Quantum®

Empire 540/1080S Product Manual

February 1994
81-103983-02

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

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

Publication Number: 81-103983-02

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

These products are covered by or licensed under one or more of the following U.S. Patents: 4,419,701; 4, 538,193 4,625,109; 4,639,798; 4,647,769; 4,647,997; 4,661,696; 4,669,004; 4,675,652; 4,703,176; 4,730,321; 4,772,974; 4,783,705; 4,819,153; 4,882,671; 4,920,442; 4,920,434; 4,982,296; 5,005,089; 5,027,241; 5,031,061; 5,084,791; 5,119,254; 5,160,865; 5,170,229; 5,177,771; 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 you might find helpful.

1.1 AUDIENCE DEFINITION

The Empire 540/1080S 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
fci	Flux changes per inch
Hz	Hertz
KB	Kilobytes
lsb	Least-significant bit
M	Meters
mA	Milliamperes
MB	Megabytes
Mbit/s	Megabits per second

MB/s	Megabytes per second
MHz	Megahertz
ms	Milliseconds
msb	Most-significant bit
mv	Millivolts
ns	Nanoseconds
tpi	Tracks per inch
µs	Microseconds
V	Volts

The following conventions are used in this manual:

- **Computer Voice**

  Computer voice refers to items you type at the computer keyboard. These items are listed in all capitals in courier font. For example:

  `FORMAT C:/S`

- **Commands and Messages**

  Firmware commands and messages sent from the drive to 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**

  Hexadecimal notation is shown by a lowercase "h." For example:

  30h.

- **Signal Negation**

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

  RD–

- **Host**

  In general, the system in which the drive resides is referred to as the host.

1.4 REFERENCE

For additional information about the SCSI interface, refer to the ANSI Small Computer System Interface-2 (SCSI-2) Specification, ANSI X3T9.2/86-109, Revision 10h, October 17, 1991.

Chapter 2
GENERAL DESCRIPTION

This chapter summarizes general functions and key features of the Empire 540/1080S hard disk drive, as well as the standards and regulations it meets.

2.1 PRODUCT OVERVIEW

Quantum's Empire 540/1080S 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. ProDrive LPS™ hard disk drives use nonremovable, 3 1/2-inch hard disks.

Empire 540/1080S hard disk drives feature an embedded Small Computer System Interface (SCSI) drive controller and use SCSI commands. The drive manages media defects and error recovery internally, so these operations are transparent to the user.

2.2 KEY FEATURES

Empire 540/1080S hard disk drives include the following key features:

- Formatted storage capacities of 540 or 1080 MB
- Industry-standard, 3 1/2-inch form factor
- Low-profile, 1-inch height
- SCSI-2 Compliance
- Embedded fast SCSI controller
- Wide SCSI (optional)
- Embedded servo design
- Data transfer rate of up to 5.0 MB/s asynchronous and 10.0 MB/s synchronous
- Average seek time of 10 ms
- Average rotational latency of 5.6 ms
- Proprietary 512K "look ahead" DisCache® with continuous prefetch and WriteCache™ write-buffering capabilities
- 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
- 500,000 hour Mean Time Between Failure (MTBF)
- Automatic retry on read errors
- Transparent media-defect mapping

- ECC on-the-fly
- 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
- 1:1 interleave on read/write operations

2.3 STANDARDS AND REGULATIONS

Empire 540/1080S 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 (UL): 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

Empire 540/1080S hard disk drives are compatible with host computers and controllers that provide a SCSI or SCSI-2 interface.

Chapter 3
INSTALLATION

This chapter explains how to unpack, configure, mount, and connect the Empire 540/1080S hard disk drive prior to operation. It also explains how to start up and operate the drive.

3.1 SPACE REQUIREMENTS

Figure 3-1 shows the external dimensions of the drive.

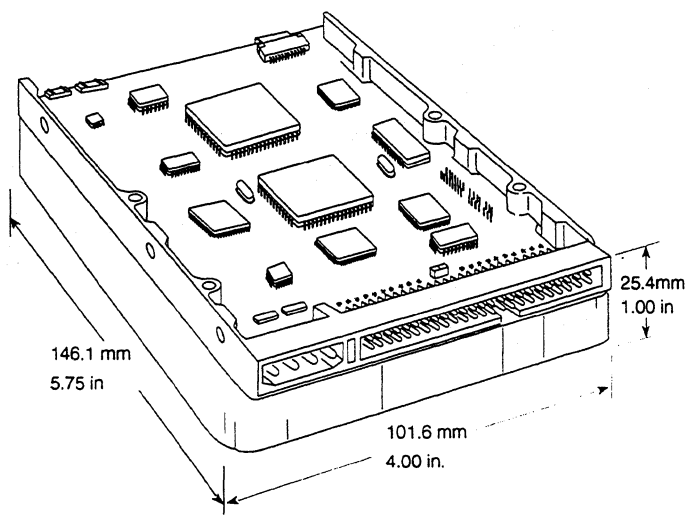

Figure 3-1 *Mechanical Dimensions*

3.2 UNPACKING INSTRUCTIONS

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

   Figure 3-2 shows the packaging for the Empire 540/1080S hard disk drive in a 1-pack shipping container. (A 12-pack shipping container is available for multiple drive shipments and has a similar type of protective packaging.)

3. Remove the drive from the box.

CAUTION:	During shipment and handling, the antistatic electrostatic discharge (ESD) bag prevents 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 packing materials for possible future use.

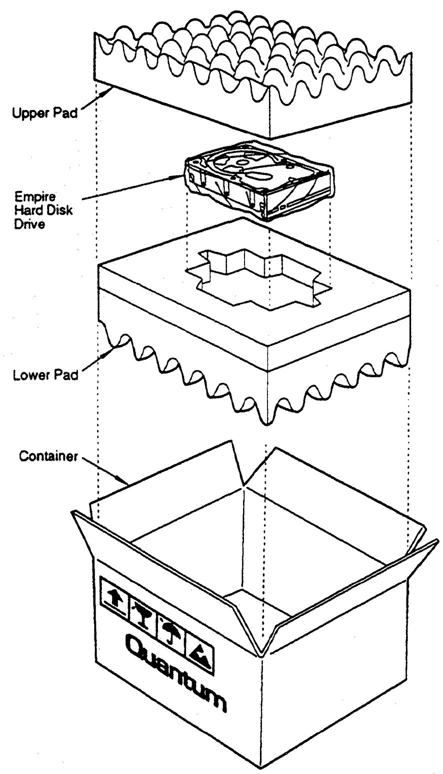

Figure 3-2 *Packing Assembly*

3.3 JUMPER SETTINGS

The configuration of a Empire 540/1080S hard disk drive depends on the host system into which it is to be installed. This section describes the hardware options that you must take into account prior to installation. Figure 3-3 shows the printed circuit board assembly (PCBA), indicating the jumpers that control some of these options.

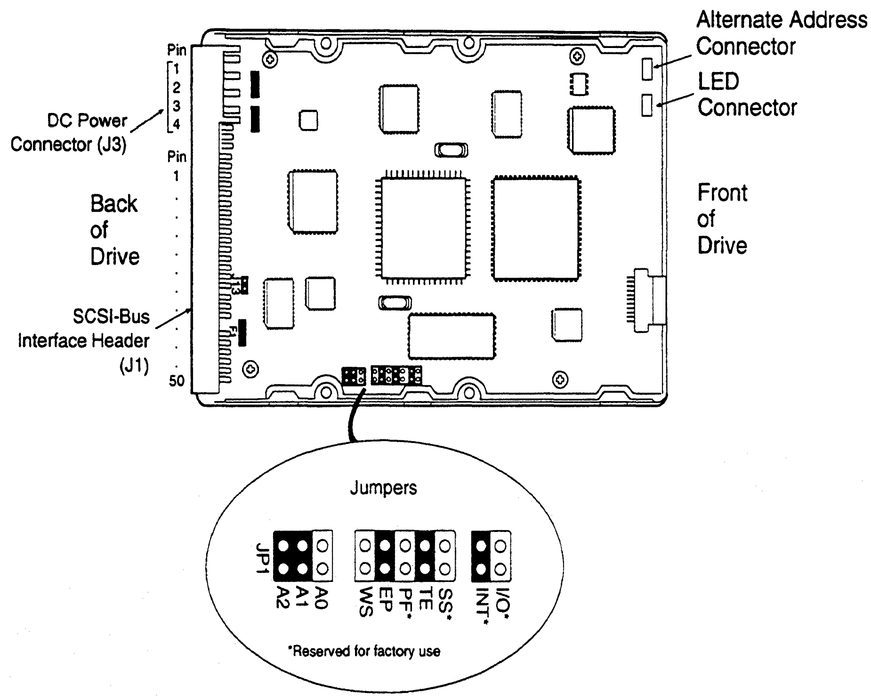

Figure 3-3 *Jumper Pin and Terminator Locations on the Drive PCBA*

On the Drive's PCBA, there are eight pairs of jumper pins, usually called *jumpers*, that control the drive's operation. (In addition, numerous options controlled by SCSI commands affect the operation of the drive.) Four of these jumper pins, PF, SS, INT, and I/O, are reserved for factory use. However, you might need to change some of the other six jumpers to configure the drive properly. These user-configurable jumpers include:

- TE—Termination Enable Jumper
- EP—Enable Parity Jumper
- WS—Wait/Spin Jumper
- A2, A1, A0—SCSI-Bus Identification Jumpers
- J13—SCSI Termination Power Jumper

When you change any jumper setting but J13, the new setting is in effect the next time the drive is powered on.Table 3-1 shows the default and optional settings for the drive. J13 immediately affects the SCSI terminator power source.

Table 3-1 *Jumper Settings*

JUMPER	SETTING	DESCRIPTION
TE	OFF ON*	Termination disabled Termination enabled
EP	OFF ON*	Parity checking disabled Parity checking enabled
WS	OFF* ON	Wait/spin disabled Wait/spin enabled
A2, A1, A0	—	These jumpers are used in combination to determine the drive's SCSI ID. (See Table 3-2.)
J13	ON* OFF	SCSI terminator power supplied by the host SCSI terminator power supplied by the drive

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

3.3.1 Termination Enable (TE) Jumper

One of the requirements of the SCSI bus is that it be terminated with proper termination resistors. If more than one hard disk drive is connected to the host, the last one must be terminated. The Termination Enable (TE) jumper, when installed, connects terminating resistors to the SCSI bus. With the TE jumper off, no resistors are connected.

At the factory, Quantum configures the Empire 540/1080S hard disk drive with termination enabled; that is, with a jumper installed across the pins labeled TE. To disable termination, remove the jumper from the TE pins.

3.3.2 Enable Parity (EP) Jumper

The Enable Parity (EP) jumper controls parity checking of data across the SCSI bus. With parity checking enabled, the drive performs parity checking. With parity checking disabled, the drive continues to generate parity information, but does not perform parity checking.

At the factory, Quantum configures the Empire 540/1080S hard disk drive with parity checking enabled; that is, with a jumper installed across the pins labeled EP. To disable parity checking, remove the jumper from the EP pins.

3.3.3 Wait/Spin (WS) Jumper

With the Wait/Spin (WS) jumper disabled, at a power on or reset, the drive motor starts to spin automatically, and the drive prepares itself to perform read or write operations without the need for a start command. At a power on or reset, with the Wait/Spin jumper enabled, the host must send a START/STOP UNIT command to activate the drive motor for read or write operations. In addition, with the W/S jumper enabled, the drive performs power sequencing to prevent an overload of the host's power supply due to simultaneous, peak start-up current demands from multiple devices.

At the factory, Quantum configures the Empire 540/1080S hard disk drive with the Wait/Spin jumper disabled; that is, without a jumper installed across the pins labeled WS. To enable the Wait/Spin jumper, install a jumper across the WS pins.

3.3.4 SCSI-Bus Device Identification (A2, A1, and A0) Jumpers

Used in combination, the jumper settings across pins A2, A1, and A0 determine the Empire 540/1080S hard disk drive's SCSI-bus device identification (SCSI ID). By default, Quantum configures the drive with a SCSI ID of 6; that is, with jumpers installed across the pins labeled A2 and A1, and no jumper installed across the pins labeled A0.

To assign a different SCSI ID, change the jumper settings as shown in Table 3-2.

Table 3-2 *Jumper Settings for SCSI ID*

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: An 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 MOUNTING

Drive mounting, orientation, clearance, and ventilation requirements are described in the subsections that follow. 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.

3.4.1 Orientation

The mounting holes on the Empire 540/1080S hard disk drive allow the drive to be mounted in any orientation. Figure 3-4 shows the location of the three mounting holes on each side of the Empire 540/1080S hard disk drive.

± 5° of plane?

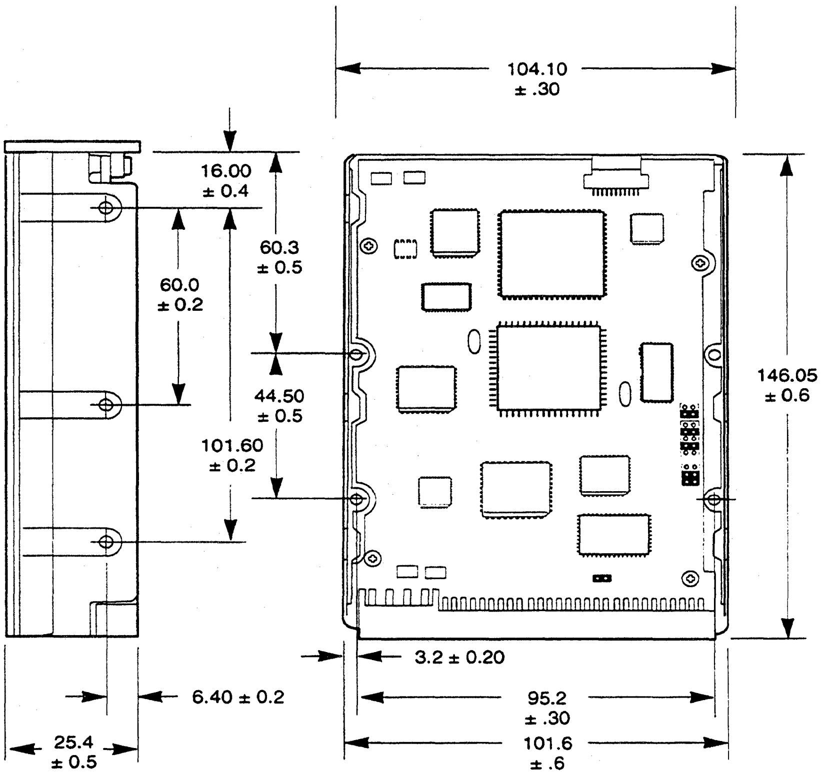

Figure 3-4 *Mounting Dimensions (in Millimeters)*

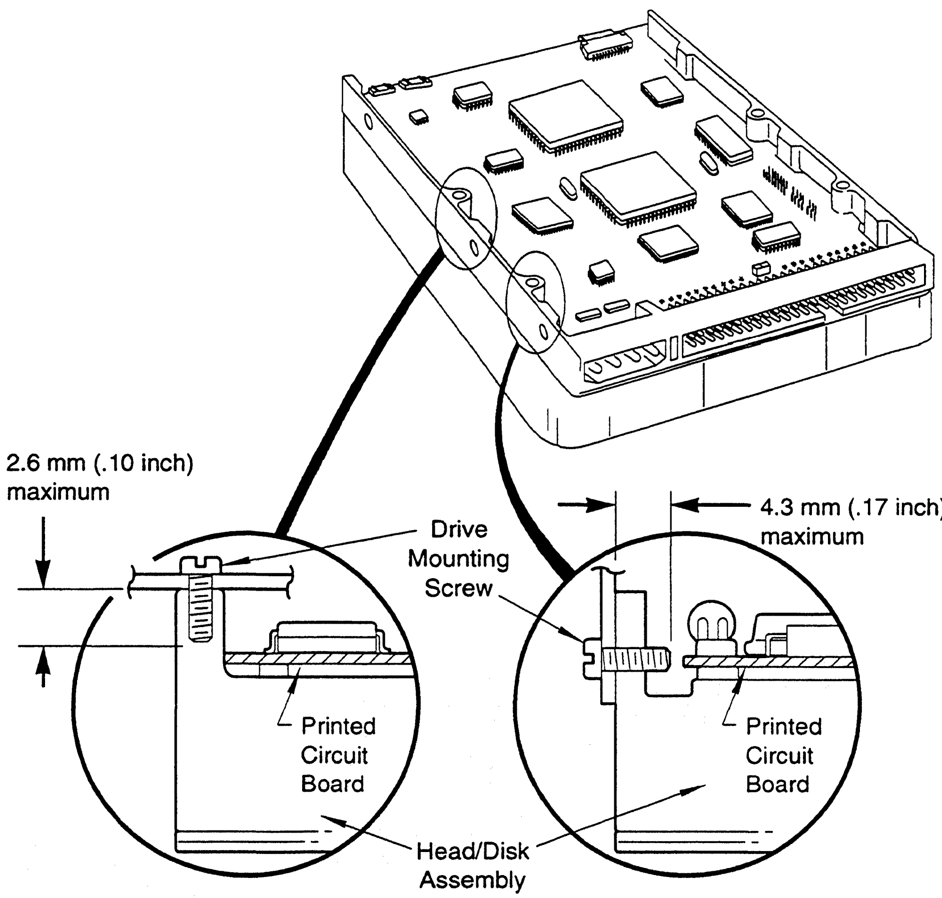

Figure 3-5 *Mounting Screw Clearance*

CAUTION:	Use of mounting screws longer than the maximum lengths specified in Figure 3-5 voids the warranty on the Empire 540/1080S.

3.4.2 Clearance

The PCBA is very close to the mounting holes. Clearance from the drive to any other surface—except mounting surfaces—must be 2.6 millimeters (0.10 inches) minimum. Figure 3-5 specifies the clearance between the screws in the mounting holes and the PCBA. Do not use mounting screws longer than the maximum lengths shown. The specified screw length allows full use of the mounting-hole threads, while avoiding damaging or placing unwanted stress on the PCBA.

3.4.3 Ventilation

The Empire 540/1080S hard disk drive operates without a cooling fan, provided the ambient air temperature does not exceed 131°F (55°C).

where measured.

3.5 CABLE CONNECTORS

The cable connectors for the Empire 540/1080S are mounted on the back edge of the PCB, as shown in Figure 3-6.

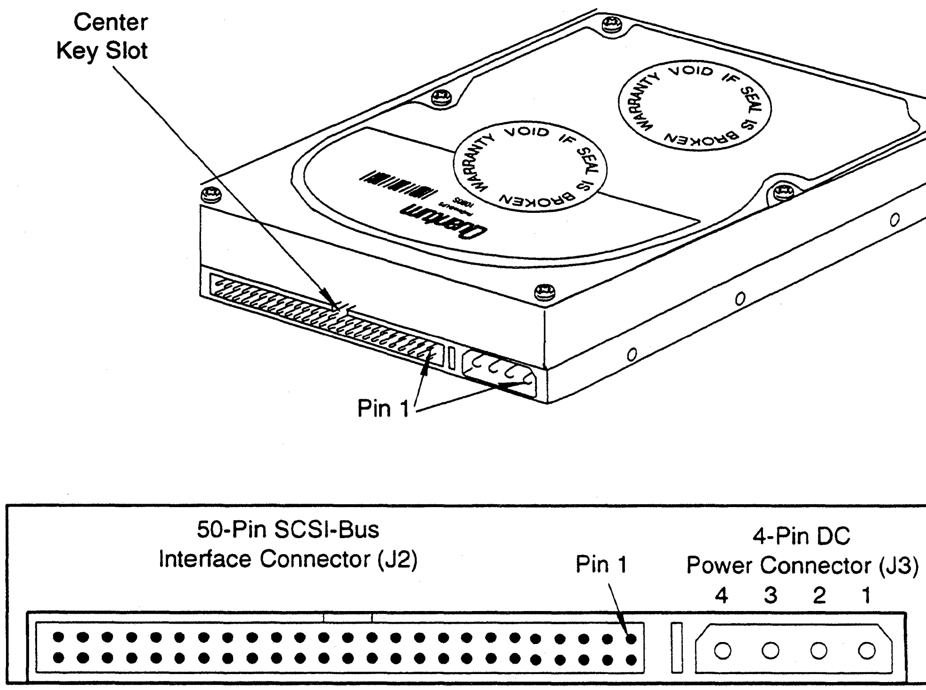

Figure 3-6 *Cable Connectors*

3.5.1 SCSI Connector (J1)

The SCSI cable connector (J1) on the Empire 540/1080S hard disk drive is a 50-pin universal header, as shown in Figure 3-6. 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.

The recommended header connector and strain relief for the SCSI connector include:

- 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

3.5.2 DC Power Connector (J3)

The drive's DC power connector (J3) is mounted on the back edge of the printed circuit board (PCB). The recommended mating connectors and pins for the DC power connector include:

- 4-pin connector housing (P3) AMP P/N 1-480424-0 or equivalent
- Loose-piece contacts AMP P/N VS 60619-4 or equivalent (20–14 ga.)
- Strip contacts AMP P/N VS 61117-4 or equivalent (20–14 ga.)

Table 3-3 shows the power line designations for the pins on the DC power connector (J3).

Table 3-3 *J3 Pin Assignments*

PIN NUMBER	POWER LINE DESIGNATION
1	+12V DC
2	+12V Return (Ground)
3	+5V Return (Ground)
4	+5V DC

Note: Labels indicate the pin numbers on the connector. Pins 2 and 3 (ground returns) are connected on the drive PCB.

Chapter 4
SPECIFICATIONS

This chapter gives a detailed description of the physical, electrical, and environmental characteristics of the Empire 540/1080S hard disk drive.

4.1 SPECIFICATION SUMMARY

Table 4-1 *Specifications*

DESCRIPTION	VALUE	
	540S	1080S
Formatted capacity (bytes)	540,000,256	1,080,000,512
Number of 512-byte user sectors	1,054,688	2,109,376
Nominal rotation speed (rpm)	5,400 ± 0.3%	5,400 ± 0.3%
Maximum recording density (bpi)	61,300	61,300
Maximum flux density (fci)	46,000	46,000
Track density (tpi)	3,014	3,014
Data tracks per disk	2,874	2,874
Data sectors per track	65 to 107	65 to 107
Data tracks per cylinder	4	8
User data cylinders	2,874	2,874
Spare sectors per cylinder	2	4
Read/write heads	4	8
Disks	2	4
Encoding scheme	RLL 1,7	RLL 1,7

4.2 FORMATTED CAPACITY

At the factory, the Empire 540/1080S receives a low-level format that creates the actual tracks and sectors on the drive. Table 4-1 shows the capacity resulting from this process. Formatting done at the user level for operation with DOS, UNIX, or other operating systems results in less capacity than the physical capacity shown in the table.

4.3 DATA TRANSFER RATES

Data is transferred from the read buffer to the SCSI bus at a rate of up to 5.0 MB/s in the asynchronous mode and 10.0 MB/s in the synchronous mode.

4.4 TIMING SPECIFICATIONS

Table 4-2 lists the timing specifications of the Empire 540/1080S hard disk drive.

Table 4-2 *Timing Specifications*

PARAMETERS	TYPICAL NOMINAL[4]	MAXIMUM NOMINAL[4]	MAXIMUM WORST CASE[5]
Single-Track Seek[1]	1.8 ms	3 ms	TBD
Average Seek[2]	10 ms	11 ms	TBD
1/3-Stroke Seek[1]	11 ms	13 ms	TBD
Full Stroke Seek[1]	19ms	22 ms	TBD
Sequential Head-Switch Time[3]	TBD	TBD	3 ms
Average Write	11ms	12ms	TBD
Average Rotational Latency	5.6 ms	5.6 ms	5.6 ms

Notes:

1 Seek time is defined as the time required for the actuator to seek and settle on-track. It is measured by averaging 5,000 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.

2 Average seek time is the average of 5,000 random seeks. When a seek error occurs, recovery for that seek can take up to seven seconds.

3 Sequential head-switch time is the time required for the head to move from the end of the last sector on a track to the beginning of the next sequential sector, which is on the same track in the same cylinder. Track skewing determines the sequential head-switch time.

4 Typical nominal conditions are as follows:
- Nominal temperature (25° C)
- Nominal supply voltages (12.0V, 5.0V)
- No applied shock or vibration

5 Worst-case conditions are as follows:
- Worst-case temperature extremes (4°C to 50°C)
- Worst-case supply voltages (12V +10% or 12V –7%, 5V ±5%)

4.5 POWER MODE TIMES

Table 4-3 shows the timing for completing the various power modes.

Table 4-3 *Power Mode Times*

PARAMETERS	TYPICAL NOMINAL[2]	MAXIMUM WORST CASE[3]
Power On to Drive Ready[1]	12 sec	20 sec
Shutdown to Drive Ready	18 sec	25 sec

Notes: [1] In the Drive Ready state, the disks are spinning at the rated speed, and the drive is able to accept and execute commands requiring disk access.

[2] Typical nominal conditions are as follows:
- Nominal temperature (25° C)
- Nominal supply voltages (12V, 5V)
- No applied shock or vibration

[3] Worst-case conditions are as follows:
- Worst-case temperature (50° C)
- Worst-case supply voltages (12V +10% or 12V –7%, 5V ±5%)

4.6 POWER

The Empire 540/1080SS hard disk drive operates from two supply voltages:

- +12V +7%
- +5V ± 5%

The allowable ripple and noise for each voltage is 200 mV peak-to-peak for the +12 Vdc supply and 100 mV peak-to-peak for the +5 Vdc supply.

4.6.1 Power Sequencing

You can apply the power in any order or manner or you can short or open either the power or power-return line with no loss of data or damage to the disk drive. However, data might 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.

4.6.2 Power Reset Limits

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

Table 4-4 *Power Reset Limits*

DC VOLTAGE	THRESHOLD
+5 V	4.4 V to 4.7 V
+12 V	9.6 V to 10.4 V

4.6.3 Power Requirements

Table 4-5 lists the voltages and corresponding current for the various modes of operation of the Empire 540/1080S hard disk drive.

Table 4-5 *Typical Power and Current Consumption*

MODE OF OPERATION	TYPICAL AVERAGE CURRENT[4] (mA avg)		TYPICAL AVERAGE POWER[5] (WATTS)	MAXIMUM POWER (WATTS)
	+12 V	+5 V		
Startup	1,500 peak	500 peak	–	–
Idle[1]	330	410	6.0 W	7.0 W
Random Read/Write[2]	420	500	7.5 W	8.5 W
Random Seek[6]	420	450	7.3 W	8.3 W
Noise and Ripple[3]	200 mv	100 mv	–	–

Notes:

1 Idle mode is in effect when the drive is not reading, writing, seeking, or executing any commands. A portion of the R/W circuitry is powered down, the motor is up to speed, and the Drive Ready condition exists. The actuator resides on the last track accessed.

2 Read/write mode is defined as:
- 30% random seeks
- 40% read/write (1 write plus 10 reads)
- 30% ready (idle)

3 Maximum (peak-to-peak)

4 Current is rms (except for Startup)

5 Reflects nominal values for +12V and +5V power supplies and does not include power required for the SCSI termination resistors.

6 30% seek duty cycle

4.7 ACOUSTICS

Table 4-6 specifies the acoustical characteristics of the Empire 540/1080S hard disk drive.

Table 4-6 *Acoustical Characteristics*

OPERATING MODE	MEASURED NOISE	DISTANCE
Idle On Track	35 dbA (mean) 38 dbA (maximum)	39.4 in (1 m)
Seeking (Random)	TBD	39.4 in (1 m)

4.8 MECHANICAL CHARACTERISTICS

Table 4-7 specifies the mechanical dimensions of the Empire 540/1080S hard disk drive.

Table 4-7 *Mechanical Dimensions*

CHARACTERISTIC	US	METRIC
Height	1.00 in	25.4 mm
Width	4.00 in	101.6 mm
Depth	5.75 in	146.1 mm
Weight	1.25 lb	0.55 kg

4.9 ENVIRONMENTAL CONDITIONS

Table 4-8 summarizes the environmental specifications for the Empire 540/1080S hard disk drive.

Table 4-8 *Environmental Specifications*

PARAMETER	OPERATING	NONOPERATING
Temperature	0° to 55°C (32° to 131°F)	–40° to 75°C (–40° to 167°F)
Temperature Gradient	24°C/hr maximum	48°C/hr maximum
Humidity[1]	5% to 95% rh	5% to 95% rh
Maximum Wet Bulb	40°C (104°F)	50°C (122°F)
Altitude[2]	–61 m to 3 km (–200 to 10,000 ft.)	–300 m to 12 km (–1,000 to 40,000 ft.)

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

4.10 SHOCK AND VIBRATION

The Empire 540/1080S 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 4-9. 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 12-pack shipping container, Empire 540/1080S drives can withstand a drop from 30 inches onto a concrete surface on any of its surfaces, six edges, or three corners.

Table 4-9 *Shock and Vibration Specifications*

	OPERATING	NONOPERATING
Shock 1/2 sine wave, 11 ms duration (10 hits maximum)	10 G	60 G
Vibration 5-400 Hz sine wave (zero-to-peak) 1.0 octave/minute sweep	0.5 G	1.0 G

Note: No unrecoverable errors allowed.

4.11 RELIABILITY

Table 4-10 summarizes the reliability measurements for the Empire 540/1080S hard disk drive.

Table 4-10 *Reliability Measurements*

RELIABILITY FACTOR	MEASUREMENT
Mean Time Between Failures (MTBF)	500,000 Power On Hours (POH), typical usage
Preventive Maintenance (PM) Required	None
Mean Time To Repair (MTTR)	30 minutes
Start/Stop Cycles	Minimum 20,000 cycles

Note: The Quantum MTBF numbers represent an estimation based on field data, vendor data, and Bell-Core MTBF predictions.

4.12 DISK ERRORS

Table 4-11 provides the error rates for the Empire 540/1080S hard disk drive.

Table 4-11 *Error Rates*

ERROR TYPE	MAXIMUM NUMBER OF ERRORS
Recovered read errors[1]	10 errors per 10^{11} bits read
Reallocated data errors:	
Correctable read errors[2]	10 errors per 10^{13} bits read
Uncorrectable read errors[3]	10 errors per 10^{15} bits read
Transferred errors[4]	11 errors per 10^{25} bits read
Seek Errors[5]	10 errors per 10^{7} seeks

Notes:

1 Recovered read errors are corrected by retries. The drive does not automatically reallocate the sectors.

2 Correctable read errors are read errors that are not recovered by retries but are recovered by application of the double-burst error correction algorithm.

3 Uncorrectable read errors are errors that are not correctable using ECC or retries. The drive terminates retry reads either when a repeating error pattern occurs, or after eight unsuccessful retries and the application of double-burst error correction.

4 Transferred errors are errors that are not detected and subsequently not corrected by the drive.

5 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 explains the basic operation of the Empire 540/1080S hard disk drive, which consists of various mechanical assemblies, drive electronics, and firmware features.

5.1 DRIVE MECHANISM

The Empire 540/1080S hard disk drive conforms to the industry-standard, 3 1/2-inch form factor. It consists of the head/disk assembly (HDA) and a printed-circuit board (PCB). The HDA contains the mechanical subassemblies of the drive, which are sealed under a metal cover. The HDA consists of the following components:

- Base casting assembly
- Motor assembly
- Disk stack assembly
- Headstack assembly
- Rotary positioner assembly
- Automatic actuator lock
- Air filter

The drive is assembled in a Class-100 clean room.

CAUTION: These subassemblies are neither adjustable nor field-repairable. To ensure the air in the HDA remains free of contamination, never remove or adjust its cover and seals. Tampering with the HDA voids the warranty.

In the Empire 540S, the HDA contains 2 magnetic disks and 4 read/write heads. In the Empire 1080S, the HDA contains 4 magnetic disks and 8 read/write heads. Figure 5-1 shows an exploded view of the Empire 1080S. The Empire 540S is the same, except it has 2 disks instead of 4.

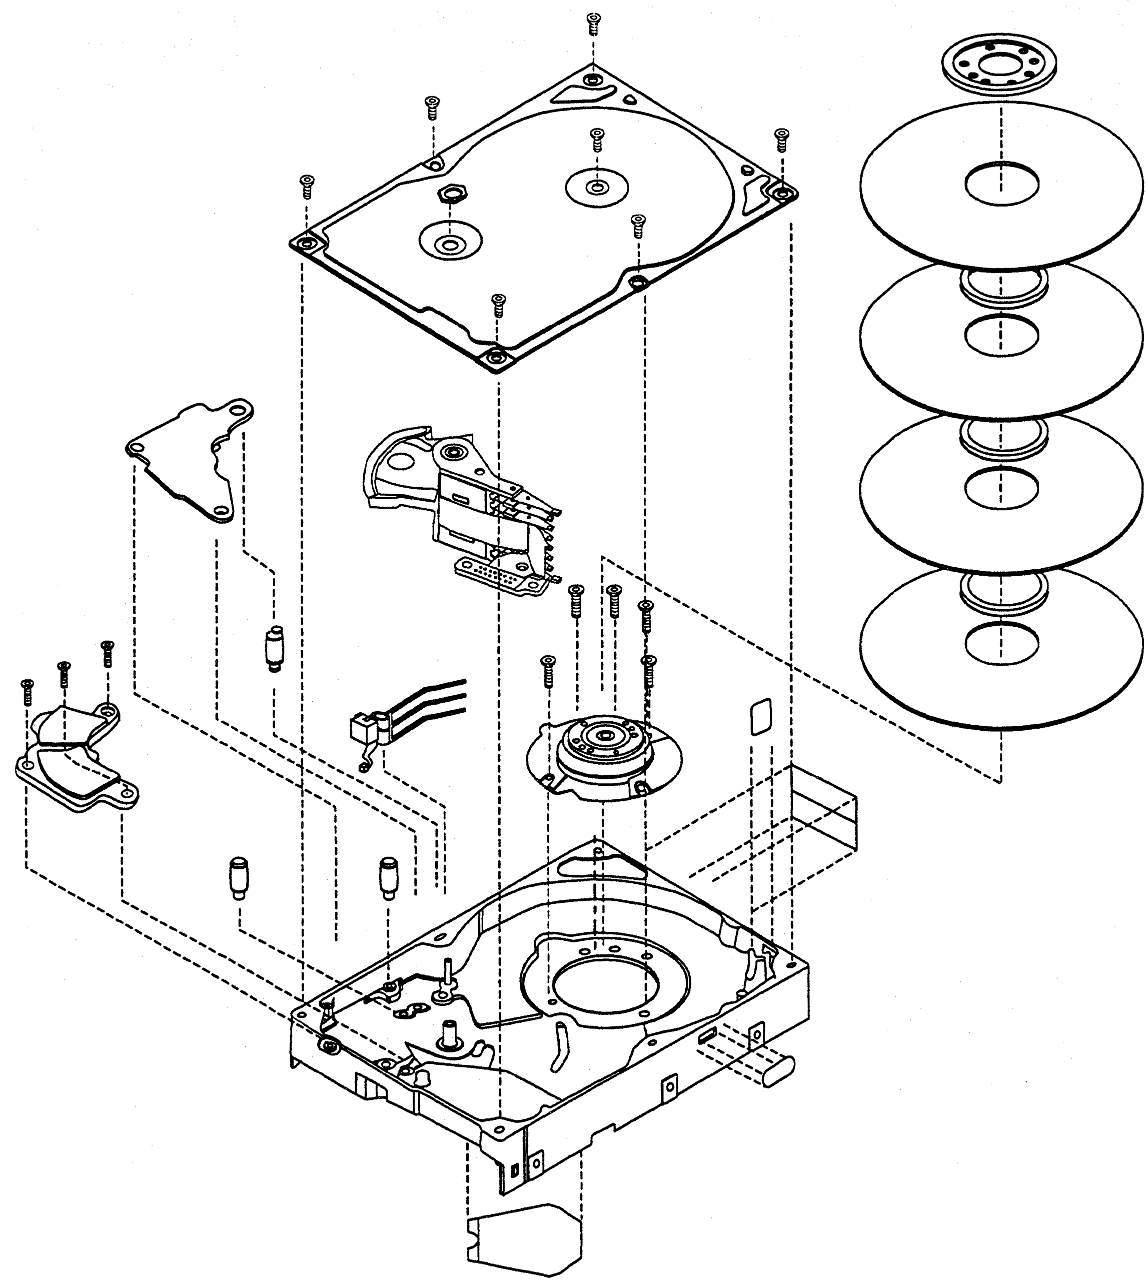

Figure 5-1 *Empire 1080S Exploded View*

5.1.1 Base Casting Assembly

A single-piece, aluminum-alloy base casting provides a mounting surface for the drive mechanism and PCB. To ensure a contamination-free environment for the HDA, a gasket provides a seal between the base casting and the metal cover that encloses the drive mechanism. An E-coat seals the surface of the base casting to provide additional contamination protection.

5.1.2 Motor Assembly

Assembled to the base casting, the DC motor assembly consists of a fixed-shaft DC spindle motor that drives the counterclockwise rotation of the disks (as viewed from the cover), the spindle/bearing assembly, and disk mounting hub. One conductive ferrofluid seal and one labyrinth prevent outside air and contaminants from entering the HDA at the bearing core or along the bearing shaft. The ferrofluid seals also serve as a conductive ground.

5.1.3 Disk Stack Assembly

The disk stack assembly consists of 2 disks in the Empire 540S and 4 disks in the Empire 1080S. Disk spacers are 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 the data area, when the disks are not rotating at full speed.

Table 5-1 summarizes the use of each disk surface.

5.1.3.1 Track Formatting

The Empire 540/1080S hard disk drive is shipped from the factory with all of the physical sector addresses prewritten on the surface of the disk. As a result, the number of bytes per sector and the number of sectors per track are not field-selectable by the user.

Each track on the disk surface can be visualized as a thin magnetic concentric ring that is divided into identical sectors. There can be from 63 to 106 sectors on a single track, depending on where the track is located. The tracks located toward the outer perimeter of the disk contain more sectors than those toward the inner perimeter. Each sector contains a set of predefined fields. These fields are grouped into two categories: ID and Data. The ID fields contain addressing information so that the firmware can locate a sector of interest for reading or writing. The Data fields contain the data stored on the disk, most of which consists of user data files. Byte values for and brief descriptions of the individual fields that comprise a sector are shown in Table 5-2.

Table 5-1 *Cylinder Contents*

ZONE	NUMBER OF TRACKS	NUMBER OF SECTORS	DATA RATE (Mbits/s)
0	189	107	47.61
1	179	107	47.61
2	179	107	47.61
3	179	107	47.61
4	179	107	47.61
5	179	106	47.27
6	179	102	45.71
7	179	98	43.81
8	179	94	42.05
9	179	88	40.00
10	179	84	37.94
11	179	81	36.66
12	179	79	35.75
13	179	74	33.84
14	179	70	31.88
15	179	65	29.84

Note: ***Zone 15 is the innermost zone, and zone 0 is the outermost zone.***

Table 5-2 *Typical Disk Format*

	FIELD TYPE	NAME	NUMBER OF BYTES	BYTE VALUES
1 SECTOR 570 to 590 Bytes	ID	Sync Field	9	66h
		Address Mark	2	71h 8Eh
		Data Field	6	Varies
		Data CRC	3	Varies
		Pad	3	77h
	Data	Sync Field	16	66h
		Address Mark	2	71h A3h
		User Data Field	512 to 532	Varies
		ECC	14	Varies
		Pad	3	77h
INTERSECTOR GAP				

ID Sync Field

Contains 9 bytes of 66h. The read channel uses this pattern to achieve synchronization so that the drive can read the rest of the ID information correctly.

ID Address Mark (Illegal Pattern)

Contains a uniquely encoded pattern that the drive recognizes as the pattern just preceding the ID Data field.

ID Data Field

Contains 6 bytes that specify the exact address, or ID, of the sector, so that the disk controller can determine where and when to read or write data.

ID Data Cyclic Redundancy Check (CRC)

Contains CRC bytes that provide the firmware with information that allows the drive to detect an erroneous reading of the ID fields.

ID Pad Bytes

Contains 3 bytes of 77h so that there is a definite change in the ID pattern before the Data Sync Field begins. The change in the ID pattern minimizes the chance of incorrect synchronization.

Data Sync Field

Contains 16 bytes of 66h. The read channel uses this pattern to achieve synchronization so that the drive correctly reads the other Data fields (Data Address Mark, Data User Data Field, Data ECC) that follow.

Data Address Mark (Illegal Pattern)

Contains a uniquely encoded pattern that the drive recognizes as the pattern just preceding the Data User Data Field.

Data User Data Field

Contains a fixed number of data bytes in which the drive stores user data files. At customer request, the factory can configure the Empire 540/1080S with a data field from 512 to 532 bytes.

Data Error Correction Code (ECC)

Contains 14 ECC bytes that allows the drive to detect and correct an erroneous reading of the Data fields.

Data Pad Bytes

Contains 3 bytes of 77h so that there is a definite change in the data pattern before the Intersector Gap field begins.

5.1.4 Headstack Assembly

The headstack assembly consists of read/write heads, the actuator E-block assembly, and a flex circuit. The head arms and rotor coil counterbalance each other to place the center of the headstack mass at the headstack mounting hub. Miniature thin-film, slider-type heads mounted to spring-steel flexures are swaged onto the head arms. The flex circuit exits the head/disk enclosure between the cover and the base. The flex circuit contains a Read Preamplifier/Write Driver integrated circuit (IC).

5.1.5 Rotary Positioner Assembly

In the Empire 540/1080S hard disk drive, 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, a rotary single-phase coil, and a bearing shaft. The magnets consist of two alternating poles bonded to the magnet plates. Resilient crash stops mounted in the base prevent the heads from being driven into the spindle or off 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.

5.1.6 Automatic Actuator Lock

To ensure data integrity and prevent damage during shipment, the Empire 540/1080S 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 airvane mounted near the perimeter of the disk stack and a locking arm that restrains the actuator arm assembly.

When 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 airvane increases with disk rotation, the locking arm pivots away from the actuator arm, enabling the headstack to move out of the landing zone. When 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.

5.1.7 Air Filter

Empire hard disk drives are Winchester-type drives. The heads fly very close to the media surface, at a nominal flying height of 3.0 microinches. 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, sealing the drive mechanism under a metal cover. When the drive is in use, the rotation of the disks forces the air inside the drive through an internal filter. Figure 5-2 shows the airflow in the enclosed HDA.

When the disks rotate, the lowest air pressure inside the HDA is above the center of the spindle. A 0.3-micron breather filter bonded to the cover near the spindle allows outside air to flow into the sealed HDA, equalizing internal and external pressures.

The highest air pressure in the HDA is at the outer perimeter of the disks. A constant stream of air flows through a 0.3-micron circulation filter bonded to the base casting. Because the air pressure at the side of the filter away from the perimeter of the disks is lower than that near the perimeter of the disks, air flows through the circulation filter in the direction of the disk rotation, providing a continuous flow of filtered air when the disks rotate.

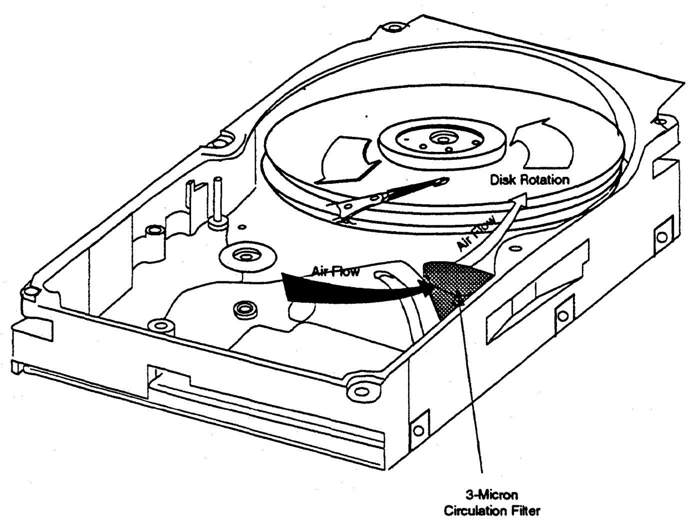

Figure 5-2 *Air Filtration*

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

5.2 DRIVE ELECTRONICS

The Empire 540/1080S hard disk drive is a highly integrated electronic system, designed to provide maximum functionality and reliability in a compact package. The use of surface-mount components, two powerful microprocessors, and several proprietary ASICs allows the drive electronics to be mounted on a PCB. The electrical connection between the PCB and the HDA is made through the flex circuit, a flexible PCB mounted inside the sealed HDA.

The drive electronics include the following elements:

- Read/write preamplifier
- Read channel IC
- Disk controller ASIC
- Digital signal processor
- Servo control
- Spindle motor
- SCSI interface IC

A block diagram of the Empire 540/1080S drive electronics is shown in Figure 5-3.

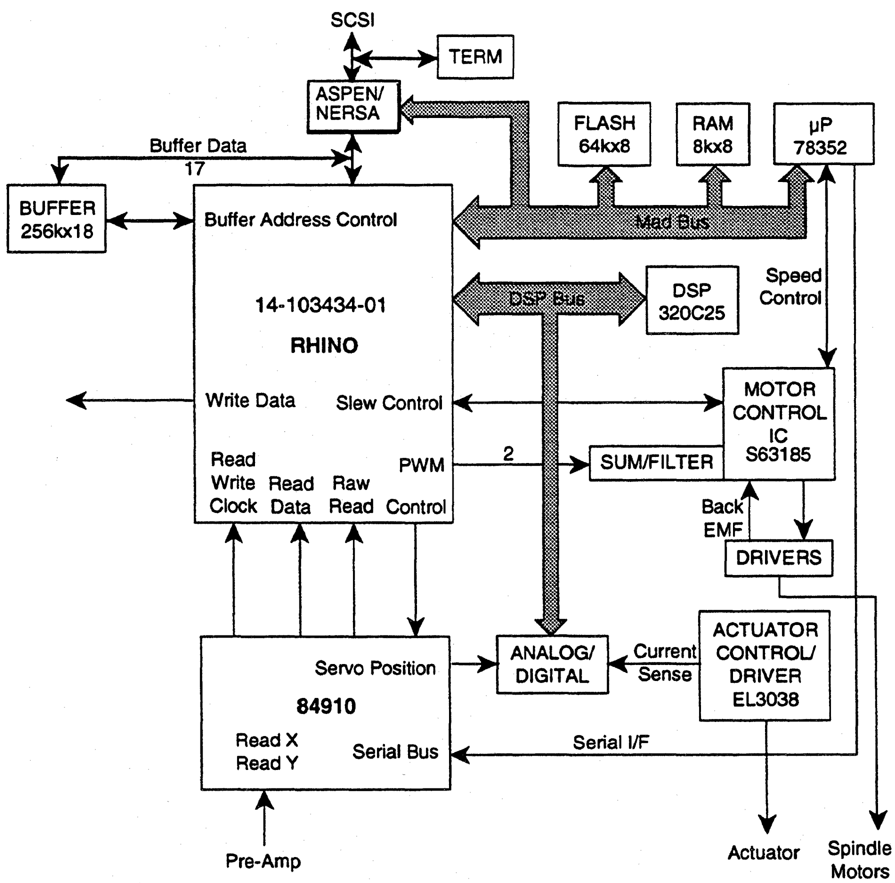

Figure 5-3 *Drive Electronics Block Diagram*

5.2.1 Read/Write Preamplifier

The read/write preamplifier provides write current and data protection circuitry for the recording heads used in the Empire 540/1080S hard disk drive. During the time when the power supply is powering up or down, protection circuitry disables write current. This feature eliminates the possibility of writing erroneous data during power cycles.

The preamplifier addresses one head at a time by means of the Head Select bits, HS0, HS1, and HS2, which are controlled by the microprocessor. In the write mode, the preamplifier provides switched write current to the applicable head. In the read mode, the preamplifier provides data amplification.

5.2.2 Read Channel IC

The read channel IC receives data from the thinfilm read/write preamplifier, peak detects the read pulses for both data and embedded servo information, and resynchronizes the resulting data. The chip operates in the range of 29 to 48 Mbit/s and performs zoned data recording using an on-chip pulse detector, filter, equalizer, frequency synthesizer, and synchronizer. The chip is configured for zoning operation by means of a four-bank control register. A control register allows selective powering down independent portions of the IC, thus allowing necessary operations to be maintained while saving power consumption on unneeded functions. No control register information is lost on selective power down. The chip is fabricated using a bipolar/CMOS process, which is suited well to high-performance analog and digital functions. The main functional blocks of the chip are described in the following sections.

5.2.2.1 Pulse Detector

The pulse detector section uses programmable equalization (pulse slimming) and filtering for zoned data recording band selection. Two gated servo detectors are used to recover quadrature embedded servo information. Rapid head-switch settling and embedded servo normalization is made possible with a pattern-insensitive, fast-responding AGC with hold function. Rapid recovery from write mode is achieved through a delayed, programmable, low-impedance switch at the pulse detector input. A 4-step programmable delay in the qualification channel provides additional baseline noise rejection.

5.2.2.2 Data Synchronizer

The data synchronizer utilizes zero-phase start (ZPS) and a digitally controlled window strobe function which allows 0.5 ns step increments about the window center. The VCO requires no external components and has a wide dynamic range, which is needed for zoned data rate applications. Data windowing is based on precise VCO duty cycle symmetry, which contrasts with delay-line based centering. An internal solid-state delay line establishes the phase detector retrace angle, automatically tracks zoned data recording data rate changes, and can derive its delay reference from either the synchronizer or synthesizer VCO period. The charge pump output and VCO input are provided as separate pins, which allows great flexibility in the design of the external loop filter. Charge pump (phase detector) gain can be selected to remain constant or to vary either by a factor of two or four and can be set to switch at the internal detection of lock or manually as instructed by the CPGAIN pin and the control register.

5.2.2.3 Frequency Synthesizer

The frequency synthesizer is able to create a large number of frequencies from a single external source. In this manner, it produces all the necessary clocks for data encoding and synchronizer reference. A 10-bit control word in the control register allows programmability of the synthesizer frequency. Full control is available over the input divider and feedback divider programming. The only external components needed are those for the RC loop filter. This filter allows full control over the PLL dynamics.

5.2.3 Disk Controller ASIC

The disk controller ASIC performs or participates in the following functions:

- Microprocessor interface
- DSP interface
- Disk read/write sequencer
- Buffer management
- Servo control
- Motor control

5.2.3.1 Microprocessor Interface

The ASIC-microprocessor communication is controlled by the following registers:

- Control and map
- Clocks control
- Version register (read-only)

5.2.3.2 DSP Interface

The disk controller ASIC communicates with the DSP through several read/write registers, which interface with the servo control and buffer sections.

5.2.3.3 Data Sequencer

The disk controller data sequencer consists of the following:

- Writable control store (WCS)
- Address mark detector
- FIFO
- Encoder/decoder for RLL 1,7 data
- Error detection/correction

Writable Control Store

The writable control store (WCS) is an area of RAM used to store executable microcode programs that are downloaded from the microprocessor. The WCS may be reprogrammed "on-the-fly" to execute disk reads or writes and to execute various commands that control he movement of data to and from the disk, including Address Marks and ECC data. The WCS is a flexible means of offloading lower-level processing tasks from the microprocessor. The WCS executes commands from the microprocessor once the program has been downloaded and reports back completion and error status.

Address Mark Detector

The address mark detector does three basic things:

- Finds Address Marks
- Determines Address Mark type—ID, Data, or Data/Apple
- Resynchronizes byte clock with data

FIFO

The FIFO allows both the host and the disk to have an interface to the buffer DRAM. Disk transfers are initiated whenever there are 5 words (or spaces) in the disk FIFO.

Encoder/Decoder

The Run-Length-Limited (RLL) 1,7 coding scheme specifies that each pair of NRZ bits written to the disk be encoded first into three bits. This means that if the data destined for the disk is at a 1.25 MB/s rate (10 Mbit/s), the RLL 1,7 encoding has to increase the data rate to 15Mbit/s in order to accommodate the extra encoded bit and keep the overall usable data at the same transfer rate. The 1,7 encoding guarantees that there will be at least one zero (and never more than seven zeros) between any two ones. The write circuitry encodes each one as a flux change on the disk.

Error Detection/Correction

A 14-bit ECC pattern is used, which includes a 2-byte cross-check and three 4-byte Reed-Solomon interleaves. The Reed-Solomon bytes are used to perform correction on-the-fly for single burst errors, and the cross check bytes are used to double check the correction. The Reed-Solomon code will correct any single-error burst, as long as it does not extend to two bytes in the same interleave. Thus, all 17-bit errors are detected as single-bursts and are corrected on-the-fly. Double bursts still are correctable, but re-reading is attempted before a double-burst correction is attempted.

5.2.3.4 Buffer Management

The disk controller controls access to the disk controller DRAM. The DRAM is accessible through arbitration from the disk, the DSP, the microprocessor, the refresh controller, or the host. The microprocessor communicates with the buffer controller through various read/write registers.

5.2.3.5 Servo Control

The disk controller IC provides the servo decoding hardware for the embedded servo operation. It works with the DSP, read channel, and voice coil motor drive to execute the seek- and track-following functions of the drive. It detects the Servo Address Mark (SAM), track number, and servo positional information from each servo sample and the timing between the samples. It is asynchronous (not phase-locked) to the disk data and is designed to interface with the read channel chip. It receives non-synchronized pulse detected read data from the read channel IC and in turn controls the AGC and peak detection functions.

The main functional blocks of the servo hardware within the disk controller are as follows:

- Sync
- Sync & 14T detector (T refers to one 26.67 Mhz clock cell)
- Data reader
- Delay timer
- Sector timer
- Master state machine
- Pulse width modulators

5.2.3.6 Motor Control

The disk controller IC controls the disk drive spindle motor through various read/write registers.

5.2.4 Digital Signal Processor

The Empire 540/1080S drive uses a DSP to create adequate margin, principally by incorporating an adaptive embedded servo control system. The older method of a dedicated servo, wherein one surface is dedicated to servo information, assumes that if a head is on-track on that surface, all other heads are also on their tracks. Higher track densities can make a dedicated servo scheme hypersensitive to temperature changes, especially when combined with shock or vibration. In this case, the drive experiences high error rates, and might not be able to retrieve data at all if conditions have changed since the servo data was recorded.

The Empire 540/1080S embedded servo control scheme provides feedback as 58 prerecorded servo bursts per track. These are placed on adjacent surfaces and are skewed with respect to each other by one 116th of a revolution. Using the DSP allows adaptive positioning of the head over the track that is actually being read, which compensates in real time for both external disturbances such as shock and vibration and internal changes such as the aging of shock mounts and creep of materials. The DSP can analyze bursts of servo data and make quick positional corrections.

5.2.4.1 Chip Architecture

The DSP implements a Harvard-type architecture that boosts processing power by using two separate program and data bus structures. The instruction set allows transferring data between the program and data spaces. Externally, the DSP allows program and data to be multiplexed on a common bus, which maximizes the address space for both and keeps pin count at a minimum.

5.2.4.2 On-Chip Memory

The DSP has a total RAM space of 544 16-bit words. This space is configured as two separate RAM blocks, one of which can be set up as either program or data space. A 4Kbyte ROM can be used to hold on-board code, and it can be expanded to a total of 64Kbytes through the addition of external ROM or EPROM.

5.2.4.3 Arithmetic Logic Unit

A 32-bit ALU and 32-bit accumulator allows the DSP to perform 2s-complement arithmetic The ALU is a general-purpose arithmetic unit that operates upon 16-bit words taken from data RAM or derived from immediate arithmetic instructions or from the 32-bit result of the multiplier's product register. The ALU can perform Boolean operations in addition to the normal arithmetic operations, and can perform powerful bit manipulations. The accumulator is used as an input to or output from the ALU, is 32 bits long, and is divided into high-order and low-order words of 8 bits each.

5.2.4.4 Multiplier

The DSP multiplier is able to perform a 16 X 16 2s complement multiplication with a 32 bit result all in a single instruction cycle. There are three elements in the multiplier:

- A "T" register that is 16 bits long and temporarily stores the multiplicand
- A "P" register that stores the 32-bit product
- A multiplier array

Multiplier values come into the multiplier from three sources:

- Data memory
- Program memory when using the MAC/MACD instructions
- Immediately from the MPYK (multiply immediate) instruction word

A scaling shifter allows the DSP to perform numerical scaling, bit extraction, extended arithmetic, and overflow prevention. The fast multiplication ability allows the DSP to perform fundamental operations such as convolution, correlation, and filtering.

5.2.4.5 Memory Interface

The DSP local memory interface consists of a 16-bit parallel data bus (D0—15), a 16-bit address bus (A0–15), three pins to select data, program, or I/O space, and several system control signals. The direction of data transfer (read or write) is controlled by the R/W pin, and a strobe signal times the actual transfer. The DSP normally executes without wait states, although a READY signal allows communication with slower external devices, such as off-chip memory.

The program counter is saved during the execution of interrupts and subroutine calls by a hardware stack that may be up to eight layers deep. In addition to the program counter, the complete context of the DSP may be saved by PUSH and POP instructions, with nesting being limited only by the amount of RAM available.

All of the control operations are supported by:

- An on-chip memory-mapped 16-bit timer
- A repeat counter
- Three external maskable user interrupts
- Internal interrupts that are generated by the need to service the serial port or the timer

5.2.4.6 Serial Port

The DSP serial port provides a full-duplex communications path with external serial devices. There are two memory-mapped registers for the serial port: one for the transmit function, and one for the receive function.

5.2.5 Servo Control

The servo system or, disk head positioning system, is based on an embedded servo architecture that provides position data from the actual head being used for reading and writing. At the factory, a special positioning format is written on each track. User data is interleaved with this information, which provides absolute head position down to small fractions of a track. This servo information is separated from the user data, synchronized, and sampled to provide position data on a regular basis. This sampled data stream is processed by the servo software to generate a series of control outputs to seek to and/or keep the active head positioned over the desired track.

An embedded servo architecture interleaves user data and position information so the actual head being used for read/write is the same one being used to gather position data. This embedded architecture maximizes surface utilization, minimizes tracking errors, and allows for high track densities. The embedded servo information is written in equally spaced intervals along each track, forming radial "wedges" of servo information on each disk surface. Some features of the embedded format are:

- Wedges staggered between surfaces to provide fast servo write and verify.
- Wedge information occupies less than 13% of each disk surface.
- Each wedge contains the following fields: Write-to-read recovery, AGC, Sync, SAM, Index, Track Number, and Servo Bursts.

The Empire 540/1080S embedded servo control scheme provides feedback as 58 prerecorded servo bursts per track. These are placed on adjacent surfaces and are skewed with respect to each other by one 116th of a revolution. Using a DSP chip allows adaptive positioning of the head over the track that actually is being read, which compensates in real time for both external disturbances such as shock and vibration and internal changes such as the aging of shock mounts and creep of materials. The DSP can analyze bursts of servo data and make quick positional corrections.

5.2.5.1 Servo Type

The servo type used in Empire 540/1080S is embedded, consisting of 58 samples per revolution, with the servo bursts staggered between surfaces. The servo sample rate is constant. The constant rate is achieved by using split data fields. With a rotational speed of 5400 rpm, the 58 samples per revolution yields a sample rate of 5.22 Khz.

5.2.5.2 Servo Control

The function of the servo system is to generate a series of control outputs to seek to and/or keep the active head positioned over the desired track. The servo control mechanism is based on an estimator and controller. The estimator is a single rate, 4th order current estimator, updated by periodic samples of the head position and voice coil motor (VCM) current. The estimator gives the head position, velocity, bias force, and VCM current. The controller generates an output based on the estimator's state and the location of the target track. The two main operating modes of the servo are tracking and seeking. The estimator and controller are used for both modes, although parameters of each may change.

The servo control mechanism is a state estimator combined with a velocity control loop. A state-space model is used to keep track of the actuator and the coil driver by periodically measuring the coil current and head position. The same estimator is used for both seeking and tracking, although parameters of the estimator and controller can change from one mode to the other. Each time the model states are updated, a new control output is computed based upon the values of those states and the location of the target track.

When the primary head is near to track center of the target track, the servo uses linear control based on the estimated states of the actuator and coil driver. For large tracking errors (> 0.25 Tks), the servo uses velocity control, with the target velocity varying according to a non-linear function of the tracking error (roughly as the square root of the tracking error), up to a predetermined maximum velocity.

5.2.5.3 Microcomputer/DSP Interface

The microcomputer and DSP operate in a master/slave relationship. The microcomputer is the master, and is able to read any location in the DSP's data or program memory, and can write to any RAM-based data or program memory locations. A single exchange, which can occur once per control interval (twice per servo sample) involves either a read or a write of a single DSP memory location. The DSP code is structured such that a single write can initiate a seek.

5.2.6 Spindle Motor

The Empire 540/1080S spindle motor is a three-phase motor and is responsible for spinning the disks at a rotational speed of 5400 ±0.3%.

The spindle motor is Y-wound, with the center tap (neutral) brought out. The motor contains eight poles and operates as a three-phase device.

5.2.7 SCSI Interface IC

Quantum's proprietary SCSI interface IC implements the SCSI interface. The drive's interface with the host system is through a 50-pin SCSI-bus cable. A microprocessor controls the SCSI interface IC through its control lines. The SCSI interface IC writes data to or reads data from the buffer RAM over eight data lines at a rate of 10 MB/s. This high data transfer rate allows the SCSI interface IC to communicate over the SCSI bus at an asynchronous rate of 5.0 MB/s and a synchronous rate of 10.0 MB/s while it simultaneously controls disk-to-RAM transfers. The advanced SCSI interface IC was designed for target mode applications only and has a buffer memory interface consisting of 16 data bits and 1 parity bit.

The SCSI interface IC automates many of the SCSI sequences described in the SCSI-2 specification. With the SCSI interface IC, it is possible to respond to selection, arbitrate, reselect an initiator, send or receive messages, receive command bytes, send or receive data, and send status automatically. Under the proper conditions, the SCSI interface IC also can handle read and write commands from responding to SELECTION to COMMAND COMPLETE without microprocessor intervention.

The SCSI electronics consists of six major blocks, as follows:

- Respond to Selection-Arbitration State Machine (RASM)
- Auto SCSI State Machine (ASSM)
- Asynchronous REQ State Machine (ARSM)
- Synchronous REQ State Machine (SRSM)
- 8-word X 16-bit FIFO
- Byte and block counters

All major blocks are interconnected and communicate with each other as necessary to carry out automated SCSI sequences.

5.3 FIRMWARE FEATURES

The section describes the following firmware features:

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

5.3.1 Disk Caching

The Empire 540/1080S hard disk drive incorporates DisCache, a 512 KB disk cache, to enhance drive performance. 256 KB are allocated to the Read Cache and 128 KB are allocated to the write cache. This integrated feature can significantly improve system throughput. Read and write caching can be enabled or disabled by using the Mode Select command.

5.3.1.1 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 store the subsequent data from the disk into high-speed RAM. If the host subsequently requests this 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 saves 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 make up 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 non-caching 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 256Kbyte DRAM buffer allocated to hold the data, which can be directly accessed by the host by means of the READ and WRITE commands.

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

5.3.1.2 Write Cache

When a write command is executed with write caching enabled, the drive stores the data to be written in a 128Kbyte 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.

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 during sequential data transfers.

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

5.3.1.3 Performance Benefits

For both read and write operations, a cache scan is performed. On read operations, if the data is in the buffer, it is transferred from the cache to the host, eliminating the need for a disk access. On write operations, if the logical block address that is being written to is already in the cache, it is overwritten in the cache and then written to the disk at a later time.

In a drive without DisCache, there is a delay during sequential reads because of rotational latency even if the disk actuator is already positioned at the desired cylinder. DisCache eliminates this rotational latency time (5.6 milliseconds on average) when requested data resides in the cache.

Moreover, the disk must often service requests from multiple processes in a multitasking or multiuser environment. In these instances, while each process may 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 it most impressive performance gains, saving both seek and latency time (15.6 milliseconds on average) when desired data resides in the cache.

5.3.2 Track and Cylinder Skewing

Track and cylinder skewing in the Empire 540/1080S minimizes latency time and, thus, increases data throughput.

5.3.2.1 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 Empire 540/1080S (1.5 ms), the sector addresses can be optimally positioned across track boundaries to minimize the latency time during a head switch.

5.3.2.2 Cylinder Skewing

Cylinder skewing is also 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 Empire 540/1080S (< 2 ms), the sector addresses can be optimally positioned across cylinder boundaries to minimize the latency time associated with a single-cylinder seek.

5.3.3 Error Detection and Correction

An error occurs whenever the drive fails to read data correctly—because either the read does not complete its execution, or the drive read data checksum is different than the recorded checksum. During a read operation an error can occur while a sector's header or data is being read. During a write operation, an error can occur while reading a sector's header before the data is written.

5.3.3.1 Error Correction Code and ECC On-The-Fly

The Empire 540/1080S hard disk drive implements a 14-byte, three-way interleaved Reed-Solomon error detection and correction algorithm. 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 drive does not report a recovered error unless it needs more than two disk revolutions to correct the error.

When errors cannot be corrected on-the-fly, an automatic retry and a more rigorous double burst error correction algorithm enable the correction of any sector with two bursts of three incorrect bytes each. In addition to this error correction capability, the drive's implementation of an additional cross-checking code and algorithm double checks the main ECC correction and greatly reduces the probability of miscorrection.

In addition to other fields (see Table 5-2), each sector contains 512 bytes of data, 2 cross-check bytes, and 12 ECC bytes. The two cross-check bytes follow the 512 bytes of data to double check the main ECC correction and reduce the probability of miscorrection. The cross-check bytes are followed by 12 ECC bytes. The cross-check and ECC bytes are computed and appended to the user data when the sector is first written.Figure 5-4 shows the format of part of a typical sector, and shows the location of the ECC and cross-check bytes.

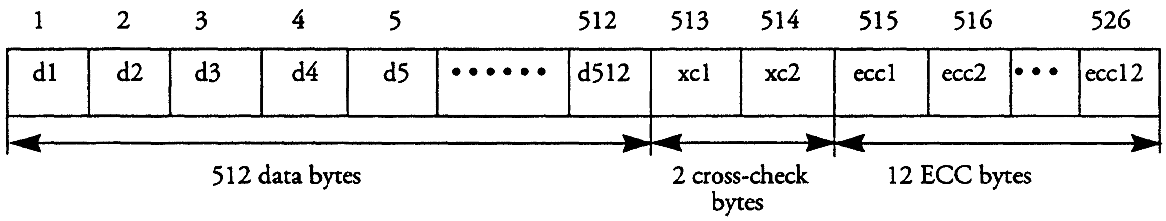

Figure 5-4 *Sector Data Field with Cross-Check and ECC Bytes*

Each byte within a sector, including the cross-check and ECC bytes, is interleaved into three groups, where the first byte is in interleave group 1, the second byte is in interleave group 2, the third in group 3, the fourth in group 1, the fifth in group 2, and so on. Four ECC bytes are computed for the data bytes in each interleave group. These 12 ECC bytes appear at the end of the sector. The 2 cross-check bytes are computed for the entire 512 data bytes, and appear between the 512 data bytes and the ECC bytes. This method of organizing the ECC and cross-check bytes, and the nature of the ECC formulas enable the drive to locate an error.Figure 5-5 illustrates the interleaving concept.

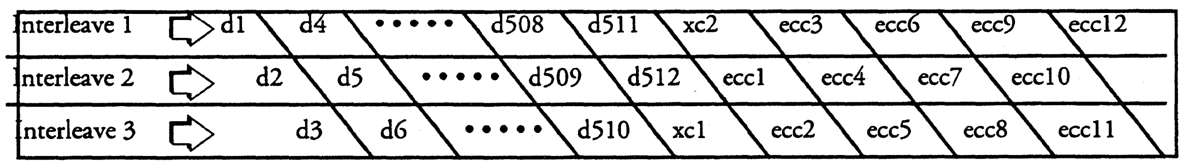

Figure 5-5 *Example of Byte Interleaving*

The power and sophistication of the error algorithm also provides for correction of errors within the ECC bytes. Errors in the cross-check or ECC bytes are corrected on-the-fly if their characteristics fall into the guidelines of a single burst correctable error.

Each time the drive reads a sector of data, it generates a new set of ECC and cross-check bytes based on the 512 bytes of data contained within the sector. The new set of cross-check and ECC bytes is compared with the corresponding bytes originally written in that particular sector. This comparison process results in bytes that are known as syndromes. There are 2 cross-check syndromes and 12 ECC syndromes. If all of the syndrome values are zero, the data has been read with no errors, and the sector of data is transferred to the host. If any on the syndromes are non-zero, an error has occurred. The type of correction applied by the drive then depends on the nature and extent of the error.

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 5-6, 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.

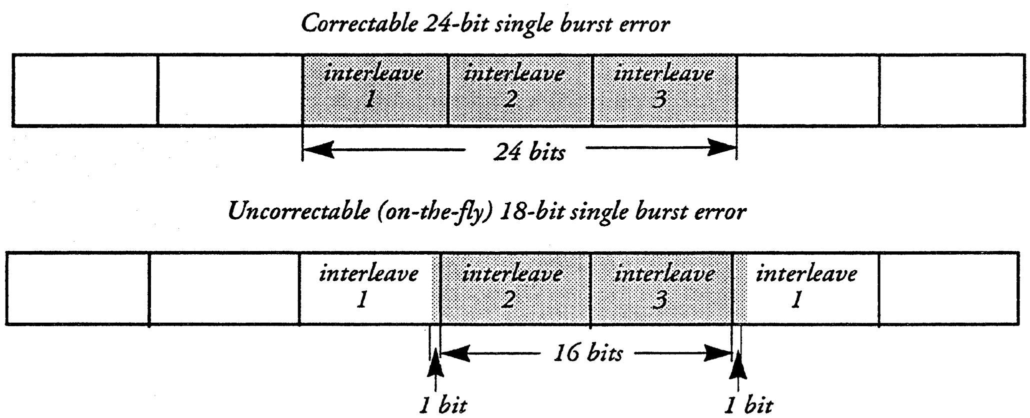

Figure 5-6 *Single Burst Error Correctability*

Note: In Figure 5-6, 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.

Correction of Double Burst Errors

The Empire 540/1080S hard disk drive also has the capability to correct double burst errors, even though the likelihood of this type of error is very low. Double burst errors may be viewed as two spans of errors within a sector. More specifically, correctable double burst errors are characterized by error bytes that occur in two or fewer bytes in an interleave, and are not corrected on-the-fly.

Error bursts up to 48 bits long can be corrected, provided that the error consists of two or fewer bytes residing in each of the interleaves. In Figure 5-7, the 42-bit error is uncorrectable, while the other two 48-bit errors are correctable. The reason the 42-bit error is uncorrectable is that occupies two interleave 2s, and two interleave 3s, but occupies three interleave 1s, whereas the limit is two bytes in each interleave. Any 41-bit error burst is corrected using double burst error correction because no more than two bytes occupy each interleave.

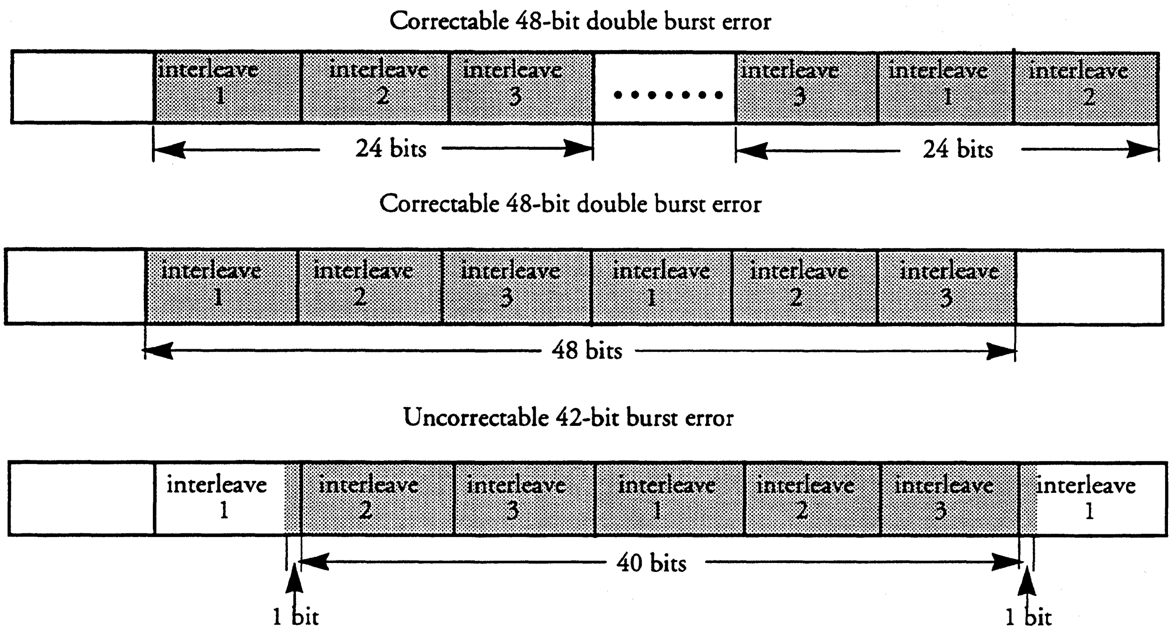

Figure 5-7 Double Burst Error Correctability

Note: In Figure 5-7, the shaded portions represent data containing errors.

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.

Multiple Random Burst Errors

The Empire 540/1080S hard disk drive can correct up to 48 bits of multiple random errors provided the errored bytes follow the guidelines for correctable double burst errors; that is, if more than two bytes in any one interleave are in error, the sector cannot be corrected. Up to 24 bits of multiple random errors can be corrected on-the-fly if only one byte per interleave contains an error.Figure 5-8 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.

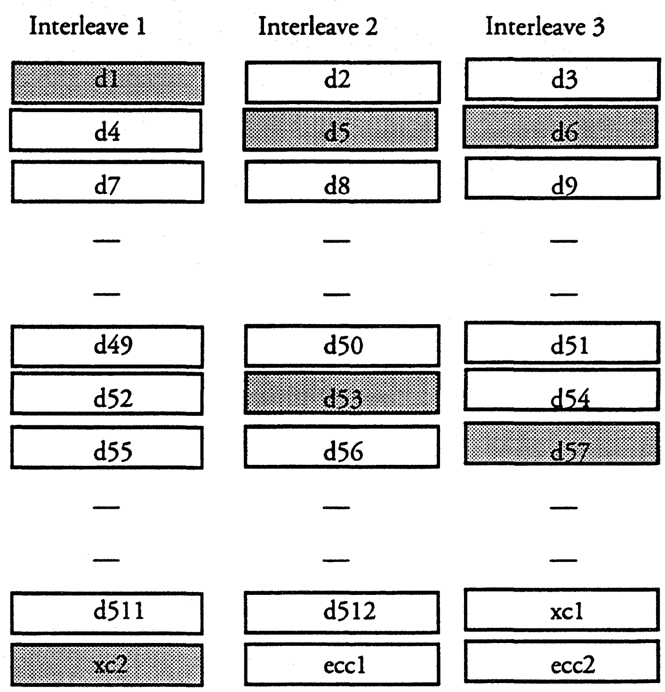

Figure 5-8 A Correctable Random Burst Error

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

Error Rates

Due to the error-correction capabilities mentioned previously, the Empire 540/1080S hard disk drive experiences an unrecoverable read error rate of less than one error in 10^{14} bits transferred.

Error Recovery Control

The CONFIGURATION command, explained fully in Chapter 6, allows you to control the following error recovery parameters:

- Automatic Write Reallocation -automatically reallocates blocks that have become unavailable for writing.
- Automatic Read Reallocation -automatically reallocates bad sectors when two consecutive read retries result in the same bad ECC syndrome, indicating a hard error.
- Read Continuous - disables error-recovery procedures, thus allowing data to be read at a higher rate from the drive.
- Enable Early Correction -enables the ECC algorithm when the drive detects two consecutive non-zero syndromes, and no rereads will occur before applying correction, unless the error is uncorrectable.
- Disable Correction - data is transferred without correction, even if correction is possible. ECC on-the-fly is not affected.
- Number of Retries - specify the number of times the drive will re-read the data, in attempting to recover from data errors.

ECC Error Handling

When a data error occurs, the Empire 540/1080S hard disk drive checks to see if the error is correctable on-the-fly. This process takes about 200 µs. If the error is correctable on the fly, the error is corrected and the data is transferred to the host system. During this entire process, data continues to be read from the disk. Thus, if an error is recoverable on-the-fly, no loss of transfer rate is observed.

If the data is not correctable on-the-fly, the sector is re-read up to 8 times in an attempt to read the data correctly without applying double burst ECC correction. Before invoking the complex double burst ECC algorithm, the drive will always try to recover from an error by attempting to re-read the data correctly. This strategy prevents invoking correction on soft errors. Each time a sector in error is re-read, a set of ECC syndromes is computed. If all of the syndrome values are 0, the data was read with no errors, and the sector is transferred to the host system. If any of the syndromes are not 0, an error has occurred. The syndrome values are then retained, and another re-read is attempted.

When the sets of syndromes from two consecutive reads compare, a “stable syndrome” has been achieved. This event can be significant depending on whether automatic read reallocation or early correction has been enabled. If early correction is enabled, and a stable syndrome has been achieved, double burst correction is applied.

If the automatic read reallocation feature is enabled, the drive attempts up to 24 re-reads for double burst errors. In this case, the ECC correction algorithm is divided into three parts:

1. When an error is encountered, the drive attempts to re-read the correct data 8 times, aborting the re-reads if the correct data was successfully read. On-the-fly correction may be applied. The drive always tries to re-read the correct data before applying double burst ECC correction and automatic sector reallocation.
2. If the first 8 re-reads are unsuccessful, the drive attempts to read the sector 8 more times. If the drive does not read the data correctly, it looks for a stable syndrome on two consecutive re-reads. If the drive finds a stable syndrome, it tries to correct the sector and automatically reallocates the data to a spare sector. If the sector is uncorrectable, the sector is re-read until the second set of 8 re-reads has been completed.
3. If the second 8 re-reads are unsuccessful, the drive attempts to read the sector yet 8 more times. If the drive finds a stable syndrome from two consecutive reads, it reallocates the data to a spare sector whether the data is ECC correctable or not. If at any time during this set of re-reads the drive discovers that the data is correctable, it corrects the error and reallocates the data to a spare sector.

When all 24 re-reads have been completed, and the error is uncorrectable, the sector is marked as bad and the data is reallocated to a spare sector. When the Disable Early Correction bit is set to 1 for testing purposes, the drive attempts to re-read. Re-reads are disabled when the Number of Retries byte in the Configuration command is set to 0. The default setting for the number of retries is set at 8.

5.3.4 Defect Management

The Empire 540S allocates two sectors per cylinder as a spare; the Empire 1080S allocates four sectors 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 which would result from offline defect sparing.

If more than four sectors are found defective on a cylinder, the inline sparing technique is applied only to the four sectors. 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 which would be otherwise needed to find a replacement sector. The result in minimal reduction of data throughput.

Defects that occur in the field are known as grown defects. Inline sparing is not performed on these grown defects.

Sectors are considered to contain grown defects if the double burst ECC algorithm must be applied in order 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 new sector, and a flag is set in the data ID field which 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 Small Computer System Interface (SCSI) is an industry standard for connecting peripheral equipment to personal computers or workstations. Unlike many interface standards, SCSI is not intended for just one type of peripheral device. Instead, it uses a generic bus protocol and hardware interface, in conjunction with a flexible command protocol, to allow a wide variety of device types to communicate across the same bus.

6.1.1 Bus Architecture

From two to eight physical devices can be connected to the SCSI bus. A SCSI device can be either an initiator or a target. The SCSI bus can include any combination of initiators and targets, provided at least one initiator and one target are present. Certain functions are assigned to either an initiator or a target:

- An initiator can arbitrate for the bus and select a target.
- A target can request the transfer of command, data, status, or other information to or from the initiator, and in some cases a target can arbitrate for the bus and reselect an initiator to continue an I/O process.

  A target can support from one to eight physical or virtual devices called logical units. A device address consists of the SCSI ID of the target and the Logical Unit Number (LUN) of the device. A physical device that does not support additional logical units–such as the Empire 540/1080S hard disk drive–comprises one logical unit. In this case, LUN is set to zero.

From this point onward, the goal of this chapter will be to familiarize the reader with how SCSI pertains to the Empire 540/1080S hard disk drive. In other words, SCSI will be analyzed and presented from the drive's perspective as a target. Later sections in this chapter are intended to familiarize the user with the Message Codes, Synchronous Data Transfer Protocol, Commands, and Status Bytes supported by the drive.

6.1.2 Bus Transactions

In SCSI, a bus transaction is defined by the protocol as an I/O process. An I/O process begins with the establishment of a logical link called a nexus, which defines the logical path between an initiator and the drive. An I/O process ends with the return of the bus to the BUS FREE phase. A successful I/O process ends following the transfer of a COMMAND COMPLETE message. An unsuccessful I/O process ends after the transfer of an ABORT, ABORT TAG, BUS DEVICE RESET, or CLEAR QUEUE message; when an unexpected disconnect occurs; or if a RESET condition is detected.

6.1.2.1 The Initiator-Target Nexus (I_T)

A nexus is a logical link between an initiator and the drive, represented by the SCSI ID of the initiator and the SCSI ID of the drive. An I_T nexus is established during the SELECTION phase.

6.1.2.2 The Initiator-Target-Logical Unit Nexus (I_T_L)

The IDENTIFY message is used to further define the nexus, by specifying a LUN. The Empire 540/1080S hard disk drive houses a single logical unit. The LUN for this logical unit is zero. Due to this fact, the LUN specified must be zero.

6.1.2.3 The Initiator-Target-LUN-Queue Tag Nexus (I_T_L_Q)

The queue tag messages are used to attach an identifier, called a queue tag, to an I/O process, so that upon reconnection the I/O process can be specified to the initiator. The I_T_L_Q nexus is established during the first MESSAGE OUT phase following SELECTION, when the initiator sends the queue tag to the drive.

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 SCSI device can only perform the tasks allowed for that phase. The bus can only be in one phase at any given time. Transition between phases is described in the explanation 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 phases.

6.2.1 Bus Phase Sequences

A SCSI device transits from one bus phase to another in a pre-defined manner. The sequence begins with the BUS FREE phase, and is followed by the ARBITRATION phase if the SCSI device actively engages in arbitration. Following this, if the device is part of the nexus established during either the SELECTION or RESELECTION phase, the device participates in a combination of Information Transfer phases beginning with either the COMMAND or MESSAGE phase. The sequence ends with the return of the bus to the BUS FREE phase, at which point the SCSI device can begin the sequence again.

In all cases, a RESET condition aborts any active phase. A RESET condition is always followed by the BUS FREE phase.

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 do not change.
- ATN, RST, MSG, C/D, I/O, and DATA BUS signals can change.

Note: When switching the direction of the signal path on the DATA BUS from out to in–that is, from an initiator-driven signal to a drive-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. The initiator releases the bus within a DATA RELEASE DELAY, measured from the transition of I/O to true.

When switching the direction from in to out–that is, from a drive-driven signal to an initiator-driven signal–after negating I/O, the drive releases the bus within a DESKEW DELAY.

6.2.3 BUS FREE Phase

The BUS FREE phase occurs 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, a SCSI device must release all bus lines.

During BUS FREE, a device can only attempt to gain control of the bus. A SCSI device attempting to get the bus must move into the ARBITRATION phase.

6.2.4 ARBITRATION Phase

SCSI systems that implement arbitration are designed to support multiple initiators, and support reselection of an initiator by the drive. 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.

The ARBITRATION phase is active when a device asserts BSY and its SCSI ID on the bus. A device's SCSI ID corresponds to a unique bit on the DATA BUS (SCSI ID #0 to DB(0), SCSI ID #1 to DB(1), etc.).

The arbitration sequence for SCSI is:

1. An arbitrating device detects the BUS FREE phase when both BSY and SEL are continuously false for at least a BUS SETTLE DELAY.

2. After detecting the BUS FREE phase, a device waits for at least a BUS FREE DELAY before driving any signals.

3. Following a BUS FREE DELAY, a device arbitrates for the SCSI bus by asserting both BSY and its SCSI ID. A device does not arbitrate if more than a BUS SET DELAY has passed since it last detected the BUS FREE phase.

Note: When asserting its SCSI ID on the DATA BUS, a device releases the other seven DATA BUS bits. Parity is invalid and DB(P) remains undriven by the device.

4. After waiting for at least an ARBITRATION DELAY–measured from its assertion of BSY–the device examines the DATA BUS. If the device finds that its SCSI ID is the highest asserted SCSI ID on the DATA BUS–DB(7) is the highest–the device wins the arbitration and asserts SEL. Otherwise, the device loses the arbitration, releases all signals within a BUS CLEAR DELAY after SEL becomes true, and repeats this procedure from Step 1.

5. When the SCSI device wins arbitration, after asserting SEL, it waits for at least (BUS CLEAR DELAY)+(BUS SETTLE DELAY) before changing any signals.

6.2.5 SELECTION Phase

The SELECTION phase allows an initiator to select the Empire 540/1080S hard disk drive to initiate a function, such as a READ or WRITE command. The drive is never an initiator.

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

6.2.5.1 Selection Of The Drive

The device that won arbitration:

1. Has BSY, SEL, and its SCSI ID asserted.

2. Has waited for at least 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 and must:

1. Assert the following on the DATA BUS:
   - The drive's SCSI ID
   - If parity is implemented, the parity bit

2. Wait for at least two DESKEW DELAYS before it can release BSY.

Note: If the initiator raises ATN before BSY is released, the drive will go to MESSAGE OUT phase following SELECTION.

3. Wait for at least a BUS SETTLE DELAY, after its release of BSY, before it looks for a response from the drive.

The Empire 540/1080S hard disk drive:

1. Becomes selected when SEL and its SCSI ID are true, and BSY and I/O are false for at least a BUS SETTLE DELAY.
2. Examines the DATA BUS to determine the SCSI ID of the selecting initiator.
3. Asserts BSY within a SELECTION ABORT TIME of its most recently detected selection to ensure the correct operation of the timeout procedure.

After the initiator detects BSY is true, it must wait for at least two DESKEW DELAYS before it releases SEL and can change the DATA BUS.

The drive waits for the initiator to release SEL before changing to an Information Transfer phase. Once drive selection is complete, the LED on the drive lights.

6.2.5.2 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 timeout procedures to clear the SCSI bus. The initiator can:

- Assert RST. (See Section 6.2.8.2, "RESET Condition.")
- Continue to assert SEL (and ATN, if ATN was asserted) and release the DATA BUS. If the initiator does not detect BSY as true after waiting for at least a SELECTION ABORT TIME plus two DESKEW DELAYS, it releases SEL (and ATN), allowing the SCSI bus to go to the BUS FREE phase.

Note: All SCSI devices must ensure that they respond to a selection attempt within a SELECTION ABORT TIME. If a device doesn't respond within a SELECTION ABORT TIME, it must test the bus again to determine whether or not the selection attempt is still valid. Failure to follow this requirement can result in bus malfunction, 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 Phase

The RESELECTION phase allows the Empire 540/1080S hard disk drive to reconnect to an initiator. By doing so, it can continue an operation previously started by the initiator, but suspended by the drive when the drive disconnected, allowing the BUS FREE phase to occur before operation was complete.

Simply stated, RESELECTION is exactly the same as SELECTION, except that the roles of the drive and the initiator are reversed, and I/O is asserted to differentiate between the two phases.

6.2.6.1 Reselection By The Drive

If the drive wins the ARBITRATION phase, the drive:

1. Becomes a target by asserting I/O.
2. Asserts the following on the DATA BUS:
   - The initiator's SCSI ID.
   - If parity is implemented, the parity bit.
3. Waits for at least two DESKEW DELAYS, then releases BSY.
4. Waits for at least a BUS SETTLE DELAY before it looks for a response from the initiator.

The initiator:

1. Becomes reselected when SEL, I/O, and its SCSI ID are true, and BSY is false for at least a BUS SETTLE DELAY.
2. Examines the DATA BUS to determine the SCSI ID of the reselecting drive.
3. Asserts BSY within a SELECTION ABORT TIME of its most recently detected reselection to ensure the correct operation of the timeout procedure.
   - After the drive detects BSY is true, it asserts BSY and must wait for at least two DESKEW DELAYS before releasing SEL. The drive continues asserting BSY until it is ready to relinquish the SCSI bus.
   - After the reselected initiator detects SEL as false, it must release BSY.

Note: If parity is implemented and the initiator detects bad parity, or if more than two SCSI IDs appear on the DATA BUS, the initiator must not respond to reselection.

6.2.6.2 Reselection Timeout Procedure

If the drive waits for at least a SELECTION TIMEOUT DELAY without receiving a BSY response from the initiator, it continues to assert SEL and I/O, and releases the DATA BUS. If the drive does not detect BSY as true after waiting for at least a SELECTION ABORT TIME plus two DESKEW DELAYS, it releases SEL and I/O, allowing the SCSI bus to go to the BUS FREE phase.

Note: All SCSI devices must ensure that they respond to a reselection attempt within a SELECTION ABORT TIME. If a device doesn't respond within a SELECTION ABORT TIME, it must test the bus again to determine whether or not the reselection attempt is still valid. Failure to follow this requirement can result in bus malfunction.

If reselection is unsuccessful, the drive will wait 250 ms and then attempt to reselect again. This process continues indefinitely until the drive is successful, or the drive is reset or told to abort the process.

6.2.7 Information Transfer Phases

Once selection of the drive or reselection by the drive is complete, there is only one initiator connected to the drive. This initiator can be any of the initiators connected to the SCSI bus and will now be referred to as the initiator for the remainder of this chapter.

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." Each of the Information Transfer phases is distinguished from the others by the combination of MSG, C/D, and I/O. The four Information Transfer phases are:

- COMMAND
- DATA
- STATUS
- MESSAGE

The state of MSG, C/D, and I/O determines the active phase, and whenever it changes the new phase becomes valid when REQ is asserted. MSG, C/D, and I/O are controlled by the drive. The state of these three signals and the corresponding Information Transfer phase are shown in Table 6-1.

Table 6-1 *MSG, C/D, and I/O Codes*

SIGNAL STATE			PHASE	DIRECTION OF TRANSFER
MSG	C/D	I/O		
0	0	0	DATA OUT	Initiator to Drive
0	0	1	DATA IN	Drive to Initiator
0	1	0	COMMAND	Initiator to Drive
0	1	1	STATUS	Drive to Initiator
1	0	0	RESERVED	–
1	0	1	RESERVED	–
1	1	0	MESSAGE OUT	Initiator to Drive
1	1	1	MESSAGE IN	Drive to Initiator

Note: 0 = False, 1 = True.

In SCSI, for all transfer-type operations, the flow of data is viewed from the perspective of the initiator. "In" refers to information (data, messages) being transferred to the initiator, while "out" refers to information being transferred from the initiator.

A description of each of the phases mentioned in Table 6-1 can be found later in this chapter. The DATA IN and DATA OUT phases are explained in the section devoted to command descriptions, since they are only used during the execution of certain commands.

When these specific commands are described, the function and use of the DATA IN and DATA OUT phases will also be expressed.

Communication with the Empire 540/1080S hard disk drive takes place on the SCSI data bus using a defined REQ/ACK handshake protocol. Transfers across the bus cancan be synchronous or asynchronous for the DATA IN and DATA OUT phases, and must be asynchronous for all other phases. BSY remains asserted and SEL remains released during all Information Transfer phases. MSG, C/D, and I/O are set by the drive a BUS SETTLE DELAY prior to the first handshake of a phase and remain valid until after the negation of ACK on the last handshake.

6.2.7.1 Pointers

In the SCSI architecture, when an initiator establishes an I/O process, it creates a set of three pointers to accompany that I/O process. These three pointers remain in existence until the I/O process is finished, and are therefore called the saved pointers. In addition, each initiator contains a set of three pointers known as the active pointers. Each set of pointers consists of a command pointer, a data pointer, and a status pointer.

When an I/O process connects to the bus, the initiator copies the values of the I/O process's 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 drive.

A saved command pointer indicates the beginning of the Command Descriptor Block (CDB) for the I/O process. A saved status pointer references the beginning of the status area for the I/O process. A saved data pointer points to the beginning of the data area for the I/O process until the drive sends a SAVE DATA POINTER message. In response to this message, the initiator copies the value of its active data pointer into the saved data pointer for that I/O process. When the I/O process disconnects from the SCSI bus, the initiator retains only the saved pointer values.

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. When the drive reconnects, a RESTORE DATA POINTER message is implied.

Note: The drive does not issue a SAVE DATA POINTER message on a READ if the drive disconnects prior to transferring any data, or on a WRITE if the drive disconnects after transferring all the data.

6.2.7.2 Asynchronous Transfer Mode

In this mode, the drive controls the direction of the transfer by asserting or negating I/O. When I/O is asserted, information is transferred from the drive to the initiator. When I/O is negated, information is transferred from the initiator to the drive.

To transfer information to the initiator with I/O true:

1. The drive sets DB(0-7, P), waits for at least one DESKEW DELAY plus a CABLE SKEW DELAY, then asserts REQ.
2. After REQ goes true, the initiator reads DB(0-7, P) and then acknowledges receipt of the data by asserting ACK.
3. When the drive detects ACK as true, it negates REQ and can change DB(0-7, P).
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(0-7, P) and asserting REQ.

To transfer information from the initiator with I/O false:

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

6.2.7.3 Synchronous Transfer Mode

Synchronous mode is available only for the DATA IN and DATA OUT phases. Because synchronous transfers require an exchange of extended messages to establish the REQ/ACK offset and transfer period, only devices that support extended messages can use synchronous transfers.

In synchronous transfers, bytes are transferred at a rate negotiated between an initiator and the drive, with the REQ/ACK handshake used to control the transfer rate, rather than validate bytes on the bus. The factors that define the parameters of a synchronous transfer are hardware dependent. These parameters are established between the drive and an initiator by exchanging a SYNCHRONOUS DATA TRANSFER REQUEST (SDTR) message. (See Section 6.3.6.1, “SYNCHRONOUS DATA TRANSFER REQUEST (01h).”)

The initiator asserts one ACK pulse for each REQ pulse received. A successful synchronous transfer is completed when the number of REQ pulses and the number of ACK pulses are equal at the end of the transfer.

On all synchronous data transfers:

1. The drive asserts REQ for at least an ASSERTION PERIOD whenever it is required for the drive to do so in the individual transfer descriptions listed below. Before asserting REQ again, the drive waits for either a TRANSFER PERIOD, measured from the last transition of REQ to true, or a NEGATION PERIOD, measured from the last transition of REQ to false, depending on which extends further in time.
2. As soon as it receives the leading edge of a REQ pulse, the initiator can assert ACK. ACK remains asserted for at least an ASSERTION PERIOD whenever the initiator asserts it. Before asserting ACK again, the initiator waits for either a TRANSFER PERIOD, measured from the last transition of ACK to true, or a NEGATION PERIOD, measured from the last transition of ACK to false, whichever extends further in time.
3. For data transfers from the drive to the initiator, data is valid when REQ is asserted. For data transfers from the initiator to the drive, data is valid when ACK is asserted.

To transfer information to the initiator with I/O true:

1. The drive sets DB(0-7, P), waits for at least one DESKEW DELAY plus one CABLE SKEW DELAY, then asserts REQ.
2. Following the assertion of REQ, DB(0-7, P) remain valid for at least one DESKEW DELAY plus one CABLE SKEW DELAY plus one HOLD TIME.
3. The initiator reads the value on DB(0-7, P) within one HOLD TIME of the transition of REQ to true.

To transfer information from the initiator with I/O false:

1. The initiator transfers one byte for each REQ pulse received. On receiving each REQ pulse, the initiator might immediately set DB(0-7, P), wait for at least one DESKEW DELAY plus one CABLE SKEW DELAY, then assert ACK. If this is not done immediately, it must be done eventually. The REQ pulses must be serviced in the order received.
2. After the assertion of ACK, DB(0-7, P) remain valid for at least one DESKEW DELAY plus one CABLE SKEW DELAY plus one HOLD TIME.
3. The drive reads the value on DB(0-7, 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–an ATTENTION condition and a RESET condition. These conditions cause a SCSI device to act and might alter the phase sequence.

6.2.8.1 ATTENTION Condition

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

The initiator can enter an ATTENTION condition at any time by asserting ATN, except during the BUS FREE or ARBITRATION phase. The drive responds with the MESSAGE OUT 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 the 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

A RESET condition (RST) clears all SCSI devices from the bus, and takes precedence over all other phases and conditions. The BUS FREE phase always follows a RESET condition. SCSI devices create a RESET condition by asserting RST for at least a RESET HOLD TIME. While a RESET condition is in effect, the states of all SCSI bus signals other than RST remain undefined.

Note: The Empire 540/1080S hard disk drive never asserts RST, but responds to RST asserted by another SCSI device.

Upon detecting a RESET condition, the drive releases all SCSI bus signals within a BUS CLEAR DELAY of the transition of RST to true and performs a hard reset. When the drive executes a hard reset, it:

- Clears all I/O processes.
- Releases all SCSI device reservations.
- Sets a UNIT ATTENTION condition.

Operating mode parameters are restored to their last saved values if saved values have been established. If operating mode parameters have not been saved or cannot be saved, they will be restored to their default values.

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

6.2.9 Logical SCSI Conditions

6.2.9.1 UNIT ATTENTION Condition

The drive generates a UNIT ATTENTION condition for each initiator whenever it is reset by a BUS DEVICE RESET message, a RESET condition, or by a power-on reset. This condition is also generated for an initiator when the operating mode parameters in effect for this initiator have been changed by another initiator, tagged commands queued for this initiator were cleared by another initiator, or any other event occurs that requires the initiator's attention.

For each initiator that the drive has a UNIT ATTENTION condition for, the drive terminates the first command sent to it with CHECK CONDITION status, unless the command is an INQUIRY or REQUEST SENSE command. The sense key is set to UNIT ATTENTION. If the first command sent is a REQUEST SENSE command, the UNIT ATTENTION condition is reported in the sense data. An INQUIRY command will execute as if the UNIT ATTENTION condition did not exist, and cause the next command to become the "first command." Once the condition is reported, it is cleared for that initiator.

6.2.9.2 CONTINGENT ALLEGIANCE Condition

The CONTINGENT ALLEGIANCE condition exists following the return of CHECK CONDITION status by the drive. The drive clears the condition when it detects a RESET condition, receives a BUS DEVICE RESET message, or receives an ABORT message or any subsequent command from the initiator it is saving the sense data for. (See Section 6.8.3, "REQUEST SENSE (Opcode=03h).")

While the CONTINGENT ALLEGIANCE condition exists, the drive responds with BUSY status to any other initiator that requests access to it. Execution of all tagged I/O processes in the drive are halted until the CONTINGENT ALLEGIANCE condition is cleared.

6.3 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. This is accomplished by using a system of pre-defined software messages. Messages can be one-byte, two-byte, or extended. The first byte of the message determines the format of the message, as shown below in Table 6-2.

Table 6-2 *Message Format*

BYTE(S)	MESSAGE FORMAT
00h	One-byte message (COMMAND COMPLETE)
01h	Extended messages
02h–1Fh	One-byte messages
20h–2Fh	Two-byte messages
30h–7Fh	Reserved
80h–FFh	One-byte message (IDENTIFY)

6.3.1 MESSAGE IN Phase

The MESSAGE IN phase allows the drive to send messages to the initiator. During the REQ/ACK handshakes of this phase, the drive asserts MSG, C/D, and I/O.

6.3.2 MESSAGE OUT Phase

The MESSAGE OUT phase allows the drive to request that messages be sent from the initiator to the drive. The drive invokes this phase in response to an ATTENTION condition created by the initiator. During the REQ/ACK handshakes of this phase, the drive asserts MSG and C/D, and negates I/O. The drive handshakes bytes until ATN is negated 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 the drive goes to the BUS FREE phase–for example, in response to an ABORT or BUS DEVICE RESET message.

See Section 6.3.7, "Message Error Handling" for a description on the drive's handling of messages with parity errors.

Table 6-3 *SCSI Messages Supported*

CODE	SUPPORT		MESSAGE NAME	DIRECTION	NEGATE ATN BEFORE LAST ACK
	INITIATOR	TARGET			
00h	M	M	COMMAND COMPLETE	In	–
01h	O	O	EXTENDED MESSAGE, SYNC OR WIDE[2]	In/Out	–/Yes
02h	O	O	SAVE DATA POINTER	In	–
04h	O	O	DISCONNECT	In	–
05h	M	M	INITIATOR DETECTED ERROR	Out	Yes
06h	O	M	ABORT	Out	Yes
07h	M	M	MESSAGE REJECT	Out	Yes
08h	M	M	NO OPERATION	In/Out	–/Yes
09h	M	M	MESSAGE PARITY ERROR	Out	Yes
0Ch	O	M	BUS DEVICE RESET	Out	Yes
0Dh	O	O	ABORT TAG[1]	Out	Yes
0Eh	O	O	CLEAR QUEUE[1]	Out	Yes
80h	M	O	IDENTIFY	In	Yes
80h, C0h	M	M	IDENTIFY	Out	No
Queue Tag Messages (Two Bytes)					
20h	O	O	SIMPLE QUEUE TAG	In/Out	–/No
21h	O	O	HEAD OF QUEUE TAG	Out	No
22h	O	O	ORDERED QUEUE TAG	Out	No

Notes: [1] The ABORT TAG and CLEAR QUEUE messages are required with command queuing.
[2] Extended message.
M = Mandatory
O = Optional
In = Drive to Initiator
Out = Initiator to Drive
Yes = Initiator negates ATN while REQ is asserted on the message's last REQ/ACK handshake
No = Initiator might not negate ATN before the last ACK of the message
– = Not applicable

6.3.3 Message Protocol

All SCSI devices implement the COMMAND COMPLETE message. Until an initiator informs the drive that it supports other messages, the drive sends only COMMAND COMPLETE messages to that initiator. By asserting ATN during SELECTION, an initiator communicates to the drive that it supports messages other than COMMAND COMPLETE.

When the drive detects ATN asserted by the initiator during the SELECTION phase, it enters the MESSAGE OUT phase to receive a message. The drive expects the initiator to send an IDENTIFY message to establish an I_T_L nexus. This message should specify LUN=0. If the initiator supports disconnect/reconnect, it should also set bit 6 of the IDENTIFY message. Alternatively, the drive can receive an ABORT or BUS DEVICE RESET message. If the drive receives any other message, it sends a MESSAGE REJECT message and goes to BUS FREE. Following the MESSAGE OUT phase in which the IDENTIFY message is received, the drive enters the COMMAND phase and requests command bytes from the initiator.

6.3.4 One-Byte Messages

One-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 one-byte messages, the hex value of the byte to send corresponds to the code assigned to the message, as shown in Table 6-3, and as explained in the following paragraphs.

6.3.4.1 COMMAND COMPLETE (00h)

The COMMAND COMPLETE message indicates that the drive has finished execution of a command and passed status information to the initiator. It indicates only that the command was ended; 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. After successfully sending this message, the drive goes to the BUS FREE phase by releasing BSY.

6.3.4.2 SAVE DATA POINTER (02h)

The SAVE DATA POINTER message is sent from the drive to direct the initiator to copy the active data pointer to the saved data pointer for the current I/O process. (See Section 6.2.7.1, "Pointers.") If the initiator rejects this message, the drive discontinues attempts to disconnect.

During a DATA IN or DATA OUT phase, when breaking a long data transfer into two or more shorter transfers using DISCONNECT messages, the drive issues a SAVE DATA POINTER message immediately before transferring the DISCONNECT message.

6.3.4.3 DISCONNECT (04h)

The drive sends the DISCONNECT message to inform the 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.

If the initiator detects the BUS FREE phase not preceded by a DISCONNECT or COMMAND COMPLETE message, and not resulting from a RESET condition, or an ABORT, ABORT TAG, BUS DEVICE RESET, or CLEAR QUEUE message, a catastrophic error condition exists.

Note: The DISCONNECT message does not cause the initiator to save the data pointer.

6.3.4.4 INITIATOR DECTECTED ERROR (05h)

The INITIATOR DETECTED ERROR message informs the drive that an error has occurred. The source of the error can be related either to previous activities on the SCSI bus or can be internal to the initiator and unrelated to any previous SCSI bus activity. The drive proceeds with the current command and returns CHECK CONDITION status at its completion. The sense key is set to ABORTED COMMAND and the additional sense code is set to INITIATOR DETECTED ERROR MESSAGE RECEIVED.

6.3.4.5 ABORT (06h)

If an I_T_L or I_T_L_Q nexus exists, the ABORT message tells the drive to abort all I/O processes, active or queued, for the initiator that sends this message. Previously established conditions, including reservations and operating mode parameters are not changed. The drive goes to the BUS FREE phase following successful receipt of this message.

See also Section 6.3.4.9, "BUS DEVICE RESET (0Ch)," Section 6.3.4.10, "ABORT TAG (0Dh)," and Section 6.3.4.11, "CLEAR QUEUE (0Eh)" for information on the BUS DEVICE RESET, ABORT TAG, and CLEAR QUEUE messages.

6.3.4.6 MESSAGE REJECT (07h)

The MESSAGE REJECT message is sent from either the initiator or the drive to indicate that the last message or message byte it received was inappropriate or has not been implemented.

To indicate its intentions of sending this message, the initiator asserts ATN prior to its release of ACK for the REQ/ACK handshake of the message byte that is to be rejected. If the drive receives this message under any other circumstance, it rejects this message.

When the drive sends this message, it changes to the MESSAGE IN phase and sends this message prior to requesting additional message bytes from the initiator. This provides an interlock so that the initiator can determine which message byte is rejected.

6.3.4.7 NO OPERATION (08h)

The NO OPERATION message is sent from the initiator in response to the drive's request for a message when the initiator does not currently have any other valid message to send.

6.3.4.8 MESSAGE PARITY ERROR (09h)

The MESSAGE PARITY ERROR message is sent from the initiator to the drive to indicate that it received a message byte with a parity error.

To indicate its intentions of sending this message, the initiator asserts ATN prior to its release of ACK for the REQ/ACK handshake of the message byte that has the parity error.

When the drive receives a MESSAGE PARITY ERROR message, it retries the operation by re-sending the original message. If the message still cannot be sent successfully, the drive goes to the BUS FREE phase and aborts the current command. The sense key is set to ABORTED COMMAND and the additional sense code is set to SCSI PARITY ERROR.

6.3.4.9 BUS DEVICE RESET (0Ch)

The BUS DEVICE RESET message orders the drive to abort all I/O processes. This message forces a RESET condition on the drive. A UNIT ATTENTION condition is generated for all initiators. The drive goes to the BUS FREE phase following successful receipt of this message.

See also Section 6.3.4.5, "ABORT (06h)," Section 6.3.4.10, "ABORT TAG (0Dh)," and Section 6.3.4.11, "CLEAR QUEUE (0Eh)" for information on the ABORT, ABORT TAG, and CLEAR QUEUE messages.

6.3.4.10 ABORT TAG (0Dh)

If an I_T_L_Q nexus exists, the ABORT TAG message forces the drive to abort the current I/O process. If the drive has already started execution of the I/O process, the execution halts. The medium contents might have been modified before execution was halted. In either case, any pending status or data for the I/O process is cleared and no status or ending message is sent to the initiator. Previously established conditions, including reservations and operating mode parameters are not changed. The drive goes to the BUS FREE phase following successful receipt of this message.

See also Section 6.3.4.5, "ABORT (06h)," Section 6.3.4.9, "BUS DEVICE RESET (0Ch)," and Section 6.3.4.11, "CLEAR QUEUE (0Eh)" for information on the ABORT, BUS DEVICE RESET, and CLEAR QUEUE messages.

6.3.4.11 CLEAR QUEUE (0Eh)

The CLEAR QUEUE message informs the drive to abort all I/O processes, from all initiators, in the queue and terminate any active I/O processes. The drive performs an action equivalent to receiving an ABORT message from each initiator. The medium might have been altered by partially executed commands. All pending status and data for all initiators are cleared. No status or message is sent for any of the I/O processes. A UNIT ATTENTION condition is generated for all other initiators with I/O processes that were either active or queued in the drive. Previously established conditions, including reservations and operating mode parameters are not changed. The drive goes to the BUS FREE phase following successful receipt of this message.

See also Section 6.3.4.5, "ABORT (06h)," Section 6.3.4.9, "BUS DEVICE RESET (0Ch)," and Section 6.3.4.10, "ABORT TAG (0Dh)" for information on the ABORT, BUS DEVICE RESET, and ABORT TAG messages.

6.3.4.12 IDENTIFY (80h–FFh)

The IDENTIFY message is sent by either the initiator or the drive to further define the I_T nexus. The resulting I_T_L nexus must have LUN=0.

The identify bit (Bit 7) is set to one to specify that this is an IDENTIFY message.

A disconnect privilege bit (Bit 6) of one specifies that the initiator has granted the drive the privilege of disconnection. A disconnect privilege bit of zero notifies the drive that it should not disconnect. This bit is not defined and is set to zero when an IDENTIFY message is sent by the drive.

The logical unit target bit (Bit 5) is set to zero to indicate that the number in the LUNTRN field is a LUN.

Bits 4 and 3 are reserved and are set to zero.

The initiator must set the LUNTRN field (Bits 2-0) to zero when it sends this message to the drive. When this message is sent to the initiator, the drive sets the LUNTRN field to zero.

The drive always sends an IDENTIFY message of 80h, whereas the initiator can send an IDENTIFY message of C0h or 80h depending on whether disconnection is allowed or not allowed, respectively. When the drive sends an IDENTIFY message following RESELECTION, the initiator will automatically restore pointers. (See Section 6.2.7.1, "Pointers.")

6.3.5 Two-Byte Messages

Tagged two-byte messages, shown in Table 6-4, allow the drive to handle multiple data streams for an initiator. The drive keeps the messages and associated commands in an ordered list known as a queue. The messages supported are HEAD OF QUEUE TAG, ORDERED QUEUE TAG, and SIMPLE QUEUE TAG.

These messages specify an identifier, known as a queue tag, for an I/O process that establishes an I_T_L_Q nexus. The queue tag field is an 8-bit unsigned integer assigned by the initiator during initial connection. The numeric value of a queue tag does not influence the order of execution, and the queue tag is available for reassignment when the I/O process ends.

Upon connection to the drive, the initiator sends the appropriate queue tag message immediately following the IDENTIFY message and within the same MESSAGE OUT phase.

When the drive reconnects to an initiator to continue a tagged I/O process, the SIMPLE QUEUE TAG message is sent following the IDENTIFY message and within the same MESSAGE IN phase.

Table 6-4 *Queue Tag Message Format*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	MESSAGE CODE (20h, 21h, 22h)							
1	QUEUE TAG (00h-FFh)							

6.3.5.1 SIMPLE QUEUE TAG (20h)

The SIMPLE QUEUE TAG message, when sent to the drive, specifies that the I/O process be placed in the drive's command queue. The command is executed in an order determined by the drive within the constraints of the algorithm responsible for queue management. This algorithm is specified in the Control Mode Page (Page Code 0Ah).

When sent to the initiator, this message just indicates the value of the queue tag so that the initiator can determine which I/O process is reconnecting.

6.3.5.2 HEAD OF QUEUE TAG (21h)

The HEAD OF QUEUE TAG message specifies that the I/O process be placed first in the drive's command queue. No preemption will be done IF an I/O process IS already being executed. All I/O processes subsequently received with a HEAD OF QUEUE TAG message are placed at the head of the command queue in a last-in, first-out order.

6.3.5.3 ORDERED QUEUE TAG (22h)

The ORDERED QUEUE TAG message indicates that the I/O process be placed in the drive's command queue for execution in the order received. All queued I/O processes for the drive received prior to this I/O process will be executed before this I/O process. All queued I/O processes received after this I/O process will be executed following this I/O process, except for I/O processes received with a HEAD OF QUEUE TAG message.

6.3.6 Extended Messages

Extended messages provide more advanced methods of bus control, command execution management, and vendor unique functions. For extended messages, byte 0 is always set to 01h, the code for an "extended message."

The basic structure of an extended message is shown in Table 6-5, and the extended message codes are given in Table 6-6.

Table 6-5 *Extended Message Data Structure*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	EXTENDED MESSAGE=01h							
1	EXTENDED MESSAGE LENGTH							
2	EXTENDED MESSAGE CODE							
3+	EXTENDED MESSAGE ARGUMENTS							

EXTENDED MESSAGE (Byte 0)

This field contains the code for an extended message, and is therefore set to 01h.

EXTENDED MESSAGE LENGTH (Byte 1)

This parameter specifies the number of extended message bytes to follow. The length of the entire message is equal to this value plus two bytes. A value of zero in this field indicates that 256 bytes follow.

EXTENDED MESSAGE CODE (Byte 2)

This field contains the code defining the message, as shown in Table 6-6.

EXTENDED MESSAGE ARGUMENTS (Bytes 3+)

This group of fields contains the execution arguments associated with the action ordered by the extended message. There can be several argument bytes in this group.

Table 6-6 *Extended Message Codes*

CODE	MESSAGE
00h	Not supported (MODIFY DATA POINTER)
01h	SYNCHRONOUS DATA TRANSFER REQUEST
02h	Reserved
03h	WIDE DATA TRANSFER REQUEST (always negotiate for an 8-bit bus)
04h–7Fh	Reserved
80h–FFh	Not supported (Vendor unique)

6.3.6.1 SYNCHRONOUS DATA TRANSFER REQUEST (01h)

The SYNCHRONOUS DATA TRANSFER REQUEST (SDTR) message is used to establish the timing parameters for synchronous data transfers between two SCSI devices capable of performing such transfers. It is the responsibility of the initiator and the drive to establish an agreement when one is required or to negotiate a new data transfer agreement if one is needed.

An agreement becomes invalid after a RESET condition, a BUS DEVICE RESET message, or a power cycle. The agreement only applies to the DATA IN and DATA OUT phases. The default transfer mode is asynchronous data transfer mode. The default data transfer mode is entered at power on, after a BUS DEVICE RESET message, or after a RESET condition.

Table 6-7 *Synchronous Data Transfer Request Data Structure*

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

TRANSFER PERIOD FACTOR (Byte 3)

This parameter times four is the value of the transfer period in nanoseconds. The transfer period is the minimum time between leading edges of successive REQ pulses and of successive ACK pulses. Table 6-8 shows the transfer rates the Empire 540/1080S hard disk drive will use based on the requested transfer period factor.

Table 6-8 *Transfer Rates as a Function of Requested Transfer Period Factor*

REQUESTED TRANSFER PERIOD FACTOR	TRANSFER RATE USED (MB/s)
0 – 25	10.00
26 – 31	8.00
32 – 37	6.70
38 – 43	5.70
44 – 50	5.00
51 – 62	4.00
63 – 75	3.30
76 – 86	2.90
87 – 100	2.50
101 – 112	2.20
113 – 125	2.00
126 – 147	1.70
148 – 178	1.40
179 – 200	1.25
201 – 227	1.10

REQ/ACK OFFSET (Byte 4)

This parameter states the maximum number of REQ pulses allowed to be outstanding at the drive before the leading edge of another ACK pulse is received. A value of zero indicates asynchronous data transfer mode; a value of FFh indicates that no limit exists.

The originating device (the device that sends the first of the pair of SDTR messages) sets its values according to the rules above to permit it to receive data successfully. If the responding device can also receive data successfully with these values, it returns the same values in its SDTR message. If it requires a larger transfer period, a smaller REQ/ACK offset, or both in order to receive data successfully, the responding device substitutes values in its SDTR message as required, returning unchanged any value that does not need to be altered. Each device when transmitting data must respect the limits set by the exchange, but it is permitted to transfer data with larger transfer periods, smaller REQ/ACK offsets, or both. Table 6-9 shows the implied agreement from the successful completion of an exchange of SDTR messages.

Table 6-9 *Implied Agreement from SDTR Messages*

RESPONDING DEVICE SDTR MESSAGE	IMPLIED AGREEMENT
1. Non-zero REQ/ACK offset.	Each device transmits data with a transfer period equal to or greater than, and a REQ/ACK offset equal to or less than the values in the responding device's SDTR message.
2. REQ/ACK offset equal to zero.	Asynchronous transfer.
3. MESSAGE REJECT message.	Asynchronous transfer.

If the initiator recognizes that negotiation is required, it asserts ATN and sends a SDTR message to begin the negotiating process. After successfully completing the MESSAGE OUT phase, the drive responds with the proper SDTR message. If an abnormal condition prevents the drive from returning an appropriate response, both devices go into asynchronous data transfer mode for data transfers between the two devices.

During the MESSAGE IN phase in which the drive sends a SDTR response, the implied agreement for synchronous operation is considered to be negated by both the initiator and the drive if the initiator asserts ATN and the first message out is either MESSAGE PARITY ERROR or MESSAGE REJECT.

In this case, both devices go to asynchronous data transfer mode for data transfers between the two devices. For the MESSAGE PARITY ERROR case, the implied agreement is reinstated if the drive can successfully re transmit its SDTR response.

If the drive recognizes that negotiation is required and the send synchronous message bit on the Quantum-Unique Control Parameters page (Page Code 37h) is set to one, it sends an SDTR message to the initiator. Prior to releasing ACK on the last byte of the SDTR message from the drive, the initiator asserts ATN and respond with either its SDTR response or a MESSAGE REJECT message. If an abnormal condition prevents the initiator from returning an appropriate response, both devices go to asynchronous data transfer mode for data transfers between the two device.

If, immediately following the initiator's responding SDTR message, the drive shifts to MESSAGE IN phase and the first message in is MESSAGE REJECT, the implied agreement is considered to be negated and both devices go to asynchronous data transfer mode for data transfers between the two devices.

6.3.6.2 WIDE DATA TRANSFER REQUEST (03h)

A WIDE DATA TRANSFER REQUEST (WDTR) message exchange is initiated by a SCSI device whenever a previously arranged transfer width agreement may have become invalid. The agreement becomes invalid after any condition which may leave the data transfer agreement in an indeterminate state such as:

1) after a hard reset condition;

2) after a BUS DEVICE RESET message and;

3) after a power cycle.

In addition, a SCSI device may initiate an WDTR message exchange whenever it is appropriate to negotiate a new transfer width agreement. The Empire 540/1080S hard disk drive always negotiates for an eight-bit wide bus. Table 6-10 shows the structure of the Wide Data Transfer Request.

Table 6-10 *Wide Data Transfer Request Data Structure*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	EXTENDED MESSAGE=01h							
1	EXTENDED MESSAGE LENGTH=02h							
2	WIDE DATA TRANSFER REQUEST CODE=01h							
3	TRANSFER WIDTH EXPONENT							

TRANSFER WIDTH EXPONENT (Byte 3)

The transfer width is two to the transfer width exponent bytes wide. The transfer width that is established applies to all logical units on both SCSI devices. Valid SCSI transfer widths are 8 bits (transfer width exponent = 00h), 16 bits (transfer width exponent = 01h), and 32 bits (transfer width exponent = 02h). Values of transfer width exponent greater than 02h are reserved. The only valid value for the Empire 540/1080S hard disk drive is 00h.

6.3.7 Message Error Handling

If the drive detects one or more parity errors in the message bytes it receives, it requests the initiator to send them again. The drive does this by asserting REQ after ATN goes false, but before changing to another phase. On detecting REQ, the initiator sends all of the previous message bytes in the same order as previously sent during this phase.

When re-sending multiple message bytes, the initiator must assert 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 remainder of the MESSAGE OUT phase. When the initiator re-sends a series of messages, the drive does not act on any message already received prior to the parity error.

6.4 COMMAND IMPLEMENTATION

This section defines the SCSI command implementation supported by the Empire 540/1080S 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 is used to indicate that additional information is available. The initiator can issue a REQUEST SENSE command to retrieve that information.

The following subsections describe the general format of a command. A list also is provided in Table 6-11 to show the commands supported by the Empire 540/1080S hard disk drive.

6.4.1 Command Descriptor Block (CDB)

SCSI commands are issued from an initiator by transferring a Command Descriptor Block (CDB) to a target device. For some commands, a parameter list sent during the DATA OUT phase accompanies the request. A CDB contains an opcode, a Logical Unit Number (LUN), fields containing command parameters, and a control byte.

6.4.1.1 Opcode

The first byte of a CDB is the opcode. It 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 command groups. The command code field contains the code for one of the 32 commands in each group. The two fields together provide 256 possible code combinations. The Empire 540/1080S hard disk drive supports the following command groups:

- Group 0, which is six-byte commands.
- Groups 1 and 2, which are ten-byte commands.

The opcode for each command indicates its format–six-byte commands have opcodes between 00h and 1Fh; ten-byte commands have opcodes between 20h and 5Fh. The opcodes for the commands supported by the drive are given in Table 6-11.

Table 6-11 *SCSI Commands Supported*

OPCODE	COMMAND	OPCODE	COMMAND
00h	TEST UNIT READY	25h	READ CAPACITY
01h	REZERO UNIT	28h	READ EXTENDED
03h	REQUEST SENSE	2Ah	WRITE EXTENDED
04h	FORMAT UNIT	2Bh	SEEK EXTENDED
07h	REASSIGN BLOCKS	2Eh	WRITE AND VERIFY
08h	READ	2Fh	VERIFY
0Ah	WRITE	35h	SYNC CACHE
0Bh	SEEK	37h	READ DEFECT DATA
12h	INQUIRY	3Bh	WRITE BUFFER
15h	MODE SELECT	3Ch	READ BUFFER
16h	RESERVE	3Eh	READ LONG
17h	RELEASE	3Fh	WRITE LONG
1Ah	MODE SENSE	40h	CHANGE DEF
1Bh	START/STOP UNIT	55h	MODE SELECT EXTENDED
1Dh	SEND DIAGNOSTIC	5Ah	MODE SENSE EXTENDED

6.4.1.2 Logical Unit Number (LUN)

The LUN for the Empire 540/1080S hard disk drive is zero. If any command other than INQUIRY or REQUEST SENSE specifies LUN not equal to zero, the drive terminates that command with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense code is set to LOGICAL UNIT NOT SUPPORTED.

6.4.1.3 Command Parameters

Command parameters can include Logical Block Address (LBA), transfer length, parameter list length, and allocation length.

Logical Block Address (LBA)

The LBA specifies the first or starting block of an operation. Group 0 CDBs contain 21-bit LBAs. Groups 1 and 2 CDBs contain 32-bit LBAs. LBAs begin with block zero and reference contiguous blocks up to the last logical block on the drive. The last logical block on the drive can be determined by issuing a READ CAPACITY command.

A logical block is 512 bytes in length for the Empire 540/1080S hard disk drive.

Transfer Length

The transfer length specifies the amount of data to be transferred to or from an initiator. It is translated as the number of blocks to be transferred. For six-byte commands, up to 256 blocks can be transferred. A value between 1 and 255 (01h-FFh) indicates the number of logical blocks to be transferred. The value zero causes 256 logical blocks to be transferred. For ten-byte commands, up to 65,535 blocks can be transferred. A value between 0 and 65,535 (0000h-FFFFh) indicates the number of logical blocks to be transferred.

Parameter List Length

The parameter list length specifies the number of bytes to be sent during the DATA OUT phase. A parameter list length of zero notifies the drive to transfer no data. This CDB field allows an initiator to send a list of parameters to the drive.

Allocation Length

The allocation length specifies the maximum number of bytes allocated by an initiator for returned data. An allocation length of zero indicates that no bytes are to be transferred by the drive. The drive terminates the DATA IN phase when it has transferred to the initiator either the number of bytes specified by the allocation length or all available data, whichever is less.

6.4.1.4 Control Byte

Table 6-12 *Control Byte*

BYTE	BIT							
	7	6	5	4	3	2	1	0
1	VU=0		RESERVED=0				F	L

F–Flag (Bit 1)

The flag bit should be set to zero if the link bit is zero. If the link bit is zero and the flag bit is one, the drive terminates the command with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST.

If the flag bit is zero and the link bit is one, and if the drive finishes execution of the command, the drive sends a LINKED COMMAND COMPLETE message. If the flag bit is one and the link bit is one, the drive sends a LINKED COMMAND COMPLETE [WITH FLAG] message upon command completion.

L–Link (Bit 0)

The link bit is used to continue the I/O process across multiple commands. A link bit of one indicates that the initiator requests the drive to enter the COMMAND phase upon successful completion of the current command. A link bit of zero indicates that the initiator requests the drive to enter the BUS FREE phase upon completion of the current command.

The Empire 540/1080S hard disk drive does not support linked commands. Therefore, the flag and link bits must be set to zero by the initiator.

6.5 SCSI STATUS

At the completion of each command, a status byte is sent from the drive to the initiator during the STATUS phase, unless one of the following events terminates the command:

- An ABORT message
- An ABORT TAG message
- A BUS DEVICE RESET message
- A CLEAR QUEUE message
- A RESET condition
- An unexpected disconnect

Table 6-13 *Status Byte*

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

6.5.1 Good (00h)

The drive successfully completed a command, without encountering an error condition.

6.5.2 Check Condition (02h)

An error, exception, or abnormal condition occurred during the drive's execution of a command, which might not have completed successfully. To determine the nature of the abnormal condition, the initiator must issue a REQUEST SENSE command.

6.5.3 Busy (08h)

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 accept an additional command at this time. The initiator must reissue the command.

6.5.4 Intermediate (10h)

The drive successfully completed a linked command, without encountering an error condition. Typically, the drive issues this status code for each successful command in a series of linked commands except for 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.

This status is not sent by the Empire 540/1080S hard disk drive.

6.5.5 Reservation Conflict (18h)

The drive has been reserved with a conflicting reservation type for another SCSI device. The initiator must reissue the command.

6.5.6 Queue Full (28h)

The command queue is full and the drive cannot accept another command from an initiator until another slot opens in the queue. This status is returned whenever the drive receives a HEAD OF QUEUE TAG, SIMPLE QUEUE TAG, or ORDERED QUEUE TAG message and the command queue is full. The I/O process is not placed in the command queue.

6.6 READING AND WRITING

6.6.1 Logical Block

The Empire 540/1080S 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. Each logical block has a unique address–its Logical Block Address (LBA). 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.

6.6.2 Logical Block Address (LBA) Ranges

Table 6-14 *LBA Ranges*

DRIVE CAPACITY	BLOCK SIZE	ADDRESS RANGE (DECIMAL)	ADDRESS RANGE (HEXADECIMAL)
540S	512 (200h)	0 – 1,054,687	0–1017DF
1080S	512 (200h)	0 – 2,109,375	0–202FBF

If an initiator attempts to access a logical block at an address outside the valid address range, the drive terminates the command with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense code is set to INVALID LOGICAL BLOCK ADDRESS.

6.6.3 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, an initiator specifies the LBA of the first logical block to be transferred and the total number of logical blocks to be transferred. For the number of logical blocks to be transferred, the drive interprets the value zero differently, depending on the form of READ or WRITE command used. (See Table 6-15.)

Table 6-15 *Interpretation of Zero*

COMMAND	OPCODE	GROUP	TRANSFER LENGTH OF ZERO
READ	08h	0	Drive transfers 256 blocks of data.
WRITE	0Ah	0	Drive transfers 256 blocks of data.
READ EXTENDED	28h	1	Drive transfers no data.
WRITE EXTENDED	2Ah	1	Drive transfers no data.

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

Number of data bytes transferred = 512 bytes * transfer length.

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

6.6.3.1 RAM Buffer

The drive contains a 512Kbyte DRAM buffer that is used for temporary storage of data. All read and write operations, involving data transfer, channel data through this buffer.

The buffer is divided into five segments:

1. 64Kbytes to store variables for the servo DSP
2. 32Kbytes to store "irregular write" data, such as Write Long, Write Physical, and Head of Queue write commands
3. 32Kbytes to store system data, temporary data, and queue commands from the host
4. 128Kbytes to store "regular write" data (write cache data)
5. 256Kbytes to store read data (read cache data)

6.6.3.2 Read Operation

On a normal read operation, the requested data is located on the disk, transferred to the read buffer, and then sent to the host interface via the SCSI bus. Two independent procedures occur during this process. The first part involves data transfer from the disk to the read buffer. The second part deals with data transfer from the read buffer to the SCSI bus.

The drive transfers data from the disk to the buffer, until the amount of data transferred from the disk matches the transfer length. The drive ensures that the portion of the buffer dedicated to this read operation does not overflow. If the area allocated 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.

Data received from the disk is simultaneously transferred out of the buffer to the SCSI bus, after it has been validated. This action continues until the amount of data transferred from the buffer matches the transfer length. If disconnection is allowed by the initiator, the drive disconnects whenever the portion of the buffer reserved for this read operation becomes empty, provided that all requested data has not been transferred. When the percentage of this space that is full matches the buffer full ratio on the Disconnect-Reconnect Page (Page Code 02h), the drive initiates a reconnection. The drive also attempts to reconnect if all the data required by the initiator has been transferred to the buffer. Once the drive reconnects to the bus, it resumes data transfer from the buffer to the SCSI bus. If the initiator does not allow disconnection, the drive temporarily halts data transfer to the SCSI bus until data becomes available.

A caching system is provided to improve performance. The new features introduced by this system, that effect the normal read operation, are explained in the next two subsections.

6.6.3.3 Read Cache

On all read operations, before accessing the disk, the drive checks to see if read caching is enabled. For read caching, if the drive locates the desired data in the buffer, the drive immediately transfers the data to the initiator without accessing the disk. This bypasses the procedure of transferring the data from the disk to the read buffer that normally occurs on execution of a read operation, significantly reducing command overhead.

6.6.3.4 Pre-Fetch

On normal read operations, the drive transfers data from the disk to the read buffer, until the amount of data transferred from the disk matches the transfer length. Pre-fetch enables the drive to continue reading beyond the requested data in anticipation of sequential reading, until it fills up the read buffer.

6.6.3.5 Write Operation

On a normal write operation, the data to be written is sent from the initiator to the drive, stored in the write buffer, and then transferred to the disk. Two independent procedures occur during this process. The two procedures described reverse the actions of the two independent procedures that happen on a read operation. The first part involves data transfer from the SCSI bus to the write buffer. The second part deals with data transfer from the write buffer to the disk.

The drive transfers data from the SCSI bus to the buffer, until the amount of data transferred matches the transfer length. If disconnection is allowed by the initiator, the drive disconnects whenever the portion of the buffer dedicated to this write operation is full or the amount of data transferred equals the transfer length. For the first condition, when the percentage of this space that is empty matches the buffer empty ratio on the Disconnect-Reconnect Page (Page Code 02h), the drive initiates a reconnection. The drive also attempts to reconnect if the remaining number of blocks to be transferred will fit into the write buffer or if all the transferred data has been written to the disk. Once the drive reconnects to the bus, it either resumes data transfer from the SCSI bus to the write buffer or completes the command by sending status information accompanied by the appropriate message.

Data received from the SCSI bus is simultaneously transferred out of the write buffer to the disk. This action continues until the amount of data transferred matches the transfer length. The drive ensures that the portion of the buffer reserved for this write operation does not underflow. If the area allocated becomes empty, because data transfer from the write buffer to the disk is faster than from the SCSI bus to the write buffer, the drive temporarily halts data transfer to the disk until data becomes available.

The same caching system provided to improve performance on read operations, contains a feature to improve performance on write operations. The new feature introduced by this system, that effects the normal write operation, is explained in the next subsection.

6.6.3.6 Write Cache

When a write operation is executed with write caching enabled, the drive stores the data to be written in the write buffer and immediately sends a COMMAND COMPLETE message to the initiator, before the data is actually written to the disk. The initiator 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.

The drive immediately writes the cached data to the disk after issuing COMMAND COMPLETE. With write cache, a single-block random write, for example, requires only about 3 ms of initiator time. Without write cache, the same operation would occupy the initiator for about 27 ms.

Write cache allows sequential 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 buffer, 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.

6.6.4 Transfer Rate

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

If the initiator'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 the disk to match their data transfer rate with the initiator's rate. The 512K buffer in the Empire 540/1080S hard disk drive makes this unnecessary. The drive channels all READ and WRITE transfers through the buffer.

The 512K buffer also serves as an onboard cache, called DisCache. (See Section 5.3.1, "Disk Caching.") The MODE SELECT command can enable or disable caching, and configure cache parameters.

6.7 CONFIGURING THE HARD DISK DRIVE

6.7.1 Operating Modes

Users can change some of the drive's operating mode parameters to match their requirements. The MODE SELECT command allows the user to set the operating mode of the drive. The drive's current operating mode is determined by issuing the MODE SENSE command. The MODE SENSE command is also useful for accessing information when setting the operating mode.

As an example, a description of changing a particular operating mode parameter is given. The number of retries attempted by the drive when it detects a read error is a parameter commonly adjusted by the user. First, the initiator issues a MODE SENSE command to determine the drive's current configuration. During the DATA IN phase of the MODE SENSE command, the drive returns the current operating mode parameters to the initiator. This data includes 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 set to a new value.

6.7.2 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 these parameters, the Current Mode Table resides in RAM.

The Current Mode Table in RAM is lost at power off or whenever the drive is reset. Therefore, the drive can save a copy of the Current Mode Table, called the Saved Mode Table, in a reserved area on the disk. Because the Saved Mode Table does not reside in the user data-storage area, a user cannot inadvertently erase it. The drive loads the parameters in the Saved Mode Table into the Current Mode Table at power on or after a RESET condition.

The MODE SELECT command can copy the parameters in the Current Mode Table to the Saved Mode Table. If the MODE SELECT command's save pages bit is set to one, the drive copies all parameters in the Current Mode Table to the Saved Mode Table, after it makes any changes ordered by the MODE SELECT command to the Current Mode Table. Thus, users can permanently configure the drive to their requirements. This eliminates the need for a device driver to send a vendor unique MODE SELECT command following every reset.

The Empire 540/1080S 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 and can be accessed only by the MODE SENSE command. These two tables cannot be modified by the user contrary to the Current Mode Table and the Saved Mode Table.

The Changeable Parameters Table informs an initiator about the drive's operating mode parameters that can be changed. If a parameter in the Current Mode Table can be modified, all bits in its field are set to one. If a parameter cannot be altered, all bits in its field are set to zero.

Note: A software driver should verify that the drive 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.

6.8 COMMAND DESCRIPTIONS

The following subsections provide detailed descriptions of the SCSI commands implemented in the Empire 540/1080S hard disk drive.

6.8.1 TEST UNIT READY (Opcode=00h)

The TEST UNIT READY command, shown in Table 6-16, verifies that the drive is up to speed and prepared to accept commands requiring disk access.

Table 6-16 *TEST UNIT READY Command*

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

The drive returns GOOD status when it is ready to accept a command that requires disk access. If the drive is not ready to accept such a command, or requires initiator action–such as a START STOP UNIT command–to become ready, the drive returns CHECK CONDITION status with the sense key set to NOT READY.

After a maximum delay of one second following power on, the drive can execute commands that do not require disk access. The following commands do not cause the drive to return CHECK CONDITION status with the sense key set to NOT READY:

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

Note: If the drive receives an INQUIRY command before it is ready to access the disk, it returns the default INQUIRY data without reading the disk. (See Section 6.8.9, "INQUIRY (Opcode=12h).")

All other commands might 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 operating mode parameters from reserved cylinders, and calibrated its actuator and servo parameters.

6.8.2 REZERO UNIT (Opcode=01h)

The REZERO UNIT command, shown in Table 6-17, notifies the drive to position the actuator to cylinder zero and head zero.

Table 6-17 *REZERO UNIT Command*

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

6.8.3 REQUEST SENSE (Opcode=03h)

The REQUEST SENSE command, shown in Table 6-18, requests that the drive send sense data to the initiator.

When the drive detects an error during execution of a command, it creates sense data for the initiator and sends CHECK CONDITION status. At this point, a CONTINGENT ALLEGIANCE condition exists, and the drive saves the sense data until this condition is cleared. (See Section 6.2.9.2, "CONTINGENT ALLEGIANCE Condition.") The initiator retrieves the sense data by issuing a REQUEST SENSE command, simultaneously clearing the CONTINGENT ALLEGIANCE condition.

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

6.8.3.1 Command Structure

Table 6-18 *REQUEST SENSE Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=03h							
1	LUN=0			RESERVED=0				
2	RESERVED=0							
3	RESERVED=0							
4	ALLOCATION LENGTH							
5	VU=0		RESERVED=0				F=0	L=0

ALLOCATION LENGTH (Byte 4)

This parameter specifies the maximum number of bytes allocated by the initiator for returned sense data. Acceptable values are 00h-FFh. The drive terminates the 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. An allocation length of 18 (12h), or greater, is recommended for the Empire 540/1080S hard disk drive.

If the allocation length specified prevents the drive from sending all available sense data to the initiator, the remaining data is lost.

6.8.3.2 Extended Sense Data Format

After the drive processes the command bytes, it transitions to the DATA IN phase to send the sense data. The Empire 540/1080S hard disk drive reports the sense data in the Extended Sense Data Format, as shown in Table 6-19.

Table 6-19 *Extended Sense Data Format*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	VLD	ERROR CODE						
1	SEGMENT NUMBER=0							
2	R=0	R=0	ILI	R=0	SENSE KEY			
3	(MSB)	INFORMATION						
4		INFORMATION						
5		INFORMATION						
6		INFORMATION						(LSB)
7	ADDITIONAL SENSE LENGTH=0Ah							
8	(MSB)	COMMAND SPECIFIC INFORMATION						
9		COMMAND SPECIFIC INFORMATION						
10		COMMAND SPECIFIC INFORMATION						
11		COMMAND SPECIFIC INFORMATION						(LSB)
12	ADDITIONAL SENSE CODE							
13	ADDITIONAL SENSE CODE QUALIFIER=0							
14	FIELD REPLACEABLE UNIT CODE=0							
15	SKSV	C/D	R=0	R=0	BPV	BIT POINTER		
16	(MSB)	FIELD POINTER						
17		FIELD POINTER						(LSB)

VLD–VALID (Byte 0, Bit 7)

A valid bit of one indicates that the information field contains good data. A valid bit of zero identifies that the information field is undefined.

ERROR CODE (Byte 0, Bits 6-0)

This parameter reveals if the error being reported is a current error (applies to the current command) or a deferred error (applies to a previous command). If the error is a current error, the drive sets this field to 70h. If it is a deferred error, the drive sets this field to 71h.

Deferred errors result during a command if a previously issued write command encounters an error when writing data to the disk. (See Section 6.6.3.6, "Write Cache.") The drive notifies the initiator of this error in the current command, and points out that it applies to the write command by setting the error code to reflect a deferred error.

ILI–ILLEGAL LENGTH INDICATOR (Byte 2, Bit 5)

Normally, the drive sets this bit to zero, indicating that the requested logical block length matches the logical block length on the disk.

However, on execution of a READ LONG or WRITE LONG command, an incorrect byte transfer length sets this bit to one. (See Section 6.8.29, "MODE SELECT EXTENDED (Opcode=55h)" and Section 6.8.27, "WRITE LONG (Opcode=3Fh).")

SENSE KEY (Byte 2, Bits 3-0)

This parameter provides generic categories in which error and exception conditions can be reported. (See Table 6-19.)

INFORMATION (Bytes 3-6)

Unless specified otherwise, this field holds the Logical Block Address (LBA) associated with the sense key. If the drive issues a sense key of MEDIUM ERROR, this parameter communicates to the initiator the value that should be used when issuing a REASSIGN BLOCKS command to map out the defective sector.

ADDITIONAL SENSE LENGTH (Byte 7)

The drive sets this field to 0Ah to state the remaining number of sense data bytes, not including itself. Therefore, the total number of sense data bytes is 18 (12h).

COMMAND SPECIFIC INFORMATION (Bytes 8-11)

If the drive is unable to successfully complete a REASSIGN BLOCKS command, it issues CHECK CONDITION status with the appropriate sense information. The LBA of the first defect descriptor not reassigned is placed in this field. If information about that defect descriptor is not available, or if all defects have been reassigned, this field is set to FFFFFFFFh.

ADDITIONAL SENSE CODE (Byte 12)

This parameter supplies more detail describing the sense key. (See Table 6-19.)

SKSV–SENSE KEY SPECIFIC VALID (Byte 15, Bit 7)

When SKSV is one, the field pointer, C/D, and BPV are valid. When SKSV is zero, the field pointer, C/D, and BPV are invalid.

FIELD POINTER (Bytes 16&17)

This parameter represents an index into either the Command Descriptor Block (CDB) bytes or the DATA OUT phase bytes for which the drive issued an ILLEGAL REQUEST sense key. In the case of multiple byte fields, the field pointer indexes the most significant byte of the field.

C/D–COMMAND DATA (Byte 15, Bit 6)

A command data bit of one informs the initiator that the value reported for the field pointer is the number of the CDB byte for which the drive issued an ILLEGAL REQUEST sense key.

A command data bit of zero notifies the initiator that the field pointer reflects the number of a data byte received during the DATA OUT phase.

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

When BPV is one, the information in the bit pointer field is correct. When BPV is zero, the bit pointer is null and void.

BIT POINTER (Byte 15, Bits 2-0)

This parameter is a pointer to the erroneous bit of the byte specified by the field pointer.

Table 6-20 *Error Sense Keys and Additional Sense Codes*

SENSE KEY	ADDITIONAL SENSE CODE	DESCRIPTION
0h		NO SENSE. Indicates that there is no specific sense key information to be reported.
	00h	NO ADDITIONAL SENSE INFORMATION
1h		RECOVERED ERROR. Indicates that the command completed successfully with some recovery action performed by the drive. Details might be determinable by examining the additional sense bytes and the information field.
	2h	NO SEEK COMPLETE
	6h	NO REFERENCE POSITION FOUND
	10h	ID CRC OR ECC ERROR
	12h	ADDRESS MARK NOT FOUND FOR ID FIELD
	13h	ADDRESS MARK NOT FOUND FOR DATA FIELD
	14h	RECORDED ENTITY NOT FOUND
	15h	RANDOM POSITIONING ERROR
	17h	RECOVERED DATA WITH NO ERROR CORRECTION APPLIED
	18h	RECOVERED DATA WITH ERROR CORRECTION APPLIED
	86h	UNEXPECTED SEQUENCER ERROR
	96h	WRITE FAULT DUE TO BUMP
	97h	UNDERRUN ERROR
	A8h	UNSUCCESSFUL COMMUNICATION WITH DSP
	AAh	DSP HAS QUIT, OR IS UNABLE TO START UP
	ABh	REQUESTED FORMAT NOT AVAILABLE
2h		NOT READY. Indicates that the drive cannot be accessed. (See Section 6.8.1, "TEST UNIT READY (Opcode=00h).") Operator intervention might be required to correct this condition.
	04h	LOGICAL UNIT NOT READY, CAUSE NOT REPORTABLE
	B0h	DRIVE IS RECALIBRATING
	B1h	DRIVE IS SPINNING UP
	B2h	DRIVE HAS NOT BEEN TOLD TO SPIN UP
	B3h	DRIVE FAILED TO SPIN UP
	B6h	DRIVE HAS A FATAL SERVO ERROR

Table 6-20. *Error Sense Keys and Additional Sense Codes (continued)*

SENSE KEY	ADDITIONAL SENSE CODE	DESCRIPTION
3h		MEDIUM ERROR. Indicates that the command terminated with a non-recovered error condition that was probably caused by a flaw in the medium or an error in the recorded data.
	10h	ID CRC OR ECC ERROR
	11h	UNRECOVERED READ ERROR
	12h	ADDRESS MARK NOT FOUND FOR ID FIELD
	13h	ADDRESS MARK NOT FOUND FOR DATA FIELD
	14h	RECORDED ENTITY NOT FOUND
	19h	DEFECT LIST ERROR
	1Ch	DEFECT LIST NOT FOUND
	32h	NO DEFECT SPARE LOCATION AVAILABLE
	80h	ERROR IN WRITING A SYSTEM SECTOR
	81h	ERROR IN READING A SYSTEM SECTOR
	95h	SEQUENCER TIMEOUT
	97h	UNDERRUN ERROR
	98h	TIMEOUT IN SETTLING
	A3h	FAILURE IN READING SECTOR
	AAh	REALLOCATED UNCORRECTABLE DATA READ
4h		HARDWARE ERROR. Indicates that the drive detected a non-recoverable hardware failure while performing the command.
	2h	NO SEEK COMPLETE
	3h	PERIPHERAL DEVICE WRITE FAULT
	6h	NO REFERENCE POSITION FOUND
	15h	RANDOM POSITIONING ERROR
	32h	NO DEFECT SPARE LOCATION AVAILABLE
	40h	RAM FAILURE
	43h	MESSAGE ERROR
	84h	FAIL WRITING TO SEQUENCER FORMAT RAM
	85h	REJECT OF MESSAGE NOT SENT
	86h	UNEXPECTED SEQUENCER ERROR
	87h	LOGICAL ASSERTION ERROR, FIRMWARE INCONSISTENCY

Table 6-20. *Error Sense Keys and Additional Sense Codes (continued)*

SENSE KEY	ADDITIONAL SENSE CODE	DESCRIPTION
4h *(continued)*	8Ah	HEAD MISCOMPARE
	90h	SYNCHRONOUS ACKNOWLEDGE ERROR
	91h	FIFO UNLOAD ERROR
	92h	FIFO LOAD ERROR
	9Dh	MOTOR NEVER GETS UP TO SPEED
	A1h	SEQUENCER ROLLOVER REGISTER FAILURE
	A2h	EXTERNAL RAM FAILURE
	A8h	UNSUCCESSFUL COMMUNICATION WITH DSP
	AAh	DSP HAS QUIT, OR IS UNABLE TO START UP
	C0h	RECAL FAILURE DURING DSP DAC OFFSET TUNE
	C1h	RECAL FAILURE DURING INITIALIZATION
	C2h	RECAL FAILURE DURING NEC DAC OFFSET TUNE
	C3h	RECAL FAILURE DURING PES GAIN/BIAS CALIBRATION
	C4h	RECAL FAILURE DURING KT/J CALIBRATION
	C5h	RECAL FAILURE DURING ONCE AROUND CALIBRATION
	C6h	RECAL FAILURE DURING HEAD OFFSET CALIBRATION
	C7h	DSP FAILED TO ASSERT SERVO READY AT START UP
	C8h	CANNOT ACQUIRE OR SERVO ON ONE OR MORE HEADS
5h		ILLEGAL REQUEST. Indicates that there was an illegal parameter in the Command Descriptor Block (CDB) or in the additional parameters supplied as data. If the drive detects an invalid parameter in the CDB, it terminates the command without altering the data on the disk. If the drive detects an invalid parameter in the additional parameters supplied as data, it might have already altered the medium.
	1Ah	PARAMETER LIST LENGTH ERROR
	20h	INVALID COMMAND OPERATION CODE
	21h	LOGICAL BLOCK ADDRESS OUT OF RANGE
	24h	INVALID FIELD IN CDB
	25h	LOGICAL UNIT NOT SUPPORTED
	26h	INVALID FIELD IN PARAMETER LIST
	AEh	INVALID PARAMETER IN MODE PAGE

Table 6-20. *Error Sense Keys and Additional Sense Codes (continued)*

SENSE KEY	ADDITIONAL SENSE CODE	DESCRIPTION
6h		UNIT ATTENTION. Indicates that a condition has occurred that requires initiator awareness. (See Section 6.2.8.1, "ATTENTION Condition.")
	29h	POWER ON, RESET, OR BUS DEVICE RESET OCCURRED
	2Ah	PARAMETERS CHANGED
	2Fh	COMMANDS CLEARED BY ANOTHER INITIATOR
	3Fh	TARGET OPERATING CONDITIONS HAVE CHANGED
	8Eh	UNEXPECTED SIC INTERRUPT OCCURRED
Bh		ABORTED COMMAND. Indicates that the drive aborted the command. The initiator might be able to recover by trying the command again.
	40h	RAM FAILURE
	47h	SCSI PARITY ERROR
	48h	INITIATOR DETECTED ERROR MESSAGE RECEIVED
	4Eh	OVERLAPPED COMMANDS ATTEMPTED
Eh		MISCOMPARE. Indicates that the source data did not match the data read from the disk.
	1Dh	MISCOMPARE DURING VERIFY OPERATION

6.8.4 FORMAT UNIT (Opcode=04h)

The FORMAT UNIT command, shown in Table 6-21, assigns logical blocks to physical sectors for optimum sequential access, within the limitation of available spare sectors. When the drive encounters a defect, it performs in-line sparing. If the drive is unable to do so, the current address to be assigned is allocated to an alternate sector reserved as a spare. The Empire 540/1080S hard disk drive is formatted at the factory with data undefined.

Note: Issuing the FORMAT UNIT command typically causes loss of data. All data should be backed up prior to formatting.

6.8.4.1 In-Line Sparing

On each cylinder, there is a certain number of sectors reserved as spares. When the drive finds a bad sector on a cylinder, it checks to see if any of these spare sectors is unused. If so, the drive is able to shift its assignments of Logical Block Addresses (LBAs) over by one for all remaining sectors on that cylinder. This preserves the property of sequential logical blocks on the cylinder.

6.8.4.2 Command Structure

Table 6-21 *FORMAT UNIT Command*

BYTE	BIT 7	6	5	4	3	2	1	0
0	OPCODE=04h							
1	LUN=0			FMT	CMP	DEFECT LIST FMT=0		
2	DATA PATTERN							
3	(MSB)	INTERLEAVE=XX						
4		INTERLEAVE=XX						(LSB)
5	VU=0		RESERVED=0				F=0	L=0

FMT–FORMAT DATA (Byte 1, Bit 4)

A format data bit of one instructs the drive to go to the DATA OUT phase after receiving all of the command bytes. During the DATA OUT phase, the initiator transfers a defect list to the drive. The initiator provides these defects in terms of LBAs. The drive maps the LBAs specified into the corresponding physical addresses and stores these values in the defect map. Table 6-22 and Table 6-23 show the format of the defect list.

A format data bit of zero informs the drive that the DATA OUT phase will not occur–that is, the initiator supplies neither a defect list header nor defect data.

CMP–COMPLETE LIST (Byte 1, Bit 3)

A complete list bit of one notifies the drive that the defect list transferred during the DATA OUT phase is a comprehensive list of known initiator-specified or field-replacement defects. Initiator-specified or field-replacement defects are kept in a list called the G LIST, which stands for grown list. The drive erases the previous G LIST and builds a new G LIST, using the defect list.

A complete list bit of zero specifies that the defect list supplied by the initiator contains defective block data to supplement the existing G LIST.

Note: The drive also keeps a list of factory known defects, called the P LIST or primary list. Together, these two lists form a working list, which the drive can use to locate defective areas when formatting the disk.

DATA PATTERN (Byte 2)

This parameter identifies 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 on the Quantum-Unique Drive-Control Page (Page Code 39h).

INTERLEAVE (Bytes 3&4)

The drive ignores any parameter value placed in this field. The Empire 540/1080S hard disk drive performs formats using a 1:1 interleave.

6.8.4.3 Defect List Format

The FORMAT UNIT defect list contains a four-byte header, shown in Table 6-22, which might be followed by defect descriptors, shown in Table 6-23. Each defect descriptor carries the four-byte address of a logical block that contains a defect.

Table 6-22 *FORMAT UNIT Defect List – Header*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	RESERVED=0							
1	FOV	DPRY	DC=X	SF=X	IP=0	DS=0	IM=0	VS=0
2	(MSB)	DEFECT LIST LENGTH						
3		DEFECT LIST LENGTH						(LSB)

Note: The Empire 540/1080S supports FOV and DPRY. The remaining six bits of byte 1 are all unsupported. X denotes that any value can be given for this bit. The drive ignores this bit, and any parameter value placed in this bit has no bearing on the command.

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

A format options valid bit of one indicates that the initiator authorizes the setting of bit 6–DPRY. A format options valid bit of zero tells the drive to use its default format scheme for the function defined by bit 6. When FOV is zero, the initiator must set DPRY to zero.

DPRY–DISABLE PRIMARY (Byte 1, Bit 6)

When DPRY is one, the drive excludes the factory defect map–primary list or P LIST–from its list of defects to be managed on formatting. When DPRY is zero, the drive includes the P LIST.

By default, the Empire 540/1080S hard disk drive includes the factory defect map when formatting.

DC, SF, IP, DS, IM, AND VS

The Empire 540/1080 supports FOV and DPRY. The remaining six bits of byte 1 are all unsupported.

DEFECT LIST LENGTH (Bytes 2&3)

This parameter defines the defect descriptors' total length, in bytes. Consequently, the defect list length is equal to four times the number of defect descriptors. The value of this parameter can be zero.

Table 6-23 *FORMAT UNIT Defect List – Block Descriptor(s)*

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

DEFECT LBA (Bytes 0-3)

This parameter conveys to the drive the Logical Block Address (LBA) of a defect. Once the defect is identified to the drive in this manner, the drive is able to mark the physical sector as bad and use this LBA to access a different sector.

6.8.4.4 Application of the FORMAT UNIT Command

The Empire 540/1080S hard disk drive offers eight formatting options. Table 6-24 shows the settings of the FORMAT UNIT CDB and defect list header required to implement each of these formatting options.

Table 6-24 *FORMAT UNIT Command Variations*

FMT	CMP	FOV	DPRY	DEFECT LIST LENGTH	DESCRIPTION
1	1	1	1	0	Format with absolutely no defects.
1	1	0	0	0	Format with original factory defects only.
1	0	1	1	0	Format with grown, or field-found, defects only—disregarding factory defects
0	0	X	X	X	Format with existing defects—factory and grown defects.
1	1	1	1	> 0	Format with provided defects only—disregarding factory and existing grown defects.
1	1	0	0	> 0	Format with provided defects and factory defects—disregarding existing grown defects.
1	0	1	1	> 0	Format with provided defects and grown defects—disregarding factory defects.
1	0	0	0	> 0	Format with provided defects and existing defects.
Note: X denotes that these settings are irrelevant, since the DATA OUT phase will not occur (FMT=0).					

6.8.5 REASSIGN BLOCKS (Opcode=07h)

The REASSIGN BLOCKS command, shown in Table 6-25, directs the drive to reassign defective logical blocks to sectors reserved as spares. The drive records the physical locations of the defective logical blocks to the G LIST. The contents of the P LIST are not altered.

If specified, a logical block that has previously been reassigned can be reassigned again. The physical address of the defective logical block is added to the drive's list of field-found defects.

Note: Because the drive must access spare locations, block reassignment might degrade its performance. The FORMAT UNIT command performs in-line sparing, which increases performance. However, the FORMAT UNIT command erases the disk's contents.

6.8.5.1 Command Structure

Table 6-25 *REASSIGN BLOCKS Command*

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

6.8.5.2 Defect List Format

The initiator sends a list of the defective logical blocks, called the defect list, during the DATA OUT phase that immediately follows the COMMAND phase. If recoverable under the parameters located in the Read-Write Error Recovery Page (Page Code 01h), the data in the logical blocks will be preserved.

The REASSIGN BLOCKS defect list contains a four-byte header, shown in Table 6-26, followed by one or more defect descriptors, shown in Table 6-27. Each defect descriptor specifies the four-byte address of a logical block that contains a defect.

Table 6-26 *REASSIGN BLOCKS Defect List – Header*

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

DEFECT LIST LENGTH (Bytes 2&3)

This parameter defines the defect descriptors' total length, in bytes. Consequently, the defect list length is equal to four times the number of defect descriptors. The value of this parameter can be zero.

Table 6-27 *REASSIGN BLOCKS Defect List – Block Descriptor(s)*

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

DEFECT LBA (Bytes 0-3)

This parameter conveys to the drive the Logical Block Address (LBA) of a defect. Once the defect is identified to the drive in this manner, the drive is able to mark the physical sector as bad and use this LBA to access a spare sector.

If the drive has insufficient capacity to reassign all defective logical blocks, the command concludes with CHECK CONDITION status. The sense key is set to HARDWARE ERROR and the additional sense code is set to NO DEFECT SPARE LOCATION AVAILABLE.

Anytime the drive returns CHECK CONDITION status, the LBA of the first logical block not reassigned is placed in the command specific information bytes of the sense data.

6.8.6 READ (Opcode=08h)

The READ command, shown in Table 6-28, requests that the drive transfer data to the initiator from the disk.

Table 6-28 *READ Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=08h							
1	LUN=0			(MSB)	LBA			
2		LBA						
3		LBA						(LSB)
4	TRANSFER LENGTH							
5	VU=0		RESERVED=0				F=0	L=0

LBA (Byte 1, Bits 4-0; Bytes 2&3)

This parameter tells the drive the logical block to start reading from.

TRANSFER LENGTH (Byte 4)

The drive reads contiguous logical blocks, starting with the logical block whose LBA matches the LBA in the Command Descriptor Block (CDB), until the number read equals the transfer length. A transfer length of zero directs the drive to send 256 logical blocks to the initiator. Any other value represents the specific number of logical blocks to read and transfer.

6.8.7 WRITE (Opcode=0Ah)

The WRITE command, shown in Table 6-29, instructs the drive to write the data transferred by the initiator to the disk.

Table 6-29 *WRITE Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=0Ah							
1	LUN=0			(MSB)	LBA			
2		LBA						
3		LBA						(LSB)
4	TRANSFER LENGTH							
5	VU=0		RESERVED=0				F=0	L=0

LBA (Byte 1, Bits 4-0; Bytes 2&3)

This parameter identifies the logical block at which to begin writing.

TRANSFER LENGTH (Byte 4)

The drive writes the transferred data to the disk, starting at the logical block whose LBA matches the LBA in the Command Descriptor Block (CDB), until the number of contiguous logical blocks written equals the transfer length. A transfer length of zero points out that the initiator will be transferring 256 logical blocks of data to the drive. Any other value represents the specific number of logical blocks the initiator will transfer and the drive will write.

6.8.8 SEEK (Opcode=0Bh)

The SEEK command, shown in Table 6-30, communicates to the drive that it should seek to the Logical Block Address (LBA) specified.

Table 6-30 *SEEK Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=0Bh							
1	LUN=0			(MSB)	LBA			
2		LBA						
3		LBA						(LSB)
4	RESERVED=0							
5	VU=0		RESERVED=0				F=0	L=0

LBA (Byte 1, Bits 4-0; Bytes 2&3)

This parameter identifies the logical block, from which to calculate the cylinder and head, used by the drive to position the actuator.

6.8.9 INQUIRY (Opcode=12h)

The INQUIRY command, shown in Table 6-31, asks the drive to send its identification information to the initiator.

6.8.9.1 Command Structure

Table 6-31 *INQUIRY Command*

<table>
<tr><th rowspan="2">BYTE</th><th colspan="8">BIT</th></tr>
<tr><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="8">OPCODE=12h</td></tr>
<tr><td>1</td><td colspan="3">LUN=0</td><td colspan="4">RESERVED=0</td><td>EPD=0</td></tr>
<tr><td>2</td><td colspan="8">PAGE CODE=0</td></tr>
<tr><td>3</td><td colspan="8">RESERVED=0</td></tr>
<tr><td>4</td><td colspan="8">ALLOCATION LENGTH</td></tr>
<tr><td>5</td><td colspan="2">VU=0</td><td colspan="4">RESERVED=0</td><td>F=0</td><td>L=0</td></tr>
</table>

ALLOCATION LENGTH (Byte 4)

This parameter indicates the maximum number of bytes allocated by the initiator for returned INQUIRY data. Acceptable values are 00h-FFh. The drive completes the DATA IN phase when it has transferred to the initiator either the number of bytes specified by the allocation length or all available INQUIRY data, whichever is less. An allocation length of 132 (84h), or greater, is recommended for the Empire 540/1080S hard disk drive.

6.8.9.2 INQUIRY Data Format

The leading bytes of the INQUIRY data, shown in Table 6-32, give detailed information about the type of medium contained in the drive, the standards that the drive complies to, and specific capabilities of the drive.

The medium in the drive is not removable and the drive does not support asynchronous event notifications or the TERMINATE I/O PROCESS message.

Table 6-32 *INQUIRY Data Format*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	PER QUAL			PERIPHERAL DEVICE TYPE				
1	RM=0	DEVICE TYPE MODIFIER=0						
2	ISO VER=0		ECMA VER=0			ANSI VER=2		
3	AN=0	TI=0	RESERVED=0			RSPNSE DATA FMT=2		
4	ADDITIONAL LIST LENGTH							

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

Typically, the drive sets this field to 00h, notifying the initiator that it is a direct-access device. However, if the drive is not capable of supporting a device at the Logical Unit Number (LUN) established by the IDENTIFY message, this field is set to 1Fh–unknown device type.

PER QUAL–PERIPHERAL QUALIFIER (Byte 0, Bits 7-5)

Normally, the drive sets this field to 000b to indicate that a direct-access device is currently connected to the logical unit of the IDENTIFY message. If the drive is not capable of supporting a device on this logical unit, this field is set to 011b.

ANSI VER–ANSI-APPROVED VERSION (Byte 2, Bits 2-0)

The drive sets this field to 2h, notifying compliance with the ANSI SCSI-2 specification.

RSPNSE DATA FMT–RESPONSE DATA FORMAT (Byte 3, Bits 2-0)

The drive sets this field to 2h to point out that the data returned is in the format specified by the ANSI SCSI-2 specification.

ADDITIONAL LIST LENGTH (Byte 4)

This parameter is the length, in bytes, of the drive's identification information. Table 6-33 shows the identification information of the Empire 540/1080S hard disk drive.

If the allocation length specified prevents the drive from sending all available INQUIRY data to the initiator, this value is not adjusted to reflect the truncation.

Table 6-33 *Identification Information*

BYTE(S)	INFORMATION
5–6	Reserved=0
7	The drive sets this byte to 12h to tell the initiator that it supports synchronous data transfers and command queuing.
8–15	Vendor identification='QUANTUM '
16–22	Product identification (model)
23–31	Product identification (part number)=XXXXXXXXX
32–35	Microcode revision level='VVVV' VVVV Denotes the appropriate revision level
36–43	Microcode date='MM/DD/YY' MM Month DD Day YY Year
44–55	Drive serial number='PTCYDDDNNNNN' P Place of manufacture T Drive type (fixed at 4) C Drive capacity (1=540 MB, 2=1080 MB) Y Last digit of the year the drive was built in (1=1991, 2=1992, etc.) DDD Julian day of manufacture (139=139th day of the year) NNNNN Drive sequence on day of manufacture (2=2nd drive built)
56–95	Reserved=0
96–131	Vendor unique

Note: Bytes 8–55 are ASCII values. If the drive is not ready when it receives an INQUIRY command, it cannot return the correct data for bytes 44–55. In this case, the drive transmits the default data of 000000000000.

6.8.10 Mode Select (Opcode=15h)

The MODE SELECT command, shown in Table 6-34, provides a means for the initiator to specify medium, logical unit, or peripheral device parameters to the target.
Targets that implement the MODE SELECT command also implement the MODE SENSE command. Initiators should issue MODE SENSE prior to each MODE SELECT to determine supported pages, page lengths, and other parameters.

6.8.10.1 Command Structure

Table 6-34 *MODE SELECT Command*

Bit / Byte	7	6	5	4	3	2	1	0
0	OPCODE=15h							
1	LUN=0			PF=0	RESERVED=0			SP
2	RESERVED=0							
3	RESERVED=0							
4	MODE PARAMETER LIST LENGTH							
5	VU=0		RESERVED=0				F=0	L=0

When the drive motor spins up at power on, the drive reads the operating mode parameters most recently saved, from a reserved cylinder, and sets a UNIT ATTENTION condition for all initiators. The additional sense code accompanying the UNIT ATTENTION sense key is determined as follows:

- If the drive cannot successfully read the parameter values from the Saved Mode Table, the drive accesses the Default Mode Table to set the state of the operating mode parameters. The additional sense code is set to PARAMETERS CHANGED, as though another initiator had altered the parameters.
- If the drive is able to read the parameter values, the additional sense code is set to POWER ON, RESET, OR BUS DEVICE RESET OCCURRED.

Note: If the disable unit attention bit on the Quantum-Unique Drive-Control Page (Page Code 39h) is set to one, a UNIT ATTENTION condition does not occur.

SP–SAVE PAGES (Byte 1, Bit 0)

When SP is one, the drive copies the operating mode parameters in the Current Mode Table to the Saved Mode Table, after updating the Current Mode Table if pages were sent. When SP is zero, the drive performs the specified MODE SELECT operation and does not alter the contents of the Saved Mode Table.

PF–PAGE FORMAT (Byte 1, Bit 4)

When PF is zero, the MODE SELECT parameters are as specified in SCSI-1, (i.e., all parameters after the block descriptors are vendor-specific). A PF bit of one indicates that the MODE SELECT parameters following the header and block descriptor(s) are structured as pages of related parameters and are as specified in SCSI-2.

MODE PARAMETER LIST LENGTH (Byte 4)

The mode parameter list length field specifies the length in bytes of the mode parameter list that is transferred from the initiator to the target during the DATA OUT phase. A parameter list length of zero indicates that no data is transferred. This condition is not considered as an error.

The drive uses this parameter's value to determine the length of the parameter list, and then transfers that many bytes.

6.8.10.2 Mode Parameter List Format

The MODE PARAMETER list, shown in Table 6-35, contains a four-byte header, followed by zero or more block descriptors, followed by zero or more variable-length pages. The pages sent correspond to unique divisions of the Current Mode Table. The drive determines which pages are included in the mode parameter list, accesses the Current Mode Table, and places the parameter values from these pages into the correct divisions. Pages are discussed separately in Section 6.9, Mode Pages, because they apply to both the MODE SELECT and MODE SENSE commands. Before using this command, or the MODE SENSE command, you should be familiar with the information provided regarding pages.

Table 6-35 *Mode Parameter List*

Bit Byte	7	6	5	4	3	2	1	0
0–3	MODE PARAMETER HEADER							
0–n	BLOCK DESCRIPTOR(S)							
0–n	PAGE(S)							

Table 6-36 shows the four-byte mode parameter header

Table 6-36 *Mode Parameter Header*

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED=0							
1	MEDIUM TYPE=0							
2	RESERVED=0							
3	BLOCK DESCRIPTOR LENGTH							

MEDIUM TYPE (Byte 1)

This byte has a value of zero to indicate a hard disk.

BLOCK DESCRIPTOR LENGTH (Byte 3)

This parameter informs the drive of the block descriptor's length, in bytes. It is equal to the number of block descriptors times eight and does not include pages or vendor-specific parameters, if any, that may follow the last block descriptor. A block descriptor length of zero indicates that there are no block descriptors in the parameter list. The Empire 540/1080S hard disk drive uses a single eight-byte block descriptor. Therefore, this parameter must be set to either zero or one by the initiator. The format of the mode parameter block descriptor is shown in Table 6-37.

Table 6-37 *Mode Parameter Block Descriptor*

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED=0							
1	(MSB)	NUMBER OF BLOCKS						
2		NUMBER OF BLOCKS						
3		NUMBER OF BLOCKS						(LSB)
4	RESERVED=0							
5	(MSB)	BLOCK LENGTH						
6		BLOCK LENGTH						
7		BLOCK LENGTH						(LSB)

NUMBER OF BLOCKS (Bytes 1-3)

This parameter conveys to the drive the number of logical blocks on the media that match the block length specified in this block descriptor. A value of zero notes that all remaining logical blocks on the drive have the media characteristics specified by the block descriptor. Any non zero value is ignored.

BLOCK LENGTH (Bytes 5-7)

The initiator must set this field to 200h, because the drive is only capable of a block length of 512 bytes.

The mode page format is shown in Table 6-38. Each mode page contains a page code, a page length, and a set of mode parameters. The mode parameters take up a variable number of bytes depending on the specific mode page. Pages are optional. They can be included in any order, immediately following the block descriptor. To avoid the specification of all mode parameters 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.

Table 6-38 *Mode Page Format*

Bit / Byte	7	6	5	4	3	2	1	0
0	RESERVED=0		PAGE CODE					
1	PAGE LENGTH							
2 and up	SEE INDIVIDUAL SECTION ON PAGES							

PAGE CODE (Byte 0, Bits 5–0)

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 6-39

Table 6-39 *Mode Pages Supported by the Empire 540/1080S*

PAGE	DESCRIPTION
01H	Read/Write Error-Recovery Parameters
02H	Disconnect/Reconnect Control Parameters
03H	Direct-Access Device Format Parameters*
04H	Rigid Disk-Drive Geometry Parameters*
08H	Cache-Control Parameters
0AH	Control Mode
0CH	Notch and Partition Parameters
32H	Quantum-Unique Automatic Shutdown Control Parameters
37H	Quantum-Unique Control Parameters
38H	Quantum--Unique Cache Control Parameters
39H	Quantum-Unique Drive-Control Parameters

* Read only. Can be accessed only by means of 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.

6.8.10.3 Error-Recovery Parameters, Page Code 01H

Table 6-40 *Error-Recovery Parameters*

BYTE	BIT 7	BIT6	BIT5	BIT4	BIT3	BIT2	BIT1	BIT0
0	RESERVED = 0		PAGE CODE = 01_H					
1	PAGE LENGTH = 6_H							
2	AWRE	ARRE	TB	RC	EEC	PER	DTE	DCR
3	RETRY COUNT							
4	CORRECTION SPAN							
5	HEAD OFFSET COUNT = 0							
6	DATA STROBE OFFSET COUNT = 0							
7	RECOVERY TIME LIMIT = 0							

DCR—DISABLE CORRECTION (Byte 2, Bit 0)

When set to one, DCR indicates that the data will be transferred without correction, whether or not correction is possible. When set to zero, this bit indicates that the data will be corrected, if possible. The default is zero.

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 (bit 2 of byte 2) 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 TB is set to one and the DTE bit (bit 1 of byte 2) is also set to one, recovered or unrecoverable data in the block in error will be transferred to the initiator. When set to zero and the DTE bit is set to one, the block in error will not be transferred to the initiator. When DTE is set to zero or RC is set to one, the drive ignores this bit. On unrecoverable errors, the drive checks TB, regardless of the state of DTE. 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 one.

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 because the drive is unable to read the sector ID of the sector targeted to be written. When set to zero, the drive will not automatically reallocate bad blocks. The default is one.

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. Values range from eight to sixteen. The default is sixteen.

Note: The head-offset count, data-strobe offset, and recovery-time limit (bytes 5–7) are not supported by the Empire 540/1080S hard disk drive.

Table 6-41 summarizes the valid modes of operation for Empire 540/1080S hard disk drive.

Table 6-41 *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
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

6.8.10.4 Disconnect/Reconnect Control Parameters, Page Code 02H

Table 6-42 *Disconnect/Reconnect Control Parameters*

BYTE	BIT 7	BIT6	BIT5	BIT4	BIT3	BIT2	BIT1	BIT0
0	RESERVED = 0		PAGE CODE = 02_H					
1	PAGE LENGTH = A_H							
2	BUFFER FULL RATIO							
3	BUFFER EMPTY RATIO							
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. A buffer full ratio of 255 (FFH) represents 100% full, 128 (80H) represents 50% full, and so on. Each bit represents 1/256th of the maximum buffer size. The default value is 217 (D9H). 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. A buffer empty ratio of 255 (FFH) represents 100% empty, 128 (80H) represents 50% empty, and so on. Each bit represents 1/256th of the maximum buffer size. The default value is 217 (D9H). The buffer must be completely empty before reconnection.

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.

6.8.10.5 Direct-Access Device Format, Page Code 03H

As explained in Table 6-39, the Direct-Access Device Format page is read only, and can be accessed by the MODE SENSE command. Please refer to the MODE SENSE command section, later in this chapter, for details concerning the information in this page.

6.8.10.6 Rigid Disk Drive Geometry Parameters, Page Code 04H

As explained in Table 6-39, the Rigid Disk Drive Geometry Parameters page is read only, and can be accessed by the MODE SENSE command. Please refer to the MODE SENSE command section, later in this chapter, for details concerning the information in this page.

6.8.10.7 Cache Control, Page Code 08H

Table 6-43 *Cache Control*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED = 0		PAGE CODE = 08_H					
1	PAGE LENGTH = $0A_H$							
2	RESERVED = 0					WCE	MS=0	RCD=0
3	RESERVED = 0							
4-5	RESERVED = 0							
6-11	RESERVED = 0							

The Cache Control Parameters page specifies the parameters that control the operation of the cache. This page and page 37H can be used interchangeably 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)
When set to one, the drive will activate caching on all writes. When WCE is set to zero, the drive will disable caching on writes. The default is one.

MS—MULTIPLE SELECTION (Byte 2, Bit 1)
Not supported.

RCD—READ CACHE DISABLE (Byte 2, Bit 0)
Set to zero by default, indicating that the drive can return some or all of the data requested by a READ command by accessing the cache, rather than the disk. When the RCD bit is set to one, the drive must read all requested data from disk and cannot return any data by accessing the cache. Setting the RCD bit to one causes the drive to automatically set the PE and CE bits to zero in page 37H. Setting the CE bit in page 37H causes the drive to turn off the RCD bit in this page.

6.8.10.8 Control Mode, Page Code 0AH

The control mode page, shown in Table 6-44, provides controls over several SCSI-2 features which are applicable to all device types such as tagged queuing, extended contingent allegiance, asynchronous event notification, and error logging.

Table 6-44 *Control Mode Page*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	PS	RESERVED	PAGE CODE = $0A_H$					
1	PAGE LENGTH = 06_H							
2	RESERVED=0							RLEC=0
3	QUEUE ALGORITHM MODIFIER				RESERVED=0		QERR	DQUE
4	EECA	RESERVED				RAENP	UAAENP	EAENP
5	RESERVED = 0							
6	READY AEN HOLDOFF PERIOD (MSB)							
7	READY AEN HOLDOFF PERIOD (LSB)							

RLEC (Byte 2, Bit 0)

A report log exception condition (RLEC) bit of one specifies that the drive reports log exception conditions. A RLEC bit of zero specifies that the target does not report log exception conditions. The default is zero, because RLEC is not supported.

QUEUE ALGORITHM MODIFIER (Byte 3, Bits 4–7)

The queue algorithm modifier field specifies restrictions on the algorithm used for reordering commands that are tagged with the SIMPLE QUEUE TAG message.

A value of zero in this field specifies that the target orders the actual execution sequence of the commands with a SIMPLE QUEUE tag such that data integrity is maintained for that initiator. This means that, if the transmission of new commands was halted at any time, the final value of all data observable on the medium will have exactly the same value as it would have if the commands had been executed in the same received sequence without tagged queuing. The restricted reordering value of zero is the default value, and is not changeable.

QERR (Byte 3, Bit 1)

A queue error management (QERR) bit of zero specifies that remaining suspended I/O process resume after the contingent allegiance condition or extended contingent allegiance condition.

A QERR bit of one specifies all remaining suspended I/O processes are to be aborted after the contingent allegiance condition or extended contingent allegiance condition. A UNIT ATTENTION condition is generated for each initiator that had a suspended I/O process aborted except for the initiator that had the contingent allegiance condition or extended contingent allegiance condition. The target sets the additional sense code to TAGGED COMMANDS CLEARED BY ANOTHER INITIATOR.

DQUE (Byte 3, Bit 0)

A disable queuing (DQUE) bit of zero specifies that tagged queuing is to be enabled if the target supports tagged queuing. A DQUE bit of one specifies that tagged queuing be disabled. Any queued commands for that I_T_xnexus is aborted. Any subsequent queue tag message received is to be rejected with a MESSAGE REJECT message and the I/O process is to be executed as an untagged command.

EECA (Byte 4, Bit 7)

An enable extended contingent allegiance (EECA) bit of one specifies that extended contingent allegiance is enabled. An EECA bit of zero specifies that extended contingent allegiance is disabled. The default value for this bit is zero, and is not changeable.

RAENP, UAAENP, EAENP (Byte 4, Bits 2,1, and 0)

The RAENP, UAAENP, and EAENP bits enable specific events to be reported via the asynchronous event notification protocol. When all three bits are zero, the target does not create asynchronous event notifications. The default value for these bits is zero, and is not changeable.

RAENP (Byte 4, Bit 2)

A ready AEN permission (RAENP) bit of one specifies that the target may issue an asynchronous event notification upon completing its initialization sequence instead of generating a UNIT ATTENTION condition. A RAENP bit of zero specifies that the target not issue an asynchronous event notification upon completing its initialization sequence. The default value for this bit is zero, and is not changeable.

UAAENP (Byte 4, Bit 1)

A unit attention AEN permission (UAAENP) bit of one specifies that the target may issue an asynchronous event notification instead of creating a UNIT ATTENTION condition upon detecting an event which would cause a UNIT ATTENTION condition (other than upon completing an initialization sequence). A UAAENP bit of zero specifies that the target not issue an asynchronous event notification instead of creating a UNIT ATTENTION condition. The default value for this bit is zero, and is not changeable.

EAENP (Byte 4, Bit 0)

An error AEN permission (EAENP) bit of one specifies that the target may issue an asynchronous event notification upon detecting a deferred error condition instead of waiting to report the deferred error on the next command. An EAENP bit of zero specifies that the target not report deferred error conditions via an asynchronous event notification. The default value for this bit is zero, and is not changeable.

READY AEN HOLDOFF PERIOD (Bytes 6 and 7)

The ready AEN holdoff period field specifies the minimum time in milliseconds after the target starts its initialization sequence that it delays before attempting to issue an asynchronous event notification. This value may be rounded up. The default value for these bytes is zero, and is not changeable.

6.8.10.9 Notch and Partition, Page Code 0CH

The Notch and Partition Page, shown in Table 6-45, 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 6-45 *Notch and Partition*

BYTE	BIT7	BIT6	BIT5	BIT4	BIT3	BIT2	BIT1	BIT0
0	PS=0	RESERVED=0	PAGE CODE = $0C_H$					
1	PAGE LENGTH = 16_H							
2	ND=1	PLN	RESERVED = 0					
3	RESERVED = 0							
4-5	MAXIMUM NUMBER OF NOTCHES							
6-7	ACTIVE NOTCH							
8-11	STARTING BOUNDARY							
12-15	ENDING BOUNDARY							
16-23	PAGES NOTCHED							

PS—PARAMETERS SAVEABLE (Byte 0, Bit 7)

The MODE SENSE command reserves this bit for its use. In the Empire 540/1080S 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 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 Empire 540/1080S hard disk drive use 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 drive supports a maximum of eight 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 Empire 540/1080S hard disk drive, valid notch numbers range from zero to seven. 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. Only the MODE SENSE command can set this parameter. When active notch is set to zero, this parameter defines the starting address of the logical unit. For drives using physical notch boundaries (PLN=0), bytes 8–10 define the cylinder and byte 11 defines the head. For drives using logical notch boundaries (PLN=1), bytes 8–11 define the logical block address. The drive uses physical notch boundaries. 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 drives using physical notch boundaries (PLN=0), bytes 12–14 define the cylinder and byte 15 defines the head. For drives using logical notch boundaries (PLN=1), bytes 12–15 define the logical block address. The drive uses physical notch boundaries. 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 3FH; the least significant bit, to page 00H. 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 Empire 540/1080S hard disk drive, the bits corresponding to pages 03H and 0CH are set to one, and are notch dependent. All other bits are set to zero. This is a read-only parameter.

6.8.10.10 Quantum Unique Automatic Shutdown Control, Page Code 32H

Table 6-46 *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 = 2_H							
2	AUTO STANDBY TIME (Seconds)							
3	AUTO SHUTDOWN TIME (Minutes)							

The Automatic Shutdown option is intended to save power by parking the heads and spinning down the spindle if the drive remains unselected for the Auto Shutdown time. After the Auto Shutdown option has been triggered, the drive should still respond to the next selection, spin up to motor and complete the requested command as soon as possible. The fact that the drive was shut down should be transparent to the initiator (except for added delay required for spin up). The Empire 540/1080S hard disk drive does not support automatic shutdown—the Automatic Shutdown Control page is supported for compatibility reasons only. This means that the vendor equipment may read the page, but all parameters are set to zero.

AUTO STANDBY TIME (Byte 3

The maximum time, in seconds, that the drive can remain selected before it will enter the Standby mode. The heads are parked (in the landing zone) 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 unselected before entering power-shutdown mode. On entering this mode, the heads are automatically parked in the landing zone and the spindle motor is shut off. When this byte is set to zero (the default), auto shutdown is disabled.

6.8.10.11 Quantum-Unique Control Parameters, Page Code 37H

Table 6-47 *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 CODE = 37_H					
1	PAGE LENGTH = $0E_H$							
2	RESERVED = 0		PSM	SSM	RESERVED = 0		PE	CE
3	NUMBER OF CACHE SEGMENTS							
4-15	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.

The following parameters control the operation of Quantum's DisCache.

Mode Select page 8H can also be used to control the operation of DisCache—see Section 6.8.10.7, "Cache Control, Page Code 08H". When these parameters are set in either page, the drive automatically sets the corresponding parameters in the other page.

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

8H is set to zero. The default setting for PE is one.

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 RAM only as a transfer buffer. The default is one.

The drive automatically sets this bit when the RCD bit in page 8H is set to zero. When the CE bit is set, the drive automatically turns off the RCD bit in page 8H.

NUMBER OF CACHE SEGMENTS (Byte 3)

This parameter indicates the number of segments the drive can index in the cache—that is, the number of entries in the cache table. Valid entries are 1 and 2. Any other entry will result in a CHECK CONDITION status, with the sense key ILLEGAL REQUEST. The default setting for this byte is one.

6.8.10.12 Quantum Unique Cache Control Parameters, Page Code 38H

Table 6-48 *Power Control Parameters*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
0	RESERVED = 0		PAGE CODE = 38_H					
1	PAGE LENGTH = $0E_H$							
2	CACHE CONTROL							
3	PREFETCH THRESHOLD							
4	MAXIMUM PREFETCH ABSOLUTE=0							
5	MAXIMUM PREFETCH MULTIPLE=3							
6	MINIMUM PREFETCH ABSOLUTE=0							
7	MINIMUM PREFETCH MULTIPLE=0							
8-10	NUMBER OF CACHE HITS=0							
11-13	NUMBER OF CACHE MISSES=0							
14-15	RESERVED=0							

CACHE CONTROL (byte 2)
This byte is set to 5CH.

PREFETCH THRESHOLD (byte 3)
This byte is set to 10H.

MAXIMUM PREFETCH ABSOLUTE (byte 4)
This byte is set to 0.

MAXIMUM PREFETCH MULTIPLE (byte 5)
This byte is set to 3.

MINIMUM PREFETCH ABSOLUTE (byte 6)
This byte is set to 0.

MINIMUM PREFETCH MULTIPLE (byte 7)
This byte is set to 0.

NUMBER OF CACHE HITS (bytes 8–10)
These bytes are set to 0.

NUMBER OF CACHE MISSES (bytes11–13)
These bytes are set to 0.

RESERVED (bytes 14–15)
These bytes are set to 0.

The Empire 540/1080S does not control nor respond to this mode page. It has been retained to maintain compatibility with earlier drives.

6.8.10.13 Quantum-Unique Drive Control Parameters, Page Code 39H

Table 6-49 *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 = 6_H							
2	DIO	DII	~~FDB~~	RUEE	FDPE	R=0	DUA	DRT
3	DDIS	DELDIS	RESERVED	DPC	SSID	SCSIADR		
4-7	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 neither used nor required by applications using the drive. It is included to maintain software compatibility with Quantum's Q200 Series™ hard disk drives.

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 3H and the sense code to AAH. 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.

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 one.

R:—RESERVED (Bit 2 of byte 2)
This bit 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 neither used nor required by applications using the drive. It is included to maintain software compatibility with Quantum's Q200 Series hard disk drives. The default is zero.

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

When set to zero, DCP 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 set to zero, SSID indicates the drive SCSI ID will be determined by the value set using the hardware ID select (see Chapter 3). When SSID is set to one, the drive SCSI ID is determined by the value set in SCSIADR (Byte 3, bits 2 - 0). The default is zero.

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

When SSID (bit 3 of byte 3) is set to one, SCSIADR determines the drive's SCSI ID. When changing the SCSI ID with the MODE SELECT command, the SCSI ID will change immediately. 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. See the description of SSID, above. The default is zero.

6.8.11 RESERVE (Opcode=16h)

The RESERVE and RELEASE commands form the basic mechanism for contention resolution in multiple-initiator systems. The RESERVE command, shown in Table 6-50, captures the drive for exclusive use by either the initiator issuing the command or another specified SCSI device.

Table 6-50 *RESERVE Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=16h							
1	LUN=0			3RD	3RD PARTY DEVICE ID			XT=0
2	RESERVATION IDENTIFICATION=XX							
3	(MSB)	EXTENT LIST LENGTH=XX						
4		EXTENT LIST LENGTH=XX						(LSB)
5	VU=0		RESERVED=0				F=0	L=0

Note: XX denotes that any value can be given for this field. The drive ignores this field, and any parameter value placed in this field has no bearing on the command.

3RD–THIRD PARTY (Byte 1, Bit 4)

When 3RD is one, the drive performs a Third Party Reservation. When 3RD is zero, the drive executes a Logical Unit Reservation.

6.8.11.1 Logical Unit Reservation

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

This field is ignored, and can therefore be set to any value by the initiator.

This type of reservation instructs the drive to reserve itself for the exclusive use of the initiator.

6.8.11.2 Third Party Reservation

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

This parameter is the SCSI ID of the device for which the drive reserves itself.

This type of reservation allows an initiator to reserve the drive for another SCSI device. This is intended for use in multiple-initiator systems that use the COPY command.

6.8.11.3 All Reservations

In both cases, the reservation remains valid until:

- The reservation is superseded by another valid RESERVE command or released by a RELEASE command. The RESERVE or RELEASE command must have been sent from the initiator that made the reservation.
- A BUS DEVICE RESET message is received from any initiator.
- A RESET condition is detected.
- A power cycle occurs.

If the drive is reserved for an initiator, the drive passes RESERVATION CONFLICT status to any other initiator that attempts to perform a command on it other than an INQUIRY, REQUEST SENSE, or RELEASE command. The drive ignores any attempt to release a reservation made by another initiator.

An initiator that holds a current reservation can modify that reservation by issuing another RESERVE command to the drive. The superseding RESERVE command releases the previous reservation state when the new reservation request is granted. The current reservation is not modified if the superseding reservation cannot be granted.

6.8.12 RELEASE (Opcode=17h)

The RELEASE command, shown in Table 6-51, frees the drive from a previous reservation. A reservation can be released only by the initiator that made it. It will not cause an error if an initiator attempts to release a reservation that is not currently valid, or is held by another initiator.

Table 6-51 *RELEASE Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=17h							
1	LUN=0			3RD	3RD PARTY DEVICE ID			XT=0
2	RESERVATION IDENTIFICATION=XX							
3	RESERVED=0							
4	RESERVED=0							
5	VU=0		RESERVED=0				F=0	L=0

Note: XX denotes that any value can be given for this field. The drive ignores this field, and any parameter value placed in this field has no bearing on the command.

3RD–THIRD PARTY (Byte 1, Bit 4)

When 3RD is one, the drive performs a Third Party Release. When 3RD is zero, the drive executes a Logical Unit Release.

6.8.12.1 Logical Unit Release

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

This field is ignored, and can therefore be set to any value by the initiator.

This type of release causes the drive to terminate a Logical Unit Reservation that is active from the initiator.

6.8.12.2 Third Party Release

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

This parameter is the SCSI ID of the device from which the drive releases itself.

This type of release allows an initiator to release the drive from a previously made Third Party Reservation. This is intended for use in multiple-initiator systems that use the COPY command. The drive is only released if the reservation was made using a Third Party Reservation, by the initiator that is requesting the release.

The SCSI device, from which the drive is discharged, must also be exactly the same as the one specified during the RESERVE command.

6.8.13 Mode Sense (Opcode=1Ah)

The MODE SENSE command, shown in Table 6-52, tells the drive to report its operating mode parameters to the initiator. This command complements the MODE SELECT command. Before preceding further, it is advisable to read Section 6.8.10.2, "Mode Parameter List Format".

6.8.13.1 Command Structure

Table 6-52 *MODE SENSE Command*

Bit Byte	7	6	5	4	3	2	1	0
0	OPCODE=1Ah							
1	LUN=0			R=0	DB=0	RESERVED=0		
2	PC		PAGE CODE					
3	RESERVED=0							
4	ALLOCATION LENGTH							
5	VU=0		RESERVED=0				F=0	L=0

PC–PAGE CONTROL (Byte 2, Bits 7 and 6)

This parameter defines the type of page parameter values the drive reports. There are four possible types, or formats, for the fields and bits contained in each page:

1. Current Values (00b)

   The drive sends the page, or pages, defined by the page code, with fields and bits set to the current operating mode values. The current values are:

   - Those set by the last MODE SELECT command successfully completed.
   - Identical to the saved values–if there has been no MODE SELECT command since the last power on or RESET condition.

2. Changeable Values (01b)

   The drive sends the page, or pages, defined by the page code, with fields and bits that can be modified set to one. Fields and bits that cannot be altered by the initiator are set to zero.

3. Default Values (10b)

   The drive sends the page, or pages, defined by the page code, with fields and bits set to the default operating mode values–that is, as shipped from the factory. Fields and bits not supported by the drive are set to zero.

   Note: To determine whether a returned value of zero indicates a default parameter or an unsupported parameter, the initiator can examine the changeable values.

4. Saved Values (11b)

   The drive sends the page, or pages, defined by the page code, with fields and bits set to the saved operating mode values. The saved values are:

   - Those set by the last MODE SELECT command which had the save pages bit set to one.

- Identical to the default values–if there has been no MODE SELECT command that had the save pages bit set to one.

Note: For all page control parameters, the value returned in the page length field indicates the number of bytes the drive supports within each page. This value must be the value entered for the page length when issuing a MODE SELECT command.

PAGE CODE (Byte 2, Bits 5-0)

This parameter notifies the drive of the specific page, or pages, the initiator wants. Table 6-53 shows the page codes for the pages the drive supports.

Table 6-53 *Page Codes as a Function of Pages Supported*

Page Code	Page(s) Returned
01h	Read-Write Error Recovery Page
02h	Disconnect-Reconnect Page
03h	Direct-Access Device Format Parameters Page
04h	Rigid Disk Drive Geometry Page
08h	Cache Control Parameters Page
0Ah	Control Mode Page
0Ch	Notch and Partition Page
32h	Quantum-Unique Automatic-Shutdown Control Parameters Page
37h	Quantum-Unique Control Parameters Page
38h	Quantum- Unique Cache Control Page
39h	Quantum-Unique Drive-Control Parameters Page
3Fh	All Pages (Order=01h, 02h, ... , 39h, 00h)

ALLOCATION LENGTH (Byte 4)

This parameter specifies the maximum number of bytes allocated by the initiator for returned MODE SENSE data. Acceptable values are 00h-FFh. The drive terminates the DATA IN phase when it has transferred to the initiator either the number of bytes specified by the allocation length or all available MODE SENSE data, whichever is less.

6.8.13.2 Mode Sense Data Format

The MODE SENSE data returned by the drive contains a four-byte header, shown in Table 6-54, followed by one or more pages. Pages are discussed in Section 6.8.10 as well as in this section because they apply to both the MODE SELECT and MODE SENSE commands. Before using the MODE SENSE or the MODE SELECT command, be familiar with the information provided regarding mode pages.

Table 6-54 *Mode Sense Header*

Bit / Byte	7	6	5	4	3	2	1	0
0	MODE DATA LENGTH							
1	MEDIUM TYPE=0							
2	WP=0	RSVD=0		DF=0	RESERVED=0			
3	BLOCK DESCRIPTOR LENGTH=08h							

Note: The medium is not write protected.

MODE DATA LENGTH (Byte 0)

The drive transfers the MODE SENSE data during the DATA IN phase of the command. The first byte it sends is the mode data length. This parameter informs the initiator of the number of available MODE SENSE data bytes. The mode data length does not include itself.

BLOCK DESCRIPTOR LENGTH (Byte 3)

The Empire 540/1080S hard disk drive uses a single eight-byte block descriptor. This parameter gives the initiator the block descriptor's length, in bytes. Therefore, the drive sets this field to 08h.

The block descriptor is shown in Table 6-55.

Table 6-55 *Mode Sense Block Descriptor*

Bit / Byte	7	6	5	4	3	2	1	0
0	RESERVED=0							
1	(MSB)	NUMBER OF BLOCKS						
2		NUMBER OF BLOCKS						
3		NUMBER OF BLOCKS						(LSB)
4	RESERVED=0							
5	(MSB)	BLOCK LENGTH						
6		BLOCK LENGTH						
7		BLOCK LENGTH						(LSB)

NUMBER OF BLOCKS (Bytes 1-3)
This parameter conveys to the initiator the number of logical blocks on the media that match the block length specified in this block descriptor. A value of zero notes that all remaining logical blocks on the drive have the media characteristics specified by the block descriptor.

BLOCK LENGTH (Bytes 5-7)
The drive sets this field to 200h to indicate that it supports 512 bytes per logical block.

The MODE SENSE page descriptor is shown in Table 6-56.

Table 6-56 *Mode Sense Page Descriptor*

Bit / Byte	7	6	5	4	3	2	1	0
0	PS	R = 0	PAGE CODE					
1	PAGE LENGTH							
2 and up	REFER TO PAGE DESCRIPTIONS							

6.8.13.3 MODE SENSE Data

Table 6-57 shows the header, block descriptor, and page descriptors all in one table, for clarity. Following the table, you will find explanations of the various fields and bits.

Table 6-57 *Mode Sense Data*

BYTE	BIT 7	BIT 6	BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0
HEADER								
0	SENSE DATA LENGTH							
1	MEDIUM TYPE = 0							
2	WP=0	RESERVED = 0						
3	BLOCK DESCRIPTOR LENGTH = 8							
BLOCK DESCRIPTOR								
0	DENSITY CODE = 00							
1-3	NUMBER OF BLOCKS = 0							
4	RESERVED = 0							
5-7	BLOCK DESCRIPTOR LENGTH							
PAGE DESCRIPTORS								
0	PS	R=0	PAGE CODE					
1	PAGE LENGTH							
2 TO N	REFER TO PAGE DEFINITION							

HEADER

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 itself.

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 the block descriptor, in bytes, and is set to eight for the Empire 540/1080S hard disk drives.

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 DESCRIPTORS

PS—PARAMETERS SAVEABLE (Byte 0, Bit 7)

When PS is set to zero in a 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.

6.8.13.4 Direct-Access Device Format Parameters, Page Code 3H

Table 6-58 *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	TRACKS 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=1	HSEC=0	RMB=0	SURF=0	INS = 0	RESERVED = 0		
21-23	RESERVED = 0							

DEFECT-HANDLING FIELDS

TRACKS PER ZONE (Byte 2–3)

The value in this field indicates the number of tracks per defect zone. For the Empire 540/1080S, a defect zone is defined as one cylinder. Empire 540/1080S hard disk drives use six 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 defect zone from the initiator-addressable blocks on execution of the FORMAT UNIT command. These sectors are available as replacement sectors for defective sectors. The default value is one.

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. This field is set to 92, which is the number of sectors per track on the outside zone. The innermost zone has 42 sectors per track.

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, so this field is set to 512.

INTERLEAVE (Bytes 14–15)
The drive has an interleave factor of one, so this field is set to 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. This field is set to 19.

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. This field is set to 25.

SSEC—SOFT SECTOR (Byte 20, Bit 7)
SSEC is always set to one, indicating that Empire 540/1080S hard disk drive uses soft sector formatting.

HSEC—HARD SECTOR (Byte 20, Bit 6)
HSEC is always set to zero, indicating that the Empire 540/1080S hard disk drive does not use hard sector formatting. The HSEC and SSEC bits are mutually exclusive.

RMB—REMOVABLE (Byte 20, Bit 5)
Set to zero for the 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.

INHIBIT SAVE (INS) (Byte 20, Bit 3)
Set to zero, indicating that the drive will save the page parameters defined by page code 4H as a read-only page. The Mode Page is available through Mode Sense only.

6.8.13.5 Rigid Disk Drive Geometry Parameters, Page Code 4H

Table 6-59 *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 = 4_H					
1	PAGE LENGTH = 12_H							
2-4	NUMBER OF CYLINDERS							
5	NUMBER OF HEADS							
6-8	STARTING CYLINDER - WRITE PRECOMPENSATION							
9-11	STARTING CYLINDER - REDUCED WRITE CURRENT							
12-13	DRIVE STEP RATE							
14-16	LANDING ZONE CYLINDER = 0							
17-19	RESERVED = 0							

NUMBER OF CYLINDERS. (Bytes 2–4)
This field is set to B3Ah., indicating that the number of cylinders is 2,874.

NUMBER OF HEADS. (Byte 5)
This field is set to four for the ProDrive LPS 540S, indicating that the number of heads is four. This field is set to eight for the ProDrive LPS 1080S.

STARTING CYLINDER - WRITE PRECOMPENSATION (bytes 6–8)
The value of this field is set to zero.

STARTING CYLINDER - REDUCED WRITE CURRENT (bytes 9–11)
The value of this field is set to zero.

DRIVE STEP RATE (bytes 12–13)
The value of this field is set to zero.

Note: The embedded SCSI controller on the drive handles the write precompensation starting cylinders, reduced write current starting cylinders, and drive step rate fields.

LANDING ZONE CYLINDER (Bytes 14–16)
This parameter is set to zero, because the drive automatically parks the heads in the landing zone at power off. This field applies only to drives that do not automatically seek to the landing zone, before stopping the spindle motor.

6.8.14 START STOP UNIT (Opcode=1Bh)

The START STOP UNIT command, shown in Table 6-60, informs the drive to either enable itself or disable itself for media access operations.

Table 6-60 *START STOP UNIT Command*

BYTE	BIT 7	6	5	4	3	2	1	0
0	OPCODE=1Bh							
1	LUN=0		RESERVED=0					IMED
2	RESERVED=0							
3	RESERVED=0							
4	RESERVED=0						LE=0	STRT
5	VU=0		RESERVED=0				F=0	L=0

IMED–IMMEDIATE (Byte 1, Bit 0)

An immediate bit of one causes the drive to return status information as soon as it validates the Command Descriptor Block (CDB). An immediate bit of zero tells the drive to return status after it completes the operation.

STRT–START (Byte 4, Bit 0)

A start bit of one requests that the drive be made ready for use. A start bit of zero asks that the drive be stopped (media cannot be accessed by the initiator).

When the drive is stopped, both the spindle motor and the actuator are disabled and do not draw appreciable currents. Starting the drive requires up to twenty seconds before commands requiring disk access can be executed. (See Section 6.8.1, "TEST UNIT READY (Opcode=00h)".)

6.8.15 SEND DIAGNOSTIC (Opcode=1Dh)

The SEND DIAGNOSTIC command, shown in Table 6-61, makes the drive perform diagnostic tests on itself.

Table 6-61 *SEND DIAGNOSTIC Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=1Dh							
1	LUN=0			PF=0	R=0	ST=1	DO=0	UO=0
2	RESERVED=0							
3	(MSB)	PARAMETER LIST LENGTH=0						
4		PARAMETER LIST LENGTH=0						(LSB)
5	VU=0		RESERVED=0				F=0	L=0

If the self test successfully passes, the command concludes with GOOD status; otherwise, CHECK CONDITION status is issued with the sense key set to HARDWARE ERROR.

6.8.16 READ CAPACITY (Opcode=25h)

The READ CAPACITY command, shown in Table 6-62, allows the initiator to gain access to information regarding the drive's capacity.

6.8.16.1 Command Structure

Table 6-62 *READ CAPACITY Command*

BYTE	BIT 7	6	5	4	3	2	1	0
0	OPCODE=25h							
1	LUN=0			RESERVED=0				RA=0
2	(MSB)	LBA						
3		LBA						
4		LBA						
5		LBA						(LSB)
6	RESERVED=0							
7	RESERVED=0							
8	RESERVED=0							PMI
9	VU=0		RESERVED=0				F=0	L=0

PMI–PARTIAL MEDIUM INDICATOR (Byte 8, Bit 0)

When PMI is one, the drive returns the Logical Block Address (LBA) and block length, in bytes, of the logical block after which a substantial delay in data transfer will occur. The LBA returned is greater than or equal to the LBA specified in the Command Descriptor Block (CDB).

When PMI is zero, the drive reports the last logical block's LBA and block length.

LBA (Bytes 2-5)

If the partial medium indicator bit is one, this field should contain the LBA from which the drive will measure the data transfer delay. If the partial medium indicator bit is zero, the initiator must set this field to zero. If this parameter reflects a non-zero LBA when PMI is zero, the drive terminates the command with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense code is set to INVALID FIELD IN CDB.

6.8.16.2 READ CAPACITY Data

The eight bytes of READ CAPACITY data, shown in Table 6-63, are sent during the command's DATA IN phase.

Table 6-63 *READ CAPACITY Data*

BYTE	BIT 7	6	5	4	3	2	1	0
0	(MSB)	RETURNED LBA						
1		RETURNED LBA						
2		RETURNED LBA						
3		RETURNED LBA						(LSB)
4	(MSB)	BLOCK LENGTH IN BYTES						
5		BLOCK LENGTH IN BYTES						
6		BLOCK LENGTH IN BYTES						
7		BLOCK LENGTH IN BYTES						(LSB)

6.8.17 READ EXTENDED (Opcode=28h)

The READ EXTENDED command, shown in Table 6-64, requests that the drive transfer data to the initiator from the disk.

Table 6-64 *READ EXTENDED Command*

BYTE	BIT 7	6	5	4	3	2	1	0
0	OPCODE=28h							
1	LUN=0			DP=0	FA=0	RESERVED=0		RA=0
2	(MSB)	LBA						
3		LBA						
4		LBA						
5		LBA						(LSB)
6	RESERVED=0							
7	(MSB)	TRANSFER LENGTH						
8		TRANSFER LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

LBA (Bytes 2-5)

This parameter tells the drive the logical block to start reading from.

TRANSFER LENGTH (Bytes 7&8)

The drive reads contiguous logical blocks, starting with the logical block whose LBA matches the LBA in the Command Descriptor Block (CDB), until the number read equals the transfer length. See Section 6.6.3, "Transferring Data." A transfer length of zero results in no data being transferred.

6.8.18 WRITE EXTENDED (Opcode=2Ah)

The WRITE EXTENDED command, shown in Table 6-65, instructs the drive to write the data transferred by the initiator to the disk.

Table 6-65 *WRITE EXTENDED Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=2Ah							
1	LUN=0			DP=0	FA=0	RESERVED=0		RA=0
2	(MSB)	LBA						
3		LBA						
4		LBA						
5		LBA						(LSB)
6	RESERVED=0							
7	(MSB)	TRANSFER LENGTH						
8		TRANSFER LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

LBA (Bytes 2-5)

This parameter identifies the logical block at which to begin writing.

TRANSFER LENGTH (Bytes 7&8)

The drive writes the transferred data to the disk, starting at the logical block whose LBA matches the LBA in the Command Descriptor Block (CDB), until the number of contiguous logical blocks written equals the transfer length. See Section 6.6.3, "Transferring Data." A transfer length of zero results in no data being transferred.

6.8.19 SEEK EXTENDED (Opcode=2Bh)

The SEEK EXTENDED command, shown in Table 6-66, communicates to the drive that it should seek to the Logical Block Address (LBA) specified.

Table 6-66 *SEEK EXTENDED Command*

<table>
<tr><th rowspan="2">BYTE</th><th colspan="8">BIT</th></tr>
<tr><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="8">OPCODE=2Bh</td></tr>
<tr><td>1</td><td colspan="3">LUN=0</td><td colspan="5">RESERVED=0</td></tr>
<tr><td>2</td><td>(MSB)</td><td colspan="6">LBA</td><td></td></tr>
<tr><td>3</td><td></td><td colspan="6">LBA</td><td></td></tr>
<tr><td>4</td><td></td><td colspan="6">LBA</td><td></td></tr>
<tr><td>5</td><td></td><td colspan="6">LBA</td><td>(LSB)</td></tr>
<tr><td>6</td><td colspan="8">RESERVED=0</td></tr>
<tr><td>7</td><td colspan="8">RESERVED=0</td></tr>
<tr><td>8</td><td colspan="8">RESERVED=0</td></tr>
<tr><td>9</td><td colspan="2">VU=0</td><td colspan="4">RESERVED=0</td><td>F=0</td><td>L=0</td></tr>
</table>

LBA (Bytes 2-5)

This parameter identifies the logical block, from which to calculate the cylinder and head, used by the drive to position the actuator.

6.8.20 WRITE AND VERIFY (Opcode=2Eh)

The WRITE AND VERIFY command, shown in Table 6-67, instructs the drive to write the data transferred by the initiator to the disk, and then verify that the data was written correctly. A medium verification is performed with no data comparison. The verification process reads the data written, checks this data against the ECC bytes read, and then transfers either GOOD or CHECK CONDITION status to the initiator.

Table 6-67 *WRITE AND VERIFY Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=2Eh							
1	LUN=0			DP=0	R=0	R=0	BC=0	RA=0
2	(MSB)	LBA						
3		LBA						
4		LBA						
5		LBA						(LSB)
6	RESERVED=0							
7	(MSB)	TRANSFER LENGTH						
8		TRANSFER LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

LBA (Bytes 2-5)

This parameter identifies the logical block at which to begin writing. It also is the address of the first logical block read and verified during the verification process.

TRANSFER LENGTH (Bytes 7&8)

The drive writes the transferred data to the disk, starting at the logical block whose LBA matches the LBA in the Command Descriptor Block (CDB), until the number of contiguous logical blocks written equals the transfer length. It then goes back and reads the data written to the disk, verifying that it is correct. The drive verifies logical blocks until the number of contiguous logical blocks verified equals the transfer length. A transfer length of zero results in no data being transferred.

6.8.21 VERIFY (Opcode=2Fh)

The VERIFY command, shown in Table 6-68, makes the drive check the data written on the disk. A medium verification is performed with no data comparison. The verification process reads the data written, checks this data against the ECC bytes read, and then transfers either GOOD or CHECK CONDITION status to the initiator.

Table 6-68 *VERIFY Command*

BYTE	BIT 7	6	5	4	3	2	1	0
0	OPCODE=2Fh							
1	LUN=0			DP=0	R=0	R=0	BC=0	RA=0
2	(MSB)	LBA						
3		LBA						
4		LBA						
5		LBA						(LSB)
6	RESERVED=0							
7	(MSB)	VERIFICATION LENGTH						
8		VERIFICATION LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

LBA (Bytes 2-5)

This parameter is the address of the first logical block read and verified during the verification process.

VERIFICATION LENGTH (Bytes 7&8)

The drive reads the data on the disk, starting at the logical block whose LBA matches the LBA in the Command Descriptor Block (CDB), verifying that it is correct. The drive verifies logical blocks until the number of contiguous logical blocks verified equals the verification length. A verification length of zero results in no verify operation.

6.8.22 SYNC CACHE (Opcode=35h)

Information not yet available.

6.8.23 READ DEFECT DATA (Opcode=37h)

The READ DEFECT DATA command, shown in Table 6-69, requests that the drive transfer the media-defect data to the initiator.

6.8.23.1 Command Structure

Table 6-69 *READ DEFECT DATA Command*

BYTE	BIT 7	6	5	4	3	2	1	0
0	OPCODE=37h							
1	LUN=0			RESERVED=0				
2	RESERVED=0			P LST	G LST	DEFECT LIST FORMAT		
3	RESERVED=0							
4	RESERVED=0							
5	RESERVED=0							
6	RESERVED=0							
7	(MSB)	ALLOCATION LENGTH						
8		ALLOCATION LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

P LST–PRIMARY LIST (Byte 2, Bit 4)

When P LST is one, the drive returns the primary list of defects as part of the defect data. When P LST is zero, the drive excludes the Primary List from the defect data sent to the initiator.

The primary list of defects, also called the P LIST, is a record of defective sectors found on the medium at the factory, before the drive was shipped.

G LST–GROWN LIST (Byte 2, Bit 3)

When G LST is one, the drive returns the grown defect list as part of the defect data. When G LST is zero, the drive excludes the Grown List from the defect data sent to the initiator.

The grown defect list, also called the G LIST, is a record of defective sectors found on the medium during the drive's operation.

Note: If both P LST and G LST are one, the drive returns both lists, with the primary list of defects preceding the grown defect list. If both P LST and G LST are zero, the drive sends only the defect data header.

DEFECT LIST FORMAT (Byte 2, Bits 2-0)

This field is used by the initiator to indicate to the drive the preferred format for the defect list. If the initiator requests a format not supported by the drive, the drive will default to Physical Sector Format for the returned defect data. The command completes with CHECK CONDITION status. The sense key is set to RECOVERED ERROR and the additional sense code is set to REQUESTED FORMAT NOT AVAILABLE.

The drive supports two defect data formats:

- **Physical Sector Format (101b)**
  This format returns the cylinder, head, and sector of the media defect.
- **Bytes Offset from Index Format (100b)**
  This format reports the media defect in terms of its cylinder and head, and the number of bytes from the index to the defect. The offset number given is the first byte of the defective sector.

ALLOCATION LENGTH (Bytes 7&8)

This parameter indicates the maximum number of bytes allocated by the initiator for returned defect data. Acceptable values are 00h-FFh. The drive completes the DATA IN phase when it has transferred to the initiator either the number of bytes specified by the allocation length or all available defect data, whichever is less.

6.8.23.2 READ DEFECT DATA Format

The defect data, shown in Table 6-70, Table 6-71, and Table 6-72, contains a four-byte header, followed by a defect list composed of zero or more defect descriptors. Each defect descriptor carries the eight-byte location of a media defect.

Table 6-70 *READ DEFECT DATA – Header*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	RESERVED=0							
1	RESERVED=0			P LST	G LST	DEFECT LIST FORMAT		
2	(MSB)	DEFECT LIST LENGTH						
3		DEFECT LIST LENGTH						(LSB)

P LST–PRIMARY LIST (Byte 1, Bit 4)

A P LST of one indicates that the data returned contains the primary list of defects. A P LST of zero indicates that the data returned does not contain the Primary List.

G LST–GROWN LIST (Byte 1, Bit 3)

A G LST of one indicates that the data returned contains the grown defect list. A G LST of zero indicates that the data returned does not contain the Grown List.

DEFECT LIST FORMAT (Byte 1, Bits 2-0)

This parameter specifies the actual format of the defect list–Physical Sector Format (101b) or Bytes Offset from Index Format (100b). The defect data is presented in one format only; there is no intermixing of Physical Sector Format defect descriptors with Bytes Offset from Index defect descriptors.

DEFECT LIST LENGTH (Bytes 2&3)

The drive communicates to the initiator the length of the defect list, in bytes, through the defect list length.

Since the defect list consists entirely of eight-byte block descriptors, this parameter's value is equal to the number of block descriptors times eight.

If the allocation length is insufficient to allow the transfer of all defect descriptors, the defect list length is not adjusted to reflect their truncation.

Table 6-71 *READ DEFECT DATA – Defect Descriptor(s) in Physical Sector Format*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	(MSB)	CYLINDER NUMBER						
1		CYLINDER NUMBER						
2		CYLINDER NUMBER						(LSB)
3	HEAD NUMBER							
4	(MSB)	SECTOR NUMBER						
5		SECTOR NUMBER						
6		SECTOR NUMBER						
7		SECTOR NUMBER						(LSB)

Table 6-72 *READ DEFECT DATA – Defect Descriptor(s) in Bytes-Offset-from-Index Format*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	(MSB)	CYLINDER NUMBER						
1		CYLINDER NUMBER						
2		CYLINDER NUMBER						(LSB)
3	HEAD NUMBER							
4	(MSB)	BYTES OFFSET						
5		BYTES OFFSET						
6		BYTES OFFSET						
7		BYTES OFFSET						(LSB)

6.8.24 WRITE BUFFER (Opcode=3Bh)

The WRITE BUFFER command, shown in Table 6-73, is used in conjunction with the READ BUFFER command to test the drive's memory and the integrity of the SCSI bus. This command does not alter the media.

Table 6-73 *WRITE BUFFER Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=3Bh							
1	LUN=0			RESERVED=0		MODE		
2	BUFFER ID = 0							
3	(MSB)	BUFFER OFFSET						
4		BUFFER OFFSET						
5		BUFFER OFFSET						(LSB)
6	(MSB)	PARAMETER LIST LENGTH						
7		PARAMETER LIST LENGTH						
8		PARAMETER LIST LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

MODE (Byte 1, Bits 2-0)

The function of this command and the meaning of the fields within the Command Descriptor Block (CDB) depend on the contents of this field. (See Table 6-74.) In all cases, a DATA OUT phase occurs after the requested mode is determined. The use of the DATA OUT phase is mode specific. To determine the initiator's role and the meaning of the bytes crossing the bus, the following subsections are given.

Table 6-74 *WRITE BUFFER Mode Field*

BIT 2	BIT 1	BIT 0	MODE
0	0	0	Combined header and data
0	0	1	Not supported (Vendor unique)
0	1	0	Data
0	1	1	Reserved
1	0	0	Not supported (Download microcode)
1	0	1	Download microcode and save
1	1	0	Reserved
1	1	1	Reserved

6.8.24.1 Combined Header and Data Mode (000b)

In this mode, a four-byte header that contains all reserved bytes precedes the data to be transferred. The drive writes the data supplied by the initiator to its data buffer beginning at the buffer's first byte.

BUFFER ID (Byte 2)

The initiator must set this field to zero.

BUFFER OFFSET (Bytes 3–5)

These bytes are reserved and must be set to zero.

PARAMETER LIST LENGTH (Bytes 6-8)

This parameter tells the drive the number of bytes to transfer from the initiator during the DATA OUT phase. This number includes four bytes of header, so the data length to be stored in the drive's buffer is parameter list length minus four.

The initiator should ensure that the parameter list length is not greater than four plus the buffer capacity. The buffer capacity can be found in the header sent by the drive in response to a READ BUFFER command having mode=000b. (See Section 6.8.24.1, "Combined Header and Data Mode (000b).") Therefore, since the value returned in the header is 60000h, the maximum parameter list length is 60004h. If the parameter list length exceeds this, the drive terminates the command with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense code is set to INVALID FIELD IN CDB.

6.8.24.2 Data Mode (010b)

In this mode, the DATA OUT phase contains buffer data. Data is written to the drive buffer starting at the first byte of the specified segment.

BUFFER ID (Byte 2)

These bytes are reserved and must be set to zero.

BUFFER OFFSET (Bytes 3–5)

This field contains the offset within the buffer at which the write data will be loaded. The number must be an even offset.

PARAMETER LIST LENGTH (Bytes 6-8)

The parameter list length provides the drive with information regarding the number of bytes to transfer from the initiator during the DATA OUT phase.

The initiator should ensure that the parameter list length does not exceed the capacity of the chosen segment. The segment's capacity is available in the READ BUFFER descriptor. (See Section 6.8.25.3, "Descriptor (011b).") If the parameter list length specifies a transfer that would exceed the capacity, the drive terminates the command with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense code is set to INVALID FIELD IN CDB.

6.8.24.3 Download Microcode and Save Mode (101b)

In this mode, microcode or control information is transferred to the drive and, if the WRITE BUFFER command is completed successfully, saved in a non-volatile memory space. The downloaded code is then effective after each power-cycle and RESET condition, until it is replaced by new code in another download microcode and save operation. This mode will operate properly if and only if the microcode is provided by Quantum.

BUFFER ID (Byte 2)

This field is ignored, and can therefore be set to any value by the initiator.

BUFFER OFFSET (Bytes 3–5)

This field is ignored, and can therefore be set to any value by the initiator.

PARAMETER LIST LENGTH (Bytes 6-8)

The initiator must set this field to C000h for correct operation, because this is the length of the Quantum provided microcode.

When the download microcode and save command has completed successfully, the drive generates a UNIT ATTENTION condition for all initiators except the one that issued the WRITE BUFFER command.

When reporting the UNIT ATTENTION condition, the drive sets the additional sense code to TARGET OPERATING CONDITIONS HAVE CHANGED.

6.8.25 READ BUFFER (Opcode=3Ch)

The READ BUFFER command, shown in Table 6-75, is used in conjunction with the WRITE BUFFER command to test the drive's memory, and the integrity of the SCSI bus. This command does not alter the media.

Table 6-75 *READ BUFFER Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=3Ch							
1	LUN=0			RESERVED=0		MODE		
2	BUFFER ID = 0							
3	(MSB)	BUFFER OFFSET						
4		BUFFER OFFSET						
5		BUFFER OFFSET						(LSB)
6	(MSB)	ALLOCATION LENGTH						
7		ALLOCATION LENGTH						
8		ALLOCATION LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

MODE (Byte 1, Bits 2-0)

The function of this command and the meaning of the fields within the Command Descriptor Block (CDB) depend on the contents of this field. (See Table 6-76.)

Table 6-76 *READ BUFFER Mode Field*

BIT2	BIT 1	BIT 0	MODE
0	0	0	Combined header and data
0	0	1	Not supported (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

ALLOCATION LENGTH (Bytes 6-8)

This parameter specifies the maximum number of bytes allocated by the initiator for returned READ BUFFER data. Acceptable values are 000000h-FFFFFFh.

6.8.25.1 Combined Header And Data Mode (000b)

In this mode, a four-byte header followed by data bytes, from the drive's data buffer, are returned to the initiator during the DATA IN phase. The data bytes begin with the first byte of the drive's data buffer and continue to the last byte, allocation length permitting.

BUFFER ID (Byte 2)

The initiator must set this field to zero.

BUFFER OFFSET (Bytes 3–5)

The initiator must set this field to zero.

ALLOCATION LENGTH (Bytes 6-8)

The drive terminates the DATA IN phase when it has transferred to the initiator either the allocated number of bytes or all available header and buffer data, whichever is less.

Table 6-77 *READ BUFFER – Header*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	RESERVED=0							
1	(MSB)	BUFFER CAPACITY						
2		BUFFER CAPACITY						
3		BUFFER CAPACITY						(LSB)

BUFFER CAPACITY (Bytes 1-3)

This parameter is the total number of data bytes that the drive's data buffer can retain. This number is not reduced to reflect the allocation length, nor is it reduced to reflect the actual number of bytes written using the WRITE BUFFER command.

Following the READ BUFFER header, the drive transfers data from its data buffer. The number of data bytes sent after the header either equals the allocation length minus four or 416 kilobytes, whichever is less.

6.8.25.2 Data Mode (010b)

In this mode, the DATA IN phase contains buffer data. Data is transferred from the drive buffer starting at the first byte of the specified segment.

BUFFER ID (Byte 2)

These bytes are reserved and must be set to zero.

BUFFER OFFSET (Bytes 3–5)

This field contains the offset within the buffer at which the write data will be loaded. The number must be an even offset.

ALLOCATION LENGTH (Bytes 6-8)

The drive terminates the DATA IN phase when it has transferred to the initiator either the allocated number of bytes or all available buffer data from the selected segment, whichever is less.

6.8.25.3 Descriptor (011b)

In this mode, READ BUFFER descriptor information is returned. The DATA IN phase number of bytes does not exceed four.

BUFFER ID (Byte 2)

This parameter provides a way for the initiator to access information on each of the buffer's segments. The initiator communicates to the drive the desired segment's number by placing it in this field. The drive supports only buffer 0. The drive places all zeroes in any other BUFFER ID descriptor.

ALLOCATION LENGTH (Bytes 6-8)

An allocation length of 4, or greater, is recommended for the Empire 540/1080S hard disk drive. An allocation length greater than 4 results in the transfer of four bytes. Any value less than or equal to 4 causes that number of bytes to be sent.

Table 6-78 *READ BUFFER – Descriptor*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	BUFFER OFFSET=FFh							
1	(MSB)	BUFFER CAPACITY						
2		BUFFER CAPACITY						
3		BUFFER CAPACITY						(LSB)

BUFFER OFFSET (Byte 0)

The drive sets this field to 1 to indicate that it supports even buffer offsets.

BUFFER CAPACITY (Bytes 1-3)

This parameter notifies the initiator of the selected segment's size, in bytes.

6.8.26 READ LONG (Opcode=3Eh)

The READ LONG command, shown in Table 6-79, directs the drive to transfer data to the initiator. The total data transferred is 526 bytes, composed of 512 data bytes, 2 cross check bytes, and 12 ECC bytes, in that order. The logical block is read without any correction made by the drive.

Table 6-79 *READ LONG Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=3Eh							
1	LUN=0			RESERVED=0			CT=0	RA=0
2	(MSB)	LBA						
3		LBA						
4		LBA						
5		LBA						(LSB)
6	RESERVED=0							
7	(MSB)	BYTE TRANSFER LENGTH						
8		BYTE TRANSFER LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

LBA (Bytes 2-5)

This parameter is the Logical Block Address (LBA) of the data sector read.

BYTE TRANSFER LENGTH (Bytes 7&8)

This parameter specifies the number of bytes to be transferred from the drive to the initiator. If the bytes are set to zero, no bytes are transferred. If the bytes are set to 526 (20Eh), for example, then that number of bytes are transferred. If this field is set to an illegal value, the drive terminates the command with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense code is set to INVALID FIELD IN CDB. The sense data's information field is set to the difference of the requested length minus 526 bytes (for this example). Negative values are indicated by two's complement notation.

6.8.27 WRITE LONG (Opcode=3Fh)

The WRITE LONG command, shown in Table 6-80, directs the drive to write the data transferred by the initiator to the disk. The total data transferred is 526 bytes, composed of 512 data bytes, 2 cross check bytes, and 12 ECC bytes, in that order.

Table 6-80 *WRITE LONG Command*

BYTE	BIT 7	6	5	4	3	2	1	0
0	OPCODE=3Fh							
1	LUN=0			RESERVED=0				RA=0
2	(MSB)	LBA						
3		LBA						
4		LBA						
5		LBA						(LSB)
6	RESERVED=0							
7	(MSB)	BYTE TRANSFER LENGTH						
8		BYTE TRANSFER LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

LBA (Bytes 2-5)

This parameter is the Logical Block Address (LBA) of the data sector written.

BYTE TRANSFER LENGTH (Bytes 7&8)

This parameter specifies the number of bytes to be transferred from the drive to the initiator. If the bytes are set to zero, no bytes are transferred. If the bytes are set to 526 (20Eh), for example, then that number of bytes are transferred. If this field is set to an illegal value, the drive terminates the command with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense code is set to INVALID FIELD IN CDB. The sense data's information field is set to the difference of the requested length minus 526 bytes (for this example). Negative values are indicated by two's complement notation.

6.8.28 CHANGE DEF (Opcode = 40h)

Information not yet available.

6.8.29 MODE SELECT EXTENDED (Opcode=55h)

The MODE SELECT EXTENDED command, shown in Table 6-81, allows the initiator to set the drive's operating mode parameters. A MODE SELECT EXTENDED command overrides any previous selection of operating mode parameters, even by another initiator. Execution of this command creates a UNIT ATTENTION condition for all other initiators, provided that at least one parameter changes.

6.8.29.1 Command Structure

Table 6-81 *MODE SELECT EXTENDED Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=55h							
1	LUN=0			PF=0	RESERVED=0			SP
2	RESERVED=0							
3	RESERVED=0							
4	RESERVED=0							
5	RESERVED=0							
6	RESERVED=0							
7	(MSB)	PARAMETER LIST LENGTH						
8		PARAMETER LIST LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

When the drive motor spins up at power on, the drive reads the operating mode parameters most recently saved, from a reserved cylinder, and sets a UNIT ATTENTION condition for all initiators. The additional sense code accompanying the UNIT ATTENTION sense key is determined as follows:

- If the drive cannot successfully read the parameter values from the Saved Mode Table, the drive accesses the Default Mode Table to set the state of the operating mode parameters. The additional sense code is set to PARAMETERS CHANGED, as though another initiator had altered the parameters.
- If the drive is able to read the parameter values, the additional sense code is set to POWER ON, RESET, OR BUS DEVICE RESET OCCURRED.

Note: If the disable unit attention bit on the Quantum-Unique Drive-Control Page (Page Code 39h) is set to one, a UNIT ATTENTION condition does not occur.

SP–SAVE PAGES (Byte 1, Bit 0)

When SP is one, the drive copies the operating mode parameters in the Current Mode Table to the Saved Mode Table, after updating the Current Mode Table if pages were sent. When SP is zero, the drive performs the specified MODE SELECT EXTENDED operation and does not alter the contents of the Saved Mode Table.

PARAMETER LIST LENGTH (Bytes 7&8)

For this command, the initiator must be able to communicate the changes to be made in the Current Mode Table. A parameter list accomplishes this task. The drive uses this parameter's value to determine the length of the parameter list, and then transfers that many bytes.

6.8.29.2 Parameter List Format

The MODE SELECT EXTENDED parameter list, shown in Table 6-82, Table 6-83, and Table 6-84, contains an eight-byte header, followed by zero or one block descriptors, followed by zero or more pages. The pages sent correspond to unique divisions of the Current Mode Table. The drive determines which pages are included in the parameter list, accesses the Current Mode Table, and places the parameter values from these pages into the correct divisions. Pages are discussed separately in Section 6.9 Mode Pages, because they apply to both the MODE SELECT EXTENDED and MODE SENSE EXTENDED commands. Before using this command, or the MODE SENSE EXTENDED command, be familiar with the information provided regarding pages.

Table 6-82 *MODE SELECT EXTENDED Parameter List – Header*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	RESERVED=0							
1	RESERVED=0							
2	MEDIUM TYPE=0							
3	RESERVED=0							
4	RESERVED=0							
5	RESERVED=0							
6	(MSB)	BLOCK DESCRIPTOR LENGTH						
7		BLOCK DESCRIPTOR LENGTH						(LSB)

BLOCK DESCRIPTOR LENGTH (Bytes 6&7)

This parameter informs the drive of the block descriptors' length, in bytes. A block descriptor length of zero indicates that there are no block descriptors in the parameter list. The Empire 540/1080S hard disk drive uses a single eight-byte block descriptor. Therefore, this parameter must be set to either zero or eight by the initiator.

Table 6-83 *MODE SELECT EXTENDED Parameter List – Block Descriptor*

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

NUMBER OF BLOCKS (Bytes 1-3)

This parameter conveys to the drive the number of logical blocks on the media that match the block length specified in this block descriptor. A value of zero notes that all remaining logical blocks on the drive have the media characteristics specified by the block descriptor. Any non zero value is ignored.

BLOCK LENGTH (Bytes 5-7)

The initiator must set this field to a value that corresponds to the block length of the drive. For example, if the block length is 512 bytes, the block length bytes must be set to 200h.

Table 6-84 *MODE SELECT EXTENDED Parameter List – Page(s)*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0+	SEE INDIVIDUAL SECTION ON PAGES							

6.8.30 MODE SENSE EXTENDED (Opcode=5Ah)

The MODE SENSE EXTENDED command, shown in Table 6-85, tells the drive to report its operating mode parameters to the initiator. This command complements the MODE SELECT EXTENDED command. Before preceding further, it is advisable to read the introduction to Section 6.9 Mode Pages.

6.8.30.1 Command Structure

Table 6-85 *MODE SENSE EXTENDED Command*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	OPCODE=5Ah							
1	LUN=0			R=0	DB=0	RESERVED=0		
2	PC		PAGE CODE					
3	RESERVED=0							
4	RESERVED=0							
5	RESERVED=0							
6	RESERVED=0							
7	(MSB)	ALLOCATION LENGTH						
8		ALLOCATION LENGTH						(LSB)
9	VU=0		RESERVED=0				F=0	L=0

PC–PAGE CONTROL (Byte 2, Bits 7&6)

This parameter defines the type of page parameter values the drive reports. There are four possible types, or formats, for the fields and bits contained in each page:

Note: The words "MODE SELECT command" apply to both the MODE SELECT command and the MODE SELECT EXTENDED command in the following definitions.

1. **Current Values (00b)**

   The drive sends the page, or pages, defined by the page code, with fields and bits set to the current operating mode values. The current values are:

   - Those set by the last MODE SELECT command successfully completed.
   - Identical to the saved values–if there has been no MODE SELECT command since the last power on or RESET condition.

2. **Changeable Values (01b)**

   The drive sends the page, or pages, defined by the page code, with fields and bits that can be modified set to one. Fields and bits that cannot be altered by the initiator are set to zero.

3. **Default Values (10b)**

   The drive sends the page, or pages, defined by the page code, with fields and bits set to the default operating mode values–that is, as shipped from the factory. Fields and bits not supported by the drive are set to zero.

Note: To determine whether a returned value of zero indicates a default parameter or an unsupported parameter, the initiator can examine the changeable values.

4. **Saved Values (11b)**

   The drive sends the page, or pages, defined by the page code, with fields and bits set to the saved operating mode values. The saved values are:

   - Those set by the last MODE SELECT command that had the save pages bit set to one.
   - Identical to the default values–if there has been no MODE SELECT command that had the save pages bit set to one.

Note: For all page control parameters, the value returned in the page length field indicates the number of bytes the drive supports within each page. This value must be the value entered for the page length when issuing a MODE SELECT EXTENDED command. (See Section 6.9, "MODE PAGES.")

PAGE CODE (Byte 2, Bits 5-0)

This parameter notifies the drive of the specific page, or pages, the initiator wants. Table 6-86 shows the page codes for the pages the drive supports.

Table 6-86 *Page Codes as a Function of Pages Supported*

PAGE CODE	PAGES(S) RETURNED
00h	Disable Unit Attention Page
01h	Read-Write Error Recovery Page
02h	Disconnect-Reconnect Page
03h	Format Device Page
04h	Rigid Disk Geometry Page
08h	Caching Page
0Ah	Control Mode Page
0Ch	Notch and Partition Page
32h	Automatic-Shutdown Control Page
37h	Quantum-Unique Control Page
38h	Quantum-Unique Cache-Control Page
39h	Quantum-Unique Drive-Control Page
3Fh	All Pages (Order=01h, 02h, ..., 39h, 00h)

ALLOCATION LENGTH (Bytes 7&8)

This parameter specifies the maximum number of bytes allocated by the initiator for returned MODE SENSE EXTENDED data. Acceptable values are 00h-FFh. The drive terminates the DATA IN phase when it has transferred to the initiator either the number of bytes specified by the allocation length or all available MODE SENSE EXTENDED data, whichever is less.

6.8.30.2 MODE SENSE EXTENDED Data Format

The MODE SENSE EXTENDED data, shown in Table 6-87, Table 6-88, and Table 6-89, contains an eight-byte header, followed by one block descriptor, followed by one or more pages. Pages are discussed separately in Section 6.9 Mode Pages, because they apply to both the MODE SELECT EXTENDED and MODE SENSE EXTENDED commands. Before using this command, or the MODE SELECT EXTENDED command, be familiar with the information provided regarding pages.

Table 6-87 *MODE SENSE EXTENDED Data – Header*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	(MSB)	MODE DATA LENGTH						
1		MODE DATA LENGTH						(LSB)
2	MEDIUM TYPE=0							
3	WP=0	RSVD=0		DF=0	RESERVED=0			
4	RESERVED=0							
5	RESERVED=0							
6	(MSB)	BLOCK DESCRIPTOR LENGTH=00h						
7		BLOCK DESCRIPTOR LENGTH=08h						(LSB)

Note: The medium is not write protected.

MODE DATA LENGTH (Bytes 0&1)

The drive transfers the MODE SENSE EXTENDED data during the DATA IN phase of the command. The first byte it sends is the mode data length. This parameter informs the initiator of the number of available MODE SENSE EXTENDED data bytes. The mode data length does not include itself.

BLOCK DESCRIPTOR LENGTH (Bytes 6&7)

The Empire 540/1080S hard disk drive uses a single eight-byte block descriptor. This parameter gives the initiator the block descriptor's length, in bytes. Therefore, the drive sets this field to 08h.

Table 6-88 *MODE SENSE EXTENDED Data – Block Descriptor*

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

NUMBER OF BLOCKS (Bytes 1-3)

This parameter conveys to the initiator the number of logical blocks on the media that match the block length specified in this block descriptor. A value of zero notes that all remaining logical blocks on the drive have the media characteristics specified by the block descriptor.

BLOCK LENGTH (Bytes 5-7)

The initiator must set this field to a value that corresponds to the block length of the drive. For example, if the block length is 512 bytes, the block length bytes must be set to 200h.

Table 6-89 *MODE SENSE EXTENDED Data – Page(s)*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0+	SEE INDIVIDUAL SECTION ON PAGES							

6.9 MODE PAGES

To avoid having to specify all of the operating mode parameters each time an initiator wants to change one, the operating mode parameters are split into pages. A page is the smallest unit that can be specified in a MODE SELECT, MODE SELECT EXTENDED, MODE SENSE, or MODE SENSE EXTENDED command. Pages are numbered for reference, with each page containing parameters grouped by functionality.

For MODE SELECT and MODE SELECT EXTENDED commands, 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 can be included in any order.

Table 6-90 *Mode Pages Supported*

PAGE NUMBER/CODE	DESCRIPTION
00h	Disable Unit Attention Page
01h	Read-Write Error Recovery Page
02h	Disconnect-Reconnect Page
03h	Format Device Page
04h	Rigid Disk Geometry Page
08h	Caching Page
0Ah	Control Mode Page
0Ch	Notch and Partition Page
32h	Automatic-Shutdown Control Page
37h	Quantum-Unique Control Page
38h	Quantum-Unique Cache-Control Page
39h	Quantum-Unique Drive-Control Page

Each mode page contains a page code, a page length, and a set of mode parameters. In addition to this, for the MODE SENSE and MODE SENSE EXTENDED commands, an additional bit is provided to indicate if the page is savable or not. This bit is reserved when using the MODE SELECT and MODE SELECT EXTENDED commands.

Table 6-91 *Mode Page Format*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	PS	R=0	PAGE CODE					
1	PAGE LENGTH							
2+	MODE PARAMETERS							

PS–PARAMETERS SAVABLE (Byte 0, Bit 7)

When using the MODE SENSE and MODE SENSE EXTENDED commands, a parameters savable bit of one indicates that the operating mode parameters on this page can be saved by the drive to the Saved Mode Table. A parameters savable bit of zero indicates that the supported parameters cannot be saved.

When using the MODE SELECT and MODE SELECT EXTENDED commands, PS is reserved.

PAGE LENGTH (Byte 1)

The page length is the mode parameters' length, in bytes. The initiator must set this field to the value that is returned for the page by the MODE SENSE or MODE SENSE EXTENDED command.

The following subsections are the individual page descriptions. Each subsection elaborates on the mode parameters applicable to that page.

6.10 ELECTRICAL CHARACTERISTICS OF THE SCSI BUS

The following description of the SCSI electrical characteristics includes a definition of the signal lines as well as their timing characteristics.

6.10.1 Signal Lines

There are a total of 18 signal lines on the bus, of which 9 are control lines and 9 are data lines (including the parity line). Table 6-92 defines the SCSI bus signal lines, Table 6-93 shows which SCSI devices assert these signals and when they can do so, and Table 6-94 defines the conductor signal assignments for the SCSI connector.

6.10.2 Timing Characteristics

Timing characteristics for the SCSI bus are given in the following subsections.

6.10.2.1 ARBITRATION DELAY

The minimum amount of time that must elapse between the point at which a SCSI device asserts BSY for arbitration and the point at which it examines the DATA BUS (to determine if it has won arbitration).

ARBITRATION DELAY is 2.4 microseconds.

6.10.2.2 ASSERTION PERIOD

The minimum amount of time that REQ (drive) and ACK (initiator) must be asserted during synchronous data transfers.

ASSERTION PERIOD is 90 nanoseconds.

6.10.2.3 BUS CLEAR DELAY

The maximum amount of time that may elapse prior to a SCSI device releasing all signals 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 is 800 nanoseconds following a lost arbitration attempt or detection of RST on the bus, and is 1200 nanoseconds minus the time it took the SCSI device to detect BUS FREE following Case 1 above.

Table 6-92 *SCSI Bus Signal Lines*

SIGNAL LINE	USE
BSY	An "OR-tied" signal line used to indicate that the bus is occupied by a device.
SEL	An "OR-tied" signal line used by an initiator to select the drive, or by the drive to reselect an initiator.
C/D	A signal line driven by the drive to indicate whether CONTROL or DATA information is on the DATA BUS. True identifies CONTROL.
I/O	A signal line driven by the drive to indicate the direction of data movement on the DATA BUS. This signal line is also used to distinguish between the SELECTION and RESELECTION phases.
MSG	A signal line driven by the drive to indicate that a message is being transferred.
REQ	A signal line driven by the drive to initiate a REQ/ACK data transfer handshake.
ACK	A signal line driven by the initiator to set the acknowledgment portion of the REQ/ACK handshake.
ATN	A signal line driven by an initiator to indicate an ATTENTION condition.
RST	A signal line that indicates a RESET condition.
DB(7)-DB(0)	The eight data-bit signal lines that form a majority of the DATA BUS. DB(7) is the most significant bit and has the highest priority during the ARBITRATION phase.
DB(P)	The parity-bit signal line that forms the remainder of the DATA BUS. Parity is odd and is undefined during the ARBITRATION phase.

Table 6-93 *Signal Sources*

BUS PHASE	BSY	SEL	MSG, C/D, I/O, REQ	ATN, ACK	DB(7-0, P)
BUS FREE	None	None	None	None	None
ARBITRATION	All	Winner	None	None	SCSI ID
SELECTION	I&D	Initiator	None	Initiator	Initiator
RESELECTION	I&D	Drive	Drive	Initiator	Drive
COMMAND	Drive	None	Drive	Initiator	Initiator
DATA IN	Drive	None	Drive	Initiator	Drive
DATA OUT	Drive	None	Drive	Initiator	Initiator
STATUS	Drive	None	Drive	Initiator	Drive
MESSAGE IN	Drive	None	Drive	Initiator	Drive
MESSAGE OUT	Drive	None	Drive	Initiator	Initiator

Note: All = All SCSI devices that are actively arbitrating assert the signal.
SCSI ID = A unique bit on the DATA BUS is asserted by each SCSI device actively arbitrating.
I&D = The initiator, drive, or both assert the signal.
Initiator = If the signal is asserted, it is only done by the active initiator.
Winner = The signal is asserted by the SCSI device that wins arbitration.

Table 6-94 *System Cable Pin-Out*

SCSI BUS PIN	SIGNAL	SCSI BUS PIN	SIGNAL
2	DB(0)	26	Terminator Power
4	DB(1)	28	Ground
6	DB(2)	30	Ground
8	DB(3)	32	ATN
10	DB(4)	34	Ground
12	DB(5)	36	BSY
14	DB(6)	38	ACK
16	DB(7)	40	RST
18	DB(P)	42	MSG
20	Ground	44	SEL
22	Ground	46	C/D
24	–	48	REQ
25		50	I/O

Note: All odd pins, except pin 25, are grounded. Pin 25 is open.

6.10.2.4 BUS FREE DELAY

The minimum amount of time that must elapse between the point at which a SCSI device detects the BUS FREE phase and the point at which it asserts BSY in an attempt to get the SCSI bus.

BUS FREE DELAY is 800 nanoseconds.

6.10.2.5 BUS SET DELAY

The maximum amount of time that may elapse between the point at which a SCSI device detects the BUS FREE phase and the point at which it asserts BSY and its SCSI ID in an attempt to get the bus.

BUS SET DELAY is 1.8 microseconds.

6.10.2.6 BUS SETTLE DELAY

The minimum amount of time that must elapse after a SCSI device detects a transition on certain control signals before the transition is valid.

For instance, when BSY and SEL both go false, a SCSI device must wait a BUS SETTLE DELAY before “detecting” the BUS FREE phase.

BUS SETTLE DELAY is 400 nanoseconds.

6.10.2.7 CABLE SKEW DELAY

The maximum difference in propagation time between any two signals measured between any two SCSI devices resident on the same SCSI bus.

CABLE SKEW DELAY is 10 nanoseconds.

6.10.2.8 DATA RELEASE DELAY

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 the DATA BUS.

DATA RELEASE DELAY is 400 nanoseconds.

6.10.2.9 DESKEW DELAY

The minimum amount of time required to deskew certain signals.

DESKEW DELAY is 45 nanoseconds.

6.10.2.10 HOLD TIME

The minimum amount of time that must be added between the assertion of REQ or ACK and the changing of the DATA BUS during synchronous data transfers.

HOLD TIME is 45 nanoseconds.

6.10.2.11 NEGATION PERIOD

The minimum amount of time that REQ (drive) and ACK (initiator) must be negated during synchronous data transfers.

NEGATION PERIOD is 90 nanoseconds.

6.10.2.12 RESET HOLD TIME

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

The maximum amount of time that may elapse between the point at which a SCSI 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 TIMEOUT DELAY

The minimum amount of time that should elapse between the point at which a SCSI device attempts to select another SCSI device and the point at which it begins a timeout procedure, provided the other SCSI device does not respond with BSY.

SELECTION TIMEOUT DELAY should be approximately 250 milliseconds.

6.10.2.15 TRANSFER PERIOD

The minimum amount of time that must elapse between the leading edges of successive REQ and successive ACK pulses during synchronous data transfers.

TRANSFER PERIOD is set during an exchange of SYNCHRONOUS DATA TRANSFER REQUEST messages. (See Section 6.3.6.1, "SYNCHRONOUS DATA TRANSFER REQUEST (01h).")

6.10.3 Fast SCSI Timing Characteristics

Timing characteristics for the Fast SCSI option are given in the following subsections.

When SCSI devices negotiate a synchronous data transfer period of less than 200 ns they are said to be using "fast synchronous data transfers". SCSI devices which negotiate a synchronous data transfer period greater than or equal to 200 ns use timing parameters specified in Section 6.10.2. When a fast synchronous data transfer period is negotiated, those specific times redefined in this section are used; those not redefined remain the same. The minimum synchronous data transfer period is 100 nanoseconds.

6.10.3.1 FAST ASSERTION PERIOD

The minimum amount of time that REQ (drive) and ACK (initiator) must be asserted during fast synchronous data transfers.

FAST ASSERTION PERIOD is 30 nanoseconds.

6.10.3.2 FAST CABLE SKEW DELAY

The maximum difference in propagation time between any two signals measured between any two SCSI devices resident on the same SCSI bus.

FAST CABLE SKEW DELAY is 5 nanoseconds.

6.10.3.3 FAST DESKEW DELAY

The minimum amount of time required to deskew certain signals.

FAST DESKEW DELAY is 20 nanoseconds.

6.10.3.4 FAST HOLD TIME

The minimum amount of time that must be added between the assertion of REQ or ACK and the changing of the DATA BUS during fast synchronous data transfers.

FAST HOLD TIME is 10 nanoseconds.

6.10.3.5 FAST NEGATION PERIOD

The minimum amount of time that REQ (drive) and ACK (initiator) must be negated during fast synchronous data 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.)

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 – Specialized High-speed RAM used to optimize data transfers between system elements with different performance characteristics, e.g., disk to main memory or main memory to CPU.

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.

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 – 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.

LPS – Low Profile Series.

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) – A unit of measurement equal to 1,000 kilobytes, or 1,000,000 bytes. (See also *kilobyte.*)

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. *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.

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 steps 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

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
81-103983-02	Empire 540/1080S Product Manual	

How did you use this publication?	What is your occupation?	Do you wish a written reply?
☐ Reference ☐ Sales ☐ Installation ☐ Maintenance		☐ Yes ☐ No

What is your overall rating of this publication?	Is the material presented effectively?
☐ Very Good ☐ Fair ☐ Good ☐ Poor	☐ 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.