a 1, then the drive does not employ an MFM encoding method to write data on the media. If Bit 3 is a 0, an MFM encoding method is used. This bit state may be used by controllers to select Disk Data Error correction methods (The utilization of encoding other than MFM may result in larger data error burst lengths). If this bit is a 1 the controller shall ensure that the start of a PIO sync field is defined.

9.3.1.14 Bit 2: Controller Soft Sector (Address Mark)

If Bit 2 is a 1, then the drive format mode is Controller Soft Sector (Address Mark). The controller should format the drive according to the guidelines in Section 8.4. The ADDRESS MARK ENABLE and ADDRESS MARK FOUND signals are valid and are described in Section 6.2.7 and Section 6.3.7 respectively. Configuration Data of Unformatted Bytes per Sector and Sectors per Track are not applicable (and are undefined). The Set Bytes Per Sector command is an Invalid command when the drive is operating in the Soft Sector mode.

Bits 0, 1 and 2 of the Configuration Response bits returned from the Request Configuration command are mutually exclusive.

9.3.1.15 Bit 1: Drive Hard Sector (Sector Pulses)

If Bit 1 is a 1, then the drive format mode is Drive Hard Sector (Sector Pulses) and the controller should format the drive according to the guidelines in Section 8.3. Request Configuration commands with command modifiers 0101 (unformatted bytes per sector) and 0110 (sectors per track) are valid commands and the drive should respond with the current sector size information. The sector size may be modified with the Set Bytes Per Sector command if implemented by the drive.

Bits 1 and 2 of the Configuration Response bits returned from the Request Configuration command are mutually exclusive.

9.3.1.16 Bit 0: Reserved

This bit shall always be set to zero.

9.3.2 Magnetic Disk Commands Determined Invalid by Config Data

In addition to defining the drive configuration which allows the controller to self configure for the specific drive attached, the contents of the General Configuration response should also be utilized to define certain commands as invalid for this drive. These invalid commands should not be issued to the drive by the controller.

Commands which may be invalid for a drive depending on the state of General Configuration response bits are defined as follows:

- a) Magnetic disk specific commands are not valid for an optical drive, and visa versa, Status and Configuration Request commands are valid for all drives.
- b) Data Strobe Offset and/or Track Offset commands are not valid if the drive does not implement the applicable option.
- c) The Start and Stop Spindle commands are not valid if spindle motor control ability is not implemented by the drive. Even if spindle motor control is available in the drive, the Stop Spindle command may or may not be a valid command. In this case the controller can only determine the validity of the Stop Motor command by attempting to issue this command to the drive.
- d) The Set Bytes Per Sector command is not valid if the drive's sector mechanism is Address Marks. The Set Bytes Per Sector command may or may not be valid if the drive's sector mechanism is Sector Pulses. In this case, the controller can only determine the validity of this command by attempting to issue the Set Bytes Per Sector command.
- e) The use of subscripts other than zero is valid only if the drive has reported that it supports subscripting. All features requiring subscripted data are invalid if subscripting is not supported.

9.4. Recommended Controller Responses to Status Conditions

Status data should be requested by the Controller whenever the ATTENTION line is asserted. This section covers the general philosophy for handling faults, I/O control signal status, status bit interpretation, and command utilization for error recovery.

9.4.1 General Error Recovery Philosophy

- As indicated in Figure 10-18, Figure 10-19 and Figure 10-26, the normal procedure to recover from a fault condition detected by the drive (indicated by ATTENTION being asserted), is for the controller to first issue a Request Status Command, read the Standard Status response and then issue a Reset Attention Command to reset the fault condition in the drive.

The only notable exception to this rule occurs when the Vendor Unique Status Available bit is set and the controller desires to read the Vendor Unique Status. In this case, the controller should read the Vendor Unique Status responses prior to issuing the Reset Attention Command because the Reset Attention Command shall also reset the Vendor Unique Status responses (these responses are reset by virtue of resetting Standard Status bits 11-0 with the Reset Attention Command).

The general philosophy is that a drive should recover from a fault state by attempting recovery techniques when the Reset Attention Command is received (i.e. the drive shall attempt to become operational if possible). In some cases, depending on the status bits which were set, additional recovery commands may be required by the controller (e.g. a command to eliminate a track offset condition, turn on the spindle motor). In addition, the drive should utilize an asserted ATTENTION line to inhibit reads or writes if an internal fault condition exists.

Once the controller has determined that a fault condition exists (i.e. ATTENTION asserted) the controller should terminate the current operation and properly read and reset the Standard Status. After receiving a valid Standard Status response the controller may then proceed to decode the status bits and perform error recovery based on status bits set.

If the controller cannot recover from the fault detected after three retries, the controller should proceed by reading Vendor Unique Status (if available) and reporting this status to the system. When reading Vendor Unique Status, the controller should read all available responses, as specified in configuration data, in sequence starting with the first location.

If a Parity Error, Interface Fault, or Invalid Command fault should occur while reading Vendor Unique Status, the controller should recover from the error and continue reading at the location where the error occurred.

9.4.2 I/O Status for Error Recovery

Table 9-2 shows all possible states of the I/O control signals and the subsequent valid commands which may be sent by the controller for each of these states.

Table 9-2 reveals that no command transfer may be initiated by the

controller if either TRANSFER REQ or TRANSFER ACK is asserted, or if COMMAND COMPLETE is negated.

Table 9-2 also reveals that the READY line is utilized to indicate that the spindle motor is up to speed. The READY line being negated does not necessarily indicate a fault condition in the drive. For example, The spindle may have been stopped by a previous command.

The applicability of the potentially valid commands listed in Table 9-2 are explained more fully in Section 9.5.

COMMAND		XFER		XFER	
READY	COMPLETE	ATTN	REQ	ACK	
SIGNAL CONDITION					VALID COMMANDS
x	x	x	0	1	None
x	x	x	1	0	None
x	x	x	1	1	None
0	0	1	0	0	None
0	1	0	0	0	Any Status Request command *, Start Spindle **, Stop Spindle **, or Initiate Diagnostics
0	1	1	0	0	Any Status Request command *, Reset Attention, or Initiate Diagnostics
1	0	0	0	0	None
1	0	1	0	0	None
1	1	0	0	0	Any command *, **
1	1	1	0	0	Any Status Request command *, Reset Attention, or Initiate Diagnostics

x - May be either 0 or 1

* Vendor Unique status is only valid if Vendor Unique Status Available bit in the Standard Status Response is set to 1.

** If appropriate Configuration Data bit is asserted.

NOTE: A Seek, Read or Write cannot be initialized:
if READY (defined only as the spindle motor is up to speed) or
COMMAND COMPLETE is negated or
if ATTENTION is asserted.

TABLE 9-2 I/O CONTROL SIGNAL STATES FOR ERROR RECOVERY

9.5 STANDARD STATUS

Standard Status bits 15-12 are status bits which do not cause ATTENTION to be asserted when set. The controller should check these bits as a potential reason for a fault being set by the drive in response to a controller attempted operation or invalid read data.

Standard status bits 11-0 are set to 1 and cause ATTENTION to be asserted to advise when a fault condition or change of status occurs that the controller may not be aware of.

9.5.1 Interpreting Magnetic Disk Standard Status

9.5.1.1 Bit 15: Reserved

This bit shall always be set to zero.

9.5.1.2 Bit 14: Removable Media Not Present

This bit shall be set by the drive to indicate that the removable media has been removed. The controller may not seek, read, or write to this device until removable media is in place. If the controller attempts to perform a write operation on removable media while this bit is set, Write Fault should occur. A seek attempt shall result in an Invalid Command status.

9.5.1.3 Bit 13: Write Protected Removable Media

This bit should be set by the drive to indicate that the removable media has been write protected. The controller may not write to this device if this bit is set. If the controller attempts to perform a write on removable media while this bit is set, a Write Fault should occur.

9.5.1.4 Bit 12: Write Protected, Fixed Media

This bit shall be set by the drive to indicate that the fixed media has been write protected. The controller may not write to this device if this bit is set. If the controller attempts to perform a write on fixed media while this bit is set, a Write Fault should occur.

9.5.1.5 Bit 11: Reserved

This bit shall always be set to zero.

9.5.1.6 Bit 10: Reserved

This bit shall always be set to zero.

9.5.1.7 Bit 9: Spindle Motor Stopped

This bit is set to indicate that the spindle motor is not up to speed. A change of state on this bit due to a Start Spindle Motor Command or a Stop Spindle Motor Command should not cause ATTENTION to be asserted (i.e. any command which changes the state of this bit should not cause ATTENTION to be asserted if the command is successfully executed). The controller should depend on COMMAND COMPLETE and READY with ATTENTION negated to determine when the Start Spindle Motor or Stop Spindle Motor command has been executed.

properly.

If this bit is set, the controller should verify that the spindle motor control option is implemented before attempting to send the Start Spindle Motor Command. Once the controller has sent the Start Spindle Motor Command, the controller should wait for COMMAND COMPLETE and READY to be asserted with ATTENTION negated to indicate that the spindle motor is up to speed. A timing diagram for a successful spindle start sequence is shown in Figure 10-27.

9.5.1.8 Bit 8: Power On Reset Conditions Exist

This bit is set after a Power On Sequence occurs or when an internal drive fault condition occurs which would cause an internal Power On Reset operation to occur (e.g. a momentary loss of a DC voltage). This bit is intended to notify the controller that the drive has reinitialized itself to the default drive characteristics (e.g. default number of bytes per sector).

If this bit is set, the controller should read the configuration responses to determine if the default drive characteristics are the desired operating conditions and/or send the required commands to the drive to set up the desired drive operating conditions. The controller should also check the Spindle Motor Stopped bit in the status response and perform a Motor-On sequence if needed.

9.5.1.9 Bit 7: Command Data Parity Fault

This bit should be set by the drive when a parity error is detected on the command data received from the controller. The drive shall not attempt execution of a command if a parity error is detected.

9.5.1.10 Bit 6: Interface Fault

This bit shall be set when an interface protocol timeout is detected by the drive.

9.5.1.11 Bit 5: Invalid or Unimplemented Command

This bit is set by the drive when any of the following commands are received:

- A Command Function (bits 15-12 of the command) defined as Reserved by the standard.
- A legal Command Function plus a Command Modifier defined as reserved in the standard.
- A legal command which is not implemented by the drive.
- A legal command with an invalid command parameter (e.g. illegal seek address)
- A legal command which the drive does implement, but is not able to execute at this time (e.g. a Seek command while the spindle motor is stopped).

If Bit 5 is set, the controller should verify configuration for the command being used.

9.5.1.12 Bit 4: Seek Fault

This bit should be set for any reason the drive suspects that a seek function has not been executed properly. If this bit is set, the controller should reset Standard Status, perform a calibrate and reissue the Seek command a minimum of three times prior to defining that the drive may require operator intervention. This bit is intended to indicate that a transient (or potentially solid) problem exists with the drive positioning system. This bit should not be set for an invalid seek address received from the controller.

9.5.1.13 Bit 3: Write Gate With Track Offset

This bit should be set when WRITE GATE is asserted with a non-zero Track Offset set in the drive. If this bit is set, the controller should restore the Track Offset to zero and re-perform the write operation.

9.5.1.14 Bit 2: Vendor Unique Status Available

This bit should be set by the drive to indicate that vendor unique status is now valid. The controller is not required to read this vendor unique status. Vendor Unique Status is only intended to be utilized by trained field maintenance personnel and is not intended to be interpreted by disk controllers or operating systems. The Vendor Unique Status responses should be read by the Controller and sent back to the system for a print out of these codes which may then be utilized by the field maintenance personnel.

9.5.1.15 Bit 1: Write Fault

This bit is set to indicate that a Write Fault condition has occurred in the drive and that further writing shall be inhibited until the fault is cleared. The drive should assert the ATTENTION line in real time when the Write Fault condition is detected. The controller should inhibit writing once ATTENTION is detected in order to ascertain on which sector the Write Fault occurred.

9.5.1.16 Bit 0: Removable Media Changed

This bit should be set by the drive to indicate to the controller that the removable media may have been changed.

9.5.2 Error Recovery Commands for Magnetic Disk

Table 9-3 lists the applicable commands and a recommended order for the controller to issue these recovery commands. It is assumed that a single status bit is set in the Standard Status response. Table 9-4 lists the priority in interpreting multiple status bits set in the Standard Status response.

In Table 9-3, numbers enclosed in circles denote the order in which the indicated recovery command should be issued. In Table 9-3, numbers identified by asterisks are notes and not associated with the order in which recovery commands are issued. The Standard Status Bits 12, 13 and 14 indicate conditions in the drive that are probably not changeable via the

Condition III is basically the same as Condition II, except that ATTENTION was asserted not as a result of a command transfer to the drive, thus a communications fault was not the reason for the ATTENTION line being asserted.

		15	14	13	12	11	10	9	8
Op Code	Command	STANDARD STATUS RESPONSE							
0	Seek								
1	Recalibrate								
4	Select Head Group								
5(0)	Reset Attention							#1	#1
5(2)	Stop Spindle								
5(3)	Start Spindle							#2	
6(0,1)	Data Strobe Offset-0								
6(Y)	Data Strobe Offset								
7(0,1)	Track Offset-0								
7(Y)	Track Offset								
8(x)	Initiate Diagnostics							#3	
9	Set Bytes per Sector								#4
	Status								
2(0)	General Status								
2(x)	Vendor Unique Status								
3(0)	General Configuration								#2
3(x)	Specific Configuration								#3
	Other Actions		#1	#1	#1			#2	#3

- #n - Priority of command usage.
- *1 - Manual intervention probably required
- *2 - Manual intervention required if command not supported by drive.
- *3 - Controller may require reinitialization and/or special interpretation.
- *4 - Manual intervention probably required if controller cannot reset fault.
- *5 - Any one of these commands will restore Track Offset to zero.

TABLE 9-3A COMMANDS USED TO RECOVER FROM REPORTED STATUS

		Removable Media Changed:							
		Write Fault:							
		Vendor Unique Status:							
		Write Gate with Track Offset Fault:							
		Seek Fault:							
		Invalid or Unimplemented Command Fault:							
		Interface Fault:							
		Command Data Parity Fault:							
		7	6	5	4	3	2	1	0
Op Code	Command	STANDARD STATUS RESPONSE							
0	Seek				#3	#2*5			
1	Recalibrate				#2	#2*5			
4	Select Head Group								
5(0)	Reset Attention	#1	#1	#1	#1	#1	#2	#1	#1
5(2)	Stop Spindle								
5(3)	Start Spindle								
6(0,1)	Data Strobe Offset-0								
6(Y)	Data Strobe Offset								
7(0,1)	Track Offset-0					#2*5			
7(Y)	Track Offset								
8(x)	Initiate Diagnostics				#4	#3		#2	
9	Set Bytes per Sector								
Status									
2(0)	General Status								
2(x)	Vendor Unique Status					#1			
3(0)	General Configuration								
3(x)	Specific Configuration								
Other Actions				*3	*4	*3		*4	*3

- #n - Priority of command usage.
 *1 - Manual intervention probably required
 *2 - Manual intervention required if command not supported by drive.
 *3 - Controller may require reinitialization and/or special interpretation.
 *4 - Manual intervention probably required if controller cannot reset fault.
 *5 - Any one of these commands will restore Track Offset to zero.

TABLE 9-3B COMMANDS USED TO RECOVER FROM REPORTED STATUS

		Power On Condition:							
		Spindle Motor Stopped by Stop command:							
		Removable Media Changed:							
		Write Fault:							
		Vendor Unique Status:							
		Write Gate with Track Offset Fault:							
		Seek Fault:							
		Invalid or Unimplemented Command Fault:							
		Interface Fault:							
		Command Data Parity Fault:							
		9	8	7	6	5	4	3	2
		1	0						
Attention Set During Command Sequence									
Status Read Successfully On First Attempt		5	4	1	2	3	6	n	n
Attention Set During Command Sequence									
Status Read Successfully On First Attempt		2	1	*8	*9	*10	3	n	n
Attention Set While COMMAND COMPLETE Asserted		2	1	n	n	n	3	4	n

NOTE: Standard Status bits 15-10 do not cause ATTENTION to be asserted.
 1 = Highest Priority 10 = Lowest Priority
 n = Bits that may be tested in any order after those with priority.
 * = Bits that should be tested with the priorities indicated after all other bits have been determined NOT to be the cause of ATTENTION being asserted.

TABLE 9-4 STANDARD STATUS BIT PRIORITY

9.6 MISCELLANEOUS

The controller should guarantee that a drive is deselected and WRITE GATE is negated if the controller is powered up or down while power is still applied to the drive.

The drive should not destroy data on the media if the drive is powered up or down while power is applied to the controller and the controller is not activating WRITE data signals.

The drive need not implement special logic to detect the deselection of a drive prior to completion of a Command or Configuration/Status transfer. The defined communication fault protocol resynchronization features should detect this condition if the drive is adversely affected.

THIS

PAGE

INTENTIONALLY

LEFT

BLANK

160



10. TIMING

10.1 DESKEWING

The controller shall provide cable deskewing for all signals originating from the device. The device shall provide for cable deskewing for all signals originating in the device.

10.2 SYMBOLS

Certain symbols are used in the timing diagrams. These symbols and their respective definitions are listed below.

- / or \ - signal transition (asserted or negated)
- < or > - data transition (asserted or negated)
- XXXXXX - undefined but not necessarily released
- . . . - the "other" condition if a signal is shown with no change
- #n - used to number the sequence in which events occur e.g. #a, #b
-  - a degree of uncertainty as to when a signal may be asserted
-  - a degree of uncertainty as to when a signal may be negated
- T - Nominal Clock Period
- 1/T - Bit Time

Within each figure the timing terms i.e. tA, tB etc are repeated. There is no continuity of definition of tA from one figure to another.

10.3 TERMS

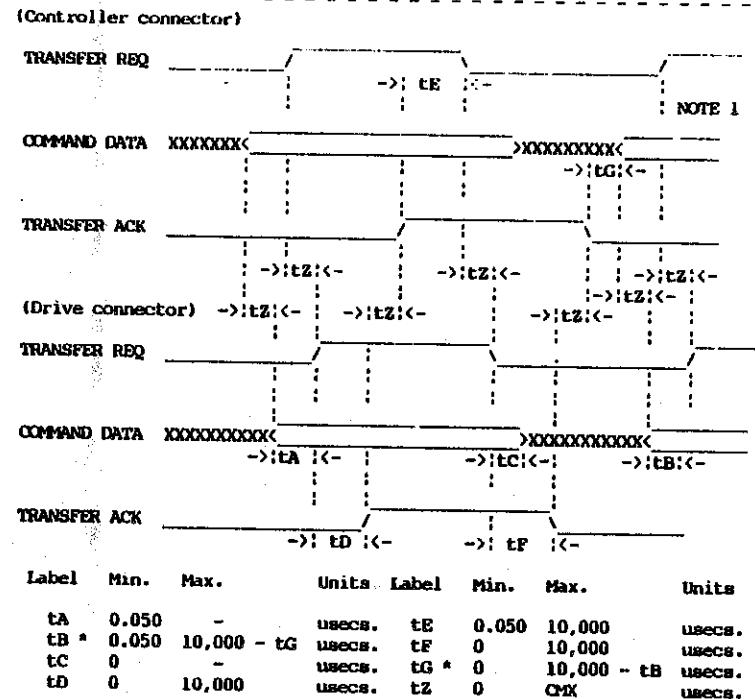
The interface uses a mixture of open collector TTL signals for control and differential signals for data. The terms asserted and negated are used for consistency and to be independent of electrical characteristics.

Asserted - the negative logic signal is true i.e. zero.
- the differential signal is positive i.e. one.

Negated - the negative logic signal is false i.e. one.
- the differential signal is negative i.e. zero

Other terms may be used in lieu of actual numbers so that timing may be determined separately for different drives and their supported cable types.

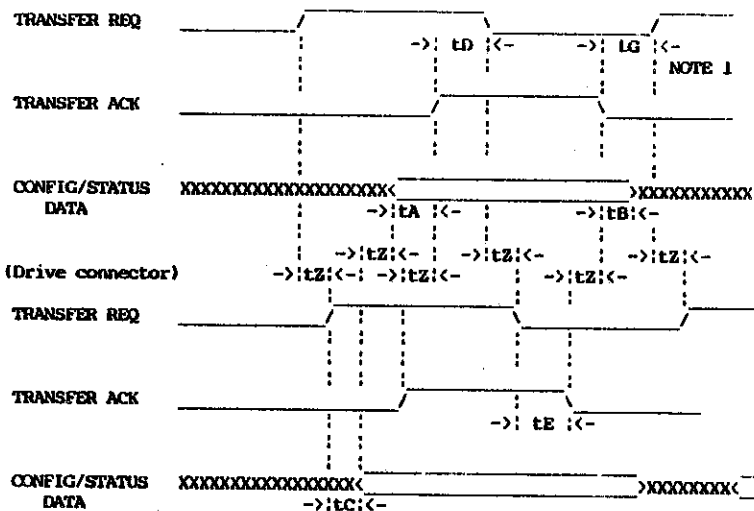
CXX = Cable Delay Maximum - The cable propagation delay based on the maximum length of cable in the configuration in which this interface is used.



NOTE 1: tG timing applies on all except last bit.
* tG + tB must be less than or equal to 10 milliseconds.

FIGURE 10-1 ONE BIT TRANSFER TIMING (TO DRIVE)

(Controller connector)



Label	Min.	Max.	Units	Label	Min.	Max.	Units
t_A *	0.050	10,000 - t_C	usecs.	t_D	0.050	10,000	usecs.
t_B	0	-	usecs.	t_E	0	10,000	usecs.
t_C *	0	10,000 - t_A	usecs.	t_G	0	10,000	usecs.
t_Z	0	CMX					

NOTE 1: t_G timing applies on all except last bit.* $t_A + t_C$ must be less than or equal to 10 milliseconds.

FIGURE 10-2 ONE BIT TRANSFER TIMING-FROM DRIVE

(Drive connector)

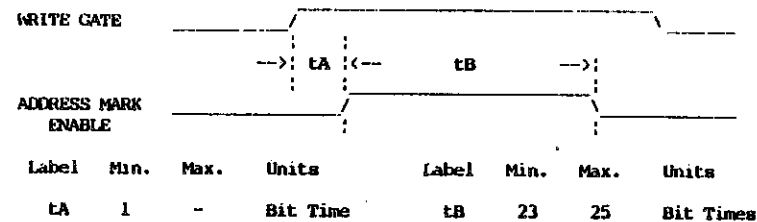


FIGURE 10-3 WRITE ADDRESS MARK TIMING

(Drive connector)

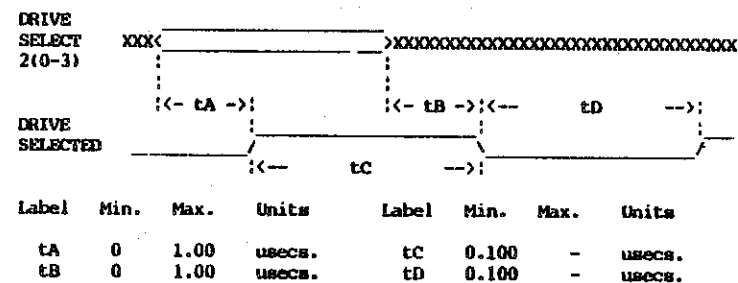
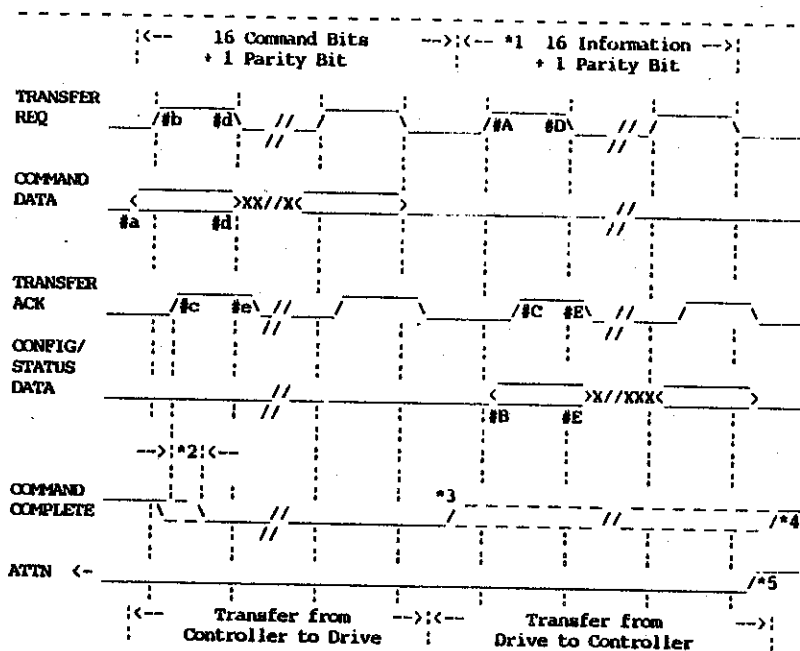


FIGURE 10-4 DRIVE SELECT TIMING



- *1 Applicable for all Request Status and Configuration commands.
 *2 COMMAND COMPLETE is negated for all commands to the drive. COMMAND COMPLETE shall be negated following the assertion of TRANSFER REQ and the maximum time by which COMMAND COMPLETE shall be negated is 100 nsecs after TRANSFER ACK is asserted.
 *3 COMMAND COMPLETE is asserted to signify completion of execution of a command. Applicable for all commands.
 *4 COMMAND COMPLETE is asserted to signify completion of the requested Configuration/Status transfer.
 *5 If an error was encountered during the current command, ATTENTION shall be asserted at least 100 nsec before COMMAND COMPLETE is asserted.

FIGURE 10-5 TYPICAL OPERATION(S)

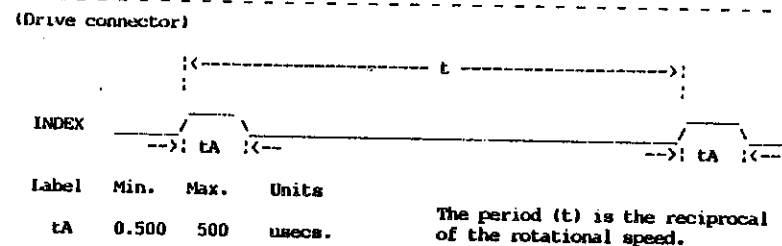


FIGURE 10-6 INDEX PULSE TIMING

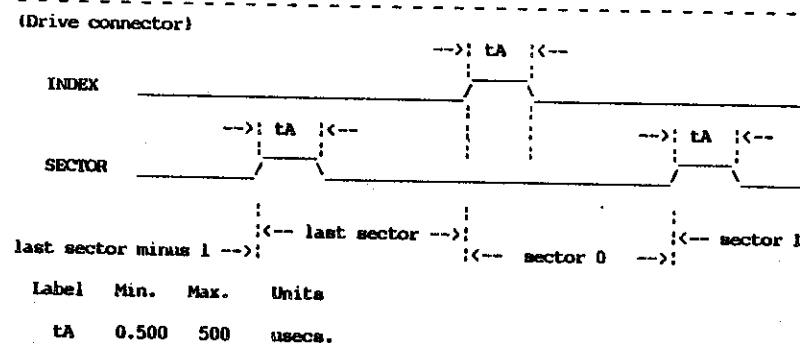
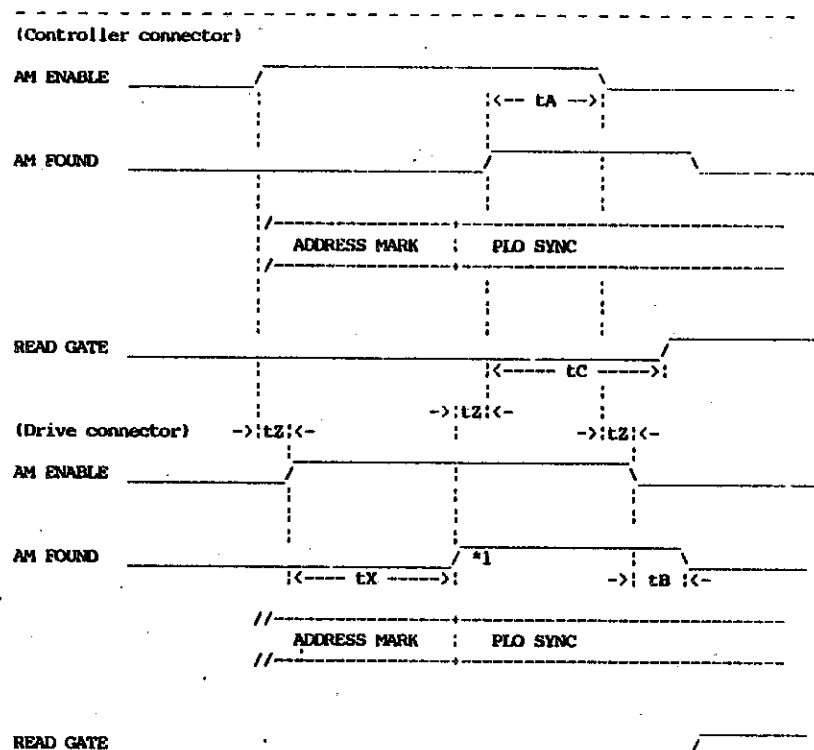


FIGURE 10-7 SECTOR PULSE TIMING

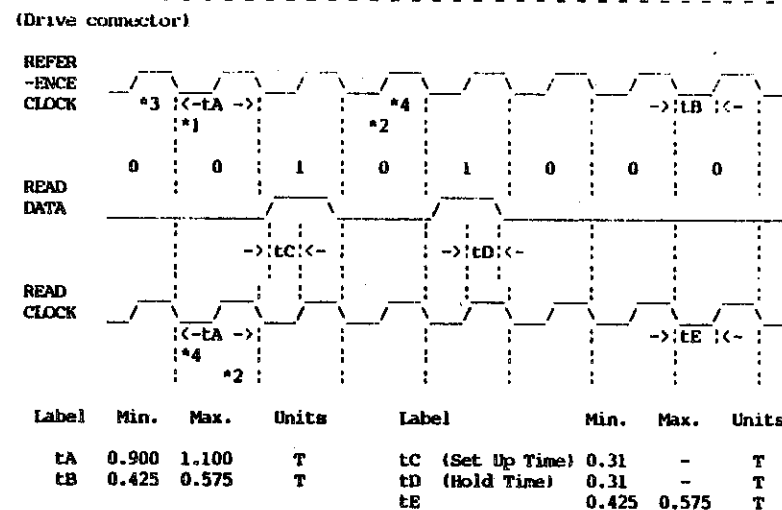


Label	Min.	Max.	Units	Label	Min.	Max.	Units
tA	0	-	usecs.	tX	Determined by Drive		
tB	0	0.100	usecs.		Manufacturer		
tZ	0	CMK	usecs.				

*1 Leading edge indicates the location of the end of an Address Mark

NOTE: There shall be no more than an 8 bit delay between Address Mark detection at the drive and the assertion of ADDRESS MARK FOUND at the controller connector.

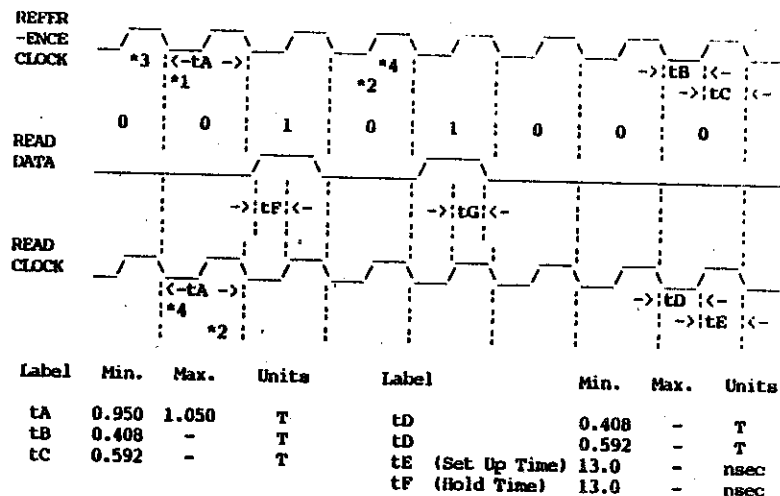
FIGURE 10-8 READ ADDRESS MARK TIMING



- *1 All times are measured in fractions of the clock period. T is the nominal period of the clock signals and is the inverse of the REFERENCE or READ CLOCK frequency.
- *2 Phase relationship between REFERENCE CLOCK and NRZ WRITE DATA or WRITE CLOCK is not defined.
- *3 REFERENCE CLOCK is valid when READ GATE is inactive. READ CLOCK is valid when READY is asserted.
- *4 Similar period symmetry shall be +/- 4 nsecs between any two adjacent cycles during reading or writing.

FIGURE 10-9 LOW SPEED PORT (0-15Mhz) DISK NRZ READ DATA TIMINGS

(All signals are measured at the drive connector)

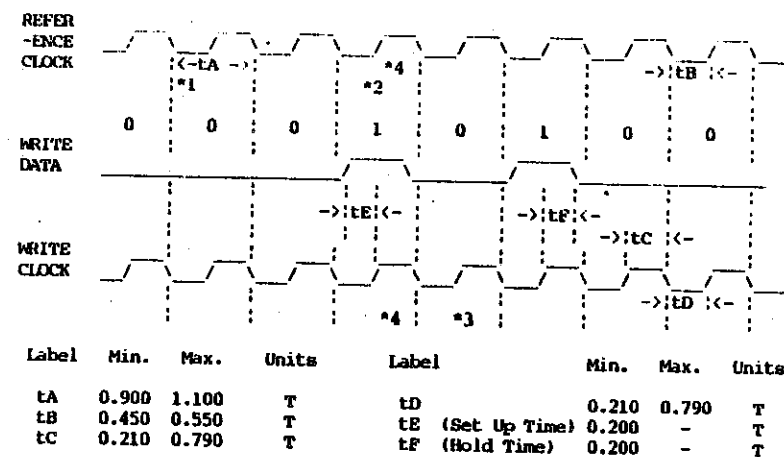


- *1 Some times are measured in fractions of the clock period. T is the nominal period of the clock signals and is the inverse of the REFERENCE or READ CLOCK frequency.
- *2 Phase relationship between REFERENCE CLOCK and NRZ WRITE DATA or WRITE CLOCK is not defined.
- *3 REFERENCE CLOCK is valid when READ GATE is inactive. READ CLOCK is valid when READY is asserted.
- *4 Similar period symmetry shall be +/- 2 nsecs between any two adjacent cycles during reading or writing.

NOTE: When operating at 24 MHz the following values shall be used:
tB = 17 nsec tD = 17 nsec

FIGURE 10-10 HIGH SPEED PORT (10-21MHz) DISK NRZ READ DATA TIMINGS

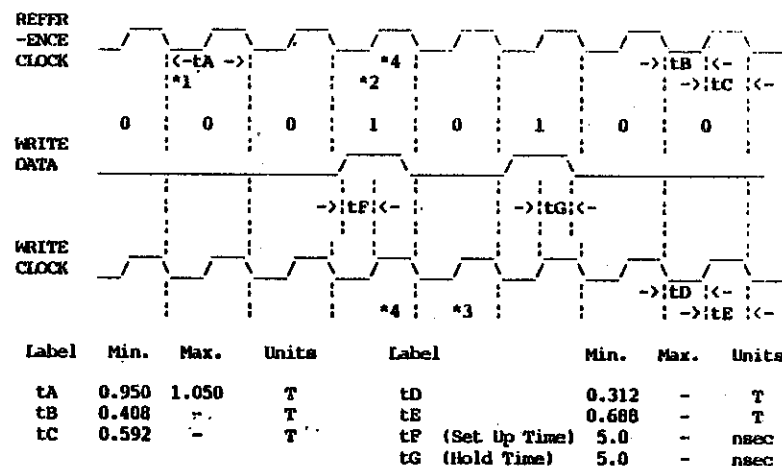
(Drive connector)



- *1 All times are measured in fractions of the clock period. T is the nominal period of the clock signals and is the inverse of the REFERENCE or READ CLOCK frequency.
- *2 Phase relationship between REFERENCE CLOCK and NRZ WRITE DATA or WRITE CLOCK is not defined.
- *3 The WRITE CLOCK shall be the same frequency as the drive supplied REFERENCE CLOCK (i.e. the WRITE CLOCK is the controller received and retransmitted DRIVE REFERENCE CLOCK).
- *4 Similar period symmetry shall be +/- 4 nsecs between any two adjacent cycles during reading or writing.

FIGURE 10-11 LOW SPEED PORT (0-15MHz) DISK NRZ WRITE DATA TIMINGS

(REFERENCE CLOCK is measured at the drive connector;
WRITE DATA and WRITE CLOCK are measured at the controller connector)



*1 Some times are measured in fractions of the clock period. T is the nominal period of the clock signals and is the inverse of the REFERENCE or READ CLOCK frequency.

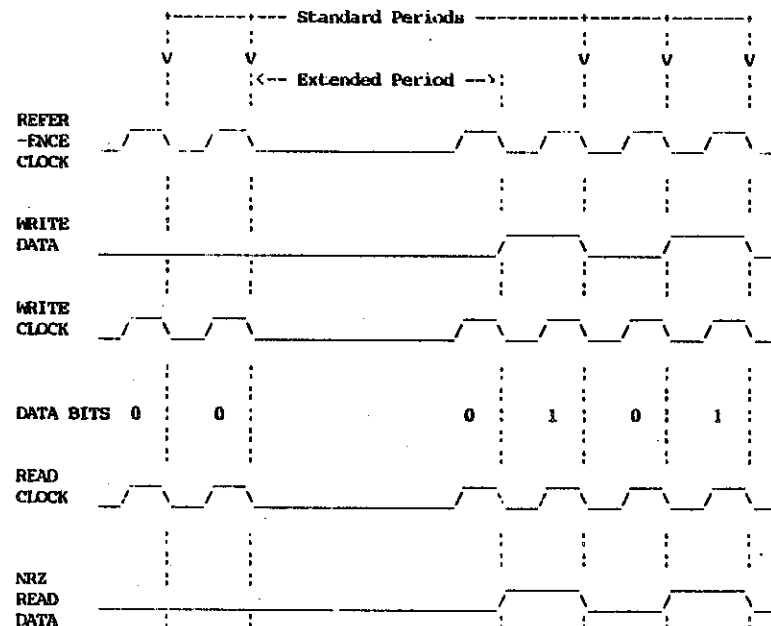
*2 Phase relationship between REFERENCE CLOCK and NRZ WRITE DATA or WRITE CLOCK is not defined.

*3 The WRITE CLOCK shall be the same frequency as the drive supplied REFERENCE CLOCK (i.e. the WRITE CLOCK is the controller received and retransmitted DRIVE REFERENCE CLOCK).

*4 Similar period symmetry shall be +/- 2 nsecs between any two adjacent cycles during reading or writing.

NOTE: When operating at 24 MHz the following values shall be used:
tB = 17 nsec tD = 13 nsec

FIGURE 10-12 HIGH SPEED PORT (10-21MHZ) DISK NRZ WRITE DATA TIMINGS



Standard Period: A symmetrical wave form which is the reciprocal of the steady state clock period.

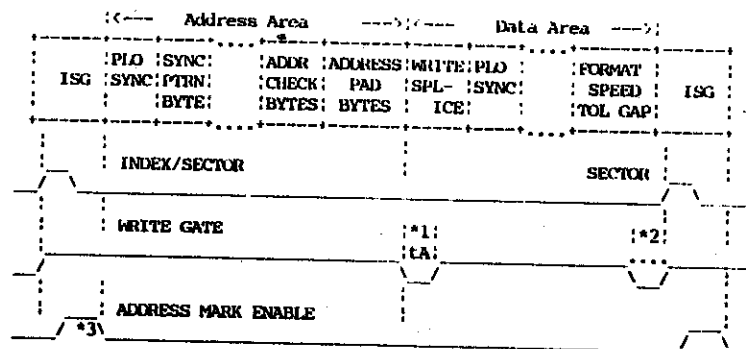
Extended Period: A wave form in which the negated level is longer than the asserted level by more than the specified clock tolerances. Extended periods are drive specific.

There shall be no clock periods presented which are less than the specified minimum periods i.e. no slivers.

There shall be no clock periods which extend more than 15 minimum periods.

Data is latched on the leading edge of Clock Period.

FIGURE 10-13 HIGH SPEED PORT AND OPTICAL NRZ TIMING PERIODS

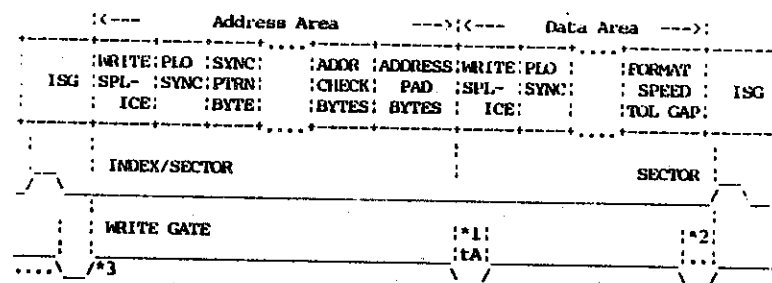


- *1 Transition required only if the disk is read after a format and prior to a data field write update.
 *2 Controller must re-initialize timing for each sector relative to SECTOR pulse (need not negate WRITE GATE).
 *3 Trailing edge of ADDRESS MARK ENABLE (negation) signifies the start of Header PLO Sync Field. Drive will not write an Address Mark on the disk.

Label	Min.	Max.	Units	Label	Min.	Max.	Units
tA	2	-	Bit Times				

NOTE: The use of Address Mark Enable allows writing of an Address Area without a write splice between the ISG Gap and the Sync field.

FIGURE 10-14 FIXED SECTOR FORMAT TIMING - ALTERNATIVE 1

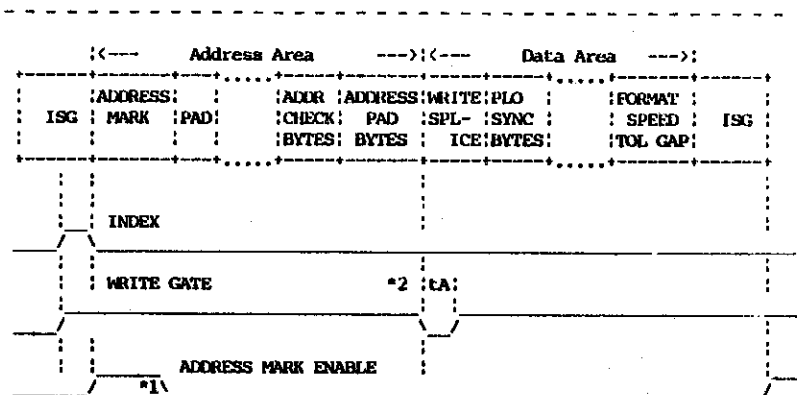


- *1 Transition required only if the disk is read after a format and prior to a data field write update.
 *2 Controller must re-initialize timing for each sector relative to SECTOR pulse (need not negate WRITE GATE at this point).
 *3 Leading edge of WRITE GATE (assertion) defines a Write Splice and the start of the Header PLO Sync Field.

Label	Min.	Max.	Units	Label	Min.	Max.	Units
tA	2	-	Bit Times				

NOTE: Alternative 2 is required on disks implemented with a high speed port.

FIGURE 10-15 FIXED SECTOR FORMAT TIMING - ALTERNATIVE 2



Label	Min.	Max.	Units	Label	Min.	Max.	Units
ta	2	-	Bit Times				

- *1 Trailing edge of ADDRESS MARK ENABLE signifies the start of the Header PLO Sync Field
- *2 Transition required only if the disk is read after a format and prior to a data field write update.

NOTE: Not required on disks implemented with a high speed port.

FIGURE 10-16 SOFT SECTOR FORMAT TIMING

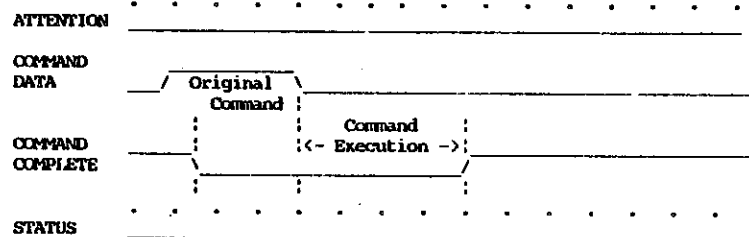


FIGURE 10-17 NORMAL OPERATION - NO ERRORS OR FAULTS

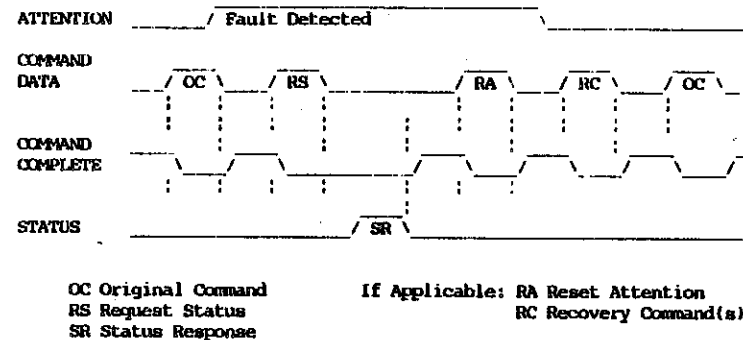


FIGURE 10-18 ABNORMAL OPERATION - FAULT DETECTED BY DRIVE WHILE ATTEMPTING TO EXECUTE A COMMAND

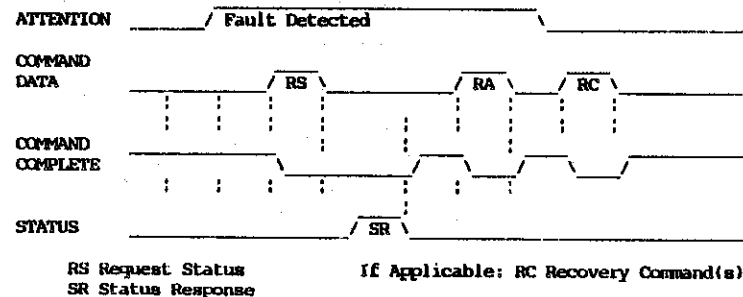
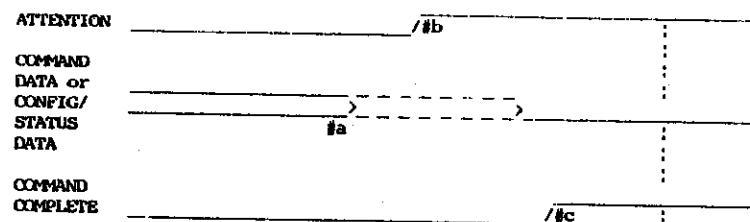


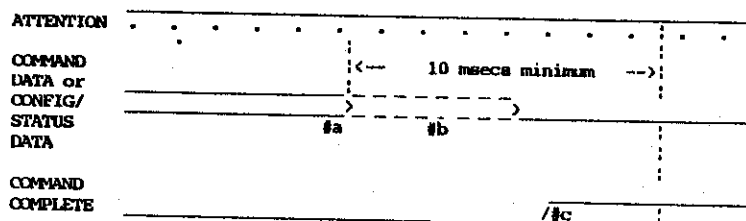
FIGURE 10-19 ABNORMAL OPERATION - POWER TURN ON OR FAULT DETECTED BY DRIVE WITHOUT A COMMAND OR CONFIG/STATUS DATA TRANSFER



Drive and controller are synchronized to "Command Transfer State" for subsequent recovery

- #a Serial communications fault detected.
- #b Drive asserts ATTENTION to inhibit further bit transfers and notify the controller that a fault has been detected.
- #c Drive asserts COMMAND COMPLETE after "Orderly Termination."

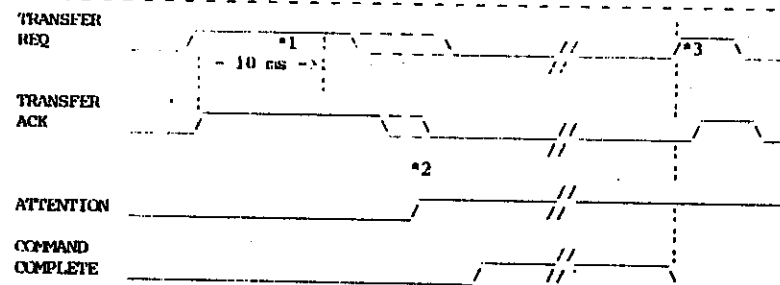
FIGURE 10-20 COMMUNICATIONS TRANSFER FAULT, ATTENTION PREVIOUSLY NEGATED



Drive and controller are synchronized to "Command Transfer State" for subsequent recovery

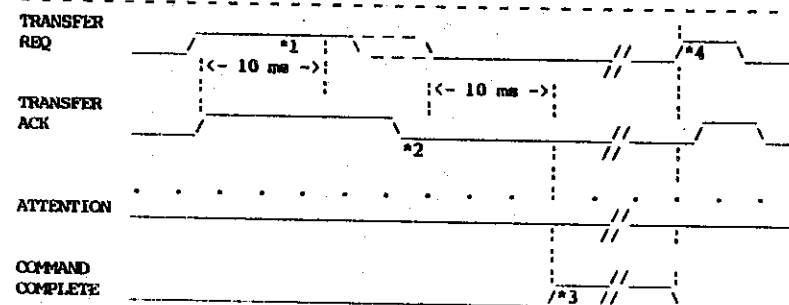
- #a Serial communications fault detected.
- #b Drive (or controller) generates an "Interface Fault" on the next request for a bit transfer (if applicable) to inhibit further bit transfers and notify the controller that a fault has been detected.
- #c Drive asserts COMMAND COMPLETE after "Orderly Termination."

FIGURE 10-21 COMMUNICATIONS TRANSFER FAULT, ATTENTION PREVIOUSLY ASSERTED



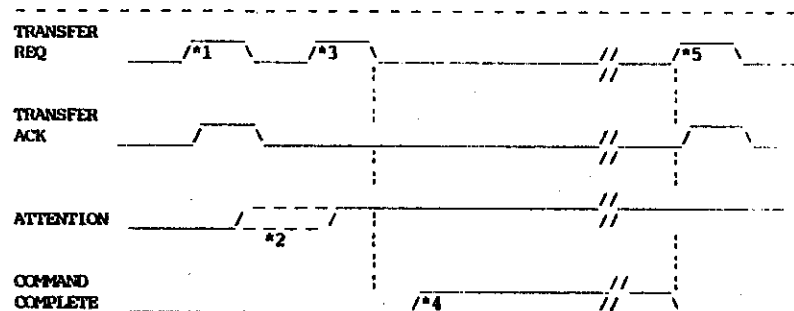
- *1 Interface fault detected by the drive.
- *2 TRANSFER ACK is negated and ATTENTION is asserted.
- *3 Controller initiates Request Status command.

FIGURE 10-22 INTERFACE FAULT (<17 BITS TRANSFERRED, TRANSFER REQ=1) WITH ATTENTION PREVIOUSLY NEGATED



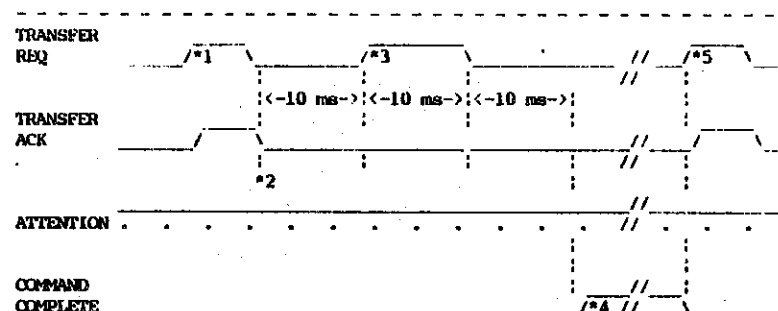
- *1 Interface fault detected by the drive.
- *2 TRANSFER ACK is negated and ATTENTION is asserted.
- *3 Drive sets COMMAND COMPLETE after TRANSFER REQ is negated continuously for 10 msec (required in case controller attempts to send another transfer as a part of this transfer so that the drive will create an interface fault to stop the controller).
- *4 Controller initiates Request Status command.

FIGURE 10-23 INTERFACE FAULT WITH ATTENTION PREVIOUSLY ASSERTED



- *1 Last bit of Command Data transfer.
- *2 Attention set by drive could occur before or after TRANSFER ACK is negated. If it occurs after TRANSFER ACK is asserted, it could occur before or after TRANSFER REQ is negated.
- *3 First TRANSFER REQ of Config/Status Data Response can occur before ATTENTION is asserted.
- *4 TRANSFER REQ is negated and ATTENTION is asserted. Drive asserts COMMAND COMPLETE after TRANSFER REQ is negated.
- *5 Controller initiates Request Status command.

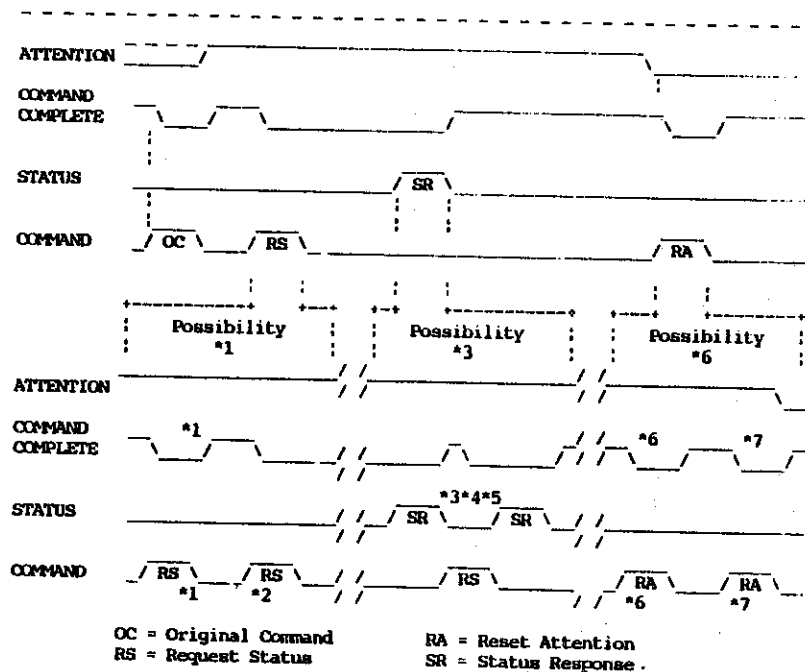
FIGURE 10-24 PARITY ERROR OR INVALID COMMAND FAULT WITH STATUS DATA RESPONSE EXPECTED BY CONTROLLER, AND ATTENTION PREVIOUSLY NEGATED



- *1 Last bit of Command Data transfer.
- *2 Drive discovers error and starts looking for TRANSFER REQ negated continuously for 10 msecs.
- *3 Controller asserts first TRANSFER REQ for Config/Status Data Transfer but removes it after interface fault created by the drive to stop controller.
- *4 TRANSFER ACK is negated and ATTENTION is asserted. Drive asserts COMMAND COMPLETE after TRANSFER REQ is negated continuously for 10 msecs.
- *5 Controller initiates Request Status command.

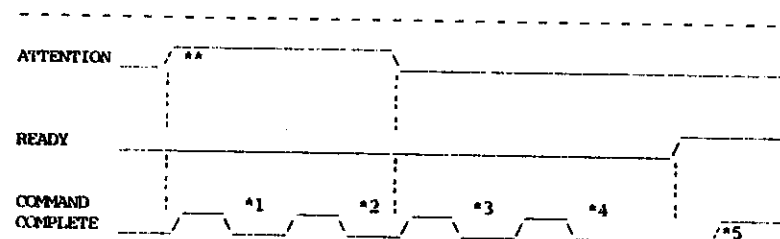
FIGURE 10-25 PARITY ERROR OR INVALID COMMAND FAULT WITH STATUS DATA RESPONSE EXPECTED BY CONTROLLER, AND ATTENTION PREVIOUSLY ASSERTED

170



- *1 Drive detects serial communication faults and sets appropriate status bit.
- *2 Controller reissues Request Status command which is accepted by the drive.
NOTE: Controller should issue Request Status command a minimum of three times.
- *3 Drive detects interface fault, sets Interface Fault Status bit and creates an interface fault at controller OR controller detects interface fault OR parity error on STATUS
- *4 Controller re-requests status. NOTE: Conditions *1 and *2 may occur.
- *5 Good status received at controller.
NOTE: Controller should perform this sequence a minimum of three times.
- *6 Drive detects *1 OR is unable to reset ATTENTION.
- *7 Controller reissues Reset Attention command (Condition *5 may occur).
NOTE: Controller should issue the Reset Attention command a minimum of three times. After third time, if ATTENTION is still asserted, controller should Read Status to find out why ATTENTION is asserted.

FIGURE 10-26 COMMUNICATIONS FAULTS RECOVERY PROCEDURE



- *1 Request Standard Status.
- *2 Request Status and ATTENTION.
- *3 Request Configuration (Optional).
- *4 Start spindle motor.
- *5 READY and COMMAND COMPLETE asserted with ATTENTION negated indicates motor is up to speed.
- ** ATTENTION due to a power on sequence which results in the spindle not started or a spindle which was previously started is stopped.

FIGURE 10-27 TIMING DIAGRAM FOR A SUCCESSFUL SPINDLE START SEQUENCE VIA A MOTOR ON COMMAND

THIS

PAGE

INTENTIONALLY

LEFT

BLANK

172

APPENDIX A: MAGNETIC DISK DEFECT LIST

The drive's defect list is capable of being read in modified soft sectoral operation as well as hard sectoral.

The defect list shall reside on Sector 0 of the maximum cylinder and be repeated on two other cylinders; maximum cylinder minus 8 and cylinder 4095 (x'FFF'). This allows for redundancy should an error occur on the maximum cylinder.

All ESDI disks shall ship with copies of the defect list on three cylinders. The cylinder at the address of 4095 is a drive-unique location. It may be located anywhere on the drive that the manufacturer chooses.

The manufacturer shall guarantee that either maximum cylinder or the cylinder at 4095 is error free.

The sector 0 of each surface will contain the defects for that surface only.

The format for the data field portion (see Figure A-1) of this sector is normally 256 bytes with 2 bytes of CRC (x(16) + x(12) + x(5) + 1):

- o Defect locations are 5 bytes long and the bytes are defined in Figure A-1.
- o The start of the actual defect may be off by up to 7 bits due to the one byte resolution. This resolution is drive dependent.
- o The end of the defect list for each surface will be indicated by 5 bytes of ones in the defect location field or the end of the sector.
- o The Byte Count field is the number of bytes from INDEX.
- o The CRC check bytes should be used if that capability exists but may be ignored if multiple reads are a more desirable approach.

The following rules apply to the reading (and writing) of the Defect List:

- o The CRC seed shall be zero (initialized state)
- o The Sync Byte will be included in the CRC calculation (applies to both header and data fields).
- o Header and Data fields are recorded (and read) most significant byte first.

Provision has been made for future growth in the number of bytes in a Defect List data sector by defining the following fields in the two least significant bits of the Header field Flag byte. Bits 7-2 shall be zero.

Bit 1	Bit 0	
0	0	256 Byte Defect List Data Field
0	1	512 Byte Defect List Data Field
1	0	Reserved for 1,024 Byte Defect List Data Field
1	1	Reserved for 2,048 Byte Defect List Data Field

Only one size of Defect List shall be recorded per drive and must be the minimum sector size which will contain the locations of the maximum number of defects allowed per track.

A 256 byte Defect List data field can list up to 50 defects, and a 512 byte data field can list up to 101 defects.

NOTE: On drives developed prior to this standard being defined there may be a deviation from the formatting rules as not all drives provide a write splice between the header and data fields. The controller may first try recovering the data field assuming that a write splice is formatted between the header and the data field. If the data field cannot be read without error then the controller may then assume that there is no write splice formatted between the header and the data field, and retry the read function.

INDEX/SECTOR										
*1										
<--->										
Address Area										
<--->										
WRITE:PLO SYNC SEC:FLAG:ADDR ADDRESS:										
ISG	SPL	SYNC	PTRN	CYLINDER	HEAD	TOR	STA	CHECK	PAD	
ICE:BYTES:BYTE:msb lsb: TUS BYTES BYTES										
<--->										
Hex	00	00	00	FE	*3	*3	*3	*3	*4	00
<--->										
Size	*2	1	*2	1	2	1	1	1	2	2
<--->										
...										
<--->										
Data Area										
<--->										
WRITE:PLO SYNC DATA DATA										
SPL SYNC PTRN DEFECT LIST CHECK PAD ISG										
ICE:BYTES:BYTE BYTES BYTES										
<--->										
Hex	00	00	FE	*5			*4	00	00	
<--->										
Size	1	*2		256			2	2	n	
<--->										
...										

FIGURE A-1 DEFECT LIST FORMAT

- *1 Drives with a High Speed Port require that a write splice be formatted at the start of the PLO Sync field.
- *2 PLO Sync Field and ISG are as reported in response to Request Configuration commands.
- *3 Cylinder = Maximum Cylinder or (Maximum Cylinder minus 8) or 4095
Head = 0, 1, etc Sector = 0 Flag = x'00' or x'01'
- *4 $X(16) + X(12) + X(5) + 1$

Field	<div> <div>Date</div> <div>Defect</div> <div>Defect</div> <div>Defect</div> <div>//</div> <div>Defect</div> <div>Defect</div> </div>					
	<div> <div>Location</div> <div>Location</div> <div>Location</div> <div>//</div> <div>Location</div> <div>Location</div> </div>					
	<div> <div>MDMPP</div> <div>CCBBL</div> <div>CCBBL</div> <div>CCBBL</div> <div>//</div> <div>CCBBL</div> <div>CCBBL</div> </div>					
	<div> <div>6</div> <div>5</div> <div>5</div> <div>5</div> <div>//</div> <div>5</div> <div>5</div> </div>					
Size	6	5	5	5	//	5
//						

The Date consists of six one-byte fields:

M=Month: x'01-0C' D=Day: x'01-1F' Y=Year: x'52-xx'

H=Head: x'00-xx' P=Pad: x'00' P=Pad: x'00'

The Defect Location consists of three fields:

CC = MSB and LSB of cylinder address.

BB = MSB and LSB of byte count from index.

L = Length of error in bits (NOTE: resolution is within 7 bit cells of start of the flaw)

The last Defect Location entry is followed by x'FF' to end of sector.

The controller should recognize either 5 bytes of x'FF' or end of sector as terminating the list.

Header and Data Field bytes are recorded with Most Significant Byte and Most Significant Bit first.

NOTE: The manufacturer may choose to pad the PLO Sync area by up to four bytes beyond those reported in Configuration Data.

APPENDIX B: DISK DEFECTS

B.1 Manufacturer Certified Optical WORM Media

If the manufacturer formats and certifies the media, each sector with defects large enough to cause unrecoverable errors is typically identified in the header address. This identification is used by the drive to prevent the controller attempting to write into such a sector.

The controller can build its own Flaw Map by scanning the media and reading all the Header Addresses in order to locate those which are flagged as defective by the manufacturer. Alternatively, it could find the defective sectors during normal operations, and dynamically assign alternate sectors at that time. Algorithms to manage defective sectors are controller dependant and there are several alternative methods which may be used.

If the controller is using the media in a strictly sequential application, there is no need to assign alternates, but simply use the following sector to record the data.

B.2 Optical WORM Disk Considerations

In applications which use optical WORM disk as a randomly accessible media a comprehensive defect management scheme must be provided. The following are considerations that may be followed when assigning alternate sectors.

When a new cartridge is loaded, read into memory the contents of the Flaw Map (which is maintained by the controller). If this is a "virgin" cartridge i.e. never been loaded before then there is no Flaw Map, so this should be reported back to the host. It is assumed that the host will then load another cartridge, or request the controller to Scan the new one and build a Flaw Map.

If the Header Address identifies that the sector is defective then its address shall be listed in the Flaw Map and an alternate sector allocated. It is recommended that the alternate sector area be allocated in close proximity to the Flaw Map so that seeking to the Flaw Map will place the controller only a short seek away from the requested sector.

If the Flaw Map is loaded into controller memory when a cartridge is loaded, then alternates could be assigned first to the spare on-track sector, and if already assigned, to the next closest available. Note: If a disk drive does not provide the optional Post Field there is no means to record the address of the alternate sector other than in the Flaw Map. On a high capacity optical disk the Flaw Map is likely to exceed the RAM capacity of many controllers.

Re-allocating defective sectors could be done dynamically, but then each sector of the Flaw Map would contain only one entry, being the address of the defective sector, and the address of its alternate (since the media is not re-writable). A more efficient use of space is to scan a cartridge when initially loaded and build a Flaw Map that has as many entries per sector as can be fitted. The first sector of the Flaw Map should be the last one written by the controller, and should contain the address of the last sector in the Flaw Map that holds factory-identified sectors. This address should be loaded into controller memory.

Not all defects can be factory identified. If the controller detects a defective sector during operation (this is possible on the next revolution if the host requires that each sector written be verified as correctly recorded and readable), then an alternate must be assigned dynamically.

The controller shall begin reading at the first sector in the Flaw Map following the entries for factory identified defects. When the first non-recorded sector is read, then the address of the defective sector and its assigned alternate is used by the controller to write the data in the alternate sector.

If the data is written successfully in the alternate sector then the sector containing the entry is written to the Flaw Map. If the data could not be successfully written to the alternate sector, then a new alternate is assigned and until data is successfully recorded in an alternate, no entry is written to the Flaw Map.

APPENDIX C: DISK ERROR CORRECTION

C.1 ERROR CORRECTION REQUIREMENTS

The requirements for defect management and error correction for disk drives, especially those utilizing optical technology are established by initial and end-of-life defect rates. With optical disks, allowing for margin, an initial defect rate as high as $1.0E-4$ must be accommodated, and the defect rate at end-of-life is expected to be twice the initial defect rate.

C.1.1 INITIAL DEFECTS

One of the most powerful methods for handling initial defects is retiring bad sectors.

On optical disks it is not practical to retire all sectors containing a defect when the disk has an initial defect rate as high as $1.0E-4$. In this case the area to be retired must be smaller than a sector or small defects meeting some criteria must be left for ECC to handle. Large defects must still be retired.

Magnetic disks have a much lower defect rate, and it is practical to retire all sectors which contain a defect. The ECC is used as an error recovery method for occasional errors and is not a strategy for data integrity.

C.1.2 END-OF-LIFE DEFECTS

With optical disks the end-of-life defects must be handled by ECC. The selected code must handle initial defects that were left for ECC at the time of writing plus new defects that have surfaced between initial writing and end-of-life.

Code selection is influenced by defect rates and length distributions and the drive's ability to clock over various defect lengths.

C.1.3 ERROR TOLERANCE REQUIRED THROUGHOUT

When architecting a drive and controller one must strive for error tolerance in all areas of the architecture. In the case of high error rate media such as optical disk, all special fields must be error tolerant. This includes but is not limited to sync patterns and header fields.

C.1.4 ERROR-TOLERANT SYNC PATTERNS

There are two ways to handle sync pattern errors. First an error-tolerant sync pattern can be used. Error-tolerant sync patterns have been used in the past that can be detected at the proper time even if several random bits of the pattern are in error. Other error-tolerant sync patterns have been used that can be detected at the proper time even if there is a burst error anywhere in the pattern. A large defect can completely wipe out even an error-tolerant sync pattern.

Another way to handle sync pattern errors is to architect the drive and controller such that no single sync pattern is required to be detected. If the bit timing is correct and the sync pattern is not detected, just keep reading. This is possible only if periods of timing uncertainty between sectors,

analogous to those caused by a write splice in magnetic recording, are avoided.

The drive/controller architecture should provide for several consecutive sync patterns to be in error without causing a loss of data. Even if the architecture avoids timing uncertainty between sectors, there may still be some timing uncertainty between the last sector of a track and the first sector of a track. In this case, special care must be taken so that a defect in the first sync pattern of a track does not cause a loss of data.

It is also possible to use a combination of the above techniques.

If a 16-bit sync pattern is used and the defect rate is $1.0E-4$ then the probability that a sync pattern contains a defect is $1.6E-3$. Under these conditions and assuming 50 sync patterns (header and data) per track, roughly one of ten tracks would contain a defective sync pattern.

C.1.6 ERROR TOLERANT HEADER FIELDS

It is possible to use ECC on the header fields as well as data fields. However, the header field is such a small field that a large defect could completely wipe it out. Another approach is to use an error detection code on the header and store in it track and sector address information only. The idea is that address information can be generated from track orientation.

Assuming eight-byte header fields, a defect rate of $1.0E-4$, and 25 sectors per track, roughly one of six tracks would contain a defective header. Whether or not a lost revolution can be tolerated in defective header handling depends on performance requirements.

C.1.7 DEFECT CONSIDERATIONS

For optical media, the initial defect rate that the ECC must handle is that associated with defects that were intentionally left for ECC at the time of writing. The defect rate will increase over a period of time toward the end-of-life defect rate.

For magnetic media, the ECC is typically used to handle errors that result from "grown defects" (i.e. latent defects not identified in the factory scan or field certification). Automatic reallocation of affected sectors is one strategy that a controller may adopt.

It is possible to monitor the error correction thresholds and pass flags to the user on marginal sectors before data becomes uncorrectable. One flag would request that a particular sector be retired, while another would request that the entire piece of media be retired. Obviously, this type of defect prognosis is effective only if the media is read periodically.

176

C.2 ERROR CORRECTION CODES

Controller manufacturers may choose to use any ECC that suits their purpose.

On removable media it is desirable to use a common error correcting code that may help to achieve format compatibility and minimize the development efforts of controller companies.

C.3 DIAGNOSTIC CONSIDERATIONS

In many implementations hardware computes the syndromes but is not involved in the correction algorithm, which is totally contained in software. In this case, it is easy to distinguish between hardware and software failures by testing the software first. First test all tables contained in software, then supply syndromes to the software for which proper responses are known.

There are several approaches to diagnosing the hardware. The approach used in controllers for magnetic media devices requires the implementation of "read long" and "write long" commands within the controller which bypass the ECC circuitry. The "read long" command is identical to the normal read command except that redundancy bytes are read as if they were data bytes. The "write long" command is identical to the normal write command except that redundancy bytes to be written are supplied, not generated by hardware.

The "read long" command is used to read a known defect-free data record and its redundancy bytes. A simulated error condition is XOR'd into the record and the modified data record plus redundancy bytes are written back, using the "write long" command. On readback, using the normal read command, an ECC error should be detected by the controller circuitry and the correction routines should generate correct responses for the error condition simulated. The test is repeated for several simulated error conditions, correctable and uncorrectable.

The approach described above can be used for optical WORM disks but has the disadvantage of using up write-once media. Also, the corrupted data record read back is more likely to contain real errors along with the simulated ones than is the case with magnetic disks.

Another approach would be to put the controller in a special diagnostic mode where write is inhibited. The controller is in a write mode for accepting data, but is in a read mode for ECC processing, thus supplying to the ECC circuits a data record plus redundancy bytes which contains known error conditions. Correct syndrome generation can then be verified.

C.4 BUFFER CONSIDERATIONS

Correction times for the code are established by the implementation alternatives selected and the speed of the processor performing the correction algorithm. The number of errors correctable without losing a revolution is dependent on correction times, data transfer rate, and buffer strategy. Several of the many possible strategies are outlined below. With each strategy, data are delayed by one sector time, plus the time required for correction.

Correction times must be considered when selecting a processor and buffer strategy. In many cases, three or more buffer will be required.

C.4.1 SINGLE BUFFER STRATEGY WITH SECTOR INTERLEAVING

Using a single buffer with sector interleaving, the amount of time available for correction is established by the interleave factor. An interleave factor of two would allow only one sector time for both correction and buffer unload. Buffer unload to the host would have to be faster than buffer load from the device in order to have any time for error correction. If an interleave factor of three is used and buffer unload is at least as fast as buffer load, error correction would be allowed a minimum of one full sector time.

C.4.2 TWO BUFFER STRATEGY WITH NO SECTOR INTERLEAVING

The two buffer strategy with no sector interleaving requires that data transfer to the host be faster than data transfer from the device. The only time available for correction is the difference in time between loading a buffer from the device and unloading a buffer to the host.

C.4.3 THREE BUFFER STRATEGY WITH NO SECTOR INTERLEAVING

The three buffer strategy with no sector interleaving allows one full sector time for correction. One buffer could be receiving data from the device, a second buffer could be undergoing correction, and a third buffer could be transferring data to the host.

C.4.4 BUFFERING OF SYNDROMES AS WELL AS DATA

The strategies above assume that syndromes are not buffered i.e. syndromes for one sector have to be used before syndromes for the next sector are generated. It is possible to implement multiple buffering such that on multiple sector reads, both data and syndromes are buffered. Software starts performing error correction when buffering of data and syndromes is complete for the first sector. Only when a buffer has been corrected is it released to the host. Orientation is not lost in severe error cases provided that the number of sector times of accumulated delay does not exceed the number of buffers. After an error correction delay, the lost time can be made up if data transfer to the host is faster than data transfer from the device.

C.5 ERASURE POINTERS

The correction capability of an EDC scheme can be improved if additional external information can be provided to indicate where the errors are located in the data. This information is called erasure pointers. Erasure

pointers mark a byte as being bad according to some criteria. This criteria may include marginal amplitude, run length violations, or marginal timing.

Within ESDI, a special read command can be initiated by which erasure pointers are transferred across the ESDI interface in place of normal read data on the NRZ read data lines. The pointers are coincident with the "bad bytes". Pointers are transferred for the ECC check bytes, CRC bytes, and associated data field only. All other fields associated with the sector in question are transferred across the interface as in an ordinary read operation.

Upon processing a sector in the buffer and finding it bad despite retry procedures, the processor can reread the sector using the erasure read command. This will place the erasure pointer information into the buffer for use in the EDAC algorithm.

APPENDIX D: OPTICAL DISK FORMATS

There is a wide variety of optical formats which may be supported by ESDI.

D.1 ESDI FORMAT

Figure D-1 illustrates a recommended format showing both drive and controller characteristics which consists of five functional areas:

- 1) ISG - Intersector Gap
- 2) Address Field
- 3) Data
- 4) Post Field Gap (Optional)
- 5) Format Speed Tolerance Gap

The Address area is used to locate and verify the track and sector location on the Disk where the Data areas are to be recorded. The Data area is used to record the system's data files. The optional Post Field area may be used to provide the address of alternate sectors.

These fields are used as described in the standard.

178

ESDI Header Address									
*3	*3	*3	*3	*3	*3	*3	*3	*3	*3
ISG	PLO	SYNC	WORLD WIDE	TRACK	SECTOR	FLAG	ADDRESS	SECTOR	ADRSS
SYNC	PATRN	ID NUMBER	ADDRESS	ADDRESS	STATUS	CHECK	WRITE	PAD	
BYTES	BYTES	(optional)	sub	lab	(opt1)	(opt1)	STATUS	BYTES	
2-n	2-n	1-2	4	2	1	1	2	1	2-n

Data									
*1	*3	*3	*3	*3		*3	*3	*3	
WRITE	PLO	RECRDG		CHECK					
SPL	SYNC	PATRN	DATA FIELD	BYTES		DATA FIELD	CHECK	DATA	
ICE	BYTES	BYTES	1	1		n	n	PAD	
1-2	2-n	1-2	128n	8n		128n	8n	1-n	

Optional Post Field Area						*4	
*1	*3	*3	*1	*1	*2	*1	
WRITE	PLO	RECRDG		POST	POST		
SPL	SYNC	PATRN	POST FIELD	CHECK	PAD		
ICE	BYTES	BYTES		BYTES	BYTES		
1-2	2-n	1-2	n	n	1-n		

- *1 These fields may be structured to suit individual customer requirements.
 *2 The number of check bytes is user-defined.
 *3 Established by device and may be reported in response to Request Configuration Commands.
 *4 This area may not be present.
 *5 Motor Speed tolerance gap, if any, after last sector only.

FIGURE D-1 RECOMMENDED FIXED SECTOR OPTICAL FORMAT

D.1.1 Address Field

o World Wide ID (optional) - This field may be used by the manufacturer of the media to identify every address with the 32 bit identification or serial number of the media. Optical disk cartridges are a storage and archival medium which have many library applications. The presence of this field

allows a system to ensure that all reading and/or writing is done to the correct cartridge.

o Track Address - This field is used to define the track that the head(s) are to be positioned to.

o Sector Address - This field defines the address of the sector on the track.

o Flag Status - This field contains flag information recorded by the controller and pertinent to the sector. If Bit 0 is set to one it indicates the second surface of a cartridge.

D.1.2 Fixed Sector Timing

The timing provided in Figure D-2 supports drives that are pre-formatted with header information.

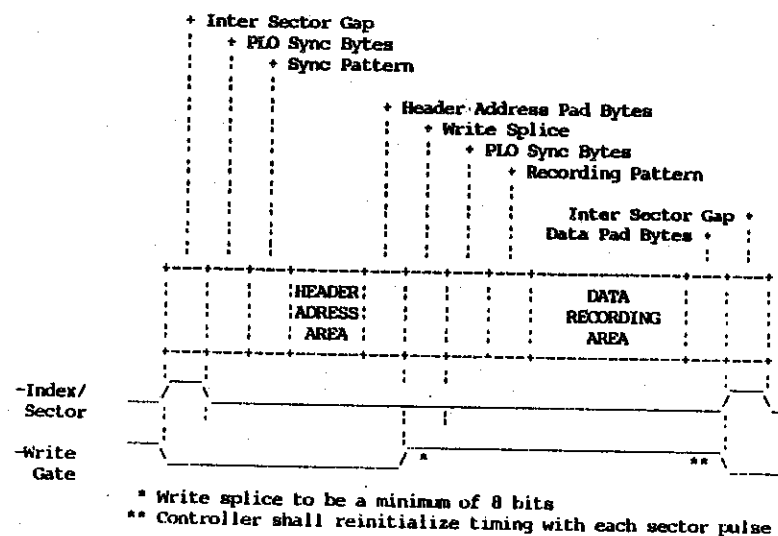


FIGURE D-2 FIXED SECTOR OPTICAL DISK TIMING

D.2 OPTIMEM FORMAT

Figure D-3 and Figure D-4 illustrate the relationship between physical format and timing at the interface.

***** Insert Optimem Figure Here *****

180
FIGURE D-3 OPTIMEM ESDI FORMAT READ TIMING (PRELIMINARY)

FIGURE D-4 OPTIMEM ESDI FORMAT WRITE TIMING (PRELIMINARY)

APPENDIX E: TRANSMISSION LINES AND THEIR EFFECTS ON LINE DRIVERS

Most line drivers have output impedance of 100 ohms or less and are capable of driving 100 ohm lines. When driving long lines the resistance of the transmission line must be considered. If the IR drop of the line becomes significant the voltage swing at the receiver input will be dramatically reduced.

It is important to terminate the transmission line in its characteristic impedance. In general, receivers are designed to have high impedance inputs relative to the transmission line. Input impedances of 10 K ohms or greater insure that the receivers do not load the transmission line.

The output signals from drivers are capable of rise and fall times of 1 nsec/volt or faster i.e. high frequency components of 1 GHz or greater. These frequencies are attenuated by the transmission line and can result in significant reduction in the rise and fall time at the input of the receiver. The attenuation of the line is a function of the line losses and its length.

When the transmission line attenuates the high frequency signal, the rise and fall rate of the signal at the receiving end are reduced. This increase in rise and fall time at the input of the receiver will result in placing the input in an undetermined state for an extended period. Since noise from other sources can cause perturbations on the input signals, jitter at the output of the receiver can result. The signals will appear as variable width pulses at the output of the receiver.

It is essential that the driver and receiver be located close to the edge of the PCB (printed circuit board) to reduce the mismatch between the board and the transmission line, and to minimize reflections and hence distortion, as the board is not matched to the line. Variations in the length of the board metallization can affect delays, and efforts should be made to maintain equal lengths. This delay introduced by the PCB is approximately 0.15 nsec/inch of metallization.

APPENDIX F: HIGH SPEED DATA PORT CONSIDERATIONS USING ECL

The 41LG Line Driver and 41LF Line Receiver are high speed parts suitable for use in the high speed data port. Optimum operation of these circuits requires proper termination of the transmission line or twisted pair. The 41LG is typically a 25 ohm line driver and is specified at 50 ohms.

The 41LG is an emitter follower line driver whose output signal swing is virtually independent of output current. The output voltage, both high and low, has a logarithmic variation with current.

When using the 41LG line driver it is necessary to terminate each output with a resistor in order that a discharge path be provided as shown in Figure E-1. For proper matching the driver end termination should be $Z_e = 2 * Z_o$ and not exceed $2.5 * Z_o$. When working into a line impedance, Z_o , the termination at the far end to minimize signal distortion should be Z_o .

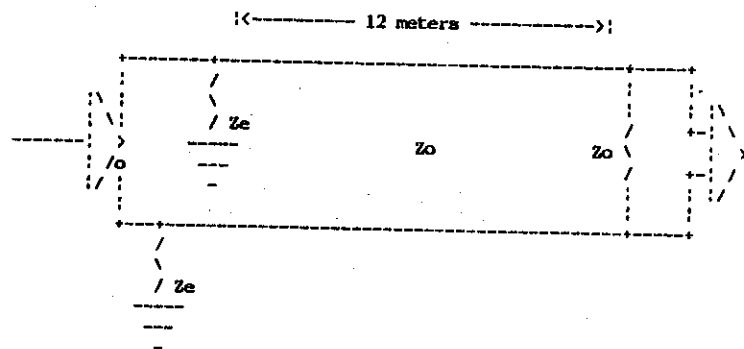


FIGURE E-1 APPLICATION ENVIRONMENT

The skew of the 41LG line driver is typically less than 0.5 nsec and guaranteed to be less than 1 nsec. The variation in skew between two identical paths in the same package is typically less than 0.2 nsec. The 41LG line driver and 41LF line receiver in combination have been shown to have less than 20% distortion of the output wave shape of the 41LF line receiver when operating with transmission lines up to 12 meters at 25 MHz as

APPENDIX G: RS-422 DIFFERENTIAL COMPONENTS

EIA Standard RS-422, commonly referred to as RS-422, defines the electrical characteristics of balanced voltage digital interface circuits. Drivers and receivers which meet the specifications of RS-422 have been in use for many years for high speed data transmission. RS-422 has been defined to ensure that transmission of data is done reliably and provides a "standard" set of characteristics for multiple vendors to design to, providing compatibility between vendor circuits.

G.1 CHARACTERISTICS

G.1.1 Driver

The balanced driver characteristics specified in RS-422 are as follows:

- o A driver circuit should result in a low impedance (100 Ohms or less) balanced voltage source that will produce a differential voltage to the interconnecting cable in the range of 2V to 6V.
- o With a test load of 2 resistors, 50 Ohms each, connected in series between the driver output terminals, the magnitude of the differential voltage (V_T) measured between the two output terminals shall be equal to or greater than 2V, or 50% of the magnitude of V_O , whichever is greater. For the opposite binary state the polarity of V_T is reversed (V_T).
- o During transitions of the driver output between alternating binary states, the differential voltage measured across 100 Ohm load shall monotonically change between 0.1 and 0.9 of VSS less than 0.1 of the unit interval or 20ns, whichever is greater. Thereafter, the signal voltage shall not change more than 10% of VSS from the steady state values until the binary state occurs.

G.1.2 Receiver

The electrical characteristics of the receiver specified in RS-422 are:

- o The receiver shall not require differential input voltage more than 200mV to correctly assume the intended binary state, over an entire common-mode voltage range of -7 to +7V. The common-mode voltage (V_{CM}) is defined as the algebraic mean of the 2 voltages appearing at the receiver input terminals with respect to the receiver circuit ground. This allows for operations where there are ground differences caused by IR drop and noise of up to +/- 7V.
- o The receiver shall maintain correct operation for a differential input signal ranging between 200mV and 6V in magnitude.
- o The maximum voltage between either receiver input terminal and receiver circuit ground shall not exceed 10V (3V_{signal} + 7V common-mode) in magnitude. Also, the receiver shall tolerate a maximum differential signal of 12V applied across its input terminals without being damaged.
- o The total load (up to 10 receivers) shall not have a resistance more than 90 Ohms at its input points.

G.2 RS-422 LINE DRIVERS/RECEIVERS AND ESDI

The ESDI specification permits any Line Driver/Receiver meeting RS-422 to meet the requirements except for the timing specifications. Selecting from available devices requires knowledge of the interface connection followed by a comparison between the data transfer timing specifications and manufacturers published data.

READ CLOCK and READ DATA are supplied by the drive when in read mode. REFERENCE CLOCK is supplied on the same pin as READ CLOCK when not in read mode.

REFERENCE CLOCK is accepted by the controller and is used to generate WRITE CLOCK and to synchronize WRITE CLOCK to WRITE DATA. READ/REFERENCE CLOCK and READ DATA are measured at the drive and WRITE CLOCK and WRITE DATA are measured at the controller.

For the controller, the READ/REFERENCE CLOCK received from the drive is guaranteed to meet the minimum/maximum timings. This signal will pass through the cable, a Line Receiver, appropriate controller circuitry, a Line Driver and the connector. It must then meet the WRITE CLOCK and WRITE DATA timings. Each element in this path causes some distortion.

The performance requirements of ESDI vary with data rate and also depends on whether the Low Speed or High Speed Port specifications are being designed to. RS-422 devices offer a range of guaranteed and typical performance levels which apply over a variety of operating conditions.

Implementors should be careful in selection of the proper components to ensure the optimum trade-off between cost and data errors due to pulse distortion. RS-422 devices tend to minimize the probability of data errors due to normal or common mode noise.

Skew is the key factor contributing to pulse distortion. It is also referred to as Differential Delay and is the difference between the circuit delay of a positive-going transition and that of a negative-going transition. It is usually designated T_{plh} - T_{phl} .

The effect of skew is to shorten or lengthen a pulse by the amount of the skew. In the ESDI document the allowable skew is the minimum (or maximum) original signal pulse width (e.g. READ/REFERENCE CLOCK) minus (or plus) the minimum (or maximum) pulse width at the end of the circuitry (WRITE CLOCK). The actual skew is the algebraic sum of the individual skews.

G.2.2 Allowable Skews

Table G-1 shows allowable skews (READ/REFERENCE CLOCK to WRITE CLOCK) for selected data rates.

		Data Rate in Megabits/second				
		5	10	15	20	24
TA (1/F)	(nsec)	200.0	100.0	66.7	50.0	41.60
Low Speed Port						
tB min. (.45tA)	(nsec)	90.0	45.0	30.0	22.5	18.70
tC min. (.21tA)	(nsec)	42.0	21.0	14.0	10.5	8.73
Total skew	(nsec)	48.0	24.0	16.0	12.0	9.96
High Speed Port						
tB min. (.408tA)	(nsec)	81.6	40.8	27.2	20.4	17.0
tD min. (.312tA)	(nsec)	62.4	31.2	20.8	15.6	13.0
Total skew	(nsec)	19.2	8.6	6.4	4.8	4.0

TABLE G-1 ALLOWABLE SKEW

G.2.2 RS-422 DEVICES

Table G-2 lists data for RS-422 devices appropriate for ESDI applications:

PART NO/TYPE	SKEW (nsec)		TEMP	PWR SUP	STATUS
	(max)	(typ)			
XX3486 Quad Receiver	NA	NA	25C	5V	ind.std
XX3487 Quad Driver	NA	2	25C	5V	ind.std
XX26LS31 Quad Driver	6	2	25C	5V	ind.std
XX26LS32 Quad Rcvr	NA	NA	25C	5V	ind.std
DS8921 Single Xcvr					
Receiver	5.0	0.5	0-70C	5V+-10%	NSC
Driver	6.0	0.5	"	"	"
DS8922/23 Dual Xcvr					
Receiver	5.0	0.5	0-70C	5V+-10%	NSC
Driver	6.0	0.5	"	"	"
DS8921A Single Xcvr					
Receiver	3.5	0.5	0-70C	5V+-10%	NSC
Driver	2.75	0.5	"	"	"
DS8922/23A Dual Xcvr					
Receiver	3.5	0.5	0-70C	5V+-10%	NSC
Driver	2.75	0.5	"	"	"

TABLE G-2. APPROPRIATE RS-422 DRIVERS/RECEIVERS/TRANSCIEVERS

G.3 Calculation of Allowable Skew

As an example of calculating allowable controller circuitry skew:

Assume: 24Mbits/Sec Data rate
0.25nsec cable skew
2X typical transceiver performance
High Speed Port specification.

Allowable skew = $4.0 - (0.25 + 1.0 + 1.0)$ nsec
= $4.0 - 2.25$ nsec
= 1.75 nsec

APPENDIX H: DISK SERVOS

Due to the very tight tolerances of high performance disk drives, the drive must actively follow the data track around the disk. This technique is called "servoing". To do this, information is permanently written on the drive which controls the servo process. This data is commonly referred to as "servo data". In addition to track following, servo data may be used to control clocking, data recovery, verify seek position, and other functions within the drive. These functions are invisible to an outside controller.

Disk servoing techniques affect the way in which data may be written on the disk and may restrict the sector sizes used or number of sectors available on a track.

Disk servo techniques may be divided into several broad classes:

- o No servo data: Some low performance hard disks do not use any recorded servo data. Instead, they depend on mechanical tolerances being tight enough to follow the track. Most floppy disks also use this technique. There is no restriction on where data is written.

- o Dedicated servo surface: One surface of the disk is used to hold servo information. There is no servo information on data tracks and data may be written anywhere on the disk. This technique is normally only used on magnetic disks, as optical disks don't have dedicated surfaces.

- o Servo embedded continuously around the disk: In this scheme servo data is frequently written around the disk. The drive makes this invisible to the controller by extending the read clock (extended periods). This technique is invisible to the controller, though the drive may place some restrictions on where data is written. This method is more common with optical disks.

- o Servo embedded every sector: There is servo information written on the disk in the inter sector gap of each sector. The sector size and number of sectors per track is typically fixed.

- o Servo embedded less frequently: The drive embeds servo information periodically around the disk but not necessarily every sector (e.g. four servo regions per track). More than one sector fits between servo regions as a rule. A simple version of this technique is a single servo area at the index pulse, in which case servo area is invisible to the controller.

- o Combinations of servo techniques: Various combinations of the above techniques are also possible, e.g. a dedicated servo may be combined with servo data embedded each sector. When combinations exist, the form most restrictive on data formatting dictates how the drive may be used.

The primary impact on the interface of the servo technique used in the drive is on head switch times and how data may be written on the disk.

Head switch times: Drives which do not use servo data embedded on data tracks can typically switch heads within the same cylinder in a very short time (< 15usec). Drives with embedded servo data typically require a longer head switch time, around 1 ms to 3 ms. ESDI reports this fact via a general configuration bit. For long head switch times the drive must negate COMMAND COMPLETE until the head switch is complete and the controller must wait for

COMMAND COMPLETE assertion before operating on the new track.

Where data may be written on the disk: This depends on the servoing technique. As a general rule, data may not be written in a servo area so the drive must prevent the controller from doing so. This is done by controlling the ISG (Inter Sector Gap) size and the Unformatted Bytes per Sector and Sectors per Track permitted by the drive. This discussion primarily applies to magnetic disks, as optical disk drives control the disk format either through pre-formatted media or through the format command.

The restrictions, if any, are generally divided between no servo data on the track, servo embedded every sector, and servo embedded less often.

o No servo on data track: When there is no servo information on the data track, the controller has complete control over disk format. Both hard and soft sector formats may be used and any sector size and sectors per track may be used e.g. if a drive of this type has 20,000 unformatted bytes per track it might be able to be formatted with any of:

sectors per track	unformatted bytes per sector	formatted bytes per sector
9	2104	2048
18	1080	1024
35	568	512
64	312	256

o Servo embedded every sector: When servo is embedded in every sector, the servo area is reported by the drive as the ISG and ISG After Index parameters. The drive prohibits writing to this area and the controller takes this into account when formatting the drive. Embedded servo drives of this type are hard sectored and typically support a single sector size e.g. a drive of this type might have 20,000 unformatted bytes per track which are divided into 32 sectors of 625 unformatted bytes each, for use with a 512 byte formatted sector size.

o Servo embedded less frequently: When servo is embedded less frequently then the drive must ensure that all servo areas are within ISGs but may generate sector pulses more frequently. Drives of this type are hard sectored.

The drive can support more than one sector between servo areas and protect the servo areas either by making all track ISGs large enough to encompass the servo information or by using an invisible gap for servo areas. Since controllers typically ignore the area between the end of a data field and the next sector or index pulse, the drive can report a small ISG between normal sectors then extend the ISG invisibly when servo information is present e.g. suppose a drive has four servo regions with 5000 bytes between each region.

Using the second method above (invisibly extended ISG), the drive will report 20,000 Unformatted Bytes per Track (space exclusive of servo regions). The unformatted sector size is computed assuming a track of 20,000 bytes with an integral number of sectors fitting within each 5000 byte area e.g.

1,125 byte Unformatted Sector Size for 1024 byte sectors.

The drive generates INDEX/SECTOR pulses every 1125 bytes, but delays the SECTOR/INDEX pulse at each servo area to the end of the servo area. This extra ISG is invisible to the controller and not reported back to it. This allows a small ISG e.g. 16 bytes, to be used with an invisible ISG of 32 or 48 bytes for servo data.

In this example, the drive might support the following standard sector sizes:

sectors per track	unformatted bytes per sector	formatted bytes per sector
4	5000	4096
8	2500	2048
16	1250	1024
32	625	512
56	357	256

Operation of the Set Unformatted Bytes per Sector command (if implemented) is different from non-embedded servo drives. The controller cannot assume that Sectors per Track equals Unformatted Bytes per Track divided by Unformatted Bytes per Sector. Instead, the drive must adjust sectors per track so an integral number of sectors fit between each servo region. In order to handle the defect list, in this situation the drive should adjust Unformatted Bytes per Sector to evenly spread the requested sectors within the servo area. Otherwise a large remainder could exist in each servo area which complicates computations of the sector each defect is part of.

Because of these restrictions, the controller should verify the Unformatted Bytes per Sector and Sectors per Track actually used by the drive after issuing a Set Unformatted Bytes per Sector. When the defect list is written it ignores servo areas when reporting defect offsets from index. This is consistent with reporting track sizes exclusive of servo areas. Defect locations may be in error by a few bytes per servo area, however the controller should be able to identify the sector each defect is contained within.

APPENDIX 1: DETERMINING READ GATE ASSERTION/NEGATION TIMES

1.1 DEFINITIONS

The following times are defined to simplify formula writing. All values are measured in bit times:

T_{rdd} = Read Data Delay: as defined in the standard.

T_{wdd} = Write Data Delay: as defined in the standard.

T_{rgw} = READ GATE Window Size: as defined in the standard.

T_{was} = Write Splice Size: as defined in the standard.

T_{nds} = Mark Detection Skew: as defined in the standard.

T_{cra} = Controller READ GATE assertion delays: Time measured from detecting a mark to asserting READ GATE.

T_{crn} = Controller READ GATE negation delays: Time measured from latching the last READ DATA bit from the interface (assertion of READ CLOCK) to the negation of READ GATE.

T_{cwa} = Controller WRITE GATE assertion delays: Time measured from detecting a mark to asserting WRITE GATE.

T_{can} = Controller ADDRESS MARK ENABLE negation delays: Time measured from detecting a mark to negating ADDRESS MARK ENABLE.

T_{cws} = Controller write splice size: Time that WRITE GATE is negated between the address area and the data area.

1.2 READ GATE NEGATION (PAD SIZE)

There is a minimum pad size that is required to be written by the controller that will ensure that READ GATE will not be asserted during the time that WRITE GATE was previously asserted or negated. This pad requirement results from the fact that drives have certain delays associated with recording and reading data:

- o When data is written to the media, there are delays between the time when the controller sent the data and the time it is actually recorded on the media.

- o There is another delay from when data is detected on the media to the time when the controller gets it across the interface.

- o The controller may have additional delays that cause it to delay negating READ GATE after receiving the last valid data bit (typically check bits).

The minimum pad size that a controller must record is defined by the formula below:

$$T_{pad} = T_{rdd} + T_{wdd} + T_{crn} \quad (\text{measured in bit times})$$

1.3 READ GATE ASSERTION FOR ADDRESS AREA READ

1.3.1 READ GATE Assertion for Alternative 1 Hard Sector Format

READ GATE assertion time is determined by ADDRESS MARK ENABLE and Mark Detection Skew times. The minimum time from the assertion of INDEX or SECTOR to the assertion of READ GATE is:

$$T_{rgmin} = T_{nds} + T_{can}$$

The maximum time, if there is a Read Gate Window time is:

$$T_{rgmax} = T_{nds} + T_{can} + T_{rgw}$$

Whether or not there is a Read Gate Window time, the controller must turn on READ GATE so that the drive reads the required number of PLO Sync bytes (as defined in configuration data) before detecting the first bit of the header or data area. The controller must also consider any detection delays (T_{cra}) in determining when to assert READ GATE.

1.3.2 READ GATE Assertion for Alternative 2 Hard Sector Format

READ GATE assertion time is determined by Mark Detection Skew, Write Splice Size, and WRITE GATE assertion times. The minimum time from the assertion of INDEX or SECTOR to the assertion of READ GATE is:

$$T_{rgmin} = T_{nds} + T_{cwa} + T_{was}$$

The maximum time, if there is a Read Gate Window time is:

$$T_{rgmax} = T_{nds} + T_{cwa} + T_{was} + T_{rgw}$$

Whether or not there is a Read Gate Window time, the controller must turn on READ GATE so that the drive reads the required number of PLO Sync bytes (as defined in configuration data) before detecting the first bit of the header or data area. The controller must also consider any detection delays (T_{cra}) in determining when to assert READ GATE.

1.3.3 READ GATE Assertion for Soft Sector Format

READ GATE assertion time is defined in Section 10 and is not discussed here.

The controller must turn on READ GATE so that the drive reads the required number of PLO Sync bytes (as defined in configuration data) before detecting the first bit of the header or data area. The controller must also consider any detection delays (T_{cra}) in determining when to assert READ GATE.

1.4 READ GATE Assertion for Data Area Read

Data area reads are keyed off the end of the address area READ DATA, rather than a mark signal. READ GATE assertion times are determined by Read and Write Data Delays, and Write Splice Size. The minimum time from the negation of READ GATE for the address area and the assertion of READ GATE for the data area is:

$$T_{rgmin} = T_{pad}(\text{actual}) + T_{cws} + T_{was}$$

The maximum time, if there is a Read Gate Window time is:

$$T_{\text{rmax}} = T_{\text{pad}} (\text{actual}) + T_{\text{CWS}} + T_{\text{WSS}} + T_{\text{rGW}}$$

Whether or not there is a Read Gate Window time, the controller must turn on READ GATE so that the drive reads the required number of PIO Sync bytes (as defined in configuration data) before detecting the first bit of the header or data area.

APPENDIX J: MAGNETIC DISK FORMATTING

After a format of the disk, either a normal format or a Defect List format, read operations before the first write operation may fail if the header field is delayed by the value specified as the Read Data Propagation Delay. This delay requires some compensation in order to reliably read data fields before the first update write operation.

There are two common solutions to the problem.

J.1 Two Pass Format

Immediately after the format operation (done with an extended PIO Sync field after the header), all the data fields are rewritten with the desired PIO Sync field size. This compensates for the read channel delay.

J.2 One Pass Format with Read Channel Compensation

During the format operation a pad is added between the header and the data field, usually while WRITE GATE is negated. This pad length is equal to the Read Data Propagation Delay. This has the effect of causing the write splice to be moved to a position consistent with that produced by an update write operation.

J.3 Defect List Format

One of these two methods should be used by the manufacturer to record the Defect List.