Quantum

Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT Product Manual

October 25, 2000

81-124430-03

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: 81-124430-03

UL/CSA/TUV/CE

UL standard 1950 recognition granted under File No. E78016

CSA standard C22.2 No. 950 certification granted under File No. LR49896

TUV Rheinland EN 60 950 granted under File No. R 9677196

Tested to FCC Rules for Radiated and Conducted Emissions, Part 15, Sub Part J, for Class-B Equipment.

SERVICE CENTERS

Quantum Service Center
160 E. Tasman
San Jose, California 95134
Phone: (408) 894-4000
Fax: (408) 894-3218
http://www.quantum.com

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

Quantum Customer Service Group
Quantum Ireland Ltd.
Finnabair Industrial Park
Dundalk
County Louth, Ireland
Tel: (353) 42-55350
Fax: (353) 45-55355

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.

© 2000 Quantum Corporation. All rights reserved. Printed in U.S.A.

Quantum, the Quantum logo, and AIRLOCK are trademarks of Quantum Corporation, registered in the U.S.A. and other countries. Capacity for the extraordinary, Quantum Fireball lct20, AutoTransfer, AutoRead, AutoWrite, DisCache, DiskWare, Defect Free Interface, and WriteCache are trademarks of Quantum Corporation. All other brand names or trademarks are the property of their manufacturers.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Quantum and its licensors, if any.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

THIS PUBLICATION IS PROVIDED “AS IS’ WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

Table of Content

Chapter 1
ABOUT THIS MANUAL

1.1 AUDIENCE 1-1
1.2 MANUAL ORGANIZATION 1-1
1.3 TERMINOLOGY AND CONVENTIONS 1-1
1.4 REFERENCES 1-3

Chapter 2
GENERAL DESCRIPTION

2.1 PRODUCT OVERVIEW 2-1
2.2 KEY FEATURES 2-1
2.3 Regulatory Compliance Standards 2-3
2.4 HARDWARE REQUIREMENTS 2-3

Chapter 3
INSTALLATION

3.1 SPACE REQUIREMENTS 3-1
3.2 UNPACKING INSTRUCTIONS 3-2
3.3 HARDWARE OPTIONS 3-4
3.3.1 Cable Select (CS) Jumper 3-5
3.3.2 Drive Select (DS) Jumper 3-6
3.3.3 Master Jumper configuration 3-6
3.3.4 Jumper Parking (PK) Position 3-6
3.3.5 Alternate Capacity 3-7
3.4 ATA BUS ADAPTER 3-8
3.4.1 40-Pin ATA Bus Connector 3-8
3.4.2 Adapter Board 3-8
3.5 MOUNTING THE DRIVE ON PERSONAL COMPUTERS 3-9
3.5.1 Orientation 3-9
3.5.2 Clearance 3-11
3.5.3 Ventilation 3-11
3.6 COMBINATION CONNECTOR (J1) 3-11
3.6.1 DC Power (J1, Section A) 3-12
3.6.2 External Drive Activity LED 3-12
3.6.3 ATA Bus Interface Connector (J1, Section C) 3-12
3.7 FOR SYSTEMS WITH A MOTHERBOARD ATA ADAPTER 3-13
3.8 FOR SYSTEMS WITH AN ATA ADAPTER BOARD 3-13
3.8.1 Adapter Board Installation 3-13
3.9 TECHNIQUES IN DRIVE CONFIGURATION 3-16

3.9.1 The 528-Megabytes Barrier 3-16
3.9.2 The 8.4-Gigabytes Barrier 3-16
3.9.3 Operating system limitations 3-17
3.10 SYSTEM STARTUP AND OPERATION 3-17
3.11 CONSIDERATIONS FOR MOUNTING AN AV DRIVE 3-19
3.12 Causes for Noise Generated by Hard Disk Drives 3-20
3.13 Cautions in Reducing the Effects of External Shock Forces 3-20
3.13.1 Cautions in Reducing the Effects of External Vibration Forces 3-20
3.13.2 Acoustical Tolerance Level 3-20
3.14 Mounting Methods to Reduce Noise and Vibration in Set-Top Box Assemblies 3-21
3.14.1 Set-Top Box Enclosure 3-23
3.14.2 Recommended Acoustic and Vibration Damping Material 3-24
3.14.3 Recommended Acoustic and Vibration Damping Material 3-24
3.15 Acoustical Measurements 3-25
3.15.1 Measuring Acoustical Emissions 3-25

Chapter 4
SPECIFICATIONS

4.1 SPECIFICATION SUMMARY 4-1
4.2 FORMATTED CAPACITY 4-3
4.3 DATA TRANSFER RATES 4-3
4.4 TIMING SPECIFICATIONS 4-4
4.5 POWER 4-5
4.5.1 Power Sequencing 4-5
4.5.2 Power Reset Limits 4-5
4.5.3 Power Requirements 4-6
4.6 ACOUSTICS 4-7
4.7 MECHANICAL 4-8
4.8 ENVIRONMENTAL CONDITIONS 4-8
4.9 SHOCK AND VIBRATION 4-9
4.10 HANDLING the DRIVE 4-10
4.11 RELIABILITY 4-10
4.12 ELECTROMAGNETIC SUSCEPTIBILITY 4-11
4.13 SPINDLE IMBALANCE 4-11
4.14 DISK ERRORS 4-11

Chapter 5
BASIC PRINCIPLES OF OPERATION

5.1 QUANTUM FIREBALL lct20 DRIVE MECHANISM 5-1
5.1.1 Base Casting Assembly 5-3
5.1.2 DC Motor Assembly 5-3
5.1.3 Disk Stack Assemblies 5-3
5.1.4 Headstack Assembly 5-4
5.1.5 Rotary Positioner Assembly 5-4
5.1.6 Automatic Actuator Lock 5-4

5.1.7 Air Filtration 5-5
5.2 DRIVE ELECTRONICS 5-5
5.2.1 Integrated μProcessor, Disk Controller and ATA Interface Electronics 5-6
5.2.2 Read/Write ASIC 5-8
5.2.3 PreAmplifier and Write Driver 5-10
5.3 FIRMWARE FEATURES 5-10
5.3.1 Disk Caching 5-10
5.3.2 Head and Cylinder Skewing 5-12
5.3.3 Error Detection and Correction 5-13

Chapter 6
ATA BUS INTERFACE AND ATA COMMANDS

6.1 INTRODUCTION 6-1
6.2 SOFTWARE INTERFACE 6-1
6.3 MECHANICAL DESCRIPTION 6-1
6.3.1 Drive Cable and Connector 6-1
6.4 ELECTRICAL INTERFACE 6-1
6.4.1 ATA Bus Interface 6-1
6.4.2 Host Interface Timing 6-9
6.5 REGISTER ADDRESS DECODING 6-21
6.6 REGISTER DESCRIPTIONS 6-23
6.6.1 Control Block Registers 6-23
6.6.2 Command Block Registers 6-24
6.7 COMMAND DESCRIPTIONS 6-29
6.7.1 Recalibrate 1xh 6-29
6.7.2 Read Sectors 20h 6-29
6.7.3 Write Sector 30h 6-30
6.7.4 Read Verify Sectors 40h 6-30
6.7.5 Seek 7xh 6-30
6.7.6 Execute Drive Diagnostic 90h 6-31
6.7.7 Initialize Drive Parameters 91h 6-31
6.7.8 Download Microcode 6-32
6.7.9 SMART B0h 6-33
6.7.10 Read Multiple C4h 6-47
6.7.11 Write Multiple C5h 6-48
6.7.12 Set Multiple Mode C6h 6-48
6.7.13 Read DMA C8h 6-49
6.7.14 Write DMA CAh 6-49
6.7.15 Read Buffer E4h 6-49
6.7.16 Flush Cache 6-49
6.7.17 Write Buffer E8h 6-50
6.7.18 Power Management Commands 6-51
6.7.19 Identify Drive – ECh 6-53
6.7.20 Set Features EFh 6-60
6.7.21 Set Features (Ultra ATA/66) 6-60
6.7.22 Read Defect List 6-61
6.7.23 Configuration 6-64
6.7.24 Host Protected Mode Feature 6-68
6.8 ERROR REPORTING 6-71

List of Figures

Figure 3-1 Mechanical Dimensions of Quantum Fireball lct20 Hard Disk Drive 3-1
Figure 3-2 Drive Packing Assembly 3-2
Figure 3-3 Drive Packing Assembly of a Polypropylene 20-Pack Container 3-3
Figure 3-4 Jumper Locations for the Quantum Fireball lct20 Hard Disk Drive 3-4
Figure 3-5 Jumper Locations on the Interface Connector 3-4
Figure 3-6 AT Connector and Jumper Location 3-7
Figure 3-7 Mounting Dimensions for the Quantum Fireball lct20 Hard Disk Drives 3-9
Figure 3-8 Mounting Screw Clearance for the Quantum Fireball lct20 Hard Disk Drives 3-10
Figure 3-9 J1 DC Power and ATA Bus Combination Connector 3-11
Figure 3-10 Drive Power Supply and ATA Bus Interface Cables 3-14
Figure 3-11 Completing the Drive Installation 3-15
Figure 3-12 Drive Mounted in a Set-Top Box 3-21
Figure 3-13 Set-Top Enclosure (Container Box) Assembly 3-22
Figure 3-14 Mounting an AV Drive 3-23
Figure 3-15 Measuring Acoustical Emissions 3-25
Figure 5-1 Quantum Fireball lct20 AT Hard Disk Drive Exploded View 5-2
Figure 5-2 Quantum Fireball lct20 AT Hard Disk Drive Block Diagram 5-5
Figure 5-3 Block Diagram 5-6
Figure 5-4 Sector Data Field with ECC Check Bytes 5-14
Figure 5-5 Byte Interleaving 5-14
Figure 5-6 Correctable and Uncorrectable Double-Burst Errors 5-15
Figure 5-7 Correctable and Uncorrectable Quadruple-Burst Errors 5-16
Figure 5-8 Twelve Correctable Random Burst Errors 5-17
Figure 6-1 PIO Interface Timing 6-10
Figure 6-2 Multiword DMA Bus Interface Timing 6-11
Figure 6-3 Initiating a Data In Burst 6-15
Figure 6-4 Sustained Data In Burst 6-15
Figure 6-5 Host Pausing a Data In Burst 6-16
Figure 6-6 Device Terminating a Data In Burst 6-16
Figure 6-7 Host Terminating a Data In Burst 6-17
Figure 6-8 Initiating a Data Out Burst 6-18
Figure 6-9 Sustained Data Out Burst 6-18
Figure 6-10 Device Pausing a Data Out Burst 6-18
Figure 6-11 Host Terminating a Data Out Burst 6-19
Figure 6-12 Device Terminating a Data out Burst 6-20
Figure 6-13 Host Interface RESET Timing 6-21

List of Tables

Table 3-1 AT Jumper Options 3-5
Table 3-2 J1 Power Connector, Section A 3-12
Table 3-3 Logical Addressing Format 3-18
Table 4-1 Specifications 4-1
Table 4-2 Formatted Capacity 4-3
Table 4-3 Timing Specifications 4-4
Table 4-4 Power Reset Limits 4-5
Table 4-5 Typical Power and Current Consumption 4-6
Table 4-6 Acoustical Characteristics—Sound Power 4-7
Table 4-7 Environmental Specifications 4-8
Table 4-8 Shock and Vibration Specifications 4-9
Table 4-9 Error Rates 4-11
Table 5-1 Cylinder Contents 5-3
Table 5-2 Skew Offsets 5-13
Table 6-1 Drive Connector Pin Assignments (J1, Section C) 6-2
Table 6-2 Series Termination for Ultra ATA/66 6-6
Table 6-3 Signal Line Definitions 6-8
Table 6-4 Interface Signal Name Assignments 6-8
Table 6-5 PIO Host Interface Timing 6-10
Table 6-6 Multiword DMA Host Interface Timing 6-11
Table 6-7 Ultra DMA Data Transfer Timing Requirements 6-12
Table 6-8 Ultra DMA Data Burst Timing Descriptions 6-13
Table 6-9 Host Interface RESET Timing 6-21
Table 6-10 I/O Port Functions and Selection Addresses 6-22
Table 6-11 Command Block Register Initial Values 6-22
Table 6-12 Device Control Register Bits 6-23
Table 6-13 Drive Address Register Bits 6-24
Table 6-14 Error Register Bits 6-25
Table 6-15 Drive Head Register Bits 6-26
Table 6-16 Status Register Bits 6-27
Table 6-17 Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT Command Codes and Parameters 6-28
Table 6-18 Diagnostic Codes 6-31
Table 6-19 Device attribute thresholds data structure 6-38
Table 6-20 Individual threshold data structure 6-38
Table 6-21 Device attributes data structure 6-40
Table 6-22 Individual attribute data structure 6-41
Table 6-23 Valid Count Range 6-52
Table 6-24 Identify Drive Parameters 6-54
Table 6-25 Transfer/Mode Values 6-60
Table 6-26 READ DEFECT LIST LENGTH Command Bytes 6-61
Table 6-27 AT READ DEFECT LIST Command Bytes 6-62

Table 6-28 DEFECT LIST DATA FORMAT 6-63
Table 6-29 DEFECT ENTRY DATA FORMAT 6-63
Table 6-30 Accessing the READ CONFIGURATION Command 6-64
Table 6-31 Accessing the SET CONFIGURATION Command 6-65
Table 6-32 Accessing the SET CONFIGURATION WITHOUT SAVING TO DISK Command 6-66
Table 6-33 Configuration Command Format 6-67
Table 6-34 Command Errors 6-71

Chapter 1
ABOUT THIS MANUAL

This chapter gives an overview of the contents of this manual, including the intended audience, how the manual is organized, terminology and conventions, and references.

1.1 AUDIENCE

The Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT Product Manual is intended for several audiences. These audiences include: the end user, installer, developer, consumer electronics and personal computer original equipment manufacturer (CE/PC OEM), and distributor. The manual provides information about installation, principles of operation, interface command implementation, and maintenance.

1.2 MANUAL ORGANIZATION

This manual 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 – *ATA Bus Interface and ATA Commands*

1.3 TERMINOLOGY AND CONVENTIONS

In the Glossary at the back of this manual, you can find definitions for many of the terms used in this manual. In addition, the following abbreviations are used in this manual:

- ASIC application-specific integrated circuit
- ATA advanced technology attachment
- bpi bits per inch
- dB decibels
- dBA decibels, A weighted

- DPS Data Protection System
- SPS Shock Protection System
- ECC error correcting code
- fci flux changes per inch
- Hz hertz
- KB kilobytes
- LSB least significant bit
- mA milliamperes
- MB megabytes (1 MB = 1,000,000 bytes when referring to disk storage and 1,048,576 bytes in all other cases)
- 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 typographical and naming conventions used in this manual are listed below. Conventions that are unique to a specific table appear in the notes that follow that table.

Typographical Conventions:

- **Names of Bits:** Bit names are presented in initial capitals. An example is the Host Software Reset bit.
- **Commands:** Interface commands are listed in all capitals. An example is WRITE LONG.
- **Register Names:** Registers are given in this manual with initial capitals. An example is the Alternate Status Register.
- **Parameters:** Parameters are given as initial capitals when spelled out, and are given as all capitals when abbreviated. Examples are Prefetch Enable (PE), and Cache Enable (CE).
- **Hexadecimal Notation:** The hexadecimal notation is given in 9-point subscript form. An example is 30_{H}.
- **Signal Negation:** A signal name that is defined as active low is listed with a minus sign following the signal. An example is RD–.

- **Messages:** A message that is sent from the drive to the host is listed in all capitals. An example is ILLEGAL COMMAND.

Naming Conventions:

- **Host:** In general, the system in which the drive resides is referred to as the host.
- **Computer Voice:** This refers to items you type at the computer keyboard. These items are listed in 10-point, all capitals, Courier font. An example is `FORMAT C:/S.`

1.4 REFERENCES

For additional information about the AT interface, refer to:

- IBM Technical Reference Manual #6183355, March 1986.
- ATA Common Access Method Specification, Revision 5.0.

Chapter 2
GENERAL DESCRIPTION

This chapter summarizes the general functions and key features of the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives, as well as the applicable standards and regulations.

2.1 PRODUCT OVERVIEW

Quantum's Fireball lct20 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.

These hard disk drives use nonremovable, 3 1/2-inch hard disks and are available with the ATA interface.

The Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives feature an embedded hard disk drive controller, and use ATA commands to optimize system performance. Because the drive manages media defects and error recovery internally, these operations are fully transparent to the user.

The innovative design of the Quantum Fireball lct20 hard disk drives incorporate leading edge technologies such as UltraATA/100, Advanced Cache Management, Shock Protection System™ II (SPS II), and Data Protection System (DPS). These enhanced technologies enable Quantum to produce a family of high-performance, high-reliability drives.

Additionally, the Quantum Fireball lct20 family of drives are CE-class drives designed for the consumer electronics market. Acoustics is one of the most important features in the CE arena. These hard disk drives are available with very quiet seek codes resulting in very low acoustical emissions from the drive. This makes the Quantum Fireball lct20 drives an attractive choice for most CE users.

2.2 KEY FEATURES

The Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives include the following key features:

General

- Formatted storage capacity of 10.0 GB (1 disk, 1 heads), 20.0 GB (1 disks, 2 heads), 30.0 (2 disks, 3 heads) and 40.0 (2 disks, 4 heads)
- Low profile, 1-inch height
- Industry standard 3 1/2-inch form factor
- Emulation of IBM® PC AT® task file register, and all AT fixed disk commands
- Windows NT and 9X Certification

Performance

- Average seek time of 12 ms
- Average rotational latency of 6.8 ms
- New Ultra ATA interface with Quantum-patented UltraATA/100 protocol supporting burst data transfer rates of 100 MB/s.
- Extra quiet seek code available (Quantum Quiet Drive Technology)
- 128 K buffer with 64 K (approximately) Advance Cache Management (ACM).
- Look-ahead DisCache feature with continuous prefetch and WriteCache write-buffering capabilities
- Read promotion and Write Reordering features
- AutoTask Register update, Multi-block AutoRead, and Multi-block AutoWrite features in a custom ASIC
- Read-on-arrival firmware
- Quadruple-burst ECC, and double burst ECC on-the-fly
- 1:1 interleave on read/write operations
- Support of all standard ATA data transfer modes with PIO mode 4 and multiword DMA mode 2, and Ultra DMA modes 0, 1, 2, 3, 4 and 5
- Adaptive cache segmentation

Reliability

- 625,000 hours mean time between failure (MTBF) in the field
- Automatic retry on read errors
- 288-bit, interleaved Reed-Solomon Error Correcting Code (ECC), with cross checking correction up to four separate bursts of 32 bits each totalling up to 128 bits in length
- S.M.A.R.T. 4 (Self-Monitoring, Analysis and Reporting Technology)
- Patented Airlock® automatic shipping lock, magnetic actuator retract, and dedicated landing zone
- Transparent media defect mapping
- High performance, in-line defective sector skipping
- Reassignment of defective sectors discovered in the field, without reformatting
- Shock Protection System II to enhance protection against both operating and non-operating shock
- Data Protection System to verify drive integrity

Versatility

- Power saving modes
- Downloadable firmware
- Cable select feature
- Ability to daisy-chain two drives on the interface

2.3 Regulatory Compliance Standards

Quantum Corporation's disk drive products meet all domestic and international product safety regulatory compliance requirements. Quantum's disk drive products conform to the following specifically marked Product Safety Standards:

- Underwriters Laboratories (UL) Standard 1950. This certificate is a category certification pertaining to all 3.5-inch drive models in the U.S.A. and Canada.
- TUV Rheinland Standard EN60 950. This certificate is a category certification pertaining to all 3.5-inch drive models.

Product EMI/EMS Qualifications:

- CE Mark authorization is granted by TUV Rheinland in compliance with our qualifying under EN 55022:1994 and EN 50082-1:1997.
- C-Tick Mark is an Australian authorization marked noted on Quantum's disk drive products. The mark proves conformity to the regulatory compliance document AS/NZS 3548: 1995 and BS EN 55022: 1995.
- Approval from Taiwan BSMI. Number: 3892A638

Quantum's disk drives are designed as a separate subassembly that conforms to the FCC Rules for Radiated and Conducted emissions, Part 15 Subpart J; Class B when installed in a given computer system.

2.4 HARDWARE REQUIREMENTS

The Quantum Fireball lct20 hard disk drives are compatible with the IBM PC AT, and other computers that are compatible with the IBM PC AT. It connects to the PC either by means of a third-party IDE-compatible adapter board, or by plugging a cable from the drive directly into a PC motherboard that supplies an ATA interface.

Chapter 3
INSTALLATION

This chapter explains how to unpack, configure, mount, and connect the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drive prior to operation. It also explains how to start up and operate the drive.

3.1 SPACE REQUIREMENTS

The Quantum Fireball lct20 hard disk drives are shipped without a faceplate. Figure 3-1 shows the external dimensions of the Quantum Fireball lct20 10.0/ 20.0/30.0/40.0 GB AT drives.

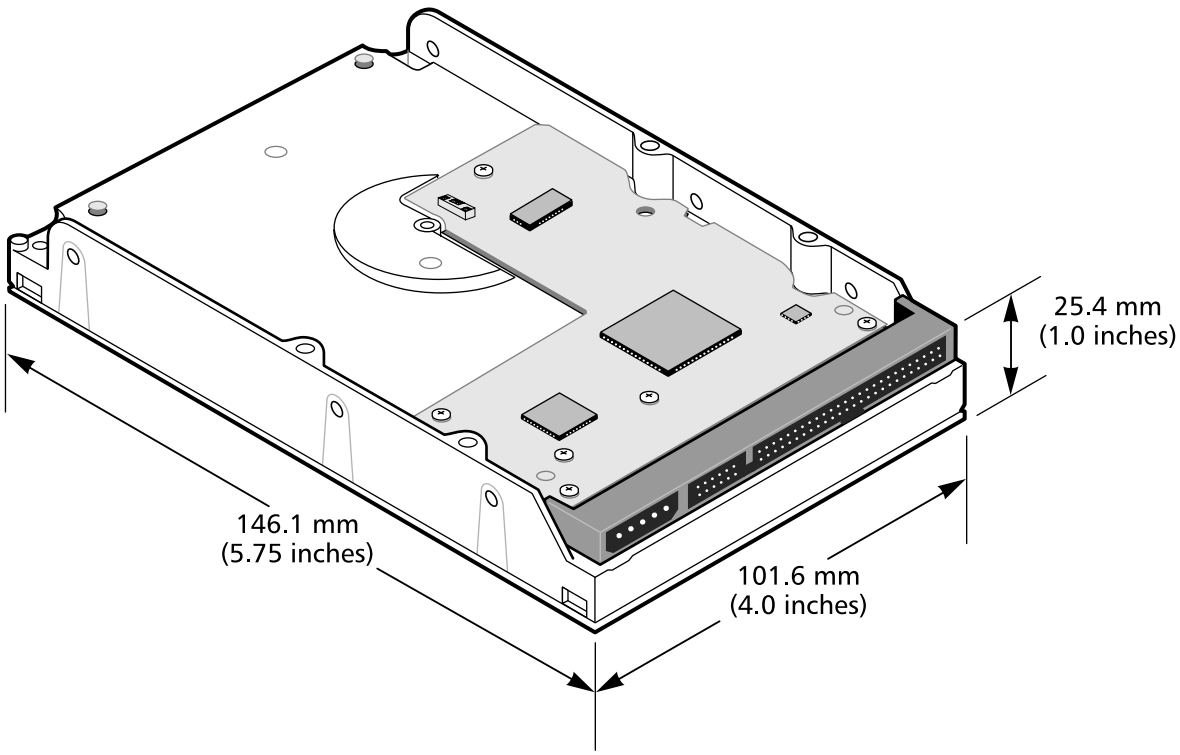

Figure 3-1 *Mechanical Dimensions of Quantum Fireball lct20 Hard Disk Drive*

3.2 UNPACKING INSTRUCTIONS

CAUTION: The maximum limits for physical shock can be exceeded if the drive is not handled properly. Special care should be taken not to bump or drop the drive. It is highly recommended that Quantum Fireball lct20 drives are not stacked or placed on any hard surface after they are unpacked. Such handling could cause media damage.

1. Open the shipping container and remove the packing assembly that contains the drive.
2. Remove the drive from the packing assembly.

CAUTION: During shipment and handling, the antistatic electrostatic discharge (ESD) bag prevents electronic component damage due to electrostatic discharge. To avoid accidental damage to the drive, do not use a sharp instrument to open the ESD bag and do not touch PCB components. Save the packing materials for possible future use.

3. When you are ready to install the drive, remove it from the ESD bag.

Figure 3-2 shows the packing assembly for a single Quantum Fireball lct20 hard disk drive. A 20-pack shipping container is available for multiple drive shipments.

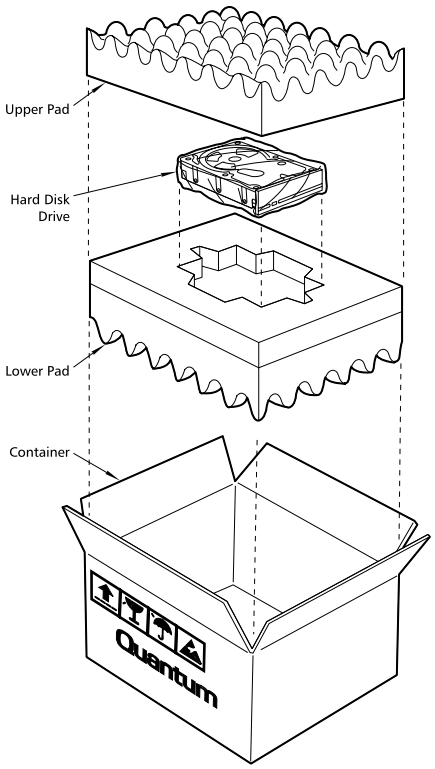

Figure 3-2 *Drive Packing Assembly*

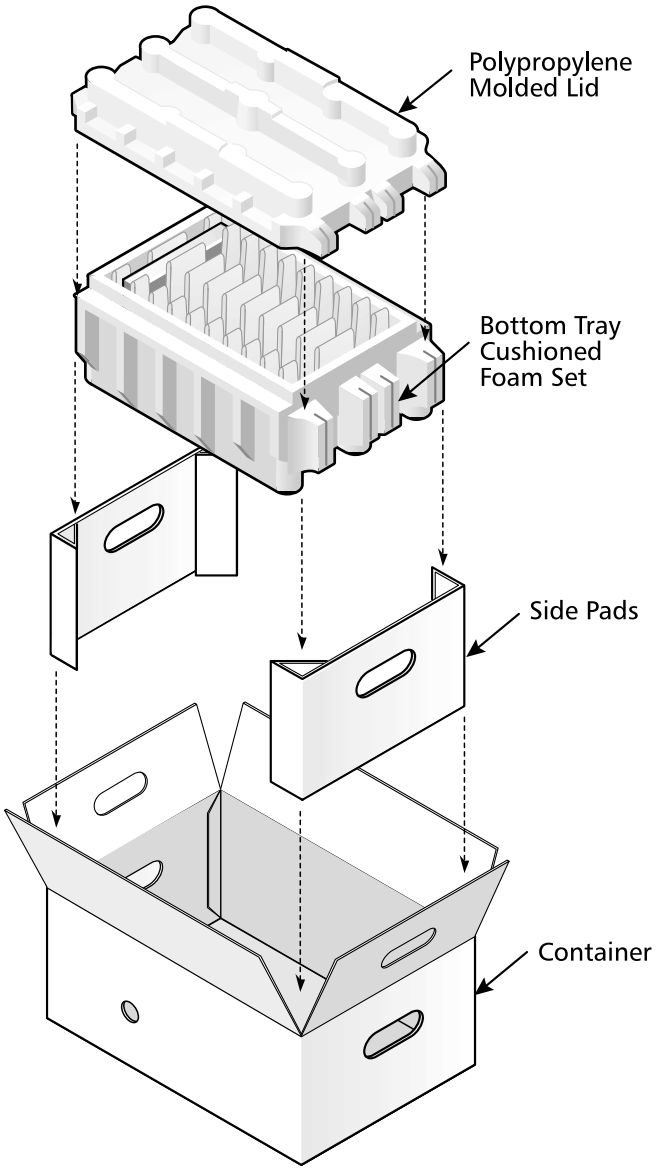

Figure 3-3 *Drive Packing Assembly of a Polypropylene 20-Pack Container*

Note: The 20-pack container should be shipped in the same way it was received from Quantum. When individual drives are shipped from the 20-pack container then it should be appropriately packaged (not supplied with the 20-pack) to prevent damage.

3.3 HARDWARE OPTIONS

The configuration of a Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drive depends on the host system in which it is to be installed. This section describes the hardware options that you must take into account prior to installation. Figure 3-4 shows the printed circuit board (PCB) assembly, indicating the jumpers that control some of these options.

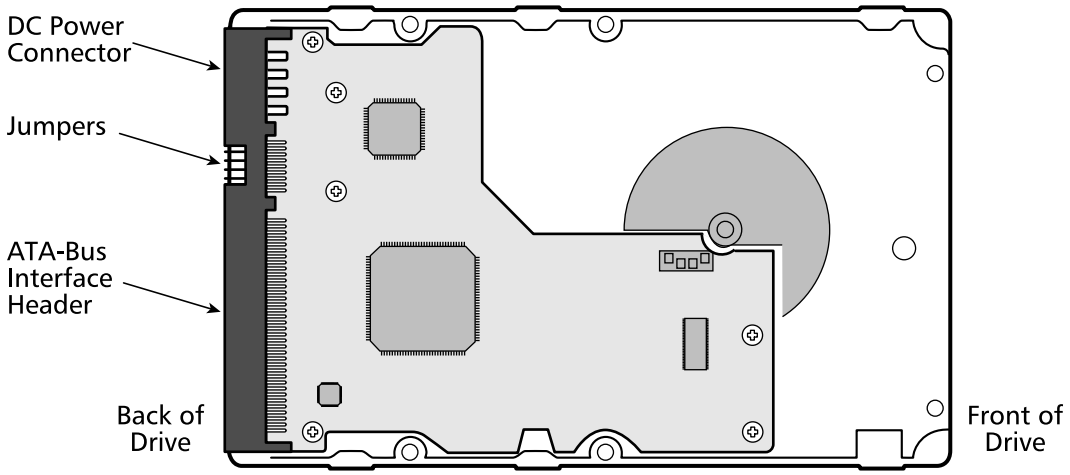

Figure 3-4 *Jumper Locations for the Quantum Fireball lct20 Hard Disk Drive*

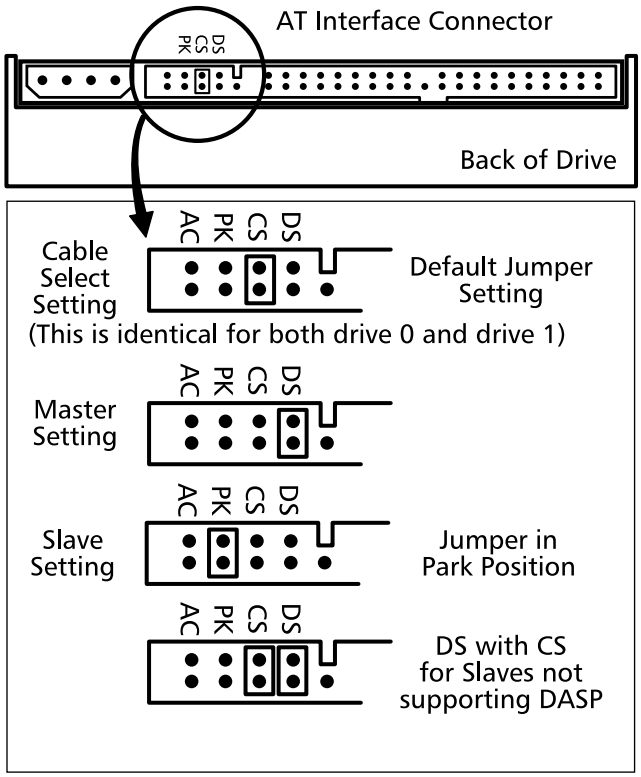

Figure 3-5 *Jumper Locations on the Interface Connector*

The configuration of the following four jumpers controls the drive's modes of operation:

- CS – Cable Select
- DS – Drive Select
- PK– Jumper Parking Position (Slave mode)
- AC– Alternate Capacity

The AT PCB has two jumper locations provided to configure the drive in a system. The default configuration for the drive as shipped from the factory is with a jumper across the CS location, and open positions in the DS, PK and AC positions.

Table 3-1 defines the operation of the master-slave jumpers and their function relative to pin 28 on the interface. *1* indicates that the specified jumper is installed; *0* indicates that the jumper is not installed.

Table 3-1 *AT Jumper Options*

CS	DS	PK	PIN 28	DESCRIPTION
0	0	X	X	Drive is configured as a slave
1	0	X	Gnd	Drive is configured as Master (Device 0) when attached to the end of a 80 conductor Ultra ATA cable
0	1	X	X	Drive is configured as a Master
1	0	X	Open	Drive is configured as a Slave (Device 1) when attached to the middle of a 80 conductor Ultra ATA cable
1	1	X	X	Drive is configured as a Master with an attached slave that does not support DASP

Note: In Table 3-1, a *0* indicates that the jumper is removed, a *1* indicates that the jumper is installed, and an X indicates that the jumper setting does not matter.

3.3.1 Cable Select (CS) Jumper

When a Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drive and another ATA hard disk drive are daisy-chained together, they can be configured as Master or Slave either by the CS or DS jumpers. To configure the drive as a Master or Slave with the CS feature, the CS jumper is installed (1). The drive's position on the 80 conductor Ultra ATA data cable then determines whether the drive is a Master (Device 0) or a Slave (Device 1). If the drive is connected to the end of the Ultra (cable Select) data cable the drive is a Master. If the drive is connected to the middle connection it is set as a Slave.

Once you install the CS jumper, the drive is configured as a Master or Slave by the state of the Cable Select signal: pin 28 of the ATA bus connector. Please note that pin 28 is a vendor-specific pin that Quantum is using for a specific purpose. More than one function is allocated to CS, according to the ATA CAM specification (see reference to this specification in Chapter 1). If pin 28 is a *0* (grounded), the

drive is configured as a Master. If it is a *1* (high), the drive is configured as a Slave. In order to configure two drives in a Master/Slave relationship using the CS jumper, you need to use a cable that provides the proper signal level at pin 28 of the ATA bus connector. This allows two drives to operate in a Master/Slave relationship according to the drive cable placement.

The Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives are shipped from the factory as a Master (Device 0 - CS jumper installed). To configure a drive as a Slave (Device 1- DS scheme), the CS jumper must be removed. In this configuration, the spare jumper removed from the CS position may be stored on the PK jumper pins.

3.3.2 Drive Select (DS) Jumper

You can also daisy-chain two drives on the ATA bus interface by using their Drive Select (DS) jumpers. To use the DS feature, the CS jumper must not be installed.

To configure a drive as the Master (Device 0), a jumper must be installed on the DS pins.

Note: The order in which drives are connected in a daisy chain has no significance.

3.3.3 Master Jumper configuration

In combination with the current DS or CS jumper settings, the Slave Present (SP) jumper can be implemented if necessary as follows:

Note: The CS position doubles as the Slave present on this drive.

- When the drive is configured as a Master (DS jumper installed **or** CS jumper installed, and the Cable Select signal is set to (*0*), adding an additional jumper (both jumpers DS and CS now installed) will indicate to the drive that a Slave drive is present. This Master with Slave Present jumper configuration should be installed on the Master drive only if the Slave drive does *not* use the Drive Active/Slave Present (DASP–) signal to indicate its presence.

3.3.4 Jumper Parking (PK) Position

The PK position is used as a holding place for the jumper for a slave drive in systems that do not support Cable Select. The pins used for the parking position are vendor unique.

3.3.5 Alternate Capacity (AC)

For user capacities below 66,055,248 sectors (32 GB), inserting the AC jumper limits the Number of Cylinders field 1 to a value of 4,092, as reported in IDENTIFY DEVICE data word. This allows software drivers to determine that the actual capacity is larger than indicated by the maximum CHS, requiring LBA addressing to use the full capacity.

For larger user capacities above 66,055,248 sectors (32 GB), inserting the AC jumper limits the Total Number of User Addressable Sectors field to a value of 66,055,248, as reported in IDENTIFY DEVICE data words 60 and 61. This protects hosts which are not able to address beyond this capacity.

A summary of these effects for the Quantum Fireball lct20 drives is shown in the following table:

	AC JUMPER OUT	AC JUMPER IN
10 GB	C=16,383 H=16 S=63 LBA=20,044,080	C=4,092 H=16 S=63 LBA=20,044,080
20 GB	C=16,383 H=16 S=63 LBA=39,876,480	C=4,092 H=16 S=63 LBA=39,876,480
30 GB	C=16,383 H=16 S=63 LBA=58,633,344	C=4,092 H=16 S=63 LBA=58,633,344
40 GB	C=16,383 H=16 S=63 LBA=78,177,792	C=4,092 H=16 S=63 LBA= 66,055,248

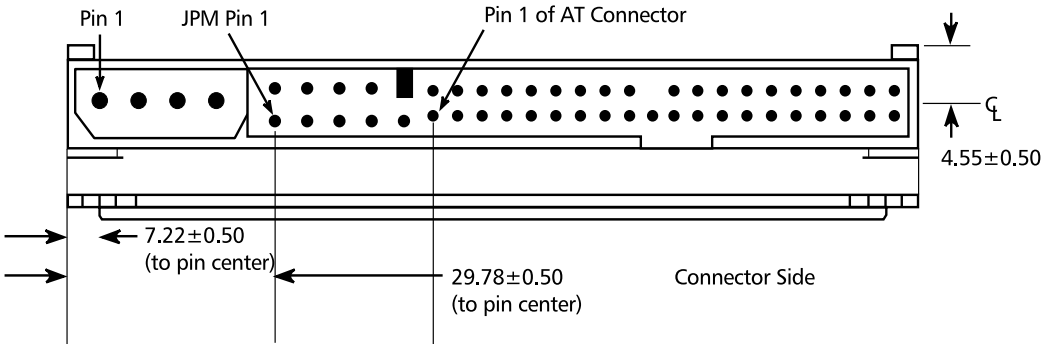

Figure 3-6 *AT Connector and Jumper Location*

3.4 ATA BUS ADAPTER

There are two ways you can configure a system to allow the Quantum Fireball lct20 hard disk drives to communicate over the ATA bus of an IBM or IBM-compatible PC:

1. Connect the drive to a 40-pin ATA bus connector (if available) on the motherboard of the PC.
2. Install an IDE-compatible adapter board in the PC, and connect the drive to the adapter board.

3.4.1 40-Pin ATA Bus Connector

Most PC motherboards have a built-in 40-pin ATA bus connector that is compatible with the 40-pin ATA interface of the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives. If the motherboard has an ATA connector, simply connect a 40-pin ribbon cable between the drive and the motherboard.

You should also refer to the motherboard instruction manual, and refer to Chapter 6 of this manual to ensure signal compatibility.

3.4.2 Adapter Board

If your PC motherboard does not contain a built-in 40-pin ATA bus interface connector, you must install an ATA bus adapter board and connecting cable to allow the drive to interface with the motherboard. Quantum does not supply such an adapter board, but they are available from several third-party vendors.

Please carefully read the instruction manual that comes with your adapter board, as well as Chapter 6 of this manual to ensure signal compatibility between the adapter board and the drive. Also, make sure that the adapter board jumper settings are appropriate.

3.5 MOUNTING THE DRIVE ON PERSONAL COMPUTERS

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

3.5.1 Orientation

The mounting holes on the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives allow the drive to be mounted in any orientation. Figure 3-6 and Figure 3-7 show the location of the three mounting holes on each side of the drive. The drive can also be mounted using the four mounting hole locations on the PCB side of the drive.

Note: It is highly recommended that the drive is hard mounted on to the chassis of the system being used for general operation, as well as for test purposes. Failure to hard mount the drive can result in erroneous errors during testing.

Drives can be mounted in any orientation. Normal position is with the PCB facing down.

All dimensions are in millimeters. For mounting, #6-32 UNC screws are recommended.

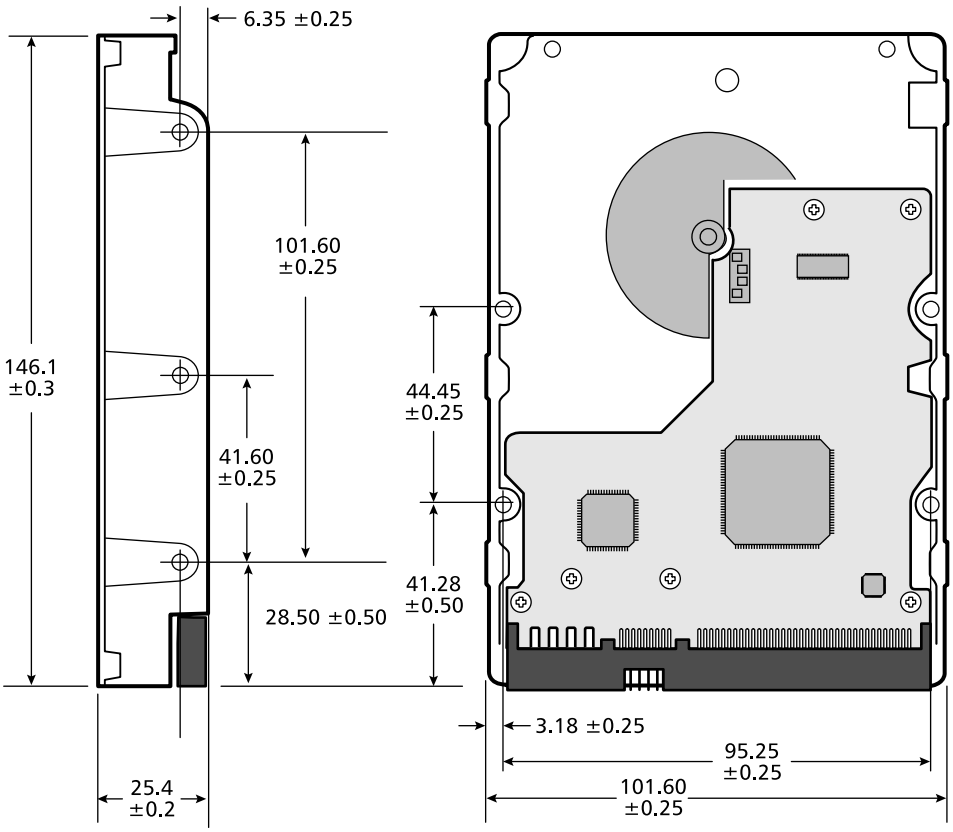

Figure 3-7 *Mounting Dimensions for the Quantum Fireball lct20 Hard Disk Drives*

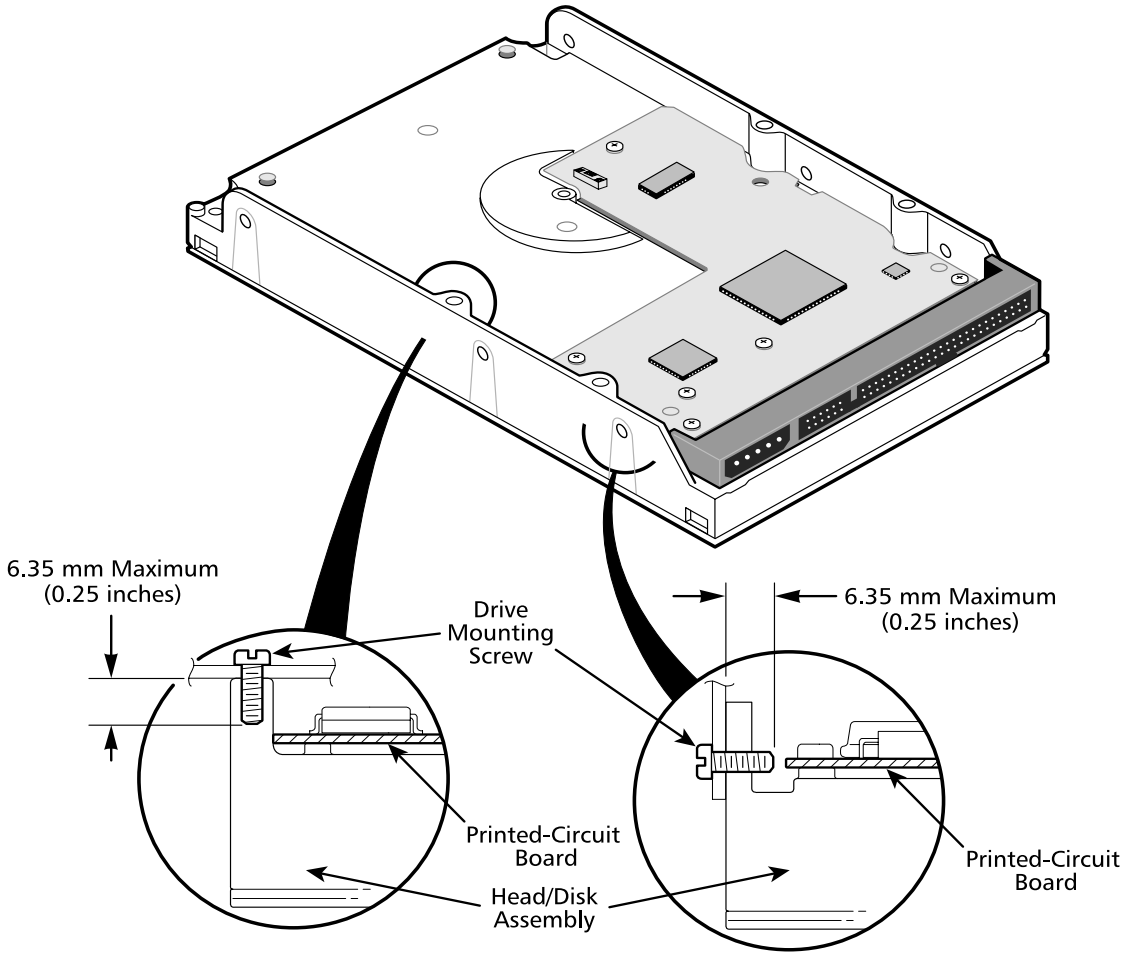

Figure 3-8 *Mounting Screw Clearance for the Quantum Fireball lct20 Hard Disk Drives*

CAUTION: The PCB is very close to the mounting holes. Do not exceed the specified length for the mounting screws. The specified screw length allows full use of the mounting hole threads, while avoiding damaging or placing unwanted stress on the PCB. Figure 3-8 specifies the minimum clearance between the PCB and the screws in the mounting holes. To avoid stripping the mounting hole threads, the maximum torque applied to the screws must not exceed 8 inch-pounds. A maximum screw length of 0.25 inches may be used.

3.5.2 Clearance

Clearance from the drive to any other surface (except mounting surfaces) must be a minimum of 1.25 mm (0.05 inches).

3.5.3 Ventilation

The Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives operate without a cooling fan, provided the ambient air temperature does not exceed 131°F (55°C) at any point along the drive form factor envelope.

3.6 COMBINATION CONNECTOR (J1)

J1 is a three-in-one combination connector. The drive's DC power can be applied to section A. The ATA bus interface (40-pin) uses section C. The connector is mounted on the back edge of the printed-circuit board (PCB), as shown in Figure 3-9.

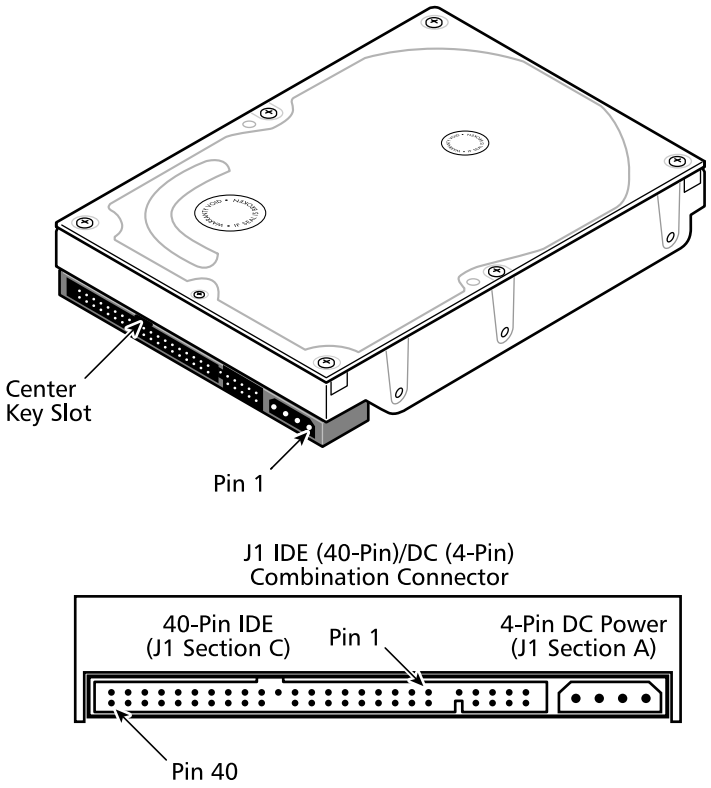

Figure 3-9 *J1 DC Power and ATA Bus Combination Connector*

3.6.1 DC Power (J1, Section A)

The recommended mating connectors for the +5 VDC and +12 VDC input power are listed in Table 3-2.

Table 3-2 *J1 Power Connector, Section A*

PIN NUMBER	VOLTAGE LEVEL	MATING CONNECTOR TYPE AND PART NUMBER (OR EQUIVALENT)
J1 Section A (4-Pin):		
1	+12 VDC	4-Pin Connector: AMP P/N 1-480424-0 Loose piece contacts: AMP P/N VS 60619-4 Strip contacts: AMP P/N VS 61117-4
2	Ground Return for +12 VDC	
3	Ground Return for +5 VDC	
4	+5 VDC	

Note: Labels indicate the pin numbers on the connector. Pins 2 and 3 of section A are the +5 and +12 volt returns and are connected together on the drive.

3.6.2 External Drive Activity LED

An external drive activity LED may be connected to the DASP-I/O pin 39 on J1. For more details, see the pin description in Table 6-1.

3.6.3 ATA Bus Interface Connector (J1, Section C)

On the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives, the ATA bus interface cable connector (J1, section C) is a 40-pin Universal Header, as shown in Figure 3-9.

To prevent the possibility of incorrect installation, the connector has been keyed by removing Pin 20. This ensures that a connector cannot be installed upside down.

See Chapter 6, “ATA Bus Interface and ATA Commands,” for more detailed information about the required signals. Refer to Table 6-1 for the pin assignments of the ATA bus connector (J1, section C).

3.7 FOR SYSTEMS WITH A MOTHERBOARD ATA ADAPTER

You can install the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives in an AT-compatible system that contains a 40-pin ATA bus connector on the motherboard.

To connect the drive to the motherboard, use a 40 conductor ribbon cable (80 conductor ribbon cable if using UltraATA/66 drive) 18 inches in length or shorter. Ensure that pin 1 of the drive is connected to pin 1 of the motherboard connector.

3.8 FOR SYSTEMS WITH AN ATA ADAPTER BOARD

To install the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drive in an AT-compatible system without a 40-pin ATA bus connector on its motherboard, you need a third-party IDE-compatible adapter board.

3.8.1 Adapter Board Installation

Carefully read the manual that accompanies your adapter board before installing it. Make sure that all the jumpers are set properly and that there are no address or signal conflicts. You must also investigate to see if your AT-compatible system contains a combination floppy and hard disk controller board. If it does, you must disable the hard disk drive controller functions on that controller board before proceeding.

Once you have disabled the hard disk drive controller functions on the floppy/ hard drive controller, install the adapter board. Again, make sure that you have set all jumper straps on the adapter board to avoid addressing and signal conflicts.

Note: For Sections 3.7 and 3.8, power should be turned off on the computer before installing the drive.

3.8.1.1 Connecting the Adapter Board and the Drive

Use a 40-pin ribbon cable to connect the drive to the board. See Figure 3-10. To connect the drive to the board:

1. Insert the 40-pin cable connector into the mating connector of the adapter board. Make sure that pin 1 of the connector matches with pin 1 on the cable.
2. Insert the other end of the cable into the header on the drive. When inserting this end of the cable, make sure that pin 1 of the cable connects to pin 1 of the drive connector.
3. Secure the drive to the system chassis by using the mounting screws, as shown in Figure 3-11.

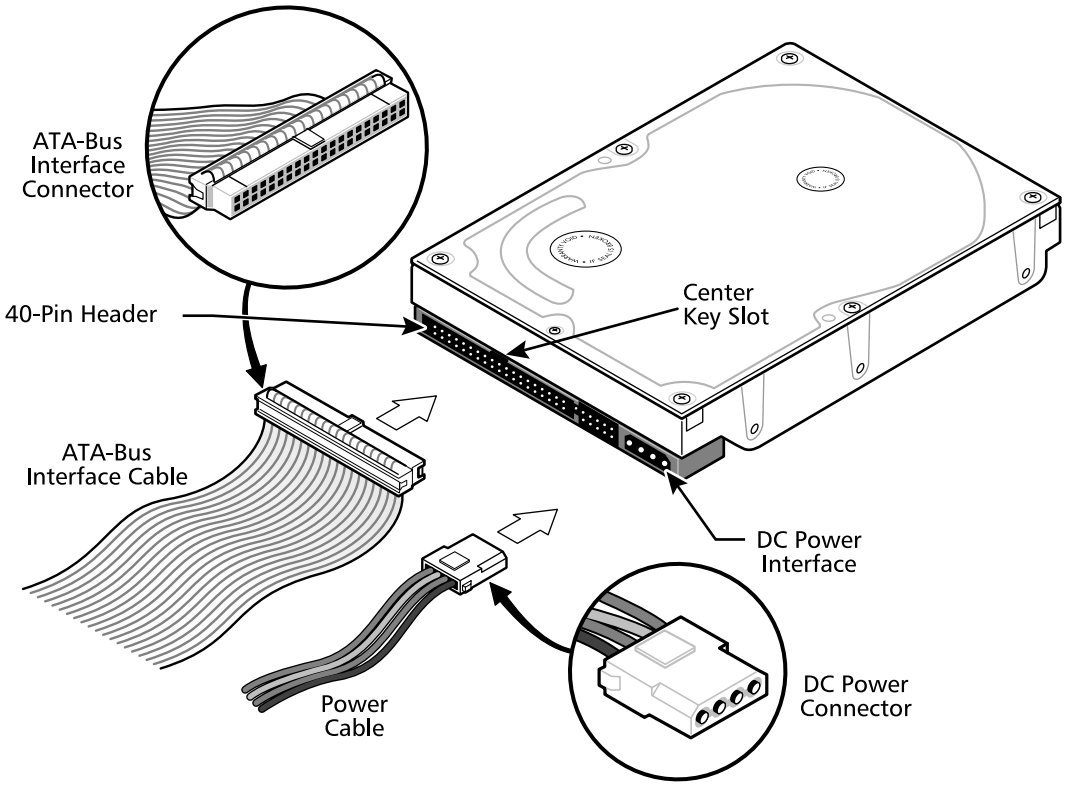

Figure 3-10 *Drive Power Supply and ATA Bus Interface Cables*

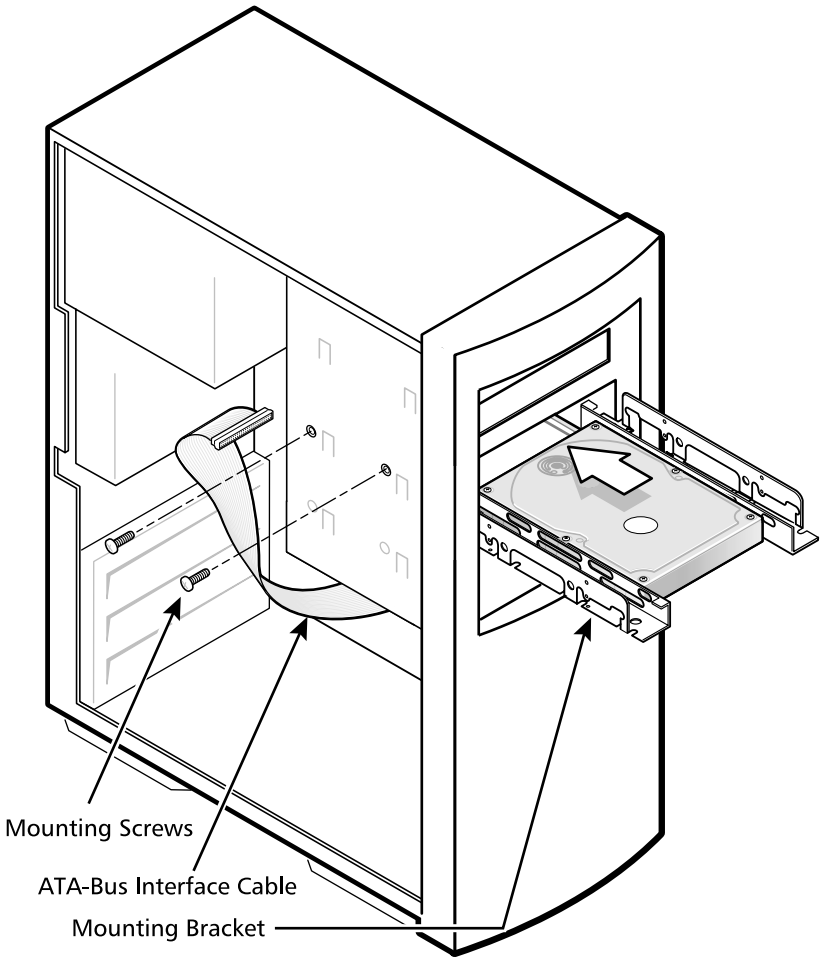

Figure 3-11 *Completing the Drive Installation*

3.9 TECHNIQUES IN DRIVE CONFIGURATION

3.9.1 The 528-Megabytes Barrier

Older BIOS that only support Int 13 commands for accessing ATA drives through DOS based operating systems will be limited to use only 1024 cylinders. This will reduce the effective capacity of the drive to 528 Mbytes.

Whenever possible the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT drive should be used on systems that support LBA translation to ensure the use of the entire capacity of the disk drive. If that is not possible the following are some techniques that can be used to overcome this barrier.

- Use a third party software program that translates the hard drive parameters to an acceptable configuration for MS-DOS.
- Use a hard disk controller that translates the hard drive parameters to an appropriate setup for both MS-DOS and the computer system's ROM-BIOS.
- Insert the Alternate Capacity jumper on the drive (see Section 3.3.5).

3.9.2 The 8.4-Gigabytes Barrier

Newer BIOS's allow users to configure disk drives to go beyond the 528 MB barrier by using several BIOS translation schemes. However, while using these translations the BIOS using Int 13 functions are limited to 24 bits of addressing which results in another barrier at the 8.4 GB capacity.

To overcome this barrier a new set of Int 13 extensions are being implemented by most BIOS manufacturers. The new Int 13 extension allows for four words of addressing space (64 bits) resulting in 9.4 terabytes of accessible space.

Whenever possible the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT drive should be used on systems with BIOS that support Int 13 extensions. If that is not possible the following are some techniques that can be used to overcome this barrier:

- Use a third party software that supplements the BIOS and adds Int 13 extension support.
- Obtain a BIOS upgrade from the system board manufacturer. Many system board manufacturers allow their BIOS to be upgraded in the field using special download utilities. Information on BIOS upgrades can be obtained on the System Board Customer Service respective web sites on the Internet.
- Insert the Alternate Capacity jumper on the drive (see Section 3.3.5).

3.9.3 Operating system limitations

Most popular operating systems available today have additional limitations which affect the use of large capacity drives. However, these limitations can not be corrected on the BIOS and it is up to the operating system manufacturers to release improved versions to address these problems.

The most popular operating systems available today, DOS and Win 95, use a File Allocation Table (FAT) size of 16 bits which will only support partitions up to 2.1 GB. A newer release of Win 95 called OSR2 with a 32 bit FAT has been released to system manufacturers only. This new FAT size table will support partitions of up to 2.2 terabytes.

3.10 SYSTEM STARTUP AND OPERATION

Once you have installed the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drive, and adapter board (if required) in the host system, you are ready to partition and format the drive for operation. To set up the drive correctly, follow these steps:

1. Power on the system.
2. Run the SETUP program. This is generally on a Diagnostics or Utilities disk, or within the system's BIOS. Some system BIOS have an auto-detecting feature making SETUP unnecessary.
3. Enter the appropriate parameters.

The SETUP program allows you to enter the types of optional hardware installed—such as the hard disk drive type, the floppy disk drive capacity, and the display adapter type. The system's BIOS uses this information to initialize the system when the power is switched on. For instructions on how to use the SETUP program, refer to the system manual for your PC.

During the AT system CMOS setup, you must enter the drive type for the Quantum Fireball lct20 hard disk drives. The drive supports the translation of its physical drive geometry parameters such as cylinders, heads, and sectors per track to a logical addressing mode. The drive can work with different BIOS drive-type tables of the various host systems.

You can choose any drive type that does not exceed the capacity of the drive. Table 3-3 gives the logical parameters that provide the maximum capacity on the Quantum Fireball lct20 family of hard disk drives.

Table 3-3 *Logical Addressing Format*

	QUANTUM FIREBALL lct20			
	10.0	**20.0**	**30.0**	**40.0**
LBA Capacity	10,262 MB	20,416 MB	30,020 MB	40,027 MB
CHS Capacity	8,455 MB* (10,262 MB)	8,455 MB* (20,416 MB)	8,455 MB* (30,020 MB)	8,455 MB* (40,027 MB)
Logical Cylinders	16,383*	16,383*	16,383*	16,383*
Logical Heads	16	16	16	16
Logical Sectors/ Track	63	63	63	63
Total Number Logical Sectors	20,044,080	39,876,480	58,633,344	78,177,792

Note: *Capacity may be restricted to 8.4 GB (or less) due to system BIOS limitations. Check with your system manufacturer to determine if your BIOS supports LBA Mode for hard drives greater than 8.4 GB. Default logical cylinders is limited to 16,383 as per the ATA-4 specifications.

To match the logical specifications of the drive to the drive type of a particular BIOS, consult the system's drive-type table. This table specifies the number of cylinders, heads, and sectors for a particular drive type.

You must choose a drive type that meets the following requirements:

For the 10.0 GB, 15.0 GB, 20.5 GB, 30.0 GB AT:

Logical Cylinders x Logical Heads x Logical Sectors/Track x 512 = 8,455,200,768

4. Boot the system using the operating system installation disk—for example, MS-DOS—then follow the installation instructions in the operating system manual.

3.11 CONSIDERATIONS FOR MOUNTING AN AV DRIVE

For CE applications in which quiet operation of the drive is very critical, we recommend that the drive be mounted on shock and vibration isolators. Refer to the sections below when mounting an AV drive in a CE box.

There are many different ways to mount the hard drive in a set-top box (Figure 3-12 shows a drive mounted in a set-top box). The method that you should choose depends on many objectives of the designer of the set-top box. Some of these objectives are:

1. The acoustical tolerance level of the design. For example, a company that is designing a set-top box for games is not as concerned about noise emissions as a company that is designing a box for video recording and playback in the home.
2. Shock and vibration tolerance level. Users of game boxes would be expected to treat the equipment more harshly than users for other applications.
3. Cost objectives. The methods that achieve no sound emissions will cost more than those that reduce the sound emissions by 50 percent, and so on.
4. Packaging constraints. For example, the size of the set-top box and the components that go into the box have some influence on the mounting of the hard drive.
5. Thermal conditions in the set-top box and environmental ambient temperature requirements. The higher the operational temperature, the greater the problems that will exist from a life expectancy view of the hard drive. Designers need to determine whether to include a fan and, if not, how much ventilation to provide and what effects ventilation will have on acoustics.

Four different sources of acoustic emissions must be dealt with in a hard drive:

1. Spindle motor airborne noise
2. Spindle vibration induced acoustics
3. Head-seeking airborne noise
4. Head-seeking vibration-induced acoustics

The vibrations from the spindle balance and from the head positioner seeking can be amplified many times by an incorrectly designed set-top box and drive mounting method. The designer should be aware of the frequencies of these vibrations and try to avoid designing the box with nearby resonance frequencies, if possible.

In addition, note that signal cable routing and EMI problems can be created if care is not given to the placement of cables and the drive in the set-top box.

3.12 Causes for Noise Generated by Hard Disk Drives

Noise generated by a hard disk drive can be caused by three sources:

1. Head positioner moving from track to track during the seeking process
2. Spindle motor switching phases during the spinning process
3. Out of balance condition of the disks as they spin

These noise problems are in two forms:

1. Mechanical vibrations
2. Airborne acoustical noise

3.13 Cautions in Reducing the Effects of External Shock Forces

The normal shock mount isolation methods that are applied to machinery must be used with caution in the case of a hard disk drive. The problem with shock mounting a hard drive is that motion is produced to the drive as a result of the elastomeric shock mounts. This motion causes the head positioner to move off the data tracks and thus produce read error problems and time-consuming read retry attempts by the hard drive. As a result, the data buffer can empty and video problems can result.

3.13.1 Cautions in Reducing the Effects of External Vibration Forces

The normal vibration mount isolation methods that are applied to machinery must be used with caution in the case of a hard disk drive. The problem with shock mounting a hard drive is that motion is produced to the drive as a result of the elastomeric vibration mounts. This motion causes the head positioner to move off the data tracks and thus produce read and write error problems and time-consuming read retry attempts by the hard drive. As a result, the data buffer can empty and video problems can result.

The best results for transferring data during vibrations is to directly mount the hard drive to the chassis, which however, will result in an increase in noise. The limitation on the data transfer rate depends on buffer size and the frequency band of the vibration.

3.13.2 Acoustical Tolerance Level

The acoustical reduction hardware required for a product depends on the acoustical tolerance level for a given product. The cost of the hardware required will also be a factor in the choice of how much acoustical reduction is needed for a given design. Almost complete silence, for example, in the set-top box application will be much more expensive to accomplish than a 50 percent reduction in sound level. The choice of hardware is a trade-off between cost of the hardware and acoustical reduction needed.

At a minimum, a method for isolating the spin motor and head seek vibration from the customer chassis should be considered. Two ways of accomplishing this goal are by using:

- A sheet of vibration isolation foam material such as parts 4 and 5 of Figure 3-13.

- Vibration isolation grommets as shown with part 6 of Figure 3-13.

At the other end of the scale, with complete acoustical and vibration isolation, is the enclosed method shown with parts 1, 2, 3, 4, and 5 of Figure 3-13.

3.14 Mounting Methods to Reduce Noise and Vibration in Set-Top Box Assemblies

The application of the following four methods eliminate noise and vibration sources produced by hard disk drives:

1. Application of damping materials on the mechanical vibration generators
2. Application of acoustical sound absorption materials for airborne noise
3. Application of physical noise barrier containment
4. Isolation of the drives from the set-top box with shock absorbing materials

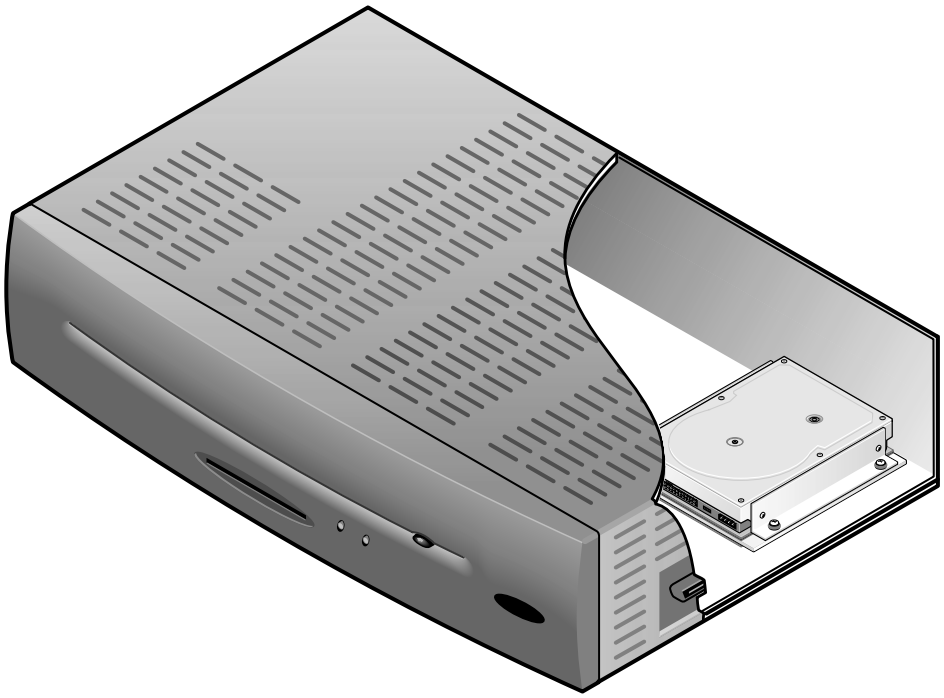

Figure 3-12 *Drive Mounted in a Set-Top Box*

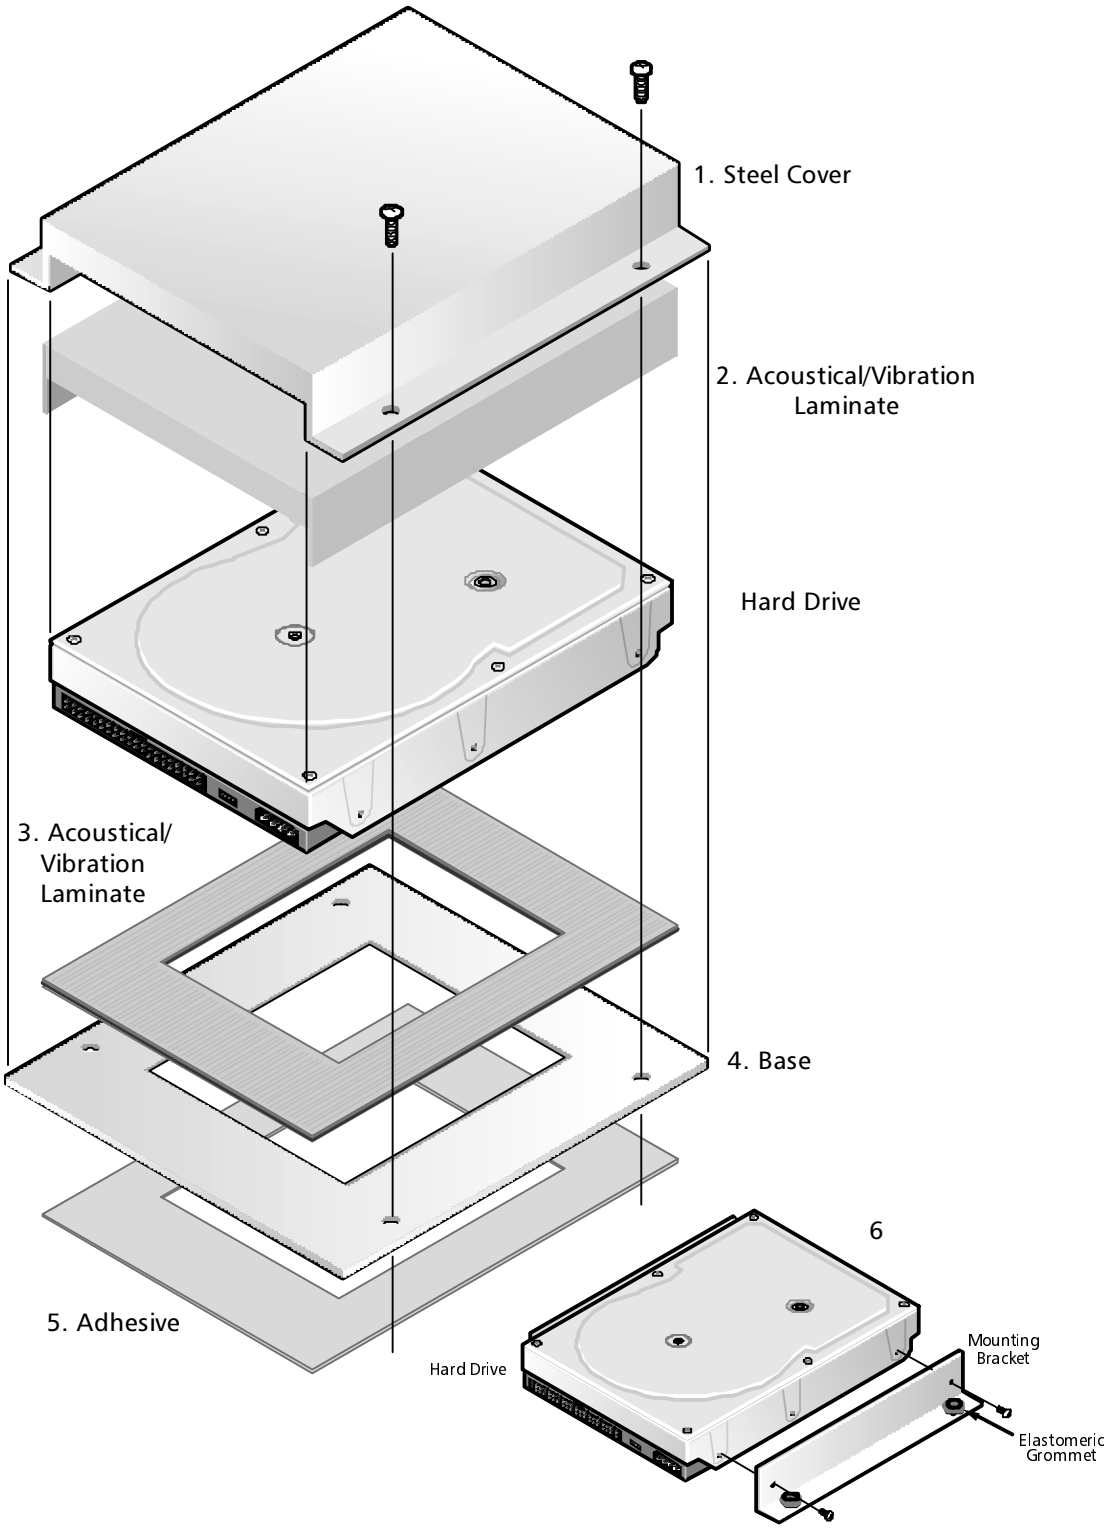

Figure 3-13 *Set-Top Enclosure (Container Box) Assembly*

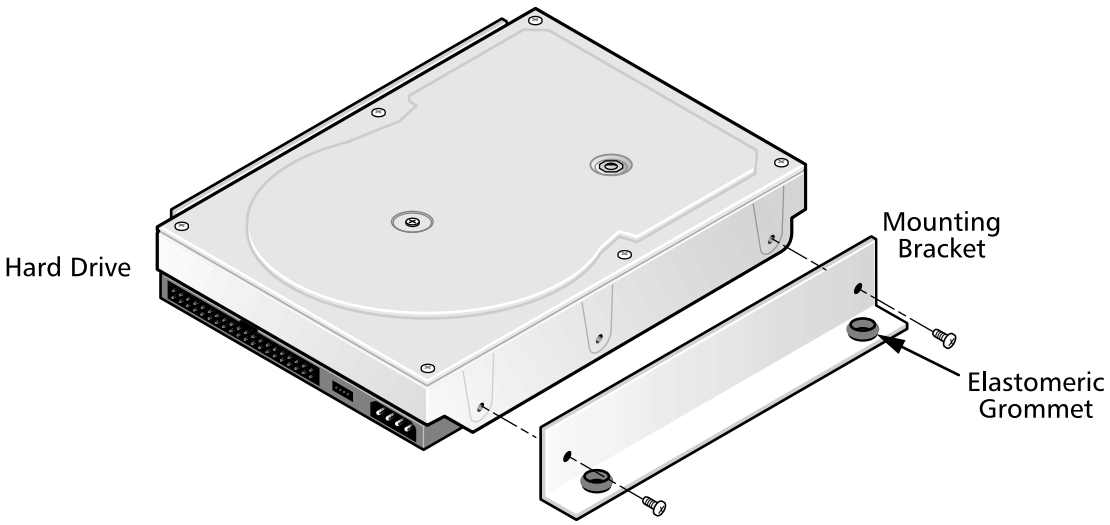

Figure 3-14 *Mounting an AV Drive*

3.14.1 Set-Top Box Enclosure

1. The drive or drives should be mounted, if possible, in such a way that there is a 0.25-inch air gap (6 mm) between the bottom of the drive and the top of the surface beneath the hard drive (the container bottom or acoustical foam, for example). This space will allow air circulation to the components on the hard drive PCB.
2. An acoustical foam lined box, lined both on the top and sides of the inside of the box, should be designed to fit over the hard drives. The foam should be about 5 mm thick, compressed to about 3 mm, and have a layer of vibration damping material bonded to it that is about 1 mm thick. For best results, this sandwich should be bonded to the inside of the steel cover. This type of material can be obtained in a one piece prefab sandwich form from acoustical materials suppliers. The box should be sheet metal steel, minimum of 1 mm thick, as shown in Figure 3-13.
3. The drive mounting base should be mounted in such a way that there is vibration isolation between the drives and the set-top box. This design prevents vibration transmission to the walls of the set-top box. This vibration isolation can be accomplished with elastomeric vibration damping materials or felt types of materials, or with standard vibration isolation grommets.

3.14.2 Recommended Acoustic and Vibration Damping Material

The most intense acoustical output of the drive comes while the head positioner is seeking, as the drive is moving the heads from one track to another. Two major frequency points contribute to the acoustical emissions of the disk drive during this seek process, in the audible range of 30 Hz to 18 KHz. One is located at about 200 Hz and one at about 1.1 KHz.

The most cost-effective way to reduce these audible frequencies is to make an acoustical and vibration barrier using acoustical and vibration blocking foam. These acoustical barriers are available from several manufacturers such as Roush Industries, EAR Corp., and Sonex Inc.

The most effective material tested is a layered combination of 1 mm EAR C2003 vibration damping material bonded to a 10 mm EAR E50 acoustical absorbing foam. This material is applied to the entire top and bottom of the hard drive and is evenly compressed to a thickness of about 4 mm for each layer. This compression is best done by using a steel sheet metal about 1.5 mm thick. Similar results were obtained by using E N Murray Vibrasorb damping sheet, 1 mm thick, with Hyfonic 1 acoustical foam, 3 mm thick.

Air circulation around the PCB area of the hard drive is also important. The life expectancy of the PCB components is related to their temperature. Care should be taken when designing the acoustical treatments to keep this area clear of obstructions.

It is important to be aware of the container box vibration amplification problem when mounting a hard drive into a set-top box. Because of the large flat areas of the container walls, a speaker diaphragm effect occurs that causes small drive vibrations to be amplified. **The enclosure must be designed to be ridged and have resonance frequencies well above and away from the spindle rotation and multiples of the spindle rotation frequency.** Because of the type of design required for the box, this goal can be very difficult to achieve. Using damping methods to isolate the drive usually gives the best results.

Another problem to be aware of when using vibration isolators with a hard drive is a possible increase in the data transfer throughput time. This increase is caused when the hard drive is free to move around on the vibration isolators as the heads are moving onto and locking on a data track. Also, external vibration can have a negative effect on data throughput time when mounted on soft shock isolators if they are not selected properly.

3.14.3 Recommended Acoustic and Vibration Damping Material

A sandwich of the following is recommended:

1. Vibrasorb Damping Sheet, 1 mm thick. (E N Murray or equivalent)
2. Hyfonic 1 acoustic foam, 3 mm thick. (E N Murray or equivalent)
3. Sheet adhesive, 0.13 mm thick (E N Murray or equivalent)

or

1. EAR C2003, 1 mm thick (EAR Corp.)
2. EAR E50 acoustic absorbing foam (EAR Corp.)
3. Sheet adhesive, 0.13 mm thick

3.15 Acoustical Measurements

3.15.1 Measuring Acoustical Emissions

Two methods are used to measure acoustical emissions:

1. Sound power
2. Sound pressure

Since sound power measurements require extremely complex and expensive measurement equipment, all measurements have to been taken in sound pressure. This method requires that the measurement conditions be carefully specified. Sound pressure is a measurement of a single point around the acoustic generator. Microphone distance, noise background, the side of the hard drive being measured, and even the test measurement instrument, all affect the measurement.

Measurements are made using the A weighted decibel scale. This scale limits the frequency range from 20 Hz to 20 KHz and the amplitude is adjusted to match the human ear response to the noise.

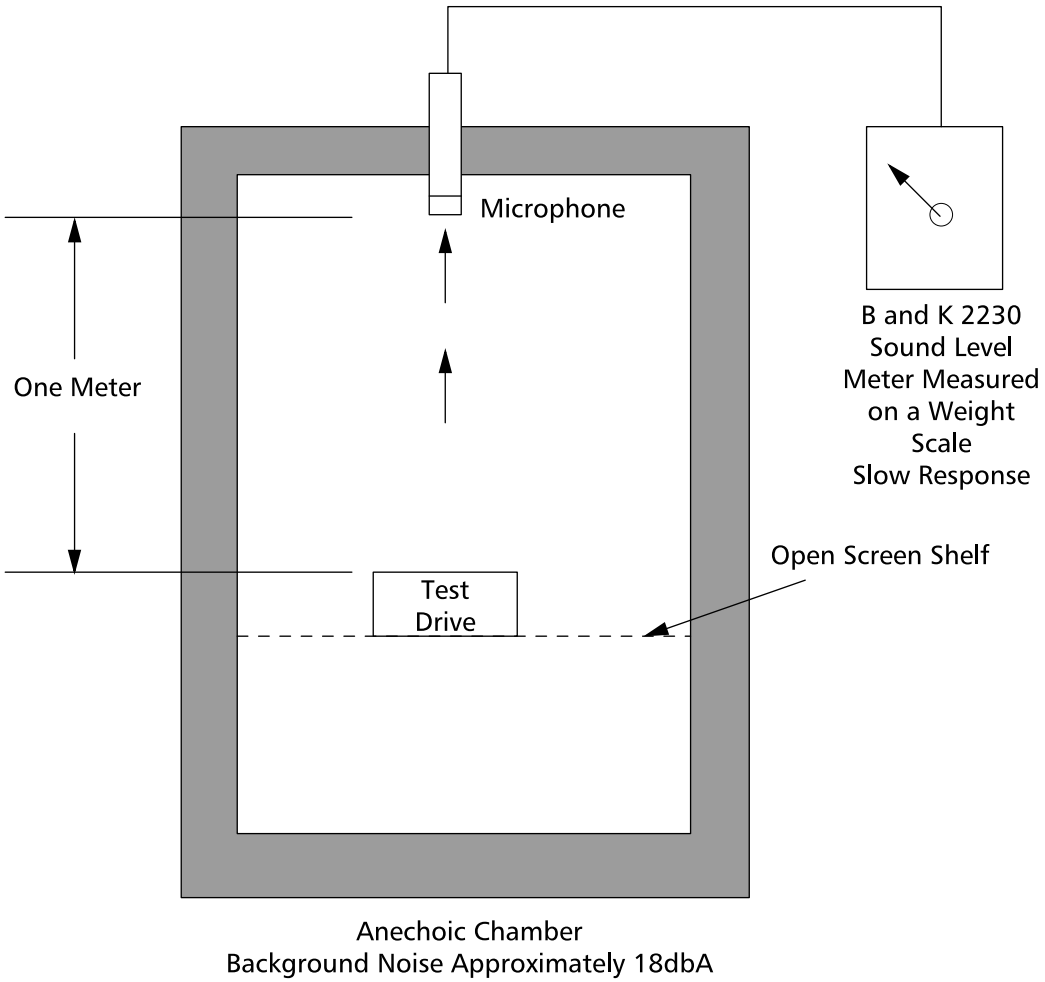

Figure 3-15 *Measuring Acoustical Emissions*

For more details, refer to the "CEBU DE AV Mounting Guide."

Chapter 4
SPECIFICATIONS

This chapter gives a detailed description of the physical, electrical, and environmental characteristics of the Quantum Fireball lct20 hard disk drives.

4.1 SPECIFICATION SUMMARY

Table 4-1 gives a summary of the Quantum Fireball lct20 hard disk drives.

Table 4-1 *Specifications*

DESCRIPTION	QUANTUM FIREBALL lct20			
	10.0 GB AT	**20.0 GB AT**	**30.0 GB AT**	**40.0 GB AT**
Capacity	10,262 MB	20,416 MB	30,020 MB	40,027 MB
Nominal rotational speed (rpm)[1]	4502	4502	4502	4502
Number of Disks	1	1	2	2
Number of R/W heads	1	2	3	4
Data Organization:				
Zones per surface	15	15	15	15
Tracks per surface	33,518	33,518	33,518	33,518
Total tracks	33,518	67,036	100,554	134,072
Sectors per track:				
Inside zone	385	385	385	385
Outside zone	721	721	721	721
Total User Sectors	20,044,080	39,876,480	58,633,344	78,177,792
Bytes per sector	512	512	512	512
Number of tracks per cylinder	1	2	3	4
Recording:				
Recording technology	Multiple Zone	Multiple Zone	Multiple Zone	Multiple Zone
Maximum linear density	401,472 fci	401,472 fci	401,472 fci	401,472 fci
Encoding method	24/25 NPRML	24/25 NPRML	24/25 NPRML	24/25 NPRML
Interleave	1:1	1:1	1:1	1:1
Track density	34,000 tpi	34,000 tpi	34,000 tpi	34,000 tpi

DESCRIPTION	QUANTUM FIREBALL lct20			
	10.0 GB AT	20.0 GB AT	30.0 GB AT	40.0 GB AT
Maximum effective areal density	13,104 Mbits/in^2	13,104 Mbits/in^2	13,104 Mbits/in^2	13,104 Mbits/in^2
Performance:				
Seek times:				
Average Read	12 ms typ. 15 ms max.	12 ms typ. 15 ms max.	12 ms typ. 15 ms max.	12 ms typ. 15 ms max.
Track-to-track	2 ms typical	2 ms typical	2 ms typical	2 ms typical
Average write	15 ms typ. 18 ms max.	15 ms typ. 18 ms max.	15 ms typ. 18 ms max.	15 ms typ. 18 ms max.
Full stroke	24 ms typ. 28 ms max.	24 ms typ. 28 ms max.	24 ms typ. 28 ms max.	24 ms typ. 28 ms max.
Data transfer Rates:				
Disk to Read Once a Revolution[2, 3]	115.12 Mb/ sec. min. 216.79 Mb/ sec. max.	115.12 Mb/ sec. min. 216.79 Mb/ sec. max.	115.12 Mb/ sec. min. 216.79 Mb/ sec. max.	115.12 Mb/ sec. min. 216.79 Mb/ sec. max.
Disk to Read Instantaneously[2]	151.47 Mb/ sec. min. 283.20 Mb/ sec. max.	151.47 Mb/ sec. min. 283.20 Mb/ sec. max.	151.47 Mb/ sec. min. 283.20 Mb/ sec. max.	151.47 Mb/ sec. min. 283.20 Mb/ sec. max.
Read Buffer to ATA Bus (PIO Mode with IORDY)	16.7 MB/sec. maximum	16.7 MB/sec. maximum	16.7 MB/sec. maximum	16.7 MB/sec. maximum
Read Buffer to ATA Bus (Ultra ATA Mode)	66 MB/sec. maximum	66 MB/sec. maximum	66 MB/sec. maximum	66 MB/sec. maximum
Buffer Size	128 KB	128 KB	128 KB	128 KB
Reliability:				
Seek error rate[3]	1 in 10^6	1 in 10^6	1 in 10^6	1 in 10^6
Unrecoverable error rate[3]	1 in 10^{14}	1 in 10^{14}	1 in 10^{14}	1 in 10^{14}
Error correction method (with cross check)	288-bit Reed Solomon	288-bit Reed Solomon	288-bit Reed Solomon	288-bit Reed Solomon
Projected MTBF[4]	625,000 hrs	625,000 hrs	625,000 hrs	625,000 hrs
Contact Start/Stop Cycles[4] (Ambient temperature)	50,000 min.	50,000 min.	50,000 min.	50,000 min.
Auto head-park method				

1. Spindle speed is 4400 + 0.43%/-1.47%. The design point is 4403 with 40 MHz reference. Actual typical reference frequency is 39.69 MHz which results in typical RPM of 4369. This is the result of the use of a creamic resonator reference.

2. Disk to read buffer transfer rate is zone-dependent, instantaneous

3. Refer to Section 4.14, “DISK ERRORS” for details on error rate definitions.

4. CSS specifications assumes a duty cycle of one power off operation for every one idle spin down.

4.2 FORMATTED CAPACITY

At the factory, the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives receive a low-level format that creates the actual tracks and sectors on the drive. Table 4-2 shows the capacity resulting from this process. Formatting done at the user level, for operation with DOS, UNIX, or other operating systems, may result in less capacity than the physical capacity shown in Table 4-2.

Table 4-2 *Formatted Capacity*

	QUANTUM FIREBALL lct20			
	10.0 GB AT	**20.0 GB AT**	**30.0 GB AT**	**40.0 GB AT**
Capacity	10,262 MB	20,416 MB	30,020 MB	40,027 MB
Number of 512-byte sectors available	20,044,080	39,876,480	58,633,344	78,177,792

4.3 DATA TRANSFER RATES

Data is transferred from the disk to the read buffer at a rate of up to 283.20 Mb/s in bursts. Data is transferred from the read buffer to the ATA bus at a rate of up to 16.7 MB/s using programmed I/O with IORDY, or at a rate of up to 66 MB/s using Ultra ATA/100. For more detailed information on interface timing, refer to Chapter 6.

4.4 TIMING SPECIFICATIONS

Table 4-3 illustrates the timing specifications of the Quantum Fireball lct20 hard disk drives.

Table 4-3 *Timing Specifications*

PARAMETER	TYPICAL NOMINAL[1]	WORST CASE[2]
Sequential Cylinder Switch Time[3]	2 ms	3 ms
Sequential Head Switch Time[4]	2 ms	3 ms
Random Average (Read or Seek)[9]	12 ms	15 ms
Random Average (Write)[9]	15 ms	18 ms
Full-Stroke Seek	24 ms	28 ms
Average Rotational Latency	6.66 ms	—
Power On[5] to Drive Ready[6]	8 seconds	12 seconds
Standby[7] to Interface Ready	6 seconds	12 seconds
Spindown Time, Standby Command	6 seconds	15 seconds[8]
Spindown Time, Power loss	12 seconds	30 seconds[8]

1. Nominal conditions are as follows:
   - Nominal temperature 77°F (25°C)
   - Nominal supply voltages (12.0V, 5.0V)
   - No applied shock or vibration
2. Worst case conditions are as follows:
   - Worst case temperature extremes 41°F to 131°F (5°C to 55°C)
   - Worst case supply voltages (12.0V ±10%, 5.0 V ±5%)
3. Sequential Cylinder Switch Time is the time from the conclusion of the last sector of a cylinder to the first logical sector on the next cylinder (no more than 6% of cylinder switches exceed this time).
4. Sequential Head Switch Time is the time from the last sector of a track to the beginning of the first logical sector of the next track of the same cylinder (no more than 6% of head switches exceed this time).
5. Power On is the time from when the supply voltages reach operating range to when the drive is ready to accept any command.
6. Drive Ready is the condition in which the disks are rotating at the rated speed, and the drive is able to accept and execute commands requiring disk access without further delay at power or start up. Error recovery routines may extend the time to as long as 45 seconds for drive ready.
7. Standby is the condition at which the microprocessor is powered, but not the HDA. When the host sends the drive a shutdown command, the drive parks the heads away from the data zone, and spins down to a complete stop.
8. After this time it is safe to move the disk drive
9. Average random seek is defined as the average seek time between random logical block addresses (LBAs).

4.5 POWER

The Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives operate from two supply voltages:

- +12V ±10%
- +5V ±5%

The allowable ripple and noise is 250 mV peak-to-peak for the +12 Volt supply and 150 mV peak-to-peak for the +5 Volt supply. The allowable ripple noise for each voltage is measured at the power supply with the drive not attached to it.

4.5.1 Power Sequencing

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

4.5.2 Power Reset Limits

When powering up, the drive remains reset (inactive) until both rising voltage thresholds reset limits are exceeded for ≥20 ms. When powering down, the drive becomes reset when either supply voltage drops below the falling voltage threshold for ≥ 500 ns.

Table 4-4 *Power Reset Limits*

DC VOLTAGE	THRESHOLD	HYSTERESIS
+5 V	V_{LT}= 3.7 V minimum V_{HT}= 4.5 V maximum	55 mV (typical)
+12 V	V_{LT}= 9.3 V maximum 8.7 V minimum V_{HT}= 9.55 V maximum 8.85 V minimum	200 mV (typical)

4.5.3 Power Requirements

Table 4-5 lists the voltages and typical average corresponding currents for the various modes of operation of the Quantum Fireball lct20 hard disk drives.

Table 4-5 *Typical Power and Current Consumption*

MODE OF OPERATION	TYPICAL AVERAGE CURRENT[2] (mA RMS unless otherwise noted)			
	+12 V		+5V	
MODEL NUMBER	10.0/20.0 GB (1-Disk)	30.0/40.0 GB (2-Disk)	10.0/20.0 GB (1-Disk)	30.0/40.0 GB (2-Disk)
Startup[1] (peak)	1141	1141	397	397
Idle[3]	157	178	320	321
Maximum Seeking[4]	592	575	329	329
Standby[5]	20	24	68	68
Read/Write On Track[6]	162	185	363	364

MODE OF OPERATION	TYPICAL AVERAGE POWER[2] (WATTS)	
MODEL NUMBER	10.0/20.0 GB (1-Disk)	30.0/40.0 GB (2-Disk)
Startup[1] (peak)	16	16
Idle[3]	3	4
Maximum Seeking[4]	9	9
Standby[4]	0.6	0.6
Read/Write On Track[6]	4	4

1. Current is rms except for startup. Startup current is the typical peak current of the peaks greater than 10 ms in duration. This power will be required for less than 5 seconds.

2. Power requirements reflect nominal for +12V and +5V power.

3. 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.

4. Maximum seeking is defined as continuous random seek operations with typical host command overhead.

5. Standby mode is defined as when the motor is stopped, the actuator is parked, and all electronics except the interface control are in low power state. Standby occurs after a programmable time-out after the last host access. Drive ready and seek complete status exist. The drive leaves standby upon receipt of a command that requires disk access or upon receiving a spinup command.

6. Read/Write On Track is defined as 50% read operations and 50% write operations on a single physical track.

4.6 ACOUSTICS

Table 4-6 specifies the acoustical characteristics of the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drive. The acoustics are measured in an anechoic chamber with background noise at least <10dBA less than the expected sound pressure Lp(A). To distinguish between sound power and sound pressure standards, sound power Lw(A) is specified in Bels. The relationship between bels and dBA for sound power is 1 bel = 10dBA.

Table 4-6 *Acoustical Characteristics—Sound Power*

OPERATING MODE	MEASURED SOUND POWER (TYPICAL)	MEASURED SOUND POWER (MAXIMUM)
Idle On Track		
10.0/20.0 GB AT (1-Disk)	2.8 Bels	3.4 Bels
30.0/40.0 GB AT (2-Disk)	2.9 Bels	3.5 Bels
Seeking Random		
10.0/20.0 GB AT (1-Disk)	3.0 Bels	3.5 Bels
30.0/40.0 GB AT (2-Disk)	3.1 Bels	3.6 Bels

4.7 MECHANICAL

Height: 1.0 in.(25.4 mm)

Width: 4.0 in.(101.6 mm)

Depth: 5.75 in.(146.1 mm)

Weight: Quantum Fireball lct20 40.0 GB AT (2-Disk)1.3 lb

4.8 ENVIRONMENTAL CONDITIONS

Table 4-7 summarizes the environmental specifications of the Quantum Fireball lct20 hard disk drives.

Table 4-7 *Environmental Specifications*

PARAMETER	OPERATING	NON-OPERATING
Temperature[1] (Non-condensing)	5° to 55°C (41° to 131°F)	-40° to 65°C (-40° to 149°F)
Temperature Gradient (Non-condensing)	24°C/hr maximum (75.2°F/hr)	48°C/hr maximum (118.4°F/hr)
Humidity[2] (Non-condensing) Maximum Wet Bulb Temperature	5% to 85% RH 37°C (98.6°F)	5% to 95% RH 46°C (114.8°F)
Humidity Gradient	30% / hour	30% / hour
Altitude[3, 4]	–200 m to 3,000 m (–650 to 10,000 ft.)	–200 m to 12, 000 m (–650 to 40,000 ft.)
Altitude Gradient	1.5 kPa/min	8 kPa/min

1. Maximum operating temperature must not exceed the drive at any point along the drive form factor envelope. Airflow or other means must be used as needed to meet this requirement.

2. The humidity range shown is applicable for temperatures whose combination does not result in condensation in violation of the wet bulb specifications.

3. Altitude is relative to sea level.

4. The specified drive uncorrectable error rate will not be exceeded over these conditions.

4.9 SHOCK AND VIBRATION

The Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives can withstand levels of shock and vibration applied to any of its three mutually perpendicular axes, or principal base axis, as specified in Table 4-8. 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 loss of user data at power on.

When packed in its 1-pack shipping container, the Quantum Fireball lct20 drives can withstand a drop from 30 inches onto a concrete surface on any of its surfaces, six edges, or three corners. The 20-pack shipping container can withstand a drop from 30 inches onto a concrete surface on any of its surfaces, six edges, or three corners.

Table 4-8 *Shock and Vibration Specifications*

SHOCK[1]	OPERATING	NONOPERATING
Translational 1/2 sine wave 11 ms duration 2 ms duration Trapezoidal	 10.0 Gs (write) 20.0 Gs (write) 63.0Gs (read) —	 70 Gs 300 Gs 110Gs (1ms) 80 Gs, 18 in/sec velocity change
Rotational 2 ms applied at actuator pivot point	—	25,000 rad/sec[2]
Vibration[1] Random Vibration (G^2/Hz) Sine wave (peak to peak) 1/2 octave per minute sweep	 0.003 (5 – 300 Hz) 1 G P-P 5-300 Hz	 0.05 (10 – 300 Hz) 2 Gs P-P 5–400 Hz

1. The specified drive unrecovered error rate will not be exceeded over these conditions.

4.10 HANDLING the DRIVE

Before handling the Quantum hard disk drive some precautions must to be taken to ensure that the drive is not damaged. Use both hands while handling the drive and hold the drive by its edges. Quantum drives are designed to withstand normal handling, however, hard drives can be damaged by electrostatic discharge (ESD), dropping the drive, rough handling, and mishandling. Use of a properly grounded wrist strap to the earth is strongly recommended. Always keep the drive inside its special antistatic bag until ready to install.

Note: To avoid causing any damage to the drive do not touch or remove the Printed Circuit Board (PCB) or any of its components when handling the drive. Doing so will cause ESD damage to the drive.

4.11 RELIABILITY

Mean Time Between Failures (MTBF):	The projected field MTBF is 625,000 hours. The Quantum MTBF numbers represent Bell-Core TR-332 Issue #6, December 1997 MTBF predictions and represent the minimum MTBF that Quantum or a customer would expect from the drive.
Component Life:	5 years
Preventive Maintenance (PM):	Not required
Start/Stop:	50,000 cycles @ ambient

Note: Drive reliability is closely related to the temperature the drive is exposed to. The operational design environment is 35°C at 71% utilization or 6240 POH/year.

4.12 ELECTROMAGNETIC SUSCEPTIBILITY

5 Volts/meter over a range of 10 KHz to 1 GHz.

4.13 SPINDLE IMBALANCE

0.5 g-mm maximum

(This is approximately equivalent to 0.05 g emitted vibrations)

4.14 DISK ERRORS

Table 4-9 provides the error rates for the Quantum Fireball lct20 hard disk drives.

Table 4-9 *Error Rates*

ERROR TYPE	MAXIMUM NUMBER OF ERRORS
Retry recovered read errors[1]	1 event per 10^9 bits read
Unrecovered data errors[2]	1 event per 10^{14} bits read
Seek errors[3]	1 error per 10^6 seeks

1. Retry recovered read errors are errors which require retries for data correction. Errors corrected by ECC on-the-fly are not considered recovered read errors. Read on arrival is disabled to meet this specification. Errors corrected by the thermal asperity correction are not considered recovered read errors.

2. Unrecovered 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 the programmed limit for unsuccessful retries and the application of quadruple-burst error correction.

3. Seek errors occur when the actuator fails to reach (or remain) over the requested cylinder and the drive requires the execution of a full recalibration routine to locate the requested cylinder.

Note: Error rates are for worst case temperature and voltage.

Chapter 5
BASIC PRINCIPLES OF OPERATION

This chapter describes the operation of Quantum Fireball lct20 AT hard disk drives' functional subsystems. It is intended as a guide to the operation of the drive, rather than a detailed theory of operation.

5.1 QUANTUM FIREBALL lct20 DRIVE MECHANISM

This section describes the drive mechanism. Section 5.2 describes the drive electronics. The Quantum Fireball lct20 hard disk drives consist of a mechanical assembly and a PCB as shown in Figure 5-1.

The head/disk assembly (HDA) contains the mechanical subassemblies of the drive, which are sealed under a metal cover. The HDA consists of the following components:

- Base casting
- DC 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: To ensure that 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.

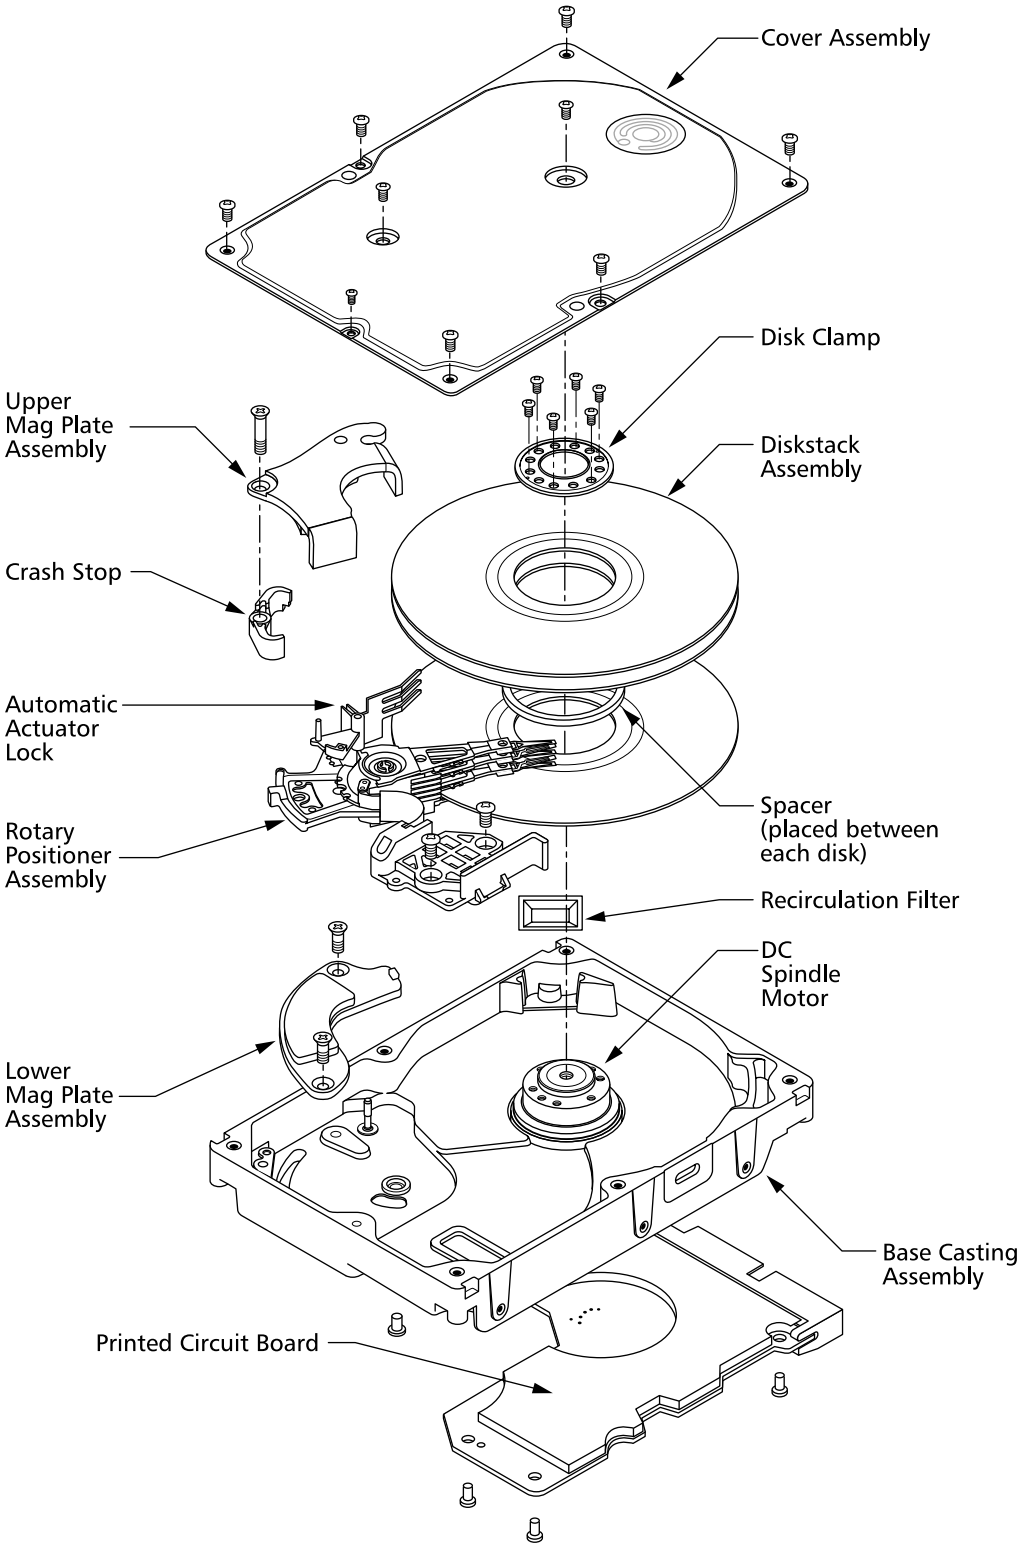

Figure 5-1 *Quantum Fireball lct20 AT Hard Disk Drive Exploded View*

5.1.1 Base Casting Assembly

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

5.1.2 DC Motor Assembly

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

5.1.3 Disk Stack Assemblies

The disk stack assembly in the Quantum Fireball lct20 hard disk drives consist of disks secured by a disk clamp. The aluminum-alloy disks have a sputtered thin-film magnetic coating.

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

Table 5-1 *Cylinder Contents*

CYLINDER CONTENTS	ZONE[1]	NUMBER OF TRACKS	SECTORS PER TRACK	DATA RATE (Mb/S)
System Data	15		385	149.76
	14		410	157.64
	13		437	168.53
	12		459	178.04
	11		492	187.48
	10		517	200.53
	9		546	211.20
	8		574	221.54
	7		601	232.66
	6		623	241.69
	5		650	251.35
	4		679	264.00
	3		693	271.06
	2		710	276.48
	1		721	280.80
	0		492	187.48

1. For user data, zone 15 is the innermost zone and zone 1 is the outermost zone.

5.1.4 Headstack Assembly

The headstack assembly consists of read/write heads, head arms, and a coil joined together by insertion molding to form a rotor subassembly, bearings, and a flex circuit. Read/write heads mounted to spring-steel flexures are swage mounted onto the rotary positioner assembly arms.

The flex circuit connects the headstack assembly to the PCB through a compression connector. The flex circuit contains a read preamplifier/write driver IC.

5.1.5 Rotary Positioner Assembly

The rotary positioner, or rotary voice-coil actuator, is a Quantum-proprietary design that consists of upper and lower permanent magnet plates, a rotary single-phase coil molded around the headstack mounting hub, and a bearing shaft. The single bi-polar magnet consists of two alternating poles and is bonded to the magnet plate. A resilient crash stop prevents 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 magnet 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 drive uses a dedicated landing zone, an actuator magnetic retract, and Quantum's patented Airlock®. The Airlock holds the headstack in the landing zone whenever the disks are not rotating. It consists of an air vane mounted near the perimeter of the disk stack, and a locking arm that restrains the actuator arm assembly.

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

5.1.7 Air Filtration

The Quantum Fireball lct20 AT hard disk drives are Winchester-type drives. The heads fly very close to the media surface. Therefore, it is essential that the air circulating within the drive be kept free of particles. Quantum assembles the drive in a Class-100 purified air environment, then seals the drive with a metal cover. When the drive is in use, the rotation of the disks forces the air inside of the drive through an internal 0.3 micron filter. The internal HDA cavity pressure equalizes to the external pressure change by passing air through a 0.3 micron, carbon impregnated breather filter.

5.2 DRIVE ELECTRONICS

Advanced circuit (Very Large Scale Integration) design and the use of miniature surface-mounted devices and proprietary VLSI components enable the drive electronics, including the ATA bus interface, to reside on a single printed circuit board assembly (PCBA).

Figure 5-2 contains a simplified block diagram of the Quantum Fireball lct20 hard disk drive electronics.

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

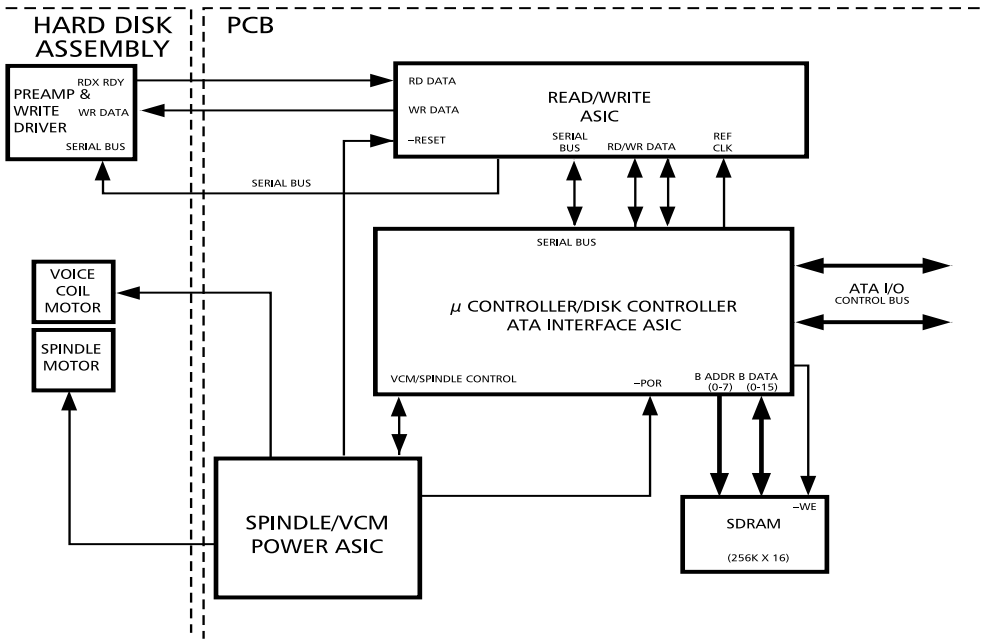

Figure 5-2 *Quantum Fireball lct20 AT Hard Disk Drive Block Diagram*

5.2.1 Integrated μProcessor, Disk Controller and ATA Interface Electronics

The μProcessor, Disk Controller, and ATA Interface electronics are contained in a proprietary ASIC developed by Quantum, as shown below in Figure 5-3.

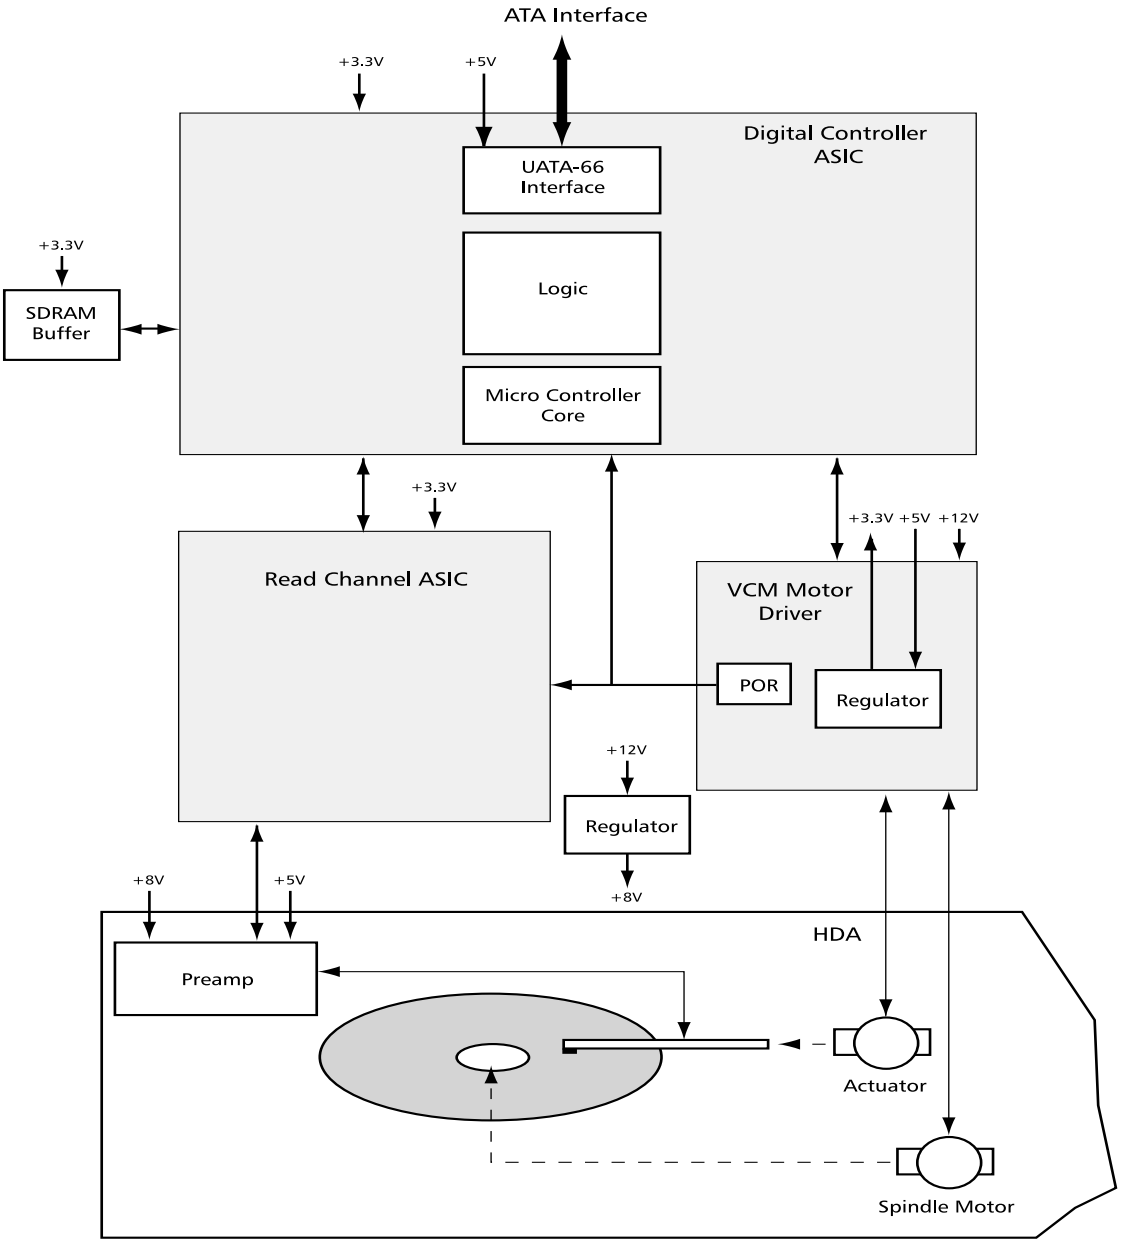

Figure 5-3 *Block Diagram*

The integrated μProcessor, Disk Controller, and ATA Interface Electronics have nine functional modules (described below):

- μProcessor
- Digital Synchronous Spoke (DSS)
- Error Correction Code (ECC) Control
- Formatter
- Buffer Controller
- Servo Controller, including PWM
- Serial Interface
- ATA Interface Controller
- Motor Controller

5.2.1.1 μProcessor

The μProcessor core provides local processor services to the drive electronics under program control. The μProcessor manages the resources of the Disk Controller, and ATA Interface internally. It also manages the Read/Write ASIC (Application Specific Integrated Circuit), and the Spindle/VCM driver externally.

5.2.1.2 Digital Synchronous Spoke

The DSS decodes servo information written on the drive at the factory to determine the position of the read/write head. It interfaces with the read/write channel, process timing and position information, and stores it in registers that are read by the servo firmware.

5.2.1.3 Error Correction Code (ECC) Control

The Error Correction Code (ECC) Control block utilizes a Reed-Solomon encoder/decoder circuit that is used for disk read/write operations. It uses a total of 36 redundancy bytes organized as 32 ECC (Error Correction Code) bytes with four interleaves, and four cross-check bytes. The ECC uses eight bits per symbol and four interleaves. This allows quadruple-burst error correction of at least 96, and as many as 128 bits in error.

5.2.1.4 Formatter

The Formatter controls the operation of the read and write channel portions of the ASIC. To initiate a disk operation, the μProcessor loads a command into the ASIC registers.

The Formatter also directly drives the read and write gates (RG, WG) and Command Mode Interface of the Read/Write ASIC and the R/W Preamplifier, as well as passing write data to the Precompensator circuit in the Read/Write ASIC.

5.2.1.5 Buffer Controller

The Buffer Controller supports a 512 Kbyte buffer, which is organized as 256 K x 16 bits. The 16-bit width implementation provides a 60 MB/s maximum buffer bandwidth. This increased bandwidth allows the μProcessor to have direct access to the buffer, eliminating the need for a separate μProcessor RAM IC.

The Buffer Controller supports both drive and host address rollover and reloading, to allow for buffer segmentation. Drive and host addresses may be separately loaded for automated read/write functions.

The Buffer Controller operates under the direction of the μProcessor.

5.2.1.6 Servo Processor

The Servo Processor in the Read Write Channel ASIC provides servo data recovery and burst demodulation to extract the actuator position information. This information is processed in the controller ASIC/microprocessor, and a control signal is output to the VCM in the Power ASIC. This controls the current in the actuator coil which controls the position of the actuator.

5.2.1.7 Read/Write Interface

The Read/Write interface allows the integrated μprocessor, disk controller to communicate with the Read/Write chip.

5.2.1.8 ATA Interface Controller

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

5.2.1.9 Motor Controller

The Motor Controller controls the spindle and voice coil motor (VCM) mechanism on the drive.

5.2.2 Read/Write ASIC

The Read/Write ASIC integrates an Advanced Partial Response Maximum Likelihood (PRML) processor, a selectable code rate Encoder-Decoder (ENDEC), and a Servo Processor with data rates up to 270 MHz (259 Mb/s). Programming is done through a fast 40 MHz serial interface. The controller and data interface through an 8-bit wide data interface. The Read/Write ASIC is a low power 3.3 Volts, single supply, with selective power down capabilities.

The Read/Write ASIC comprises 12 main functional modules (described below):

- Pre-Compensator
- Variable Gain Amplifier (VGA)
- Butterworth Filter
- FIR Filter
- Flash A/D Converter
- Viterbi Detector
- ENDEC
- Servo Processor

- Clock Synthesizer
- PLL
- Serial Interface
- TA Detection and Correction

5.2.2.1 Pre-Compensator

The pre-compensator introduces pre-compensation to the write data received from the sequencer module in the DCIIA. The pre-compensated data is then passed to the R/W Pre-Amplifier and written to the disk. Pre-compensation reduces the write interference from adjacent write bit.

5.2.2.2 Variable Gain Amplifier (VGA)

Digital and analog controlled AGC function with input attenuator for extended range.

5.2.2.3 Butterworth Filter

Continuous time data filter which can be programmed for each zone rate.

5.2.2.4 FIR (Finite Impulse Response) Filter

Digitally controlled and programmable filter for partial response signal conditioning.

5.2.2.5 Flash A/D Converter

Provides very high speed digitization of the processed read signal.

5.2.2.6 Viterbi Detector

Decodes ADC result into binary bit stream.

5.2.2.7 ENDEC

Provides 16/17 or 24/25 code conversion to NRZ. Includes preamble and sync mark generation and detection.

5.2.2.8 Servo Processor

Servo processor with servo data recovery and burst demodulation.

5.2.2.9 Clock Synthesizer

Provides programmable frequencies for each zone data rate.

5.2.2.10 PLL

Provides digital read clock recovery.

5.2.2.11 Serial Interface

High speed interface for digital control of all internal blocks.

5.2.2.12 TA Detection and Correction

Detects thermal asperities' defective sectors and enables thermal asperity recoveries.

5.2.3 PreAmplifier and Write Driver

The PreAmplifier and Write Driver provides write driver and read pre-amplifier functions, and R/W head selection. The write driver receives precompensated write data from the PreCompensator module in the Read/Write ASIC. The write driver then sends this data to the heads in the form of a corresponding alternating current. The read pre-amplifier amplifies the low-amplitude voltages generated by the R/W heads, and transmits them to the VGA module in the Read/Write ASIC. Head select is determined by the controller. The preamp also contains internal compensation for thermal asperity induced amplitude variation.

5.3 FIRMWARE FEATURES

This section describes the following firmware features:

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

5.3.1 Disk Caching

The Quantum Fireball lct20 AT hard disk drives incorporate DisCache, a 380 K (minimum) disk cache, to enhance drive performance. This integrated feature is user-programmable and can significantly improve system throughput. Read and write caching can be enabled or disabled by using the Set Configuration command.

5.3.1.1 Adaptive Caching

The cache buffer for the Quantum Fireball lct20 drives features adaptive segmentation for more efficient use of the buffer's RAM. With this feature, the buffer space used for read and write operations is dynamically allocated. The cache can be flexibly divided into several segments under program control. Each segment contains one cache entry.

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

5.3.1.2 Read Cache

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

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

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

The cache memory consists of a 369 K (approximately) DRAM buffer allocated to hold the data, which can be directly accessed by the host by means of the READ and WRITE commands. The memory functions as a group of segments with rollover points at the end of cache memory. The unit of data stored is the logical block (that is, a multiple of the 512 byte sector). Therefore, all accesses to the cache memory must be in multiples of the sector size. All non-read/write commands force emptying of the cache:

5.3.1.3 Write Cache

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

While the host is preparing data for the next transfer, the drive immediately writes the cached data to the disk, usually completing the operation in less than 20 ms after issuing GOOD STATUS. With WriteCache, a single-block random write, for example, requires only about 3 ms of host time. Without WriteCache, the same operation would occupy the host for about 20 ms.

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

The WriteCache algorithm fills the cache buffer with new data from the host while simultaneously transferring data to the disk that the host previously stored in the cache.

5.3.1.4 Performance Benefits

In a drive without DisCache, there is a delay during sequential reads because of the rotational latency, even if the disk actuator already is positioned at the desired cylinder. DisCache eliminates this rotational latency time (5.56 ms 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 might request data sequentially, the disk drive must share time among all these processes. In most disk drives, the heads must move from one location to another. With DisCache, even if another process interrupts, the drive continues to access the data sequentially from its high-speed memory. In handling multiple processes, DisCache achieves its most impressive performance gains, saving both seek and latency time when desired data resides in the cache.

The cache can be flexibly divided into several segments under program control. Each segment contains one cache entry. A cache entry consists of the requested read data plus its corresponding prefetch data.

The requested read data takes up a certain amount of space in the cache segment. Hence, the corresponding prefetch data can essentially occupy the rest of the space within the segment. The other factors determining prefetch size are the maximum and minimum prefetch. The drive's prefetch algorithm dynamically controls the actual prefetch value based on the current demand, with the consideration of overhead to subsequent commands.

5.3.2 Head and Cylinder Skewing

Head and cylinder skewing in the Quantum Fireball lct20 AT hard disk drives minimize latency time and thus increases data throughput.

5.3.2.1 Head Skewing

Head skewing reduces the latency time that results when the drive must switch read/write heads to access sequential data. A head 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 Quantum Fireball lct20 drives, the sector addresses can be optimally positioned across track boundaries to minimize the latency time during a head switch. See Table 5-2.

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 Quantum Fireball lct20 drives, the sector addresses can be optimally positioned across cylinder boundaries to minimize the latency time associated with a single-cylinder seek. See Table 5-2.

5.3.2.3 Skewing with ID-less

In the ID-less environment, the drive's track and cylinder skewing will be based in unit of wedges instead of the traditional sectors. The integrated μprocessor, disk controller and ATA interface contains a "Wedge Skew Register" to assist in the task of skewing, where the skew offset must now be calculated with every read/write operation. The firmware will program the skew offset into this register every time the drive goes to a new track. The integrated μprocessor, disk controller and ATA interface will then add this value to the wedge number in the ID calculator, effectively relocating the "first" sector of the track away from the index. For example, if without skew, sector 0 is to be found following wedge 0, then if the skew register is set to 10, sector 0 will be found following wedge 10.

Since the wedge-to-wedge time is constant over the entire disk, a single set of head and cylinder skew off-sets will fulfill the requirement for all recording zones.

5.3.2.4 Skew Offsets

Table 5-2 *Skew Offsets*

	TYPICAL SWITCH TIME	WEDGE OFFSET
Head Skew	3.3 ms	41
Cylinder Skew	3.3 ms	41

Note: The switch time is calculated on the basis of the normal wedge-to-wedge time with the spindle speed variation accounted for. Wedge offsets are rounded to the closest whole number.

5.3.2.5 Runtime Calculation

Since the wedge-to-wedge time is constant over the entire disk, a single set of head and cylinder skew offsets will fulfill the requirement for all recording zones. The formula used to compute the wedge skew for a given cylinder and head is:

Wedge skew = [C* ((# of heads – 1) * TS + CS) + H * TS] MOD 180

Where:C = Cylinder number

H = Head number

TS = Head Skew Offset

CS Cylinder Skew Offset

(wedges/track = 180)

5.3.3 Error Detection and Correction

As disk drive areal densities increase, obtaining extremely low error rates requires a new generation of sophisticated error correction codes. Quantum Fireball lct20 hard disk drive series implement 288-bit quadruple-burst Reed-Solomon error correction techniques to reduce the uncorrectable read block error rate to less than one bit in 1 x 10^{14} bits read.

When errors occur, an automatic retry, a double-burst, and a more rigorous quadruple-burst correction algorithm enable the correction of any sector with four bursts of four incorrect bytes each, or up to sixteen multiple random one-byte burst errors. In addition to these advanced error correction capabilities, the drive uses an additional cross-checking code and algorithm to double check the main ECC correction. This greatly reduces the probability of a miscorrection.

5.3.3.1 Background Information on Error Correction Code and ECC On-the-Fly

A sector on the Quantum Fireball lct20 AT drive is comprised of 512 bytes of user data, followed by four cross-checking (XC) bytes (32 bits), followed by 32 ECC check bytes (256 bits). The four cross-checking bytes are used to double check the main ECC correction and reduce the probability of miscorrection. Errors of up to 64 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 on 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.

When errors cannot be corrected on-the-fly, an automatic retry, and a more rigorous quadruple-burst error correction algorithm enables the correction of any sector with four bursts of four incorrect bytes each (up to 16 contiguous bytes), or up to 16 multiple random one-byte burst errors. 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 decreases the likelihood of miscorrection.

The 32 ECC check bytes shown in Figure 5-4 are used to detect and correct errors. The cross-checking and ECC data is computed and appended to the user data when the sector is first written.

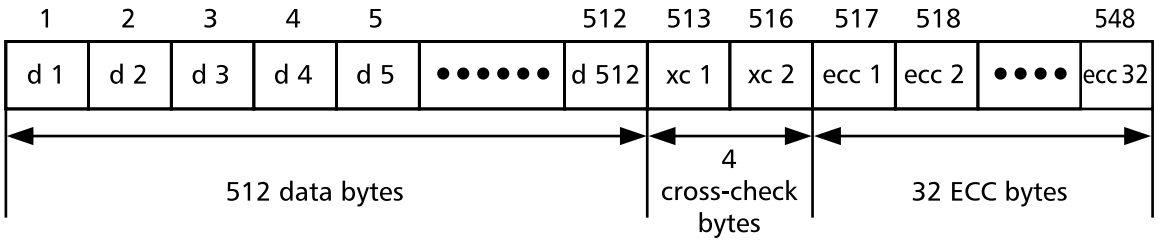

Figure 5-4 *Sector Data Field with ECC Check Bytes*

To obtain the ECC check byte values, each byte (including cross-checking and ECC bytes) within the sector is interleaved into one of four groups, where the first byte is in interleave 1, the second byte is in interleave 2, the third byte is in interleave 3, the fourth byte is in interleave 4, the fifth byte is in interleave 1, and so on, as shown in Figure 5-5.

Interleave 1 ⇨	d1	d5	••••	d509	xc1	ecc1	ecc5	ecc9	ecc13	ecc17	ecc21
Interleave 2 ⇨	d2	d6	••••	d510	xc2	ecc2	ecc6	ecc10	ecc14	ecc18	ecc22
Interleave 3 ⇨	d3	d7	••••	d511	xc3	ecc3	ecc7	ecc11	ecc15	ecc19	ecc23
Interleave 4 ⇨	d4	d8	••••	d512	xc4	ecc4	ecc8	ecc12	ecc16	ecc20	ecc24

Figure 5-5 *Byte Interleaving*

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

Each of the four interleaves is encoded with 8 ECC bytes, resulting in the 32 ECC bytes at the end of the sector. The four cross checking bytes are derived from all 512 data bytes. The combination of the interleaving, and the nature of the ECC formulas enable the drive to know where the error occurs.

Because the ECC check bytes follow the cross checking bytes, errors found within the cross-checking bytes can be corrected. Due to the power and sophistication of the code, errors found within the ECC check bytes can also be corrected.

Each time a sector of data is read, the Quantum Fireball lct20 drives will generate a new set of ECC check bytes and cross-checking bytes from the user data. These

new check bytes are compared to the ones originally written to the disk. The difference between the newly computed and original check bytes is reflected in a set of 32 *syndromes* and three cross checking syndromes, which correspond to the number of check bytes. If all the ECC syndrome values equal zero, and cx syndrome value equals 0 or 0FF, the data was read with no errors, and the sector is transferred to the host system. If any of the syndromes do not equal zero, an error has occurred. The type of correction the drive applies depends on the nature and the extent of the error.

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

Correction of Single-, or Double-Burst Errors On-the-Fly

Single-burst errors may have up to four erroneous bytes (32 bits) within a sector, provided that each of the four bytes occur in a different interleave.

The Quantum Fireball lct20 drives have the capability to correct double-burst errors on-the-fly as well. Double-burst errors can be simply viewed as two spans of errors within one sector. More specifically, correctable double-burst errors must have two or fewer erroneous bytes per interleave.

The drive's Reed-Solomon ECC corrects double-burst errors up to 64 bits long, (provided that the error consists of two or fewer bytes residing in each of the interleaves).

Double-Burst Error Examples

In the example shown in Figure 5-6 C, the 58-bit error is uncorrectable since it occupies more than two erroneous bytes per interleave.

The other two 64-bit errors, shown in Figure 5-6 A and B, are correctable because no more than two error bytes of the entire error reside in any one of the interleaves.

Note: Any 57-bit error burst can be corrected on-the-fly using double-burst error correction because no more than two bytes can occupy each interleave.

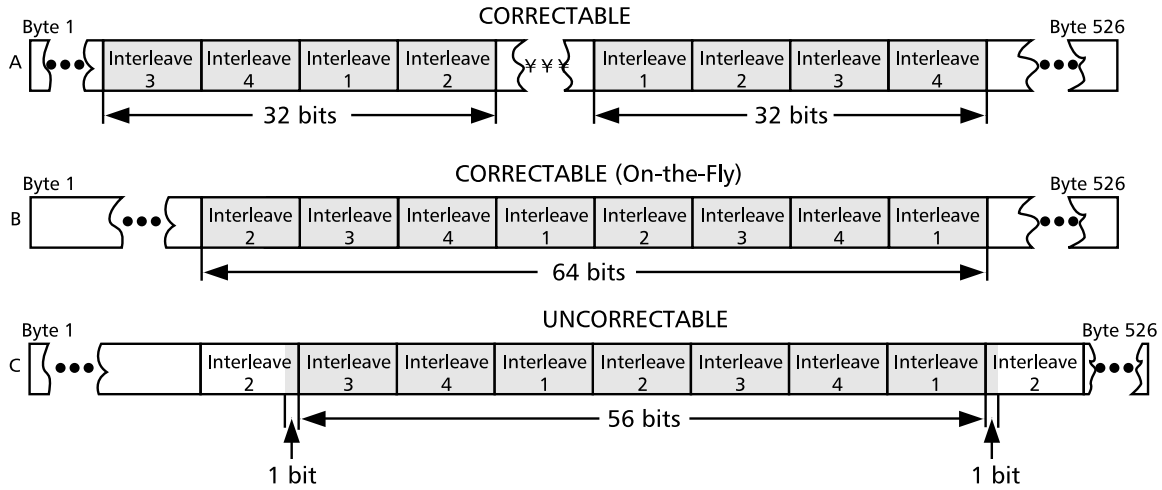

Figure 5-6 *Correctable and Uncorrectable Double-Burst Errors*

Correction of Quadruple-Burst Errors

Through sophisticated algorithms, Quantum Fireball lct20 AT drives have the capability to correct quadruple-burst errors, even though the probability of their occurrence is low. Quadruple-burst errors can be simply viewed as four spans of errors within one sector. More specifically, correctable quadruple-burst errors must have four or fewer erroneous bytes per interleave, and will not be corrected on-the-fly.

The drive's Reed-Solomon ECC corrects quadruple-burst errors up to 128 bits long, (provided that the error consists of four or fewer bytes residing in each of the interleaves).

If the quadruple-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.

Quadruple-Burst Error Examples

In the example shown in Figure 5-7 C, the 120-bit error is uncorrectable since it occupies more than four erroneous bytes per interleave.

The other two 128-bit errors, shown in Figure 5-7 A and B, are correctable because no more than four error bytes of the entire error reside in any one of the interleaves.

Note: Any 121-bit error burst can be corrected using quadruple-burst error correction because no more than four bytes can occupy each interleave.

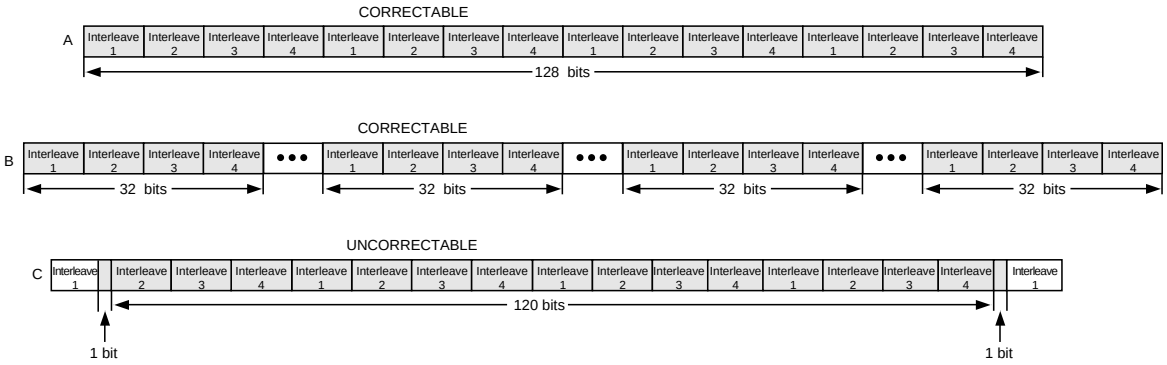

Figure 5-7 *Correctable and Uncorrectable Quadruple-Burst Errors*

Multiple Random Burst Errors

The drive's ECC can correct up to 128 bits of multiple random errors, provided that the incorrect bytes follow the guidelines for correctable quadruple-burst errors. Up to 64 bits of multiple random errors can be corrected on-the-fly, provided that the incorrect bytes follow the guidelines for correctable double-burst errors. Up to 64 bits of multiple random errors can be corrected on-the-fly if two bytes per interleave contains an error. If more than four bytes in any one interleave are in error, the sector cannot be corrected. Figure 5-8 shows an example of a correctable random burst error consisting of 16 bytes (128 bits). This random burst error is correctable because no more than four bytes within each interleave are in error.

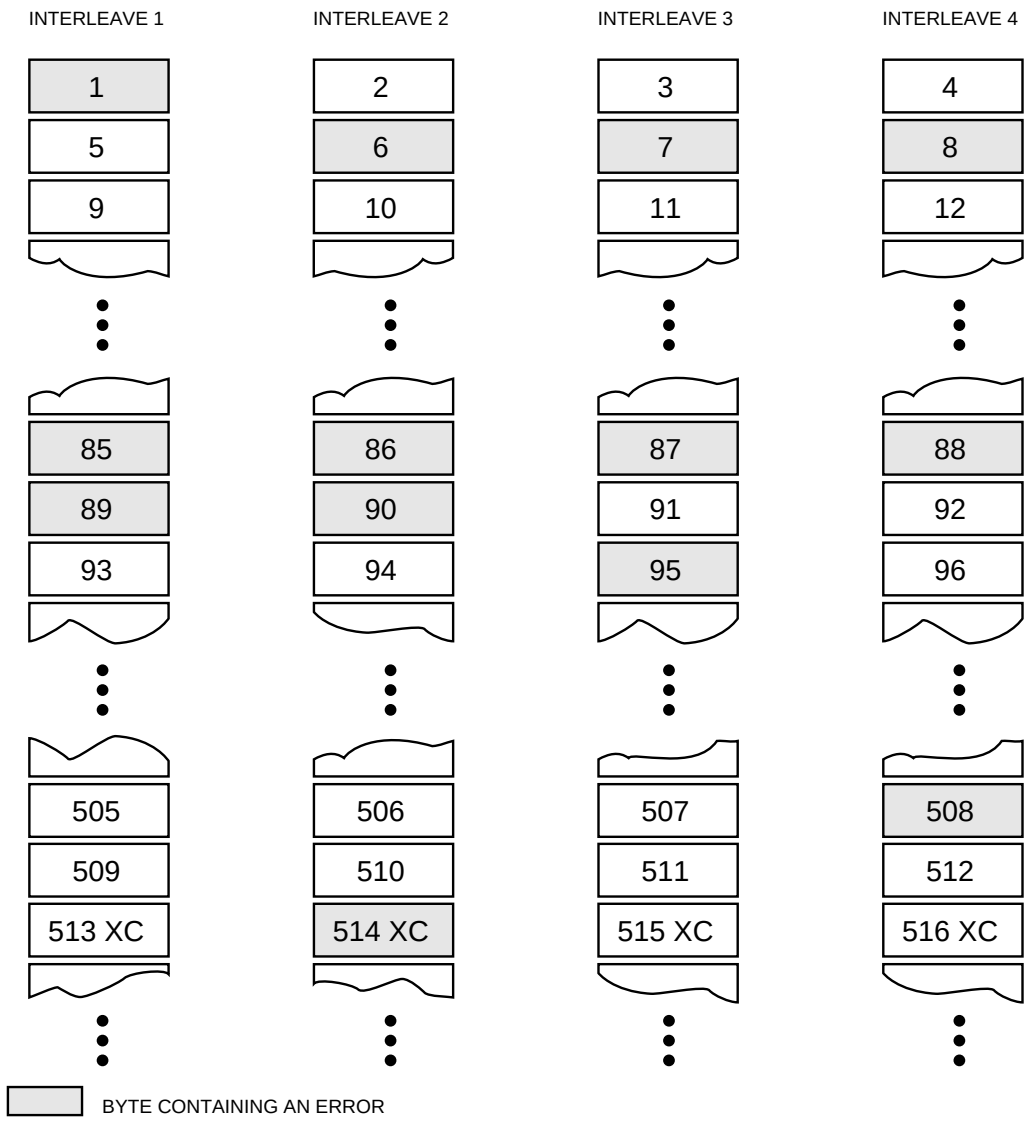

Figure 5-8 *Twelve Correctable Random Burst Errors*

5.3.3.2 ECC Error Handling

When a data error occurs, the Quantum Fireball lct20 hard disk drives check to see if the error is correctable using hardware on-the-fly ECC correction. If the error is correctable using hardware on-the-fly ECC correction, the error is corrected via hardware and the data is transferred to the host system. The hardware on-the-fly ECC correction process is very quick and does not stop the disk transfer.

If the data is not correctable using hardware on-the-fly ECC correction, the sector is re-read in an attempt to read the data correctly without applying firmware triple, or quadruple-burst ECC correction. Before invoking the firmware triple, or quadruple-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 non-repeatable errors. Each time a sector in error is re-read a set of ECC syndromes is computed. If all of the ECC syndrome values equal zero, and XC syndrome value equals to 0 or 0FF, the data was read with no errors, and the sector is transferred to the host system. If any of the syndrome values do not equal zero, an error has occurred, the syndrome values are retained, and another re-read is invoked. If the retry algorithm reaches the last step of the retry table, or the last pass through the retry table, and the sector still has a data error, then the drive will apply firmware triple, or quadruple-burst ECC correction to recover the data.

Note: Non-repeatable errors are usually related to the signal to noise ratio of the system. They are not due to media defects.

If the automatic read reallocation feature is enabled, and the drive used firmware triple, or quadruple-burst ECC correction to recover the data, then the drive will perform a media test on the sector to determine if the error was caused by a grown media defect. The media test consists of writing and reading the sector four times. If any one of the writes or reads fails, or if any read requires firmware quadruple-burst ECC correction to recover the data, then the sector is automatically reassigned.

Note: The Quantum Fireball lct20 AT drives are shipped from the factory with the automatic read reallocation feature enabled so that any new defective sectors can be easily and automatically reallocated for the average AT end user.

5.3.3.3 Defect Management

The Quantum Fireball lct20 drives use two techniques for replacing defective sectors - inline replacement and offline replacement.

During manufacturing, if a sector on a cylinder is found to be defective, the address of the sector is added to the drive's defect list. The defective sector is "skipped," and is replaced by the next immediate sector, which maintains a sequential ordering of logical blocks. This inline sparing technique is employed in an attempt to eliminate slow data transfer that would result from requiring a seek to another cylinder to access a replacement sector.

The Quantum Fireball lct20 drives are divided into 30 logical zones for purposes of defect management. Each zone has a pool of 32 spare sectors, which are reserved for offline replacement.

Defects that occur in the field are known as grown defects. If such a defective sector is found in the field, the sector is reallocated to a spare sector from the nearest available pool of spares. The defect list supports a maximum of 500 grown defects.

Sectors are considered to contain grown defects if the quadruple-burst ECC algorithm must be applied to recover the data. If this algorithm is successful, the corrected data is stored in the newly allocated sector. If the algorithm is not successful, a pending defect will be added to the defect list. Any subsequent read to the original logical block will return an error if the read is not successful. A host command to over-write the location will result in several write/read/verifies of the suspect location. If any of these operations fail, the new data will be written to a spare sector, and the original location will be added to the permanent defect list. If all of the operations are successful, data will be written to the location, and the pending defect will be removed from the list.

Chapter 6
ATA BUS INTERFACE AND ATA COMMANDS

This chapter describes the interface between Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives and the ATA bus. The commands that are issued from the host to control the drive are listed, as well as the electrical and mechanical characteristics of the interface.

6.1 INTRODUCTION

Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives use the standard IBM PC ATA bus interface, and are compatible with systems that provide an ATA interface connector on the motherboard. It may also be used with a third-party adapter board in systems that do not have a built-in ATA adapter. The adapter board plugs into a standard 16-bit expansion slot in an AT-compatible computer. A cable connects the drive to the adapter board. On newer systems a PCI 32-bit card may be allowed.

6.2 SOFTWARE INTERFACE

The Quantum Fireball lct20 AT drives are controlled by the Basic Input/Output System (BIOS) program residing in an IBM PC AT, or IBM compatible PC. The BIOS communicates directly with the drive's built-in controller. It issues commands to the drive and receives status information from the drive.

6.3 MECHANICAL DESCRIPTION

6.3.1 Drive Cable and Connector

The hard disk drive connects to the host computer by means of a cable. This cable has a 40-pin connector that plugs into the drive, and a 40-pin connector that plugs into the host computer. At the host end, the cable plugs into either an adapter board residing in a host expansion slot, or an on-board ATA adapter.

6.4 ELECTRICAL INTERFACE

6.4.1 ATA Bus Interface

A 40-pin ATA interface connector on the motherboard or an adapter board provides an interface between the drive and a host that uses an IBM PC AT bus. The ATA interface contains bus drivers and receivers compatible with the standard AT bus. The AT-bus interface signals D8–D15, INTRQ, and IOCS16– require the ATA adapter board to have an extended I/O-bus connector.

The ATA interface buffers data and controls signals between the drive and the AT bus of the host system, and decodes addresses on the host address bus. The Command Block Registers on the drive accept commands from the host system BIOS.

Note: Some host systems do not read the Status Register after the drive issues an interrupt. In such cases, the interrupt may not be acknowledged. To overcome this problem, you may have to configure a jumper on the motherboard or adapter board to allow interrupts to be controlled by the drive's interrupt logic. Read your motherboard or adapter board manual carefully to find out how to do this.

6.4.1.1 Electrical Characteristics

All signals are transistor-transistor logic (TTL) compatible—with logic 1 greater than 2.0 volts and less than 5.25 volts; and logic 0 greater than 0.0 volts and less than 0.8 volts.

6.4.1.2 Drive Signals

The drive connector (J1, section C) connects the drive to an adapter board or onboard ATA adapter in the host computer. J1, section C is a 40-pin shrouded connector with two rows of 20 pins on 100-mil centers. J1 has been keyed by removing pin 20. The connecting cable is a 40-conductor (80-conductor for UDMA modes 3 and 4 operation) flat ribbon cable with a maximum length of 18 inches.

Note: Table 6-1 describes the signals on the drive connector (J1, section C). The drive does not use all of the signals provided by the ATA bus. Table 6-4 shows the relationship between the drive connector (J1, section C) and the ATA bus.

Note: In Table 6-1 the following conventions apply:

A minus sign follows the name of any signal that is asserted as active low.
Direction (DIR) is in reference to the drive.
IN indicates input to the drive.
OUT indicates output from the drive.
I/O indicates that the signal is bidirectional.

Table 6-1 *Drive Connector Pin Assignments (J1, Section C)*

SIGNAL	NAME	DIR	PIN	DESCRIPTION
Reset	RESET–	IN	1	Drive reset signal from the host system, inverted on the adapter board or motherboard. This signal from the host system will be asserted beginning with the application of power, and held asserted until at least 25 µs after voltage levels have stabilized within tolerance during power on. It will be negated thereafter unless some event requires that the device(s) be reset following power on. ATA devices will not recognize a signal assertion shorter than 20 ns as a valid reset signal. Devices may respond to any signal assertion greater than 20 ns, and will recognize a signal equal to or greater than 25 µs. The drive has a 10kΩ pull-up resistor on this signal.
Ground	Ground	—	2	Ground between the host system and the drive.
Data Bus		I/O	3–18	An 8/16-bit, bidirectional data bus between the host and the drive. D0–D7 are used for 8-bit transfers, such as registers and ECC bytes.

Table 6-1 *Drive Connector Pin Assignments (J1, Section C) (Continued)*

SIGNAL	NAME	DIR	PIN	DESCRIPTION
	DD0		17	Bit 0
	DD1		15	Bit 1
	DD2		13	Bit 2
	DD3		11	Bit 3
	DD4		9	Bit 4
	DD5		7	Bit 5
	DD6		5	Bit 6
	DD7		3	Bit 7
	DD8		4	Bit 8
	DD9		6	Bit 9
	DD10		8	Bit 10
	DD11		10	Bit 11
	DD12		12	Bit 12
	DD13		14	Bit 13
	DD14		16	Bit 14
	DD15		18	Bit 15
Ground	Ground	—	19	Ground between the host system and the drive.
Keypin	KEYPIN	—	20	Pin removed to key the interface connector.
DMA Request	DMARQ	OUT	21	Asserted by the drive when it is ready to exchange data with the host. The direction of the data transfer is determined by DIOW– and DIOR–. DMARQ is used in conjunction with DMACK–. The drive has a 10kΩ pull-down resistor on this signal.
Ground	Ground	—	22	Ground between the host system and the drive.
I/O Write	DIOW–	IN	23	The rising edge of this write strobe provides a clock for data transfers from the host data bus (DD0–DD7 or DD0–DD15) to a register or to the drive's data port.
Ground	Ground	—	24	Ground between the host system and the drive.
I/O Read	DIOR–	IN	25	The rising edge of this read strobe provides a clock for data transfers from a register or the drive's data port to the host data bus (DD0–DD7 or DD0–DD15). The rising edge of DIOR– latches data at the host.
Ground	Ground	—	26	Ground between the host system and the drive.
I/O Channel Ready	IORDY	OUT	27	When the drive is not ready to respond to a data transfer request, the IORDY signal is asserted active low to extend the host transfer cycle of any host register read or write access. When IORDY is deasserted, it is in a high-impedance state and it is the host's responsibility to pull this signal up to a high level (if necessary).

Table 6-1 *Drive Connector Pin Assignments (J1, Section C) (Continued)*

SIGNAL	NAME	DIR	PIN	DESCRIPTION
Cable Select		—	28	This is a signal from the host that allows the drive to be configured as drive 0 when the signal is 0 (grounded), and as drive 1 when the signal is 1 (high). The drive has a 10kΩ pull-up resistor on this signal.
DMA Acknowledge	DACK1–	IN	29	Used by the host to respond to the drive's DMARQ signal. DMARQ signals that there is more data available for the host.
Ground	Ground	—	30	Ground between the host system and the drive.
Interrupt Request	INTRQ	OUT	31	An interrupt to the host system. Asserted only when the drive microprocessor has a pending interrupt, the drive is selected, and the host clears nIEN in the Device Control Register. When nIEN is a 1 or the drive is not selected, this output signal is in a high-impedance state, regardless of the presence or absence of a pending interrupt. INTRQ is deasserted by an assertion of RESET–, the setting of SRST in the Device Control Register, or when the host writes to the Command Register or reads the Status Register. When data is being transferred in programmed I/O (PIO) mode, INTRQ is asserted at the beginning of each data block transfer. Exception: INTRQ is not asserted at the beginning of the first data block transfer that occurs when any of the following commands executes: FORMAT TRACK, Write Sector, WRITE BUFFER, or WRITE LONG.
16-Bit I/O	IOCS16–	OUT	32	An open-collector output signal. Indicates to the host system that the 16-bit data port has been addressed, and that the drive is ready to send or receive a 16-bit word. When transferring data in PIO mode, if IOCS16– is not asserted, D0–D7 are used for 8-bit transfers; if IOCS16– is asserted, D0–D15 are used for 16-bit data transfers.
Drive Address Bus				A 3-bit, binary-coded address supplied by the host when accessing a register or the drive's data port.
Bit 1	DA1	IN	33	
Bit 0	DA0	IN	35	
Bit 2	DA2	IN	36	

Table 6-1 *Drive Connector Pin Assignments (J1, Section C) (Continued)*

SIGNAL	NAME	DIR	PIN	DESCRIPTION
Passed Diagnostics	PDIAG–	I/O	34	Drive 0 (Master) monitors this Drive 1 (Slave) open-collector output signal, which indicates the result of a diagnostics command or reset. The drive has a 10K pull-up resistor on this signal. Following the receipt of a power-on reset, software reset, or RESET– drive 1 negates PDIAG– within 1 ms. PDIAG– indicates to drive 0 that drive 1 is busy (BSY=1). Then, drive 1 asserts PDIAG– within 30 seconds, indicating that drive 1 is no longer busy (BSY=0) and can provide status information. Following the assertion of PDIAG–, drive 1 is unable to accept commands until drive 1 is ready (DRDY=1)—that is, until the reset procedure for drive 1 is complete. Following the receipt of a valid EXECUTE DRIVE DIAGNOSTIC command, drive 1 negates PDIAG– within 1 ms, indicating to drive 0 that it is busy and has not yet passed its internal diagnostics. If drive 1 is present, drive 0 waits for drive 1 to assert PDIAG– for up to 5 seconds after the receipt of a valid EXECUTE DRIVE DIAGNOSTIC command. Since PDIAG– indicates that drive 1 has passed its internal diagnostics and is ready to provide status, drive 1 clears BSY prior to asserting PDIAG–. If drive 1 fails to respond during reset initialization, drive 0 reports its own status after completing its internal diagnostics. Drive 0 is unable to accept commands until drive 0 is ready (DRDY=1)—that is, until the reset procedure for drive 0 is complete.
Chip Select 0	CS1FX–	IN	37	Chip-select signal decoded from the host address bus. Used to select the host-accessible Command Block Registers.
Chip Select 1	CS3FX–	IN	38	Chip select signal decoded from the host address bus. Used to select the host-accessible Control Block Registers.

Table 6-1 *Drive Connector Pin Assignments (J1, Section C) (Continued)*

SIGNAL	NAME	DIR	PIN	DESCRIPTION
Drive Active/Slave Present	DASP–	I/O	39	A time-multiplexed signal that indicates either drive activity or that drive 1 is present. During power-on initialization, DASP– is asserted by drive 1 within 400 ms to indicate that drive 1 is present. If drive 1 is not present, drive 0 asserts DASP– after 450 ms to light the drive-activity LED. An open-collector output signal, DASP– is deasserted following the receipt of a valid command by drive 1 or after the drive is ready, whichever occurs first. Once DASP– is deasserted, either hard drive can assert DASP– to light the drive-activity LED. Each drive has a 10K pull-up resistor on this signal. If an external drive-activity LED is used to monitor this signal, an external resistor must be connected in series between the signal and a +5 volt supply in order to limit the current to 24 mA maximum.
Ground	Ground	—	40	Ground between the host system and the drive.

Series termination resistors are required at both the host and the device for operation in any of the Ultra ATA/66 modes. Table 6-2 describes recommended values for series termination at the host and the device.

Table 6-2 *Series Termination for Ultra ATA/66*

SIGNAL	HOST TERMINATION	DEVICE TERMINATION
DIOR–/HDMARDY–/HSTROBE	33 Ω	82 Ω
DIOW–/STOP	33 Ω	82 Ω
CS0–, CS1–	33 Ω	82 Ω
DA0, DA1, DA2	33 Ω	82 Ω
DMACK–	33 Ω	82 Ω
DD 15 through DD0	33 Ω	33 Ω
DMARQ	82 Ω	33 Ω
INTRQ	82 Ω	33 Ω
IORDY/DDMARDY–/DSTROBE	82 Ω	33 Ω

Note: Only those signals requiring termination are listed in this table. If a signal is not listed, series termination is not required for operation in an Ultra ATA/66 mode.

6.4.1.3 Ultra ATA/66 80 Conductor Cable

The use of a 80 conductor cable is suggested in order to successfully meet the new Ultra ATA/66 mode 3 and 4 faster timing requirements. The 80 conductor cable is used with the same connector configuration as the standard 40 conductor cable. There is no new signal and the 40 additional lines are ground paths tied together to all 7 original ground conductors. Both, the host and the device (drive) may detect the type of cable being used.

Host Based Cable detection (preferred method)

This detection scheme is already defined in the ATA/ATAPI-4 document. To detect the type of cable being used the host must sample the PDIAG-/CBLID- signal. After device 0/1 handshaking and a command has been sent to device 1 to cause it to release the PDIAG- signal, the host detects the state of CBLID-.

Host detects CBLID- below V_{il} = 80 Conductor

Host detects CBLID- above V_{ih} = 40 Conductor

Device Based Cable detection

Following the issuing of an ID command by the host the device will respond by:

- Asserts PDIAG-/CBLID- (drives it low) for 30 ms minimum.
- Releases PDIAG-/CBLID-
- Measures level of PDIAG-/CBLID- 2 to 13 ms after releasing it.

The detected electrical level of PDIAG-/CBLID- will be stored in ID word 93 bit 13 (refer to Table 6-24).

- PDIAG-/CBLID- less than V_{il} = 0
- PDIAG-/CBLID- greater than V_{ih} = 1

The host can then read ID data and use information in word 93 only if the device supports Ultra DMA modes higher than 2, otherwise, it shall be ignored.

- 0 = 40 Conductor if both device and host support method
- 1 = 80 Conductor if both device and host support method.

6.4.1.4 ATA Bus Signals

See Table 6-4 for the relationship between the drive signals and the ATA bus.

Signal Line Definitions (Ultra ATA/66)

Several existing ATA signal lines are redefined during the Ultra ATA/66 protocol to provide new functions. These lines change from old to new definitions the moment the host decides to allow a DMA burst, if the Ultra ATA/66 transfer mode was previously chosen via Set Features. The drive becomes aware of this change upon assertion of the –DMACK line. These lines revert back to their original definitions upon the deassertion of –DMACK at the termination of the DMA burst.

Table 6-3 *Signal Line Definitions*

NEW DEFINITION	OLD DEFINITION
DMARQ	= DMARQ
–DMACK	= –DMACK
(These two signals remain unchanged to ensure backward compatibility with PIO modes)	
–DMARDY	= IORDY on WRITE commands = –DIOR on READ commands
STROBE	= –DIOR on WRITE commands = IORDY on READ commands
STOP	= –DIOW
–CBLID	–PDIAG

Table 6-4 *Interface Signal Name Assignments*

J1 PIN NUMBER	DESCRIPTION	HOST	DIR	DEV	ACRONYM
28	CABLE SELECT	—>			CSEL
37	CHIP SELECT 0	—>			CS0–
38	CHIP SELECT 1	<—>			CS1–
17	DATA BUS BIT 0	<—>			DD0
15	DATA BUS BIT 1	<—>			DD1
13	DATA BUS BIT 2	<—>			DD2
11	DATA BUS BIT 3	<—>			DD3
9	DATA BUS BIT 4	<—>			DD4
7	DATA BUS BIT 5	<—>			DD5
5	DATA BUS BIT 6	<—>			DD6
3	DATA BUS BIT 7	<—>			DD7
4	DATA BUS BIT 8	<—>			DD8
6	DATA BUS BIT 9	<—>			DD9
8	DATA BUS BIT 10	<—>			DD10
10	DATA BUS BIT 11	<—>			DD11
12	DATA BUS BIT 12	<—>			DD12
14	DATA BUS BIT 13	<—>			DD13
16	DATA BUS BIT 14	<—>			DD14
18	DATA BUS BIT 15	<—>			DD15
39	DEVICE ACTIVE OR SLAVE (DEVICE 1) PRESENT	(See Note 1)			DASP–
35	DEVICE ADDRESS BIT 0	—>			DA0
33	DEVICE ADDRESS BIT 1	—>			DA1
36	DEVICE ADDRESS BIT 2	—>			DA2
29	DMA ACKNOWLEDGE	—>			DMACK–
21	DMA REQUEST	<—			DMARQ
31	INTERRUPT REQUEST	<—			INTRQ

J1 PIN NUMBER	DESCRIPTION	HOST	DIR	DEV	ACRONYM
25	I/O READ DMA ready on data in bursts (see note 2) Data strobe on data out bursts (see note 2)	—> —> —>			DIOR– HDMARDY– HSTROBE
27	I/O READY DMA ready on data out bursts (see note 2) Data strobe on data in bursts (see note 2)	<— <— <—			IORDY DDMARDY– DSTROBE
23	I/O WRITE STOP (see note 2)	—><br—>			DIOW– STOP
34	PASSED DIAGNOSTCS/CABLE DETECTION	(See Notes 1 & 5)			PDIAG–/CBLID–
1	RESET	—>			RESET
32	I/O CS16	<—			IOCS16

Note:

1. See signal descriptions for information on source of these signals.
2. Used during Ultra DMA protocol only.
3. Pins numbered 2, 19, 22, 24, 26, 30, and 40 are ground.
4. Pin number 20 is the Key Pin.
5. CBLID– during device based cable detection on devices supporting Ultra DMA modes higher than 2.

6.4.2 Host Interface Timing

6.4.2.1 Programmed I/O (PIO) Transfer Mode

The PIO host interface timing shown in Table 6-5 is in reference to signals at 0.8 volts and 2.0 volts. All times are in nanoseconds, unless otherwise noted. Figure 6-1 provides a timing diagram.

Table 6-5 *PIO Host Interface Timing*

SYMBOL	DESCRIPTION	MIN/MAX	MODE 4[1] (local bus)	Quantum Fireball lct20 AT
t0	**Cycle Time**	**min**	**120**	**120**
t1	Address Valid to DIOW–/DIOR–Setup	min	25	25
t2	DIOW–/DIOR– Pulsewidth (8- or 16-bit)	min	70	70
t2i	DIOW–/DIOR– Negated Pulsewidth	min	25	25
t3	DIOW–Data Setup	min	20	20
t4	DIOW– Data Hold	min	10	10
t5	DIOR– Data Setup	min	20	20
t5a	DIOR– to Data Valid	max	—	—
t6	DIOR– Data Hold	min	5	5
t6z	DIOR– Data Tristate	max	30	30
t7	Address Valid to IOCS16– Assertion	max	N/A	N/A
t8	Address Valid to IOSC16– Deassertion	max	N/A	N/A
t9	DIOW–/DIOR– to Address Valid Hold	min	10	10
tA	IORDY Setup Time	min	35	35
tB	IORDY Pulse Width	max	1250	1250
tR	Read Data Valid to IORDY active (if IORDY is initially low after tA)	min	0	0

1. ATA Mode 4 timing is listed for reference only.

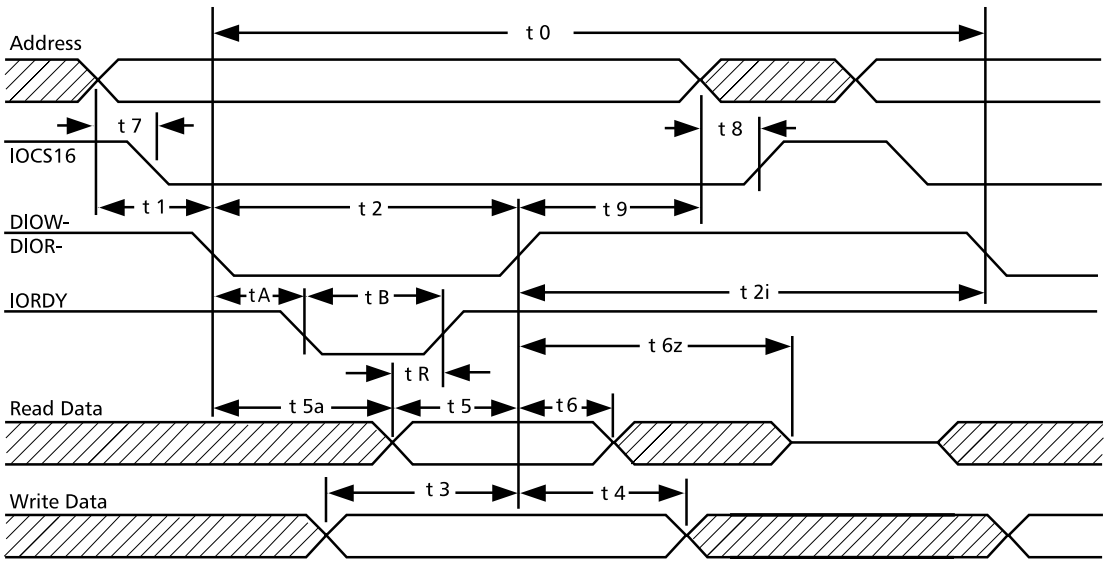

Figure 6-1 *PIO Interface Timing*

6.4.2.2 Multiword DMA Transfer Mode

The multiword DMA host interface timing shown in Table 6-6 is in reference to signals at 0.8 volts and 2.0 volts. All times are in nanoseconds, unless otherwise noted. Figure 6-2 provides a timing diagram.

Table 6-6 *Multiword DMA Host Interface Timing*

SYMBOL	DESCRIPTION	MIN/MAX	MODE 2[1] (local bus)	Quantum Fireball lct20 AT
t0	**Cycle Time**	**min**	120	120
tD	DIOR–/DIOW– Pulsewidth	min	70	70
tE	DIOR– to Data Valid	max	–	–
tF	DIOR– Data Hold	min	5	5
tFz	DIOR– Data Tristate[2]	max	20	20
tG	DIOW– Data Setup	min	20	20
tH	DIOW– Data Hold	min	10	10
tI	DMACK to DIOR–/DIOW– Setup	min	0	0
tJ	DIOR–/DIOW– to DMACK– Hold	min	5	5
tK	DIOR–/DIOW– Negated Pulsewidth	min	25	25
tL	DIOR–/DIOW– to DMARQ Delay	max	35	35
tz	DMACK– Data Tristate[3]	max	25	25

1. ATA Mode 2 timing is listed for reference only.
2. The Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT drive tristates after each word transferred.
3. Symbol tz only applies on the last tristate at the end of a multiword DMA transfer cycle.

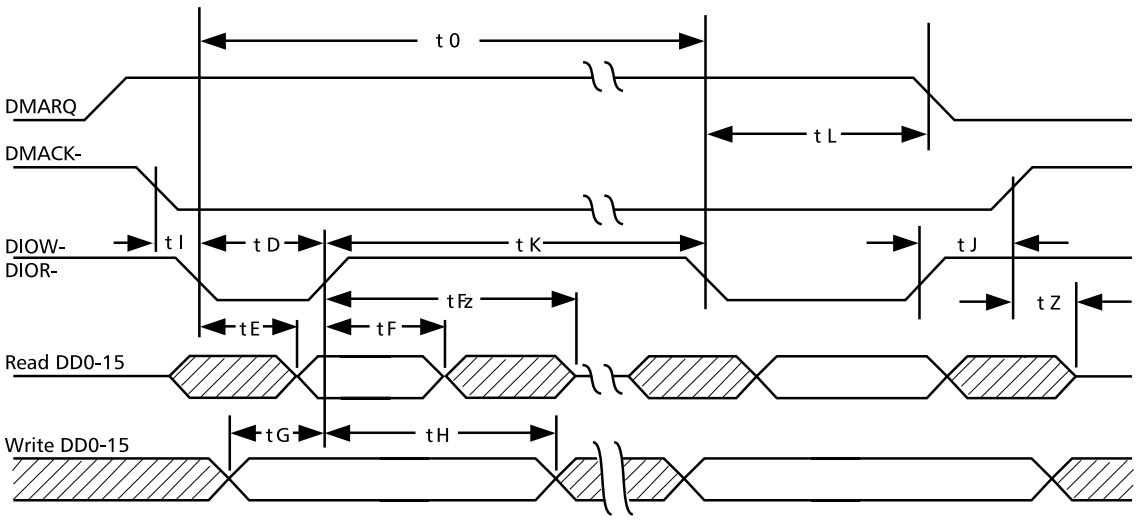

Figure 6-2 *Multiword DMA Bus Interface Timing*

Table 6-7 contains the values for the timings for each of the Ultra DMA modes. All timing measurement switching points (low to high and high to low) shall be taken at 1.5V. Table 6-8 contains descriptions and comments for each of the timing values in Table 6-7.

Table 6-7 *Ultra DMA Data Transfer Timing Requirements*

Name	Mode 0 (ns)		Mode 1 (ns)		Mode 2 (ins)		Mode 3 (ns)		Mode 4 (ns)		Mode 5 (ns)		Measurement Location
	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	
$t_{2CYCTYP}$	240		160		120		90		60		40		Sender
t_{CYC}	112		73		54		39		25		16.8		Recipient
t_{2CYC}	230		153		115		86		57		38		Sender
t_{DS}	15.0		10.0		7.0		7.0		5.0		4.0		Recipient
t_{DH}	5.0		5.0		5.0		5.0		5.0		4.6		Recipient
t_{DVS}	70.0		48.0		31.0		20.0		6.7		4.8		Sender
t_{DVH}	6.2		6.2		6.2		6.2		6.2		4.8		Sender
t_{CS}	15.0		10.0		7.0		7.0		5.0		5.0		Device
t_{CH}	5.0		5.0		5.0		5.0		5.0		5.0		Device
t_{CVS}	70.0		48.0		31.0		20.0		6.7		10.0		Host
t_{CVH}	6.2		6.2		6.2		6.2		6.2		10.0		Host
t_{ZFS}	0		0		0		0		0		35		Device
t_{DZFS}	70.0		48.0		31.0		20.0		6.7		25		Sender
t_{FS}		230		200		170		130		120		90	Device
t_{LI}	0	150	0	150	0	150	0	100	0	100	0	75	Note 2
t_{MLI}	20		20		20		20		20		20		Host
t_{UI}	0		0		0		0		0		0		Host
t_{AZ}		10		10		10		10		10		10	Note 3
t_{ZAH}	20		20		20		20		20		20		Host
t_{ZAD}	0		0		0		0		0		0		Device
t_{ENV}	20	70	20	70	20	70	20	55	20	55	20	50	Host
t_{RFS}		75		70		60		60		60		50	Sender
t_{RP}	160		125		100		100		100		85		Recipient
t_{IORDYZ}		20		20		20		20		20		20	Device
t_{ZIORDY}	0		0		0		0		0		0		Device
t_{ACK}	20		20		20		20		20		20		Host
t_{SS}	50		50		50		50		50		50		Sender

Notes:

1. All signal transitions for a timing parameter will be measured at the connector specified in the measured location column. For example, in the case of t_{RFS}, both STROBE and DMARDY- transitions are measured at the sender connector.

2. The parameter t_{LI} shall be measured at the connector of the sender or recipient that is responding to an incoming transition from the recipient or sender respectively. Both the incoming signal and the outgoing response will be measured at the same connector.

3. The parameter t_{AZ} will be measured at the connector of the sender or recipient that is releasing the bus.

Table 6-8 *Ultra DMA Data Burst Timing Descriptions*

Name	Comment
$t_{2CYCTYP}$	Typical sustained average two cycle time
t_{CYC}	Cycle time allowing for asymmetry and clock variations (from STROBE edge to STROBE edge)
t_{2CYC}	Two cycle time allowing for clock variations (from rising edge to next rising edge or from falling edge to next falling edge of STROBE)
t_{DS}	Data setup time at recipient (from data valid until STROBE edge) (see notes 2,5)
t_{DH}	Data hold time at recipient (from STROBE edge until data may become invalid) (see note 2,5)
t_{DVS}	Data valid setup time at sender (from data valid until STROBE edge) (see note 3)
t_{DVH}	Data valid hold time at sender (from STROBE edge until data may become invalid) (see note 3)
t_{CS}	CRC word setup time at device (see note 2)
t_{CH}	CRC word hold time device (see note 2)
t_{CVS}	CRC word valid setup time at host (from CRC valid until DMACK- negation) (see note 3)
t_{CVH}	CRC word valid hold time at sender (from DMACK- negation until CRC may become invalid) (see note 3)
t_{ZFS}	Time from STROBE output released-to-driving until the first transition of critical timing.
t_{DZFS}	Time from data output released-to-driving until the first transition of critical timing.
t_{FS}	First STROBE time (for device to first negate DSTROBE from STOP during a data in burst)
t_{LI}	Limited interlock time (see note 1)
t_{MLI}	Interlock time with minimum (see note 1)
t_{UI}	Unlimited interlock time (see note 1)
t_{AZ}	Maximum time allowed for output drivers to release (from asserted or negated)
t_{ZAH}	Minimum delay time required for output
t_{ZAD}	drivers to assert or negate (from released)
t_{ENV}	Envelope time (from DMACK- to STOP and HDMARDY- during data in burst initiation and from DMACK to STOP during data out burst initiation)
t_{RFS}	Ready-to-final-STROBE time (no STROBE edges shall be sent this long after negation of DMARDY-)

t_{RP}	Ready-to-pause time (that recipient shall wait to pause after negating DMARDY-)
t_{IORDYZ}	Maximum time before releasing IORDY
t_{ZIORDY}	Minimum time before driving IORDY (see note 4)
t_{ACK}	Setup and hold times for DMACK- (before assertion or negation)
t_{SS}	Time from STROBE edge to negation of DMARQ or assertion of STOP (when sender terminates a burst)

Notes:

1. The parameters t_{UI}, t_{MLI} and t_{LI} indicate sender-to-recipient or recipient-to-sender interlocks, i.e., one agent (either sender or recipient) is waiting for the other agent to respond with a signal before proceeding. t_{UI} is an unlimited interlock that has no maximum time value. t_{MLI} is a limited time-out that has a defined minimum. t_{LI} is a limited time-out that has a defined maximum.

2. 80-conductor cabling will be required in order to meet setup (t_{DS}, t_{CS}) and hold (t_{DH}, t_{CH}) times in modes greater than 2.

3. Timing for t_{DVS}, t_{DVH}, t_{CVS} and t_{CVH} will be met for lumped capacitive loads of 15 and 40 pf at the connector where the Data and STROBE signals have the same capacitive load value. Due to reflections on the cable, these timing measurements are not valid in a normally functioning system.

4. For all modes the parameter t_{ZIORDY} may be greater than t_{ENV} due to the fact that the host has a pull up on IORDY- giving it a known state when released.

5. The parameters t_{DS}, and t_{DH} for mode 5 are defined for a recipient at the end of the cable only in a configuration with one device at the end of the cable.

Figures 6-3 through 6-12 define the timings associated with all phases of Ultra DMA bursts.

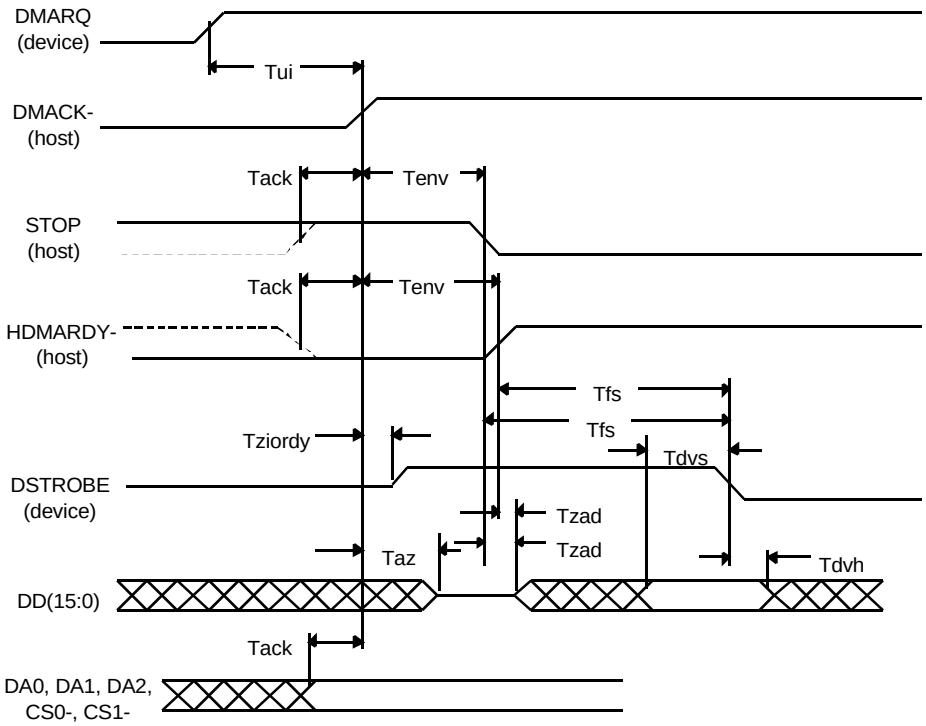

Figure 6-3 *Initiating a Data In Burst*

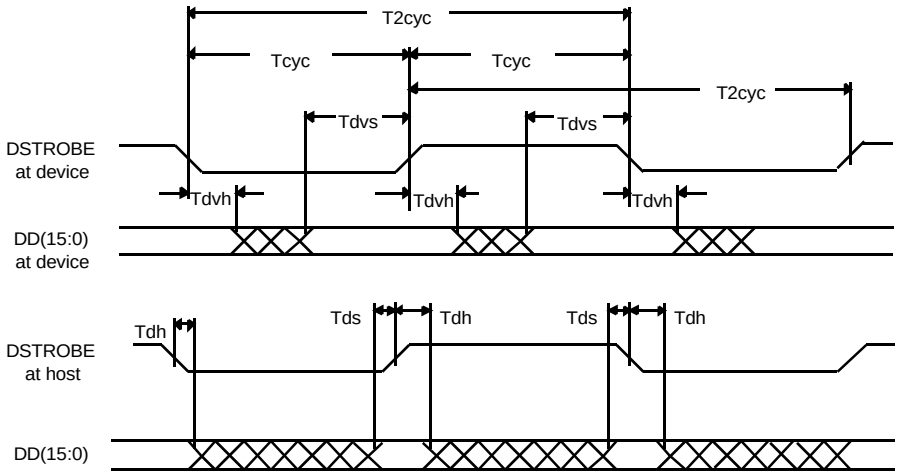

Figure 6-4 *Sustained Data In Burst*

Note: DD(15:0) and DSTROBE signals are shown at both the host and the device to emphasize that cable settling time as well as cable propagation delay shall not allow the data signals to be considered stable at the host until well after they are driven by the device.

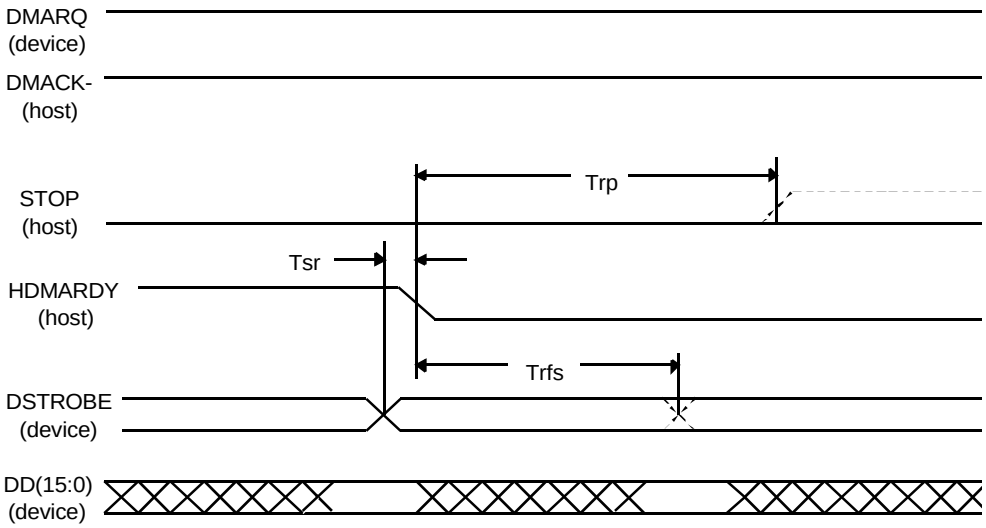

Figure 6-5 *Host Pausing a Data In Burst*

Note: The host knows the burst is fully paused Trp ns after HDMARDY- is negated and may then assert STOP to terminate the burst. Tsr timing need not be met for an asynchronous pause.

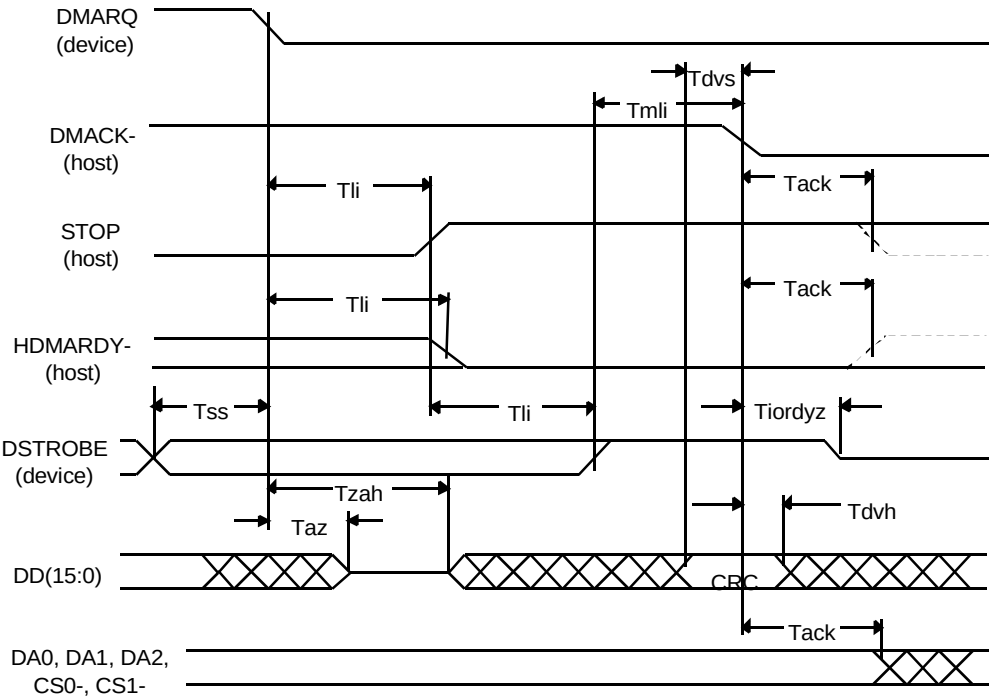

Figure 6-6 *Device Terminating a Data In Burst*

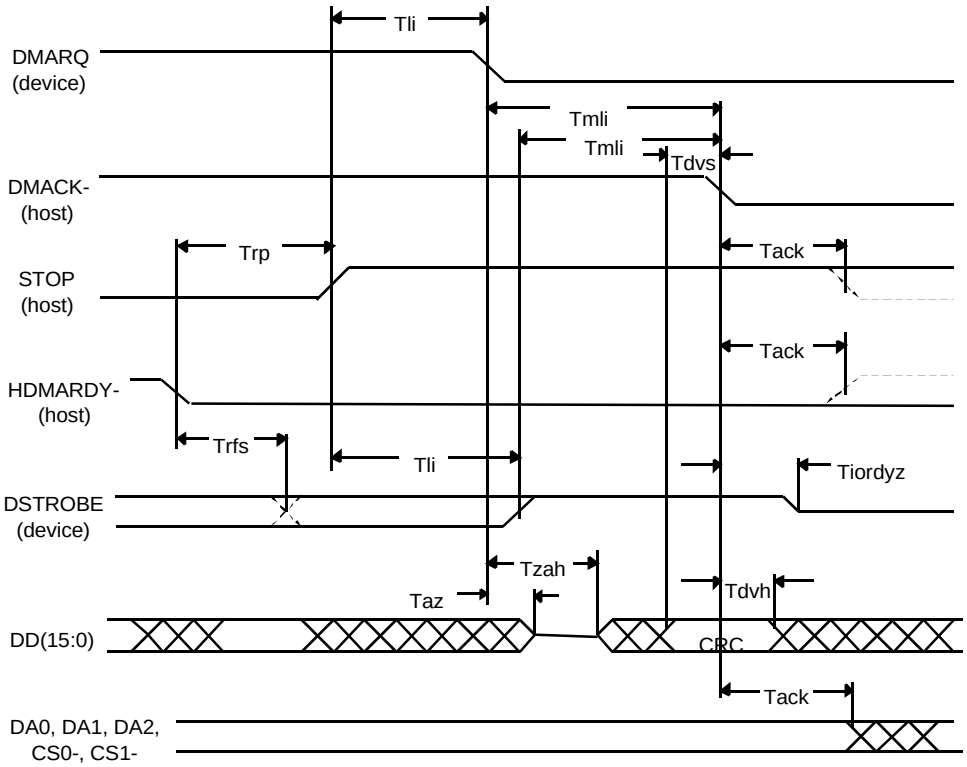

Figure 6-7 *Host Terminating a Data In Burst*

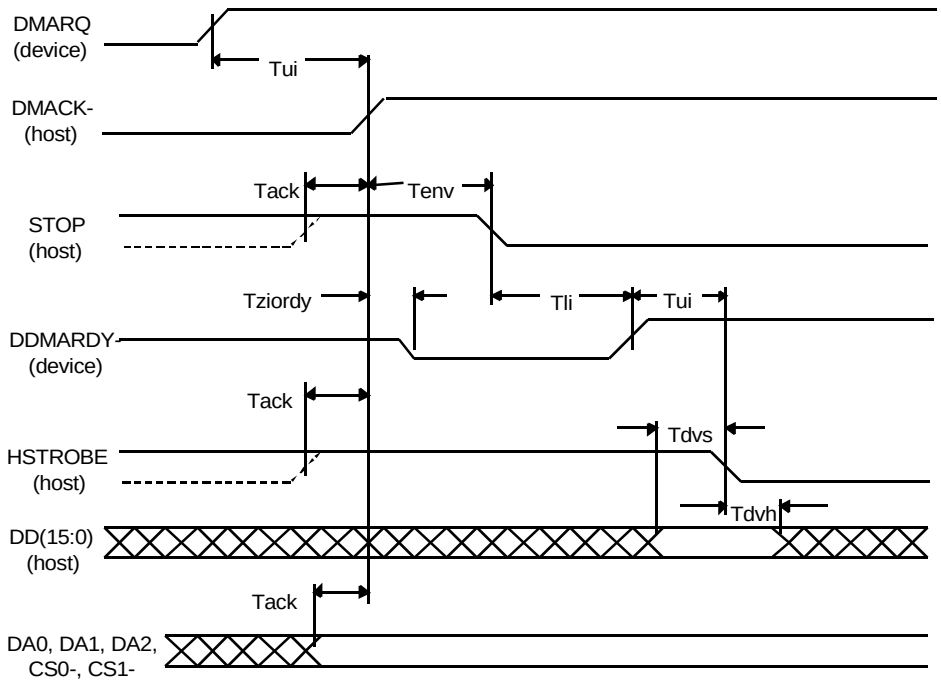

Figure 6-8 *Initiating a Data Out Burst*

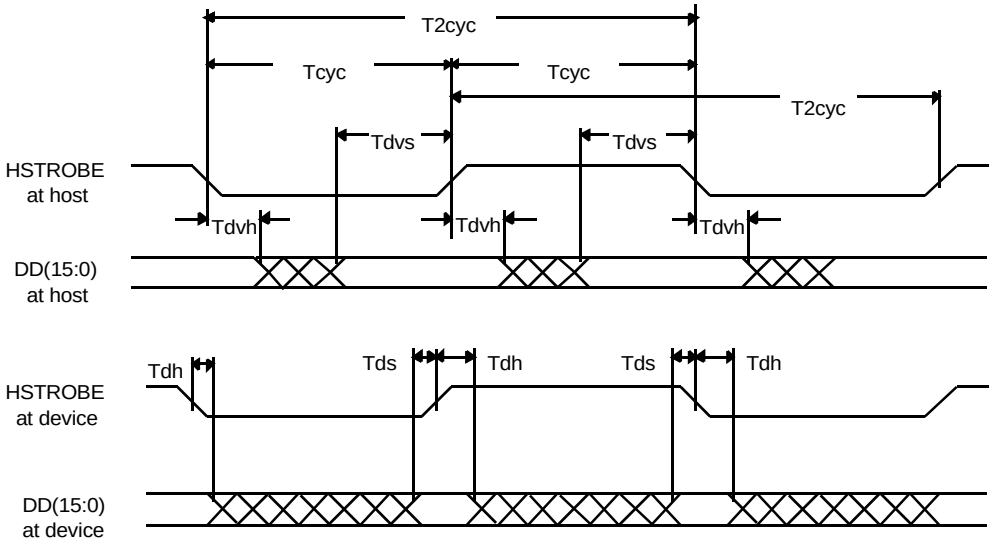

Figure 6-9 *Sustained Data Out Burst*

Note: DD(15:0) and HSTROBE signals are shown at both the device and the host to emphasize that cable settling time as well as cable propagation delay shall not allow the data signals to be considered stable at the device until well after they are driven by the host.

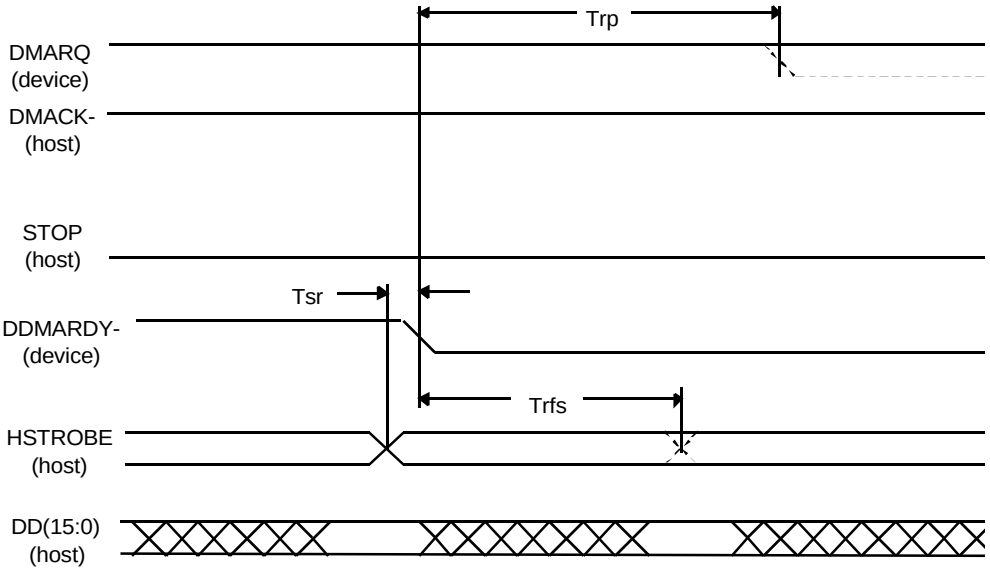

Figure 6-10 *Device Pausing a Data Out Burst*

Note: The device knows the burst is fully paused Trp ns after

DDMARDY- is negated and may then negate DMARQ to terminate the burst. Tsr timing need not be met for an asynchronous pause.

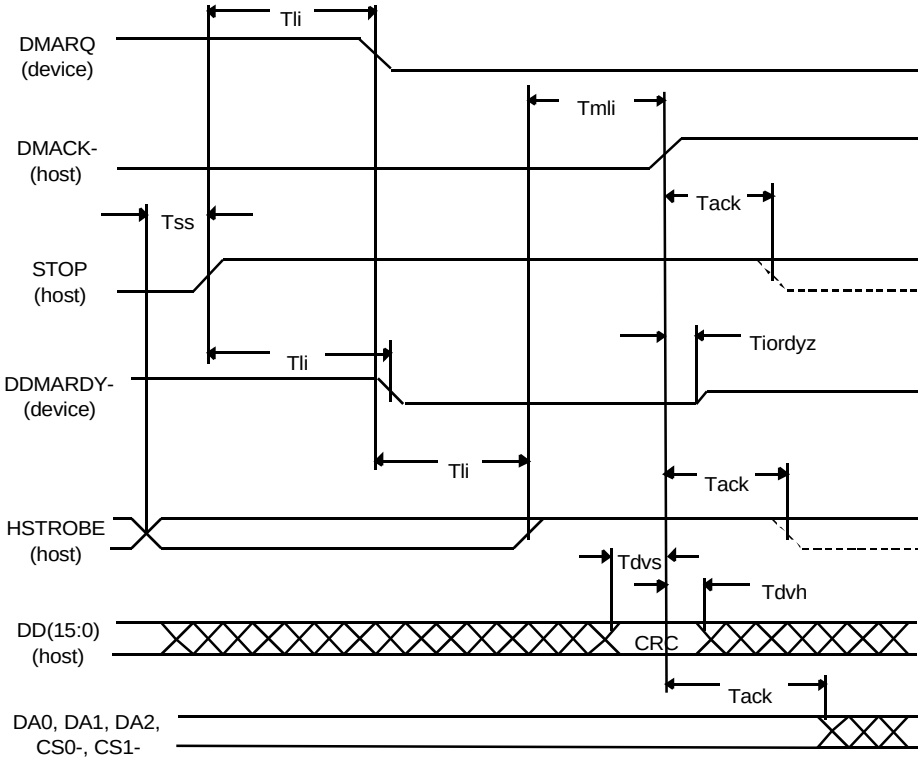

Figure 6-11 *Host Terminating a Data Out Burst*

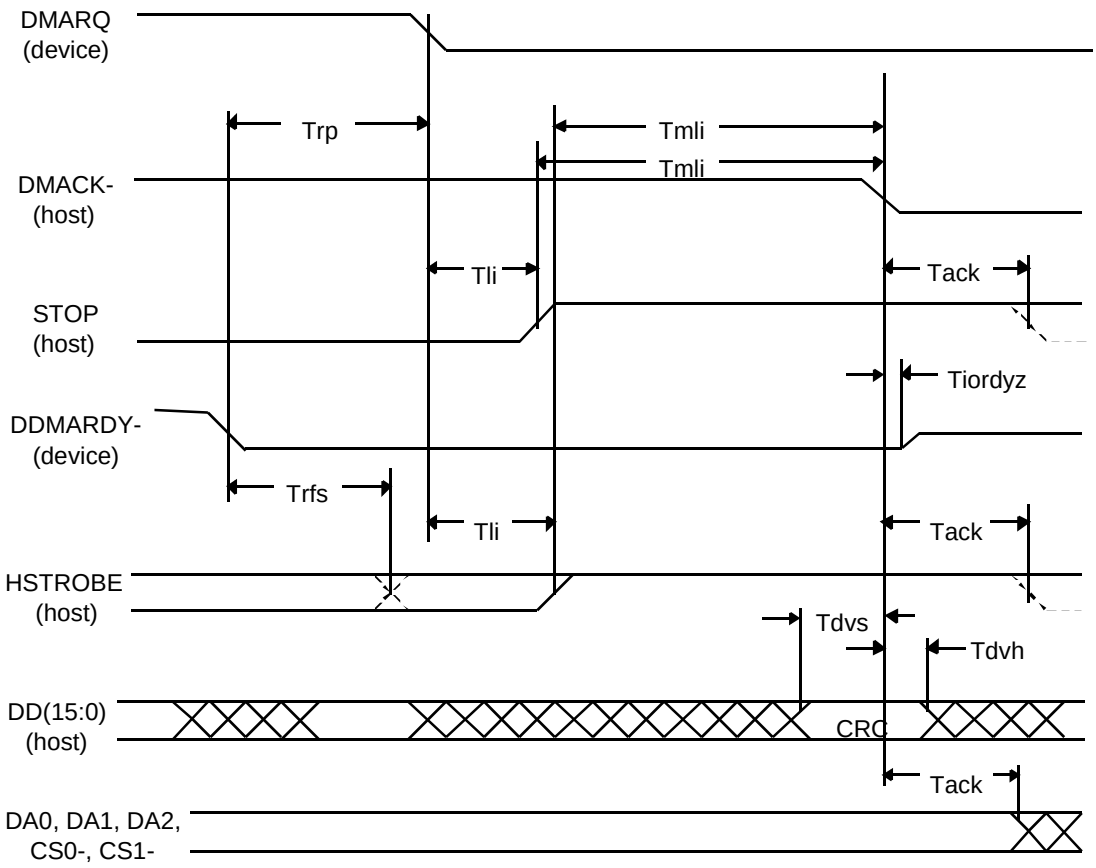

Figure 6-12 *Device Terminating a Data out Burst*

6.4.2.3 Host Interface RESET Timing

The host interface RESET timing shown in Table 6-9 is in reference to signals at 0.8 volts and 2.0 volts. All times are in nanoseconds, unless otherwise noted. Figure 6-13 provides a timing diagram.

Table 6-9 *Host Interface RESET Timing*

SYMBOL	DESCRIPTION	MINIMUM	MAXIMUM
tM	RESET– Pulse width	25	—

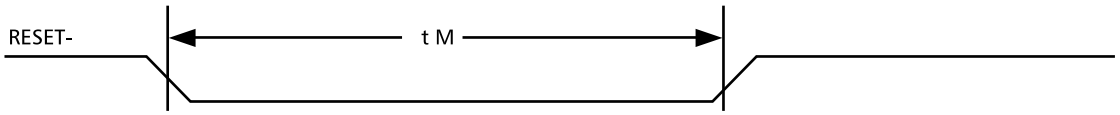

Figure 6-13 *Host Interface RESET Timing*

6.5 REGISTER ADDRESS DECODING

The host addresses the drive by using programmed I/O. Host address lines A0–A2, chip-select CS1FX– and CS3FX–, and IOR– and IOW– address the disk registers. Host address lines A3–A9 generate the two chip-select signals, CS1FX– and CS3FX–.

- Chip Select CS1FX– accesses the eight Command Block Registers.
- Chip Select CS3FX– is valid during 8-bit transfers to or from the Alternate Status Register.

The drive selects the primary or secondary command block addresses by setting Address bit A7.

Data bus lines 8–15 are valid only when IOCS16– is active and the drive is transferring data. The drive transfers ECC information only on data bus lines 0–7. Data bus lines 8–15 are invalid during transfers of ECC information.

I/O to or from the drive occurs over an I/O port that routes input or output data to or from selected registers, by using the following encoded signals from the host: CS1FX–, CS3FX–, DA2, DA1, DA0, DIOR–, and DIOW–. The host writes to the Command Block Registers when transmitting commands to the drive, and to the Control Block Registers when transmitting control, like a software reset. Table 6-10 lists the selection addresses for these registers.

Table 6-10 *I/O Port Functions and Selection Addresses*

FUNCTION		HOST SIGNALS				
CONTROL BLOCK REGISTERS		**CS1FX–**	**CS3FX–**	**DA2**	**DA1**	**DA0**
READ (DIOR–)	**WRITE (DIOW–)**					
Data Bus High Impedance	**Not Used**	**N[1]**	**N**	**X[2]**	**X**	**X**
Data Bus High Impedance	Not Used	N	A[3]	0	X	X
Data Bus High Impedance	Not Used	N	A	1	0	X
Alternate Status	Device Control	N	A	1	1	0
Drive Address	Not Used	N	A	1	1	1
COMMAND BLOCK REGISTERS						
READ (DIOR–)	**WRITE (DIOW–)**					
Data Port	Data Port	A	N	0	0	0
Error Register	Features	A	N	0	0	1
Sector Count	Sector Count	A	N	0	1	0
Sector Number[4]	Sector Number	A	N	0	1	1
LBA Bits 0–7[5]	LBA Bits 0–7	A	N	0	1	1
Cylinder Low[4]	Cylinder Low	A	N	1	0	0
LBA Bits 8–15[5]	LBA Bits 8–15	A	N	1	0	0
Cylinder High[4]	Cylinder High	A	N	1	0	1
LBA Bits 16–23[5]	LBA Bits 16–23	A	N	1	0	1
Drive/Head[4]	Drive/Head	A	N	1	1	0
LBA Bits 24–27[5]	LBA Bits 24–27	A	N	1	1	0
Status	Command	A	N	1	1	1
Invalid Address	Invalid Address	A	A	X	X	X

1. N = signal deasserted
2. X = signal either asserted or deasserted
3. A = signal asserted
4. Mapping of registers in CHS mode
5. Mapping of registers in LBA mode

After power on or following a reset, the drive initializes the Command Block Registers to the values shown in Table 6-11.

Table 6-11 *Command Block Register Initial Values*

REGISTER	VALUE
Error Register	01
Sector Count Register	01
Sector Number Register	01
Cylinder Low Register	00
Cylinder High Register	00
Drive/Head Register	00

6.6 REGISTER DESCRIPTIONS

The Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives emulate the ATA Command and Control Block Registers. Functional descriptions of these registers are given in the next two sections.

6.6.1 Control Block Registers

6.6.1.1 Alternate Status Register

The Alternate Status Register contains the same information as the Status Register in the command block. Reading the Alternate Status Register does not imply the acknowledgment of an interrupt by the host or clear a pending interrupt. See the description of the Status Register in section 6.6.2.8 for definitions of bits in this register.

6.6.1.2 Device Control Register

This write-only register contains two control bits, as shown in Table 6-12.

Table 6-12 *Device Control Register Bits*

BIT	MNEMONIC	DESCRIPTION
7	Reserved	–
6	Reserved	–
5	Reserved	–
4	Reserved	–
3	1	Always 1
2	SRST[1]	Host software reset bit
1	nIEN[2]	Drive interrupt enable bit
0	0	Always 0

1. SRST = Host Software Reset bit. When the host sets this bit, the drive is reset. When two drives are daisy-chained on the interface, this bit resets both drives simultaneously.

2. nIEN = Drive Interrupt Enable bit. When nIEN equals 0 or the host has selected the drive, the drive enables the host interrupt signal INTRQ through a tristate buffer to the host. When nIEN equals 1 or the drive is not selected, the host interrupt signal INTRQ is in a high-impedance state, regardless of the presence or absence of a pending interrupt.

6.6.1.3 Drive Address Register

The Drive Address Register returns the head-select addresses for the drive currently selected. Table 6-13 shows the Drive Address bits.

Table 6-13 *Drive Address Register Bits*

BIT	MNEMONIC	DESCRIPTION
7	HiZ[1]	High Impedance bit
6	nWTG[2]	Write Gate bit
5	nHS3[3]	Head Address msb
4	nHS2	–
3	nHS1	–
2	nHS0	Head Address lsb
1	nDS1[4]	Drive 1 Select bit
0	nDS0	Drive 0 Select bit

1. HiZ = High Impedance bit. When the host reads the register, this bit will be in a high impedance state.
2. nWTG = Write Gate bit. When a write operation to the drive is in progress, nWTG equals 0.
3. nHS0–nHS3 = Head Address bits. These bits are equivalent to the one's complement of the binary-coded address of the head currently selected.
4. nDS0–nDS1 = Drive Select bits. When drive 1 is selected, nDS1 equals 0. When drive 0 is selected, nDS0 equals 0.

6.6.2 Command Block Registers

6.6.2.1 Data Port Register

All data transferred between the device data buffer and the host passes through the Data Port Register. The host transfers the sector table to this register during execution of the FORMAT TRACK command.

6.6.2.2 Error Register

The Error Register contains status information about the last command executed by the drive. The contents of this register are valid only when the Error bit (ERR) in the Status Register is set to 1. The contents of the Error Register are also valid at power on, and at the completion of the drive's internal diagnostics, when the register contains a status code. When the error bit in the Status Register is set to 1, the host interprets the Error Register bits as shown in Table 6-14.

Table 6-14 *Error Register Bits*

MNEMONIC	BIT	DESCRIPTION
#	7	
#	6	
#	5	
#	4	
#	3	
ABRT	2	Requested command aborted due to a drive status error, such as Not Ready or Write Fault, or because the command code is invalid.
#	1	
#	0	

6.6.2.3 Sector Count Register

The Sector Count Register defines the number of sectors of data to be transferred across the host bus for a subsequent command. If the value in this register is 0, the sector count is 256 sectors. If the Sector Count Register command executes successfully, the value in this register at command completion is 0. As the drive transfers each sector, it decrements the Sector Count Register to reflect the number of sectors remaining to be transferred. If the command's execution is unsuccessful, this register contains the number of sectors that must be transferred to complete the original request.

When the drive executes an INITIALIZE DRIVE PARAMETERS or Format Track command, the value in this register defines the number of sectors per track.

6.6.2.4 Sector Number Register

The Sector Number Register contains the ID number of the first sector to be accessed by a subsequent command. The sector number is a value between one and the maximum number of sectors per track. As the drive transfers each sector, it increments the Sector Number Register. See the command descriptions in section 6.7 for information about the contents of the Sector Number Register after successful or unsuccessful command completion.

In LBA mode, this register contains bits 0 to 7. At command completion, the host updates this register to reflect the current LBA bits 0 to 7.

6.6.2.5 Cylinder Low Register

The Cylinder Low Register contains the eight low-order bits of the starting cylinder address for any disk access. On multiple sector transfers that cross cylinder boundaries, the drive updates this register when command execution is complete, to reflect the current cylinder number. The host loads the least significant bits of the cylinder address into the Cylinder Low Register.

In LBA mode, this register contains bits 8 to 15. At command completion, the drive updates this register to reflect the current LBA bits 8 to 15.

6.6.2.6 Cylinder High Register

The Cylinder High Register contains the eight high-order bits of the starting cylinder address for any disk access. On multiple sector transfers that cross cylinder boundaries, the drive updates this register at the completion of command

execution, to reflect the current cylinder number. The host loads the most significant bits of the cylinder address into the Cylinder High Register.

In LBA mode, this register contains bits 16 to 23. At command completion, the host updates this register to reflect the current LBA bits 16 to 23.

6.6.2.7 Drive/Head Register

The Drive/Head Register contains the drive ID number and its head numbers. At command completion this register is updated by the drive to reflect the current head.

In LBA mode, this register contains bits 24 to 27. At command completion, the drive updates this register to reflect the current LBA bits 24 to 27.

Table 6-15 shows the Drive/Head Register bits.

Table 6-15 *Drive Head Register Bits*

MNEMONIC	BIT	DESCRIPTION
Reserved	7[1]	Always *1*
L	6[2]	*0* for CHS mode *1* for LBA mode
Reserved	5	Always *1*
DRV	4[3]	*0* indicates the Master drive is selected *1* indicates the Slave drive is selected
HS3	3[4]	Most significant Head Address bit in CHS mode Bit 24 of the LBA Address in LBA mode
HS2	2	Head Address bit for CHS mode Bit 25 of the LBA Address in LBA mode
HS1	1	Head Address bit for CHS mode Bit 26 of the LBA Address in LBA mode
HS0	0	Least significant Head Address bit in CHS mode Bit 27 of the LBA Address in LBA mode

1. Bits 5–7 define the sector size set in hardware (512 bytes).
2. Bit 6 is the binary encoded Address Mode Select. When bit 6 is set to 0, addressing is by CHS mode. When bit 6 is set to 1, addressing is by LBA mode.
3. Bit 4 (DRV) contains the binary encoded drive select number. The Master is the primary drive; the Slave is the secondary drive
4. In CHS mode, bits 3–0 (HS0–HS3) contain the binary encoded address of the selected head. At command completion, the host updates these bits to reflect the address of the head currently selected. In LBA mode, bits 3–0 (HS0–HS3) contain bits 24–27 of the LBA Address. At command completion, the host updates this register to reflect the current LBA bits 24 to 27.

6.6.2.8 Status Register

The Status Register contains information about the status of the drive and the controller. The drive updates the contents of this register at the completion of each command. When the Busy bit is set (BSY=1), no other bits in the Command Block Registers are valid. When the Busy bit is not set (BSY=0), the information in the Status Register and Command Block Registers is valid.

When an interrupt is pending, the drive considers that the host has acknowledged the interrupt when it reads the Status Register. Therefore, whenever the host reads this register, the drive clears any pending interrupt. Table 6-16 defines the Status Register bits.

Table 6-16 *Status Register Bits*

MNEMONIC	BIT	DESCRIPTION
BSY	7	Busy bit. Set by the controller logic of the drive whenever the drive has access to and the host is locked out of the Command Block Registers. BSY is set under the following conditions: • Within 400 ns after the deassertion of RESET- or after SRST is set in the Device Control Register. Following a reset, BSY will be set for no longer than 30 seconds. • Within 400 ns of a host write to the Command Block Registers with a Read, READ LONG, READ BUFFER, SEEK, RECALIBRATE, INITIALIZE DRIVE PARAMETERS, Read Verify, Identify Drive, or EXECUTE DRIVE DIAGNOSTIC command. • Within 5 μsec after the transfer of 512 bytes of data during the execution of a Write, Format Track, or WRITE BUFFER command, or 512 bytes of data and the appropriate number of ECC bytes during the execution of a WRITE LONG command. When BSY=1, the host cannot write to a Command Block Register and reading any Command Block Register returns the contents of the Status Register.
DRDY	6	Drive Ready bit. Indicates that the drive is ready to accept a command. When an error occurs, this bit remains unchanged until the host reads the Status Register, then again indicates that the drive is ready. At power on, this bit should be cleared, and should remain cleared until the drive is up to speed and ready to accept a command.
#	5	
#	4	
DRQ	3	Data Request bit. When set, this bit indicates that the drive is ready to transfer a word or byte of data from the host to the data port.
Obsolete	2	
Obsolete	1	
ERR	0	Error bit. When set, this bit indicates that the previous command resulted in an error. The other bits in the Status Register and the bits in the Error Register contain additional information about the cause of the error.

Note: The content of # bit is command dependent.

Note: Bits 2 and 1 are obsolete according to the ATA-4 specification.

6.6.2.9 Command Register

The host sends a command to the drive by means of an 8-bit code written to the Command Register. As soon as the drive receives the command in its Command Register, it begins execution of the command. Table 6-17 lists the hexadecimal command codes and parameters for each executable command. The code F0h is common to all of the extended commands. Each of these commands is distinguished by a unique subcode. For a detailed description of each command, see Section 6.7, "COMMAND DESCRIPTIONS," found later in this chapter.

Table 6-17 *Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT Command Codes and Parameters*

COMMAND		PARAMETER					
NAME	CODE	SC Ex. Sub Code	SN	CY	DS	HD	FR
RECALIBRATE	1Xh				V		
READ SECTORS	20h	V	V	V	V	V	
WRITE SECTORS	30h	V	V	V	V	V	
READ VERIFY SECTORS	40h	V	V	V	V	V	
SEEK	7Xh		V	V	V	V	
EXECUTE DRIVE DIAGNOSTIC	90h						
INITIALIZE DRIVE PARAMETERS	91h	V			V	V	
DOWNLOAD MICROCODE	92h	V	V	V	V	V	
SMART	B0h			V			V
READ MULTIPLE	C4h	V	V	V	V	V	
WRITE MULTIPLE	C5h	V	V	V	V	V	
SET MULTIPLE MODE	C6h	V			V		
READ DMA	C8h	V	V	V	V	V	
WRITE DMA	CAh	V	V	V	V	V	
STANDBY IMMEDIATE	E0h					V	
IDLE IMMEDIATE	E1h					V	
STANDBY MODE (AUTO POWER-DOWN)	E2h	V				V	
IDLE MODE (AUTO POWER-DOWN)	E3h	V				V	
READ BUFFER	E4h				V		
CHECK POWER MODE	E5h	V				V	
SLEEP MODE	E6h					V	
FLUSH CACHE—optional cmnd.	E7h				V		
WRITE BUFFER	E8h				V		
IDENTIFY DRIVE	ECh				V		
READ DEFECT LIST—extended cmnd.	F0h	O0h	V	V	V		
READ CONFIGURATION—extended cmnd.	F0h	O1h	V	V	V		
SET CONFIGURATION—extended cmnd.	F0h	FEh/ FFh	V	V	V		
SCAN VERIFY—extended cmnd.	F0h	FCh	V	V	V		
GET SCAN VERIFY STATUS—extended cmnd.	F0h	FDh	V	V	V		
READ NATIVE MAX ADDRESS	F8						
SET MAX ADDRESS	F9		V	V	V	V	

Note: The following information applies to Table 6-17:
SC = Sector Count Register
SN = Sector Number Register
CY = Cylinder Low and High Registers
DS = Drive Select bit (Bit 4 of Drive/Head Register)
HD = 4 Head Select Bits (Bits 0–3 of Drive Head Register)
V = Must contain valid information for this command Head
FR = Features Register

6.7 COMMAND DESCRIPTIONS

The Quantum Fireball lct20 AT hard disk drives support all standard ATA drive commands. The drive decodes, then executes, commands loaded into the Command Block Register. In applications involving two hard drives, both drives receive all commands. However, only the selected drive executes commands—with the exception of the EXECUTE DRIVE DIAGNOSTIC command, as explained below. The procedure for executing a command on the selected drive is as follows:

1. Wait for the drive to indicate that it is no longer busy (BSY=0).
2. Activate the Interrupt Enable (–IEN) bit.
3. Wait for the drive to set RDY (RDY=1).
4. Load the required parameters into the Command Block Register.
5. Write the command code to the Command Register.

Execution of the command begins as soon as the drive loads the Command Block Register. The remainder of this section describes the function of each command. The commands are listed in the same order they appear in Table 6-17.

6.7.1 Recalibrate 1xh

The RECALIBRATE command moves the read/write heads from any location on the disk to cylinder 0. On receiving this command, the drive sets the BSY bit and issues a seek command to cylinder 0. The drive then waits for the seek operation to complete, updates status, negates BSY, and generates an interrupt. If the drive cannot seek to cylinder 0, it posts the message TRACK 0 NOT FOUND.

6.7.2 Read Sectors 20h

The Read Sectors command reads from 1 to 256 sectors, beginning at the specified sector. As specified in the Command Block Register, a sector count equal to 0 requests 256 sectors. When the drive accepts this command, it sets BSY and begins execution of the command.

6.7.2.1 Multiple Sector Reads

Multiple sector reads set DRQ. After reading each sector, the drive generates an interrupt when the sector buffer is full, and the drive is ready for the host to read the data. Once the host empties the sector buffer, the drive immediately clears DRQ and sets BSY.

If an error occurs during a multiple sector read, the read terminates at the sector in which the error occurred. The Command Block Register contains the cylinder, head, and sector numbers of the sector in which the error occurred. The host can then read the Command Block Register to determine what kind of error has occurred, and in which sector. Whether the data error is correctable or uncorrectable, the drive loads the data into the sector buffer.

6.7.3 Write Sector 30h

The WRITE SECTOR command writes from 1 to 256 sectors, beginning at the specified sector. As specified in the Command Block Register, a sector count equal to 0 requests 256 sectors. When the drive accepts this command, it sets DRQ and waits for the host to fill the sector buffer with the data to be written to the drive. The drive does not generate an interrupt to start the first buffer-fill operation. Once the buffer is full, the drive clears DRQ, sets BSY, and begins execution of the command.

6.7.3.1 Multiple Sector Writes

The MULTIPLE SECTOR WRITES command sets DRQ. The drive generates an interrupt whenever the sector buffer is ready to be filled. When the host fills the sector buffer, the drive immediately clears DRQ and sets BSY.

If an error occurs during a multiple sector write operation, the write operation terminates at the sector in which the error occurred. The Command Block Register contains the cylinder, head, and sector numbers of the sector in which the error occurred. The host can then read the Command Block Register to determine what kind of error has occurred, and in which sector.

6.7.4 Read Verify Sectors 40h

The execution of the READ VERIFY SECTORS command is identical to that of the READ SECTORS command. However, the Read Verify command does not cause the drive to set DRQ, the drive transfers no data to the host, and the Long bit is invalid. On receiving the READ VERIFY command, the drive sets BSY. When the drive has verified the requested sectors, it clears BSY and generates an interrupt. On command completion, the Command Block Registers contain the cylinder, head, and sector numbers of the last sector verified.

If an error occurs during a multiple sector verify operation, the read operation terminates at the sector in which the error occurred. The Command Block Registers contain the cylinder, head, and sector numbers in which the error occurred.

6.7.5 Seek 7xh

The SEEK command causes the actuator to seek the track to which the Cylinder and Drive/Head registers point. When the drive receives this command in its Command Block Registers, it performs the following functions:

1. Sets BSY
2. Initiates the seek operation
3. Resets BSY
4. Sets the Drive Seek Complete (DSC) bit in the Status Register

The drive does not wait for the seek to complete before it sends an interrupt. If the

BSY bit is *not* set in the Status Register, the drive can accept and queue subsequent commands while performing the seek. If the Cylinder registers contain an illegal cylinder, the drive sets the ERR bit in the Status Register and the IDNF bit in the Error Register.

6.7.6 Execute Drive Diagnostic 90h

The EXECUTE DRIVE DIAGNOSTIC command performs the internal diagnostic tests implemented on the drive. Drive 0 sets BSY within 400 ns of receiving of the command.

If Drive 1 is present:

- Both drives execute diagnostics.
- Drive 0 waits up to six seconds for drive 1 to assert PDIAG–.
- If drive 1 does not assert PDIAG– to indicate a failure, drive 0 appends 80h with its own diagnostic status.
- If the host detects a drive 1 diagnostic failure when reading drive 0 status, it sets the DRV bit, then reads the drive 1 status.

If Drive 1 is not present:

- Drive 0 reports only its own diagnostic results.
- Drive 0 clears BSY and generates an interrupt.

If drive 1 fails diagnostics, drive 0 appends 80h with its own diagnostic status and loads that code into the Error Register. If drive 1 passes its diagnostics or no drive 1 is present, drive 0 appends 00h with its own diagnostic status and loads that code into the Error Register.

The diagnostic code written to the Error Register is a unique 8-bit code. Table 6-18 lists the diagnostic codes.

Table 6-18 *Diagnostic Codes*

DIAGNOSTIC CODE	DESCRIPTION
01h	No Error Detected
02h	Formatter Device Error
03h	Sector Buffer Error
04h	ECC Circuitry Error
05h	Controlling Microprocessor Error
8Xh	Drive 1 Failed

6.7.7 Initialize Drive Parameters 91h

The INITIALIZE DRIVE PARAMETERS command enables the host to set the logical number of heads and the logical number of sectors per track. On receiving the command, the drive sets the BSY bit, saves the parameters, clears BSY, and generates an interrupt.

The only two register values used by this command are the Sector Count Register, which specifies the number of sectors; and the Drive/Head Register, which specifies

the number of heads, minus 1. The DRV bit assigns these values to drive 0 or drive 1, as appropriate.

This command does not check the sector count and head values for validity. If these values are invalid, the drive will not report an error until another command causes an illegal access.

6.7.8 Download Microcode

COMMAND CODE - 92h

TYPE - Optional

PROTOCOL - PIO data out

INPUTS - The head bits of the device/head register will always be set to zero. The cylinder high and low registers will be set to zero. The sector number and the sector count are used together as a 16-bit sector count value. The feature register specifies the subcommand code.

Register	7	6	5	4	3	2	1	0
Features	Subcommand code							
Sector Count	Sector count (low order)							
Sector Number	Sector count (high order)							
Cylinder Low	00h							
Cylinder High	00h							
Device/Head	1		1	D	0	0	0	0
Command	92h							

NORMAL OUTPUTS- None. required.

ERROR OUTPUTS- Aborted command if the device does not support this command or did not accept the microcode data. Aborted error if subcommand code is not a supported value.

Status Register				Error Register					
DRDY	DF	CORR	ERR	BBK	UNC	IDNF	ABRT	TK0NF	AMNF
V	V		V				V		

PREREQUISITES - DRDY set equal to one.

DESCRIPTION - This command enables the host to alter the device's microcode. The data transferred using the DOWNLOAD MICROCODE command is vendor specific.

All transfers will be an integer multiple of the sector size. The size of the data transfer is determined by the contents of the Sector Number register and the Sector Count register. The Sector Number register will be used to extend the Sector Count register, to create a sixteen bit sector count value. The Sector Number register will be the most significant eight bits and the Sector Count register will be the least significant eight bits. A value of zero in both the Sector Number register and the Sector Count register will indicate no data is to transferred. This allows transfer sizes from 0 bytes to 33, 553, 920 bytes in 512 byte increments.

The Features register will be used to determine the effect of the DOWNLOAD

MICROCODE sub command. The values for the Feature Register are:

01h — download is for immediate, temporary use

07h — save downloaded code for immediate and future use.

Either or both values may be supported. All other values are reserved.

6.7.9 SMART B0h

SMART DISABLE OPERATIONS

COMMAND CODE - B0h

TYPE - Optional - SMART Feature set. If the SMART feature set is implemented, this command shall be implemented.

PROTOCOL - Non-data command

INPUTS - The Features register shall be set to D9h. The Cylinder Low register shall be set to 4Fh. The Cylinder High register shall be set to C2h.

Register	7	6	5	4	3	2	1	0
Features	D9h							
Sector Count								
Sector Number								
Cylinder Low	4Fh							
Cylinder High	C2h							
Device/Head	1		1	D				
Command	B0h							

NORMAL OUTPUTS - None

ERROR OUTPUTS - If the device does not support this command, if SMART is not enabled or if the values in the Features, Cylinder Low or Cylinder High registers are invalid, an Aborted command error is posted.

Status Register				Error Register					
DRDY	DF	CORR	ERR	BBK	UNC	IDNF	ABRT	TK0NF	AMNF
V			V				V		

PREREQUISITES - DRDY set equal to one. SMART enabled.

DESCRIPTION - This command disables all SMART capabilities within the device including any and all timer functions related exclusively to this feature. After receipt of this command the device will disable all SMART operations. Attribute values will no longer be monitored or saved by the device. The state of SMART (either enabled or disabled) will be preserved by the device across power cycles.

Upon receipt of the SMART DISABLE OPERATIONS command from the host, the device sets BSY, disables SMART capabilities and functions, clears BSY and asserts INTRQ.

After receipt of this command by the device, all other SMART commands, with the exception of SMART ENABLE OPERATIONS, are disabled and invalid and shall be aborted by the device (including SMART DISABLE OPERATIONS commands), returning the Aborted command error.

6.7.9.1 SMART ENABLE/DISABLE ATTRIBUTE AUTOSAVE

COMMAND CODE - B0h

TYPE - Optional - SMART Feature set. If the SMART feature set is implemented, this command is optional and not recommended.

PROTOCOL - Non-data command

INPUTS - The Features register shall be set to D2h. The Cylinder Low register shall be set to 4Fh. The Cylinder High register shall be set to C2h. The Sector Count register is set to 00h to disable attribute autosave and a value of F1h is set to enable attribute autosave.

Register	7	6	5	4	3	2	1	0
Features	D2h							
Sector Count	00h or F1h							
Sector Number								
Cylinder Low	4Fh							
Cylinder High	C2h							
Device/Head	1		1	D				
Command	B0h							

NORMAL OUTPUTS - None

ERROR OUTPUTS - If the device does not support this command, if SMART is disabled or if the values in the Features, Cylinder Low or Cylinder High registers are invalid, an Aborted command error is posted.

Status register				Error register					
DRDY	DF	CORR	ERR	BBK	UNC	IDNF	ABRT	TK0NF	AMNF
V			V				V		

PREREQUISITES - DRDY set equal to one. SMART enabled.

DESCRIPTION - This command enables and disables the optional attribute autosave feature of the device. Depending upon the implementation, this command may either allow the device, after some vendor specified event, to automatically save its updated attribute values to non-volatile memory; or this command may cause the autosave feature to be disabled. The state of the attribute autosave feature (either enabled or disabled) will be preserved by the device across power cycles.

A value of zero written by the host into the device's Sector Count register before issuing this command will cause this feature to be disabled. Disabling this feature does not preclude the device from saving attribute values to non-volatile memory during some other normal operation such as during a power-on or power-off sequence or during an error recovery sequence.

A value of F1h written by the host into the device's Sector Count register before issuing this command will cause this feature to be enabled. Any other meaning of this value or any other non-zero value written by the host into this register before issuing this command is vendor specific. The meaning of any non-zero value written to this register at this time will be preserved by the device across power cycles.

If the SMART ENABLE/DISABLE ATTRIBUTE AUTOSAVE command is supported by the device, upon receipt of the command from the host, the device sets BSY, enables or disables the autosave feature (depending on the implementation), clears BSY and asserts INTRQ.

If this command is not supported by the device, the device shall abort the command upon receipt from the host, returning the Aborted command error.

During execution of the autosave routine the device shall not assert BSY nor deassert DRDY. If the device receives a command from the host while executing its autosave routine it must respond to the host within two seconds.

6.7.9.2 SMART ENABLE OPERATIONS

COMMAND CODE - B0h

TYPE - Optional - SMART Feature set. If the SMART feature set is implemented, this command shall be implemented.

PROTOCOL - Non-data command

INPUTS - The Features register shall be set to D8h. The Cylinder Low register shall be set to 4Fh. The Cylinder High register shall be set to C2h.

Register	7	6	5	4	3	2	1	0
Features	D8h							
Sector Count								
Sector Number								
Cylinder Low	4Fh							
Cylinder High	C2h							
Device/Head	1		1	D				
Command	B0h							

NORMAL OUTPUTS - None

ERROR OUTPUTS - If the device does not support this command or if the values in the Features, Cylinder Low or Cylinder High registers are invalid, an Aborted command error is posted.

Status register				Error register					
DRDY	DF	CORR	ERR	BBK	UNC	IDNF	ABRT	TK0NF	AMNF
V			V				V		

PREREQUISITES -DRDY set equal to one.

DESCRIPTION - This command enables access to all SMART capabilities within the device. Prior to receipt of this command attribute values are neither monitored nor saved by the device. The state of SMART (either enabled or disabled) will be preserved by the device across power cycles. Once enabled, the receipt of subsequent SMART ENABLE OPERATIONS commands shall not affect any of the attribute values.

Upon receipt of this command from the host, the device sets BSY, enables SMART capabilities and functions, clears BSY and asserts INTRQ.

6.7.9.3 SMART READ ATTRIBUTE THRESHOLDS

COMMAND CODE - B0h

TYPE - Optional - SMART Feature set. If the SMART feature set is implemented, this command is optional and not recommended.

PROTOCOL - PIO data in

INPUTS - The Features register shall be set to D1h. The Cylinder Low register shall be set to 4Fh. The Cylinder High register shall be set to C2h.

Register	7	6	5	4	3	2	1	0
Features	D1h							
Sector Count								
Sector Number								
Cylinder Low	4Fh							
Cylinder High	C2h							
Device/Head	1		1	D				
Command	B0h							

NORMAL OUTPUTS - None

ERROR OUTPUTS - If the device does not support this command, if SMART disabled or if the values in the Features, Cylinder Low or Cylinder High registers are invalid, an Aborted command error is posted.

Status register				**Error register**					
DRDY	DF	CORR	ERR	BBK	UNC	IDNF	ABRT	TK0NF	AMNF
V			V				V		

PREREQUISITES - DRDY set equal to one. SMART enabled.

DESCRIPTION - This command returns the device's attribute thresholds to the host. Upon receipt of this command from the host, the device sets BSY, reads the attribute thresholds from non-volatile memory, sets DRQ, clears BSY, asserts INTRQ, and then waits for the host to transfer the 512 bytes of attribute threshold information from the device via the Data register.

The following defines the 512 bytes that make up the attribute threshold information.

The sequence of active attribute thresholds must appear in the same order as their corresponding attribute values (see 6.7.9.5).

The data structure revision number shall be the same value used in the device attribute values data structure.

Table 6-19 defines the twelve bytes that make up the information for each threshold entry in the device attribute thresholds data structure. Attribute entries in the individual threshold data structure must be in the same order and correspond to the entries in the individual attribute data structure.

The attribute ID numbers are vendor specific. Any non-zero value in the attribute ID number indicates an active attribute.

Attribute threshold values are to be set at the factory and are not changeable in the field.

The data structure checksum is the two's compliment of the result of a simple eight-bit addition of the first 511 bytes in the data structure.

Table 6-19 *Device attribute thresholds data structure*

Description	Bytes	Format	Type
Data structure revision number = 0x0004h for this revision	2	binary	Rd only
1st attribute threshold	12		Rd only
.....			
.....			
.....			
30th attribute threshold	12		Rd only
reserved (0x00)	18		Rd only
Vendor specific	131		Rd only
Data structure checksum	1		Rd only
Total bytes	512		

Table 6-20 *Individual threshold data structure*

Description	Bytes	Format	Type
Attribute ID number	1	binary	Rd only
Attribute threshold (for comparison with attribute values from 0x00 to 0xFFh)	1	binary	Rd only
0x00 “always passing” threshold value to be used for			
code test purposes			
0x01 minimum value for normal operation			
0xFD maximum value for normal operation			
0xFE invalid for threshold value - not to be used			
0xFF “always failing” threshold value to be used for code test purposes			
Reserved	10		Rd only
Total bytes	12		

6.7.9.4 SMART READ ATTRIBUTE VALUES

COMMAND CODE - B0h

TYPE - Optional - SMART Feature set. If the SMART feature set is implemented, this command is optional and not recommended.

PROTOCOL - PIO data in

INPUTS - The Features register shall be set to D0h. The Cylinder Low register shall be set to 4Fh. The Cylinder High register shall be set to C2h.

Register	7	6	5	4	3	2	1	0
Features	D0h							
Sector Count								
Sector Number								
Cylinder Low	4Fh							
Cylinder High	C2h							
Device/Head	1		1	D				
Command	B0h							

NORMAL OUTPUTS - None

ERROR OUTPUTS - If the device does not support this command, if SMART is disabled or if the values in the Features, Cylinder Low or Cylinder High registers are invalid, an Aborted command error is posted.

Status register				Error register					
DRDY	DF	CORR	ERR	BBK	UNC	IDNF	ABRT	TK0NF	AMNF
V			V				V		

PREREQUISITES - DRDY set equal to one. SMART enabled.

DESCRIPTION - This command returns the device's attribute values to the host. Upon receipt of this command from the host, the device sets BSY, saves any updated attribute values to non-volatile memory, sets DRQ, clears BSY, asserts INTRQ, and then waits for the host to transfer the 512 bytes of attribute value information from the device via the Data register.

The following defines the 512 bytes that make up the attribute value information.

The number of active attributes and, therefore, number of active attribute values is determined independently by the device manufacturer for each individual device. All active attribute entries should be concatenated together directly after the data structure revision number. If there are fewer than thirty active attributes implemented on a device, the excess locations in the data structure are reserved for future attribute implementations and are designated as blanks containing the value

0x00h. Thus the first reserved byte following the attribute entries shall be the 363rd byte in the structure, the first SMART capability byte shall be the 369th byte in the structure, etc.

The data structure revision number identifies which version of this data structure is implemented by a device. Upon initial release of this specification, the revision number will be set to 0x0004h. Later revisions, if any, will increment the revision number by one for each new revision. The revision number will be the same for both the attribute value and attribute threshold structures.

Table 6-21 defines the twelve bytes that make up the information for each attribute entry in the device attributes data structure.

Table 6-21 *Device attributes data structure*

Description	Bytes	Format	Type
Data structure revision number = 0x0004 for this specification revision	2	binary	Rd only
1st device attribute	12		Rd/Wrt
...			
...			
...			
30th device attribute	12		Rd/Wrt
reserved (0x00)	6		Rd only
SMART capability	2		Rd only
reserved (0x00)	16		Rd/Wrt
Vendor specific	125		Rd only
Data structure checksum	1		Rd only
Total bytes	512		

Table 6-22 *Individual attribute data structure*

Description	Bytes	Format	Type
Attribute ID number (0x01 to 0xFFh)	1	binary	Rd only
Status flags	2	bit flags	Rd only
Pre-failure/advisory bit			
Vendor specific (5 bits)			
reserved (10 bits)			
Attribute value (valid values from 0x01 to 0xFEh)	1	binary	Rd only
0x00 invalid for attribute value - not to be used			
0x01 minimum value			
0x64 initial value for all attributes prior to any data			
collection			
0xFD maximum value			
0xFE value is not valid			
0xFF invalid for attribute value - not to be used			
Vendor specific	8	binary	Rd only
Total bytes	12		

The attribute ID numbers and their definitions are vendor specific. Any non-zero value in the attribute ID number indicates an active attribute. Valid values for this byte are from 0x01 through 0xFFh.

Status flag

- Bit 0 -Pre-failure/advisory - If the value of this bit equals zero, an attribute value less than or equal to its corresponding attribute threshold indicates an advisory condition where the usage or age of the device has exceeded its intended design life period. If the value of this bit equals one, an attribute value less than or equal to its corresponding attribute threshold indicates a pre-failure condition where imminent loss of data is being predicted.
- Bit 1 Reserved for future use.
- Bits 3 - 6 - Vendor specific.
- Bits 7 - 15 - Reserved for future use.

The range and meaning of the attribute values is described in Table 20. Prior to the monitoring and saving of attribute values, all values are set to 0x64h. The attribute values of 0x00h and 0xFFh are reserved and should not be used by the device.

SMART capability

- Bit 0 - Pre-power mode attribute saving capability - If the value of this bit equals one, the device will save its attribute values prior to going into a power saving mode (Idle, Standby or Sleep modes).
- Bit 1 - Attribute autosave after event capability - If the value of this bit is equal to one, the device supports the SMART ENABLE/DISABLE ATTRIBUTE AUTOSAVE command.
- Bits 2-15 - Reserved for future use.

The data structure checksum is the two's compliment of the result of a simple eight-bit addition of the first 511 bytes in the data structure.

6.7.9.5 SMART RETURN STATUS

COMMAND CODE - B0h

TYPE - Optional - SMART Feature set. If the SMART feature set is implemented, this command shall be implemented.

PROTOCOL - Non-data command.

INPUTS - The Features register shall be set to DAh. The Cylinder Low register shall be set to 4Fh. The Cylinder High register shall be set to C2h.

Register	7	6	5	4	3	2	1	0
Features	DAh							
Sector Count								
Sector Number								
Cylinder Low	4Fh							
Cylinder High	C2h							
Device/Head	1		1	D				
Command	B0h							

NORMAL OUTPUTS - If the device has not detected a threshold exceeded condition, the device sets the Cylinder Low register to 4Fh and the Cylinder High register to C2h. If the device has detected a threshold exceeded condition, the device sets the Cylinder Low register to F4h and the Cylinder High register to 2Ch.

ERROR OUTPUTS - If the device does not support this command, if SMART is disabled or if the values in the Features, Cylinder Low or Cylinder High registers are invalid, an Aborted command error is posted.

Status register				Error register					
DRDY	DF	CORR	ERR	BBK	UNC	IDNF	ABRT	TK0NF	AMNF
V			V				V		

PREREQUISITES - DRDY set equal to one. SMART enabled.

DESCRIPTION - This command is used to communicate the reliability status of the device to the host at the host's request. Upon receipt of this command the device sets BSY, saves any updated attribute values to non-volatile memory and compares the updated attribute values to the attribute thresholds.

6.7.9.6 SMART SAVE ATTRIBUTE VALUES

COMMAND CODE - B0h

TYPE - Optional - SMART Feature set. If the SMART feature set is implemented, this command is optional and not recommended.

PROTOCOL - Non-data command

INPUTS - The Features register shall be set to D3h. The Cylinder Low register shall be set to 4Fh. The Cylinder High register shall be set to C2h.

Register	7	6	5	4	3	2	1	0
Features	D3h							
Sector Count								
Sector Number								
Cylinder Low	4Fh							
Cylinder High	C2h							
Device/Head	1		1	D				
Command	B0h							

NORMAL OUTPUTS - None

ERROR OUTPUTS - If the device does not support this command, if SMART is disabled or if the values in the Features, Cylinder Low or Cylinder High registers are invalid, an Aborted command error is posted.

Status register				Error register					
DRDY	DF	CORR	ERR	BBK	UNC	IDNF	ABRT	TK0NF	AMNF
V			V				V		

PREREQUISITES - DRDY set equal to one. SMART enabled.

DESCRIPTION - This command causes the device to immediately save any updated attribute values to the device's non-volatile memory regardless of the state of the attribute autosave timer. Upon receipt of this command from the host, the device sets BSY, writes any updated attribute values to non-volatile memory, clears BSY and asserts INTRQ.

6.7.9.7 SMART EXECUTE OFF-LINE COLLECTION IMMEDIATE

COMMAND CODE - B0h

PROTOCOL - Non-data command

INPUTS - The Features register shall be set to D4h. The Cylinder Low register shall be set to 4Fh. The Cylinder High register shall be set to C2h.

Register	7	6	5	4	3	2	1	0
Features	D4h							
Sector Count	na							
Sector Number	Self-Test Select (See Table below)							
Cylinder Low	4Fh							
Cylinder High	C2h							
Device/Head	obs	ns	obs	DEV	na	na	na	na
Command	B0h							

Task File register setup Initiate Self-Test

Register	Function	SMART
TF.0		NA
TF.1	Features	Function subcode: 0xD4 – Initiate Self Test
TF.2	Sector count	0 = no data transfer.
TF.3	Sector number:	Self-Test Select (see 9.2.2 below)
TF.4	Cylinder Low	Password low (= 4Fh)
TF.5	Cylinder High	Password High (= C2h)
TF.6	Drive /head	Drive select, head doesn't matter
TF.7	Command	0B0H (The SMART generic access command)

Self-Test Selection

The Sector Number register (TF.3) is used to select the particular test to be run: Short or Extended, Off-line Raw Read, and the mode- either Offline or Captive.

Value	Self- Test
0	Execute Off-line Full Scan
1	Execute SMART Short self-test routine immediately in Off-line mode
2	Execute SMART Extended self-test routine immediately in Off-line mode
3-63	Reserved
64-126	Vendor specific
127	Abort off-line mode self-test routine
128	Reserved
129	Execute SMART Short self-test routine immediately in captive mode
130	Execute SMART Extended self-test routine immediately in captive mode
131-191	Reserved
192-255	Vendor specific

Self-Test Registers

The following registers are embedded in the Read Attribute data Structure. This section describes the registers in the Read Attribute data structure that pertain to Self-Test and how they are used by the Self-Test Process.

Self-Test Execution Status

(1 byte- #363) Variable.

This register contains a “percent time remaining” nibble and an Execution Status nibble.

7	6	5	4	3	2	1	0
Self-test status values (0-15)				Percent of self-test remaining (0-9)			

Bits 0-3 (Percent Self-Test Remaining) The value in these bits indicate an approximation of the percent of the self-test routine remaining until completion in ten percent increments. Valid values are 0 through 9. A value of 9 indicates 90% of the total test time remaining. A value of 0 indicates the self-test routine is complete.

Bits 4-7 (Self-Test Execution Status) The value in these bits (See Table below) indicate the current Self-Test Execution Status.

Value	Description
0	The previous self-test routine completed without error or no self-test has ever been run
1	The self-test routine was aborted by the host
2	The self-test routine was interrupted by the host with a hard or soft reset
3	A fatal error or unknown test error occurred while the device was executing its self-test routine. The device was unable to complete the self-test routine.
4	The previous self-test completed having a test element that failed. The test element that failed is not known.
5	The previous self-test completed having a test element that failed. The read/write or electrical element of the test failed.
6	The previous self-test completed having a test element that failed. The servo (and/or seek) test element of the test failed.
7	The previous self-test completed having a test element that failed. The sequential read scan element of the test failed.
8-14	Reserved.
15	Self-test routine in progress

Self-Test Elements

Element	Description
1	Buffer Ram Test
2	Head Test
3	Seek Test
4	Partial Media Scan -or- Full Media Scan
5	Check SMART Threshold

ERROR OUTPUTS - If the device does not support this command, if SMART is disabled or if the values in the Features, Cylinder Low or Cylinder High registers are invalid, an Aborted command error is posted.

Register	7	6	5	4	3	2	1	0
Error	na	na	na	IDNF	na	ABRT	na	na
Sector Count	na							
Sector Number	na							
Cylinder Low	na							
Cylinder High	na							
Device/Head	obs	ns	obs	DEV	Head number or LBA			
Command	BSY	DRDY	DF	na	DRQ	na	na	ERR

PREREQUISITES - DRDY set equal to one. SMART enabled.

DESCRIPTION - Upon issuing of the SMART command sub-code D4 the drive will begin an off-line testing. Any ECC correctable errors encountered during the off-line testing of the drive are counted, and upon completion the attribute list is updated accordingly. Shown below are the off-line data collection status byte values

Value	Definition
00h or 80h	Off-line data collection was never started
02h or 82h	Off-line data collection was completed without error
04h or 84h	Off-line data collection was suspended by an interrupting command from the host
05h or 85h	Off-line data collection was aborted by an interrupting command from the host
06h or 86h	Off-line data collection activity was aborted by the device with a fatal error

6.7.10 Read Multiple C4h

The execution of the READ MULTIPLE command is identical to that of the Read Sectors command. However, the READ MULTIPLE command:

- Transfers blocks of data to the host without intervening interrupts
- Requires DRQ qualification of the transfer only at the start of the block—*not* at each sector
- Invalidates the Long bit

The SET MULTIPLE MODE command specifies the block count, or the number of sectors to be transferred as a block. This command executes prior to the READ MULTIPLE command. When the host issues a READ MULTIPLE command, the Sector Count Register contains the number of sectors requested—not the number of blocks or the block count.

If this sector count is not evenly divisible by the block count, the drive transfers as many full blocks as possible to the host, followed by a final partial-block transfer. The partial-block transfer is for *n* sectors, where:

n = (sector count) module (block count)

If the drive attempts execution of a READ MULTIPLE command before executing the SET MULTIPLE MODE command, or if READ MULTIPLE commands are disabled, an abort command error occurs.

The drive reports disk errors encountered during READ MULTIPLE commands at the beginning of a block or partial-block transfer. However, the drive still sets DRQ and transfers the data—including any corrupted data.

6.7.11 Write Multiple C5h

The execution of the WRITE MULTIPLE command is identical to that of the Write Sectors command. However, the WRITE MULTIPLE command:

- Causes the controller to set BSY within 400 ns of accepting the command
- Causes the drive to transfer multiple-sector blocks of data to the drive without intervening interrupts
- Requires DRQ qualification of the transfer only at the start of the block, not at each sector
- Invalidates the Long bit

The SET MULTIPLE MODE command specifies the block count, or the number of sectors to be transferred as a block. This command executes prior to the WRITE MULTIPLE command. When the host issues a WRITE MULTIPLE command, the Sector Count Register contains the number of sectors requested—not the number of blocks or the block count.

If this sector count is not evenly divisible by the block count, the drive transfers as many full blocks as possible, followed by a final partial-block transfer. The partial-block transfer is for n sectors, where:

$$n = \text{(sector count) module (block count)}$$

If the drive attempts to execute a WRITE MULTIPLE command before executing the SET MULTIPLE MODE command, or while WRITE MULTIPLE commands are disabled, an Abort Command error occurs.

During the execution of a WRITE MULTIPLE command, the drive reports all disk errors encountered, following an attempted disk write of the block or partial block. When an error occurs, the WRITE MULTIPLE command ends at the sector that contains the error—even if it is in the middle of a block. The drive generates interrupts by setting DRQ at the beginning of each block or partial block.

6.7.12 Set Multiple Mode C6h

The SET MULTIPLE MODE command enables the controller to perform READ MULTIPLE and WRITE MULTIPLE operations, and establishes the block count for these commands.

Prior to issuing a command, the host should load the Sector Count Register with the number of sectors per block. On receiving this command, the drive sets BSY and checks the contents of the Sector Count Register.

If the Sector Count Register contains a valid value, and the controller supports block count, the controller loads the values for all subsequent READ MULTIPLE and WRITE MULTIPLE commands, and enables execution of these commands. Any unsupported block count in the register causes an Aborted Command error, and disables execution of READ MULTIPLE and WRITE MULTIPLE commands.

If the Sector Count Register contains a zero value when the host issues the command, READ MULTIPLE and WRITE MULTIPLE commands are disabled. Any unsupported block count in the register causes an aborted command error, and disables READ MULTIPLE and WRITE MULTIPLE commands. After the

command is executed, the controller clears BSY. At power on, the default mode for the READ MULTIPLE and WRITE MULTIPLE commands is disabled.

6.7.13 Read DMA C8h

The READ DMA command allows the host to read data using the DMA data transfer protocol. The host should not use the C9h value.

6.7.14 Write DMA CAh

The WRITE DMA command allows the host to write data using the DMA data transfer protocol. The host should not use the CBh value.

6.7.15 Read Buffer E4h

The READ BUFFER command enables the host to read the current contents of the drive's sector buffer. When the host issues this command, the drive sets BSY, sets up the sector buffer for a read operation, sets DRQ, class BSY, and generates an interrupt. The host then reads up to 512 bytes of data from the buffer.

The drive can synchronize READ BUFFER and WRITE BUFFER commands from the host; that is, sequential READ BUFFER and WRITE BUFFER commands can access the same 512 bytes within the buffer.

6.7.16 Flush Cache

Feature Set General

The Flush Cache command is used by the host to request the device to flush the write cache. When the host issues this command, the drive sets BSY and proceeds to write all the cached data to the media. If an error occurs during the write operation the drive will try the error recovery routines to ensure the data is written successfully. The flushing of write cache may take several seconds to complete depending upon the amount of data to be flushed and the success of the operation.

INPUTS

Register	7	6	5	4	3	2	1	0
Features	na							
Sector Count	na							
Sector Number	na							
Cylinder Low	na							
Cylinder High	na							
Device/Head	obs	na	obs	DEV	na	na	na	na
Command	E7h							

ERROR OUTPUTS – If the command is not supported, the device posts an Aborted command error. An unrecoverable error encountered during execution of writing data results in the termination of the command and the Command Block registers contain the sector address of the sector where the first unrecoverable error occurred. The sector is removed from the cache. Subsequent FLUSH CACHE commands continue the process of flushing the cache.

Register	7	6	5	4	3	2	1	0
Error	na	na	na	na	na	ABRT	na	na
Sector Count	na							
Sector Number	na							
Cylinder Low	na							
Cylinder High	na							
Device/Head	obs	na	obs	DEV	HEAD number or LBA			
Command	BSY	DRDY	DF	na	DRQ	na	na	ERR

6.7.17 Write Buffer E8h

The WRITE BUFFER command allows the host to write 512 bytes of the drive's buffer. On receiving this command in its Command Block Register, the drive sets BSY and prepares for a write operation. When ready, the drive sets DRQ, resets BSY, and generates INTRQ, allowing the host to the buffer.

6.7.18 Power Management Commands

The Quantum Fireball lct20 AT hard disk drive provides numerous power management options. Two important options center around a count down counter known as the automatic power down counter or APD. This counter can trigger one of two power saving events depending on which of the two commands was most recently issued.

- **Standby:** Once a standby command is issued, the drive enters the standby mode. Further, each time the APD counter reaches zero in the future, the drive enters the standby mode, the spindle and actuator motors are off and the heads are parked in the landing zone. Receipt of any command that requires media access causes the drive to exit the standby command and service the host request. Each time the drive executes the standby command, the drive will reenter the standby mode when the APD counter reaches zero.
- **Idle:** Once an idle command is issued, each time the APD counter reaches zero, the drive enters the standby mode. In the standby mode, the actuator and spindle motors are off with the heads locked in the landing area. This is the default setting.

Three commands are available which are not dependent upon the APD counter reaching zero:

- **Sleep:** When a sleep command is received, the drive enters the sleep mode. In the sleep mode, the spindle and actuator motors are off and the heads are latched in the landing zone. Receipt of a reset causes the drive to transition from the sleep to the standby mode.
- **Standby Immediate:** When a standby immediate command is received, the drive immediately enters the standby mode.
- **Idle Immediate:** When an idle immediate command is received, after the first decrement of the APD counter, the drive enters the idle mode.

The sleep, standby immediate, and idle immediate commands differ in a significant way from the standby and idle commands. Specifically, sleep, standby immediate, and idle immediate have a one-time effect and must be reissued each time their effect is desired. In contrast, standby and idle operate in conjunction with the APD counter and stay in effect continually, becoming non-effectual only upon issuance of the other of these two commands. Thus, for example, once the standby command is issued just one time, each time the APD counter reaches zero the drive will enter the standby mode.

Note: The user has the ability to determine the value to which the APD counter is set upon completion of any command. This value is set by writing to the Sector Count Register a number between 12 and 255 just prior to issuance of a standby or idle command.

6.7.18.1 Standby Immediate Mode – E0h

The Standby Immediate Mode power command immediately puts the drive in the Standby Mode. Power is removed from the spindle motor (the drive's PCB power remains) and the heads are parked.

6.7.18.2 Idle Immediate Mode – E1h

The Idle Immediate Mode power command immediately puts the drive in the Idle Mode.

6.7.18.3 Standby Mode, Automatic Power-Down – E2h

The Standby Mode, Automatic Power-Down (APD) command immediately puts the drive in the Standby Mode. The Sector Count Register is then examined. If the value in this register is not zero, the Auto Power-Down feature is enabled and will take effect once the countdown timer reaches zero. The valid count range is Table 6-16. Each time the drive is accessed, the countdown timer is reset to the value originally set in the Sector Count Register at the time the Standby Mode-Auto Power Down command was issued.

Note: If the value in the Sector Count Register is zero, the Auto Power-Down feature is disabled.

Table 6-23 *Valid Count Range*

SECTOR COUNT	TIME
1 to 12	1 minute
13 to 240	(Value * 5) seconds
241 to 251	((Value – 240) * 30) minutes
252 to 255	(Value * 5) minutes

6.7.18.4 Idle Mode, Automatic Power-Down – E3h

The Idle Mode, Automatic Power-Down command immediately puts the drive into the Idle Mode. The Sector Count Register is then examined. If the value in this register is not zero, the Auto Power-Down feature is enabled and takes effect once the countdown timer reaches zero. The valid count range is listed in Table 6-16. Each time the drive is accessed, the countdown timer is reset to the value originally set in the Sector Count Register at the time the Idle Mode-Auto Power Down command was issued.

Note: If the value in the Sector Count Register is zero, the Auto Power-Down feature is disabled.

6.7.18.5 Read Buffer – E4h

The READ BUFFER command will be synchronized within the device such that sequential READ BUFFER commands access the same 512 bytes within the buffer.

6.7.18.6 Check Power Mode – E5h

The CHECK POWER MODE command writes FFh into the Sector Count Register provided that the drive is in the Idle Mode, even if it is in Automatic Power-Down mode. However, if it is in Standby mode, the drive returns a value of 00h in the Sector Count Register.

6.7.18.7 Sleep Mode – E6h

The Quantum Quantum Fireball lct20 AT drive considers the Sleep Mode to be the equivalent of the Standby Mode, except that a reset is required before issuing a command requiring media access.

6.7.18.8 Flush Cache – E7h

This command is used by the host to request the device to flush the write cache. If the write cache is to be flushed, all data cached will be written to the media. The BSY bit will remain set to one until all data has been successfully written or an error occurs. The device should use all error recovery methods available to ensure the data is written successfully. The flushing of write cache may take several seconds to complete depending upon the amount of data to be flushed and the success of the operation.

Note: This command may take longer than 30 seconds to complete.

6.7.18.9 Write Buffer – E8h

The WRITE BUFFER command will be synchronized within the device such that sequential WRITE BUFFER commands access the same 512 bytes within the buffer.

6.7.19 Identify Drive – ECh

The IDENTIFY DRIVE command enables the host to receive parameter information from the drive. When the host issues this command, the drive sets BSY, stores the required parameter information in the sector buffer, sets DRQ, and generates an interrupt. The host then reads the information from the sector buffer. The Identify Drive Parameters Table, shown in Table 6-24, defines the parameter words stored in the buffer. All reserved bits should be zeros. A full explanation of the parameter words is listed below:

Default Logical Cylinders: The number of translated cylinders in the default translation mode.

Number of Logical Heads: The number of translated heads in the default translation mode.

Number of Unformatted Bytes Per Track: The number of unformatted bytes per translated track in the default translation mode.

Number of Unformatted Bytes Per Sector: The number of unformatted bytes per sector in the default translation mode.

Number of Logical Sectors Per Track: The number of sectors per track in the default translation mode.

Serial Number: The contents of this field are left aligned and padded with spaces (20h).

Buffer Type: The contents of this field are as follows:

- 0000h = Not specified
- 0001h = A single-ported, single-sector buffer capable of data transfers either to or from the host or to or from the disk
- 0002h = A dual-ported, multiple-sector buffer capable of simultaneous data transfers either to and from the host, or from the host and the disk
- 0003h = A dual-ported, multiple-sector buffer capable of simultaneous data transfers with read caching
- 0004 – FFFFh = Reserved

Firmware Revision: The contents of this field are left-aligned and padded with spaces (20h).

Model Number: The contents of this field are left-aligned and padded with spaces (20h). The low-order byte appears first in a word.

Table 6-24 *Identify Drive Parameters*

WORDS[1]			PARAMETER DESCRIPTION (Statements below are true if the bit is set to 1)
WORD	BIT	BIT VALUE	
0	15	0	0 = ATA device
	14	0	Retired
	13	0	
	12	0	
	11	0	
	10	1	
	9	0	
	8	0	
	7	0	1 = Removable media
	6	1	1 = Not removable controller and/or device
	5	0	Retired
	4	1	
	3	1	
	2	0	
	1	1	
	0	0	Reserved

WORDS[1]			PARAMETER DESCRIPTION (Statements below are true if the bit is set to 1)
WORD	BIT	BIT VALUE	
1		10.0 AT = 16,383 20.0AT = 16,383 30.0AT = 16,383 40.0AT = 16,383	Default logical cylinders
2		0	Reserved
3		10.0AT = 16 20.0AT = 16 30.0AT = 16 40.0AT = 16	Default number of logical heads
4		Zone dependent	Retired
5		N/A	Retired
6		63	Default number of logical sectors per track
7–9		5154h	Retired
10–19			Serial number (20 ASCII characters)[2]
20		3	Buffer type
21		374h	Buffer size in 512-byte increments
22		4	Number of ECC bytes passed on READ/WRITE LONG commands
23–26			Firmware revision (8 ASCII characters)
27–46		QUANTUM FIREBALL lct20 10 QUANTUM FIREBALL lct20 20 QUANTUM FIREBALL lct20 30 QUANTUM FIREBALL lct20 40	Model number (40 ASCII characters)
47	15–8 7–0	80h 10h	Vendor Unique Maximum number of sectors that can be transferred per interrupt is set to 8 for READ MULTIPLE and WRITE MULTIPLE commands.
48		0	Cannot perform double word I/O
49	15–12 11 10 9 8 7–0	0 1 1 1 1 0	Reserved 1 = I/O Ready is supported 1 = I/O Ready can be disabled 1 = LBA supported 1 = DMA supported Vendor Unique

WORDS[1]			PARAMETER DESCRIPTION (Statements below are true if the bit is set to 1)
WORD	BIT	BIT VALUE	
50		4000h	Capabilities
51	15–8 7–0	4 0	PIO data-transfer cycle timing mode Vendor Unique
52	15–8 7–0	N/A	Retired
53	15–3 2 1 0	 1 0 1 0 1 0	Reserved The fields reported in word 88 are valid The fields reported in word 88 are not valid The fields reported in words 64–70 are valid The fields reported in words 64–70 are not valid The fields reported in words 54–58 are valid The fields reported in words 54–58 are not valid
54			Number of current cylinders
55			Number of current heads
56		X	Number of current sectors per track
57–58		X	Current capacity in sectors (CHS mode only)
59	15–9 8 7–0	0 1 n[3]	Reserved Multiple sector setting is valid Current setting for number of sectors that can be transferred per interrupt on R/W Multiple commands
60–61		10.0AT = 20,044,080 20.0AT = 39,876,480 30.0AT = 58,633,344 40.0AT = 78,177,792	Total number of User Addressable Sectors (LBA Mode only)
62	15–8 7–0	N/A	Retired
63	15–8 7–0	4 7	Multiword DMA transfer mode active (Mode 2) (Default) Multiword DMA transfer modes supported (Mode 2)
64		3	Advanced PIO Mode is supported
65		120	Minimum multiword DMA transfer cycle time (ns) per word
66		120	Manufacturer's recommended multiword DMA cycle time (ns)
67		120	Manufacturer's PIO cycle time (ns) without flow control
68		120	Manufacturer's PIO cycle time (ns) with flow control

WORDS[1]			PARAMETER DESCRIPTION (Statements below are true if the bit is set to 1)
WORD	BIT	BIT VALUE	
80		1Eh	Major version number 0000h or FFFFh = device does not report version
	4		1 = supports ATA/ATAPI-4
	3		1 = supports ATA-3
	2		1 = supports ATA-2
	1		1 = supports ATA-1
	0		Reserved
81		11h	Minor version number 0000h or FFFFh = device does not report version
82		3069h	Command set supported. If words 82 and 83 = 0000h or FFFFh command set notification not supported
	15		Obsolete
	14		1 = NOP command supported
	13		1 = READ BUFFER command supported
	12		1 = WRITE BUFFER command supported
	11		Obsolete
	10		1 = Host Protected Area feature set supported
	9		1 = DEVICE RESET command supported
	8		1 = SERVICE interrupt supported
	7		1 = release interrupt supported
	6		1 = look-ahead supported
	5		1 = write cache supported
	4		1 = supports PACKET command feature set
	3		1 = supports Power Management feature set
	2		1 = supports Removable Media feature set
	1		1 = supports Security Mode feature set
	0		1 = supports SMART feature set
83		4001h	Command sets supported. If words 82 and 83 =0000h or FFFFh command set notification is not supported.
	15		Will be cleared to zero
	14		Will be set to one
	13–5		Reserved
	4		1 = Removable Media Status Notification feature set supported
	3		1 = Advanced Power Management feature set supported
	2		1 = CFA feature set supported
	1		1 = READ/WRITE DMA QUEUED supported
	0		1 = DOWNLOAD MICROCODE command supported
84		4000h	Command set/feature supported extension. If words 82, 83, and 84 = 0000h or FFFFh command set notification extension is not supported.
	15		Will be cleared to zero
	14		Will be set to one
	13–0		Reserved

WORDS[1]			PARAMETER DESCRIPTION (Statements below are true if the bit is set to 1)
WORD	BIT	BIT VALUE	
85		3069h	Command set/feature enabled. If words 85, 86, and 87 = 0000h or FFFFh command set enabled notification is not supported.
	15		Obsolete
	14		1 = NOP command supported
	13		1 = READ BUFFER command supported
	12		1 = WRITE BUFFER command supported
	11		Obsolete
	10		1 = Host Protected Area feature set supported
	9		1 = DEVICE RESET command supported
	8		1 = SERVICE interrupt enabled
	7		1 = release interrupt enabled
	6		1 = look-ahead enabled
	5		1 = write cache enabled
	4		1 = supports PACKET command feature set
	3		1 = supports Power Management feature set
	2		1 = supports Removable Media feature set
	1		1 = Security Mode feature set enabled
86		1h	Command set/feature enabled. If words 85, 86, and 87 = 0000h or FFFFh command set enabled notification is not supported.
	15–5		Reserved
	4		1 = Removable Media Status Notification feature set enabled
	3		1 = Advanced Power Management feature set enabled
	2		1 = CFA feature set supported
	1		1 = READ/WRITE DMA QUEUED command supported
	0		1 = DOWNLOAD MICROCODE command supported
87		4000h	Command set/feature default. If words 85, 86, and 87 = 0000h or FFFFh command set default notification is not supported.
	15		Will be cleared to zero
	14		Will be set to one
	13–0		Reserved

WORDS[1]			PARAMETER DESCRIPTION (Statements below are true if the bit is set to 1)
WORD	**BIT**	**BIT VALUE**	
88	15–14		Reserved
	13	1	Ultra DMA mode 5 is active
		0	Ultra DMA mode 5 is not active
	12	1	Ultra DMA mode 4 is active
		0	Ultra DMA mode 4 is not active
	11	1	Ultra DMA mode 3 is active
		0	Ultra DMA mode 3 is not active
	10	1	Ultra DMA mode 2 is active
		0	Ultra DMA mode 2 is not active
	9	1	Ultra DMA mode 1 is active
		0	Ultra DMA mode 1 is not active
	8	1	Ultra DMA mode 0 is active
		0	Ultra DMA mode 0 is not active
	7–6		Reserved
	5	1	Ultra DMA mode 5 and below are supported
	4	1	Ultra DMA mode 4 and below are supported
	3	1	Ultra DMA mode 3 and below are supported
	2	1	Ultra DMA modes 2 and below are supported
	1	1	Ultra DMA modes 1 and below are supported
	0	1	Ultra DMA mode 0 is supported
93	15–14		Reserved
	13	1	80 conductor if both device and host support method
		0	40 conductor if both device and host support method
	0-12		Reserved

1. The format of an ASCII field specifies that, within a word boundary, the low-order byte appears first.
2. The serial number has the following format: 00QTMTCYJJJLSSSSBBB
where:00 = Placeholders
   QT = Quantum
   M = Place of manufacture
   T = Drive type family (fixed at 2)
   C = Equal to the number of physical heads
   Y = Last digit of year drive built
   JJJ = Julian date
   L = Manufacturing production line
   SSSS = Sequence of manufacture
   BBB = Blanks (placeholders)
3. n is a variable from zero to 16.
4. Model Numbers are byte swapped for readability. See Note 1.

6.7.20 Set Features EFh

The SET FEATURES command is used by the host to establish certain parameters which control execution of the following drive features:

- 02h – Enable write cache feature
- 03h – Set transfer mode based on value in Sector Count Register
- 55h – Disable read look-ahead feature
- 66h – Disable reverting to power on defaults
- 82h – Disable write cache feature
- AAh – Enable read look-ahead feature
- CCh – Enable reverting to power on defaults

At power-on, or after a reset accomplished by either the hardware or software, the default mode is 4 bytes of ECC, read look-ahead, and write cache enabled.

A host can choose the transfer mechanism by Set Transfer Mode and specifying a value in the Sector Count register. The upper 5 bits define the type of transfer and the low order 3 bits encode the mode value.

6.7.21 Set Features (Ultra ATA/66)

Set Transfer Mode

A host can choose the transfer mechanism by the Set Transfer mode, subcommand code 03h, and by specifying a value in the Sector Count register. The upper 5 bits define the type of transfer and the low order 3 bits encode the mode value.

Table 6-25 *Transfer/Mode Values*

PIO Default Transfer Mode	00000	000
PIO Default Transfer Mode, Disable IORDY	00000	001
PIO Flow Control Transfer Mode x	00001	nnn
Multiword DMA Mode x	00100	nnn
Ultra DMA Mode	01000	nnn
Reserved	10000	nnn
Key: nnn = a valid mode number in binary x = a mode number in decimal for the associated transfer type		

6.7.22 Read Defect List

The READ DEFECT LIST command enables the host to retrieve the drive's defect list. Prior to issuing the Read Defect List command the host should issue the Read Defect List Length command. This command will not transfer any data. It instead, stores the length in sectors of the defect list in the Sector Count register (1F2), and the Sector Number register (1F3), with the Sector Count register containing the LSB of the 2-byte value (see Table 6-26). The defect list length is a fixed value for each Quantum product and can be calculated as follows:

length in sectors = (((maximum number of defects) * 8 + 4) + 511)/512

At the completion of the command, the task file registers 1F2 – 1F6 will contain bytes necessary to execute the Read Defect List command, and the host will only need to write the extended command code (F0h) to the Command register (1F7) to proceed with the Read Defect List command execution.

Table 6-26 *READ DEFECT LIST LENGTH Command Bytes*

ADDRESS	VALUE (Before)	DEFINITION	VALUE (After)
1F2	0	Defect List Subcode	Length in Sectors (LSB)
1F3	FFh	Password	Length in Sectors (MSB)
1F4	FFh	Password	FFh
1F5	3Fh	Password	3Fh
1F6	AXh (Drive 0)	Drive Select	AXh = Drive 0
	BXh (Drive 1)	—	BXh = Drive 1
1F7	F0h	Extended Command Code	Status Register

Note: Registers 1F2h through 1F5h must contain the exact values shown. These values function as a key. The drive issues the message ILLEGAL COMMAND if the bytes are not entered correctly.

The AT Read Defect List command is an extended AT command that enables the host to retrieve the drive's defect list. The host begins by writing to address 1F6h to select the drive. Then the host writes to addresses 1F2h – 1F5h using values indicated in Table 6-27. When the host subsequently writes the extended command code F0h to address 1F7h, the drive sets BSY, retrieves the defect list, sets DRQ, and resets BSY. The host can now read the requested number of sectors (512 bytes) of data. An INTRQ precedes each sector. Bytes 1F2h and 1F3h contain the 2-byte number of sectors that the host expects to read, with address 1F2h containing the LSB (see Table 6-27). The sector count (1F2h – 1F3h) may vary from product to product and if the wrong value is supplied for a specific product, the drive will issue the ILLEGAL COMMAND message. If the host does not know the appropriate sector count for a specific product, it can issue the Read Defect List Length command, described in the previous section to set up the task file for the Read Defect List command.

Table 6-27 *AT READ DEFECT LIST Command Bytes*

ADDRESS	VALUE	DEFINITION
1F2	Length in Sectors (LSB)	Defect List Subcode
1F3	Length in Sectors (MSB)	Defect List Subcode
1F4	FFh	Password
1F5	3Fh	Password
1F6	AXh = Drive 0	Drive Select
	BXh = Drive 1	—
1F7	F0h	Extended Command Code

Note: Registers 1F2h and 1F3h must contain the transfer length that is appropriate for the specific product, and 1F4h and 1F5h must contain the exact values shown. These values function as a key. The drive issues the message ILLEGAL COMMAND if the bytes are not entered correctly.

Pending defects will be excluded from the list, since no alternate sector is being used as their replacement, and since they may be removed from the drive's internal pending list at a later time. Table 6-28 shows the overall format of the defect list, and Table 6-29 shows the format of the individual defect entries.

Table 6-28 *DEFECT LIST DATA FORMAT*

BYTE	DESCRIPTION
0	0
1	1Dh
2	8* (Number of Defects) (MSB)
3	8* (Number of Defects) (LSB)
4–11	Defect Entry #1
12–19	Defect Entry #2
	•
	•

Table 6-29 *DEFECT ENTRY DATA FORMAT*

BYTE	DESCRIPTION
0	Defect cylinder (MSB)
1	Defect cylinder
2	Defect cylinder (LSB)
3	Defect head
4	Defect sector (MSB)
5	Defect sector
6	Defect sector
7	Defect sector (LSB)

Note: Bytes 4 – 7 will be set to FFh for bad track entries.

6.7.23 Configuration

In addition to the SET FEATURES command, the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives provide two configuration commands:

- The SET CONFIGURATION command, which enables the host to change DisCache and Error Recovery parameters
- The READ CONFIGURATION command, which enables the host to read the current configuration status of the drive

See Chapter 5 for more details about DisCache and setting cache parameters. See Chapter 5 also for more information about error detection and defect management.

6.7.23.1 Read Configuration

The READ CONFIGURATION command displays the configuration of the drive. Like the SET CONFIGURATION command, this command is secured to prevent accidentally accessing it. To access the READ CONFIGURATION command, you must write the pattern shown in Table 6-30 to the Command Block Registers. The first byte, 01h, is a subcode to the extended command code, F0h.

Table 6-30 *Accessing the READ CONFIGURATION Command*

ADDRESS	VALUE	DEFINITION
1F2h	01h	Read Configuration Subcode
1F3h	FFh	Password
1F4h	FFh	Password
1F5h	3Fh	Password
1F6h	AXh (Drive 0)	Drive Select
	BXh (Drive 1)	Drive Select
1F7h	F0h	Extended Command Code

Note: In Table 6-30:

Only the value in address 1F2h of the Command Block Registers is different from the SET CONFIGURATION command.

Registers 1F2h through 1F5h must contain the exact values shown in Table 6-30. These values function as a key. The drive issues the message ILLEGAL COMMAND if the key is not entered correctly.

To select the drive for which the configuration is to be read, set register 1F6h. For execution of the command to begin, load register 1F7h with F0h.

A 512-byte data field is associated with the READ CONFIGURATION command. A 512-byte read sequence sends this data from the drive to the host. The information in this data field represents the current settings of the configuration parameters. The format of the READ CONFIGURATION command data field is similar to that for the data field of the SET CONFIGURATION command, shown

in Table 6-31. However, in the READ CONFIGURATION command, bytes 0 through 31 of the data field are *not* KEY information, as they are in the SET CONFIGURATION command. The drive reads these bytes as *QUANTUM CONFIGURATION,* followed by eleven spaces. Users can read the configuration into a buffer, then alter the configuration parameter settings.

6.7.23.2 Set Configuration – FEh/FFh

The SET CONFIGURATION command is secured to prevent accessing it accidentally. To access the SET CONFIGURATION command, you must write the pattern shown in
Table 6-31 to the Command Block Registers. The first byte, FFh, is a subcode to the extended command code F0h.

Table 6-31 *Accessing the SET CONFIGURATION Command*

ADDRESS	VALUE	DEFINITION
1F2h	FFh	Set Configuration Subcode
1F3h	FFh	Password
1F4h	FFh	Password
1F5h	3Fh	Password
1F6h	AXh (Drive 0)	Drive Select
	BXh (Drive 1)	Drive Select
1F7h	F0h	Extended Command Code

Note: Registers 1F2h through 1F5h must contain the exact values shown above. These values function as a key. The drive issues the message ILLEGAL COMMAND if the key is not entered correctly.
To select the drive being reconfigured, register 1F6h should be set. For execution of the command to begin, load register 1F7h with F0h.

6.7.23.3 Set Configuration Without Saving to Disk

The SET CONFIGURATION WITHOUT SAVING TO DISK command is secured to prevent accidentally accessing it. To access this command, you must write the pattern shown in Table 6-32 to the Command Block Registers. The first byte, FEh, is a subcode to the extended command code F0h.

Table 6-32 *Accessing the SET CONFIGURATION WITHOUT SAVING TO DISK Command*

ADDRESS	VALUE	DEFINITION
1F2h	FEh	Set Configuration Subcode
1F3h	FFh	Password
1F4h	FFh	Password
1F5h	3Fh	Password
1F6h	AXh (Drive 0)	Drive Select
	BXh Drive 1)	Drive Select
1F7h	F0h	Extended Command Code

Note: In Table 6-32:
Registers 1F2h through 1F5h must contain the exact values shown above. These values function as a key. The drive issues the message ILLEGAL COMMAND if the key is not entered correctly.
To select the drive being reconfigured, set register 1F6h. For execution of the command to begin, load register 1F7h with F0h.

6.7.23.4 Configuration Command Data Field

A 512-byte data field is associated with this command. This data field is sent to the drive through a normal 512-byte write handshake. Table 6-33 shows the format of the data field. Bytes 0 through 31 of the data field contain additional KEY information. The drive issues the message ILLEGAL COMMAND if this information is not entered correctly. Bytes 32 through 35 control the operation of DisCache. Bytes 36 through 38 control operation of the error recovery procedure. The drive does not use bytes 40 through 511, which should be set to 0.

Table 6-33 *Configuration Command Format*

BYTE	BIT							
	7	6	5	4	3	2	1	0
0–31	QUANTUM CONFIGURATION KEY							
32	RESERVED = 0						PE	CE
33	RESERVED							
34	RESERVED = 0							
35	RESERVED = 0							
36	AWRE	ARR	N/A	RC	EEC	N/A	N/A	DCR
37	NUMBER OF RETRIES							
38	ECC CORRECTION SPAN							
39	RESERVED = 0					WCE	RUEE	0
40–511	RESERVED = 0							

Note: All fields marked RESERVED or N/A should be set to zero.

6.7.23.5 Quantum Configuration Key (Bytes 0-31)

Bytes 0–6 must contain the ASCII characters *Q, U, A, N, T, U,* and *M*; byte 7, the ASCII character *space*; and bytes 8–20 must contain the ASCII characters *C, O, N, F, I, G, U, R, A, T, I, O,* and *N*. Bytes 21–31 must contain an ASCII *space*. If this information is not entered correctly, the drive aborts the COMMAND.

6.7.23.6 DisCache Parameters

PE – Prefetch Enable (Byte 32, Bit 1): When set to 1, this bit indicates that the drive will perform prefetching. A PE bit set to 0 indicates that no prefetching will occur. The CE bit (bit 0) must be set to 1 to enable use of the PE bit. The default value is 1.

CE – Cache Enable (Byte 32, Bit 0): When set to 1, this bit indicates that the drive will activate caching on all READ commands. With the CE bit set to 0, the drive will disable caching and use the RAM only as a transfer buffer. The default setting is 1.

6.7.23.7 Error Recovery Parameters

AWRE – Automatic Write Reallocation Enabled (Byte 36, Bit 7): When set to 1, indicates that the drive will enable automatic reallocation of bad blocks. Automatic Write Reallocation is similar to the function of Automatic Read Reallocation, but is initiated by the drive when a defective block has become inaccessible for writing. An AWRE bit set to 0 indicates that the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT drives will not automatically reallocate bad blocks. The default setting is 1.

ARR – Automatic Read Reallocation (Byte 36, Bit 6): When set to 1, this bit indicates that the drive will enable automatic reallocation of bad sectors. The drive initiates reallocation when the ARR bit is set to 1 and the drive encounters a hard error—that is, if the triple-burst ECC algorithm is invoked. The default setting is 1. When the ARR bit is set to 0, the drive will not perform automatic reallocation of bad sectors. If RC (byte 36, bit 4) is 1, the drive ignores this bit. The default value

is 1.

RC – Read Continuous (Byte 36, Bit 4): When set to 1, this bit instructs the drive to transfer data of the requested length without adding delays to increase data integrity—that is, delays caused by the drive's error-recovery procedures. With RC set to 1 to maintain a continuous flow of data and avoid delays, the drive may send data that is erroneous. When the drive ignores an error, it does *not* post the error. The RC bit set to 0 indicates that potentially time-consuming operations for error recovery are acceptable during data transfer. The default setting is 0.

EEC – Enable Early Correction (Byte 36, Bit 3): When set to 1, this bit indicates that the drive will use its ECC algorithm if it detects two consecutive equal, nonzero error syndromes. The drive will not perform rereads before applying correction, unless it determines that the error is uncorrectable. An EEC bit set to 0 indicates that the drive will use its normal recovery procedure when an error occurs: rereads, followed by error correction. If the RC bit (byte 36, bit 4) is set to 1, the drive ignores the EEC bit. The default setting is 0.

DCR – Disable Correction (Byte 36, Bit 0): When set to 1, this bit indicates that all data will be transferred without correction, even if it would be possible to correct the data. A DCR bit set to 0 indicates that the data will be corrected if possible. If the data is uncorrectable, it will be transferred without correction, though the drive will attempt rereads. If RC (byte 36, bit 4) is set to 1, the drive ignores this bit. The default setting is 0. The drive will post all errors, whether DCR is set to 0 or 1.

NUMBER OF RETRIES (Byte 37): This byte specifies the number of times that the drive will attempt to recover from data errors by rereading the data, before it will apply correction. The drive performs rereads before ECC correction—unless EEC (byte 36, bit 3) is set to 1, enabling early correction. The default is eight.

ECC CORRECTION SPAN (Byte 38): This byte specifies the maximum number of bits per interleave that can be corrected using quadruple-burst ECC. The default value for this byte is 20h or 32 decimal.

6.7.23.8 Drive Parameters

WCE – Write Cache Enable (Byte 39, Bit 2): When this bit is set to1, the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives enable the Write Cache. This indicates that the drive returns GOOD status for a write command after successfully receiving the data, but before writing it to the disk. A value of zero indicates that the drive returns GOOD status for a write command after successfully receiving the data and writing it to the disk.

If the next command is another WRITE command, cached data continues to be written to the disk while new data is added to the buffer. The default setting is 1.

RUEE – Reallocate Uncorrectable Error Enables (Byte 39, Bit 1): When set to 1, this bit indicates that the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives will automatically reallocate uncorrectable hard errors, if the ARR bit (byte 36, bit 6) is set to 1. The default setting is 1.

6.7.24 Host Protected Mode Feature

These commands allow for a data storage area outside the normal operating system file area. Systems can use this area to store configuration data or save memory to the device in a location that the operating system cannot change.

6.7.24.1 Read Native Max Address.

Feature Set: Host Protected Area

This command return the native maximum address. The native maximum address is the highest address accepted by the device in the factory default condition. The native maximum address is the maximum address that is valid when using the SET MAX ADDRESS command.

Register	7	6	5	4	3	2	1	0
Features	na							
Sector Count	na							
Sector Number	na							
Cylinder Low	na							
Cylinder High	na							
Device/Head	obs	LBA	obs	DEV	na	na	na	na
Command	F8h							

If LBA is set to one, the maximum address shall be reported as an LBA value. If LBA is cleared to 0, the maximum address shall be reported as a CHS value.

6.7.24.2 Set Max Address F9h

Feature Set: Host Protected Area

This command allows the host to redefine the maximum address of the user-accessible address space in either LBA translation or the current CHS translation. The host may either set a new maximum cylinder number for CHS translation or a new LBA address for LBA translation. After successful SET MAX address the content of the IDENTIFY DEVICE parameter table will change in the following manner:

Using new max cylinder:

- The content of words 3,6,55 and 56 are unchanged.
- Word 1 = lesser of New set max cylinder + 1 or 16,383
- Word 60:61 =(Word 1XWord3XWord6)

If word 60:61 is less than 16,514,064 then word 54 will be the result of [(words60:61)/ [(Word55)X(word56)]] or 65,535 whichever is less.

If word 60:61 is greater than 16,514,064 then word 54 will be the result of [(16,514,064)/[(Word55)X(word56)]] or 65,535 whichever is less. The content of words 58:57 will be equal to [(the new content of word54)X(word55)X(word56)]

Using max LBA:

- The content of words 3,6,55 and 56 are unchanged.
- Word 60:61 new Maximum LBA +1.

If word 60:61 is less than 16,514,064 then Word1= [(new content of word60:61) / [(word3Xword6)]] or 65535 whichever is less.

If word 60:61 is greater than 16,514,064 then word 1 will be equal to 16,383.

If word 60:61 is less than 16,514,064 then word 54 will be the result of [(words60:61)/ [(Word55)X(word56)]]

If word 60:61 is greater than 16,514,064 then word54 will be equal to 16,383.

Words58:57 will be equal to [(word54)X(word55)X(word56)].

Register	7	6	5	4	3	2	1	0
Features	na							
Sector Count	na							
Sector Number	na							
Cylinder Low	na							
Cylinder High	na							
Device/Head	obs	LBA	obs	DEV	na	na	na	na
Command	F9h							

6.8 ERROR REPORTING

At the start of a command's execution, the Quantum Fireball lct20 10.0/20.0/30.0/40.0 GB AT hard disk drives check the Command Register for any conditions that would lead to an abort command error. The drive then attempts execution of the command. Any new error causes execution of the command to terminate at the point at which it occurred. Table 6-34 lists the valid errors for each command.

Table 6-34 *Command Errors*

COMMAND	ERROR REGISTER						STATUS REGISTER				
	BBK	UNC	IDNF	ABRT	TKO	AMNF	DRDY	DF	DSC	CORR	ERR
Check Power Mode				**V**			V	V	V		V
Execute Drive Diag.											V
Flush Cache				V			V	V			V
Format Track			V	V			V	V	V		V
Identify Drive				V			V	V	V		V
Initialize Parameters							V	V	V		
Invalid Cmnd. Codes				V			V	V	V		V
Read Buffer				V			V	V	V		V
Read Configuration	V	V	V	V		V	V	V	V	V	V
Read Defect List	V	V	V	V		V	V	V	V	V	V
Read DMA	V	V	V	V		V	V	V	V	V	V
Read Multiple	V	V	V	V		V	V	V	V	V	V
Read Sectors	V	V	V	V		V	V	V	V	V	V
Read Verify Sectors	V	V	V	V		V	V	V	V	V	V
Recalibrate				V	V		V	V	V		V
Seek			V	V			V	V	V		V
Set Configuration	V		V	V			V	V	V		V
Set Features				V			V	V	V		V
Set Multiple Mode				V			V	V	V		V
Write Buffer				V			V	V	V		V
Write DMA	V		V	V			V	V	V		V
Write Multiple	V		V	V			V	V	V		V
Write Sectors	V		V	V			V	V	V		V
Set Host Protected ARCA Commands											

Note: V = Valid errors for each command
ABRT=Abort command error
AMNF=Data address mark not found error
BBK =Bad block detected
CORR=Corrected data error
DRDY=Drive ready

DSC=Disk seek complete not detected
DF=Device fault detected
ERR=Error bit in the Status Register
IDNF=Requested ID not found
TK0=Track zero not found error
UNC=Uncorrectable data error

A

ACCESS – (v) Read, write, or update information on some storage medium, such as a disk. (n) One of these operations.

ACCESS TIME – The interval between the time a request for data is made by the system and the time the data is available from the drive. Access time includes the actual 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 rotary voice coil and the head mounting arms. One end of each head mounting arm attaches to the rotor with the read/write heads attached at the opposite end of each arm. As current is applied to the rotor, it rotates, positioning the heads over the desired cylinder on the media.

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 particular files. 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.

AVERAGE SEEK TIME – The average time it takes for the read/write head to move to a specific location. To compute the average seek time, you divide the time it takes to complete a large number of random seeks all over the disk 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 because of a media flaw or damaged format markings.

BAD TRACK TABLE – A label affixed to the casing of a hard disk drive that tells which tracks are flawed and cannot hold data. The listing is typed into the low-level formatting program when the drive is being installed. Because Quantum disk drive's defect-management scheme handles all such flaws automatically, there is no need to concern yourself with bad track tables.

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 1or logic 0 binary settings that make up a byte of data. See also byte.

BLOCK – A sector or group of sectors. By default, a block of data consists of 512 bytes.

BPI – Abbreviation for *bits per inch*. A measure of how densely information is packed on a storage medium. Flux changes per inch is also a term commonly used in describing storage density on a magnetic surface.

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 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 – Random-access memory used as a buffer between the CPU and a hard disk. Information more likely to be read or changed is placed in the cache, where it can be accessed more quickly to speed up general data flow.

CAPACITY – The amount of information that can be stored on a disk drive. The data is stored in bytes, and capacity is usually expressed in megabytes.

CDB – Command Descriptor Block. The SCSI structure used to communicate requests from an initiator (system) to a target (drive).

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 – A group of sectors on a disk drive that is addressed as one logical unit by the operating system.

CONTROLLER – Short form of *disk controller*. The chip or complete circuit that translates computer data and commands into a form suitable for use by the disk drive.

CONTROLLER CARD – An adapter holding the control electronics for one or more hard disks, usually installed in a slot in the computer.

CPU – Acronym for *Central Processing Unit*. The microprocessor chip that performs the bulk of data processing in a computer.

CRC – Acronym for *Cyclic Redundancy Check*. An error detection code that is recorded within each sector and is used to see whether parts of a string of data are missing or erroneous.

CYLINDER – On a disk drive that has more than one recording surface and heads that move to various tracks, the group of all tracks located at a given head position. The number of cylinders times the number of heads equals the number of tracks per drive.

D

DATA SEPARATOR – On a disk drive that stores data and timing information in an encoded form, the circuit that extracts the data from the combined data and clock signal.

DEDICATED SERVO – A surface separate from the surface used for data that contains only disk timing and positioning information and contains no data.

DEFECT MANAGEMENT – A method that is implemented to ensure long term data integrity. Defect management eliminates the need for user defect maps. This is accomplished by scanning the disk drives at the factory for defective sectors. Defective sectors are deallocated prior to shipment. In addition, during regular use, the drive continues to scan and compensate for any new defective sectors on the disk.

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 disk drives all have controllers embedded on the drive printed-circuit board.

DISKWARE – The program instructions and data stored on the disk for use by a processor.

DMA – Acronym for *direct memory access*. A process by which data moves directly between a disk drive (or other device) and system memory without passing through the CPU, thus allowing the system to continue processing other tasks while the new data is being retrieved.

DRIVE – Short form of *disk drive*.

DRIVE GEOMETRY – The functional dimensions of a drive in terms of the number of heads, cylinders, and sectors per track. See also logical format.

E

ECC – Acronym for *error correction code*. The recording of extra verifying information encoded along with the disk data. The controller uses the extra information to check for data errors, and corrects the errors when possible.

EMBEDDED SERVO – A timing or location signal placed on the disk's surface on the tracks that also store data. These signals allow the actuator to fine-tune the position of the read/write heads.

ENCODING – The protocol by which particular data patterns are changed prior to being written on the disk surface as a pattern of On and Off or 1 and 0 signals.

EXTERNAL DRIVE – A drive mounted in an enclosure separate from the PC or computer system enclosure, with its own power supply and fan, and connected to the system by a cable.

F

FAT – Acronym for *file allocation table*. A data table stored on the outer edge of a disk that tells the operating system which sectors are allocated to each file and in what order.

FCI – Acronym for *flux changes per inch*. 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.

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.

FLYING HEIGHT – The distance between the read/write head and the disk surface caused by a cushion of air that keeps the head from contacting the media. Smaller flying heights permit more dense storage of data, but require more precise mechanical designs.

FORMAT – To write onto the disk surface a magnetic track pattern that specifies the locations of the tracks and sectors. This information must exist on a disk before it can store any user data. Formatting erases any previously stored data.

FORMATTED CAPACITY – The amount of room left to store data on the disk after the required space has been used to write sector headers, boundary definitions, and timing information generated by a format operation. All Quantum drive capacities are expressed in formatted capacity.

FORM FACTOR – The physical outer dimensions of a device as defined by industry standard. For example, most Quantum 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 disk drive mounted in an IBM AT and compatible computers so that the drive easily slides into place.

H

HALF HEIGHT – Term used to describe a drive that occupies half the vertical space of the original full size 5 1/4-inch drive. 1.625 inches high.

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 repeatable error in disk data that persists when the disk is reread, usually caused by defects in the media surface.

HEAD – The tiny electromagnetic coil and metal pole piece used to create and read back the magnetic patterns (write and read information) on the media.

HIGH-CAPACITY DRIVE – By industry conventions typically a drive of 1 gigabytes or more.

HIGH-LEVEL FORMATTING – Formatting performed by the operating system's format program. Among other things, the formatting program creates the root directory and file allocation tables. See also low-level formatting.

HOME – Reference position track for recalibration of the actuator, usually the outer track (track 0).

HOST ADAPTER – A plug-in board that forms the interface between a particular type of computer system bus and the disk drive.

I

INITIALIZE – See low level formatting.

INITIATOR – A SCSI device that requests another SCSI device to perform an operation. A common example of this is a system requesting data from a drive. The system is the initiator and the drive is the target.

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.

INTERLEAVE – The arrangement of sectors on a track. A 1:1 interleave arranges the sectors so that the next sector arrives at the read/write heads just as the computer is ready to access it. See also interleave factor.

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, which is installed in one of the computer's slots).

J

JUMPER – A tiny box that slips over two pins that protrude from a circuit board. When in place, the jumper connects 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 position inside the disk's inner cylinder in a non data area reserved as a place to rest the heads during the time that power is off. Using this area prevents the heads from touching the surface in data areas upon power down, adding to the data integrity and reliability of the disk drive.

LATENCY – The period of time during which the read/write heads are waiting for the data to rotate into position so that it can be accessed. Based on a disk rotation speed of 3,662 rpm, the maximum latency time is 16.4 milliseconds, and the average latency time is 8.2 milliseconds.

LOGICAL FORMAT – The logical drive geometry that appears to an AT system BIOS as defined by the drive tables and stored in CMOS. With an installation program like Disk Manager, the drive can be redefined to any logical parameters necessary to adapt to the system drive tables.

LOOK AHEAD – The technique of buffering data into cache RAM by reading subsequent blocks in advance to anticipate the next request for data. The look ahead technique speeds up disk access of sequential blocks of data.

LOW-LEVEL FORMATTING – Formatting that creates the sectors on the platter surfaces so the operating system can access the required areas for generating the file structure. Quantum drives are shipped with the low-level formatting already done.

LOW PROFILE – Describes drives built to the 3 1/2-inch form factor, which are only 1 inch high.

M

MB – See megabyte.

MEDIA – The magnetic film that is deposited or coated on an aluminum substrate which is very flat and in the shape of a disk. The media is overcoated with a lubricant to prevent damage to the heads or media during head take off and landing. The media is where the data is stored inside the disk in the form of magnetic flux or polarity changes.

MEGABYTE (MB) – A unit of measurement equal to 1,024 kilobytes, or 1,048,576 bytes except when referring to disk storage capacity.

1 MB = 1,000,000 bytes when referring to disk storage capacity.

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. Used as a reliability rating to determine the expected life of the product expressed in power on hours (POH). There are several accepted methods for calculating this value that produce very different results and generate much confusion in the industry. When comparing numbers you should first verify which method was used to calculate the values.

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 – The processing time of a command by the controller, host adapter or drive prior to any actual disk accesses taking place.

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 devoted to a particular operating system and accessed as one logical volume by the system.

PERFORMANCE – A measure of the speed of the drive during normal operation. Factors affecting performance are seek times, transfer rate and command overhead.

PERIPHERAL – A device added to a system as an enhancement to the basic CPU, such as a disk drive, tape drive or printer.

PHYSICAL FORMAT – The actual physical layout of cylinders, tracks, and sectors on a disk drive.

PLATED MEDIA – Disks that are covered with a hard metal alloy instead of an iron-oxide compound. Plated disks can store greater amounts of data in the same area as a coated disk.

PLATTER – An disk made of metal (or other rigid material) that is mounted inside a fixed disk drive. Most drives use more than one platter mounted on a single spindle (shaft) to provide more data storage surfaces in a small package. The platter is coated with a magnetic material that is used to store data as transitions of magnetic polarity.

POH – Acronym for *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 MTBF.

POSITIONER – See actuator.

R

RAM – Acronym for *random access memory*. An integrated circuit memory chip which allows information to be stored and retrieved by a microprocessor or controller. The information may be stored and retrieved in any order desired, and the address of one storage location is as readily accessible as any other.

RAM DISK – A “phantom disk drive” for which a section of system memory (RAM) is set aside to hold data, just as if it were a number of disk sectors. The access to this data is extremely fast but is lost when the system is reset or turned off.

READ AFTER WRITE – A mode of operation that has the computer read back each sector on the disk, checking that the data read back is the same as recorded. This slows disk operations, but raises reliability.

READ VERIFY – A disk mode where the disk reads in data to the controller, but the controller only checks for errors and 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 used on some hard disks to encode data into magnetic pulses. RLL requires more processing, but stores almost 50% more data per disk than the MFM method.

ROM – Acronym for *read only memory*. Usually in the form of an ROM in the controller that contains programs that can be accessed and read but not modified by the system.

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 the average, it is half of the time needed for a full rotation (about 8 ms.).

S

SCSI – Acronym for *Small Computer System Interface*, an American National Standards Institute (ANSI) version of Shugart Associates' SASI interface between the computer and controller. SCSI has grown in popularity and is one of the most flexible and intelligent interfaces available.

SECTOR – A section of space along a track on the disk, or the data that is stored in that section. Hard disks most often have sectors that are 512 data bytes long plus several bytes overhead for error correcting codes. Each sector is preceded by ID data known as a header, which cannot be overwritten.

SEEK – A movement of the disk read/write head in or out to a specific 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 when a track to track movement of the head stops, and when the residual vibration and movement dies down to a level sufficient for reliable reading or writing.

SHOCK RATING – A rating (expressed in Gs) of how much shock a disk drive can sustain without damage.

SOFT ERROR – An error in reading data from the disk that does not recur if the same data is reread. Often caused by power fluctuations or noise spikes.

SOFT SECTORED – Disks that mark the beginning of each sector of data within a track by a magnetic pattern.

SPINDLE – The center shaft of the disk upon which the drive's platters are mounted.

SPUTTER – A type of coating process used to apply the magnetic coating to some high-performance disks. In sputtering, the disks are placed in a vacuum chamber and the coating is vaporized and deposited on the disks. The resulting surface is hard, smooth, and capable of storing data at high density. Quantum disk drives use sputtered thin film disks.

STEPPER – A type of motor that moves in discrete amounts for each input electrical pulse. Stepper motors used to be widely used for read/write head positioner, since they can be geared to move the head one track per step. Stepper motors are not as fast or reliable as the rotary voice coil actuators which Quantum disk drives use.

SUBSTRATE – The material the disk platter is made of beneath the magnetic coating. Hard disks are generally made of aluminum or magnesium alloy (or glass, for optical disks) while the substrate of floppies is usually mylar.

SURFACE – The top or bottom side of the platter 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, used for disk surfaces. Thin film surfaces allow more bits to be stored per disk.

TPI – Acronym for *tracks per inch*. The number of tracks or cylinders that are written in each inch of travel across the surface of a disk.

TRACK – One of the many concentric magnetic circle patterns written on a disk surface as a guide to where to store and read the data.

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. Drive specifications usually reference a high number that is the burst mode rate for transferring data across the interface from the disk buffer to system RAM. Sustained data transfer is at a much lower rate because of system processing overhead, head switches, and seeks.

U

UNFORMATTED CAPACITY – The total number of bytes of data that could be fit onto a disk. Formatting the disk requires some of this space to record location, boundary definitions, and timing information. After formatting, user data can be stored on the remaining disk space, known as formatted capacity. The size of a Quantum drive is expressed in formatted capacity.

V

VOICE COIL – A type of motor used to move the disk read/write head in and out to the right track. Voice-coil actuators work like loudspeakers with the force of a magnetic coil causing a proportionate movement of the head. Quantum's actuator uses voice-coil technology, and thereby eliminates the high stress wearing parts found on stepper motor type actuators.

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 – Hard disks that use a technology similar to an IBM model using Winchester as the code name. These disks use read/write heads that ride just above the magnetic surface, held up by the air flow created by the turning disk. When the disk stops turning, the heads land on the surface, which has a specially lubricated coating. Winchester disks must be sealed and have a filtration system since ordinary dust particles are large enough to catch between the head and the disk.

WRITE ONCE – In the context of optical disks, technologies that allow the drive to store data on a disk and read it back, but not to erase it.

INDEX

A

abbreviations 1-1
actuator lock 5-4
adapter board 2-3, 3-13, 6-1
Adaptive Caching 5-10
Adaptive segmentation 5-10
air filtration 5-5
AIRLOCK® 5-4
alternate status register 6-23
ARR (automatic read reallocation) 6-67
AWRE (automatic write reallocation enabled) 6-67

B

base casting 5-3
BIOS 6-1
bits, status register 6-27
block diagram 5-5
Buffer Controller 5-8
busy bit 6-27
Byte Interleaving 5-14

C

Cable Select 3-6
cable select (CS) jumper 3-5
CE (cache enable) 6-67
check bytes 5-14
Check Power Mode (E5H) 6-52
clean room 5-1
clearance 3-11
command
  identify drive parameters 6-54
command block registers 6-24
command codes and parameters 6-28
command descriptions 6-29
Command Register 6-28
command register 6-28
Configuration 6-64
configuration command data field 6-66
connector, IDE 3-12
control block registers 6-23
cooling fan requirements 3-11
correctable double burst errors 5-15, 5-16
crash stops 5-4
cross checking (XC) 5-13
cylinder high register 6-25
cylinder low register 6-25

D

daisy-chain 2-3
daisy-chained 3-5
data port register 6-24
data request bit 6-27
Data Synchronizer 5-9
Data transfer operations 5-8
DATA TRANSFER RATE 4-3
Data transfer Rates 4-2
DC motor assembly 5-3
DCR (disable early correction) 6-68
device control register 6-23
diagnostics 6-31
DisCache parameters 6-67
DISK ERRORS 4-2
disk stack assembly 5-3
DOS 4-3
double burst errors 5-15
DP8491 Read Channel Device 5-8
drive address register 6-24
drive electronics 5-5
drive mechanism 5-1
drive parameters 6-68
drive ready bit 6-27
drive select (DS) jumper 3-6
drive write fault bit 6-27
drive/head register 6-26

E

ECC 5-14
ECC Check Bytes 5-14
ECC correction span 6-68
ECC interleaving 5-14
EEC (enable early correction) 6-68
ENDEC 5-7
error bit 6-27
error detection, correction 5-13
error rates 5-13
error recovery parameters 6-67
error register 6-24
error reporting 6-71
errors, single burst 5-15
execute drive diagnostic 6-31

F

faceplate 3-1
firmware features 5-10
flex circuit. 5-4
floppy drive 3-13

H

hardware options 3-4
headstack assembly 5-4
host software reset 6-23
Humidity 4-8

I

IDE 2-3, 3-8
IDE interface controller 5-8
IDE-bus interface 6-1
IDE-bus interface connector 3-12
identify drive 6-53
identify drive parameters 6-54
Idle 4-6
Idle Immediate Mode (E1H) 6-52
Idle Mode (E3H) 6-52, 6-53
initialize drive parameters 6-31
input power connections 3-12
interface, IDE-bus 6-1

J

jumper configurations 3-5
jumper locations 3-4
jumper options 3-5

L

landing zone 5-4
logical cylinders 3-18
logical heads 3-18
logical sectors/track 3-18
low 4-3
low-level format 4-3

M

maximum screw torque 3-10
mechanical dimensions 3-1
microcontroller 5-7
motherboard 3-13
Motor Controller 5-8
mounting 3-9
mounting dimensions 3-9
mounting holes 3-9
mounting screw clearance 3-10
mounting screws 3-10
MS-DOS 3-18
MTBF 4-10
multiple random errors 5-17
multiple sector reads 6-29
multiple-sector writes 6-30

N

noise 4-5
Nominal conditions 4-4
number of retries 6-68

P

packing assembly 3-2
packing materials 3-2
PCB power 6-51
PE (prefetch enable) 6-67
Peak Detect 5-9
PM 4-10

power and AT bus connector 3-11
power and bus interface cables 3-14
power connector 3-11
Power Requirements 4-6
Power Sequencing 4-5
pre-amplifier 5-10
pre-compensator 5-9
precompensator circuit 5-7
printed circuit board 3-4
programmed I/O 6-21
Pulse Detector 5-9
PWM 5-7

Q

Quantum configuration key 6-67

R

R/W head matrix 5-10
RC (read continuous) 6-68
read buffer 6-49
read configuration 6-64
read defect list 6-61
read multiple 6-47
read preamplifier 5-4
read sectors 6-29
read verify sectors 6-30
read/write ASIC 5-10
Recalibrate 6-29
register
    error 6-24
reset 6-23
Reset Limits 4-5
ripple 4-5
rotary positioner 5-4
rotary positioner assembly 5-4
RUEE (reallocate uncorrectable error enables) 6-68

S

Sector Count Register 6-52
sector count register 6-25
sector data field 5-14
sector number register 6-25
seek 6-30
Servo Controller 5-8
set configuration 6-65
set configuration without saving to disk 6-65
set features 6-60
set multiple mode 6-48, 6-49
shipping container 3-2
SHOCK 4-9
Slave Present 3-6
slave present (SP) jumper 3-6
Sleep Mode (E6H) 6-52
Specifications 4-9
sputtered thin-film coating 5-3
Standby Immediate Mode (E0H) 6-51
status register 6-27
    busy bit 6-27
    data request bit 6-27
    drive ready bit 6-27
    drive write fault bit 6-27
    error bit 6-27
status register bits 6-27
supply voltages 4-4
syndrome values 5-15

T

tampering with the HDA 5-1
temperature 4-4
theory of operation 5-1
timebase generator 5-9
timing diagram 6-9, 6-11, 6-21
transient voltages 4-5
TTL logic 6-2

U

uncorrectable double burst errors 5-15, 5-16
UNIX 4-3

V

ventilation 3-11
ventilation requirements 3-9
vibration 4-4

W

warranty (void) 5-1
WCE (write cache enable) 6-68
write driver 5-4, 5-10
write multiple 6-48
write sector 6-30