Quantum®

Maverick 270/540S
Product Manual

August 1994
81-106845-02

Quantum reserves the right to make changes and improvements to its products, without incurring any obligation to incorporate such changes or improvements in units previously sold or shipped.

You can request Quantum publications from your Quantum Sales Representative, or order them directly from Quantum.

Publication Number: 81-106845-02

UL/CSA/VDE/TUV

UL standard 1954 recognition granted under File No. E78016
CSA standard C22.2-950 certification granted under File No. LR49896
VDE 0805 (IEC 950) certification granted under File No. 11342-3250-2003/A1A
TUV Rheinland EN 60 950
Tested to FCC Rules for Radiated and Conducted Emissions, Part 15, Sub Part J, for Class-B Equipment.

SERVICE CENTERS

Quantum Service Center
715 Sycamore Avenue
Milpitas, California 95035
Phone: (408) 894-4000
Fax: (408) 894-3218

Quantum GmbH
Genferstrasse 4B
D-6000 Frankfurt am Main 56
West Germany
Phone: (49) 69-509-1080
Fax: (49) 69-509-10891

Quantum Asia-Pacific Pte. Ltd.
50 Tagore Lane #b1-04
Singapore, 2678
Phone: (65) 450-9333
Fax: (65) 452-2544

PATENTS

Quantum products in this manual are covered or licensed under one or more of the following U.S. patents: 4,396,959; 4,419,701; 4,490,635; 4,516,177; 4,538,099; 4,538,193; 32,075; 4,593,194; 4,625,109; 4,639,798; 4,639,863; 4,647,769; 4,647,997; 4,660,106; 4,661,696; 4,669,004; 4,675,652; 4,692,829; 4,703,176; 4,712,146; 4,730,321; 4,772,974; 4,783,705; 4,814,909; 4,819,153; 4,825,321; 4,858,034; 4,860,194; 4,872,074; 4,882,671; 4,914,530; 4,920,434; 4,920,442; 4,982,296; 311,737; 5,005,089; 5,027,241; 5,031,061; 5,041,924; 5,062,023; 5,065,262; 5,084,791; 5,097,439; 5,119,254; 5,128,813; 5,160,865; 5,170,299; 5,177,771; 5,187,621; 5,199,033; 5,208,712; 5,208,713; 5,210,474; 5,216,582; 5,227,671; 5,227,930; 5,231,545; 5,241,546; 5,253,129; 5,255,136; Other U.S. and Foreign Patents Pending.

Table of Contents

List of Figures

List of Tables

Chapter 1
ABOUT THIS MANUAL

This chapter gives an overview of the contents of this manual, including the intended audience, manual organization, and terminology and conventions. In addition, it provides a list of other references that might be helpful to the reader.

1.1 AUDIENCE DEFINITION

The Quantum Maverick™ 270/540S product manual is intended for several audiences, including the original equipment manufacturer (OEM), distributor, installer, and end user.

1.2 MANUAL ORGANIZATION

This manual provides you with information about installation, principles of operation, and interface command implementation. It is organized into the following chapters:

- Chapter 1—About This Manual
- Chapter 2—General Description
- Chapter 3—Installation
- Chapter 4—Specifications
- Chapter 5—Basic Principles of Operation
- Chapter 6—SCSI Description

In addition, this manual contains a glossary of terms and an index to help you locate important information.

1.3 TERMINOLOGY AND CONVENTIONS

The abbreviations listed below are used in this manual.

- bpi Bits per inch
- dB Decibels
- dBA decibels, A weighted
- fci Flux changes per inch
- Hz Hertz
- K Kilobytes
- lsb Least-significant bit
- mA Milliampere

- MB Megabytes (1 MB = 1,000,000 bytes when referring to disk storage and 1,048,576 bytes in all other cases)
- m Meter
- Mbit/s Megabits per second
- MB/s Megabytes per second
- MHz Megahertz
- ms Millisecond
- msb Most-significant bit
- mV Millivolts
- ns Nanoseconds
- tpi Tracks per inch
- µs Microseconds
- V Volts

The following conventions are used in this manual:

- **Commands and Messages**

  Commands and messages, sent between the drive and the host, are listed in all capitals. For example:

  WRITE LONG
  ILLEGAL COMMAND

- **Parameters**

  Parameters are given as initial capitals when spelled out and as all capitals when abbreviated. For example:

  Prefetch Enable
  Cache Enable
  PE
  CE

- **Names of Bits and Registers**

  Bit names and register names are presented in initial capitals. For example:

  Host Software Reset
  Alternate Status Register

- **Hexadecimal Notation**

  The hexadecimal notation "H" is given in subscript form. For example:

  30_H

- **Signal Negation**

  A signal name that is defined as active low is listed with a minus sign following the signal. For example:

  RD–

- **Notes**

  Notes are used after tables to provide you with supplementary information.

- **Host**

  In general, the system in which the drive resides is referred to as the host. The SCSI host adapter is considered to be part of the host.

1.4 REFERENCE

For additional information about the SCSI interface, refer to:

- ANSI Small Computer System Interface-2 (SCSI-2) Specification, ANSI X3T9.2/86-109, Revision 10k.
- ANSI Small Computer System Interface-3 (SCSI-3) Specification, ANSI X3T9.2/855D, Revision 12B, June 7, 1993.

Chapter 2
GENERAL DESCRIPTION

This chapter summarizes general functions and key features of the Maverick 270/540S hard disk drive, as well as the standards and regulations it meets.

2.1 PRODUCT OVERVIEW

Quantum's Maverick 270/540S hard disk drives are part of a family of high-performance, 1-inch-high, hard disk drives manufactured to meet the highest product quality standards.Maverick drives use nonremovable, 3 1/2-inch hard disks.

The Maverick 270/540S hard disk drives feature an embedded SCSI drive controller and use SCSI commands to optimize system performance. The drive manages media defects and error recovery internally, so these operations are transparent to the user.

The innovative design of Maverick 270/540S hard disk drives enables Quantum to produce a family of low-cost, high-reliability drives.

2.2 KEY FEATURES

Maverick 270/540S hard disk drives include the following key features:

General

- Formatted storage capacity of 270.7 or 541.5 MB
- Industry-standard, 3 1/2-inch form factor
- Low-profile, 1-inch height
- Embedded SCSI controller
- Embedded servo design

Performance

- Average seek time of 14 ms
- Data transfer rate of up to 6.0 MB/s asynchronous and 10.0 MB/s synchronous
- Average rotational latency of 8.33 ms
- 1:1 interleave on read/write operations
- Proprietary 96K "look ahead" DisCache® with adaptive segmentation, continuous prefetch, and WriteCache™ write-buffering capabilities
- SCSI-bus active termination
- SCSI-bus active negation drivers

Reliability

- Projected Field Mean Time Between Failure (MTBF): 300,000 hours
- 96-bit, interleaved Reed-Solomon Error Correcting Code (ECC), with additional 16-bit cross-checking and double burst correction for bursts up to 24 bits in length
- Automatic retry on read errors
- Transparent media defect mapping
- Reassignment of defective sectors discovered in the field, without reformatting
- High performance, in-line defective sector skipping
- Patented AIRLOCK® automatic shipping lock and dedicated landing zone

Versatility

- Power-saving modes
- Downloadable firmware

2.3 STANDARDS AND REGULATIONS

Maverick 270/540S hard disk drives satisfy the following standards and regulations:

- Federal Communications Commission (FCC): FCC Rules for Radiated and Conducted Emissions, Part 15, Sub Part J, for Class B Equipment.
- Underwriters Laboratory (U.L.): Standard 1950. Information technology equipment including business equipment.
- Canadian Standards Association (CSA): Standard C22.2 No. 950-M89. Information technology equipment including business equipment.
- European Standards – Verband Deutscher Electroechnier (VDE) and Technisher Uberwachungs Verein (TUV): Standard EN 60 950. Information technology equipment including business equipment.

2.4 HARDWARE REQUIREMENTS

Maverick 270/540S hard disk drives are compatible with host computers and controllers that provide a SCSI and SCSI-2 interface.

Chapter 3
INSTALLATION

This chapter explains how to unpack, configure, mount, and connect the Maverick 270/540S hard disk drive prior to operation. It also explains how to start up and operate the drive.

3.1 SPACE REQUIREMENTS

Quantum ships the Maverick 270/540S hard disk drive without a faceplate. Figure 3-1 shows the external dimensions of the drive.

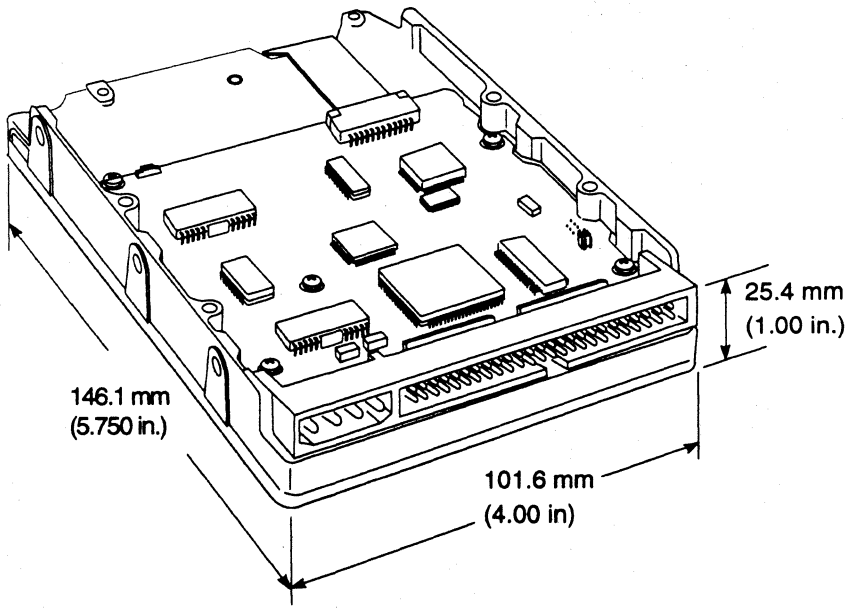

Figure 3-1 *Maverick 270/540S Mechanical Dimensions*

3.2 UNPACKING INSTRUCTIONS

CAUTION: The maximum limits for physical shock can be exceeded if the drive is not handled properly. Special care should be taken not to bump or drop the drive.

1. Open the shipping container.
2. Remove the upper protective packaging pad from the box.

   Figure 3-2 shows the packaging for the Maverick 270/540S hard disk drive in a 1-pack shipping container. (A 12-pack shipping container is available for multiple drive shipments and has the same type of protective packaging.)

3. Remove the drive from the box.

CAUTION: During shipment and handling, the drive is packed in an antistatic electrostatic discharge (ESD) bag to prevent electronic component damage due to electrostatic discharge. Remove the drive from the ESD bag only when you are ready to install it. To avoid accidental damage to the drive, do not use a sharp instrument to open the ESD bag.

4. Save the packaging materials for possible future use.

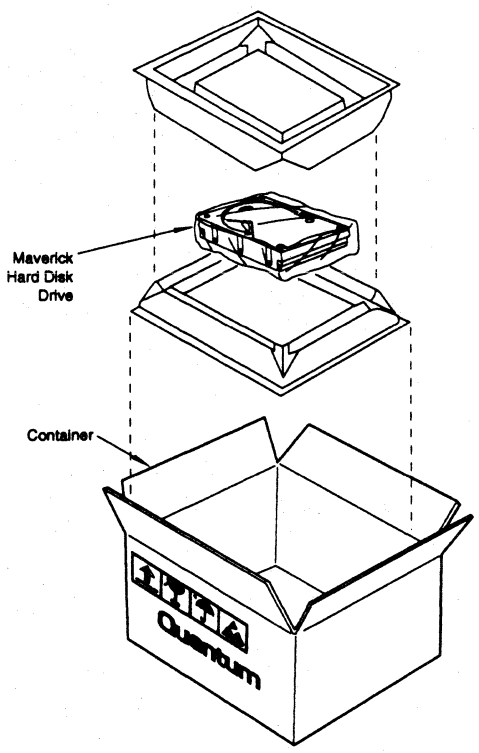

Figure 3-2 *Packaging for a 1-Pack Shipping Container*

3.3 SCSI-BUS DEVICE IDENTIFICATION

The Maverick 270/540S hard disk drive SCSI address depends on the host system in which it is to be installed. Figure 3-3 shows the printed circuit board assembly (PCBA), indicating the jumpers that control the hard drive's SCSI address.

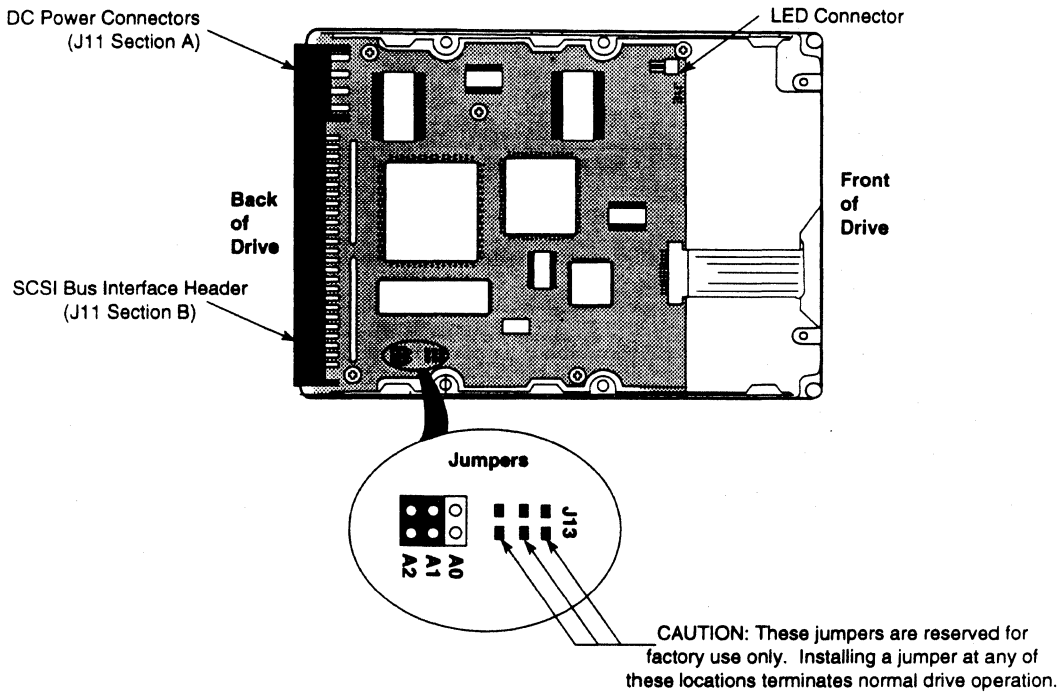

Figure 3-3 *Jumper Pin and Terminator Locations on the Drive PCBA*

On the drive's printed circuit board assembly (PCBA), there are three pairs of jumper pins that allow you to set the drive's SCSI address. When you change a jumper setting, the new setting is in effect when the drive is next powered on.

Used in combination, the jumper settings across pins A2, A1, and A0 determine the Maverick 270/540S hard disk drive's SCSI-bus device identification (SCSI ID). By default, Quantum configures the drive with a SCSI ID of six—that is, with jumpers installed across the pins labeled A2 and A1 and no jumper installed across the pins labeled A0. Table 3-1 shows the default address settings for the drive.

Table 3-1 *SCSI ID Jumper Settings*

JUMPER SETTINGS			SCSI ID
A2	A1	A0	
OFF	OFF	OFF	0
OFF	OFF	ON	1
OFF	ON	OFF	2
OFF	ON	ON	3
ON	OFF	OFF	4
ON	OFF	ON	5
ON*	ON*	OFF*	6*
ON	ON	ON	7

Note: ON indicates a jumper is installed. OFF indicates that no jumper is installed. The asterisk (*) indicates the default setting.

The SCSI bus supports up to eight devices, including the host system. A device's identification number determines its priority and must be unique in the system.

3.4 TERMINATING RESISTORS

At the factory, Quantum installs two resistor networks in sockets on the drive PCBA to terminate the SCSI bus. Figure 3-4 shows the location of the terminating resistors on the PCBA. Only the first device—usually the host—and the last device on a SCSI bus should contain terminating resistors. When installing the Maverick 270/540S hard disk drive in any other position on the SCSI bus, remove the terminating resistors.

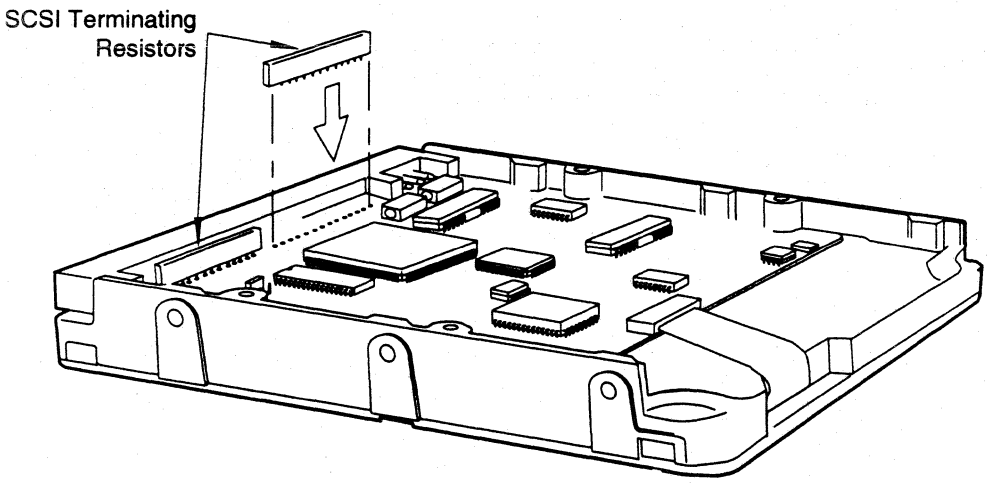

Figure 3-4 *SCSI Terminating Resistor Locations*

3.5 MOUNTING

Drive mounting orientation, clearance, and ventilation requirements are described in the following subsections. For mounting, #6-32 UNC screws are recommended. To avoid stripping the mounting-hole threads, the maximum torque applied to the screws must not exceed 8 inch-pounds.

3.5.1 Orientation

The mounting holes on the Maverick 270/540S hard disk drive allow the drive to be mounted in any orientation. Figure 3-5 shows the location of the three mounting holes on each side of the drive.

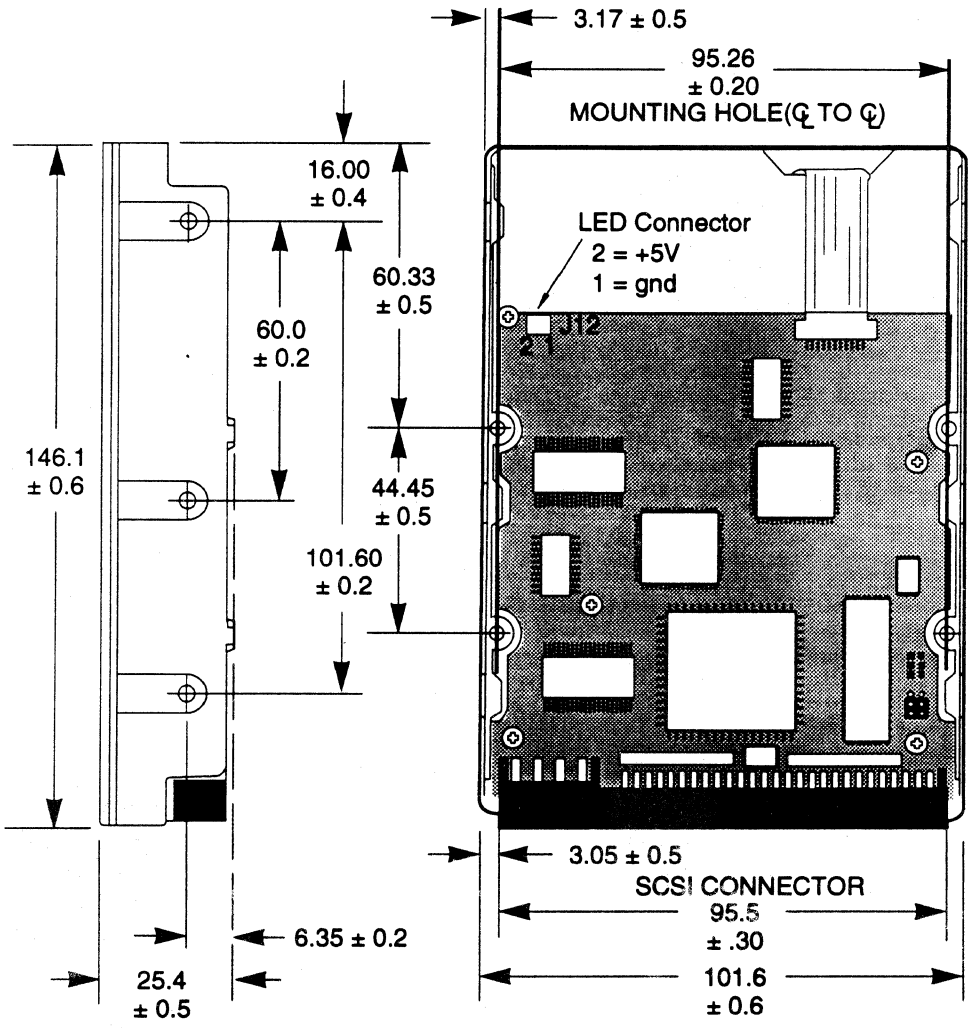

Figure 3-5 *Maverick 270/540S Mounting Dimensions (in Millimeters)*

3.5.2 Clearance

The printed-circuit board assembly (PCBA) is very close to the mounting holes. Clearance from the drive to any other surface—except mounting surfaces—must be 1.25mm (0.05) inches minimum. Figure 3-6 specifies the clearance between the screws in the mounting holes and the PCBA. Do not use mounting screws longer than the maximum lengths specified in Figure 3-6. The specified screw length allows full use of the mounting-hole threads, while avoiding damaging or placing unwanted stress on the PCBA.

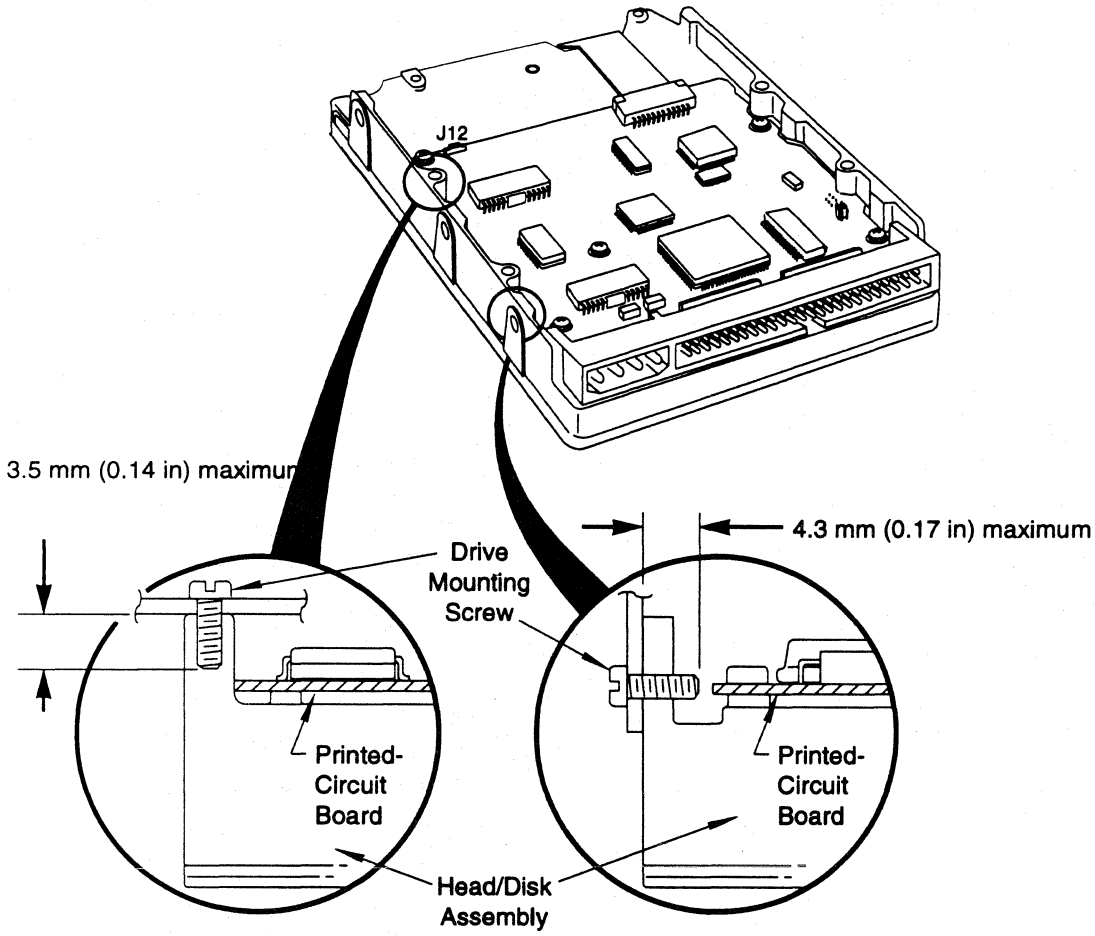

Figure 3-6 *Mounting Screw Clearance*

CAUTION: The printed-circuit board (PCB) is very close to the mounting holes. Do not exceed the specified length for the mounting screws. The specified screw length allows full use of the mounting-hole threads, while avoiding damaging or placing unwanted stress on the PCB. Figure 3-6 specifies the minimum clearance between the PCB and the screws in the mounting holes. To avoid stripping the mounting-hole threads, the maximum torque applied to the screws must not exceed 8 inch-pounds.

3.5.3 Ventilation

The Maverick 270/540S hard disk drive operates without a cooling fan, provided the ambient air temperature does not exceed 122°F (50°C).

3.6 CABLE CONNECTOR

The Maverick 270/540S has a combination DC/SCSI cable connector as shown in Figure 3-7. The drive's combination DC/SCSI connector part number is Dupont 90910-001, 2 in 1 SMT.

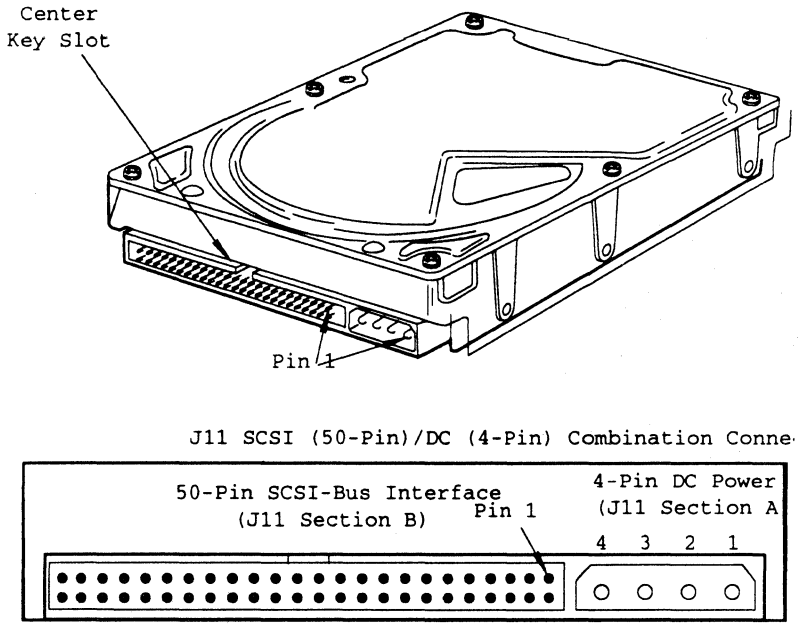

Figure 3-7 *DC Power and SCSI Bus Combination Connector (J11)*

Recommended header connector and strain relief:

50-pin connector (base part	AMP P/N 1-746195-2 or equivalent
50-pin connector strain relief	AMP P/N 499508-2, or equivalent

Recommended mating connectors and their part numbers:

4-pin connector (J11 Section A)	AMP P/N 1-480424-0 or equivalent
Loose-piece contacts	AMP P/N 61173-4 or equivalent
Strip contacts	AMP P/N 350078-4 or equivalent

Table 3-2 shows the power line designations for the pins on the DC power portion (J11 Section A) of the combination connector.

Table 3-2 *Power Connector J11 Section A Pin Assignments*

PIN NUMBER	POWER LINE DESIGNATION	MATING CONNECTOR TYPE AND PART No. (OR EQUIVALENT)
1	+12V DC	4-Pin Connector: AMP P/N 1-480424-0 Loose piece contacts: AMP P/N 61173-4 Strip Contacts: AMP P/N 350078-4
2	+12V RETURN (Ground)	
3	+5V RETURN (Ground)	
4	+5V DC	

Note: Labels indicate the pin numbers on the connector. Pins 2 and 3 are connected together on the drive.

3.6.1 SCSI Connector (J11 B)

The SCSI cable connector (J11 B) on the Maverick 270/540S hard disk drive is a 50-pin universal header, as shown in Figure 3-7. A key slot on the header prevents the incorrect installation of the header connector on the SCSI cable. To prevent the header connector on the SCSI cable from being plugged in backwards, use only a keyed header connector.

Recommended header connector and strain relief:

- 50-pin connector (base part) AMP P/N 1-746195-2, or equivalent
- 50-pin connector strain relief AMP P/N 499508-2, or equivalent

See Chapter 6, "SCSI Description," for more detailed information about the required SCSI signals.

3.6.2 LED Connector (J12)

The LED connector may be used to connect an external LED for the purpose of monitoring drive activity. A 220-ohm resistor is connected in series from the +LED pin to +5 volts. The -LED pin is connected to a switching transistor on the drive that turns on and lights the LED when the disk is performing a read or write operation. When no read or write operation is taking place, the transistor turns off and the LED is extinguished. The recommended mating connector parts are shown in Table 3-3.

Table 3-3 *J12 LED Connector*

J12 Pins		Mating Connector Type and Part Number (or equivalent)
Pin Number	Function	
1	–LED	Two-position housing: Molex P/N 51021-0200 Loose-piece contacts: Molex P/N 50058-8100
2	+LED	

Chapter 4
SPECIFICATIONS

This chapter gives a detailed description of the physical, electrical, and environmental characteristics of the Maverick 270/540S hard disk drive.

4.1 SPECIFICATION SUMMARY

Table 4-1 gives a summary of the Maverick 270/540S hard disk drive.

Table 4-1 *Maverick 270/540S Specifications*

Description	Maverick 270S	Maverick 540S
Capacity (formatted)	270.7 MB	541.5 MB
Nominal rotational speed (rpm)	3,600 ±0.2%	3,600 ±0.2%
Number of disks	1	2
Number of R/W heads	2	4
Data Organization		
Zones per surface	16	16
Tracks per surface	2,853	2,853
Total tracks	5,706	11,412
Sectors per track		
Inside zone	58	58
Outside zone	118	118
Total User Sectors	528,879	1,057,758
Bytes per sector	512	512
Number of tracks per cylinder	2	4
Recording		
Recording technology	Multiple zone	Multiple zone
Maximum linear density	45,594 fci	45,594 fci
Maximum Recording Density	60,792 (bpi)	60,792 (bpi)
Encoding method	RLL 1,7	RLL 1,7
Interleave	1:1	1:1
Track density	2,950 tpi	2,950 tpi
Maximum effective areal density	134 Mbit/in^2	134 Mbit/in^2

Table 4-2 *Maverick 270/540S Specifications (Continued)*

Description		Maverick 270S	Maverick 540S
Performance			
Typical Seek times[1]			
Average seek		14 ms	14 ms
Track-to-track		5.0 ms	5.0 ms
Full stroke		28 ms	28 ms
Sequential Head Switch		4.5 ms	4.5 ms
Rotational Latency		8.33 ms	8.33 ms
Data Transfer rate			
Disk to Read Buffer[2]		18.67 to 36.08 Mbit/s	18.67 to 36.08 Mbit/s
SCSI-2	Asynchronous	6.00 MB/s max.	6.00 MB/s max.
	Synchronous	10.0 MB/s max.	10.0 MB/s max.
Buffer size (The upper 32K is used for firmware)		128 KB	128 KB
Projected Field MTBF		300,000 hrs	300,000 hrs
Contact Start/Stop Cycles		20,000	20,000
Auto head-park method		AirLock	AirLock

1. Seek times are at nominal conditions and include settling.
2. Disk to read buffer transfer rate is zone-dependent.

4.2 FORMATTED CAPACITY

At the factory, the Maverick 270/540S receives a low-level format that creates the actual tracks and sectors on the drive. Table 4-3 shows the capacity resulting from this process. For operation with DOS, UNIX, or other operating systems, formatting done at the user level results in less capacity than the physical capacity shown in Table 4-3.

Table 4-3 *Formatted Capacity*

	Maverick 270S CAPACITY	Maverick 540S CAPACITY
Number of bytes	270,786,048	541,572,096
Number of 512-byte sectors	528,879	1,057,758

4.3 DATA TRANSFER RATES

Data is transferred from the disk to the read buffer at a rate up to 3.5 MB/s in bursts. Data is transferred from the read buffer to the SCSI bus at a rate of up to 6.0 MB/sec in the asynchronous mode, or at up to 10.0 MB/s in the synchronous mode.

4.4 TIMING SPECIFICATIONS

Table 4-4 illustrates the timing specifications of the Maverick 270/540S hard disk drive.

Table 4-4 *Maverick 270/540S Timing Specifications*

Parameter	Typical Nominal[1]	Maximum Worst Case[2]
Single Track Seek[3]	5.0 ms	6.5 ms
Sequential Cylinder Switch Time[4]	4.5 ms	6.5 ms
Sequential Head Switch Time[5]	4.5 ms	6 ms
Random Average (Read or Seek)[6]	14 ms	18 ms
Random Average (Write)	16 ms	20 ms
Full-Stroke Seek	28 ms	30 ms
Average Rotational Latency	8.33 ms	—
Power On[7] to Interface Ready[8]	10 seconds	16 seconds
Shutdown[9] to Drive Ready[10]	10 seconds	16 seconds
SCSI "Hard Reset Time"[11]	150 ms	—
Spindown Time	4.5 seconds	

1. Nominal conditions are as follows:
   - Nominal temperature (25° C)
   - Nominal supply voltages (12.0V, 5.0V)
   - No applied shock or vibration
2. Worst-case conditions are as follows:
   - Worst-case temperature extremes (4°C to 50°C)
   - Worst-case supply voltages (12V ±10%, 5V±5%)
3. Seek time is defined as the time required for the actuator to seek and settle on-track. It is measured by averaging 5000 seeks of the indicated type as shown in this table. The seek times listed include head settling time, but do not include command overhead time or rotational latency delays.
4. Sequential Cylinder Switch Time is the time from the conclusion of the last sector of a cylinder to the first logical sector on the next cylinder.
5. Sequential Head Switch Time is the time from the last sector of a track to the beginning of the first logical sector of the next track of the same cylinder.
6. Average seek time is the average of 5000 random seeks. When a seek error occurs, recovery for that seek can take up to seven seconds.
7. Power On is the time from when the supply voltages reach operating range to when the drive is ready to accept any command.
8. Interface Ready is the condition in which the drive is ready to accept any command, before the disks are rotating at rated speed.
9. Shutdown is the mode where the microprocessor is powered, but not the HDA. When the system sends the drive a shutdown command, the drive parks the heads off the data zone and spins down to a complete stop.
10. Drive Ready is the condition in which the disks are rotating at the rated speed and the drive is able to accept and execute commands requiring disk access without further delay.
11. SCSI "Hard Reset Time" is the time from Reset to Selection.

4.5 POWER

The Maverick 270/540S hard disk drive operates from two supply voltages:

- +12V ±10%
- +5V ±5%

The allowable ripple and noise for each voltage is 100 mV for the +12 Vdc supply and 50 mV for the +5 Vdc supply.

4.5.1 Power Sequencing

You can apply the power in any order or manner, or short or open either the power or power-return line with no loss of data or damage to the disk drive. However, data may be lost in the sector being written at the time of power loss. The drive can withstand transient voltages of +10% to –100% from nominal while powering up or down.

4.5.2 Power Reset Limits

When powering up, the drive remains reset until both of the reset limits in Table 4-5 are exceeded. When powering down, the drive becomes reset when either supply voltage drops below the lower threshold.

Table 4-5 *Power Reset Limits*

DC VOLTAGE	THRESHOLD
+5 V	4.25 V to 4.70 V
+12 V	8.90 V to 10.7 V

4.5.3 Power Requirements

Table 4-6 lists the voltages and corresponding current for the various modes of operation of the Maverick 270/540S hard disk drive.

Table 4-6 *Typical Power and Current Consumption*

MODE OF OPERATION	TYPICAL AVERAGE CURRENT[1](mA rms)				TYPICAL AVERAGE POWER (WATTS)	
	+12 V		+5 V			
	270S	540S	270S	540S	270S	540S
Startup[2]	1000 peak	1000 peak	400 peak	400 peak	-	–
Idle[3]	170	170	330	330	3.7 W	3.7 W
Random Read/Write[4]	290	300	300	310	5.0 W	5.2 W
Random Seek[5]	475	490	275	275	7.1 W	7.3 W
Standby[6]/Sleep[7]	8	8	190	190	1.0 W	1.0 W
Noise and Ripple	100 mv	100 mv	50 mv	50 mv	-	–

1. Reflects nominal values for +12 V and +5 V power supplies
2. Current is rms except for startup. Startup current is peak current of peaks greater than 10 ms in duration.
3. Idle mode is in effect when the drive is not reading, writing, seeking, or executing any commands. A portion of the read/write circuitry is powered down, the motor is up to speed, and the Drive Ready condition exists. The actuator resides on the last track accessed.
4. Random Read/Write mode is defined as when data is being read from or written to the disk. It is computed based on 40% seeking, 30% on-track read, and 30% on-track write.
5. Continuous random seek operations with no controller delay.
6. Standby is defined to be when the motor is stopped, the actuator is parked, and all electronics are except the interface control are in the sleep state. Standby occurs after a programmable timeout after the last host access. Drive ready and seek complete status exist. The drive leaves standby upon receipt of a command that requires disk access or upon receiving a spinup command.
7. Sleep is when the spindle and actuator motors are off with the heads latched in the landing zone. Receipt of a reset causes the drive to transition from the sleep to the standby mode.

Note: Measured power does not include power consumed by the SCSI termination resistors.

4.6 ACOUSTICS

Table 4-7 and Table 4-8 specifies the acoustical characteristics of the Maverick 270/540S hard disk drive.

Table 4-7 *Acoustical Characteristics – Sound Pressure*

OPERATING MODE	MEASURED NOISE	DISTANCE
Idle On Track	34 dbA (mean) 38 dbA (maximum)	39.3 in (1 m)
Random Seek[1]	39 dbA (mean) 42 dbA (maximum)	39.3 in (1 m)

1. Random seek is defined as 40% read, 40% write, and 20% idle.

Table 4-8 *Acoustical Characteristics—Sound Power*

Operating Mode	Measured Noise (Sound Power per ISO 7779)
Idle On Track	4.0 bels (mean) 4.5 bels (max)
Random Seek[1]	4.8 bels (mean) 5.0 bels (max)

1. Random seek is defined as 40% read, 40% write, and 20% idle.

4.7 MECHANICAL DIMENSIONS

Table 4-8 specifies the mechanical dimensions of the Maverick 270/540S hard disk drive. Dimensions measurements do not include the faceplate.

Table 4-9 *Mechanical Dimensions and Weight*

Dimension	Inches	Millimeters
Height	1.0 in	25.4 mm
Width	4.0 in	101.6 mm
Depth	5.75 in	146.1 mm

Weight

Pounds	Grams	Ounces
0.86 (1 disk)	390 (1 disk)	13.7 (1disk)
0.91 (2 disk)	413 (2 disk)	14.6 (2disk)

4.8 ENVIRONMENTAL CONDITIONS

Table 4-9 summarizes the environmental specifications for the Maverick 270/540S hard disk drive.

Table 4-10 *Environmental Specifications*

PARAMETER	OPERATING	NONOPERATING
Temperature	4° to 50°C 39° to 122°F	–40° to 65°C –40° to 149°F
Temperature Gradient	20°C/hr	40°C/hr
Humidity[1] Maximum Wet Bulb	8% to 85% rh 26°C (79°F)	5% to 95% rh 46°C (115°F)
Humidity Gradient	30% hr	30% hr
Altitude[2]	–60 m to 3 km (–200 to 10,000 ft.)	–60 m to 12 km (–200 to 40,000 ft.)
Altitude Gradient	1.5 kPa/min.	8 kPa/min.

Notes: [1] No condensation.
[2] Altitude is relative to sea level.

4.9 SHOCK AND VIBRATION

The Maverick 270/540S hard disk drive can withstand levels of shock and vibration applied to any of its three mutually perpendicular axes, or principal base axes, as specified in Table 4-10. A functioning drive can be subjected to specified *operating* levels of shock and vibration.

When a drive has been subjected to specified *nonoperating* levels of shock and vibration, with power to the drive off, there will be no change in performance at power on.

When packed in either the 1-pack or 10-pack shipping container, Maverick 270/540S drives can withstand a drop from 30 inches onto a concrete surface on any of its surfaces, six edges, or three corners.

Table 4-11 *Shock and Vibration Specifications*

	OPERATING	NONOPERATING
Shock 1/2 sine wave, 11 ms duration (10 hits maximum)	6 G (no soft errors) 10 G (1 soft error/shock)	70 G (No unrecovered errors)
Vibration 5-300 Hz sine wave (peak to peak) 1.0 octave/minute sweep	1.0 G	2.0 G

4.10 RELIABILITY

Mean Time Between Failures (MTBF): 150,000 Power On Hours (POH), typical usage

The Quantum MTBF numbers represent Bell-Core MTBF predictions at ambient operating conditions and represent the minimum MTBF that Quantum or a customer would expect from the drive. Quantum's ongoing reliability testing and field return data has historically confirmed an actual MTBF of twice the Bell-Core prediction.

Preventive Maintenance (PM): Not required

Mean Time To Repair (MTTR): 30 minutes

Start/Stop: 20,000 cycles (minimum)

Component Life 5 Years

4.11 DISK ERRORS

Table 4-11 provides the error rates for the Maverick 270/540S hard disk drive.

Table 4-12 *Error Rates*

ERROR TYPE	MAXIMUM NUMBER OF ERRORS
Reallocated data errors:	
Correctable read errors[1]	1 error per 10^{12} bits read
Uncorrectable read errors[2]	1 error per 10^{14} bits read
Transferred errors[3]	1 error per 10^{23} bits read
Seek Errors[4]	1 error per 10^{6} seeks

1. Correctable read errors are read errors that are recovered by retries or by application of the double-burst error correction algorithm.
2. Uncorrectable read errors are errors that are not correctable using ECC or retries. The drive will terminate retry reads either when a repeating error pattern occurs, or after eight unsuccessful retries and the application of double-burst error correction.
3. Transferred errors are errors that are not detected and subsequently not corrected by the drive.
4. Seek errors occur when the actuator fails to reach (or remain) over the requested cylinder or if the drive executes a recalibration routine to find the requested cylinder (a full recalibration takes about seven seconds).

Chapter 5
BASIC PRINCIPLES OF OPERATION

This chapter describes the operation of Maverick 270/540S functional subsystems. It is intended as a guide to the operation of the drive, rather than a detailed theory of operation.

5.1 MAVERICK DRIVE MECHANISM

This section describes the drive mechanism. The Maverick 270/540S hard disk drive consists of a mechanical assembly and a printed circuit board assembly (PCBA) as shown in Figure 5-1.

5.1.1 Head/Disk Assembly

The head/disk assembly (HDA) contains the mechanical subassemblies of the drive, which are sealed under a metal cover. The HDA consists of the disk assembly, rotary positioner assembly, headstack assembly, actuator latch assembly, and base casting which includes the DC motor assembly. Assembly of the HDA takes place in a Class-100 clean room. These subassemblies are not adjustable or field-repairable.

> **CAUTION: To ensure that the air in the HDA remains free of contamination, never remove or adjust its cover and seals. Tampering with the HDA voids your warranty.**

The Maverick 270S has one magnetic disk and two read/write heads, while the Maverick 540S has two disks and four read/write heads.

5.1.2 Base Casting Assembly

A single-piece, aluminum alloy base casting provides a mounting surface for the drive mechanism and PCBA. The base casting also acts as the flange for the DC motor assembly. To provide a contamination-free environment for the HDA, a gasket provides a seal between the base casting and the metal cover that enclose the drive mechanism.

5.1.3 DC Motor Assembly

Integral with the base casting, the DC motor assembly is a rotating shaft, brushless DC spindle motor that drives the counter-clockwise rotation of the disks.

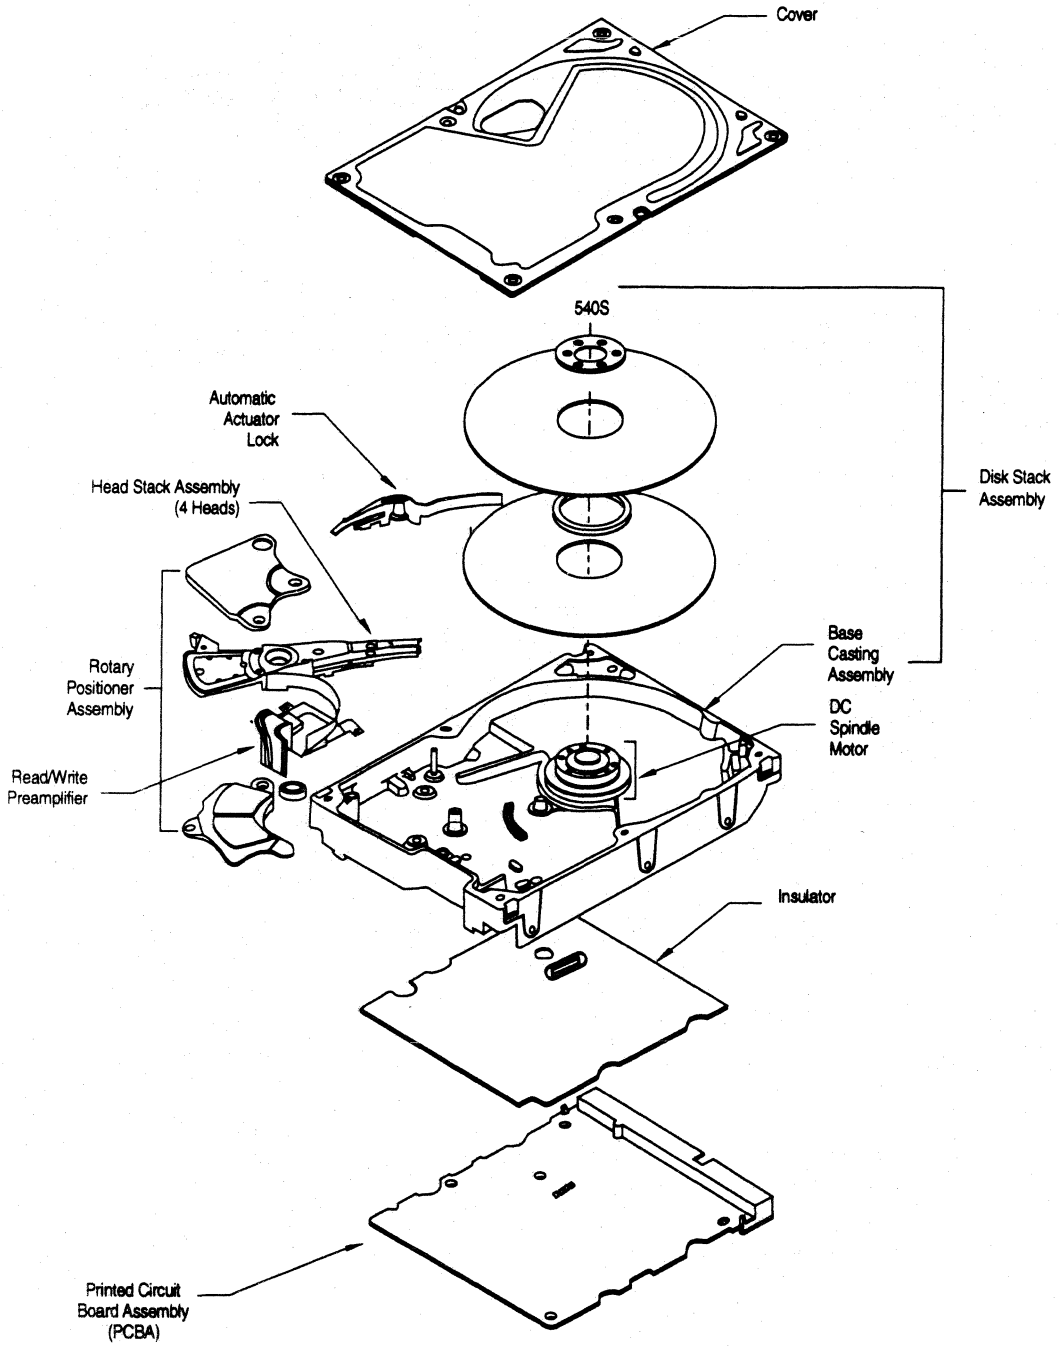

Figure 5-1 *Maverick 270/540S Exploded View*

5.1.4 Disk Stack Assemblies

The disk stack assembly in the Maverick 270/540S hard disk drive consists of one or two disks secured on the hub of the DC motor assembly by a disk clamp. The aluminum-alloy disks have a sputtered thin-film magnetic coating.

A carbon overcoat lubricates the disk surfaces, to prevent head and media wear due to head contact with the disk surfaces during head takeoff and landing. Head contact with the disk surfaces occurs only in the landing zone outside of the data area, when the disks are not rotating at full speed. The landing zone is located at the inner diameter of the disk, well beyond the last cylinder of the data area.

5.1.4.1 Surface Layout

The Maverick 270/540S hard disk drive has 2,853 tracks per recording surface. Table 5-1 details the surface layout for each platter (recording surface). The data tracks are divided into 16 recording zones. The drive uses multiple zone recording, where each data track contains between 58 and 118sectors (depending on the recording zone)

5.1.5 Headstack Assembly

The headstack assembly consists of read/write heads, an E-block/coil subassembly, bearings, and a flex circuit. (The E-block/coil subassembly consists of an E-block and coil joined together by insertion molding.) Read/write heads mounted to spring-steel flexures are swage mounted onto the rotary positioner assembly arms. The E-block is a single piece, die-cast design.

The flex circuit exits the HDA between the base casting and the cover. A cover gasket seals the gap. The flex circuit connects the headstack assembly to the PCBA. The flex circuit contains a read PreAmplifier and Write Driver IC.

5.1.6 Rotary Positioner Assembly

The rotary positioner, or rotary voice-coil actuator, is a Quantum-proprietary design that consists of upper and lower permanent magnet plates bolted to the base casting, and a rotary single-phase coil molded to the headstack E-block. The magnets consist of two alternating poles bonded to the magnet plates. Resilient crash stops mounted on the magnet plates and base casting prevent the heads from being driven into the spindle or off of the disk surface.

Current from the power amplifier induces a magnetic field in the voice coil. Fluctuations in the field around the permanent magnets cause the voice coil to move. The movement of the voice coil positions the heads over the requested cylinder.

5.1.7 Automatic Actuator Lock

To ensure data integrity and prevent damage during shipment, the drive uses a dedicated landing zone and Quantum's patented AIRLOCK®. The AIRLOCK holds the headstack in the landing zone whenever the disks are not rotating. It consists of an air vane mounted near the perimeter of the disk stack and a locking arm that restrains the actuator arm assembly.

When DC power is applied to the motor and the disk stack rotates, the rotation generates an airflow on the surface of the disk. As the flow of air across the air vane increases with disk rotation, the locking arm pivots away from the actuator arm, enabling the headstack to move out of the landing zone. When DC power is removed from the motor, an electronic return mechanism automatically pulls the actuator into the landing zone, where the AIRLOCK holds it in place.

Table 5-1 *Surface Layout*

CYLINDER CONTENTS	ZONE	STARTING CYLINDER	ENDING CYLINDER	NUMBER OF CYLINDERS	SECTORS PER TRACK
Guard band					
System Test Data Diskware Firmware Diagnostics	—	–1	–7	7	93
Data	0	0	199	200	118
	1	200	358	159	118
	2	359	596	238	118
	3	597	744	148	114
	4	745	872	128	112
	5	873	1030	158	108
	6	1031	1218	188	104
	7	1219	1396	178	97
	8	1397	1584	188	93
	9	1585	1782	198	88
	10	1783	1940	158	83
	11	1941	2178	238	78
	12	2179	2296	118	74
	13	2297	2434	138	69
	14	2435	2612	178	65
	15	2613	2852	240	58
Guard band					
Total Data Tracks = 5,706 (270S); 11,412 (540S),					

5.1.8 Air Filtration

The Maverick 270/540S hard disk drives are Winchester-type drives. The heads fly very close to the media surface. Therefore, it is essential that the air circulating within the drive be kept free of particles. Quantum assembles the drive in a Class-100, purified air environment, then seals the drive with a metal cover. When the drive is in use, the rotation of the disks forces the air inside of the drive through an internal filter.

The highest air pressure within the HDA is at the outer perimeter of the disks. A constant stream of air flows through a 0.3-micron circulation filter positioned in the base casting. As illustrated in Figure 5-2, air flows through the circulation filter in the direction of the disk rotation. This design provides a continuous flow of filtered air when the disks rotate.

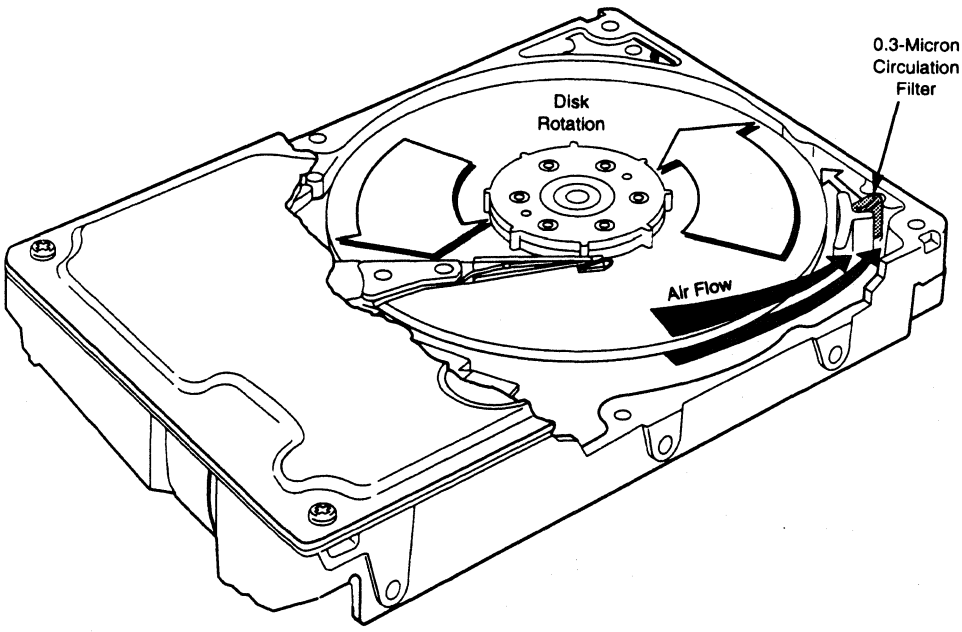

Figure 5-2 *Maverick 270/540S HDA Air Filtration*

5.2 MAVERICK 270/540S DRIVE ELECTRONICS

Advanced circuit design, and the use of miniature surface-mounted devices and proprietary VLSI components, enable the drive electronics, including the SCSI bus interface, to reside on a single printed circuit board assembly (PCBA). The components are mounted only on one side of the PCBA. Figure 5-3 contains a simplified block diagram of the Maverick 270/540S electronics.

The only electrical component not on the PCBA is the PreAmplifier and Write Driver IC. It is on the flex circuit (inside of the sealed HDA). Mounting the PreAmplifier as close as possible to the read/write heads improves the signal-to-noise ratio. The flex circuit (including the PreAmplifier and Write Driver IC) provides the electrical connection between the PCBA, the rotary positioner assembly, and read/write heads.

5.2.1 µController

The µController (microcontroller) provides local processor services to the drive electronics, under program control. The µController provides direction for the Disk Controller and SCSI Interface ASIC (DCSIA), the Read/Write ASIC, and the Actuator Driver. In addition, it controls the head selection process.

The MAD (multiplexed address and data) bus consists of 8 data bits and 16 address bits, plus control signals.

An internal 32 Kbyte ROM (within the µController) provides program code that is used to complete the drive spinup and calibration procedure. Once this is complete, the µController reads additional control code from the disk (Diskware) and stores it in the buffer DRAM.

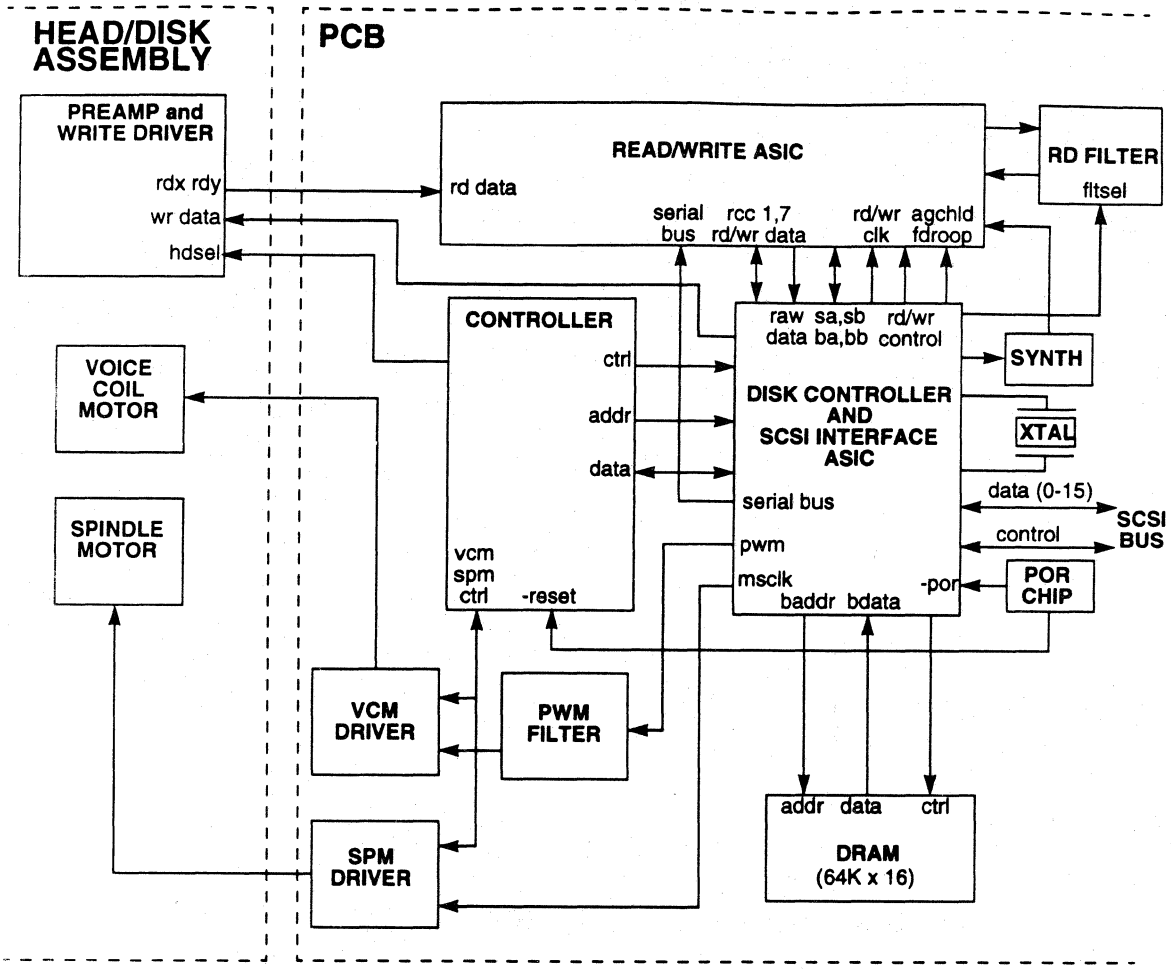

Figure 5-3 *Maverick 270/540S Block Diagram*

5.2.2 DCSIA

The DCSIA (Disk Controller And SCSI Interface) shown in Figure 5-4 provides control functions to the drive under the direction of the µController.

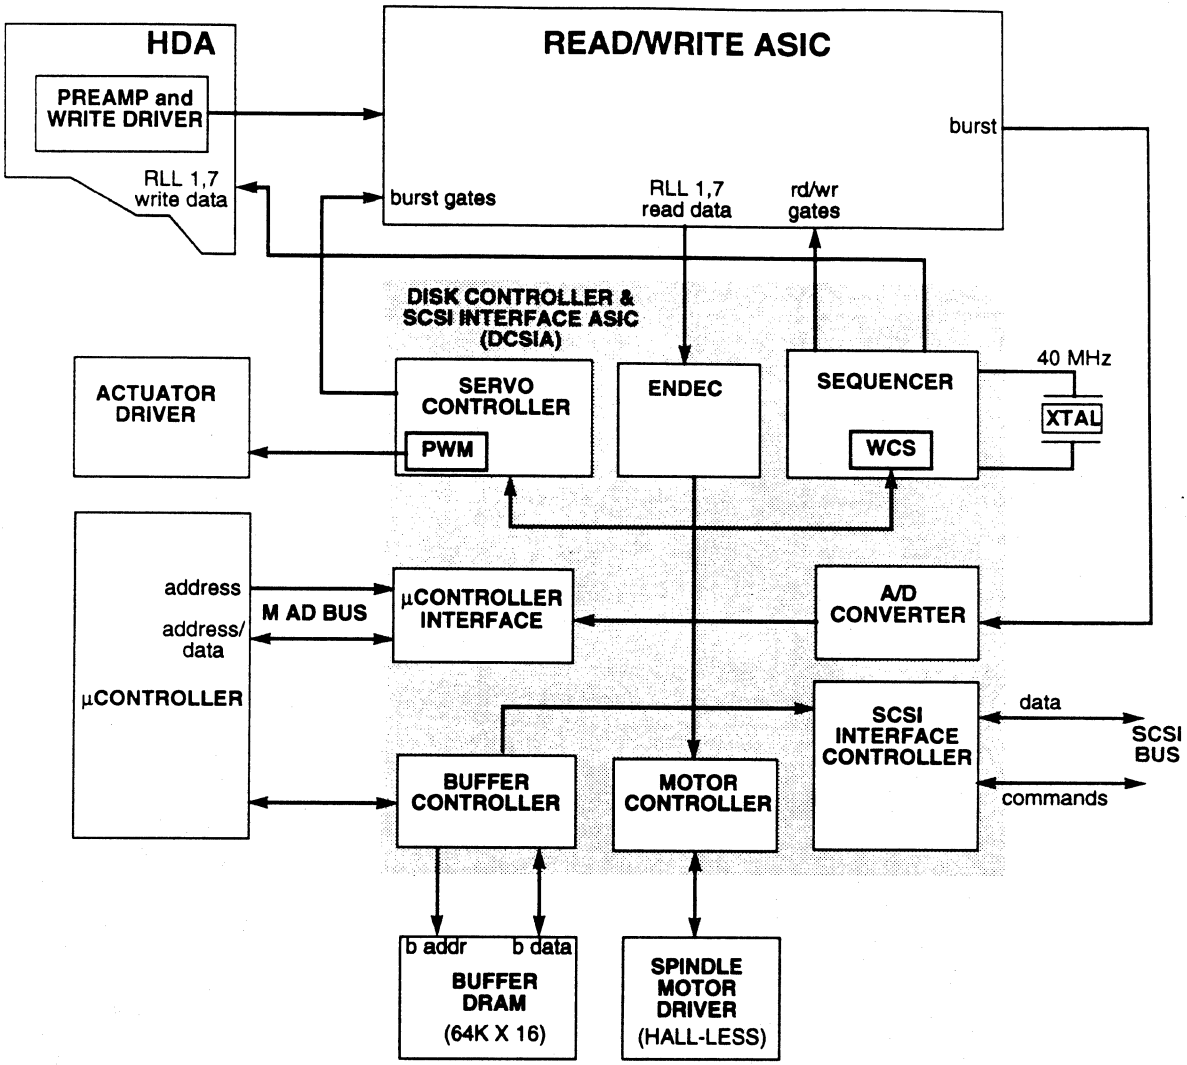

Figure 5-4 *DCSIA Block Diagram*

The DCSIA is a proprietary ASIC developed by Quantum. The DCSIA comprises eight functional modules (described below):

- Encoder-Decoder (ENDEC)
- 5-Channel 8 -Bit ADC
- Sequencer
- Buffer Controller
- µController Interface
- Motor Controller
- Servo Controller, including PWM
- SCSI Interface Controller

5.2.2.1 A/D Converter

The Analog to Digital converter (A/D) receives multiple burst analog inputs from the Read/Write ASIC. The A/D is used to sample the demodulated position information (burst inputs) and convert it to a digital position signal used by the Servo Controller to position the HDA actuator.

5.2.2.2 Encoder/Decoder (ENDEC)

The ENDEC is a 1,7 RLL encoder/decoder. The ENDEC codes and decodes data under the control of the sequencer portion of the DCSIA. The ENDEC encodes write data and decodes read data. Write data is sent to the Sequencer for transfer to the Read/Write ASIC. Read data is input directly to the ENDEC from the Read/Write ASIC.

5.2.2.3 Sequencer

The sequencer controls the operation of the read and write channel portions of the DCSIA, including the ENDEC, through commands loaded into a 28 x 24-bit writable control store (WCS) by the µController. Each command word is 28 bits in length. To initiate a disk operation, the µController loads a set of these 28-bit commands into the WCS. Up to 24 commands may be loaded. Loading and manipulating the WCS is done through the µController Interface registers.

The sequencer also directly drives the read and write gates (RG, WG) to the Read/Write ASIC and the R/W Preamplifier, as well as passing write data to the Precompensator circuit in the Read/ Write ASIC. The sequencer handles the appending and processing of the ECC bytes for the read/ write data, using a 96-bit, interleaved Reed-Solomon scheme with two cross-check bytes.

5.2.2.4 Buffer Controller

The buffer controller supports a 128Kbyte buffer, which is organized as 64K x 16 bits. The 16-bit width implementation provides a 20 MB/s maximum buffer bandwidth, which allows 6MB/s disk channel operation with a minimum 6MB/s host channel bandwidth. This increased bandwidth allows the µController to have direct access to the buffer, eliminating the need for a separate µController IC.

The buffer controller supports both drive and host address rollover and reloading, to allow for buffer segmentation. A separate access channel is provided for the µController, which allows it to directly access the buffer.

Thé Buffer Controller operates under the direction of the µController.

5.2.2.5 µController Interface

The µController Interface provides the means for the µController to read and write data to the DCSIA modules to control their operation or supply them with needed information. It consists of both physical and logical components.

The physical component of the interface comprises the eight-bit MAD bus, eight address lines, a read strobe, a write strobe, an address latch enable (ALE) signal, and a wait control line.

The logical component of the interface comprises internal control and data registers accessible to the µController. By writing and reading these registers, the µController loads the Sequencer, controls and configures the Buffer and Motor controllers, and passes coded servo information to the Servo Controller.

5.2.2.6 Servo Controller

The Servo Controller contains a 10-bit Digital to Analog converter (D/A), in the form of a Pulse Width Modulator (PWM). The PWM signal is output to the Actuator Driver to control the motion of the actuator. The Servo Controller also decodes raw data from the disk to extract the current position information. That information is read by the µController and is used to generate the actuator control signal that is sent to the PWM. The actuator driver is an analog power amplifier circuit external to the DCSIA. The Servo Controller operates under the direction of the µController.

5.2.2.7 Motor Controller

The Motor Controller provides a reference clock for speed control for the hall-less Spindle Motor Driver. The Spindle Motor Driver is an analog power amplifier circuit external to the DCSIA.

5.2.2.8 SCSI Interface

The SCSI Interface Controller portion of the DCSIA provides data handling, bus control, and transfer management services for the SCSI interface. Configuration and control of the interface is accomplished by the microcontroller across the MAD bus. Data transfer operations are controlled by the DCSIA Buffer Controller module

5.2.3 Read/Write ASIC

The Read/Write ASIC shown in Figure 5-5 provides write data precompensation and read channel processing functions for the drive. The Read/Write ASIC receives the RD GATE signal, write control, reference clock from the frequency synthesizer, serial programming, and servo burst and sample gates from the DCSIA. The Read/Write ASIC sends encoded read data and the read reference clock. Write data from the DCSIA is sent directly to the preamplifier.

The Read/Write ASIC comprises three functional modules (described below):

- Pulse Detector
- Peak Detector and Servo Demodulator
- Data Synchronizer

Frequency synthesis is done using a standalone ASIC.

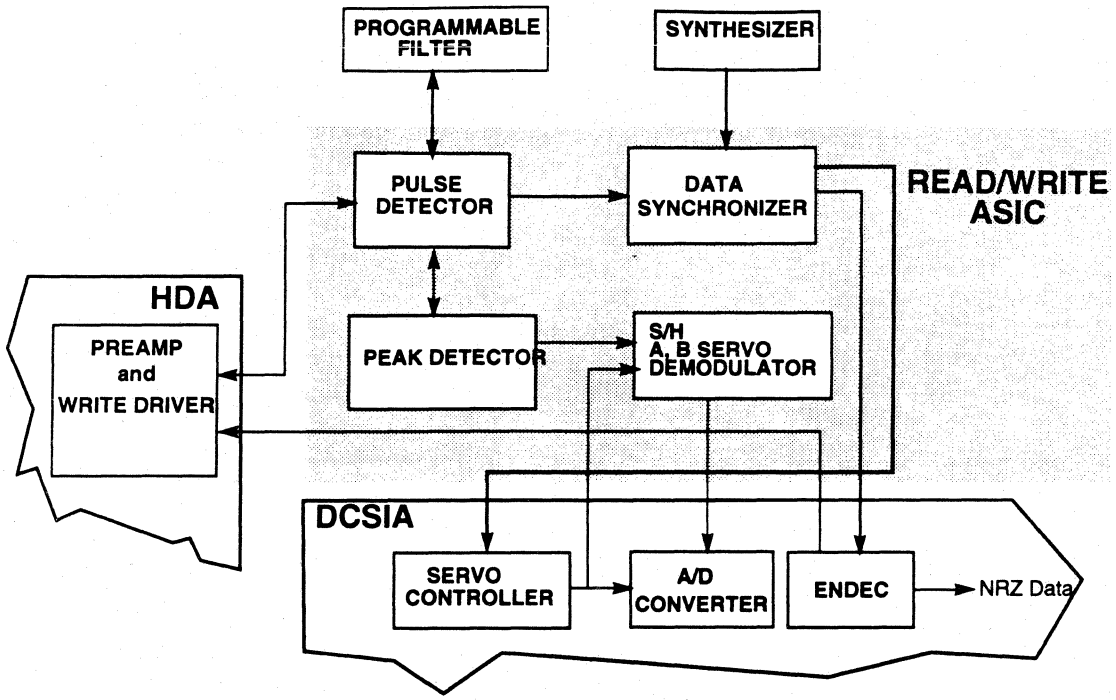

Figure 5-5 *Read/Write ASIC Block Diagram*

5.2.3.1 Synthesizer

The frequency synthesizer in the timebase generator provides the reference frequency for the Data Synchronizer and the DCSIA. The µController programs the Timebase Generator with the appropriate frequency for each zone on the disk by writing data over the serial interface.

5.2.3.2 Pulse Detector

The pulse detector filters and amplifies the read data received from the R/W Pre-Amp. The read data at this point is 1,7 RLL encoded analog signals. Once converted to an asynchronous digital data stream, the read data is passed to the Data Synchronizer module for synchronization.

5.2.3.3 Peak Detector

The peak detector extracts the servo wedge Burst A, Burst B, and Burst C information from the read data signal. The gate inputs to the peak detector are sent from the Servo Controller module in the DCSIA. The gate inputs are used to extract the burst fields from the servo wedge.

5.2.3.4 Data Synchronizer

Using a phase-lock-loop (PLL), the Data Synchronizer synchronizes the read data to the read clock (after the correct address mark is recovered) and passes both to the ENDEC module in the DCSIA. The synthesizer generates the reference frequency to the PLL for data clock recovery and resynchronization.

5.2.4 PreAmplifier and Write Driver IC

The PreAmplifier and Write Driver IC provides read preamplifier and write driver functions and the R/W head selection. The read preamplifier amplifies the low amplitude differential voltages generated by the R/W heads and transmits them to the Pulse Detector module in the Read/Write ASIC. The write driver outputs current to the read/write heads in alternating directions according to the write data received from the Read/Write ASIC. A head matrix, under control of the μController, provides a switch point for head selection.

5.3 SERVO SYSTEM

5.3.1 General Description

The servo system controls the positioning of the read/write heads and holds the read/write heads on track during read/write operations. The servo system also compensates for thermal offsets between heads on different surfaces, and shock and vibration subjected to the drive.

The Maverick 270/540S uses a sectored servo system. Positioning information is radially encoded in 78 evenly-spaced servo bursts on each track. These servo burst wedges provide radial position information for each data head, 78 times per revolution. Because the drive uses multiple zone recording, where each zone has a different bit density, split data fields are necessary to optimally utilize the non-servo area of the disk. The split data fields are achieved by special processing through the DCSIA, and their presence is transparent to the host system. The servo area remains phase coherent across the surface of the disk, even though the disk is divided into various data zones. The main advantage of the sectored servo systems is that the data heads are also servo heads, which means that sectored servo systems eliminate the problems of static and dynamic offsets between heads on different surfaces.

The Maverick 270/540S servo system is also classified as a digital servo because track following compensation is done in firmware. The bump detect, on-track, velocity profiles, and other "housekeeping" tasks are also done in firmware.

The servo system has three basic modes of operation (1) velocity mode, (2) settle mode, and (3) track following mode.

1. Velocity Mode. The acceleration and deceleration of the actuator for seek lengths of 20 or more tracks are accomplished with a velocity loop. The velocity profile is calculated as the seek is in progress, using programmable coefficients. The bandwidth of the velocity loop is kept low, and feed forward is heavily used.

2. Settle Mode.The settle mode is used for all accesses—head switches, 1- to 20-track seeks, and the end of all velocity seeks. The settle mode is a position loop with velocity damping. No feed forward is used for the settle mode.

3. The track following mode is used when "on-track". This is also a position loop, with an integrator in the compensation.

5.3.2 Servo Burst and Track Information

Positional information is encoded on all tracks on all data surfaces. All data heads are also servo heads. The areas with servo/position information are called wedge areas. These wedge areas are evenly spaced radially around the disk, like spokes on a wheel. There are 78 wedge areas per track. Since the disk rotation is 60 Hz, the position information is updated at 4680 Hz (78 x 60). This is also known as the sampling frequency *fs*. The sample period, Ts, is $1/fs = 213.6\ \mu s$. Every wedge area consists of four separate fields: (1) Automatic Gain Control (AGC)/Sync field, (2) Sector Address Mark (SAM) field, (3) Track number and (4) Burst area. Since a Phase Lock Loop (PLL) is not used in the servo wedge area, time discrimination is used. Timing for all four fields is generated from the same crystal reference. This clock has a period of 37.5 nanoseconds, and will be referred to as T_{clk}. Data is encoded in the following manner: a servo data bit 1 = 1 0 0 1 0 0 0 0 0, and a servo data bit 0 = 1 0 0 0 0 0 1 0 0, with each digit having a T_{clk} width.

1. The AGC/sync field consists of a $3T_{clk}$ pattern, and is used by the AGC to acquire the proper amplitude for the encoded track number and position bursts. It is also used by the DCSIA Servo Controller for synchronizing to the raw data pulses. The total length of the AGC/sync field is: AGC + Sync = 88 + 40 = $128T_{clk}$.

2. The SAM follows the sync field. It consists of a $14T_{clk}$ pattern repeated twice, followed by a servo data bit 0. Following the servo data bit 0 is either a servo data bit 0 or 1, which is also known as the index bit. If a one is decoded, then an index pulse is generated. The total length of the SAM field is $46T_{clk}$.

3. Following the index bit is the track number. The track number is a 12-bit Gray coded number. The Gray code to binary conversion is done in the DCSIA Servo Controller. Each Gray code bit is encoded as servo data bits. The total length of the track number is $108T_{clk}$.

4. Following the track number is the burst area. There are three bursts per servo wedge time, an A burst followed by the B burst and finally the C burst. When "on-track", the A and C burst will be at half amplitude. For even tracks the B burst will be full amplitude, and for odd tracks the B burst will be zero.

   The A, B, and C bursts are $42T_{clk}$ long. There are dc erase areas of $9T_{clk}$ before and after the burst fields, giving a total time of $144T_{clk}$ for the burst area.

5.3.3 Position and Velocity

The state of the servo system determines how the position information is derived. During the velocity and coarse settle modes of operation, the position signal is the convolution of the track number and A or B burst values and has a 1/256 of a track pitch resolution—about 0.4%. During track following and fine settle, the A and C bursts are used for the position information, and the resolution is at least 1/512th of a track pitch—about 0.2%. The velocity is the combination of two components: the difference between the present and last track, and a value that is proportional to the current.

5.4 READ AND WRITE OPERATIONS

The following sections provide descriptions of the read channel, write channel, and SCSI interface control operations.

5.4.1 The Read Channel

The drive has one read/write head for each data surface. There are two read/write heads for the Maverick 270S and four for the Maverick 540S. The signal path for the read channel begins at the read/write heads. As the magnetic flux transitions recorded on a disk pass under the head, they generate low amplitude, differential output voltages. These read signals pass from the read/write head to the flex circuit's read preamplifier, which amplifies the signal. Preamplification occurs on the flex circuit because of its close proximity to the heads (preserving the high signal-to-noise ratio).

The flex circuit transmits the preamplified signal from the HDA to the drive PCBA. On the PCBA, the Read/Write ASIC's Pulse Detector and Data Synchronizer modules further amplify, filter, and process the read signal to reduce ambiguities (e.g., drop-ins and drop-outs). In addition the Data Synchronizer converts the signal from the serial encoded head data to a synchronized data stream, with its accompanying clock. The Read/Write ASIC then sends the resynchronized data output to Quantum's proprietary Disk Controller and SCSI Interface ASIC (DCSIA).

The DCSIA manages the flow of data between the Data Synchronizer on the Read/Write ASIC and its SCSI Protocol Interface Controller. It also controls data access for the external RAM buffer. The DCSIA ENDEC decodes the RLL 1,7 (Run Length Limited) format to provide a serial bit stream. This serial data is converted to 8-bit bytes. The Sequencer module identifies the data as belonging to the target sector by interpreting the ID field.

After a full sector is read, the DCSIA will check to see if the firmware needs to apply ECC single or double-burst correction to the data. The Buffer Controller section of the DCSIA stores the data in the Cache and transmits the data to the SCSI Protocol Interface Controller module, which transmits the data to the SCSI bus.

5.4.2 The Write Channel

For the write channel, the signal path follows the reverse order of that for the read channel. The host presents data via the SCSI bus to the DCSIA SCSI Protocol Interface Controller. The Buffer Controller section of the DCSIA stores the data in the buffer. Because data is presented to the drive than the drive can write data to a disk, data is stored temporarily in the buffer. Thus, the host can present data to the drive at a rate that is independent of the rate at which the drive can write data to the disk.

Upon correct identification of the target sector (by interpreting the ID field), the data is shifted to the DCSIA Sequencer where crosscheck and ECC bytes are generated and appended. The Sequencer then converts the bytes of data to a serial bit stream. The DCSIA ENDEC encodes the data in the RLL 1,7 format, generates a preamble field, inserts address marks, and transmits the data to the PreAmplifier and Write Driver IC via the write data line.

The DCSIA Sequencer switches the PreAmplifier and Write Driver IC to write mode and the µController selects a head. Once the PreAmplifier and Write Driver IC receives a write gate signal, it transmits current reversals to the head, thereby writing magnetic transitions on the disk.

5.4.3 Interface Control

The interface with the host system is through a 50-pin SCSI interface connector. The DCSIA SCSI Protocol Interface Controller implements the SCSI bus protocol (interface logic). Operating under the control of the drive's microprocessor, the SCSI Protocol Interface Controller receives and transmits bytes of data over the bus.

The DCSIA Buffer Controller writes data to or reads data from the Cache over 16 data lines. Under µController direction, the DCSIA controls the transfer of data and handles the addressing of the Cache. The internal data transfer rate to and from the Cache is 36.08 Mb/s.

This high transfer rate allows the DCSIA to communicate over the SCSI interface at a asynchronous data transfer rate of 6.0 MB/s, or a synchronous rate of up to 10.0 MB/s while it simultaneously controls disk-to-RAM transfers and microcontroller access to control code stored in the buffer RAM.

5.5 FIRMWARE FEATURES

This section describes the following firmware features:

- Disk caching
- Track and cylinder skewing
- Error detection and correction
- Defect management

5.5.1 Disk Caching

The Maverick 270/540S hard disk drive incorporates DisCache, a 96K disk cache, to enhance drive performance. This integrated feature can significantly improve system throughput. Read and write caching can be enabled or disabled using the MODE SELECT command.

5.5.1.1 Adaptive Caching

The cache buffer for the Maverick 270/540S features adaptive segmentation for more efficient use of the buffer's RAM. With this feature, the buffer space used for read and write operations is dynamically allocated. The cache can be flexibly divided into several segments, under program control. Each segment contains one cache entry.

A cache entry consists of the requested read data plus its corresponding prefetch data. Adaptive segmentation allows the drive to make optimum use of the buffer, the amount of stored data can be increased.

5.5.1.2 Read Cache

DisCache anticipates host-system requests for data and stores that data for faster access. When the host requests a particular segment of data, the caching feature uses a prefetch strategy to "look ahead" and automatically read and store the subsequent data from the disk into high-speed RAM. If the host requests this subsequent data, the RAM is accessed rather than the disk.

Since typically 50 percent or more of all disk requests are sequential, there is a high probability that subsequent data requested will be in the cache. This cached data can be retrieved in microseconds rather than milliseconds. As a result, DisCache can provide substantial time savings during at least half of all disk requests. In these instances, DisCache could save most of the disk transaction time by eliminating the seek and rotational latency delays that dominate the typical disk transaction. For example, in a 1K data transfer, these delays take up to 90 percent of the elapsed time.

DisCache works by continuing to fill its cache memory with adjacent data after transferring data requested by the host. Unlike a noncaching controller, Quantum's disk controller continues a read operation after the requested data has been transferred to the host system. This read operation terminates after a programmed amount of subsequent data has been read into the cache segment.

The cache memory consists of a 128 KB DRAM buffer. The upper 32K contains the drive firmware, the balance (96K) is allocated to hold the data, which can be directly accessed by the host by means of read and write commands. The memory functions as a group of segments (ring buffers) with rollover points at the end of each segment (buffer). The unit of data stored is the logical block (that is, a multiple of the 512-byte sector). Therefore, all accesses to the cache memory must be in multiples of the sector size. The following commands force emptying of the cache:

- FORMAT UNIT
- INQUIRY
- MODE SELECT
- MODE SENSE
- READ CAPACITY
- READ DEFECT DATA
- READ LONG
- REASSIGN BLOCKS
- VERIFY
- WRITE LONG
- WRITE VERIFY
- WRITE BUFFER

5.5.1.3 Write Cache

When a write command is executed with write caching enabled, the drive stores the data to be written in a DRAM cache buffer and immediately sends a COMMAND COMPLETE message to the host before the data actually is written to the disk. The host is then free to move on to other tasks, such as preparing data for the next data transfer, without having to wait for the drive to seek to the appropriate track or rotate to the specified sector.

While the host is preparing data for the next transfer, the drive immediately writes the cached data to the disk, usually completing the operation in less than 23 ms. With WriteCache, a single-block, random write, for example, requires only about 3 ms of host time. Without WriteCache, the same operation would occupy the host for about 27 ms.

WriteCache allows data to be transferred in a continuous flow to the drive rather than as individual blocks of data separated by disk access delays. This functionality is achieved by taking advantage of the ability to write blocks of data sequentially on a disk that is formatted with a 1:1 interleave. This means that as the last byte of data is transferred out of the write cache and the head passes over the next sector of the disk, the first byte of the of the next block of data is ready to be transferred. Thus, there is no interruption or delay in the data transfer process. The WriteCache algorithm fills the cache buffer with new data from the host while simultaneously transferring the data to the disk that the host previously stored in the cache.

5.5.1.4 Performance Benefits of Read Cache

In a drive drive without DisCache, there is a delay during sequential reads because of the rotational latency—even if the disk actuator already is positioned at the desired cylinder. DisCache eliminates this rotational latency time (8.33 ms on average) when requested data resides in the cache.

Moreover, the disk often must service requests from multiple processes in a multitasking or multiuser environment. In these instances, while each process might request data sequentially, the disk drive must share time among all these processes. In most disk drives, the heads must move from one location to another. With DisCache, even if another process interrupts, the drive continues to access the data sequentially from its high-speed memory.

In handling multiple processes, DisCache achieves its most impressive performance gains, saving both seek and latency time (23 ms on average) when desired data resides in the cache.

For read operations, the requested read data takes up a certain amount of space in the cache segment so the corresponding prefetch data can essentially occupy the rest of the space within the segment. The drive's prefetch algorithm dynamically controls the actual prefetch value based on the current demand, with the consideration of overhead to subsequent commands.

5.5.2 Track and Cylinder Skewing

Track and cylinder skewing in the Maverick 270/540S minimizes latency time, and thus, increases data throughput. See Table 5-2 for the track and cylinder data.

5.5.2.1 Track Skewing

Track skewing reduces the latency time that results when the drive must switch read/write heads to access sequential data. A track skew is employed such that the next logical sector of data to be accessed will be under the read/write head once the head switch is made and the data is ready to be accessed. Thus, when sequential data is on the same cylinder but on a different disk surface, a head switch is needed but not a seek. Since the sequential head-switch time is well defined on the Maverick 270/540S (6 ms worst case), the sector addresses are optimally positioned across track boundaries to minimize the latency time during a head switch.

5.5.2.2 Cylinder Skewing

Cylinder skewing also is used to minimize the latency time associated with a single-cylinder seek. The next logical sector of data that crosses a cylinder boundary is positioned on the drive such that after a single-cylinder seek is performed, and when the drive is ready to continue accessing data, the sector to be accessed is positioned directly under the read/write head. Therefore, the cylinder skew takes place between the last sector of data on the last head of a cylinder and the first sector of data on the first head of the next cylinder. Since single-cylinder seeks are well defined on the Maverick 270/540S, the sector addresses can be optimally positioned across cylinder boundaries to minimize the latency time associate with a single-cylinder seek.

Table 5-2 *Track and Cylinder Skewing*

ZONE NUMBER	SECTORS PER TRACK	TRACK SKEW WEDGE	CYLINDER SKEW WEDGE
15	58	28	32
14	65	28	32
13	69	28	32
12	74	28	32
11	78	28	32
10	83	28	32
9	88	28	32
8	93	28	32
7	97	28	32
6	104	28	32
5	108	28	32
4	112	28	32
3	114	28	32
2	118	28	32
0, 1	118	28	32
System	93	28	32

5.5.3 Error Detection and Correction

As disk drive areal densities increase, obtaining extremely low error rates requires a new generation of sophisticated error-correction codes. Quantum Maverick 270/540S series hard disk drives implement 112-bit single- and double-burst Reed-Solomon error-correction techniques to reduce the unrecoverable read error rate to less than one bit in 1×10^{14} bits read.

When errors occur, an automatic retry and a more rigorous double-burst correction algorithm enable the correction of any sector with two bursts of three incorrect bytes each or up to six multiple random one-byte burst errors. In addition to these advanced error correction capabilities, the drive's additional cross-checking code and algorithm double checks the main ECC correction to greatly reduce the probability of miscorrection.

5.5.3.1 Background Information on ECC and ECC On-The-Fly

A sector for the Maverick 270/540S hard disk drives is comprised of 512 bytes of user data, followed by two cross-checking (XC) bytes (16 bits), followed by twelve ECC check bytes (96 bits). The two cross-checking bytes are used to double check the main ECC correction and reduce the probability of miscorrection. Single burst errors of up to 24 bits within one sector can be corrected "on-the-fly," in real time as they occur, allowing a high degree of data integrity with no impact to the drive's performance.

The drive does not need to re-read a sector on the next disk revolution or apply ECC for those errors that are corrected on-the-fly. Errors corrected in this manner are invisible to the host system.

The twelve ECC check bytes shown in Figure 5-6 are used to detect and correct errors. The cross-checking and ECC data is computed and appended to the user data when the sector is first written.

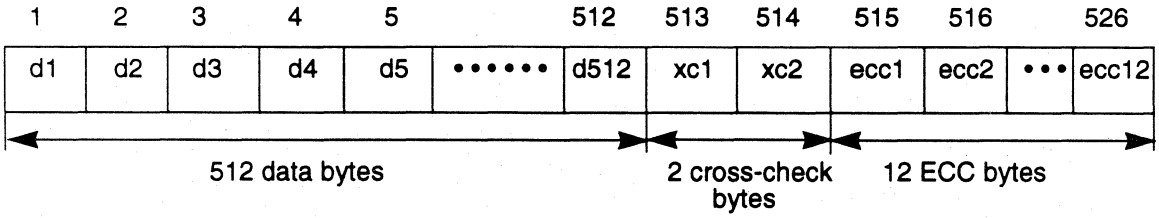

Figure 5-6 *Sector Data Field with ECC Check Bytes*

To obtain the ECC check byte values, each byte (including cross-checking and ECC bytes) within the sector is interleaved into one of three groups, where the first byte is in interleave 1, the second byte is in interleave 2, the third byte is in interleave 3, the fourth byte is in interleave 1, the fifth byte is in interleave 2, and so on, as shown in Figure 5-7.

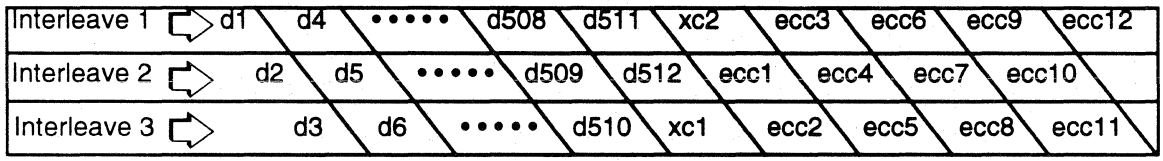

Figure 5-7 *Byte Interleaving*

Note: ECC interleaving is not the same as the sector interleaving that is done on the disk.

Each of the three interleaves is encoded with four ECC bytes, resulting in the twelve ECC bytes at the end of the sector. The two cross checking bytes are derived from all 512 data bytes. The combination of the interleaving and the nature of the ECC formulas enable the drive to know where the error occurs.

Because the ECC check bytes follow the cross checking bytes, errors found within the cross-checking bytes can be corrected. Due to the power and sophistication of the code, errors found within the ECC check bytes can also be corrected. Errors in the cross-check or ECC bytes can be corrected on-the-fly if their characteristics fall into the guidelines of a single burst correctable error.

Each time a sector of data is read, the Maverick 270/540S hard disk drives generate a new set of ECC check bytes and cross-checking bytes from the user data. These new check bytes are compared to the ones originally written to the disk. The difference between the newly computed and original check bytes is reflected in a set of twelve *syndromes* and two cross checking syndromes, which corresponds to the number of check bytes. If all of the syndrome values equal zero, the data was read with no errors, and the sector is transferred to the host system. If any of the syndromes do not equal zero, an error has occurred. The type of correction the drive applies depends on the nature and the extent of the error.

High speed on-the-fly error correction saves several milliseconds on each single burst error, because there is no need to wait for a disk revolution to bring the sector under the head for re-reading.

Correction of Single Burst Errors On-The-Fly

Single-burst errors may have up to three erroneous bytes (24 bits) within a sector, provided that each byte of the three must occur in a different interleave. In other words, if the first error bit is in interleave 1, the last error bit must occur no later than interleave 3. If the first error bit falls in interleave 1, and the last error bit falls in the next interleave 1, the error is uncorrectable on-the-fly. In Figure 5-8, the 24-bit error is correctable, because it is spread across three distinct interleaves. The 18-bit error is uncorrectable on-the-fly because it falls across four interleaves: two interleave 1s, interleave 2, and interleave 3.

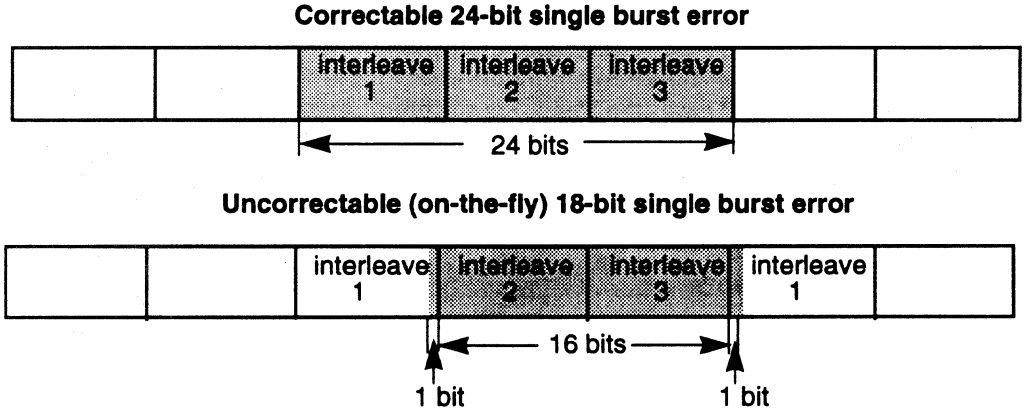

Figure 5-8 *Single Burst Error Correctability*

Note: In Figure 5-8, the shaded portions represent data containing errors

The 18-bit error can be corrected if the drive rereads the sector and applies double burst error correction techniques. Any 17-bit error can always be corrected on-the-fly because each byte is guaranteed to occupy different interleaves.

5.5.3.2 Correction Of Double-Burst Errors

Through sophisticated algorithms, Maverick 270/540S drives have the capability to correct double-burst errors, even though the probability of their occurrence is low. Double-burst errors can be simply viewed as two spans of errors within one sector. More specifically, correctable double-burst errors must have two or fewer erroneous bytes per interleave. This allows the drive's Reed-Solomon ECC to correct double-burst errors up to 48 bits long (provided that the error consists of two or fewer bytes residing in each of the interleaves).

If the double-burst correction is successful, the data from the sector can be written to a spare sector, and the logical address will be mapped to the new physical location.

Double-Burst Error Examples

Of the examples shown in Figure 5-9, examples A and B are correctable because no more than two error bytes of the entire error reside in any one of the interleaves. The other 42-bit error, example C, is uncorrectable because it occupies more than two erroneous bytes per interleave.

Note: Any 41-bit error burst can be corrected using double-burst error correction because no more than two bytes can occupy each interleave.

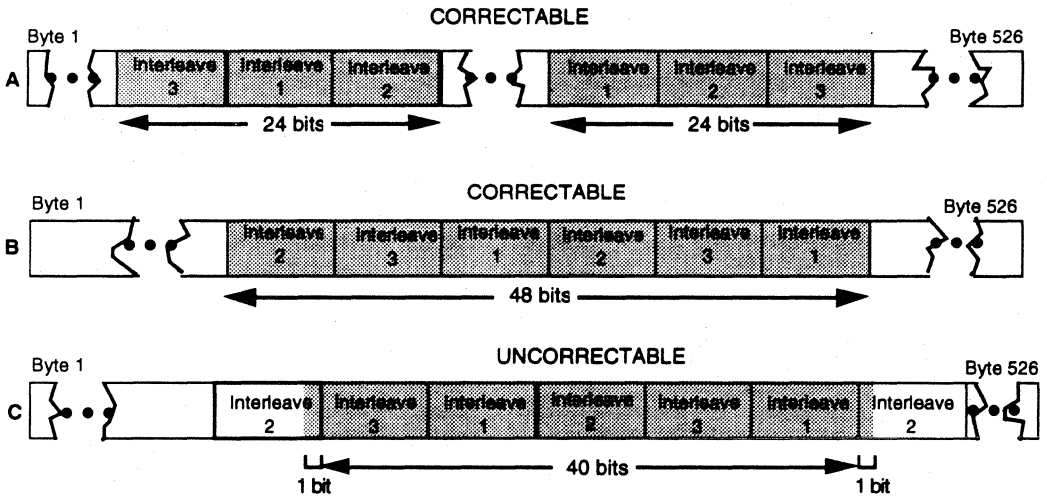

Figure 5-9 *Correctable and Uncorrectable Double-Burst Errors*

5.5.3.3 Multiple Random Burst Errors

The drive's ECC can correct up to 48 bits of multiple random errors, provided that the incorrect bytes follow the guidelines for correctable double burst errors. Up to 24 bits of multiple random errors can be corrected on-the-fly if only one byte per interleave contains an error. If more than two bytes in any one interleave are in error, the sector cannot be corrected. Figure 5-10 shows an example of a correctable random burst error consisting of 6 bytes (48 bits). This random burst error is correctable because no more than two bytes within each interleave are in error.

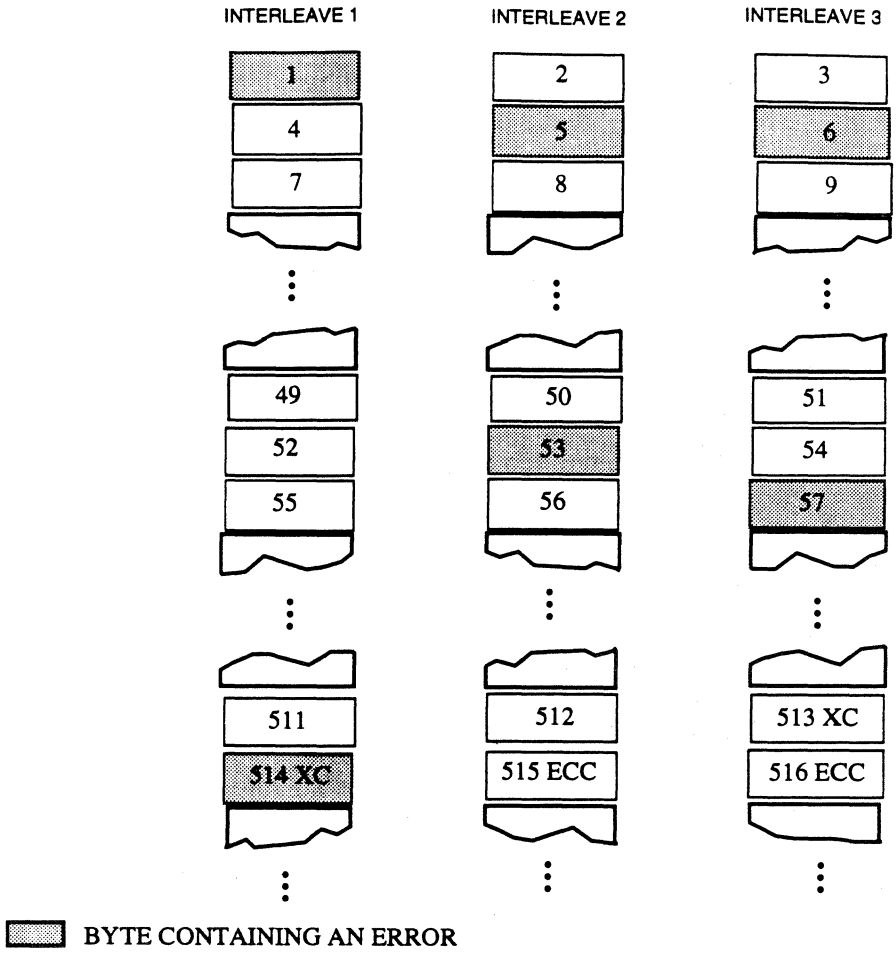

Figure 5-10 *Six Correctable Random Burst Errors*

5.5.3.4 ECC Error Handling

When an ECC error occurs, the sector is re-read up to eight times (default) in an attempt to read the data correctly without applying the complex double-burst ECC correction algorithm. This strategy prevents invoking correction on non-repeatable errors. Each time a sector in error is re-read, a set of ECC syndromes is computed. If all of the syndrome values equal zero, the data was read with no errors, and the sector is transferred to the host system. If any of the syndrome values does not equal zero, an error has occurred, the syndrome values are retained, and another re-read is invoked.

Note: Non-repeatable errors are usually related to the signal-to-noise ratio of the system. They also can represent marginal conditions of the media, heads, read/write circuitry, or data synchronizer circuits but are not due to media defects.

When the sets of syndromes from two consecutive re-reads are the same, a stable syndrome has been achieved. This event can be significant depending on whether the automatic read reallocation or early correction features have been enabled.

Note: These features can be enabled or disabled on SCSI Mode Select Page 1H through the SCSI MODE SELECT command. The EEC bit enables early ECC double-burst correction if a stable syndrome has been achieved before all of the re-reads have been exhausted. The ARRE bit enables the automatic reallocation of defective sectors.

If the early correction feature has been enabled and a stable syndrome has been achieved, double-burst ECC correction is applied, and the appropriate message is transferred to the host system (e.g., corrected data, etc.).

If the automatic read reallocation feature is enabled, the drive will attempt up to 24 re-reads (3 times the retry count set in SCSI Mode page 1H) for a double-burst error.

In this case, the ECC correction algorithm is divided into four parts.

1. When an ECC error is encountered, the drive will try to reread the correct data eight times, aborting the re-reads if the correct data was successfully read The drive will always try to re-read the correct data before applying double-burst ECC correction and automatic sector reallocation.

2. If the first eight re-reads were unsuccessful, the drive will attempt to re-read the sector eight more times. This time, if the drive does not recover from the error with a re-read, it will look for a stable syndrome from two consecutive re-reads. If it finds a stable syndrome, it will try to correct the sector and set a flag indicating that the sector is eligible for reallocation. If the sector is uncorrectable, the drive continues to re-read until the second set of eight re-reads has been exhausted.

3. In addition to the next group of eight re-reads, the drive either looks for a stable syndrome from two consecutive re-reads or for an ECC correctable error. If the drive discovers a stable syndrome from two consecutive re-reads, it will set a flag indicating that the sector is eligible for reallocation.

Note: If at any time during this third set of re-reads the drive discovers that the error is ECC correctable, it will correct the error.

4. When the ECC correction algorithm has been completed and the data was corrected by ECC or found to be uncorrectable, the reallocation algorithm is invoked.

   Reallocation is enabled if both of the following events occur:

   - The flag indicating that the sector is eligible for reallocation was set by the ECC algorithm.
   - The Automatic Read Reallocation Enabled (ARRE) bit is set.

   If the error was not ECC correctable, the Reallocate Uncorrectable Error Enable (RUEE) must also be set. The reallocation algorithm first tests if the sector is defective by performing a sequence of ten write verifies; the recovered data is written to the sector and then reread. If an error is detected during the write verify procedure, the recovered data is reallocated to a spare sector, and the bad sector location is entered in the grown defect list.

   If the recovered data was not ECC correctable, the data is written using a special data ID field. This is done so that when data is reread, a "Read failure on uncorrectable data previously reallocated" (sense key/code 3,0AA) error will be reported.

   If there is no error during the write verify procedure, the data is not reallocated to a spare sector. However, if the recovered data was not ECC correctable, the error described above will be reported when the data is reread.

Note: The RUEE bit can be enabled or disabled on SCSI Mode Select Page 39H through the SCSI Mode Select command. Setting the RUEE bit enables the automatic reallocation of uncorrectable or unrecovered (e.g., ID not found) sectors.

For testing purposes, the Disable Correction bit can be set to one, in which case the drive will continue to re-read per the value in the "Retry Count." The default value is eight. Re-reads can be disabled by setting this value to zero.

5.5.3.5 Error Types

The following subsections provide a brief description of two types of drive related data errors.

Soft Errors

Soft errors are errors that are not readily repeatable. Soft errors are usually related to the system signal to noise ratio. In addition, they can be caused by marginal conditions of the heads, read/write circuitry, or the data synchronizer circuits. They are not caused by media defects. An error recovered from a re-read (where the syndrome values for the sector are zero) is classified as a soft error.

Hard Errors

Hard errors are repeatable with high probability. They are usually caused by media defects (pits, scratches, or thin spots). Defective media errors can be detected and their associated locations can be skipped (not used for data storage). An error recovered from two identical syndromes from two consecutive re-reads is classified as a hard error (because of its repeatability). Unrecoverable errors also are classified as hard errors.

5.5.4 Defect Management

The Maverick 270/540S hard disk drives allocate one sector per cylinder as a spare. In the factory, the media is scanned for defects. If a sector on a cylinder is found to be defective, the address of the sector is added to the drive's defect list. Sectors located physically subsequent to the defective sector are assigned logical block addresses such that a sequential ordering of logical blocks results. This inline sparing technique is employed in an attempt to eliminate slow data transfer that would result from a single defective sector on a cylinder.

If more than one sector is found defective on a cylinder, the inline sparing technique is applied only to the first sector. The remaining defective sectors are replaced with the nearest available spare sectors on nearby cylinders. Such an assignment of additional replacement sectors from nearby sectors rather than having a central pool of spare sectors is an attempt to minimize the motion of the actuator and head that otherwise would be needed to find a replacement sector. The result is minimal reduction of data throughput.

Defects that occur in the field are known as *grown* defects. If such a defective sector is found in the field, the sector is reallocated according to the same algorithm used at the factory for those sectors that are found defective *after* the first defective sector on a cylinder; that is, inline sparing is not performed on these grown defects. Instead, the sector is reallocated to an available spare sector on a nearby cylinder.

Sectors are considered to contain grown defects if the double-burst ECC algorithm must be applied to recover the data. If this algorithm is successful, the corrected data is stored in the newly allocated sector. If the algorithm is not successful, the erroneous data is stored in the newly allocated sector, and a flag is set in the data ID field that causes the drive to report an ECC error each time the sector is read. This condition remains until the sector is rewritten.

Chapter 6
SCSI DESCRIPTION

6.1 GENERAL DESCRIPTION

The SCSI bus supports up to eight SCSI devices, configured as any combination of initiators and targets. The Maverick 270/540S hard disk drive has a fixed role as a target. Certain SCSI bus functions are assigned to the initiator or the target:

- The initiator can arbitrate for the SCSI bus and select a target.
- The target can request the transfer of COMMAND, DATA, STATUS, or other information on the DATA BUS and, in some cases, can arbitrate for the SCSI bus and reselect an initiator to continue an operation.

Communication with the Maverick 270/540S hard disk drive takes place on the SCSI data bus, in the asynchronous mode, using a defined REQ/ACK handshake protocol.

Later sections in this chapter are intended to familiarize the user with the Message Codes, Asynchronous Data Transfer Protocol, Status Bytes, and Commands supported by the Maverick 270/540S.

6.2 LOGICAL CHARACTERISTICS OF THE SCSI BUS

The operating environment of the SCSI bus consists of a series of logical divisions called phases. There are eight phases, during which a device only can perform the tasks allowed for that phase. The bus can be in only one phase at any given time. Transition between phases is described in the explanations of the phases later in this section. The eight phases of the SCSI bus are:

- Bus Free
- Arbitration
- Selection
- Reselection
- Command
- Data
- Status
- Message

Note: The Command, Data, Status, and Message phases are collectively known as the Information Transfer (IT) phases.

6.2.1 Bus Phase Sequences

A SCSI device transits from one bus phase to another in a predefined manner. The exact sequence of the bus phases depends on whether or not the SCSI host system implements arbitration.

In all cases, a reset condition aborts any active phase. A reset condition is always followed by the Bus Free phase.

6.2.1.1 Non-Arbitrating Systems

SCSI systems that do not implement arbitration are designed to support only one initiator, and do not support reselection of the initiator by the target. In these systems, the initiator is not required to arbitrate for the bus because there are no other devices capable of getting it.

Non-arbitrating systems transit from Bus Free directly to Selection, omitting the Arbitration phase.

6.2.1.2 Arbitrating Systems

SCSI systems that implement arbitration are designed to support multiple initiators, and support reselection of the initiator by the target. In these systems, devices contest by SCSI ID to get the bus, with the device having the highest SCSI ID winning the bus.

Arbitrating systems transit from Bus Free to Arbitration.

6.2.2 Signal Restrictions Between Phases

Between information transfer phases, the following restrictions apply to signals on the SCSI bus:

- BSY, SEL, REQ, and ACK signals do not change.
- C/D, I/Q, MSG, and DB signals can change.

  When switching the direction of the signal path on the DB lines from out to in—that is, from an initiator-driven signal to a target-driven signal—after asserting I/O, the drive delays driving the DATA BUS for at least a data release delay plus a bus settle delay; and the initiator releases the bus within a data release delay after the transition of I/O to true. When switching the direction of the signal path on the bus from in to out—that is, from a target-driven signal to an initiator-driven signal—after negating I/O, the drive releases the bus within a deskew delay.
- ATN and RST signals can change.

6.2.3 Bus Free

The Bus Free phase is active when BSY and SEL have been released for at least a bus settle delay. All devices resident on the bus must release all signal lines within a bus clear delay of Bus Free. This means that within (bus settle delay)+(bus clear delay) of the point at which BSY and SEL were released, the device must release all bus lines.

During Bus Free, a device might only attempt to get the bus. For systems that implement arbitration, a device attempting to get the bus must move into the Arbitration phase. Systems that do not implement arbitration are the initiators of single-initiator systems.

6.2.4 Arbitration

The Arbitration phase is active when a device asserts BSY and its SCSI ID bit on the bus. The ID bit is asserted across the corresponding DB line (ID #1 to DB1, and so on).

Arbitration consists of a device asserting BSY and the SCSI ID, then checking the data lines for an asserted SCSI ID with higher priority. SCSI ID numbers range from 0 to 7, lowest to highest. If no higher priority ID is asserted, then the device gets the bus, and all other devices must release all signals. If a higher priority ID is asserted, the device must release all lines and wait for the next Bus Free phase to try again. The timing constraints within the Arbitration phase are described in the following paragraphs.

To initiate arbitration (attempt to get the bus), a device must wait for at least a bus free delay after detection of Bus Free before asserting BSY and its SCSI ID. Once it has initiated arbitration, the device must wait at least an arbitration delay before testing the data bus for higher priority IDs. If no higher priority ID is asserted, the device wins arbitration and asserts SEL. The winning device must wait at least a bus clear delay plus a bus settle delay before initiating any signal line transitions.

If the device loses arbitration (higher priority ID asserted), it must release all signal lines within a bus clear delay from the point at which SEL is asserted by the winning device.

6.2.5 Selection

The SELECTION phase allows an initiator to select a target—the Maverick 270/540S hard disk drive—to initiate a target function, such as a READ or WRITE command. The drive is never the initiator.

Note: To distinguish this phase from the reselection phase, the I/O signal remains undriven.

6.2.5.1 Nonarbitrating Systems

If the ARBITRATION phase is not implemented, the initiator must:

1. Detect a BUS FREE phase.
2. Wait for at least the duration of a BUS CLEAR DELAY.
3. Except in single-initiator environments in which initiators employ the single-initiator option, assert the following on the DATA BUS:
   - Its SCSI ID
   - The target's SCSI ID
   - If parity is implemented, the parity bit
4. After two DESKEW DELAYS, assert SEL.

See Section 6.2.5.4, "Single-Initiator Option" on page 6-4.

6.2.5.2 Arbitrating Systems

If the ARBITRATION phase is implemented, the device that wins an arbitration:

1. Asserts both BSY and SEL.
2. Waits for at least the duration of a BUS CLEAR DELAY plus a BUS SETTLE DELAY before ending the ARBITRATION phase.

If the device doesn't assert I/O, it becomes an initiator. Except in single-initiator environments in which initiators employ the single-initiator option, an initiator must:

1. Assert the following on the DATA BUS:
   - Its SCSI ID
   - The target's SCSI ID
   - If parity is implemented, the parity bit is set.
2. Wait for at least two DESKEW DELAYS before it can release BSY.
3. Wait for at least a BUS SETTLE DELAY before it can receive a response from the drive.

6.2.5.3 All Systems

In all systems, the Maverick 270/540S hard disk drive:

1. Becomes selected when SEL and its SCSI ID bit are true, and BSY and I/O are false for at least the duration of a BUS SETTLE DELAY.
2. Examines the DATA BUS to determine the SCSI ID of the selecting initiator, unless the initiator employed the single-initiator option. (See Section 6.2.5.4, "Single-Initiator Option" for more information.)
3. Asserts BSY within a SELECTION-ABORT TIME of its most recently detected selection, to ensure the correct operation of the time-out procedure.

If parity is implemented and the drive detects bad parity, or if more than two SCSI ID bits appear on the DATA BUS, the drive will not respond to a selection.

After the initiator detects BSY is true, it must:

1. Wait at least two DESKEW DELAYS before it can release SEL.
2. Release the DATA BUS.

The drive waits for the initiator to release SEL before setting an information transfer phase. Once drive selection is complete, the LED on the drive lights.

6.2.5.4 Single-Initiator Option

During a SELECTION phase, only the target's SCSI ID bit can be set by an initiator that neither implements the RESELECTION phase nor operates in a multiple-initiator environment—not the initiator's own SCSI ID bit. Thus, it is impossible for a target to determine the initiator's SCSI ID.

6.2.5.5 Selection Timeout Procedure

If the initiator waits for at least a SELECTION TIMEOUT DELAY without receiving a BSY response from the drive, it executes one of two selection-time-out procedures to clear the SCSI bus. Either:

- The initiator asserts a RST signal. (For more information, see Section 6.2.8.2, "Reset Condition" on page 6-10.)
- The initiator continues to assert SEL and releases the DATA BUS. If the initiator does not detect BSY as true after waiting at least a SELECTION-ABORT TIME plus two DESKEW DELAYS, it can release SEL, allowing the SCSI bus to go to a BUS FREE phase. The drive responds to a selection only if the selection remains valid for the duration of a SELECTION TIME following its assertion of BSY. This prevents improper selections—such as two targets connected to the same initiator, the wrong target connected to an initiator, or a target connected to no initiator.

6.2.6 Reselection

A RESELECTION phase allows the Maverick 270/540S hard disk drive to reconnect to an initiator, to continue an operation previously started by the initiator, but suspended by the drive when the drive disconnected, allowing a BUS FREE phase to occur before the operation was complete.

RESELECTION can be used only in systems that implement the ARBITRATION phase. Following an ARBITRATION phase, the drive:

1. Asserts I/O to become a target.
2. Sets the DATA BUS to a value equivalent to the OR of its SCSI ID bit.
3. Sets the initiator's SCSI ID bit.
4. Sets the parity bit if parity is implemented.
5. Waits for at least the duration of two DESKEW DELAYS, then releases BSY.
6. Waits for at least a BUS SETTLE DELAY before it is ready to receive a response from the initiator.

The initiator becomes reselected when SEL, I/O, and its SCSI ID bit are true, and BSY is false for at least a BUS SETTLE DELAY. The reselected initiator examines the DATA BUS to determine the SCSI ID of the reselecting target. To ensure the correct operation of the time-out procedure, the initiator must assert BSY within a SELECTION-ABORT TIME of its most recently detected reselection. If parity is implemented and the initiator detects bad parity, or if more than two SCSI ID bits appear on the DATA BUS, the initiator must not respond to a RESELECTION.

After it detects BSY, the drive:

1. Asserts BSY
2. Waits for at least two DESKEW DELAYS, then releases SEL
3. Asserts I/O
4. Sets the DATA BUS

After the reselected initiator detects SEL as false, it must release BSY. The drive continues asserting BSY until it is ready to relinquish the SCSI bus.

6.2.6.1 Reselection Timeout Procedure

During a RESELECTION phase, if a target waits for at least a SELECTION TIMEOUT DELAY and receives no BSY response from the initiator, it releases all DATA BUS signals, but continues to assert SEL and I/O. Then, after the target waits for the duration of at least one SELECTION ABORT TIME plus two DESKEW DELAYS, if it does not detect BSY as true, it releases SEL and I/O, allowing the SCSI bus to enter the BUS FREE phase for a SELECTION-TIMEOUT DELAY.

When responding to a reselection, the initiator must verify that the reselection remains valid for a SELECTION-ABORT TIME following its assertion of BSY. Thus, no improper reselection can result—such as two initiators connected to the same target or the wrong initiator connected to a target. If reselection is unsuccessful after three attempts, the target terminates the command and returns to the BUS FREE phase.

6.2.7 Information Transfer (IT) Phases

The four phases that perform all of the work in a SCSI transfer share a common bus mechanism, and so are grouped together under the name *Information Transfer (IT) phases.* Each of the information transfer phases is distinguished from the others by the combination of the MSG, C/D, and I/O signal lines. The four information transfer phases are:

- Command
- Data
- Status
- Message

The state of the MSG, C/D, and I/O lines determines the active phase. These lines are controlled by the target. They are asserted a BUS SETTLE DELAY prior to the first information transfer handshake and remain valid until after negation of ACK in the last handshake. The state of these three lines and other transfer phase information are shown in Table 6-1. In SCSI, for all transfer-type operations, "in" refers to information (data, messages) being transferred to the initiator, while "out" refers to information being transferred from the initiator.

Table 6-1 *Information Transfer Phase Detail*

SIGNAL STATE			PHASE	DIRECTION OF TRANSFER	ACTION
MSG	C/D	I/O			
0	0	0	Data Out	Initiator to target	*The target requests data from the initiator.*
0	0	1	Data In	Target to initiator	*The target sends data to the initiator.*
0	1	0	Command	Initiator to target	*The target requests a command from the initiator.*
0	1	1	Status	Target to initiator	*The target sends status information to the initiator.*
1	0	0	Reserved	Reserved	—
1	0	1	Reserved	Reserved	—
1	1	0	Message Out	Initiator to target	*The target requests messages from the initiator.*
1	1	1	Message In	Target to initiator	*The target sends messages to the initiator.*

(0 = false, 1 = true)

All information transfer phases use the REQ/ACK handshake to control transfer. BSY remains asserted and SEL remains released during all information transfer phases. Transfers across the bus must be asynchronous for all phases.

6.2.7.1 Pointers

In the SCSI architecture, two sets of three pointers for each I/O process reside in each initiator—the active pointers and the saved pointers. Each set of pointers consists of a command pointer, a data pointer, and a status pointer.

When an I/O process becomes active, the initiator copies the values from its saved pointers into its active pointers. The active pointers represent the current state of the interface and point to the next command, data, or status byte to be transferred between the initiator and the target. The target currently connected to the initiator uses the active pointers.

A set of saved pointers exists for each command currently active—whether or not it is currently connected. A saved command pointer indicates the beginning of the command descriptor block for the current I/O process. A saved status pointer indicates the beginning of the status area for the current I/O process.

A saved data pointer indicates the beginning of the data area for the current I/O process. Its value remains constant until the target sends a SAVE DATA POINTER message to the initiator. In response to this message, the initiator copies the value of the active data pointer into the saved data pointer for the current I/O process. When a target disconnects from the SCSI bus, the initiator retains only the saved pointer values. When a target reconnects to the SCSI bus, the initiator restores the current pointer values from the saved values.

The drive does not issue a SAVE DATA POINTER message if the drive disconnects prior to transferring any data during a read operation. If the drive uses disconnect messages to break a long data transfer into two or more shorter transfers, it issues a save data pointer message before each disconnect message.

6.2.7.2 Asynchronous Transfer Mode

The drive uses the I/O signal to control the direction of information transfer. When I/O is true, the drive transfers information to the initiator. When I/O is false, the initiator transfers information to the drive.

To transfer information to the initiator if I/O is true:

1. The drive sets DB(7-0,P), waits for at least the duration of one DESKEW DELAY plus a CABLE SKEW DELAY, then asserts REQ.
2. The initiator reads DB(7-0,P) after REQ goes true, then acknowledges receipt of the data by asserting ACK.
3. When the drive detects ACK as true, it sets or releases DB(7-0,P), and releases REQ.
4. After REQ goes false, the initiator negates ACK.
5. Once ACK goes false, the drive can continue the data transfer by again setting DB(7-0, P) and asserting REQ.

To transfer information to the target, or drive if I/O is false:

1. The drive requests information by asserting REQ.
2. The initiator sets DB(7-0,P), waits for at least the duration of one DESKEW DELAY plus a CABLE SKEW DELAY, then asserts ACK.
3. When the drive detects ACK as true, it reads DB(7-0,P), then releases REQ.
4. When the initiator detects REQ as false, it sets or releases DB(7-0,P), and negates ACK.
5. The drive can continue the data transfer by again asserting REQ.

6.2.7.3 Synchronous Transfer Mode

In the data phase, synchronous data transfer can occur only when the initiator and target establish a synchronous data-transfer agreement. This agreement specifies the REQ/ACK offset and the minimum transfer period. For information about the Synchronous Data Transfer Request, see Section 6.4.5, "Extended Messages."

The REQ/ACK offset specifies the maximum number of REQ pulses, in excess of the number of ACK pulses received from the initiator, that can be sent by the target. If the number of REQ pulses exceeds the REQ/ACK offset, the target waits to assert REQ again until it receives the leading edge of the next ACK pulse. To successfully complete the data phase, the number of ACK and REQ pulses must be equal.

On synchronous data transfers:

1. The target asserts REQ for at least an ASSERTION PERIOD. Before asserting REQ again, the target must wait for at least a TRANSFER PERIOD from the last transition of REQ to true or one NEGATION PERIOD from the last transition of REQ to false, whichever is greater.

2. As soon as it receives the leading edge of a REQ pulse, the initiator asserts ACK and sends one byte of data. The ACK signal remains asserted for at least one ASSERTION PERIOD. Then, before asserting ACK again, the initiator must wait for at least a TRANSFER PERIOD from the last transition of ACK to true or one NEGATION PERIOD from the last transition of ACK to false, whichever is greater.

To transfer information to the initiator if I/O is true:

1. The drive sets DB(7-0,P), waits for at least one DESKEW DELAY plus one CABLE-SKEW DELAY, then asserts REQ.

2. The DB(7-0,P) signals remain valid for at least one DESKEW DELAY plus one CABLE-SKEW DELAY and one HOLD TIME, following the assertion of REQ.

3. The target asserts REQ for at least the duration of an ASSERTION PERIOD. It can then negate REQ, and set or release DB(7-0,P).

4. The initiator reads the value on DB(7-0,P) within one HOLD TIME of the transition of REQ to true, then responds with an ACK pulse.

To transfer information to the drive if I/O is false:

1. The initiator transfers one byte on receiving each REQ pulse, sets DB(7-0,P) to the appropriate values, waits for at least one DESKEW DELAY plus one CABLE-SKEW DELAY, then asserts ACK.

2. DB(7-0,P) remains valid for at least one DESKEW DELAY, one CABLE-SKEW DELAY, and one HOLD TIME after the assertion of ACK.

3. ACK remains asserted for a minimum of one ASSERTION PERIOD.

4. The initiator negates ACK, and changes or releases DB(7-0,P).

5. The drive reads the value of DB(7-0,P) within one HOLD TIME of the transition of ACK to true.

6.2.8 SCSI Bus Conditions

Two asynchronous conditions can exist on the SCSI bus—the ATTENTION condition and the RESET condition. These conditions cause a SCSI device to perform certain actions and might alter the phase sequence.

6.2.8.1 Attention Condition

An ATTENTION condition (ATN) allows an initiator to inform a target that it has a message ready. The drive can obtain the message by initiating a MESSAGE OUT phase.

- The initiator can enter an ATTENTION condition at any time by asserting ATN, except during an ARBITRATION or BUS FREE phase.
- The drive responds with a MESSAGE OUT phase on selection and on transitions to phases other than a DATA phase.

To transfer more than one byte, the initiator holds ATN asserted. The initiator can negate ATN at any time, except while asserting ACK during a MESSAGE OUT phase. Typically, the initiator negates ATN while REQ is true and ACK is false, during the last REQ/ACK handshake of the MESSAGE OUT phase.

6.2.8.2 Reset Condition

The RESET condition (RST) clears all SCSI devices from the bus, and takes precedence over all other phases and conditions. SCSI devices set the RESET condition by asserting RST for at least the duration of a RESET-HOLD TIME. While a RESET condition is asserted, the states of all SCSI bus signals other than RST remain undefined.

Note: The Maverick 270/540S hard disk drive never asserts RST, but responds to a RST signal asserted by another SCSI device.

When another device asserts RST, the drive releases all SCSI bus signals within a BUS CLEAR DELAY of the transition of RST to true. A BUS FREE phase always follows a RESET condition. When the drive executes a hard reset, it:

- Clears all I/O processes
- Releases all reservations of SCSI devices
- Sets an ATTENTION condition

Mode parameters retain their saved values.

Note: The drive cannot be selected for at least 50 microseconds following the release of the RST signal.

6.3 SCSI STATUS

At the termination of each command, during the STATUS phase, the drive sends a status byte to the initiator—as specified in Tables 6-2, 6-3, and 6-4—unless one of the following events terminates the command:

- An ABORT message
- A BUS DEVICE RESET message
- A hard RESET condition
- An unexpected disconnect

Table 6-2 shows the status bytes; Table 6-3 shows the status byte code bit values for the Maverick 270/540S hard disk drive.

Table 6-2 *Status Byte*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED = 0		STATUS BYTE CODE					RSVD

Table 6-3 *Status Byte Bit Values*

Status Byte Bits								Condition
7	6	5	4	3	2	1	0	
R	R	0	0	0	0	0	R	GOOD
R	R	0	0	0	0	1	R	CHECK CONDITION
R	R	0	0	1	0	0	R	BUSY
R	R	0	1	0	0	0	R	INTERMEDIATE/GOOD
R	R	0	1	1	0	0	R	RESERVATION CONFLICT

GOOD (00_H): The drive successfully completed a linked command, without encountering an error condition.

CHECK CONDITION (02_H): An error, exception, or abnormal condition occurred during the drive's execution of a linked command, which may or may not have completed successfully. To determine the nature of the abnormal condition, the initiator must issue a REQUEST SENSE command.

BUSY (08_H): The drive is busy and unable to process a command sent by an initiator. Generally, when the status for the drive is BUSY, the drive has disconnected to execute a previous command and cannot execute an additional command at that time. The initiator must reissue the command.

INTERMEDIATE/GOOD (10_H): The drive successfully completed a linked command, without encountering an error condition. Typically, the drive issues this status code for each command in a series of linked commands but the last command, for which the link bit is zero. When an error or abnormal condition causes the drive to issue a CHECK CONDITION or RESERVATION CONFLICT status code instead, execution of the series of linked commands terminates.

RESERVATION CONFLICT (18_H): The initiator attempted to access a logical unit reserved by another SCSI device that has a conflicting reservation type. The drive rejects the command. The initiator must reissue the command.

6.4 SCSI MESSAGES

Because there is no dedicated bus controller module, the two devices actively using the bus must be able to communicate with one another in order to manage the physical path between them across the bus. This is accomplished by using a system of pre-defined software messages. Messages can be single-byte or extended.

The only mandatory message is the single-byte Command Complete (00_H). However, if you implement any of the other messages, you must implement Message Reject (07_H).

6.4.1 MESSAGE IN Phase

A MESSAGE IN phase allows the drive to request that messages be sent from the drive to the initiator. During the REQ/ACK handshakes of this phase, the drive asserts I/O, C/D, and MSG.

6.4.2 MESSAGE OUT Phase

A MESSAGE OUT phase allows the drive to request that a message be sent from the initiator to the drive. The drive invokes this phase in response to an ATTENTION condition created by the initiator. (See Section 6.2.8.1, "Attention Condition" on page 6-10.) During the REQ/ACK handshakes of this phase, the drive asserts C/D and MSG, and negates I/O. In this phase, the drive handshakes bytes until ATN goes false or the drive rejects a message.

If the drive receives all of the message bytes without any parity errors, it changes to a different information-transfer phase and transfers at least one byte, or changes to the BUS FREE phase—for example, in response to an ABORT or BUS DEVICE RESET message.

If the drive detects one or more parity errors in the message bytes, it requests that the message be sent again by the initiator, by asserting REQ after ATN goes false, but before changing to another phase. On detecting REQ, the initiator again sends the message bytes. When sending multiple message bytes, the initiator asserts ATN at least two DESKEW DELAYS prior to asserting ACK on the first byte. ATN remains asserted until the initiator sends the last byte.

Following a parity error, the drive ignores messages sent during the ATN condition. When the initiator resends a series of messages following a parity error, the drive does not act on any message already received.

6.4.3 Message Protocol

All SCSI devices implement the command complete message. To maximize flexibility, the drive also supports all single-byte, nonextended messages defined in SCSI-2. Until the initiator informs the drive that it supports other messages, the drive sends only command complete messages. During a Selection phase, the initiator indicates that it will generate a message reject message if it does not understand a message from the drive, by asserting ATN before SEL goes true and BSY goes false on the SCSI bus.

When the drive detects an ATN signal asserted by the initiator during a Selection phase, it enters a Message Out phase to receive a message byte. The drive expects the initiator to send an identify message to establish the physical path to a logical unit specified by the initiator. This message should specify LUN=0. If the initiator supports disconnect/reconnect, it should set bit 6 of the identify message. Without bit 6 set, the drive will not attempt a disconnect during the execution of a command. Bit 6 should be set only if the initiator supports the save data pointer and disconnect messages. Alternatively, the drive can receive a bus device reset or abort message. If the drive receives any other message, it sends a message reject message and goes to Bus Free. When the drive receives an identify message, it enters the Command phase and requests command bytes from the initiator.

6.4.4 Single-Byte Messages

Single-byte messages provide the simplest means of managing the bus. These messages are used to indicate the state of command execution or data transfer operations, to establish a physical path, and to control the condition of the bus. For single-byte messages, the hex value of the byte corresponds to the code assigned to the message being sent, as shown in Table 6-4.

Table 6-4 *Single-Byte Messages*

CODE	MESSAGE
00_H	command complete
01_H	extended message-synchronous data request
02_H	save data pointer
03_H	not supported
04_H	disconnect
05_H	initiator detected error
06_H	abort
07_H	message reject
08_H	no operation
09_H	message parity error
$0A_H$	linked command complete
$0B_H$	linked command complete (with flag)
$0C_H$	bus device reset
$0D_H$–$7F_H$	reserved
80_H–FF_H	identify

6.4.4.1 command complete (00_H)

This message indicates that the target has terminated execution of the command and passed status information to the initiator. It indicates only that the command was terminated, it does not specify whether the command was fully executed, or terminated prior to completing execution. The status information passed to the initiator indicates the success or failure of the command.

For linked commands, command complete is sent when the last of the linked commands has been terminated.

This message always is sent from the target to the initiator.

6.4.4.2 extended message (01_H)

This message indicates that a multiple-byte extended message is to follow. It is always the first byte of the extended message.

This message may be sent by either the target or the initiator.

6.4.4.3 save data pointers (02_H)

The drive sends this message to direct an initiator to copy the active data pointer for the drive to the saved data pointer. (See Section 6.2.7.1, "Pointers" on page 6-7 for a definition of SCSI pointers.) If the initiator rejects this message, the drive discontinues attempts to disconnect. During a DATA phase, when breaking a long data transfer into two or more shorter transfers using DISCONNECT messages, the drive must issue a SAVE DATA POINTER message immediately before issuing the DISCONNECT message.

6.4.4.4 disconnect (04_H)

The drive sends this message to inform an initiator that the current connection will be broken. The drive then disconnects by releasing BSY and goes to the BUS FREE phase. To complete the I/O operation in progress at disconnection, the drive must reconnect. The initiator does not save the data pointer. If the initiator detects a BUS FREE phase that is not preceded by a DISCONNECT or COMMAND COMPLETE message and that does not result from a RESET condition, or an ABORT or BUS DEVICE RESET message, a catastrophic error condition exists.

6.4.4.5 initiator detected error (05_H)

This message indicates that some type of non-fatal error occurred during a target-to-initiator operation (such as passing the data requested by the initiator in a READ command from the target to the initiator). As the error was non-fatal, the target may retry the operation; however, the integrity of the active pointers cannot be guaranteed.

Note: In order to set the pointers to a known-good value, you can send restore pointers. However, this will set the pointers to the value active when the current target last sent save data pointers.

This message always is sent from the initiator to the target.

6.4.4.6 Abort (06_H)

This message clears any active I/O processes for the current I_T_L nexus. (For more information about I_T_L nexus, see Section 1.4, "REFERENCE" on page 1-3.) All pointers active for the target (device or logical unit), as well as any data, are cleared, and the target moves to the Bus Free phase.

In a multiple-initiator environment, only the commands, pointers and data "belonging" to the initiator are affected by abort; other initiator's commands, pointers, and data are not affected.

Previously-established conditions, including reservation status and MODE SELECT parameters, are not cleared by this message.

This message always is sent from the initiator to the target.

6.4.4.7 message reject (07_H)

Either the drive or an initiator can send this message to indicate that the last message or message byte it received was inappropriate or has not been implemented.

To indicate that it will send this message, the initiator asserts ATN prior to releasing ACK on the REQ/ACK handshake for the message being rejected. If the drive receives this message under any other condition, it rejects the message.

To send this message, the drive goes to the MESSAGE IN phase, then sends MESSAGE REJECT before requesting additional message bytes from the initiator. Thus, the initiator can determine which message the drive is rejecting. If the initiator asserts ATN during the initial SELECTION phase of a command, it must support this message.

6.4.4.8 no operation (08_H)

This message indicates that the initiator does not have any messages to send. It is sent when the target goes to the Message Out phase and the initiator has no messages to send.

This message always is sent from the initiator to the target.

6.4.4.9 message parity error (09_H)

An initiator sends this message to the drive to indicate that one or more bytes in the last message it received had a parity error. To indicate that it will send this message, the initiator asserts ATN prior to releasing ACK on the REQ/ACK handshake for the message with the parity error. Thus, the drive can determine which message has the parity error.

6.4.4.10 linked command complete ($0A_H$)

This message indicates that the target has terminated execution of one of a series of commands and passed status information on that command to the initiator. It indicates only that the most recent command of the group was terminated; it does not specify whether the command was fully executed, or terminated prior to completing execution. The status information passed to the initiator indicates the success or failure of the command.

This message differs from command complete in that it concerns linked commands, and therefore does not imply and end to the exchange between the initiator and the target. Rather, it indicates that the initiator may update its pointers to the next command in the linked group.

Note: It is advisable to check the status information prior to sending the next linked command, as this message does not provide you with any way of knowing whether or not the command was successful. This is especially important in linked groups that are dependent upon successful execution of the preceding command.

This message always is sent from the target to the initiator.

6.4.4.11 linked command complete [with flag] ($0B_H$)

This message indicates that the target has terminated execution of a linked command that had flag of the command block set to 1, and passed status information on that command to the initiator. It indicates only that the most recent command of the group was terminated; it does not specify whether the command was fully executed, or terminated prior to completing execution. The status information passed to the initiator indicates the success or failure of the command.

This message differs from command complete in that it concerns linked commands, and therefore does not imply and end to the exchange between the initiator and the target. Rather, it indicates that the initiator may begin sending the next command in the linked group.

Note: The mechanism of setting flag and returning this message provides the best means of managing linked commands in groups that are dependent upon successful execution of the preceding command, as it provides an interrupt handle to break and check the status information returned for the completed command prior to sending the next linked command.

This message always is sent from the target to the initiator.

6.4.4.12 bus device reset ($0C_H$)

An initiator sends this message to direct the drive to clear all current commands. This message causes the drive to revert to its initial state, with no operations pending for any initiator. On receiving this message, the drive goes to the BUS FREE phase. This message does not affect MODE SELECT command parameters.

6.4.4.13 identify (80_H–FF_H)

These messages identify the logical unit of the target device involved in the exchange, and indicate whether or not the initiator supports disconnection/reconnection.

Bit 7 of these messages is always set to 1, in order to distinguish them from the other messages.

Bit 6 is used by the initiator to indicate whether or not it supports disconnection/reconnection. If it is set to 1, the initiator does support disconnection/reconnection.

Bits 3–5 are reserved.

Bits 0–3 contain the LUN of the target logical unit.

When the target sends identify following the Reselection phase, the initiator automatically restores the pointers for the target logical unit.

6.4.5 Extended messages

Extended messages provide more advanced methods of bus control, command execution management, and vendor-specific functions. For extended messages, the hex value of byte 0 is always $01, the code for the single-byte message "extended message". The basic structure of an extended message is shown in Figure 6-1. The extended message codes are given in Table 6-5.

Figure 6-1 *Extended Messages Data Structure*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	EXTENDED MESSAGE (01H)							
1	EXTENDED MESSAGE LENGTH							
2	EXTENDED MESSAGE CODE							
3	EXTENDED MESSAGE ARGUMENTS							

extended message: This parameter contains the code for the single-byte message "extended message" (01_H).

extended message length: This parameter contains the number of extended message bytes to follow. The length of the entire message, then, is equal to this parameter plus two bytes. A value of 0 in this parameter indicates that 256 bytes follow.

extended message code: This parameter contains the code defining the message.

extended message arguments: This field contains the execution arguments associated with the action ordered by the extended message. There may be several argument bytes in the argument field.

Table 6-5 *Extended Messages*

Code	Message
00_H	not used
01_H	synchronous data transfer request
02_H–$7F_H$	reserved
80_H–FF_H	vendor-specific

6.4.5.1 synchronous data transfer request (01_H)

This message is used to establish the timing parameters for synchronous data transfers between two SCSI devices capable of performing such transfers. Any device which supports synchronous transfers will send this message upon first coming online, after a hard reset of the bus, or after executing a bus device reset. An initiator sends this message to all devices on the bus. A target sends this message to all initiators on the bus. Targets and initiators may also exchange this message when one or the other wishes to initiate a synchronous data transfer operation.

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	EXTENDED MESSAGE (01H)							
1	LENGTH (03H)							
2	SYNCHRONOUS DATA TRANSFER REQUEST CODE (01h)							
3	TRANSFER PERIOD							
4	REQ/ACK OFFSET							

Figure 6-2. *Synchronous Data Transfer Request Data Structure*

extended message: This parameter contains the code for the single-byte message "extended message" (01_H).

length: This parameter is set to 03_H for this message. The total length of this message is 5 bytes.

synchronous data transfer request code: This parameter is set to 01_H.

transfer period: This field contains the value of the time period between the leading edges of successive REQ pulses and successive ACK pulses in nanoseconds. Table 6-6 shows the transfer rates the Maverick 270/540S will use based on the requested period.

Table 6-6 *Transfer Rates As a Function of Requested Period*

Requested Period (ns)	Transfer Rate Used (MB/s)
25–00	10.000
31–26	8.000
37–32	6.667
43–38	5.714
50–44	5.000
62–51	4.000
75–63	3.333
86–76	2.857
100– 87	2.500
112–101	2.222
125–113	2.000
137–126	1.818
150–138	1.667
162–151	1.538
174–163	1.429
187–175	1.333
255–188	0.625

Note: In Table 6-6, multiply the requested period by 4 to get the actual transfer period. Transfer rates of 5.714 MB/s and greater use Fast SCSI timing.

REQ/ACK offset: This field contains the value of the number of REQ pulses that may be received at the initiator prior to returning an ACK pulse. A value of 00_H in this field indicates that the transfer must take place in asynchronous mode. A value of FF_H in this field indicates that no limit exists. The hard disk drive supports a maximum REQ/ACK offset of eight.

The two arguments for this message, **xfer period** and **REQ/ACK offset**, reflect the requirements of the device originating the message. These requirements may be incompatible with those of the recipient of the message. Thus, a request for synchronous operation is always made by exchanging this message. The originator specifies its requirements in the message it sends. The recipient compares those requirements with its own capabilities, and sends synchronous data transfer request back to the originator. At this point, one of three conditions exists:

1. The recipient can comply with the synchronous transfer request. In this case, the recipient will set its response message arguments as follows:

   xfer period: equal to or greater than that requested by the originator. xfer period for synchronous transfers between the two devices will be set to the value set in the recipient's response message.

   REQ/ACK offset: less than or equal to that requested by the originator. REQ/ACK offset for synchronous transfers between the two devices will be set to the value set in the recipient's response message.

2. The recipient must negotiate with the originator before the synchronous transfer can be accomplished. This occurs when the originator has set REQ/ACK offset to FF_H and the recipient cannot support the unlimited offset. In this case, the recipient will set its response message arguments as follows.

xfer period: equal to or greater than that requested by the originator. xfer period for synchronous transfers between the two devices will be set to the value set in the recipient's response message.

REQ/ACK offset: set to 00_H. This implies the transfer must be asynchronous, but the originator may respond with another synchronous data transfer request with REQ/ACK offset set to a value less than FF_H.

3. The recipient cannot comply with the request for any reason other than REQ/ACK offset set to FF_H by the originator. In this case, the recipient will send either message reject, or synchronous data transfer request with its arguments set as follows:

   xfer period: n/a

   REQ/ACK offset: set to 00_H. This indicates that transfers must take place in asynchronous mode.

Once the parameters of synchronous transfers between an originator and a recipient have been established, they must remain in effect until a reset occurs, or one of the devices modifies the parameters by sending synchronous data transfer request.

transfer period is set during a Message phase.

The originating device sends an SDTR message, including values for the TRANSFER PERIOD and REQ/ACK offset that will enable it to receive data successfully. The responding device returns the same values in its SDTR message if it can receive data successfully using:

- The same values, or
- A shorter TRANSFER PERIOD, or
- A larger REQ/ACK offset, or
- Both a shorter TRANSFER PERIOD and a larger REQ/ACK offset

If the responding device requires a longer TRANSFER PERIOD, a smaller REQ/ACK offset, or both, it substitutes the appropriate values in its SDTR message. The drive uses the nearest TRANSFER PERIOD value that is greater than or equal to the value received from the initiator. When transmitting data, each device must respect the limits set by the other's SDTR message, but can transfer data with longer TRANSFER PERIODS, smaller REQ/ACK offsets, or both. The successful completion of an exchange of SDTR messages between two SCSI devices implies agreement as shown in Table 6-7. This agreement applies only to the DATA phase.

Table 6-7 *Synchronous Data Transfer Request Responses*

Responses	Implied Agreement
Nonzero REQ/ACK offset	Each device transmits data using a TRANSFER PERIOD that is equal to or greater than the value received in the other's SDTR message, and a REQ/ACK offset equal to or less than the value received in the other's SDTR message
REQ/ACK offset equal to zero	Asynchronous transfer
MESSAGE REJECT	Asynchronous Transfer

If the initiator recognizes a need for negotiation, it asserts ATN during selection, then sends an SDTR message to indicate its REQ/ACK offset and minimum TRANSFER PERIOD. The REQ/ACK offset prevents buffer overflows in the initiator, while the minimum TRANSFER PERIOD satisfies the initiator's data-handling requirements. After successfully completing the MESSAGE OUT phase, the drive responds with the appropriate SDTR message. If an abnormal condition prevents the drive from returning an appropriate response, data transfer between the two devices is asynchronous.

If the drive recognizes a need for negotiation, because the SEND SYNCHRONOUS MESSAGE (SSM) bit is one on mode page 37H, it sends an SDTR message to the initiator, with a REQ/ACK offset value of eight and a TRANSFER PERIOD value of 25. The REQ/ACK offset prevents buffer and offset-counter overflows, while the minimum TRANSFER PERIOD satisfies the drive's data-handling requirements. Prior to releasing ACK on receiving the last byte of an SDTR message from the drive, the initiator asserts ATN and responds with its SDTR message. If an abnormal condition prevents the initiator from returning an appropriate response, data transfer between the two devices is asynchronous.

Implied agreement does not exist until the drive leaves the MESSAGE OUT phase and enters a phase other than the BUS FREE phase—that is, only if the drive detects no parity error. Implied agreement remains in effect until:

- The drive receives a BUS DEVICE RESET message
- A hard RESET condition occurs
- One of the two devices modifies the agreement

The drive enters its default data-transfer mode—asynchronous mode—at power on, or following a BUS DEVICE RESET message or a hard RESET condition. To ensure that devices on a SCSI bus agree on a data-transfer mode for successful data transfer, a SYNCHRONOUS DATA TRANSFER REQUEST message exchange can occur only following a SELECTION phase that includes the SCSI IDs for both the initiator and the drive.

6.4.6 Message Error Handling

In the event that the target detects a parity error in the received message bytes, it can request that the initiator resend them by asserting REQ after ATN has been released and prior to initiating any other bus phases or phases. When the initiator detects REQ active, it will resend all of the message bytes sent during the current Message Out phase. If more than one byte is being resent, the initiator asserts ATN prior to asserting the first ACK, so that the target will remain in the current Message Out phase until resending is complete. When resending is complete, the initiator releases ATN. The target may also indicate a parity error on some messages by initiating the Bus Free Phase.

The target indicates the no-retry condition by initiating another Information Transfer (IT) phase, other than another Message Out phase, and transferring at least one byte. The target also can indicate no retry by initiating the Bus Free phase on an Abort or Bus Device Reset Message.

6.5 COMMAND IMPLEMENTATION

This section defines the SCSI command implementation supported by the Maverick 270/540S intelligent hard disk drive. Each command consists of a group of command bytes and, in some cases, a group of associated data bytes.

On command completion, the drive returns a status byte to the initiator. Because most error and exception conditions cannot be described adequately in a single status byte, the CHECK CONDITION status code indicates that additional information is available. The initiator can issue a REQUEST SENSE (03_H) command to retrieve that information.

The following sections describe each command—including the sequence of bytes an initiator transfers to the drive during the Command phase to send a request to the drive. Where applicable, these sections also describe the format of data bytes that can be transferred to or from the initiator following a Command phase.

6.5.1 Command Descriptor Block (CDB)

SCSI commands are issued from an initiator by transferring a Command Descriptor Block (CDB) to the target device. For some commands, a parameter list sent during a Data Out phase accompanies the request. A CDB contains an opcode (opcode), a logical unit number (LUN), a set of command parameters, and a control byte.

Opcode (byte 0): The first byte of a CDB is the opcode. In a Command Descriptor Block (CDB), the opcode consists of a three-bit group code field and a five-bit command code field. The group code field contains the code for one of eight group commands. The command code field contains the code for one of the 32 command codes in each group. The two fields together provide 256 possible code combinations. The Maverick 270/540S hard disk drive supports the following command groups: Group 0, which are six-byte commands, and Groups 1 and 2, which are ten-byte commands. The operation code (opcode) for each command indicates its format—six-byte commands have opcodes between 00_H and $1F_H$; ten-byte commands, from 20_H to $5F_H$. The opcodes for the commands supported by the are given in Table 6-8.

Table 6-8 *SCSI Commands*

OPCODE	COMMAND
00_H	TEST UNIT READY
01_H	REZERO UNIT
03_H	REQUEST SENSE
04_H	FORMAT UNIT
07_H	REASSIGN BLOCKS
08_H	READ
$0A_H$	WRITE
$0B_H$	SEEK
12_H	INQUIRY
15_H	MODE SELECT
16_H	RESERVE
17_H	RELEASE
$1A_H$	MODE SENSE

OPCODE	COMMAND
$1B_H$	START/STOP UNIT
$1D_H$	SEND DIAGNOTIC
25_H	READ CAPACITY
28_H	READ EXTENDED
$2A_H$	WRITE EXTENDED
$2B_H$	SEEK EXTENDED
$2E_H$	WRITE AND VERFY
$2F_H$	VERIFY
37_H	READ DEFECT DATA
$3B_H$	WRITE BUFFER
$3C_H$	READ BUFFER
$3E_H$	READ LONG
$3F_H$	WRITE LONG

LUN: the Logical Unit Number for the Maverick 270/540S is 00_H. If any command other than INQUIRY or REQUEST SENSE specifies LUN not equal to zero, the drive rejects the command with the CHECK CONDITION status and sets sense key to ILLEGAL REQUEST—Invalid LUN.

LBA: the Logical Block Address specifies the first or starting block of an operation. Group 0 CDBs contain 21-bit LBAs. Group 1 and 2 CDBs contain 32-bit LBAs. The LBA begins with block zero and occupies contiguous blocks up to the last logical block on the drive, as determined by the READ CAPACITY command.

A logical block must be 512 bytes in length, so the MODE SELECT command configures each drive with a logical block length of 512 bytes.

Transfer Length: The transfer length specifies the amount of data to be transferred to or from an initiator. It can be specified as:

- The number of blocks to be transferred
- The number of bytes to be transferred, or parameter list length
- The maximum number of bytes that can be transferred, or allocation length

Number of Blocks: Typically, transfer length specifies the Number of Blocks (NB) to be transferred.

- For six-byte commands, up to 256 blocks can be transferred. A value between 1 and 255 (01_H–FF_H) indicates the number of blocks to be transferred. The value 00_H indicates that 256 blocks are to be transferred.
- For ten-byte commands, up to 65,535 blocks can be transferred. A value between 1 and 65,535 (0001_H–$FFFF_H$) indicates the number of blocks to be transferred. The value 0 indicates that no blocks are to be transferred.

Parameter List Length: The parameter list length specifies the number of bytes to be sent during a Data Out phase. This CDB field allows an initiator to send a list of parameters to the drive.

Allocation Length: The allocation length specifies the number of bytes allocated by an initiator for returned data. With an allocation length of zero, the initiator returns no sense data. The drive terminates a Data In phase when it has transferred to the initiator either the number of bytes specified by the allocation length or all available sense data, whichever is less.

Vendor Unique: Vendor-unique bits allow command modifications or special conditions to be selected. Vendor-unique bits are set to zero.

Reserved: Reserved bits are set to zero by the initiator.

Flag: If the link bit is zero, the flag bit must also be set to zero. If the link bit is one and the flag bit is zero, the target sends a linked command complete message when a command terminates. If both the link bit and the flag bit are one, the target sends the linked command complete (with flag) message.

Link: If the link bit is one when a command terminates, the drive sends the message defined by the flag bit and goes immediately to the Command phase for the next command. When the drive successfully completes the execution of a linked command, it issues an INTERMEDIATE status. If the link bit is zero, the drive sends the command complete message.

Reladr: The Relative Address bit is used only in ten-byte commands and is not supported by the Maverick 270/540S.

Note: If the reserved bits of the IDENTIFY message, or vendor-unique, reserved, or reladr are not set to zero, the drive returns a CHECK CONDITION status with sense key set to ILLEGAL REQUEST and sense code set to Illegal Use of Bit or Byte in the CDB or IDENTIFY message.

6.6 COMPLETION STATUS BYTE

At the termination of each command during the STATUS phase, the drive sends a status byte to the initiator (as specified in Table 6-9 and Table 6-10) unless one of the following events terminates the command:

- An ABORT message
- A BUS DEVICE RESET message
- A hard RESET condition
- An unexpected disconnect

Table 6-11 defines the status codes implemented in the drive.

Table 6-9 *Drive Status Byte Bit Layout*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED		STATUS BYTE CODE					RSVD

Table 6-10 *Drive Status Byte Code Bit Values*

STATUS BYTE BIT								STATUS REPRESENTED
7	6	5	4	3	2	1	0	
R	R	0	0	0	0	0	R	Good
R	R	0	0	0	0	1	R	Check Condition
R	R	0	0	1	0	0	R	Busy
R	R	0	1	0	0	0	R	Intermediate/Good
R	R	0	1	1	0	0	R	Reservation Conflict

Note: R = Reserved bit (0)

Table 6-11 *Drive Status Codes*

STATUS	CODE	DESCRIPTION
GOOD	00_H	The drive successfully completed the command.
CHECK CONDITION	02_H	An error, exception, or abnormal condition occurred during the drive's execution of a command, which may or may not have completed successfully. To determine the nature of the abnormal condition, the initiator must issue a REQUEST SENSE command
BUSY	08_H	The drive is busy and unable to process a command sent by an initiator. Generally, when the status for the drive is BUSY, the drive has disconnected to execute a previous command and cannot execute an additional command at that time. The initiator must reissue the command
INTERMEDIATE/GOOD	10_H	The drive successfully completed a linked command, without encountering an error condition. Typically, the drive issues this status code for each command in a series of linked commands except the last command, for which the link bit is zero. When an error or abnormal condition causes the drive to issue a CHECK CONDITION or RESERVATION CONFLICT status code instead, execution of the series of linked commands terminates.
RESERVATION CONFLICT	18_H	The initiator attempted to access a logical unit reserved by another SCSI device that has a conflicting reservation type. The drive rejects the command. The initiator must reissue the command. (See Section 6.9.11, "RESERVE Command, Opcode = 16H" on page 6-72.)

6.7 READING AND WRITING

6.7.1 Logical Block

The Maverick 270/540S hard disk drive stores and retrieves information in a series of logical blocks that consist of eight-bit bytes. The drive has 512 bytes per logical block. Table 6-12 shows the address ranges for each logical block.

Each logical block has a unique address—its Logical Block Address. The address of the first logical block on the drive is zero. The address of the last logical block on the drive depends on the drive's capacity.

Table 6-12 *Logical Block Address Ranges*

	BLOCK SIZE	VALID LOGICAL BLOCK
Maverick 270S	512 (200_H)	0–528,879 (811EFh)
Maverick 540S	512 (200_H)	0–1,057,758 (1023DEh)

Note: The value selected for tracks per zone using the MODE SELECT command, page 3, determines the number of available logical blocks.

If an initiator attempts to access a logical block at an address outside the valid address range, the drive will not execute the command. The drive returns a CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST—Invalid Logical Block Address.

6.7.2 Transferring Data

The WRITE command stores logical blocks of data on the disk. The READ command retrieves logical blocks of data from the disk. For both READ and WRITE commands, the initiator specifies the LBA (Logical Block Address) of the first logical block to be transferred and the total number of logical blocks to be transferred. For the number of blocks to be transferred, the drive interprets the value zero differently, depending on the form of READ or WRITE command used:

- When using the Group 1, or extended form of, READ or WRITE command—opcodes 28_H and $2A_H$—a transfer length of zero indicates that the drive will transfer no data.
- A standard READ or WRITE command uses only one byte to specify transfer length. When using the Group 0, or standard form of, READ or WRITE command—opcodes 08_H and $0A_H$—a transfer length of zero indicates that the drive will transfer 256 blocks of data on execution of the command.

The transfer length specified in the command and the logical-block size determine the number of bytes transferred during the DATA phase following the command bytes:

Number of data bytes = (Logical-block size in bytes) x (Transfer length in blocks)

If an error occurs, the drive might transfer fewer bytes.

The data-transfer rate to or from the disks depends on the density of data written on the disk and the speed of disk rotation, which are constant. The drive has eight data zones with eight different data transfer rates.

If the host's data-transfer rate across the SCSI bus is slower than the drive's data-transfer rate, the drive stores data in its buffer on READ transfers. Many other hard disk drives interleave sectors on disks to match their data-transfer rate with the SCSI-bus transfer rate. The 96K buffer in the Maverick 270/540S hard disk drive makes this unnecessary. The internal disk-to-buffer transfer rate is less than 2.5 megabytes per second (21.00 megabits per second) disk-sustained. This buffer also serves as an onboard cache, DisCache. The MODE SELECT command can enable or disable caching, and configure cache parameters. See Section 6.9.10, "MODE SELECT Command, Opcode = 15H" on page 6-54. Refer to Section 5.5.1, "Disk Caching" on page 5-14 for a detailed description of DisCache.

6.8 CONFIGURING THE HARD DISK DRIVE

6.8.1 Operating Modes

Users can change some of the Maverick 270/540S hard disk drive to match their requirements. The MODE SELECT command allows the user to set the operating mode of the drive. The MODE SENSE command allows the user to determine the drive's current operating mode and access information that is useful when setting the operating mode.

For example, the number of retries attempted by the drive when it detects a read error is an operating-mode parameter commonly adjusted by the user. An initiator can change the number of retries the drive attempts on detecting a read error. First, the initiator issues a MODE SENSE command to determine the drive's current configuration. During the DATA phase of the MODE SENSE command, the drive returns the current mode parameters to the initiator, including the number of retries the drive attempts on detecting a read error. Then, to change the number of retries, the initiator sends a MODE SELECT command with the Retry Count parameter set to a new value.

6.8.2 Operating Mode Tables

Parameters in the Current Mode Table set the drive's operating mode. Each time the drive executes a command, it checks the parameters in the Current Mode Table, then responds accordingly. Because the drive frequently accesses its parameters, the Current Mode Table resides in RAM and is lost at power off or whenever the drive is reset. The drive loads the parameters in the Saved Mode Table into the Current Mode Table at power on or reset.

The MODE SELECT command can copy the parameters in the Current Mode Table into the Saved Mode Table. If the MODE SELECT command's Save Parameters bit—bit 0 of byte 1—is set to one, the drive copies the selected pages in the Current Mode Table to the Saved Mode Table. Thus, users can permanently configure the drive to their requirements using the MODE SELECT command's Save Parameters bit. This eliminates the need for the device driver to send a vendor-specific MODE SELECT command following every reset.

The Maverick 270/540S hard disk drive supports two other mode tables—the Changeable Parameters Table and the Default Mode Table. Both of these tables reside in the firmware ROM, can be accessed only by the MODE SENSE command, and cannot be modified by the user.

- The Changeable Parameters Table informs an initiator about the drive operating-mode parameters that can be changed using the MODE SELECT command. If a parameter in the Current Mode Table can be modified, all bits in its field are set to one. A software driver should verify that a device supports the modification of an operating-mode parameter before attempting to change the parameter using the MODE SELECT command.
- The Default Mode Table allows the user to restore the drive to a known or reference condition. If, for any reason, the Current Mode Table and Saved Mode Table are inaccessible, the drive uses the Default Mode Table.

6.9 COMMAND DESCRIPTIONS

The following sections provide detailed descriptions of the SCSI commands, as implemented in the Maverick 270/540S.

6.9.1 TEST UNIT READY Command, Opcode = 00_H

The TEST UNIT READY command, shown in Table 6-13, verifies that the drive is up to speed and ready to accept commands requiring disk access.

Table 6-13 *TEST UNIT READY Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE = 00H							
1	LUN =0			RESERVED =0				
2-4	RESERVED =0							
5	VU =0		RESERVED =0				F	L

When the drive is ready to accept a command that requires disk access, it returns GOOD status in the completion status byte. If the drive is not ready to accept a command or requires initiator action—such as a START UNIT command—to become ready, the drive returns CHECK CONDITION status, with the sense key set to NOT READY. (See Section 6.9.14, "START/STOP UNIT Command, Opcode = 1BH" on page 6-80.)

After a maximum delay of one second following power on, the drive can execute commands that do not require disk access. Unless the drive detects an error on executing such a command, it will not return CHECK CONDITION status with the sense key set to NOT READY, even though it is not ready. The following commands do not require disk access:

- REQUEST SENSE
- INQUIRY
- RESERVE
- SEND DIAGNOSTICS
- RELEASE
- START/STOP UNIT
- READ BUFFER
- WRITE BUFFER

All other commands may require disk access. The drive is not ready until it has:

- Passed its hardware self tests
- Brought its motor up to speed
- Retrieved its defect list and mode parameters from reserved cylinders
- Calibrated its actuator and servo parameters

If the drive is unable to retrieve its defect list and mode parameters, it:

- Rejects all commands requiring disk access
- Sets the sense key to 3, MEDIUM ERROR
- Sets the additional sense code 31_H, MEDIUM FORMAT CORRUPTED

6.9.2 REZERO UNIT Command, Opcode = 01_H

The REZERO UNIT command, shown in Table 6-14, requests that the drive actuator be repositioned to cylinder zero and head zero.

Table 6-14 *REZERO UNIT Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =01H							
1	LUN =0			RESERVED =0				
2-4	RESERVED =0							
5	VU =0		RESERVED =0				F	L

6.9.3 REQUEST SENSE Command, Opcode = 03_H

6.9.3.1 Command Structure

The REQUEST SENSE command, shown in Table 6-15, requests that the drive send sense data to the initiator. The drive transmits the sense data returned in the CHECK CONDITION status for the command last executed for the initiator.

The drive saves the sense data for the initiator's previous command, until either the initiator retrieves the sense data using a REQUEST SENSE command or the drive executes another command issued by the same initiator. When the drive receives a command subsequently issued by that initiator, it clears the sense data. A command issued by one initiator does not affect the sense data saved for another initiator.

The REQUEST SENSE command returns CHECK CONDITION status only to report a fatal error. After such an error occurs, the sense data may be invalid. If a nonfatal error occurs during the execution of a REQUEST SENSE command, the drive returns the sense data with a GOOD status.

Table 6-15 *REQUEST SENSE Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =03H							
1	LUN =0			RESERVED =0				
2-3	RESERVED =0							
4	ALLOCATION LENGTH (BYTES)							
5	VU =0		RESERVED =0				F	L

ALLOCATION LENGTH

The maximum number of bytes allocated by the initiator for returned sense data. Acceptable values are 00–FF_H. A value of zero returns no sense data. The drive terminates a DATA IN phase when it has transferred either the allocated number of bytes or all available sense data to the initiator, whichever is less. An allocation length of 18 (12_H) or greater is recommended for the Maverick 270/540S hard disk drive.

If the allocation length specified prevents the drive from sending all available sense data to the initiator, the initiator cannot recover the remaining data with a subsequent REQUEST SENSE command.

6.9.3.2 Extended Sense Data Format

The Maverick 270/540S hard disk drive uses the extended sense data format, shown in Table 6-16.

Table 6-16 *Extended Sense Data Format*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	VALID	ERROR CLASS AND CODE = 70h						
1	SEGMENT NUMBER = 0							
2	FM =0	EOM =0	ILI =0	R =0	SENSE KEY			
3-6	IMFORMATION BYTE							
7	ADDITIONAL SENSE LENGHT (BYTES) = 0A h							
8-11	RESERVED =0							
12	ADDITIONAL SENSE CODE							
13	ADDITIONAL SENSE CODE QUALIFIER							
14	FRU FAILED =0							
15	FPV	C/D	VEND UNIQUE		BPV	BIT POINTER =0		
16-17	(MSB)	FIELD POINTER						(LSB)

VALID—INFORMATION VALID (Byte 0, Bit 7)

Set to one to indicate that the information bytes contain valid data.

ERROR CLASS AND CODE (Byte 0, Bits 6–0)

Set to 70_H—current errors—to specify the extended sense data format for error codes.

SEGMENT NUMBER (Byte 1)

Set to zero for the Maverick 270/540S hard disk drive.

FM—FILE MARK (Byte 2, Bit 7)

Set to zero by default. Reserved.

EOM—END-OF-MEDIUM (Byte 2, Bit 6)

Set to zero by default. Reserved.

ILI—INCORRECT LENGTH INDICATOR (Byte 2, Bit 5)

Normally, this bit is set to zero for the Maverick 270/540S hard disk drive, indicating that the requested logical block length does not match the logical block length on the disk. However, on execution of a READ BUFFER, READ LONG, or WRITE LONG command, this bit is set to one. For a READ BUFFER command, this bit is set to one to indicate that the requested transfer length is greater than the drive's buffer size. For a READ LONG or WRITE LONG command, this bit is set to one to indicate that the requested transfer length is greater than the logical block length.

R—RESERVED (Byte 2, Bit 4)

This parameter must be set to zero for the Maverick 270/540S hard disk drive.

SENSE KEY (Byte 2, Bits 3–0)

Table 6-17 provides information about the error sense keys.

Table 6-17 *Error Sense Key*

SENSE KEY (HEX)	ADDITIONAL SENSE CODE (HEX)	ADDITIONAL SENSE CODE QUALIFIER (HEX)	DESCRIPTION
0			**NOSENSE.** The drive reports ne sense key information when it executes a command successfully or issues a CHECK CONDITION status for a command, with the filemark, EOM, or ILI bit set to one.
0	00	00	No error detected at drive level
1			**RECOVERED ERROR.** After taking some error-recovery action, the drive successfully completed execution of the previous command. The initiator can determine the nature of an error by issuing a REQUEST SENSE command to examine the additional sense bytes and information field.
1	03	00	Recovered from WUS write fault (bump) Recovered from write fault

Table 6-17 *Error Sense Key (continued)*

SENSE KEY (HEX)	ADDITIONAL SENSE CODE (HEX)	ADDITIONAL SENSE CODE QUALIFIER (HEX)	DESCRIPTION
1	06	01	Recovered from RCL FLT—Coarse Slope PES calibration
	06	02	Recovered from RCL FLT—Fine Slope PES calibration at AEQBH
	06	03	Recovered from RCL FLT—Fine Slope PES calibration at AEQBL
	06	04	Recovered from RCL FLT—Can't Lock to track
	06	05	Recovered from RCL FLT—Can't detect SAM during unpark
	06	06	Recovered from RCL FLT—DAC offset calib failure
	06	06	Recovered from RCL FLT—Can't detect SAM on any head
	06	07	Recovered from RCL FLT—Can't seek to OD near Sys Cyl
	06	08	Recovered from RCL FLT—Can't seek to Fine SLP calib trk
	06	09	Recovered from RCL FLT—Seek failure during Nulli calibration
	06	0A	Recovered from RCL FLT—Seek failure during V_SCALE adaptation
	06	0B	Recovered from RCL FLT—Seek failure during KLOOP calibration
	06	0C	Recovered from RCL FLT—Seek failure RRO calibration
	06	0D	Recovered from RCL FLT—Seek failure to track 0 during rezero
	06	0E	Recovered from RCL FLT—Unable to complete KLOOP calibration
	06	0F	Recovered from RCL FLT—Unable to complete RRO calibration
1	09	00	Recovered from bump timeout
	09	04	Recovered from bad servo sync
	09	05	Recovered from bad servo address mark (SAM)
	09	06	Recovered from bad track number data
	09	07	Recovered from servo defect
	09	08	Recovered from bump
1	10	00	Recovered ID ECC error
1	12	00	Recovered from AM mark not found for ID field
	12	01	Recovered from AM not found for ID with internal continue
1	13	00	Recovered from data field sync timeout
1	14	00	Recovered from no record found
1	15	00	Recovered from seek error
	15	01	Recovered from Gray code !=desired track while ontrack
	15	03	Recovered from multi bad Sync/SAM while Settle/Ontrack
	15	04	Recovered from seek timeout with no servo fault
	15	05	Recovered from multi bad Sync/SAM during seek ISR
1	17	01	Recovered from ECC error in data field
	17	08	Recovered from marker for CRC/Continue
	17	09	Data corrected via ECC-on-the-fly algorithm
1	18	00	Recovered from data error via ECC w/2 = syndromes
	18	01	Recovered from data error via ECC on last retry
1	86	00	Recovered from unexpected sequencer error
	86	01	Recovered from unexpected SEQ error during recovery from SEQ TMO
	86	02	Recovered from read/write ID miscompare
1	95	00	Recovered from sequencer timeout
1	97	00	Recovered from underrun error

Table 6-17 *Error Sense Key (continued)*

SENSE KEY (HEX)	ADDITIONAL SENSE CODE (HEX)	ADDITIONAL SENSE CODE QUALIFIER (HEX)	DESCRIPTION
1	9E 9E	00 01	Recovered from motor unable to get up to speed Recovered from speed out of range
1	AA	00	Recovered from Read Data written on reallocation of uncorrectable data
1	AB	00	Requested format in Read Defect Data not available
2			**NOT READY.** The initiator cannot access the drive. Correcting this condition may require operator intervention.
2	04	00	Drive is up to speed and recalibrating
2	04	01	Drive is spinning up
2	04	02	Drive has not been told to spin up
2	C0	00	Burn-In test in progress
3			**MEDIUM ERROR.** Execution of a command terminated in an unrecovered error condition—probably due to a media defect, an error in the recorded data, or a hardware failure.
3	01	00	No disk index found on current track
3	03 03	00 00	Write fault Write Gate still asserted when wedge detected
3	10	00	ID ECC error
3	11 11	00 01	Uncorrectable data field ECC error Marker for CRC/Continue
3	12 12	00 01	AM mark not found for ID field AM not found for ID field with internal continue
3	13	00	Data field sync timeout
3	14 14	01 01	Bad block mark set for ID (AT) No record found
3	19	00	Bad defect list
3	31 31	00 01	Medium format corrupted FDPE write failed during format unit command
3	32 32	00 01	Defect list is full No more alternate sectors available
3	80	00	Error in writing to system sector
3	81	00	Error in reading from system sector

Table 6-17 *Error Sense Key (continued)*

SENSE KEY (HEX)	ADDITIONAL SENSE CODE (HEX)	ADDITIONAL SENSE CODE QUALIFIER (HEX)	DESCRIPTION
3	82	00	Error during reading of diskware
3	A3	00	Failure reading sector in Reassign Blocks command
3	AA	00	Read Data was written on reallocation of uncorrectable data
4			**HARDWARE ERROR.** The drive detected an unrecoverable hardware failure while executing a command.
4	03	00	WUS write fault (bump)
4	06	01	RCL FLT—Coarse Slope PES Gain calibration
	06	02	RCL FLT—Fine Slope PES Gain calibration at AEQBH
	06	03	RCL FLT—Fine Slope PES Gain calibration at AEQBL
	06	04	RCL FLT—Cannot lock to track
	06	05	RCL FLT—Cannot detect SAM during unparking
	06	06	RCL FLT—Cannot detect reliable SAM on any head
	06	07	RCL FLT—Can't seek to OD to get near Sys Cyl
	06	08	RCL FLT—Can't seek to Fine Slope PES calibrate track
	06	09	RCL FLT—Seek failure during Nulli calibration
	06	0A	RCL FLT—Seek failure during V_SCALE adaptation
	06	0B	RCL FLT—Seek failure during KLOOP calibration
	06	0C	RCL FLT—Seek failure during RRO calibration
	06	0D	RCL FLT—Seek failure to track 0 during rezero
	06	0E	RCL FLT—Unable to complete KLOOP calibration
	06	0F	RCL FLT—Unable to complete RRO calibration
4	09	00	Bump timeout
	09	04	Unrecoverable bad servo sync
	09	05	Unrecoverable bad servo address mark (SAM)
	09	06	Unrecoverable bad track number data
	09	07	Unrecoverable servo defect
	09	08	Unrecoverable bump
4	15	00	Seek error
	15	01	Unrecovered Gray code !=desired track while ontrack
	15	03	Unrecovered multi bad Sync/SAM while Settle/Ontrack
	15	04	Unrecovered seek timeout with no servo fault
	15	05	Unrecovered multi bad Sync/SAM during seek ISR
4	1B	00	Synchronous transfer error
4	40	00	RAM error (most likely found in a diagnostic)

Table 6-17 *Error Sense Key (continued)*

SENSE KEY (HEX)	ADDITIONAL SENSE CODE (HEX)	ADDITIONAL SENSE CODE QUALIFIER (HEX)	DESCRIPTION
4	42	00	Internal ROM checksum error
	42	01	Marker for resident code checksum
	42	02	Marker for resident and overlay are incompatible
	42	03	Marker for ROM and resident are incompatible
	42	04	Marker for ROM and overlay are incompatible
	42	05	Marker for overlay checksum
	42	06	Marker for diskware vector table checksum
	42	07	Invalid diskware version
	42	08	Invalid ROM version
4	43	00	Invalid message
4	49	00	FIFO unload error
4	4A	00	FIFO load error
4	4B	00	FIFO predicted full error
4	84	00	Failure in writing to sequencer format table
4	85	00	Reject of message that should never have been sent
4	86	00	Unexpected sequencer error
	86	01	Unexpected SEQ error during recovery from SEQ TMO
	86	02	Read/write ID miscompare
4	87	00	Logical assertion/firmware consistency check error
4	8A	00	Head read from ID not equal to selected head
4	90	00	Synchronous acknowledge error
4	91	00	Synchronous request error
4	94	00	SIC error
4	95	00	Sequencer timeout
4	97	00	Underrun error
4	9E	00	Motor unable to get up to speed
	9E	01	Unrecoverable speed out of range
4	A1	00	Sequencer rollover register failure
4	AC	00	Airlock stuck open

Table 6-17 *Error Sense Key (continued)*

SENSE KEY (HEX)	ADDITIONAL SENSE CODE (HEX)	ADDITIONAL SENSE CODE QUALIFIER (HEX)	DESCRIPTION
4	FF FF FF	00 01 02	Autowrite command received while host channel disabled ID in format track descriptor list not found Bad descriptor in format track descriptor list
5			**ILLEGAL REQUEST.** The drive detected an illegal parameter in the CDB or in the additional parameters for command data; or received an invalid IDENTIFY message. If the drive detects an invalid parameter in the CDB, it terminates the command without altering the data on disk. If the drive detects an invalid parameter in the additional parameters for command data, it may already have altered the data.
5	1A	00	Parameter overrun
5	20	00	Invalid command
5	21	00	Invalid LBA
5	24	00	Invalid bits set in CDB
5	25	00	Invalid LUN specified
5	26	00	Invalid field in parameters
5	8A	00	Invalid head specified
5	8B	00	Invalid cylinder specified
5	8C	00	Attempt by intruding initiator selected the drive a second time
5	8D	00	Bytes per block/bytes per sector gives a remainder
5	8F	00	Invalid sector specified
5	9B	00	Invalid period or offset in synchronous message
5	9C	00	Active initiator selected us while disconnected
5	AE	00	Bad parameter(s) found in mode pages during init
5	C2	00	Command not implemented

Table 6-17 *Error Sense Key (continued)*

SENSE KEY (HEX)	ADDITIONAL SENSE CODE (HEX)	ADDITIONAL SENSE CODE QUALIFIER (HEX)	DESCRIPTION
6			**UNIT ATTENTION.** Another initiator changed the MODE SELECT parameters or the drive has been reset. Once the drive is ready, the unit attention condition for the initiator persists until the drive receives a command from the initiator—other than INQUIRY. The drive does not execute the command issued by the initiator, but reports a CHECK CONDITION status, enabling the initiator to issue a REQUEST SENSE command to determine that a UNIT ATTENTION condition exists. The execution of an INQUIRY command does not affect a unit attention condition. If the drive receives a REQUEST SENSE command before reporting CHECK CONDITION status for the unit attention condition, it reports and clears UNIT ATTENTION status, unless sense data is available for the previous command.
6	29	00	Reset occurred
6	2A	00	Mode select parameters were changed
6	8E	00	Unexpected SIC interrupt occurred
6	9A	00	A target attempted to reselect drive
B			**ABORTED COMMAND.** The drive aborted execution of a command. The initiator may be able to recover by sending the command again.
B	00	00	Response for an ABORT message
B	29	00	Miscellaneous SIC error
B	45	00	Initiator did not reselect
B	47	00	Parity error on SCSI bus
B	48	00	Initiator detected error
E			**MISCOMPARE.** Indicates that the same data bit did not match the data read from the medium.
E	1D	00	Read buffer miscompare

INFORMATION (Bytes 3–6)

If the VALID bit is set to one, the information field is valid and contains the logical block address associated with the sense key. If the drive issues a MEDIUM ERROR sense key, it notifies the initiator to use the REASSIGN BLOCKS command to map out the defective sector.

ADDITIONAL SENSE LENGTH (Byte 7)

This parameter indicates the number of additional sense bytes to follow. If the allocation length in the Command Descriptor Block (CDB) prevents the transfer of all additional sense bytes, this value is not adjusted to reflect their truncation. The drive uses a value of $0A_H$ in byte 7, even though additional sense bytes are available.

ADDITIONAL SENSE CODE (Byte 12)

See Table 6-17.

FRU—FIELD REPLACEABLE UNIT FAILED (Byte 14)

A value of zero indicates that there is no field-replaceable unit failure to be reported.

FPV—FIELD POINTER VALID (Byte 15, Bit 7)

When FPV is set to zero, the C/D bit and BPV bit, and bytes 16 and 17 are invalid. When FPV is set to one, the field pointer bytes 16 and 17, the C/D bit, and the BPV bit are significant.

C/D—CONTROL/DATA BIT (Byte 15, Bit 6)

When C/D is set to one, the value reported in the field pointer is the CDB byte number for which the drive issued an ILLEGAL REQUEST sense key. When C/D is set to zero, the value reported in the field pointer is the byte number of the DATA phase for which the drive issued an ILLEGAL REQUEST sense key.

VENDOR UNIQUE (Byte 15, Bits 5–4)

Bits 5 and 4 of byte 15 are vendor unique.

BPV—BIT POINTER VALID (Byte 15, Bit 3)

When BPV is set to zero, the bit-pointer field is invalid. When BPV is set to one, the bit-pointer field (bits 0–2) is valid.

BIT POINTER (Byte 15, Bits 2–0)

When BPV is set to one, this field points to the bit within the byte that is in error.

FIELD POINTER (Bytes 16–17)

These bytes represent an index to the CDB fields or the DATA phase fields for which the drive issued an ILLEGAL REQUEST sense key. In the case of multiple-byte fields, the field-pointer bytes index the most significant byte of the field for which the drive issued an ILLEGAL REQUEST sense key. These bytes are valid only if the Field Pointer Valid (FPV) bit is set to one.

6.9.4 FORMAT UNIT Command, Opcode = 04_H

The FORMAT UNIT command, shown in Table 6-18, assigns logical blocks to physical sectors for optimum sequential access—within the limitation of available spare blocks—skipping known defective areas. The data pattern parameter causes the drive to write a repetitive, one-byte data pattern in every block.

The Maverick 270/540S hard disk drive is formatted at the factory, with data undefined.

Note: Issuing the FORMAT UNIT command typically causes loss of data, even when the data pattern parameter is not used. All data should be backed up prior to formatting.

6.9.4.1 Command Structure

Table 6-18 *FORMAT UNIT Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =04 H							
1	LUN =0			FMTDAT	CMPLST	DEFECT LIST FORMAT		
2	DATA PATTERN							
3-4	(MSB)			INTERLEAVE =XX				(LSB)
5	VU =0			RESERVED =0			F	L

FMTDAT—FORMAT DATA (Byte 1, Bit 4)

When FMTDAT is set to one, defect block data is transferred during the DATA OUT phase. The defect list included with this data specifies the defects to be entered in the defect map. This defect list refers to the current block length and current logical block addresses, not physical addresses. Table 6-20 shows the format of the defect list. When FMTDAT is set to zero, the DATA OUT phase will not occur—that is, the initiator will supply neither a defect list header, nor defect data.

CMPLST—COMPLETE LIST (Byte 1, Bit 3)

When CMPLST is set to one, the drive supplies a complete list of known initiator-specified or field-replacement defects. The drive erases any previous initiator-specified defect map or defect data, purging any previous initiator-specified defect list, and builds a new defect list that includes the initiator-specified list and the factory defects. When CMPLST is set to zero, the drive supplies defective block data to supplement existing defect data, using the current format.

DEFECT LIST FORMAT (Byte 1, Bits 2–0)

This parameter specifies additional information related to the defect list. See Table 6-20 for more information.

DATA PATTERN (Byte 2)

This parameter specifies the repetitive user-data pattern to be written into each sector during execution of the FORMAT UNIT command—if enabled by the Fill Data Pattern bit in page 39 of the MODE SELECT command. (See section 6.9.10 for information about the MODE SELECT command.)

INTERLEAVE FIELD (Bytes 3–4)

This parameter specifies the interleave factor for the drive. The drive will ignore any value specified for this field by the initiator. The Maverick 270/540S hard disk drive uses a 1:1 interleave.

Table 6-19 *FORMAT UNIT Command Variations*

BIT REFERENCE (CDB, BYTE 1)						DEFECT LIST BYTE 1)		DEFECT LIST LENGTH	OPTION	DESCRIPTION
Bit	4	3	2	1	0	7	6			
1		1	0	X	X	1	1	0	1	
1		1	0	X	X	0	0	0	2	
1		0	0	X	X	1	1	0	3	
0		X	X	X	X				4	No defect list (Because the DATA OUT phase does not occur, there is no defect list header.)
1		1	0	X	X	1	1		5	
1		1	0	X	X	0	0		6	
1		0	0	X	X	1	1		7	
1		0	0	X	X	0	0		8	
1		0	1	0	0,1					Bytes offset from index format
1		0	1	0	1					Physical sector format
1		X	1	1	0					Vendor-unique (not used)
1		X	1	1	1					Reserved

Note: Bit Reference (CDB, Byte 1): Bit 4 = FMTDAT; bit 3 = CMPLST; and bits 2, 1, and 0 = Defect List Format
Defect List (Byte 1): Bit 7 = FOV; bit 6 = DPRV

The value of bits indicated by an X can be either one or zero—the drive ignores these values.

The defect list shown in Table 6-20 contains a four-byte header, which can be followed by defect descriptors. The defect list length is equal to four times the number of defect descriptors and specifies the total length, in bytes, of the defect descriptors that follow. Each defect descriptor in the logical block format specifies a four-byte logical block address for the defect. The defect list length can be zero.

Refer to Table 6-21 and Table 6-22 for additional descriptor formats—Physical Sector Format and Bytes Offset from Index Format. For these formats, the defect list length is equal to eight times the number of defect descriptors and specifies the total length, in bytes, of the defect descriptors that follow. Each defect descriptor specifies an 8-byte physical address for the defect. When using the physical descriptors, all eight of the formatting parameters described for the logical descriptors apply.

Table 6-20. *Defect List—Logical Block Format*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
DEFECT LIST HEADER								
0	RESERVED							
1	FOV	DPRY	DCRT	STPF	RESERVED			
2-3	(MSB)	DEFECT LENGTH						(LSB)
DEFECT DESCRIPTOR (S):								
0-4	(MSB)	DEFECT LOGICAL BLOCK ADDRESS						(LSB)

FOV—FORMAT OPTIONS VALID (Byte 1, Bit 7)

When FOV is set to one, the initiator authorizes the setting of bits 4 through 6—STPF, DCRT, and DPRY. With FOV set to zero, the initiator requests the implementation of the drive's default format scheme for the functions defined by bits 4 through 6.

DPRY—DISABLE PRIMARY (Byte 2, Bit)

When DPRY is set to one, the drive will exclude the factory defect map—primary list or P-list—from its list of defects to be managed on formatting. When DPRY is set to zero, the drive will include the factory defect map in its list of defects to be managed on formatting. By default, the Maverick 270/540S hard disk drive includes the factory defect map.

DCRT—DISABLE CERTIFICATION (Byte 1, Bit 5)

Not supported by the Maverick 270/540S hard disk drive. This parameter must be set to zero by the initiator.

STPF—STOP FORMAT (Byte 1, Bit 4)

Not supported by the Maverick 270/540S hard disk drive. This parameter must be set to zero by the initiator.

6.9.4.2 Application of FORMAT UNIT Command

The Maverick 270/540S hard disk drive offers eight formatting options:

- Format with absolutely no defects.
- Format with original factory defects only.
- Format with grown, or field-found, defects only—disregarding factory defects.
- Format with existing defects—factory and grown defects.
- Format with provided defects only—disregarding factory and existing grown defects.
- Format with provided defects and factory defects—disregarding existing grown defects.
- Format with provided defects and existing grown defects—disregarding factory defects.
- Format with provided defects and existing defects.

Note: With any of these options, the list of grown, or factory-found, defects is always updated to reflect the drive's current condition. The original factory defects are always preserved.

The following tables show the settings of the FORMAT UNIT CDB and defect list header required to implement each of these formatting options.

Table 6-21 *Format with Absolutely No Defects*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =04 H							
1	LUN =0			1	1	0		
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE =XX						(LSB)
5	VU =0		RESERVED =0				F	L

Defect List Header

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0							
1	1	1	0	0	RESERVED =0			
2-3	(MSB)	DEFECT LENGTH =0						(LSB)

Table 6-22 *Format with Original Factory Defects Only*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =04 H							
1	LUN =0			1	1	0		
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE =XX						(LSB)
5	VU =0		RESERVED =0				F	L

Defect List Header

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0							
1	0	0	0	0	RESERVED =0			
2-3	(MSB)	DEFECT LENGTH =0						(LSB)

Note: The list of field-found defects is erased.

Table 6-23 *Format with Grown Defects Only—Disregarding Factory Defects*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =04 H							
1	LUN =0			1	0	0		
2	DATA PATTERN							
3 - 4	(MSB)	INTERLEAVE =XX						(LSB)
5	VU =0		RESERVED =0				F	L

Defect List Header

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0							
1	1	1	0	0	RESERVED =0			
2-3	(MSB)	DEFECT LENGTH =0						(LSB)

Note: The list of field-found defects is erased.

Table 6-24 *Format with Existing Defects—Factory and Grown Defects*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =04 H							
1	LUN =0			0	0	0		
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE =XX						(LSB)
5	VU =0		RESERVED =0				F	L

Note: The DATA OUT phase is *not* used.

No defect list header or defect descriptor blocks follow the command.

Table 6-25 *Format with Provided Defects Only—Disregarding Factory and Existing Grown Defects*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =04 H							
1	LUN =0			1	1	0		
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE =XX						(LSB)
5	VU =0		RESERVED =0				F	L

Defect List Header

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0							
1	1	1	0	0	RESERVED =0			
2-3	(MSB)	DEFECT LENGTH						(LSB)

Note: The defect descriptor block follows the defect list header.

Table 6-26 *Format with Provided Defects and Factory Defects (Disregarding Existing Grown Defects)*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =04 H							
1	LUN =0			1	1	0		
2	DATA PATTERN							
3 - 4	(MSB)	INTERLEAVE =XX						(LSB)
5	VU =0		RESERVED =0				F	L

Defect List Header

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0							
1	0	0	0	0	RESERVED =0			
2-3	(MSB)	DEFECT LENGTH						(LSB)

Note: The defect descriptor blocks follow the defect list header. The provided defects replace the existing field-found defects.

Table 6-27 *Format with Provided Defects and Existing Grown Defects (Disregarding Factory Defects)*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =04 H							
1	LUN =0			1	0			
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE =XX						(LSB)
5	VU =0		RESERVED =0				F	L

Defect List Header

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0							
1	1	1	0	0	RESERVED =0			
2-3	(MSB)	DEFECT LENGTH						(LSB)

Note: The defect descriptor blocks follow the defect list header. The existing list of grown defects will be combined with the provided defects to become the updated grown defects.

Table 6-28 *Format with Provided Defects and Existing Defects*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =04 H							
1	LUN =0			1	0	0		
2	DATA PATTERN							
3-4	(MSB)	INTERLEAVE =XX						(LSB)
5	VU =0		RESERVED =0				F	L

Defect List Header

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0							
1	0	0	0	0	RESERVED =0			
2-3	(MSB)	DEFECT LENGTH =0						(LSB)

Note: The defect descriptor blocks follow the defect list header. The existing list of field-found defects will be combined with the provided defects.

6.9.5 REASSIGN BLOCKS Command, Opcode = 07_H

The REASSIGN BLOCKS command, shown in Table 6-29, requests the drive to reassign defective logical blocks to sectors reserved as spares.

The initiator transfers a defect list during the DATA OUT phase that follows the command bytes. This defect list contains the logical block addresses to be reassigned by the drive. The drive will reassign the sectors that correspond to each logical block address in the list. If recoverable under the MODE SELECT error-recovery parameters, the data in the logical blocks specified by the defect list will be preserved. The data in all other logical blocks will be preserved. If specified, a logical block that has previously been reassigned can be reassigned again. The address of the defective logical block will be added to the drive's list of field-found defects.

Note: The REASSIGN BLOCKS command adds to the grown and working defects lists, but it will not alter the contents of the Primary Defects List, which indicates the location of permanent disk defects.

Because the drive must access spare locations, block reassignment may degrade its performance. The FORMAT UNIT command performs in-line sparing, which increases performance. However, the FORMAT UNIT command will erase the disk contents.

Table 6-29 *REASSIGN BLOCKS Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =07 H							
1	LUN =0			RESERVED =0				
2-4	RESERVED =0							
5	VU =0		RESERVED =0				F	L

The REASSIGN BLOCKS defect list, shown in Table 6-30, contains a four-byte header, followed by one or more defect descriptors. The length of each defect descriptor is four bytes. The defect list length field specifies the total length of the defect descriptors, in bytes. Therefore, the defect list length is equal to four times the number of defect descriptors.

Table 6-30 *REASSIGN BLOCKS Command Defect List*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0-1	RESERVED =0							
2-3	(MSB)	DEFECT LIST LENGTH						(LSB)

Defect List Header

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0-3	(MSB)							(LSB)

The defect descriptor specifies the four-byte address of the logical block that contains the defect.

If the drive has insufficient capacity to reassign all defective logical blocks, the command terminates with a CHECK CONDITION status and the drive sets the sense key to HARDWARE ERROR, with the additional sense code Defect List is Full. The the drive returns the logical block address of the first logical block not reassigned in the information bytes of the sense data.

6.9.6 READ Command, Opcode = 08_H

The READ command, shown in Table 6-31, requests the drive to transfer data to the initiator.

6.9.6.1 Command Structure

Table 6-31 *READ Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =08 H							
1	LUN =0			(MSB) LOGICAL BLOCK ADDRESS				
2-3	LOGICAL BLOCK ADDRESS (LSB)							
4	TRANSFER LENGTH							
5	VU =0		RESERVED =0				F	L

LOGICAL BLOCK ADDRESS (Byte 1, Bits 4–0; Bytes 2–3)

Specifies the logical block at which the READ operation will begin.

TRANSFER LENGTH (Byte 4)

Specifies the number of contiguous logical blocks of data to be transferred. A transfer length of zero indicates that 256 blocks will be transferred. Acceptable values are 00–FF_H.

6.9.6.2 READ Command Operation

The 96-kilobyte buffer uses variable size segments. These segments are used as a read/write buffer (cache). The read operation is based on a least recently used algorithm. When the host requests read data, the drive transfers the data to the buffer. Two independent operations occur on reads:

- The drive transfers data from the disk to the buffer, until the amount of data transferred from the disk matches the transfer length or the buffer is full.
- The drive begins transferring data from the buffer to the SCSI bus when the percentage of the buffer that is full matches the Buffer Full Ratio specified in MODE SELECT page 2, byte 2, and continues until it has transferred all data requested.

If the buffer becomes full because data transfer from the buffer to the SCSI bus is slower than from the disk to the buffer, the drive temporarily halts data transfer to the buffer until space becomes available in the buffer. The drive then continues transferring data from the disk to the buffer until the amount of data transferred equals the transfer length.

If the initiator has indicated that it supports disconnection by sending an IDENTIFY message with the Disconnect Supported bit (bit 6) set to one, the drive will disconnect whenever the buffer is empty and the amount of data transferred does not equal the transfer length. The Buffer Full Ratio controls reconnection. When the percentage of the buffer that is full matches the Buffer Full Ratio, the drive will initiate a reconnection. Once the drive reconnects to the bus, it resumes data transfer from the buffer to the SCSI bus. If reconnection to the bus is delayed—for example, because the bus is busy—the drive continues transferring data from the disk to the buffer until the buffer is full or the drive has transferred the number of blocks requested.

6.9.7 WRITE Command, Opcode = $0A_H$

The WRITE command, as shown in Table 6-32, requests that the data transferred to the target be written in the area specified by the logical block address.

6.9.7.1 Command Structure

Table 6-32 *WRITE Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =0A H							
1	LUN =0			(MSB) LOGICAL BLOCK ADDRESS				
2-3	LOGICAL BLOCK ADDRESS (LSB)							
4	TRANSFER LENGTH							
5	VU =0		RESERVED =0				F	L

LOGICAL BLOCK ADDRESS (Byte 1, Bits 4–0; Bytes 2–3)

Specifies the logical block at which the WRITE operation will begin.

TRANSFER LENGTH (Byte 4)

Specifies the number of contiguous logical blocks to be transferred. A transfer length of zero indicates that 256 blocks will be transferred. Acceptable values are $00–FF_H$.

6.9.7.2 WRITE Command Operation

The 96K read/write buffer is used for all write operations. The cache segmentation is dynamically allocated based on the current command requirements. When the host sends data to the drive, the drive stores the data in the buffer. Then, when the head is over the correct sector on the track, the drive writes the data from the buffer to the disk. Two independent operations occur on writes:

- The drive transfers data from the SCSI bus to the buffer, until the amount of data transferred to the buffer matches the transfer length or the buffer is full.
- The drive transfers data from the buffer to the disk as long as the buffer contains data and the head is over the correct sector.

If the buffer becomes empty because data transfer from the SCSI bus to the buffer is slower than from the buffer to the disk, the drive temporarily halts data transfer to the disk until new data is written in the buffer. If the buffer becomes full because data transfer from the buffer to the disk is slower than from the SCSI bus to the buffer, the drive temporarily halts data transfer from the SCSI bus to the buffer until space becomes available in the buffer.

If the initiator has indicated that it supports disconnection by sending an IDENTIFY message with the Disconnect Supported bit (bit 6) set to one, the drive will disconnect whenever the buffer is full or the amount of data transferred to the disk matches the transfer length. The Buffer Empty Ratio controls reconnection. When the percentage of the buffer that is empty matches the Buffer Empty Ratio, the drive will initiate a reconnection. Once the drive reconnects to the bus, it resumes data transfer from the SCSI bus to the buffer.

6.9.7.3 Write Cache

When a write command is executed with write caching enabled, the drive stores the data to be written in a DRAM cache buffer and immediately sends a COMMAND COMPLETE message to the host before the data is actually written to the disk. The host is then free to move on to other tasks, such as preparing data for the next data transfer, without having to wait for the drive to seek to the appropriate track or rotate to the specified sector.

While the host is preparing data for the next transfer, the drive immediately writes the cached data to the disk, usually completing the operation in less than 27 ms after issuing COMMAND COMPLETE. With WriteCache, a single-block random write, for example, requires only about 3 ms of host time. Without WriteCache, the same operation would occupy the host for about 30 ms.

WriteCache allows data to be transferred in a continuous flow to the drive rather than as individual blocks of data separated by disk access delays. This is achieved by taking advantage of the ability to write blocks of data sequentially on a disk that is formatted with a 1:1 interleave. This means that as the last byte of data is transferred out of the write cache, and the head passes over the next sector of the disk, the first byte of the next block of data is ready to be transferred; thus there is no interruption or delay in the data transfer process.

The WriteCache algorithm writes data to the cache buffer while simultaneously transferring data to the disk that was previously written to the cache.

6.9.8 SEEK Command, Opcode = $0B_H$

The SEEK command, shown in Table 6-33, requests the drive to seek to the specified logical block address.

Table 6-33 *SEEK Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =0BH							
1	LUN =0			(MSB) LOGICAL BLOCK ADDRESS				
2-3	LOGICAL BLOCK ADDRESS (LSB)							
4	RESERVED = 0							
5	VU =0		RESERVED =0				F	L

LOGICAL BLOCK ADDRESS (Byte 1, Bits 5–0; Bits 2–3)

Specifies the logical block to which the drive will reposition the actuator after completing the SEEK operation.

6.9.9 INQUIRY Command, Opcode = 12_H

The INQUIRY command, shown in Table 6-34, requests the drive to send its identification to the initiator.

6.9.9.1 Command Structure

Table 6-34 *INQUIRY Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =12 H							
1	LUN =0			RESERVED =0				
2-3	RESERVED =0							
4	ALLOCATION LENGTH (BYTES)							
5	VU =0		RESERVED =0				F	L

ALLOCATION LENGTH

Specifies the maximum number of bytes allocated by the initiator for returned INQUIRY data. A value of 00 indicates that the drive will return no data. Acceptable values are 00–FF_H. The drive will terminate the DATA IN phase when the bytes corresponding to the allocation length have been transferred or when all available INQUIRY data have been transferred to the initiator, whichever is less. An allocation length of 130 (82_H) is recommended.

6.9.9.2 INQUIRY DATA Format

Table 6-35 *INQUIRY DATA Format*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	PERIPHERAL QUALIFIER =0			PERIPHERAL DEVICE TYPE =0				
1	RMB =0		DEVICE TYPE MODIFIER =0					
2	RESERVED =0		ECMA VERSION =0			ANSI VERSION =2		
3	RESERVED =0				RESPONSE DATA FORMAT =1			
4	ADDITIONAL LIST LENGTH							

PERIPHERAL QUALIFIERS (Byte 0, Bits 5–7)

Not Supported. This value must be zero.

PERIPHERAL DEVICE TYPE (Byte 0, Bits 0–4)

Set to zero, indicating that the drive is a direct-access device, unless the LUN specified in the Command Descriptor Block is invalid. If the LUN is invalid, the peripheral-device type is set to $7F_H$—unknown device type. The other data bytes remain unaffected and are returned as though the LUN was set to zero.

RMB—REMOVABLE MEDIUM BIT (Byte 1, Bit 7)

This bit is always set to zero for the Maverick 270/540S hard disk drive, indicating nonremovable media.

DEVICE TYPE MODIFIER (Byte 1, Bits 6–0)

Not supported by the Maverick 270/540S hard disk drive. This parameter is set to zero.

ECMA VERSION (Byte 2, Bits 5–3)

Set to zero, indicating the Maverick 270/540S hard disk drive does not claim compliance to the European Computer Manufacturers Association version of SCSI.

ANSI VERSION (Byte 2, Bits 2–0)

Set to two, indicating compliance with the ANSI SCSI-2 specification.

RESPONSE DATA FORMAT (Byte 3, Bits 3–0)

Set to one, indicating compliance with the Common Command Set defined in the SCSI-2 Specification, ANSI X3T9.2/86-109, Revision 10C.

ADDITIONAL LIST LENGTH (Byte 4)

Specifies the length of the list of drive identification information, in bytes. Table 6-36 shows the Maverick 270/540S hard disk drive identification information. If the CDB allocation length is too small to allow the transfer of all vendor-unique parameters, the additional length value will not be adjusted to reflect the truncation.

Table 6-36 *Drive Identification Information*

BYTE	INFORMATION
5	Vendor Unique
6	Reserved=0
7	Set to 08H to indicate support for linked commands.
8–15	Vendor Identification = 'QUANTUM'
16–22	Product Identification (Model)= 270S or 540S
23-31	Product Identification (Part Number) = XXXXXXXXX
32–35	Microcode Revision Level = 'VV'VV is replaced by the appropriate revision level.
36–43	Microcode Date = 'MMDDYY' (2 spaces at the end) MM = month DD = day YY = year

BYTE	INFORMATION
44-55	Drive Serial Number = 'PTCYDDDLNNNN' P = place of manufacture T = drive type (fixed at 3) C = drive capacity (2=270, 5=540) Y = last digit of year drive built (example: 2=92, 3=93, etc.) DDD = julian day of manufacture (example: 139 = 139th day of the year) L = one digit line number NNNN = drive sequence on day of manufacture (example: 2 = 2nd drive built)
56–95	Reserved = 0
96–119	Vendor Unique

Note: Bytes 8 through 55 are ASCII values. If the drive is not ready when it receives the INQUIRY command, it will be unable to return the correct data for bytes 16–31 and 44–55. In this case, the drive returns the default data shown in Table 6-37.

Table 6-37 *Default Data*

BYTES	DATA
16–31	270SXXXXXXXXX or 540SXXXXXXXXX
44–55	No serial ##

6.9.10 MODE SELECT Command, Opcode = 15_H

The MODE SELECT command, shown in Table 6-38, allows the initiator to specify device parameters to the drive. A MODE SELECT command will override any previous selection of device parameters, even by another initiator. Command execution will create a UNIT ATTENTION condition for other initiators, provided that at least one parameter changes.

6.9.10.1 Command Structure

Table 6-38 *MODE SELECT Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =15 H							
1	LUN =0			PF =0	RESERVED =0			SP
2	RESERVED =0							
3	RESERVED =0							
4	PARAMETER LIST LENGTH (BYTES)							
5	VU =O		RESERVED =0				F	L

When the drive motor spins up at power on, the drive reads the set of device parameters most recently saved, from a reserved cylinder, and sets a UNIT ATTENTION condition for all initiators.

- If the drive cannot successfully read the parameter values from the Saved Mode Table, these values will revert to their defaults and the additional sense code will be set to Mode Select Parameter Changed, as though another initiator had altered the parameters.
- If the drive can successfully read the parameter values, it will set the additional sense code to Drive Reset, to indicate a reset condition.

Note: If bit 1, byte 2 of MODE SELECT, page 39_H, is set to a value of one, a UNIT ATTENTION condition will not occur. (See section 6.9.10.10.)

PF—PAGE FORMAT (Byte 1, Bit 4)

Not supported. See the ANSI Specification for a description of this parameter.

SP—SAVE PAGES (Byte 1, Bit 0)

When SP is set to one, only the pages of the Current Mode Table that can be saved will be copied to the Saved Mode Table currently selected.

PARAMETER LIST LENGTH (Byte 4)

Specifies the length of the MODE SELECT parameter list transferred during the DATA OUT phase, in bytes. A parameter list length of zero indicates that the drive will transfer no data.

6.9.10.2 MODE SELECT Parameter List

The MODE SELECT parameter list, shown in Table 6-39, contains a four-byte header, followed by block descriptors, then zero or more pages.

Table 6-39 *MODE SELECT Command Parameter List*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0							
1	MEDIA TYPE =0							
2	RESERVED =0							
3	BLOCK DESCRIPTOR LENGTH (BYTES)							

Block Descriptor

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	DENSITY CODE =0							
1-3	(MSB)	NUMBER OF BLOCKS						(LSB)
4	RESERVED =0							
5-7	BLOCK LENGTH							

Page Descriptor

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0		PAGE CODE					
1	PAGE LENGTH							
2-n	REFER TO PAGE DESCRIPTION							

MEDIA TYPE (Byte 1)

Set to zero—the current media type—by default, because the drive has nonremovable media.

BLOCK DESCRIPTOR LENGTH (Byte 3)

Specifies the length of all of the block descriptors, in bytes. This value is equal to the number of block descriptors times eight. It does not include the pages, if any. A block descriptor length of zero indicates the drive includes no block descriptors in the parameter list. The Maverick 270/540S hard disk drive uses a single block descriptor.

BLOCK DESCRIPTOR

Specifies the media characteristics of the drive, including its logical block length, density code, and number of blocks.

DENSITY CODE (Byte 0)

Set to zero. This parameter is not used by direct-access devices.

NUMBER OF BLOCKS (Byte 1-3)

Specifies the number of logical blocks on the media that match the density code and block length in the block descriptor. A value of zero indicates that all remaining logical blocks on the drive have the media characteristics specified by the block descriptor. Any nonzero value within the capacity of the drive can limit access to a specified number of blocks.

BLOCK LENGTH (Bytes 5–7)

Specifies the length of each logical block described by the block descriptor, in bytes. For the Maverick 270/540S hard disk drive, only a block length of 512 is allowed.

PAGE DESCRIPTORS

PAGE CODE (Byte 0, Bits 5–0)

Pages are optional. They can be included in any order, immediately following the block descriptor. To avoid the specification of all mode parameter each time the initiator issues a MODE SELECT command, the mode parameters are divided into pages. A page is the smallest unit that can be specified in a MODE SELECT or MODE SENSE command. Each time an initiator accesses a page, all parameters on that page must be specified. Modifiable parameters can be set to any acceptable value. Unmodifiable parameters must be set to zero. Pages are numbered for reference. Each page contains parameters grouped by functionality. For example, page 1 contains the read/write error recovery parameters that determine the drive's behavior during a data-handling error, including the retry count parameter, and bits that turn error detection on or off, and determine whether the drive reports soft errors. The drive supports the pages shown in Table 6-40.

Table 6-40 *Mode Pages Supported*

PAGE	DESCRIPTION
01_H	Read/Write Error-Recovery Parameters
02_H	Disconnect/Reconnect Control Parameters
03_H	Direct-Access Device Format Parameters*
04_H	Rigid Disk-Drive Geometry Parameters*
08_H	Cache-Control Parameters
$0C_H$	Notch and Partition Parameters
32_H	Automatic-Shutdown Control Parameters
37_H	Quantum-Unique Control Parameters
38_H	Not used
39_H	Quantum-Unique Drive-Control Parameters

Note: *Read only. Can be accessed only via the MODE SENSE command.

PAGE LENGTH (Byte 1)

Indicates the number of bytes for the page that follows, beginning with the first byte of flags or values, then continuing with consecutive bytes. The page length must be set to the value returned by the drive in the MODE SENSE page length byte. Otherwise, a CHECK CONDITION status will result, with a sense key of ILLEGAL REQUEST.

6.9.10.3 Error-Recovery Parameters, Page Code 01_H

Table 6-41 *Error-Recovery Parameters*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0		PAGE CODE =01 H					
1	PAGE LENGTH = 6H							
2	AWRE	ARRE	TB	RC	EEC	PER	DTE	DCR
3	RETRY COUNT							
4	CORRECTION SPAN							
5	RESERVED =0							
6	RESERVED =0							
7	RESERVED =0							

DCR—DISABLE CORRECTION (Byte 2, Bit 0)

When set to one, DCR indicates that the data will be transferred without double burst error correction, whether or not correction is possible. When set to zero, this bit indicates that the data will be corrected, if possible. Uncorrectable data will be transferred without attempting error-correction; however, retries will be attempted. If RC (bit 4 of byte 2) is set to one, the drive ignores this bit. The default is zero.

DTE—DISABLE TRANSFER ON ERROR (Byte 2, Bit 1)

When set to one and PER (bit 2 of byte 2) is also set to one, DTE indicates that the drive will enter CHECK CONDITION status immediately on detecting an error. The drive will terminate data transfer to the initiator. The block in error may or may not be transferred to the initiator, depending on the setting of the TB bit. A DTE bit set to zero enables data transfer for any data that can be recovered within the limits of the error-recovery flags. Errors are not posted until the transfer length is exhausted. If PER is zero or RC is one, the drive ignores this bit. The default is zero.

PER—POST ERROR (Byte 2, Bit 2)

When set to zero, PER indicates that the drive will not enter CHECK CONDITION status on errors recovered within the limits established by the other error-recovery flags. Recovery procedures that exceed the limits established by the other error-recovery flags will be posted. The data transfer may terminate prior to exhausting the transfer length, depending on the error and state of the other error-recovery flags. A PER bit set to one enables CHECK CONDITION status to be reported for detected errors, with the appropriate sense key. If multiple errors occur, the sense data will report the logical block address at which the unrecoverable error occurred. If no unrecoverable error occurred, the sense data will report the last block in which a recovered error occurred. The default is zero.

EEC—ENABLE EARLY CORRECTION (Byte 2, Bit 3)

When EEC is set to one, the drive will use its ECC algorithm if it detects two consecutive, equal, nonzero error syndromes. The drive will not perform read retries before applying correction, unless it determines that the error is uncorrectable. Seek or positioning retries, and the message system's recovery-procedure retries are not affected by the EEC bit's value. When set to zero, the drive will use its normal recovery procedures when an error occurs. If the RC bit is one, the drive ignores this bit. The default is zero.

RC—READ CONTINUOUS (Byte 2, Bit 4)

When RC is set to one, the drive transfers data of the requested length, without adding delays that would increase data integrity—that is, delays caused by the drive's error-recovery scheme. To maintain a continuous flow of data and avoid delays, the drive may send data that is erroneous. Ignored errors will not cause a CHECK CONDITION status. When set to zero, time-consuming, error-recovery operations are acceptable during data transfer. The default is zero.

TB—TRANSFER BLOCK (Byte 2, Bit 5)

When set to one, TB indicates that a data block not recovered within the recovery time limit specified will be transferred to the initiator before the CHECK CONDITION status is returned. When set to zero, this bit indicates that such a data block will not be transferred to the initiator. This bit does not affect the action taken for recovered data. The default is zero.

ARRE—AUTOMATIC READ REALLOCATION ENABLED (Byte 2, Bit 6)

When ARRE is set to one, the drive will enable automatic reallocation of bad blocks. Automatic reallocation functions similarly to the REASSIGN BLOCKS command, but is initiated by the drive when it encounters a hard error—that is, when it encounters the same nonzero ECC syndrome on two consecutive retries. When set to zero, the drive will not automatically reallocate bad blocks. When RC is one, the drive ignores this bit. The default is zero in PROM. Drive is shipped with this bit set to 1.

AWRE—AUTOMATIC WRITE REALLOCATION ENABLED (Byte 2, Bit 7)

When AWRE is set to one, the drive enables automatic reallocation of bad blocks. Automatic Write Reallocation is similar in function to Automatic Read Reallocation, but is initiated by the drive when a defective block becomes inaccessible for writes. When set to zero, the drive will not automatically reallocate bad blocks. The default is zero in PROM. Drive is shipped with this bit set to 1.

RETRY COUNT (Byte 3)

The number of times the drive will attempt to recover from a data error by rereading before it applies error correction. The default is eight.

CORRECTION SPAN (Byte 4)

Specifies the size, in bits, of the largest read data error on which correction can be attempted. This will correct as single burst of up to 48 bits in length, or one double burst of up to 24 bits each in length. The default is sixteen.

Note: The data-strobe offset and recovery-time limit (bytes 6–7) are not supported by the Maverick 270/540S hard disk drive.

Table 6-42 summarizes the valid modes of operation for the Maverick 270/540S hard disk drive.

Table 6-42 *Modes of Operation*

EEC	PER	DTE	DCR	DESCRIPTION
0	0	0	0	Normal error-recovery procedure. The drive attempts read retries until it reads good data, obtains a stable syndrome, or exhausts the retry count. When correction is possible, the drive invokes ECC. Data transfer is complete, unless the drive encounters an uncorrectable error. The drive reports only uncorrectable errors.
0	0	0	1	Same as 0,0,0,0—except the drive attempts no ECC correction. If read retries are unsuccessful, the drive stops the data transfer and reports an unrecoverable error.
0	0	1	0	Invalid Request.
0	0	1	1	Invalid Request.
0	1	0	0	Same as 0,0,0,0—except the drive reports all recoverable and unrecoverable data errors. The drive reports a recoverable error after the data transfer is complete.
0	1	0	1	Same as 0,0,0,0—except the drive reports all data errors. The drive reports a data error recovered through read retries after the data transfer is complete.
0	1	1	0	The drive attempts read retries until it reads good data, obtains stable syndrome, or exhausts the retry count. If error correction is possible, the drive invokes ECC. The drive stops data transfer on detecting an error and reports all data errors.
0	1	1	1	Same as 0,1,1,0—except the drive attempts no ECC correction. If read retries are unsuccessful, the drive reports an error as unrecoverable.
1	0	0	0	If error correction is possible, the drive immediately invokes ECC. If an error is uncorrectable, the drive attempts read retries until it reads good data, obtains a stable syndrome, or exhausts the retry count. Data transfer is complete, unless the drive encounters an unrecoverable error. The drive reports only unrecoverable errors.
1	0	0	1	Invalid Request.
1	0	1	0	Invalid Request.
1	0	1	1	Invalid Request.

EEC	PER	DTE	DCR	DESCRIPTION
1	1	0	0	Same as 1,0,0,0—except the drive reports all data errors. The drive reports a recoverable error after the data transfer is complete.
1	1	0	1	Invalid Request.
1	1	1	0	Same as 1,0,0,0—except the drive stops the data transfer on detecting a recoverable or unrecoverable error, and reports all data errors.
1	1	1	1	Invalid Request.

6.9.10.4 Disconnect/Reconnect Control Parameters, Page Code 02_H

Table 6-43 *Disconnect/Reconnect Control Parameters*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0		PAGE CODE =02 h					
1	PAGE LENGTH =Ah							
2	BUFFER FULL RATIO =0							
3	BUFFER EMPTY RATIO =0							
4-11	RESERVED =0							

BUFFER FULL RATIO (Byte 2)

On reads, the drive disconnects when the buffer contains no data. For commands that require data transfer to the initiator, the buffer full ratio represents the percentage of the buffer that must become full before the drive will reconnect—unless the buffer can hold all requested data. The buffer full ratio is automatically controlled by the drive. Any value may be written to and read back from this byte, but there will be no effect on the drive operation.The drive must transfer at least 512 bytes to the buffer before reconnection.

BUFFER EMPTY RATIO (Byte 3)

For commands that require data transfer from the initiator, the buffer empty ratio represents the percentage of the buffer that must become empty before the drive will reconnect to fetch more data—unless the buffer can hold all requested data. The buffer full ratio is automatically controlled by the drive. Any value may be written to and read back from this byte, but there will be no effect on the drive operation.

Note: For commands that require a logical block transfer, the drive rounds the buffer full ratio down and the buffer empty ratio up, to the nearest multiple of 512 bytes.

6.9.10.5 Direct-Access Device Format, Page Code 03_H

Table 6-44 *Direct-Access Device Format*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED=0		PAGE CODE =03 H					
1	PAGE LENGTH =16 H							
2-3	TRACK PER ZONE							
4-5	ALTERNATE SECTORS PER ZONE							
6-7	ALTERNATE TRACKS PER ZONE							
8-9	ALTERNATE TRACKS PER LOGICAL UNIT							
10-11	SECTORS PER TRACK							
12-13	DATA BYTES PER PHYSICAL SECTOR							
14-15	INTERLEAVE							
16-17	TRACK SKEW FACTOR							
18-19	CYLINDER SKEW FACTOR							
20	ssec=0	hsec=1	rmb=0	surf=0	RESERVED =0			
21-23	RESERVED =0							

Note: The drive uses only the default values for the fields in bytes 2–9 and ignores any other values.

DEFECT-HANDLING FIELDS

TRACKS PER ZONE (Bytes 2–3)

The value in this field indicates the number of tracks per defect zone. The Maverick 270/540S has four tracks per defect zone.

ALTERNATE SECTORS PER ZONE (Bytes 4–5)

The value in this field indicates that the drive will deallocate one sector per zone from the initiator-addressable blocks on execution of the FORMAT UNIT command. These sectors are available as replacement sectors for defective sectors.

ALTERNATE TRACKS PER ZONE (Bytes 6–7)

Set to zero. The drive does not allocate alternate tracks.

ALTERNATE TRACKS PER VOLUME (Bytes 8–9)

Set to zero. The drive does not allocate alternate tracks.

TRACK-FORMAT FIELD

SECTORS PER TRACK (Bytes 10–11)

The value in this field indicates the number of physical sectors the drive allocates per track. Set to zero, indicating that the drive has a variable number of sectors per track (52 to 91).

SECTOR-FORMAT FIELDS

DATA BYTES PER PHYSICAL SECTOR (Bytes 12–13)

This parameter indicates the number of data bytes the drive allocates per physical sector. This value may be different from the block descriptor length specified in the MODE SELECT parameters. Each physical sector on the drive contains 512 data bytes.

INTERLEAVE (Bytes 14–15)

The drive has an interleave factor of one.

TRACK SKEW FACTOR (Bytes 16–17)

This parameter indicates the number of physical sectors between the last logical block on one track and the first logical block on the next sequential track on the same cylinder. See Table 6-40 for track skews in the Maverick 270/540S hard disk drive.

CYLINDER SKEW FACTOR (Bytes 18–19)

This parameter indicates the number of physical sectors between the last logical block on one cylinder and the first logical block on the next sequential cylinder. See Table 6-45 for cylinder and track skews in the Maverick 270/540S hard disk drive.

SSEC—SOFT SECTOR (Byte 20, Bit 7)

SSEC is always set to zero, indicating that the Maverick 270/540S hard disk drive does not use soft sector formatting.

HSEC—HARD SECTOR (Byte 20, Bit 6)

HSEC is always set to one, indicating that the Maverick 270/540S hard disk drive uses hard sector formatting. The HSEC and SSEC bits are mutually exclusive.

RMB—REMOVABLE (Byte 20, Bit 5)

Set to zero for the Maverick 270/540S hard disk drive, indicating that the logical unit is not removable.

SURF—SURFACE (Byte 20, Bit 4)

Set to zero, indicating that the drive allocates successive addresses to all sectors within a cylinder, prior to allocating sector addresses to the next cylinder.

Table 6-45 *Track and Cylinder Skewing*

ZONE NUMBER	SECTORS PER TRACK	TRACK SKEW WEDGE	CYLINDER SKEW WEDGE
15	58	28	32
14	65	28	32
13	69	28	32
12	74	28	32
11	78	28	32
10	83	28	32
9	88	28	32
8	93	28	32
7	97	28	32
6	104	28	32
5	108	28	32
4	112	28	32
3	114	28	32
2	118	28	32
0,1	118	28	32
System	93	28	32

6.9.10.6 Cache Control, Page Code 08_H

Table 6-46 *Cache Control*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0		PAGE CODE =8 H					
1	PAGE LENGTH=0A H							
2	RESERVED =0					WCE	MS =0	RCD
3	DEMAND READ RETENTION PRIORITY=0				WRITE RETENTION PRIORITY =0			
4-5	DISABLE PREFETCH TRANSFER LENGTH =0							
6-7	(MSB)	MINIMUM PREFETCH						(LSB)
8-9	(MSB)	MAXIMUM PREFETCH						(LSB)
10-11	MAXIMUM PREFETCHING CEILING							

The Cache Control Parameters page specifies the parameters that control the operation of the cache. This page and page 37_H should be used to control the caching parameters. Parameters set in either page cause the drive to automatically set the corresponding parameters in the other page.

WCE—WRITE CACHE ENABLE (Byte 2, Bit 2)

Set to one by default, indicating that the drive returns GOOD status for a WRITE command after successfully receiving the data, but before writing it to the disk. A value of zero indicates that the drive returns GOOD status for a WRITE command after successfully receiving the data and writing it to the disk. The default is set to zero in PROM. Drive is shipped with this bit set to one.

MF—MULTIPLICATION FACTOR (Byte 2, Bit 1)

Not supported.

RCD—READ CACHE DISABLE (Byte 2, Bit 0)

Set to zero by default, indicating that the Maverick 270/540S hard disk drive can return some or all of the data requested by a READ command by accessing the cache or the disk. Setting the RCD bit to one causes the drive to automatically clear the PE and CE bits in page 37_H. The drive must read all requested data from disk and cannot return any data by accessing the cache. Setting the CE bit in page 37_H causes the drive to turn on the RCD bit in this page.

When prefetching data, the drive reads data not yet specifically requested by an initiator and stores it in the cache - usually in conjunction with reading requested data from disk. Prefetching always begins at the logical block immediately following the last logical block previously read and always stops before the end of the media. The drive reports no errors to the initiator during prefetching, unless it is unable to execute subsequent commands successfully due to the error. Subsequent parameters listed in Table 6-46 indicate how much additional data the drive will read and store in the cache following a READ command.

DEMAND READ RETENTION PRIORITY (Byte 3, Bits 4–7)

This field advises the drive on the retention priority to assign the data that is read into the cache that has also been transferred from the drive to the initiator. A value of zero indicates that the drive should not distinguish between retaining the indicated data and data placed into the cache by other means (such as prefetch).

WRITE RETENTION PRIORITY (Byte 3, Bits 0–3)

This field advises the drive on the retention priority to assign the data written into the cache that has also been transferred from the cache to the drive. A value of zero indicates that the drive should not distinguish between retaining the indicated data and data placed into the cache by other means (such as prefetch).

DISABLE PREFETCH TRANSFER LENGTH (Bytes 4–5)

Not supported.

MINIMUM PREFETCH (Bytes 6–7)

Not supported.

MAXIMUM PREFETCH (Bytes 8–9)

Not supported.

MAXIMUM PREFETCHING CEILING (Bytes 10–11)

Not supported.

6.9.10.7 Notch and Partition, Page Code $0C_H$

The Notch and Partition Page, shown in Table 6-47, contains the parameters that define the notch characteristics of the drive. Each notch, or zone, has a unique number of blocks per cylinder and spans a range of consecutive logical blocks. Notches cannot overlap, nor can a logical block within the boundaries of a notch be excluded from that notch.

Table 6-47 *Notching and Partitioning*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	ps =0	rsvd =0	PAGE CODE =0C H					
1	PAGE LENGTH =16 H							
2	ND=0	PLN=0	RESERVED =0					
3	RESERVED =0							
4-5	MAXIMUM NUMBER OF NOTCHES							
6-7	ACTIVE NOTCH							
8-11	STARTING BOUNDRY							
12-15	ENDING BOUNDRY							
16-23	PAGES NOTCHED							

PS—PARAMETERS SAVEABLE (Byte 0, Bit 7)

The MODE SENSE command reserves this bit for its own use. In the Maverick 270/540S hard disk drive, PS is set to zero, indicating that the drive does not save the page.

ND—NOTCHED DRIVE (Byte 2, Bit 7)

This parameter indicates whether the drive is notched or not. Each notch has a different numbers of blocks per cylinder. When set to zero, ND indicates that the device is not notched. The drive returns all other parameters in the page as zero. When set to one, ND indicates that the device is notched. On notched, direct-access devices, starting and ending boundaries on the media define the notches. The Maverick 270/540S hard disk drive is a notched device—that is, ND is set to one. This is a read-only parameter.

PLN—PHYSICAL OR LOGICAL NOTCH (Byte 2, Bit 6)

This parameter indicates whether the notch boundaries are physical or logical locations on the drive. When PLN is set to zero, the drive uses physical notch boundaries. Physical locations are defined by cylinder and head. When PLN is set to one, the drive uses logical notch boundaries. Logical locations are defined by logical block address. The Maverick 270/540S hard disk drive uses physical notch boundaries—PLN is set to zero. This is a read-only parameter.

MAXIMUM NUMBER OF NOTCHES (Bytes 4–5)

This parameter defines the maximum number of notches supported by the drive. The Maverick 270/540S hard disk drive supports a maximum of sixteen notches. This is a read-only parameter.

ACTIVE NOTCH (Bytes 6–7)

This parameter indicates the notch to which this and subsequent MODE SELECT and MODE SENSE commands refer—until a later MODE SELECT command changes this parameter. When this parameter is set to zero, this and subsequent MODE SELECT and MODE SENSE commands refer to those parameters that apply across notches. For the Maverick 270/540S hard disk drive, valid notch numbers range from zero to fifteen. This is the only Notch and Partition Page parameter that can be set by the MODE SELECT command.

STARTING BOUNDARY (Bytes 8–11)

This parameter defines the starting address of the active notch. When active notch is set to zero, this parameter defines the starting address of the logical unit. For all drives, bytes 8–10 define the cylinder and byte 11 defines the head. This is a read-only parameter.

ENDING BOUNDARY (Bytes 12–15)

This parameter defines the ending address of the active notch. Only the MODE SENSE command can set this parameter. When active notch is set to zero, this parameter defines the ending address of the logical unit. For all drives, bytes 12–14 define the cylinder and byte 15 defines the head. This is a read-only parameter.

PAGES NOTCHED (Bytes 16–23)

This parameter consists of a bit map of the mode page codes that indicates pages containing parameters that can be different for different notches. The most significant bit of this field corresponds to page $3F_H$; the least significant bit, to page 00_H. When a bit is set to one, the corresponding mode page contains parameters that can be different for different notches. When a bit is set to zero, the corresponding mode page contains parameters that are constant for all notches. For the Maverick 270/540S hard disk drive, the bits corresponding to pages 03_H and $0C_H$ are set to one, and are notch dependent. All other bits are set to zero. This is a read-only parameter.

Table 6-48 *Starting and Ending Boundaries for Active Notches*

ACTIVE NOTCH	STARTING BOUNDARY	ENDING BOUNDARY
15	2613	2852
14	2435	2612
13	2297	2434
12	2179	2296
11	1941	2178
10	1783	1940
9	1585	1782
8	1397	1586
7	1219	1396
6	1031	1218
5	873	1030
4	745	872
3	597	744
2	359	596
1	200	358
0	0	199

6.9.10.8 Automatic Shutdown Control, Page Code 32_H

Table 6-49 *Automatic Shutdown Control*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0		PAGE CODE =32 H					
1	PAGE LENGTH =2H							
2	AUTO STANDBY TIME							
3	AUTO SHUTDOWN TIME (MINUTES)							

AUTO STANDBY TIME (Byte 2)

The maximum time, in minutes, that the drive can remain selected before it will enter the Standby mode.The actuator is on and the motor remains spinning. When this byte is set to zero (the default), Auto Standby is disabled.

AUTO SHUTDOWN TIME (Byte 3)

The maximum time period, in minutes, the drive can remain deselected before entering power-shutdown mode. On entering this mode, the drive turns off its servo and motor circuitry. When this byte is set to zero, auto shutdown is disabled. The default is zero.

6.9.10.9 Quantum-Unique Control Parameters, Page Code 37_H

Table 6-50 *Quantum-Unique Control Parameters*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	Reserved=0		PAGE CIDE =37 H					
1	PAGE LENGTH =0E H							
2	Reserved =0		PSM	SSM	WIE	PO	PE	CE
3	NUMBER OF CACHE SEGMENTS							
4	MINIMUM PREFETCH							
5	MAXIMUM PREFETCH							
6-15	RESERVED =0							

PSM—PRESERVE SYNCHRONOUS MODE (Byte 2, Bit 5)

When PSM is set to one, the drive will not clear the table containing the synchronous-mode parameters for all initiators when the drive is reset. When PSM is set to zero will cause all these parameters to be cleared when the drive is reset. The default is zero.

SSM—SEND SYNCHRONOUS MESSAGE (Byte 2, Bit 4)

When SSM is set to one, the drive sends the Extended Message (01) Synchronous Data Transfer Request to the initiator. When SSM is set to zero, the initiator sends the Synchronous Data Transfer Request message. The default is zero.

Note: The following parameters control the operation of Quantum's DisCache. See Chapter 5 for a description of DisCache and more information about its options.

Mode Select page 8_H can also be used to control the operation of **DisCache**—see Section 6.9.10.5. When these parameters are set in either page, the drive automatically sets the corresponding parameters in the other page.

WIE—WRITE INDEX ENABLE (Byte 2, Bit 3)

This bit is set by default to zero and is not used.

PO—PREFETCH ONLY (Byte 2, Bit 2)

This bit is set by default to zero and is not used.

PE—PREFETCH ENABLE (Byte 2, Bit 1)

When PE is set to one, the drive prefetches data into the cache. When PE is set to zero, the drive will not prefetch data into the cache. To enable the PE bit, the CE bit must be set to one, the default. The drive automatically sets this bit when the RCD bit in page 8_H is set to zero. PE is set to one by default.

CE—CACHE ENABLE (Byte 2, Bit 0)

When CE is set to one, the drive will activate caching on all reads. When CE is set to zero, the drive will disable caching and use the 96K RAM only as a transfer buffer. The default is one. The drive automatically sets this bit when the RCD bit in page 8_H is set to zero. When the CE bit is set, the drive automatically turns on the RCD bit in page 8_H. CE is set to one by default.

NUMBER OF CACHE SEGMENTS (Byte 3)

This parameter indicates the number of segments the drive can index in the read cache—that is, the number of entries in the cache table. This byte is read-only and is dynamically adjusted by the drive. The default is one.

MINIMUM PREFETCH (Byte 4)

This byte is dynamically adjusted by the drive.

MAXIMUM PREFETCH (Byte 5)

This byte is dynamically adjusted by the drive.

6.9.10.10 Quantum-Unique Drive Control Parameters, Page Code 39_H

Table 6-51 *Quantum-Unique Drive Control Parameters*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0		PAGE CODE =39 H					
1	PAGE LENGTH =06 H							
2	DIO	DII	FDB	RUEE	FDPE	R =0	DUA	DRT
3	DDIS	DELDIS	RES=0	DPC	SSID	SCSIADR		
4	RESERVED =0							
5	MOTOR DELAY TIME							
6-7	RESERVED =0							

DIO—DISABLE IDENTIFY OUT (Byte 2, Bit 7)

When DIO is set to one, the drive does not require an IDENTIFY message from the initiator to disconnect and reconnect. When DIO is set to zero, the drive requires the initiator to send an IDENTIFY message prior to disconnecting. The default is zero.

DII—DISABLE IDENTIFY IN (Byte 2, Bit 6)

When DII is set to one, the drive does not send an IDENTIFY message after reconnecting to an initiator. When DII is set to zero, the drive sends an IDENTIFY message after reconnecting. The default is zero.

FDB—FAST DEASSERTION OF BUSY (Byte 2, Bit 5)

This parameter is not used by the drive. The default value is zero.

RUEE—REALLOCATE UNCORRECTABLE ERROR ENABLED (Byte 2, Bit 4)

When RUEE is set to one and the ARRE bit in page 1 of the error-recovery parameters is set, the drive automatically writes data containing uncorrectable, hard errors to a different logical block and appends a new ECC that is correct for the data. Reading the new block would result in a MEDIUM ERROR—Reallocated Uncorrectable Data Read, with the sense key set to 03_H and the sense code to AA_H. Any write to the new block will clear this condition. When RUEE is set to zero and the drive determines that a hard error it encountered is uncorrectable, the drive will not attempt automatic reallocation. A CHECK CONDITION status, with a sense key of MEDIUM ERROR—Uncorrectable Data Error will result. The default is zero in PROM. The drive is shipped with this bit set to one.

FDPE—FILL DATA PATTERN ENABLED (Byte 2, Bit 3)

When FDPE is set to one, the drive will write the data pattern specified in byte 2 of the FORMAT UNIT command into every user-accessible sector on the drive when it executes a FORMAT UNIT command. When FDPE is set to zero, the drive will ignore the information in byte 2 of the FORMAT UNIT command. Issuing a FORMAT UNIT command typically causes loss of data, even if FDPE is set to zero. Users should back up their data prior to formatting. The default is zero.

R—RESERVED (Byte 2, Bit 2)

Bit 2 of byte 2 is reserved and must be set to zero.

DUA—DISABLE UNIT ATTENTION (Byte 2, Bit 1)

When DUA is set to one, the drive will not issue a CHECK CONDITION status with the UNIT ATTENTION sense key on the first command following a reset. The drive will execute the first command issued after a reset condition. When DUA is set to zero, a UNIT ATTENTION condition will exist following a power on, or a reset caused by either a BUS DEVICE RESET message or a hard RESET condition. The default is zero.

DRT—DISABLE RESELECTION TIMEOUT (Byte 2, Bit 0)

When DRT is set to one, the drive will not time out during a reselection request. When DRT is set to zero and there is no BSY response from the initiator after a selection time-out period—250 milliseconds—the drive will clear the bus. The default is zero.

DDIS—DISABLE DISCONNECTION (Byte 3, Bit 7)

Setting DDIS to one prevents disconnection and reconnection during a data transfer—either a read or write. The initial disconnection due to the implied seek during the execution of a READ command is not prevented, but no other disconnect will occur. On a WRITE command, the drive will disconnect after transferring all data to the buffer. When DDIS is set to zero, disconnection is not suppressed. The default is zero.

DELDIS—DELAY OF DISCONNECTION (Byte 3, Bit 6)

This parameter is not used by the drive. The default is zero.

R—RESERVED (Byte 3, Bit 5)

Bit 5 of byte 3 is reserved and must be set to zero.

DPC—DISABLE PARITY CONTROL (Byte 3, Bit 4)

When set to zero, DPC controls parity checking of data across the SCSI bus. Setting DPC to one disables parity checking (the drive still generates parity, but does not perform parity checking). The default is zero.

SSID—SOFTWARE SELECTABLE ID (Byte 3, Bit 3)

When SSID is set to zero, the jumper settings on the drive PCB determine the drive's SCSI ID at all times. When SSID is set to one, the SCSIADR bits (bits 0–2 of byte 3) determine the drive's SCSI ID and the jumper settings will be ignored. When the SSID bit is set to one, the SCSI ID will remain unchanged until the next reset or power on. If at initial power on, the SSID bit is set to one—that is, if this bit's previous setting was saved on the disk—the SCSIADR bits will determine the drive's SCSI ID. Otherwise, the jumper settings on the drive PCB will determine its SCSI ID. The default is zero.

SCSIADR—SCSI Address (Byte 3, Bits 2–0)

If SSID (bit 3 of byte 3) is set to one, SCSIADR determines the drive's SCSI ID. If SSID is set to zero, the jumper settings on the drive PCB determine its SCSI ID. (See the description of SSID above.) The default is zero.

MOTOR DELAY TIME—(Byte 5)

The value stored in this byte specifies the motor delay time in 10 millisecond increments. The LSB (bit 0) is equal to 10 milliseconds.

6.9.11 RESERVE Command, Opcode = 16_H

The RESERVE command, shown in Table 6-52, reserves a logical unit for the use of the initiator or another specified, third-party SCSI device. This command requests that the drive be reserved for the exclusive use of the initiator or third-party SCSI device, until the reservation is superseded by another valid RESERVE command from the same initiator; or is released by a RELEASE command from the same initiator, a BUS DEVICE RESET message from any initiator, or a hard RESET condition. The reservation of the drive will not be granted if another initiator has already reserved the drive. An initiator can reserve the drive when it is already reserved for that initiator. The reservation identification and the extent list length will be ignored.

If the drive is reserved for another initiator, it will return a RESERVATION CONFLICT status.

If, after honoring the reservation, any other initiator subsequently attempts to perform a command on the reserved drive—other than a RELEASE command, which will be ignored—the drive will reject the command, with RESERVATION CONFLICT status.

Table 6-52 *RESERVE Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =16 H							
1	LUN =0			3rdprty	3RD PARTY DEVICE ID			extent
2	RESERVATION IDENTIFICATION =XX							
3-4	EXTENT LIST LENGTH =XX							
5	VU =0		RESERVED =0				F	L

EXTENT—EXTENT RESERVATION (Byte 1, Bit 0)

This parameter must be set to zero. If this bit were set to one, the drive would reject the RESERVE command with CHECK CONDITION status and the ILLEGAL REQUEST sense key.

3RD-PARTY RESERVATION (Byte 1, Bit 4)

This RESERVE command option allows an initiator to reserve a logical unit for another SCSI device. It is intended for use in multiple-initiator systems that use the COPY command.

3RD-PARTY DEVICE ID (Byte 1, Bits 1-3)

This RESERVE command option provides a SCSI device ID for a third party initiator (reserving the drive).

6.9.12 RELEASE Command, Opcode = 17_H

The RELEASE command, shown in Table 6-53, releases a logical unit that was previously reserved. It will not cause an error if an initiator attempts to release a reservation that is not currently active.

Table 6-53 *RELEASE Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =17 H							
1	LUN =0			3rdprty	3RD PARTY DEVICE ID			Extent
2	RSEERVATION IDENTIFICATION =XX							
3-4	RESERVED =0							
5	VU =0		RESERVED =0				F	L

EXTENT—EXTENT RESERVATION (Byte 1, Bit 0)

This parameter must be set to zero. If this bit were set to one, the drive would reject the RELEASE command with CHECK CONDITION status and the ILLEGAL REQUEST sense key.

3RD-PARTY RESERVATION (Byte 1, Bit 4)

This RESERVE command option allows an initiator to release a logical unit that was previously reserved using a third-party reservation. It is intended for use in multiple-initiator systems that use the COPY command. When the third-party bit is set to zero, the initiator will not request a third-party release. When the third-party bit is set to one, the drive will release the specified logical unit only if the reservation was made by the same initiator, using the third-party reservation option, with the same SCSI device specified in the third-party device-ID field.

3RD-PARTY DEVICE ID (Byte 1, Bits 1-3)

This RELEASE command option releases a SCSI device ID for a third party initiator who had reserved the drive.

6.9.13 MODE SENSE Command, Opcode = $1A_H$

The MODE SENSE command, shown in Table 6-54, allows a target to report its peripheral-device parameters to the initiator. This command complements the MODE SELECT command.

6.9.13.1 MODE SENSE Command Structure

Table 6-54 *MODE SENSE Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =1A H							
1	LUN =0			RESERVED =-0				
2	PCF		PAGE CODE					
3	RESERVED =0							
4	ALLOCATION LENGTH							
5	VU =0		RESERVED =0				F	L

PCF—PAGE CONTROL FIELD (Byte 2, Bits 7 & 6)

Defines the page parameter values to be returned. The four options—Report Current Values, Report Changeable Values, Report Default Values, and Report Saved Values—are explained below.

Report Current Values (PCF = 00)

The drive returns the page, which is defined by the page code, to the initiator. The fields and bits are set to the current values, which are either:

- Those set by the last MODE SELECT command successfully completed
- Identical to the values saved—if the initiator has issued no MODE SELECT command since the last power on or reset

Fields and bits not supported by the drive are set to zero.

Report Changeable Values (PCF = 01)

The drive returns the page defined by the page code to the initiator, with all fields and bits that can be modified by the initiator set to one. Fields and bits that cannot be changed by the initiator are set to zero.

Report Default Values (PCF = 10)

The drive returns the page defined by the page code to the initiator, with fields and bits set to the drive's default values—that is, as shipped from the factory. Fields and bits not supported by the drive are set to zero. To determine whether a returned value of zero indicates a default parameter or an unsupported parameter, the initiator can examine the modifiable values.

Report Saved Values (PCF = 11)

The drive returns the page defined by the page code to the initiator, with fields and bits set to the values saved. The values saved are either:

- Those set in the last MODE SELECT command successfully completed, with the Save Parameters bit (byte 1, bit 0) of the CDB set to one
- Identical to the default values—if no MODE SELECT command completed successfully, with the SP bit set

Fields and bits not supported by the drive are set to zero.

PAGE CODE (Byte 2, Bits 5–0)

For all page-control parameters, the value returned in the page-length byte indicates the number of bytes the drive supports within each page. This value must be specified in the page-length field (byte 1 of each page descriptor) when issuing the MODE SELECT command. The page code specifies the specific page information to be returned to the initiator in the MODE SENSE data. If a page code of zero is specified, command is rejected with an error code 5/24. The drive supports the pages shown in Table 6-55.

Table 6-55 *Pages Supported*

PAGE	DESCRIPTION
01_H	Read/Write Error-Recovery Parameters.
02_H	Disconnect/Reconnect Control Parameters.
03_H	Direct-Access Device Format Parameters (MODE SENSE only).
04_H	Rigid Disk Drive Geometry Parameters—valid for the MODE SENSE command only.
08_H	Cache-Control Parameters.
$0C_H$	Notch and Partitioning Page.
32_H	Automatic-Shutdown Control Parameters.
37_H	Quantum-Unique Control Parameters.
38_H	Not used.
39_H	Quantum-Unique Drive-Control Parameters.
$3F_H$	Return all pages to the initiator—valid for the MODE SENSE command only.

The MODE SENSE data provides information to the initiator about the drive's format parameters, if page 03_H is specified, and its drive-geometry parameters, if page 04_H, is specified. If page $3F_H$ is specified, the drive returns information from all pages implemented to the initiator in the MODE SENSE data. This page code is valid only in the MODE SENSE command. Specifying page code 04_H in a MODE SELECT command results in a CHECK CONDITION status, with the Illegal Request sense key.

ALLOCATION LENGTH (Byte 4)

This parameter specifies the number of bytes the initiator has allocated for returned MODE SENSE data. An allocation length of zero indicates that the drive will transfer no MODE SENSE data. Any other value indicates the maximum number of bytes that the drive will transfer. The drive will terminate the DATA IN phase when it has transferred either the allocation-length bytes or all available MODE SENSE data to the initiator, whichever is less.

MODE SENSE DATA

The MODE SENSE data, shown in Table 6-56, contain a four-byte header, followed by an eight-byte block descriptor, then zero or more pages. The meaning and organization of these data are the same as for the corresponding MODE SELECT data—see Section 6.9.10—which are modified by the option specified in the page-control field of the CDB, byte 2.

6.9.13.2 MODE SENSE Data

Table 6-56 *MODE SENSE Command Data*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	SENSE DATA LENGTH							
1	MEDIUM TYPE =0							
2	WP =0	RESERVED =0						
3	BLOCK DESCRIPTOR LENGTH =8							

Block Descriptor

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	DENSITY CODE =0							
1-3	NUMBER OF BLOCKS =0							
4	RESERVED =0							
5-7	BLOCK LENGTH							

Page Descriptor

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	PS	R =0	PAGE CODE					
1	PAGE LENGTH							
2-n	REFER TO PAGES DEFINITION AND SECTION 6.9.10							

SENSE DATA LENGTH (Byte 0)

This parameter specifies the length, in bytes, of the MODE SENSE data to be transferred during the DATA IN phase. The sense data length does not include its own length.

WP—WRITE PROTECTED (Byte 2, Bit 7)

WP is always set to zero, indicating that the drive is write enabled.

BLOCK DESCRIPTOR LENGTH (Byte 3)

This parameter specifies the length of all of the block descriptors, in bytes, and is set to eight for the Maverick 270/540S hard disk drive.

Block Descriptor

The block descriptor specifies the media characteristics of the drive—in its density code, number of blocks, and block length. These characteristics are the same as those in the corresponding fields in the MODE SELECT parameter list.

Page Descriptor

PS—PARAMETERS SAVEABLE (Byte 0, Bit 7)

When PS is set to zero in each page header, the drive cannot save the supported parameters on that page. When PS is set to one, the drive can save the supported parameters on that page. The drive can save all pages with parameters that can be modified by the initiator.

6.9.13.3 Direct-Access Device Format Parameters, Page Code 03_H

Table 6-57 shows the direct-access device format parameters.

Table 6-57 *Direct Access Device Format Parameters*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0		PAGE CODE =03 H					
1	PAGE LENGTH =16 H							

Handling of Defects Field

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
2-3	TRACK PER ZONE *							
4	ALTERNATE SECTORS PER ZONE =0							
5	ALTERNATE SECTORS PER ZONE =1 H							
6	ALTERNATE TRACKS PER ZONE =0							
7	ALTERNATE TRACKS PER ZONE *							
8	ALTERNATE TRACKS PER VOLUME =0							
9	ALTERNATE TRACKS PER VOLUME =0							

Track Format Fields

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
10	SECTORS PER TRACK =0							
11	SECTORS PER TRACK =0 **							

Sector Format Fields

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
12	DATA BYTES PER PHYSICAL SECTOR =02 H							
13	DATA BYTES PER PHYSICAL SECTOR =0							
14	INTERLEAVE =0							
15	INTERLEAVE =1 H							
16	TRACK SKEW FACTOR =0							
17	TRACK SKEW FACTOR =**							
18	CYLINDER SKEW FACTOR =0							
19	CYLINDER SKEW FACTOR = **							

Drive Type Field

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
20	ssec=1	hsec =0	rmb =0	surf =0	ins =0	RESERVED =0		
21	RESERVED =0							

* The value for alternate tracks per zone is one for the Maverick 270/540S hard disk drive.

** These fields contain parameters that are zone dependent. Their values vary depending on the active zone.

6.9.13.4 Rigid Disk Drive Geometry Parameters, Page Code 04_H

Table 6-58 *Rigid Disk Drive Geometry Parameters*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0		PAGE CODE =04 H					
1	PAGE LENGTH =12 H							
2	(MSB)	NUMBER OF CYLINDERS =0						
3	NUMBER OF CYLINDERS =0B H							(MSB)
4	NUMBER OF CYLINDERS =25 H							(LSB)
5	NUMBER OF HEADS (04 H or 02H))							
6-8	STARTING CYLINDER - WRITE RECOMPENSATION =0							
9-11	STARTING CYLINDER - REDUCED WRITE CURRENT							
12-13	DRIVE STEP RATE =0							
14-16	LANDING ZONE CYLINDER =0							
17-19	RESERVED =0							

Note: The Maverick 270S hard disk drive has 2,853 data cylinders and 2 heads, while the 540S has 2,853 data cylinders and 4 heads.

The embedded SCSI controller on the Maverick 270/540S hard disk drive handles the write precompensation starting cylinders, reduced write current starting cylinders, and drive step rate fields.

STARTING CYLINDER WRITE PRECOMPENSATION (Bytes 6–8)

This field contains the number of the physical cylinder at which the drive begins to perform write precompensation (set to zero).

STARTING CYLINDER REDUCED WRITE CURRENT (Bytes 9–11)

This field contains the number of the physical cylinder at which the drive begins to write with reduced write current (set to zero). It is set by default to zero (0_H) which disables this function.

DRIVE STEP RATE (Bytes 12–13)

This field contains the step rate in 100 nanosecond increments. The drive uses a step rate greater than or equal to the rate specified in this field (set to zero).

LANDING ZONE CYLINDER (Bytes 14–16)

This parameter is set to zero, because the drive automatically parks the heads in the landing zone, using AIRLOCK, at power off. This field applies only to drives that do not automatically seek to the landing zone before stopping the spindle motor (set to zero).

6.9.14 START/STOP UNIT Command, Opcode = $1B_H$

The START/STOP UNIT command, shown in Table 6-59, requests that the logical unit be enabled or disabled for further operations.

Table 6-59 *START/STOP UNIT Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =1B H							
1	LUN =0			RESERVED =0				IMMED
2-3	RESERVED =0							
4	RESERVED =0							
5	VU =0		RESERVED =0				F	L

IMMED—IMMEDIATE (Byte 1, Bit 0)

When IMMED is set to one, the drive will return status as soon as it initiates an operation. When IMMED is set to zero, the drive will return status after completing an operation.

START (Byte 4, Bit 0)

When set to one, this parameter requests that the drive be started. When set to zero, this parameter requests that the drive be stopped. A request to start or stop the drive can be repeated without causing an error. If the state requested is the same as the drive's current state, this parameter will have no effect on the drive.

Note: When the drive has stopped, due to a START/STOP UNIT command, both the spindle motor and actuator are disabled and do not draw appreciable currents. Starting the drive requires up to 12 seconds before commands requiring disk access can be executed. See TEST UNIT READY.

6.9.15 SEND DIAGNOSTIC Command, Opcode = $1D_H$

The SEND DIAGNOSTIC command, shown in Table 6-60, requests the drive to perform diagnostic tests on itself.

6.9.15.1 Command Structure

Table 6-60 *SEND DIAGNOSTIC Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	opcode =1D h							
1	LUN =0			RESERVED =0		selftst	devofl	unitofl
2	RESERVED =0							
3-4	(MSB)	PARAMETER LIST LENGTH =0						(LSB)
5	VU =0		RESERVED =0				F	L

SELF TEST (Byte 1, Bit 2)

For the Maverick 270/540S hard disk drive, this parameter always should be set to one, to directs the drive to complete its default self test. If the self test is successful, the command will terminate with a GOOD status. Otherwise, the command will terminate with a CHECK CONDITION status, with the sense key set to HARDWARE ERROR. When set to zero, the command will terminate with a CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST.

DEVOFL—SCSI DEVICE OFF-LINE (Byte 1, Bit 1)

Not supported. If this bit is set to one, the command will terminate in a CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the sense code Invalid Field.

UNITOFL—LOGICAL UNIT OFF-LINE (Byte 1, Bit 0)

Not supported. If this bit is set to one, the command will terminate in a CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the sense code Invalid Field.

PARAMETER LIST LENGTH (Bytes 3–4)

Not supported. Must be set to zero.

6.9.16 READ CAPACITY Command, Opcode = 25_H

The READ CAPACITY command, shown in Table 6-61, allows the initiator to request information regarding the capacity of the drive. The drive reports the total number of logical blocks and the block length, in bytes.

Table 6-61 *READ CAPACITY Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =25 H							
1	LUN =0			RESERVED =0				reladr= 0
2-5	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
6-7	RESERVED =0							
8	RESERVED =0							
9	VU =0		RESERVED =0				F	L

PMI—PARTIAL MEDIUM INDICATOR (Byte 8, Bit 0)

When PMI is set to zero, the drive returns its last valid logical block address and the block length for that logical block in the read capacity data. The logical block address in the command descriptor block is zero.

When PMI is set to one, the drive returns the logical block address of the next logical block following which a substantial delay in data transfer may be encountered, as well as the block length for that logical block, in the read capacity data. This logical block address will be greater than or equal to the logical block address specified in the command descriptor block.

The eight bytes of READ CAPACITY data shown in Table 6-62 will be sent during the DATA IN phase of the command.

Table 6-62 *READ CAPACITY Command Data*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0-3	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
4-7	(MSB)	BLOCK LENGTH						(LSB)

6.9.17 READ EXTENDED Command, Opcode = 28_H

The READ EXTENDED command, shown in Table 6-63, requests that the drive transfer data to the initiator.

Table 6-63 *READ EXTENDED Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	opcode =28 h							
1	LUN =0			RESERVED =0				reladr= 0
2-5	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
6	RESERVED =0							
7-8	(MSB)	TRANSFER LENGTH						(LSB)
9	VU =0		RESERVED =0				F	L

LOGICAL BLOCK ADDRESS (Bytes 2-5)

This parameter specifies the logical block at which the read operation is to begin.

TRANSFER LENGTH (Bytes 7-8)

This parameter specifies the number of contiguous logical blocks of data to be transferred. A transfer length of zero indicates that no logical blocks will be transferred. Any other value indicates the specific number of logical blocks to be transferred.

6.9.18 WRITE EXTENDED Command, Opcode = $2A_H$

The WRITE EXTENDED command, shown in Table 6-64, requests that the drive write the data transferred by the initiator to the disk.

Table 6-64 *WRITE EXTENDED Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =2A H							
1	LUN =0			RESERVED =0				reladr=0
2-5	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
6	RESERVED =0							
7-8	(MSB)	TRANSFER LENGTH						(LSB)
9	VU =0			RESERVED =0			F	L

LOGICAL BLOCK ADDRESS (Bytes 2–5)

This parameter specifies the logical block at which the write operation is to begin.

TRANSFER LENGTH (Bytes 7–8)

This parameter specifies the number of contiguous logical blocks of data to be transferred. A transfer length of zero indicates that no logical blocks will be transferred and no data will be written to the disk. Any other value indicates the specific number of logical blocks to be transferred.

6.9.19 SEEK EXTENDED Command, Opcode= $2B_H$

The SEEK EXTENDED command, shown in Table 6-65, requests that the drive seek to the logical block address specified.

Table 6-65 *SEEK EXTENDED Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =2B H							
1	LUN =0			RESERVED =0				reladr=0
2-5	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
6-8	RESERVED =0							
9	VU=0			RESERVED =0			F	L

6.9.20 WRITE AND VERIFY Command, Opcode = $2E_H$

The WRITE AND VERIFY command, shown in Table 6-66, instructs the drive to write the data transferred from the initiator to disk, then verify that the data was written correctly.

Table 6-66 *WRITE AND VERIFY Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =2E H							
1	LUN =0			RESERVED =0			Bytchk	reladr=0
2-5	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
6	RESERVED =0							
7-8	(MSB)	TRANSFER LENGTH						(LSB)
9	VU =0		RESERVED =0				F	L

BYTCHK—BYTE CHECK (Byte 1, Bit 1)

This parameter must be set to zero. With BYTCHK set to zero, the drive supports only media verification that uses ECC. It does not support a byte-to-byte comparison of data on the disk with the data sent by the initiator, as with BYTCHK set to one.

If this bit is set to one, the command will be terminated with a CHECK CONDITION status, with the ILLEGAL REQUEST sense key.

LOGICAL BLOCK ADDRESS (Bytes 2–5)

This parameter specifies the logical block at which the write operation is to begin.

TRANSFER LENGTH (Bytes 7–8)

This parameter specifies the number of contiguous logical blocks of data to be transferred. A transfer length of zero indicates that no logical blocks will be transferred and no data will be written.

6.9.21 VERIFY Command, Opcode = $2F_H$

The VERIFY command, shown in Table 6-67, requests the drive to verify the data written on the disk.

Table 6-67 *VERIFY Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =2F H							
1	LUN =0		RESERVED =0				Bytchk	reladr
2-5	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
6	RESERVED =0							
7-8	(MSB)	VERIFICATIONLENGTH						(LSB)
9	VU =0		RESERVED =0				F	L

BYTCHK—BYTE CHECK (Byte 1, Bit 1)

This parameter must be set to zero. For an explanation of BYTCHK, see Section 6.9.20, "WRITE AND VERIFY Command." When this bit is set to one, the command will terminate in a CHECK CONDITION status, with the ILLEGAL REQUEST sense key.

LOGICAL BLOCK ADDRESS (Bytes 2–5)

This parameter specifies the logical block at which the verify operation is to begin.

VERIFICATION LENGTH (Bytes 7–8)

This parameter specifies the number of contiguous logical blocks of data the drive will verify. A transfer length of zero indicates that the drive will verify no logical blocks. Any other value indicates the specific number of logical blocks to be verified.

6.9.22 READ DEFECT DATA Command, Opcode = 37_H

The READ DEFECT DATA command, shown in Table 6-68, requests that the drive transfer the media-defect data to the initiator.

Table 6-68 *READ DEFECT DATA Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =37 H							
1	LUN =0			RESERVED =0				
2	RESERVED =0			P	G	DEFECT LIST FORMAT		
3-6	RESERVED =0							
7-8	(MSB)	ALLOCATION LENGTH						(LSB)
9	VU =0		RESERVED =0				F	L

P—PRIMARY (Byte 2, Bit 4)

When set to one, this parameter indicates that the initiator will request the return of the primary, or factory, defect list.

G—GROWN (Byte 2, Bit 3)

When set to one, this parameter indicates that the initiator will request the return of the field-found defect list.

Note: With bits P and G both set to one, the initiator requests the drive to return both the primary and field-found defect lists. If the P and G bits were both set to zero, the drive would return only the defect-list header bytes.

DEFECT LIST FORMAT (Byte 2, Bits 0–2)

This parameter specifies the preferred format for the defect list to be returned by the drive. If the initiator requests a format not supported by the drive, the drive will return the defect data in Physical Sector Format; then enter CHECK CONDITION status, with the RECOVERED ERROR sense key. The drive supports two defect-data formats:

- PHYSICAL SECTOR FROM INDEX FORMAT (101_B) —This format returns the physical cylinder, head, and sector location from index of the media defects.
- BYTES OFFSET FROM INDEX (100_B)— This format returns the physical cylinder, head, and the bytes offset from index of the defect location. The offset number reported represents the first byte of the defective sector.

ALLOCATION LENGTH (Bytes 7–8)

This parameter specifies the number of bytes the initiator has allocated for returned defect data. An allocation length of zero indicates that no defect data are to be transferred. Any other value indicates the maximum number of bytes to be transferred. The drive will terminate the DATA IN phase when the number of bytes specified by the allocation length or all available defect data have been transferred to the initiator, whichever is less.

As shown in Table 6-69, the defect-data list contains a four-byte header, followed by zero or more defect descriptors.

Table 6-69 *Defect List Header*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0							
1	RESERVED =0			P	G	DEFECT LIST FORMAT		
2-3	(MSB)	DEFECT LIST LENGTH						(LSB)

P AND G (Byte 1, Bits 3–4)

These parameters indicate the defect list—primary or field-found—that the drive returns to the initiator.

DEFECT LIST FORMAT (Byte 1, Bits 0–2)

This parameter specifies the actual format of the defect list—Physical Sector Format or Bytes Offset From Index Format. These defect-descriptor formats are shown in Tables 6-70 and 6-71, respectively.

DEFECT LIST LENGTH (Bytes 2–3)

This parameter specifies the total length, in bytes, of the defect descriptors that follow. The defect list length is equal to eight times the number of defect descriptors. If the allocation length of the CDB is insufficient to allow the transfer of all of the defect descriptors, the defect list length will not be adjusted to reflect their truncation.

Table 6-70 *Physical Sector Format*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0 - 2	(MSB)			CYLINDER NUMBER				(LSB)
13				HEAD NUMBER				
4 - 7	(MSB)			SECTOR NUMBER				(LSB)

Table 6-71 *Bytes Offset From Index Format*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0 - 2	(MSB)			CYLINDER NUMBER				(LSB)
13				HEAD NUMBER				
4 - 7	(MSB)			BYTES OFFSET				(LSB)

6.9.23 WRITE BUFFER Command, Opcode = $3B_H$

The WRITE BUFFER command, shown in Table 6-72, is used in conjunction with the READ BUFFER command to test the target memory and integrity of the SCSI bus. This command will not alter the media.

Table 6-72 *WRITE BUFFER Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =3B H							
1	LUN =0			RESERVED =0		MODE		
2	BUFFER ID =0							
3-5	RESERVED =0							
6-8	(MSB)			TRANSFER LENGTH				(LSB)
9	CONTROL BYTE							

MODE (Byte 1, Bits 0–2)

Table 6-73 *Mode Field — WRITE BUFFER Command*

BIT2	BIT1	BIT0	MODE
0	0	0	Write combined header and data
0	0	1	Vendor-unique
0	1	0	Write data
0	1	1	Reserved
1	0	0	Not supported
1	0	1	Diskware Load
1	1	0	Reserved
1	1	1	Reserved

BUFFER ID (Byte 2)

This parameter must be set to zero. Any other buffer ID will result in a check condition status switch sense key Illegal Request (05_H).

TRANSFER LENGTH (Bytes 6–8)

This parameter specifies the maximum number of bytes to be transferred during the DATA OUT phase.

The function of the WRITE BUFFER command and the information in the transfer-length field depend on the information in the mode field. Table 6-73 defines the mode field.

6.9.23.1 COMBINED HEADER AND DATA Mode (000_B)

In this mode, a four-byte header that contains all reserved bytes precedes the data to be transferred. The specified transfer length (bytes 6–8) must include both header and data bytes. The amount of data to be stored in the drive's buffer equals the transfer length minus four. The initiator should ensure that the transfer length is not greater than four plus the available-length field returned in the header of the READ BUFFER command (mode 000_B). If the transfer length exceeds this amount, the drive will return CHECK CONDITION status, with the sense key ILLEGAL REQUEST (05_H).

6.9.23.2 VENDOR UNIQUE Mode (001_B)

Not used by the Maverick 270/540S hard disk drive.

6.9.23.3 DATA Mode (010_B)

In this mode, the DATA OUT phase contains buffer data. The transfer length (bytes 6–8) specifies the maximum number of bytes that the drive will transfer during the DATA OUT phase to be stored in the buffer. The initiator should ensure that the transfer length does not exceed the capacity of the buffer. The buffer capacity can be determined by the available-length field in the READ BUFFER header. If the transfer-length field specifies a transfer that would exceed the buffer capacity, the drive will return CHECK CONDITION status, with the sense key ILLEGAL REQUEST (05_H).

6.9.24 READ BUFFER Command, Opcode = $3C_H$

Used in conjunction with the WRITE BUFFER command, the READ BUFFER command, shown in Table 6-74, tests the memory in the drive and the integrity of the SCSI bus. This command does not alter the media.

Table 6-74 *READ BUFFER Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =3C H							
1	LUN =0			RESERVED =0		MODE		
2	BUFFER ID =0							
3-5	RESERVED =0							
6-8	(MSB)	TRANSFER LENGTH						(LSB)
9	CONTROL BYTE							

BUFFER ID (Byte 2)

This parameter is not used.

BUFFER OFFSET (Bytes 3–5)

This parameter specifies the address offset (from location zero) into the buffer.

allocation length (Bytes 6–8)

This parameter specifies the maximum number of bytes the initiator has allocated for returned data. An allocation length of zero indicates that the drive will transfer no header or data bytes. Any other value indicates the maximum number of bytes to be transferred.

The function of the READ BUFFER command and the information in the allocation-length field depend on the information in the mode field. Table 6-75 defines the mode field.

Table 6-75 *Mode Field — READ BUFFER Command*

BIT2	BIT1	BIT0	MODE
0	0	0	Read combined header and data
0	0	1	Vendor-unique
0	1	0	Data
0	1	1	Descriptor
1	0	0	Reserved
1	0	1	Reserved
1	1	0	Reserved
1	1	1	Reserved

6.9.24.1 COMBINED HEADER AND DATA Mode (000_B)

In this mode, the drive returns header and data bytes to the initiator during the DATA IN phase. The four-byte READ BUFFER header, shown in Table 6-76, is followed by data bytes from the drive data buffer. The specified allocation length (bytes 6–8) must include both header and data bytes. The drive terminates the DATA IN phase when it has transferred the allocated number of bytes—header plus data—or all available header and data bytes to the initiator, whichever occurs first.

Table 6-76 *READ BUFFER Command Header*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED =0							
1-3	(MSB)	AVAILABLE LENGTH						(LSB)

AVAILABLE LENGTH (Bytes 1–3)

This parameter specifies the total number of data bytes available in the drive's data buffer. This number is reduced neither to reflect the allocation length, nor the actual number of bytes written by the WRITE BUFFER command. Following the READ BUFFER header, the drive will transfer data from its data buffer. The number of data bytes transferred following the READ BUFFER header will equal the allocation length minus four.

6.9.24.2 VENDOR-UNIQUE Mode (001_B)

Not used by the Maverick 270/540S hard disk drive.

6.9.24.3 DATA Mode (010_B)

In this mode, the DATA IN phase contains only buffer data. The drive terminates the DATA IN phase when it has transferred the allocated number of bytes of data or all data to the initiator, whichever occurs first.

6.9.25 READ LONG Command, Opcode = $3E_H$

The READ LONG command, shown in Table 6-77, requests that the drive transfer data to the initiator. For each sector read, the drive will transfer 512 data bytes, followed by two cross check bytes and 12 ECC bytes.

Table 6-77 *READ LONG Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =3E H							
1	LUN =0			RESERVED =0			correct=0	reladr=0
2-5	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
6	RESERVED =0							
7-8	(MSB)	BYTE TRANSFER LENGTH						(LSB)
9	RESERVED =0						F	L

CORRCT

This parameter is always set to zero, indicating that the drive will read the logical block, without any correction, setting this bit to one, the command will terminate in CHECK CONDITION status, with Illegal Request Sense Key.

LOGICAL BLOCK ADDRESS

This parameter specifies the logical block at which the READ LONG operation will begin.

BYTE TRANSFER LENGTH

This parameter specifies the exact number of bytes to be transferred. Since read long only accepts single-sector requests, this length must be set at 526 (512 data bytes plus 14 ECC information bytes). If this value does not match the available data length, the drive will terminate the command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and an additional sense code set to INVALID FIELD IN CDB.

6.9.26 WRITE LONG Command, Opcode = $3F_H$

The WRITE LONG command, shown in Table 6-78, requests that the drive write data transferred by the initiator in an area specified by the logical block address. The data transferred includes 512 data bytes, followed by two cross check bytes and twelve ECC bytes for each sector. The WRITE LONG command is usually preceded by a READ LONG command. The number and order of the bytes in the WRITE LONG command should be the same as those in the READ LONG command.

Table 6-78 *WRITE LONG Command*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	OPCODE =3F H							
1	LUN =0			RESERVED =0				reladr=0
2-5	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
6	RESERVED =0							
7-8	(MSB)	BYTE TRANSFER LENGTH						(LSB)
9	RESERVED =0						F	L

LOGICAL BLOCK ADDRESS

This parameter specifies the logical block at which the WRITE LONG operation begins.

BYTE TRANSFER LENGTH

This parameter specifies the exact number of bytes to be transferred. If this value does not match the available data length, the drive will terminate the command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and an additional sense code set to INVALID FIELD IN CDB. A transfer length of zero indicates that no logical blocks will be transferred. This length must be set to 526, same as READ LONG.

6.10 ELECTRICAL CHARACTERISTICS OF THE SCSI BUS

The following description of the SCSI electrical characteristics includes an definition of the signal lines as well as the timing characteristics.

6.10.1 Signal Lines

There are a total of 18 signal lines on the SCSI bus, of which nine are control lines and nine are data lines (including the parity line). All SCSI bus signals on the Maverick 270/540S are actively driven true, and are negated by releasing the drivers. When the drivers are released, the bias circuitry of the bus terminators pulls the signals to a false state. Table 6-79 defines the SCSI bus signal lines.

Table 6-79 *SCSI Bus Signal Lines*

SIGNAL LINE	USE
BSY	An OR-tied line used to indicate that the bus is in use by a device.
SEL	An OR-tied line used by an initiator to select a target device. Used by a target to reselect an initiator.
C/D	Indicates whether Control or Data information is on the bus. Driven by the target.
I/O	Indicates direction of data transfer with respect to the initiator. Driven by the target. This signal is also asserted by the target during Reselection.
MSG	Used to indicate a message is being transferred. Driven by the target.
REQ	Used by the target to set the Request portion of the bus handshake.
ACK	Used by the initiator to set the Acknowledge portion of the bus handshake.
ATN	Used by the initiator to set the Attention condition.
RST	An OR-tied line used to indicate the RESET condition has been set.
DB7–DB0	The 8-bit parallel data bus. DB7 has the highest priority during arbitration.
P	The parity bit is also considered as part of the data bus. Parity is odd. It is not valid during arbitration.

Table 6-80 *Signal Sources*

BUS PHASE	SIGNALS				
	BSY	**SEL**	**C/D, I/O, REQ, MSG**	**ACK, ATN**	**DB (7-0, P)**
Bus Free	None	None	None	None	None
Arbitration	All	Winner	None	None	SCSI ID
Selection	I & T	Initiator	None	Initiator	Initiator
Reselection	I & T	Target	Target	Initiator	Target
Command	Target	None	Target	Initiator	Target

BUS PHASE	SIGNALS				
	BSY	**SEL**	**C/D, I/O, REQ, MSG**	**ACK, ATN**	**DB (7-0, P)**
Data In	Target	None	Target	Initiator	Target
Data Out	Target	None	Target	Initiator	Initiator
Status	Target	None	Target	Initiator	Target
Message In	Target	None	Target	Initiator	Target
Message Out	Target	None	Target	Initiator	Initiator

Note: **All:** All SCSI devices that are actively arbitrating drive the signal.

SCSI ID: Any SCSI device that is actively arbitrating can drive this unique data bit. The arbitrating SCSI device can release, but not drive, the other seven data bits. The DP(P) parity bit can be released or driven to the true state, but can never be driven to the false state during an ARBITRATION phase.

I & T: Initiator and Target The initiator, target, or both drive the signal—as specified during a SELECTION or RESELECTION phase.

INITIATOR: Only the active initiator can drive the signal. The Maverick 270/540S hard disk drive is never an initiator.

NONE: The signal is released—that is, not driven by any SCSI device. The bias circuitry of the bus terminators pulls the signal to the false state.

WINNER: The SCSI device that wins arbitration drives the signal.

TARGET: Only the active target can drive the signal.

Table 6-81 defines the conductor signal assignments for the Quantum DB50 SCSI connector.

Table 6-81 *System Cable Pin-Out*

SCSI Bus Pin	Signal	SCSI Bus Pin	Signal
2	DB0	28	Ground
4	DB1	30	Ground
6	DB2	32	ATN
8	DB3	34	Ground
10	DB4	36	BSY
12	DB5	38	ACK
14	DB6	40	RST
16	DB7	42	MSG
18	DBP	44	SEL
20	Ground	46	C/D
22	Ground	48	REQ
24	Ground	50	I/O
26	Terminator Power		

Note: All odd pins, except pin 25, are grounded.

6.10.2 Timing Characteristics

Timing characteristics for the SCSI bus are given in the following sections.

6.10.2.1 Arbitration Delay

This period is the minimum amount of time that must elapse between the point at which the device asserts BSY to try and get the bus and the point at which it examines DB0–DB7 (to determine if it has won arbitration).

Arbitration delay is 2.4 microseconds.

6.10.2.2 assertion period

This period is the minimum amount of time that REQ (target) and ACK (initiator) must be asserted during synchronous bus transfers.

assertion period is 90 nanoseconds.

6.10.2.3 Bus Clear Delay

This period is the maximum amount of time that may elapse prior to the device releasing all signal lines when:

1. The Bus Free phase is detected.
2. An arbitration attempt is lost (another device asserts SEL).
3. RST is detected on the bus.

Bus clear delay following a lost arbitration attempt or detection of RST on the bus is 800 nanoseconds.

6.10.2.4 Bus Free Delay

This period is the minimum amount of time that must elapse between the point at which the device detects the Bus Free phase and the point at which it asserts BSY in an attempt to get the bus.

Bus free delay is 800 nanoseconds.

6.10.2.5 Bus Set Delay

This period is the maximum amount of time that may elapse between the point at which the device detects the Bus Free phase and the point at which it asserts BSY and its SCSI ID bit in an attempt to get the bus.

Bus set delay is 1.8 microseconds.

6.10.2.6 Bus Settle Delay

This period is the minimum amount of time that must elapse after the device detects a signal line transition on certain control signals before the transition is valid.

For instance, when BSY and SEL both go false, the device must wait bus settle delay prior to "detecting" the Bus Free phase.

Bus settle delay is 400 nanoseconds.

6.10.2.7 Cable Skew Delay

This period is the maximum difference in propagation delay between any two signal lines measured between any two SCSI devices resident on the same bus.

Cable skew delay is 10 nanoseconds.

6.10.2.8 Data Release Delay

This period is the maximum amount of time that may elapse between the point at which I/O transits from false to true and the point at which the initiator releases DB0–DB7.

Data release delay is 400 nanoseconds.

6.10.2.9 Deskew Delay

This period is the minimum amount of time required to deskew signal lines.

Deskew delay is 45 nanoseconds.

6.10.2.10 Hold Time

This period is the minimum amount of time that must be added between the assertion of REQ or ACK and changing the data lines during synchronous bus transfers.

Hold time is 45 nanoseconds.

6.10.2.11 negation period

This period is the minimum amount of time that REQ (target) and ACK (initiator) may be negated during synchronous bus transfers.

negation period is 90 nanoseconds.

6.10.2.12 RESET Hold Time

This period is the minimum amount of time that RST must be asserted to be valid.

RESET hold time is 25 microseconds.

6.10.2.13 Selection Abort Time

This period is the maximum amount of time that may elapse between the point at which the device determines that it is being selected and the point at which it asserts BSY in response.

Selection abort time is 200 microseconds.

6.10.2.14 Selection Time-out Delay

This period is the minimum amount of time that should elapse between the point at which the device attempts to select a target and the point at which it begins a time-out procedure, provided the target device does not respond with BSY.

Selection time-out delay should be approximately 250 milliseconds.

6.10.2.15 Transfer Period

This period is the minimum amount of time that must elapse between the leading edges of successive REQ or ACK pulses during synchronous bus transfers.

Transfer period is set during a Message phase.

6.10.3 Fast SCSI Timing characteristics

Timing characteristics for the Fast SCSI option are given in the following sections.

6.10.3.1 fast assertion period

This period is the minimum amount of time that REQ (target) and ACK (initiator) must be asserted during synchronous bus transfers.

fast assertion period is 30 nanoseconds.

6.10.3.2 fast cable skew delay

This period is the maximum difference in propagation delay between any two signal lines measured between any two SCSI devices resident on the same bus.

fast cable skew delay is 5 nanoseconds.

6.10.3.3 fast deskew delay

This period is the minimum amount of time required to deskew signal lines.

fast deskew delay is 20 nanoseconds.

6.10.3.4 fast hold time

This period is the minimum amount of time that must be added between the assertion of REQ or ACK and changing the data lines during synchronous bus transfers.

fast hold time is 10 nanoseconds.

6.10.3.5 fast negation period

This period is the minimum amount of time that REQ (target) and ACK (initiator) may be negated during synchronous bus transfers.

fast negation period is 30 nanoseconds.

Glossary

A

ACCESS – (v) Read, write, or update information stored on a disk or other medium. (n) The operation of reading, writing, or updating stored information.

ACCESS TIME – The interval between the time a request is made by the system and the time the data is available from the drive. Includes the seek time, rotational latency, and command processing overhead time. (See also *seek, rotational latency,* and *overhead.*)

ACTUATOR – Also known as the *positioner.* The internal mechanism that moves the read/write head to the proper track. The Quantum actuator consists of a rotor connected to head mounting arms that position the heads over the desired cylinder. Also known as rotary actuator.

AIRLOCK – A patented Quantum feature that ensures durable and reliable data storage. Upon removal of power from the drive for any reason, the read/write heads automatically park and lock in a non data area called the landing zone. AIRLOCK allows the drive to withstand high levels of non-operating shock. When power is applied to the drive, airflow created from the spinning disks causes the AIRLOCK arm to swing back and unlock the actuator, allowing the heads to move from the landing zone. Upon power down, the AIRLOCK swings back to the locked position, locking the heads in the landing zone. A park utility is not required to park the heads on drives equipped with AIRLOCK (all Quantum drives).

ALLOCATION – The process of assigning particular areas of the disk to specific data or instructions. An allocation unit is a group of sectors on the disk reserved for specified information. On hard disks for small computer systems, the allocation unit is usually in the form of a sector, block, or cluster. (See also *allocation unit.*)

ALLOCATION UNIT – An allocation unit, also known as a *cluster,* is a group of sectors on the disk that can be reserved for the use of a particular file.

ASIC – Acronym for *Application Specific Integrated Circuit.*

AVERAGE SEEK TIME – The average time it takes for the read/write head to move to a specific location. Calculated by dividing the time it takes to complete a large number of random seeks by the number of seeks performed.

B

BACKUP – A copy of a file, directory, or volume on a separate storage device from the original, for the purpose of retrieval in case the original is accidentally erased, damaged, or destroyed.

BAD BLOCK – A block (usually the size of a sector) that cannot reliably hold data due to a physical flaw or damaged format markings.

BAD TRACK TABLE – A label affixed to the casing of a hard disk drive stating which tracks are flawed and cannot hold data. This list is typed into the low-level formatting program when the drive is installed. Quantum users can ignore bad track tables since Quantum's built-in defect-management protections compensate for these flaws automatically.

BEZEL – A plastic panel that extends the face of a drive so that it covers a computer's drive bay opening. The bezel usually contains a drive-activity LED. Also known as the *faceplate.*

BIT – Abbreviation for binary digit. A binary digit may have one of two values—1 or 0. This contrasts with a decimal digit, which may have a value from 0 to 9. A bit is one of the logic 1 or logic 0 binary settings that make up a byte of data. (See also *byte.*)

BLOCK – In UNIX workstation environments, the smallest contiguous area that can be allocated for the storage of data. UNIX blocks are generally 8 Kbytes (16 sectors) in size. In DOS environments, the block is referred to as a cluster. (Note: This usage of the term block at the operating system level is different from its meaning in relation to the physical configuration of the hard drive. See *sector* for comparison.) While an operating system may have its block size set at 8K, the drive usually assigns its "logical block" size to 1 sector (1 OS block = 16 drive locical blocks).

BPI – Bits Per Inch. A measure of how densely information is packed on a storage medium. (See also *FCI.*)

BUFFER – An area of RAM reserved for temporary storage of data that is waiting to be sent to a device that is not yet ready to receive it. The data is usually on its way to or from the hard disk drive or some other peripheral device.

BUS – The part of a chip, circuit board, or interface designed to send and receive data.

BYTE – The basic unit of computer memory, large enough to hold one character of alphanumeric data. Comprised of eight bits. (See also *bit.*)

C

CACHE – High-speed RAM used as a buffer between the CPU and the hard disk. Since the CPU can get information more quickly from the cache than from main memory, the cache usually contains information that is used frequently by the system.

CAPACITY – The amount of information that can be stored on a hard drive. Also known as storage capacity. (See also *formatted capacity.*)

CLEAN ROOM – An environmentally controlled dust-free assembly or repair facility in which hard disk drives are assembled or can be opened for internal servicing.

CLUSTER – In DOS environments, the smallest contiguous area that can be allocated for the storage of data. DOS clusters are usually 2 Kbytes (4 sectors) in size.

CONTROLLER – The chip or circuit that translates computer data and commands into a form suitable for use by the hard drive. Also known as disk controller.

CONTROLLER CARD – An adapter containing the control electronics for one or more hard disks. Usually installed in a slot in the computer.

CPU – Central Processing Unit. The microprocessor chip that performs the bulk of data processing in a computer.

CRC – Cyclic Redundancy Check. An error detection procedure that identifies incomplete or faulty data in each sector.

CYLINDER – When disks are placed directly above one another along the shaft, the circular, vertical "slice" consisting of all the tracks located in a particular position.

D

DATA SEPARATOR – The circuit that extracts data from timing information on drives that store a combined data and clock signal.

DEDICATED SERVO – A positioning mechanism using a dedicated surface of the disk that contains timing and positioning information only, as compared to surfaces that are also used for data. (See also *embedded servo.*)

DEFECT MANAGEMENT – A technique ensuring long-term data integrity. Consists of scanning disk drives both at the factory and during regular use, de-allocating defective sectors before purchase and compensating for new defective sectors afterward.

DISK – In general, any circular-shaped data-storage medium that stores data on the flat surface of the platter. The most common type of disk is the magnetic disk, which stores data as magnetic patterns in a metal or metal-oxide coating. Magnetic disks come in two forms: floppy and hard. Optical recording is a newer disk technology that gives higher capacity storage but at slower access times.

DISK CONTROLLER – A plug-in board, or embedded circuitry on the drive, that passes information to and from the disk. The Quantum hard disk drives all have controllers embedded on the drive printed-circuit board. (See also *controller.*)

DMA – Direct Memory Access. A process for transferring data directly to and from main memory, without passing through the CPU. DMA improves the speed and efficiency by allowing the system to continue processing even while new data is being retrieved.

DOS – Disk Operating System. The most common operating system used in IBM PCs. Manages all access to data on the disk.

DRIVE – Short form of *disk drive.*

DRIVE GEOMETRY – The functional dimensions of a drive, including the number of heads, cylinders, and sectors per track. (See also *logical format.*)

E

ECC – Error Correction Code. The incorporation of extra parity bits in transmitted data in order to detect errors that can be corrected by the controller.

EMBEDDED SERVO – A timing or location signal placed on tracks that store data. These signals allow the actuator to fine-tune the position of the read/write heads.

ENCODING – The conversion of data into a pattern of On/Off or 1/0 signals prior to being written on the disk surface. (See also *RLL* and *MFM.*)

EPROM – Erasable Programmable Read-Only Memory. An integrated circuit memory chip that can store programs and data in a non-volatile state. These devices can be erased by ultraviolet light and reprogrammed with new data.

EXTERNAL DRIVE – A drive mounted in an enclosure separate from the computer system enclosure, with its own power supply and fan, and connected to the system by a cable.

F

FCI – Flux Changes per Inch. The number of magnetic field patterns that can be stored on a given area of disk surface, used as a measure of data density. (See also *BPI.*)

FILE SERVER – A computer that provides network stations with controlled access to shareable resources. The network operating system is loaded on the file server, and most shareable devices (disk subsystems, printers) are attached to it. The file server controls system security and monitors station-to-station communications. A dedicated file server can be used only as a file server while it is on the network. A non dedicated file server can be used simultaneously as a file server and a workstation.

FIRMWARE – Permanent instructions and data programmed directly into the circuitry of read-only memory for controlling the operation of the computer. Distinct from software, which can be altered by programmers.

FLUX DENSITY – The number of magnetic field patterns that can be stored in a given length of disk surface. The number is usually stated as flux changes per inch (FCI), with typical values in the thousands. (See also *FCI.*)

FLYING HEIGHT – The distance between the read/write head and the disk surface, made up of a cushion of air that keeps the two objects from touching. Smaller flying heights permit denser data storage but require more precise mechanical designs. Also known as fly height.

FORMAT – To write a magnetic track pattern onto a disk surface, specifying the locations of the tracks and sectors. This information must exist on a disk before it can store data.

FORMATTED CAPACITY – The amount of room left to store data on a disk after writing sector headers, boundary definitions, and timing information during a format operation. The size of a Quantum drive is always expressed in formatted capacity, accurately reflecting the usable space required.

FORM FACTOR – The industry standard that defines the physical, external dimensions of a particular device. For example, most Quantum hard disk drives use a 3 1/2-inch form factor.

G

GIGABYTE (GB) – One billion bytes (one thousand megabytes).

GUIDE RAILS – Plastic strips attached to the sides of a hard disk drive in an IBM PC/AT or compatible computer so that the drive easily slides into place.

H

HALF-HEIGHT – Standard drive size equivalent to half the vertical space of a 5 1/4-inch drive.

HARD DISK – A type of storage medium that retains data as magnetic patterns on a rigid disk, usually made of an iron oxide or alloy over a magnesium or aluminum platter. Because hard disks spin more rapidly than floppy disks, and the head flies closer to the disk, hard disks can transfer data faster and store more in the same volume.

HARD ERROR – A data error that persists when the disk is re-read, usually caused by defects in the physical surface.

HARD-SECTORED – The most common method of indicating the start of each sector on a disk, based on information located in the embedded servo. This method is more precise than soft-sectored techniques and results in lower overhead. (See also *soft-sectored.*)

HEAD – The tiny electromagnetic coil and metal pole used to create and read back magnetic patterns on the disk. Also known as read/write head.

HEAD CRASH – Damage to the read/write head, usually caused by sudden contact with the disk surface. Head crash can also be caused by dust and other particles.

HIGH-CAPACITY DRIVE – By industry conventions typically a drive of 100 megabytes or more.

HIGH-LEVEL FORMATTING – Formatting performed by the operating system to create the root directory, file allocation tables and other basic configurations. (See also *low-level formatting.*)

HOME – Reference track used for recalibration of the actuator. Usually the outermost track (track 0).

HOST ADAPTER – A plug-in board that acts as the interface between a computer system bus and the disk drive.

I

INITIALIZATION – See *low-level formatting.*

INTERFACE – A hardware or software protocol, (contained in the electronics of the disk controller and disk drive) that manages the exchange of data between the drive and computer. The most common interfaces for small computer systems are AT (also known as IDE) and SCSI.

INTERLEAVE – The arrangement of sectors on a track. The Interleave Factor is the number of sectors that pass beneath the read/write heads before the next sector arrives. For example, a 3:1 interleave factor means that the heads read a sector, then let two pass by before reading another, requiring three full revolutions of the disk to access the complete data track. Quantum drives have an interleave factor of 1:1, allowing the system to access a full track of data in a single revolution.

INTERLEAVE FACTOR – The number of sectors that pass beneath the read/write heads before the next numbered sector arrives. When the interleave factor is 3:1, a sector is read, two pass by, and then the next is read. It would take three revolutions of the disk to access a full track of data. Quantum drives have an interleave of 1:1, so a full track of data can be accessed within one revolution of the disk, thus offering the highest data throughput possible.

INTERNAL DRIVE – A drive mounted inside one of a computer's drive bays, or a hard disk on a card installed in one of the computer's expansion slots.

J

JUMPER – A tiny box that slips over two pins on a circuit board, connecting the pins electrically. Some board manufacturers use Dual In-Line Package (DIP) switches instead of jumpers.

K

KILOBYTE (K) – A unit of measure consisting of 1,024 (2^{10}) bytes.

L

LANDING ZONE – A non-data area on the disk's inner cylinder where the heads can rest when the power is off.

LATENCY – The time during which the read/write heads wait for the data to rotate into position after the controller starts looking for a particular data track. If a disk rotates at 3,600 rpm, the maximum latency time is 16.4 milliseconds, and the average latency time is 8.2 milliseconds.

LOGICAL BLOCK - See Block.

LOOK AHEAD – The process of anticipating events in order to speed up computer operations. For example, the system can buffer data into cache RAM by reading blocks in advance, preparing the system for the next data request.

LOW-LEVEL FORMATTING – The process of creating sectors on the disk surface so that the operating system can access the required areas for generating the file structure. Quantum drives are shipped with the low-level formatting already completed. Also known as *initialization.*

LOW PROFILE SERIES (LPS)– Describes drives built to the 3 1/2-inch form factor, which are only 1 inch high. The standard form factor drives are 1.625 inches high.

M

MB – See *megabyte.*

MEDIA – The magnetic film that is deposited or coated on an aluminum substrate which is very flat and in the shape of a disk. The media is overcoated with a lubricant to prevent damage to the heads or media during head take off and landing. The media is where the data is stored inside the disk in the form of magnetic flux or polarity changes.

MEGABYTE (MB) – Quantum defines megabyte as 10^6 bytes or 1,000,000 bytes.

MEGAHERTZ – A measurement of frequency in millions of cycles per second.

MHz – See *megahertz.*

MICROPROCESSOR – The integrated circuit chip that performs the bulk of data processing and controls the operation of all of the parts of the system. A disk drive also contains a microprocessor to handle all of the internal functions of the drive and to support the embedded controller.

MICROSECOND (µs) – One millionth of a second (.000001 sec.).

MILLISECOND (ms) – One thousandth of a second (.001 sec.).

MTBF – Mean Time Between Failure. Reliability rating indicating the failure rate expected of a product expressed in power on hours (POH). Since manufacturers differ in the ways they determine the MTBF, comparisons of products should always take into account the MTBF calculation method.

MTTR – Mean Time To Repair. The average time it takes to repair a drive that has failed for some reason. This only takes into consideration the changing of the major sub-assemblies such as circuit board or sealed housing. Component level repair is not included in this number as this type of repair is not performed in the field.

O

OVERHEAD – *Command overhead* refers to the processing time required by the controller, host adapter, or drive prior to the execution of a command. Lower command overhead yields higher drive performance. (See also *zero command overhead.*) *Disk overhead* refers to the space required for non-data information such as location and timing. Disk overhead often accounts for about ten percent of drive capacity. Lower disk overhead yields greater disk capacity.

OVERWRITE – To write data on top of existing data, erasing it.

OXIDE – A metal-oxygen compound. Most magnetic coatings are combinations of iron or other metal oxides, and the term has become a general one for the magnetic coating on tape or disk.

P

PARTITION – A portion of a hard disk dedicated to a particular operating system and application and accessed as a single logical volume.

PERFORMANCE – A measure of the speed of the drive during normal operation. Factors affecting performance are seek times, transfer rate and command overhead.

PERIPHERAL – A device added to a system as an enhancement to the basic CPU, such as a disk drive, tape drive or printer.

PHYSICAL FORMAT – The actual physical layout of cylinders, tracks, and sectors on a disk drive.

PLATED MEDIA – Disks that are covered with a hard metal alloy instead of an iron-oxide compound. Plated disks can store greater amounts of data than their oxide-coated counterparts.

PLATTER – Common term referring to the hard disk.

POH – Power On Hours. The unit of measurement for Mean Time Between Failure as expressed in the number of hours that power is applied to the device regardless of the amount of actual data transfer usage. (See also *MTBF.*)

POSITIONER – See *actuator.*

R

RAM – Random Access Memory. An integrated circuit memory chip that allows information to be stored and retrieved by a microprocessor or controller. The information may be stored and retrieved in any order, and all storage locations are equally accessible.

RAM DISK – A "phantom" disk drive created by setting aside a section of RAM as if it were a group of regular sectors. Access to RAM disk data is extremely fast, but is lost when the system is reset or turned off.

READ AFTER WRITE – A mode of operation requiring that the system read each sector after data is written, checking that the data read back is the same as the data recorded. This operation lowers system speed but raises data reliability.

READ VERIFY – A data accuracy check performed by having the disk read data to the controller, which then checks for errors but does not pass the data on to the system.

READ/WRITE HEAD – The tiny electromagnetic coil and metal pole piece used to create and read back the magnetic patterns (write or read information) on the disk. Each side of each platter has its own read/write head.

REMOVABLE DISK – Generally said of disk drives where the disk itself is meant to be removed, and in particular of hard disks using disks mounted in cartridges. Their advantage is that multiple disks can be used to increase the amount of stored material, and that once removed, the disk can be stored away to prevent unauthorized use.

RLL – Run Length Limited. A method of encoding data into magnetic pulses. The RLL technique permits 50% more data per disk than the MFM method, but requires additional processing.

ROM – Read-Only Memory. Integrated circuit memory chip containing programs that can be accessed and read but can not be modified.

ROTARY ACTUATOR – The rotary actuator replaces the stepper motor used in the past by many hard disk manufacturers. The rotary actuator is perfectly balanced and rotates around a single pivot point. It allows closed-loop feedback positioning of the heads, which is more accurate than stepper motors.

ROTATIONAL LATENCY – The delay between when the controller starts looking for a specific block of data on a track and when that block rotates around to where it can be read by the read/write head. On average, it is half of the time needed for a full rotation (about 8 ms.).

S

SCSI – Small Computer System Interface. An interface designed for Apple Macintosh systems and UNIX workstations.

SECTOR – On a PC hard drive, the minimum segment of track length that can be assigned to store information. On Macintosh and UNIX drives, sectors are usually grouped into blocks or logical blocks that function as the smallest data unit permitted. Since these blocks are often defined as a single sector the terms block and sector are sometimes used interchangeably in this context. (Note: The usage of the term block in connection with the physical configuration of the disk is different from its meaning at the system level. See also *block* and *cluster* for comparison.)

SEEK – A movement of the disk read/write head to a specific data track.

SERVO DATA – Magnetic markings written on the media that guide the read/write heads to the proper position.

SERVO SURFACE – A separate surface containing only positioning and disk timing information but no data.

SETTLE TIME – The interval between the arrival of the read/write head at a specific track, and the lessening of the residual movement to a level sufficient for reliable reading or writing.

SHOCK RATING – A rating, expressed in "G's", of how much shock a disk drive can sustain without damage.

SOFT ERROR – A faulty data reading that does not recur if the same data is reread from the disk, or corrected by ECC. Usually caused by power fluctuations or noise spikes.

SOFT-SECTORED – Old time-based method of indicating the start of each sector on a disk. Soft-sectored drives require that location instructions be located in the data fields. (See also *hard-sectored.*)

SPINDLE – The drive's center shaft, on which the hard disks are mounted. A synchronized spindle is a shaft that allows two disks to spin simultaneously as a mirror image of each other, permitting redundant storage of data.

SPUTTER – A special method of coating the disk that results in a hard, smooth surface capable of storing data at a high density. Quantum disk drives use sputtered thin film disks.

STEPPER – A type of motor that moves in discrete with each electrical pulse. Stepper were originally the most common type of actuator engine, since they can be geared to advance a read/write head one track per step. However, they are not as fast, reliable, or durable as the voice coil actuators found in Quantum disk drives. (See also *voice coil.*)

SUBSTRATE – The material underneath the magnetic coating of a disk. Common substrates include aluminum or magnesium alloys for hard drives, glass, for optical disks, and mylar for floppy disks.

SURFACE – The top or bottom side of a disk, which is coated with the magnetic material for recording data. On some drives one surface may be reserved for positioning information.

T

THIN FILM – A type of coating allowing very thin layers of magnetic material, used on hard disks and read/write heads. Hard disks with thin film surfaces can store greater amounts of data.

TPI – Tracks Per Inch. The number of tracks written within each inch of disk's surface, used as a measure of how closely the tracks are packed on a disk surface. Also known as *track density.*

TRACK – One of the many concentric magnetic circle patterns written on a disk surface as a guide for storing and reading data. Also known as *channel.*

TRACK DENSITY – How closely the tracks are packed on a disk surface. The number is specified as tracks per inch (TPI).

TRACK-TO-TRACK SEEK TIME – The time required for the read/write heads to move to an adjacent track.

TRANSFER RATE – The rate at which the disk sends and receives data from the controller. The sustained transfer rate includes the time required for system processing, head switches and seeks, and accurately reflects the drive's true performance. The burst mode transfer rate is a much higher figure that refers only to the movement of data directly into RAM.

U

UNFORMATTED CAPACITY – The total number of usable bytes on a disk, including the space that will be required to later to record location, boundary definitions, and timing information. (See *formatted capacity* for comparison.)

V

VOICE COIL – A fast and reliable actuator motor that works like a loud speaker, with the force of a magnetic coil causing a proportionate movement of the head. Voice coil actuators are more durable than their stepper counterparts, since fewer parts are subject to daily stress and wear. Voice coil technology is used in all Quantum drives.

W

WEDGE SERVO – The position on every track that contains data used by the closed loop positioning control. This information is used to fine tune the position of the read/write heads exactly over the track center.

WINCHESTER DISKS – Former code name for an early IBM hard disk model, sometimes still used to refer to hard drives in general.

WRITE ONCE – An optical disk technology that allows the drive to store and read back data, but prevents the drive from erasing information once it has been written.

INDEX

A

B

C

D

E

F

G

H

I

Q

R

S

T

U

V

W

X

Quantum®

QUICK FAX - READER COMMENT FORM

PLEASE FAX YOUR COMMENTS AND SUGGESTIONS FOR IMPROVING THIS PUBLICATION TO:
QUANTUM TECHNICAL PUBLICATIONS (408) 894-6159

Manual Part Number	Manual Title	Current Date
81-1006845-02	Maverick 270/540S Product Manual	

How did you use this publication?

☐ Reference ☐ Sales

☐ Installation ☐ Maintenance

What is your occupation?

Do you wish a written reply?

☐ Yes ☐ No

What is your overall rating of this publication?

☐ Very Good ☐ Fair

☐ Good ☐ Poor

Is the material presented effectively?

☐ Well organized ☐ Fully covered ☐ Clear

☐ Correct ☐ Well Illustrated

Locating the information you were looking for was?

☐ Easy ☐ Difficult

Please enter your comments and suggestions for improvement below. If you were in any way dissatisfied with this publication, we would like to know why. Be specific, if possible; give page, column, and paragraph number references where applicable.

All comments and suggestions become the property of Quantum Corporation.

Your name and FAX NUMBER:

Thank you for your interest. We will consider all comments and will incorporate changes when appropriate.

Quantum®
500 McCarthy Blvd.
Milpitas, CA 95035
(408) 894-4000