

Quantum[®]

**ProDrive LPS 210AT
Product Manual**

Preliminary

March 1994
81-106898-01

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

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

Publication Number: 81-106898-01

UL/CSA/VDE/TUV

UL standard 1954 recognition granted under File No. E78016

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

VDE 0805 (IEC 950) certification granted under File No. 11342-3250-2003/A1A

TUV Rheinland EN 60 950

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

SERVICE CENTERS

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

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

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

PATENTS

Product covered or licensed under the following U.S. Patents: 4,538,193; 4,669,004; 4,675,652; 4,730,321; 4,819,153; 4,882,671; 4,920,442. Other patents are pending.

Copyright © 1994 by Quantum Corporation. All rights reserved. Printed in U.S.A.

ProDrive®, ProDrive Series®, AIRLOCK®, WriteCache®, and DisCache® are registered trademarks of Quantum Corporation.

PC-AT is a registered trademark of International Business Machines.

MS-DOS is a registered trademark of Microsoft Corporation.

Table of Contents

Chapter 1 ABOUT THIS MANUAL

1.1	AUDIENCE DEFINITION	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	STANDARDS AND REGULATIONS	2-2

Chapter 3 INSTALLATION

3.1	SPACE REQUIREMENTS	3-1
3.2	UNPACKING INSTRUCTIONS	3-2
3.3	JUMPER SETTINGS	3-3
3.3.1	Drive Select (DS) Jumper	3-4
3.3.2	Cable Select (CS) Jumper	3-5
3.3.3	Slave Present (SP) Jumper	3-6
3.4	IDE-BUS ADAPTER	3-6
3.4.1	40-Pin IDE-Bus Connector	3-6
3.5	MOUNTING	3-6
3.5.1	Orientation	3-7
3.5.2	Clearance	3-8
3.5.3	Ventilation	3-9
3.6	COMBINATION CONNECTOR (J11)	3-10
3.6.1	DC Power (J11 A and B)	3-11
3.6.2	LED Connector (J2)	3-11
3.6.3	IDE-Bus Interface Connector (J11)	3-12
3.7	DRIVE INSTALLATION	3-12
3.7.1	For Systems With A Motherboard IDE Adapter	3-12
3.7.2	For Systems With An Ide Adapter Board	3-12
3.8	SYSTEM STARTUP AND OPERATION	3-14

Chapter 4 SPECIFICATIONS

4.1	SPECIFICATION SUMMARY	4-1
4.2	FORMATTED CAPACITY	4-2
4.3	DATA TRANSFER RATES	4-2
4.4	TIMING SPECIFICATIONS	4-3
4.5	POWER MODE TIMES	4-4
4.6	POWER	4-4
4.6.1	Power Sequencing	4-4
4.6.2	Power Reset Limits	4-4
4.6.3	Power Requirements	4-5

4.7	ACOUSTICS	4-6
4.8	MECHANICAL DIMENSIONS	4-6
4.9	ENVIRONMENTAL CONDITIONS	4-7
4.10	SHOCK AND VIBRATION	4-7
4.11	RELIABILITY	4-8
4.12	DISK ERRORS	4-8
 Chapter 5 BASIC PRINCIPLES OF OPERATION		
5.1	PRODRIVE LPS DRIVE MECHANISM	5-1
5.1.1	Head/Disk Assembly	5-1
5.1.2	Base Casting Assembly	5-1
5.1.3	DC Motor Assembly	5-1
5.1.4	Disk Stack Assemblies	5-3
5.1.5	Headstack Assembly	5-3
5.1.6	Rotary Positioner Assembly	5-3
5.1.7	Automatic Actuator Lock	5-3
5.1.8	Air Filtration	5-4
5.2	DRIVE ELECTRONICS	5-5
5.2.1	µController	5-5
5.2.2	DCIA	5-6
5.2.3	Read/Write ASIC	5-9
5.2.4	Preamplifier and Write Driver IC	5-11
5.3	SERVO SYSTEM	5-11
5.3.1	General Description	5-11
5.3.2	Servo Burst and Track Information	5-12
5.3.3	Position and Velocity	5-12
5.4	READ AND WRITE OPERATIONS	5-13
5.4.1	The Read Channel	5-13
5.4.2	The Write Channel	5-13
5.4.3	Interface Control	5-13
5.5	FIRMWARE FEATURES	5-14
5.5.1	Disk Caching	5-14
5.5.2	Track and Cylinder Skewing	5-16
5.5.3	Error Detection and Correction	5-17
5.5.4	Defect Management	5-23
 Chapter 6 THE IDE-BUS INTERFACE AND ATA COMMANDS		
6.1	SOFTWARE INTERFACE	6-1
6.2	MECHANICAL DESCRIPTION	6-1
6.2.1	Drive Cable and Connector	6-1
6.3	ELECTRICAL INTERFACE	6-2
6.3.1	IDE-Bus Interface	6-2
6.3.2	Host Interface Timing	6-8
6.4	REGISTER-ADDRESS DECODING	6-11
6.5	REGISTER DESCRIPTIONS	6-13
6.5.1	Control Block Registers	6-13
6.5.2	Command Block Registers	6-15
6.6	COMMAND DESCRIPTIONS	6-20
6.6.1	Recalibrate	6-20
6.6.2	Read Sectors	6-20

6.6.3	Write Sector	6-21
6.6.4	Read Verify Sectors	6-21
6.6.5	Format Track	6-21
6.6.6	Seek	6-21
6.6.7	Execute Drive Diagnostic	6-22
6.6.8	Initialize Drive Parameters	6-23
6.6.9	Read Multiple	6-23
6.6.10	Write Multiple	6-24
6.6.11	Set Multiple Mode	6-24
6.6.12	Read Buffer	6-25
6.6.13	Write Buffer	6-25
6.6.14	Power Management Commands	6-25
6.6.15	Identify Drive	6-27
6.6.16	Set Features	6-30
6.6.17	Read Defect List	6-30
6.6.18	Defect List Data Structure	6-31
6.6.19	Configuration	6-32
6.7	ERROR REPORTING	6-37

GLOSSARY	G-1
----------------	-----

INDEX	I-1
-------------	-----

List of Figures

FIGURE	DESCRIPTION	PAGE
3-1	Mechanical Dimensions	3-1
3-2	ProDrive LPS 210AT Drive PAcKING Assembly	3-2
3-3	Jumper Locations on the Drive PCBA	3-3
3-4	Drive Select Jumper Circuit	3-5
3-5	Cable Select Jumper Circuit	3-6
3-6	Mounting Dimensions (in Millimeters)	3-7
3-7	Mounting Screw Clearance	3-9
3-8	J11 DC Power and IDE-Bus Combination Connector	3-10
3-9	Drive Power Supply and IDE-Bus Interface Cables	3-13
3-10	Drive Installation	3-14
5-1	Exploded View	5-2
5-2	HDA Air Filtration	5-5
5-3	ProDrive LPS 210AT Block Diagram	5-6
5-4	DCIIA Block Diagram	5-7
5-5	Read/Write ASIC Block Diagram	5-10
5-6	Sector Data Field with ECC Check Bytes	5-18
5-7	Byte Interleaving	5-18
5-8	Single Burst Error Correctability	5-19
5-9	Correctable and Uncorrectable Double-Burst Errors	5-20
5-10	Six Correctable Random Burst Errors	5-21
6-1	PIO Interface Timing	6-9
6-2	Multiword DMA Bus Interface Timing	6-10
6-3	Host Interface RESET Timing	6-11
6-4	Defect List Data Structure	6-31

List of Tables

TABLE	DESCRIPTION	PAGE
3-1	ProDrive LPS 210AT Jumper Options	3-4
3-2	J11 Power (Sections A and B)	3-11
3-3	JJ2 LED Connector	3-12
3-4	Logical Addressing Format	3-15
4-1	ProDrive LPS 210AT Specifications	4-1
4-2	Formatted Capacity	4-2
4-3	Timing Specifications	4-3
4-4	Power Mode Times	4-4
4-5	Power Reset Limits	4-4
4-6	Typical Power and Current Consumption	4-5
4-7	Acoustical Characteristics – Sound Pressure	4-6
4-8	Acoustical Characteristics—Sound Power	4-6
4-9	Mechanical Dimensions	4-6
4-10	Environmental Specifications	4-7
4-11	Shock and Vibration Specifications	4-7
4-11	Error Rates	4-8
5-1	Surface Layout	5-4
5-2	Track and Cylinder Skewing	5-17
6-1	Drive-Connector Pin Assignments (J11-C)	6-3
6-2	Relationship of Drive Signals to the IDE Bus	6-7
6-3	PIO Host Interface Timing	6-8
6-4	Multiword DMA Host Interface Timing	6-10
6-5	Host Interface RESET Timing	6-11
6-6	I/O-Port Functions and Selection Addresses	6-12
6-7	Command Block Register Initial Values	6-13
6-8	Device Control Register Bits	6-13
6-9	Drive Address Register Bits	6-14
6-10	Error Register Bits	6-15
6-11	Drive/Head Register Bits	6-17
6-12	Status Register Bits	6-18
6-13	ATA Command Codes and Parameters	6-19
6-14	Diagnostic Codes	6-23
6-15	Identify Drive Parameters	6-28
6-16	READ DEFECT LIST Command Bytes	6-30
6-17	Sample Defect List	6-32
6-18	Accessing the READ CONFIGURATION Command	6-33
6-19	Accessing the SET CONFIGURATION Command	6-34
6-20	Set Configuration (Without Saving) Access Pattern	6-34
6-21	Configuration Command Format	6-35
6-22	Command Errors	6-38

Chapter 1

ABOUT THIS MANUAL

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

1.1 AUDIENCE DEFINITION

The ProDrive LPS 210AT product manual is intended for several audiences including the original equipment manufacturer (OEM), distributor, installer, and end user. The manual provides you with information about installation, principles of operation, interface command implementation, and maintenance.

1.2 MANUAL ORGANIZATION

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

- Chapter 1—About This Manual
- Chapter 2—General Description
- Chapter 3—Installation
- Chapter 4—Specifications
- Chapter 5—Basic Principles of Operation
- Chapter 6—IDE-Bus Interface and ATA Commands

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

1.3 TERMINOLOGY AND CONVENTIONS

The abbreviations listed below are used in this manual. You can find definitions of the following terms in the Glossary at the back of this manual.

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

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

The following conventions are used in this manual:

- **Commands and Messages**

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

WRITE LONG
ABORTED COMMAND

- **Parameters**

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

Prefetch Enable, PE
Cache Enable, CE

- **Names of Bits and Registers**

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

Host Software Reset
Alternate Status Register

- **Hexadecimal Notation**

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

30_H

- **Register Names:**

Registers are given in this manual with initial capitals. An example is the Alternate Status Register.

- **Signal Negation**

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

RD-

- **Notes**

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

- **Host**

In general, the system in which the drive resides is referred to as the host. The AT/IDE host adapter is considered to be part of 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 4.0, March 17, 1993
- Small Form Factor Specification for ATA Timing Extensions #8011, July 2, 1993

Chapter 2

GENERAL DESCRIPTION

This chapter summarizes general functions and key features of the ProDrive LPS™ 210AT hard disk drive, as well as the applicable standards and regulations it meets.

2.1 PRODUCT OVERVIEW

Quantum's ProDrive® LPS 210AT hard disk drives are part of a family of high-performance, 1-inch-high, hard disk drives manufactured to meet the highest product quality standards. ProDrive LPS hard disk drives use nonremovable, 3 1/2-inch hard disks. These drives can be used in IBM® PC/AT-compatible host computer systems, which provide an AT Attachment (ATA) Interface either on an adapter board or on the system motherboard.

Note: Sometimes, the ATA Interface is referred to as the *AT Bus* or *IDE Interface*.

ProDrive LPS 210AT hard disk drives feature an embedded IDE drive controller and use ATA commands. The drive manages media defects and error recovery internally, so these operations are transparent to the user.

The innovative design of ProDrive LPS 210AT hard disk drives enables Quantum to produce a family of low-cost, high-reliability drives.

2.2 KEY FEATURES

ProDrive LPS 210AT hard disk drives include the following key features:

General

- Formatted storage capacity of 210 MB (1 disk)
- Industry-standard, 3 1/2-inch form factor
- Low-profile, 1-inch height
- Emulation of IBM® PC/AT® task file register and all ATA fixed-disk commands
- Embedded servo design

Performance

- Data transfer rate of up to 6.0 MB/s, using programmed I/O (PIO), 11.1 MB/s using PIO with IORDY, 13.3 MB/s using multiword DMA
- Average seek time of 14 ms
- Average rotational latency of 8.33 ms
- 1:1 interleave on read/write operations

- 128K buffer with 96K dynamic segmentation cache. Look-ahead DisCache[®] feature with continuous prefetch, and WriteCache[™] write-buffering capabilities
- Support for all ATA data transfer modes including local bus timing
- Auto Write, and Auto Task File Update features in Quantum's customized ASIC

Reliability

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

Versatility

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

2.3 STANDARDS AND REGULATIONS

ProDrive LPS 210AT hard disk drives satisfy the following standards and regulations:

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

Chapter 3 INSTALLATION

This chapter explains how to unpack, configure, mount, and connect the ProDrive LPS 210AT hard disk drive prior to operation. It also explains how to start up and operate the drive.

3.1 SPACE REQUIREMENTS

Quantum ships the ProDrive LPS 210AT hard disk drive without a faceplate. Figure 3-1 shows the external dimensions of the drive.




Figure 3-1 *ProDrive LPS 210AT Mechanical Dimensions*

3.2 UNPACKING INSTRUCTIONS

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

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

Figure 3-2 shows the packaging for the ProDrive LPS 210AT hard disk drive in a 1-pack shipping container. (A 12-pack shipping container is available for multiple drive shipments and has the same type of protective packaging.)

3. Remove the drive from the box.

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

4. Save the packaging materials for possible future use.




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

3.3 JUMPER SETTINGS

The configuration of a ProDrive LPS 210AT hard disk drive depends on the host system into which it is to be installed. This section describes the jumper setting options you must take into account prior to installation. Figure 3-3 shows the printed circuit board assembly (PCBA), including the location of the jumpers that control some of these options.




Figure 3-3 *Jumper Locations on the Drive PCBA*

The configuration of three jumpers controls the drive's mode of operation:

- CS—Cable Select
- DS—Drive Select
- SP—Slave Present

Table 3-1 defines the jumper configurations for the drive and the operational modes that result from the configurations. *I* indicates that the specified jumper is installed; *O*, that the jumper is not installed.

Table 3-1 *ProDrive LPS 210AT Jumper Options*

CS	DS	SP	DESCRIPTION
0	0	0	Slave Drive Compatible with drives using the PDIAG– line to handle Master/Slave communications.
0	0	1	Slave Drive The PDIAG– and DASP– lines are not driven.
0*	1*	0*	Master Drive Uses DASP– to check for the presence of a Slave.
0	1	1	Master Drive Uses the SP jumper to determine whether a Slave is present, without checking PDIAG– or DASP–.
1	0	X	Slave or Master Drive, depending on the state of the Cable Select signal (pin 28) at the IDE-bus interface connector. If the signal state is set to 0 (grounded), then the drive is configured as if DS were 1, described above. If the Cable Select signal is set to 1 (high), then the drive is configured as if DS were 0, described above.

Note: In Table 3-1, a 0 indicates that the jumper is removed, a 1 indicates that the jumper is installed, and X indicates the jumper setting does not matter. An asterick (*) indicates the factory default setting.

CAUTION: The CS and DS jumpers should never be installed at the same time.

3.3.1 Drive Select (DS) Jumper

You can configure two drives on the ATA Interface as Master (Drive 0) and Slave (Drive 1) using the DS jumpers. Set the CS jumpers to 0 for each drive. Then, set the DS jumper to 1 on one drive to configure that drive as Master. Set the DS jumper to 0 on the other drive to configure that drive as Slave.

Quantum ships ProDrive LPS 210AT hard disk drives from the factory with the DS jumper installed—that is, the DS jumper is set to 1 to configure the drive as Master. To configure a drive as the Slave, set the DS jumper to 0.

Note: For drives configured using the DS jumper, the order in which drives are connected in a daisy chain has no significance.

Figure 3-4 provides information on how the DS jumper is connected on the circuit board of the Master drive when the DS jumper is used to select drive ID.

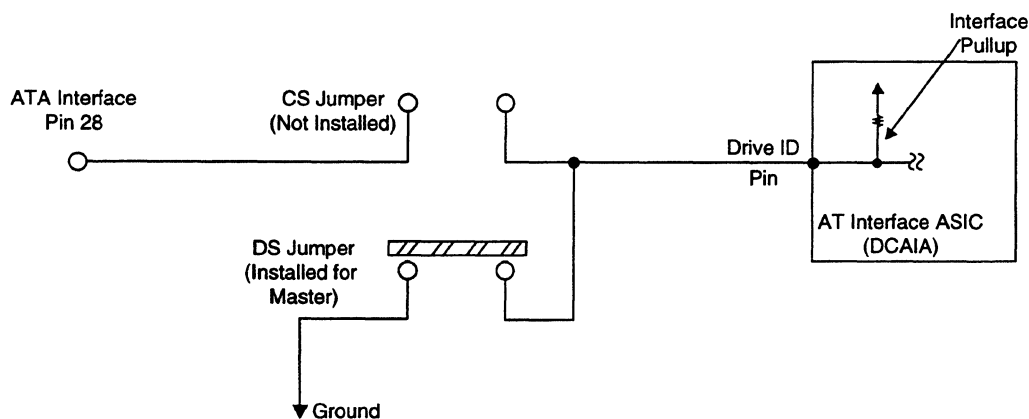


Figure 3-4 *Drive Select Jumper Circuit*

3.3.2 Cable Select (CS) Jumper

When two ProDrive LPS 210AT hard disk drives are daisy-chained together, they can be configured as Master and Slave by using the CS jumper or by using the DS jumper—but not both. To configure the two drives using the CS jumper, set the CS jumper to 1 (installed) and the DS jumper to 0 (uninstalled) on each drive.

Once you install the CS jumper, the drive is configured as a Master or Slave by the state of the Cable Select signal, which is pin 28 at the IDE-bus interface connector. According to the ATA (AT Attachment) specification referred to in Section 1.4, pin 28 may be used for either Cable Select or spindle synchronization. Quantum uses the pin 28 Cable Select function but does not implement the spindle synchronization function.

Pin 28 is grounded—that is, set to 0—on the cable coming from the host. This configures the first drive as Master. Then, pin 28 on the connector at the second drive should be made an open circuit by a cut in signal line 28 in the cable from the first drive, so it becomes high—that is, set to 1—due to a pull-up in the second drive. This configures the second drive as Slave.

Figure 3-5 provides information on how the CS jumper is connected on the circuit board of either the Master drive or the Slave drive when the interface cable is used to select drive ID.




Figure 3-5 *Cable Select Jumper Circuit*

3.3.3 Slave Present (SP) Jumper

The SP jumper normally is not needed. However, when the ProDrive LPS 210AT is configured as Master and is connected to a Slave drive that does not implement the Drive Active/Slave Present (DASP-) signal, it is necessary to set the SP jumper to 1 on the ProDrive LPS 210AT.

3.4 IDE-BUS ADAPTER

There are two ways you can configure a system to allow the ProDrive LPS 210AT hard disk drive to communicate over the IDE-bus of an IBM or IBM-compatible PC:

1. Connect the drive to a 40-pin IDE-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 IDE-Bus Connector

Many of the later design PC motherboards have a built-in 40-pin, IDE-bus connector that is compatible with the 40-pin IDE interface of the ProDrive LPS 210AT hard disk drive. If the motherboard has an IDE 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 IDE-bus interface connector, you must install an IDE-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 read over carefully the instruction manual for the adapter board you purchase 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

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

3.5.1 Orientation

The mounting holes on the ProDrive LPS 210AT hard disk drive allow the drive to be mounted in any orientation. Figure 3-6 shows the location of the three mounting holes on each side of the drive.




Figure 3-6 ProDrive 210AT Mounting Dimensions (in Millimeters)

3.5.2 Clearance

The printed-circuit board assembly (PCBA) is very close to the mounting holes. Clearance from the drive to any other surface—except mounting surfaces—must be 1.25 mm (0.05 inches) minimum. Figure 3-7 specifies the clearance between the screws in the mounting holes and the PCBA. Do not use mounting screws longer than the maximum lengths specified in Figure 3-7. The specified screw length allows full use of the mounting-hole threads, while avoiding damaging or placing unwanted stress on the PCBA.




Figure 3-7 *Mounting Screw Clearance*

CAUTION: The printed-circuit board (PCB) is very close to the mounting holes. Do not exceed the specified length for the mounting screws. The specified screw length allows full use of the mounting-hole threads, while avoiding damaging or placing unwanted stress on the PCB. Figure 3-7 specifies the minimum clearance between the PCB and the screws in the mounting holes. To avoid stripping the mounting-hole threads, the maximum torque applied to the screws must not exceed 8 inch-pounds.

3.5.3 Ventilation

The ProDrive LPS 210AT hard disk drive operates without a cooling fan, provided the ambient air temperature does not exceed 122°F (50°C).

3.6 COMBINATION CONNECTOR (J11)

As shown in Figure 3-8, J11 is a three-section combination connector mounted on the back edge of the printed-circuit board assembly (PCBA). The drive's DC power can be applied to either section A or B. The IDE-bus interface (40-pin) uses section C.




Figure 3-8 J11 DC Power and IDE-Bus Combination Connector

3.6.1 DC Power (J11 A and B)

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

Table 3-2 J11 Power (Sections A and B)

J11 SECTION	PIN NUMBER	VOLTAGE LEVEL	MATING CONNECTOR TYPE AND PART NUMBER (OR EQUIVALENT)
A	1	+12 Vdc	4-Pin Connector: AMP P/N 1-480424-0 Loose piece contacts: AMP P/N 61173-4 Strip contacts: AMP P/N 350078-4
A	2	+12 Vdc Return (Ground)	
A	3	+5 Vdc Return (Ground)	
A	4	+5 Vdc	
B	1	+5 Vdc	3-Pin Connector Molex P/N 39-01-0033 Loose piece contacts: Molex P/N 39-00-00341 Strip contacts: Molex P/N 39-00-0023
B	2	+12 Vdc	
B	3	Ground	

Note: Power can be applied to either J11 Section A or J11 section B. Labels indicate the pin numbers on the connector. On Section A, Pins 2 and 3 (the +5 and +12 Vdc returns) are connected on the drive.

3.6.2 LED Connector (J12)

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

Table 3-3 *J12 LED Connector*

J12 PINS		MATING CONNECTOR TYPE AND PART NUMBER (OR EQUIVALENT)
Pin Number	Function	
1	-LED	Two-position housing: Molex P/N 51021-0200 Loose-piece contacts: Molex P/N 50058-8100
2	+LED	

3.6.3 IDE-Bus Interface Connector (J11C)

On the ProDrive LPS 210AT hard disk drive, the IDE-bus interface cable connector (J11 section C) is a 40-pin Universal Header, as shown in Figure 3-8.

To prevent the possibility of incorrect installation, make sure the connector is keyed by removing Pin 20. Removing this pin ensures the connector cannot be installed upside down.

See Chapter 6, "IDE-Bus Interface and ATA Commands," for more detailed information about the required signals. Refer to Table 6-1 for the pin assignments of the IDE-bus connector (J11 section C).

For mating with J11 section C, recommended cable connectors include the following parts or their equivalents:

AMP receptacle with strain relief	P/N 1-499506-0
AMP receptacle without strain relief	P/N 1-746193-0

To key the cable connectors, you must plug the hole that corresponds to pin 20.

Other recommended part numbers for the mating connector include:

40-Pin Connector	3M 3417-7000 or equivalent
Strain Relief	3M 3448-2040 or equivalent
Flat Cable (Stranded 28 AWG)	3M 3365-40 or equivalent
Flat Cable (Stranded 28 AWG)	3M 3517-40 (shielded) or equivalent

3.7 DRIVE INSTALLATION

You can install the ProDrive LPS 210AT hard disk drive in an AT-compatible system in two ways:

3.7.1 For Systems With A Motherboard IDE Adapter

You can install the ProDrive LPS 210AT hard disk drive in an AT-compatible system that contains a 40-pin IDE-bus connector on the motherboard.

To connect the drive to the motherboard, use a 40-pin ribbon cable. Ensure that pin 1 of the drive is connected to pin 1 of the motherboard connector.

3.7.2 For Systems With An IDE Adapter Board

To install a ProDrive LPS 210AT hard disk drive in an AT-compatible system without a 40-pin IDE-bus connector on its motherboard, you need a third-party IDE-compatible adapter board.

3.7.2.1 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 addressing 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 the 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.

Connecting the Adapter Board and the Drive

Use a 40-pin ribbon cable to connect the drive to the board. See Figure 3-9. 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 3-10.




Figure 3-9 *Drive Power Supply and IDE-Bus Interface Cables*




Figure 3-10 Drive Installation

3.8 SYSTEM STARTUP AND OPERATION

Once you have installed the ProDrive LPS 210AT 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.

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.

3. Enter the appropriate parameters.

During the AT system CMOS setup, you must enter the drive type for the ProDrive LPS 210AT hard disk drive. This procedure allows the system to recognize the drive by translating its *physical* drive geometry parameters such as cylinders, heads, and sectors per track, into a *logical* addressing mode.

ProDrive LPS 210AT drives can work with different BIOS drive-type tables of various host systems. You can choose any drive type that does not exceed the capacity of the drive. Table 3-4 gives the logical parameters that provide the maximum capacity on ProDrive LPS 210AT hard disk drives..

Table 3-4 *Logical Addressing Format*

	ProDrive LPS 210AT
Logical Cylinders	723
Logical Heads	15
Logical Sectors/Track	38
Total Number Logical Sectors	412,110

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 ProDrive LPS 210AT

$$\text{Logical Cylinders} \times \text{Logical Heads} \times \text{Logical Sectors/Track} \times 512 \leq 211,000,320$$

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

If you are using a version of MS-DOS below 4.01:

1. Run FDISK or a third-party partitioning program.

Note: If you use DOS version 3.2 or an earlier version, the DOS partition will employ only 32 MB (33,554,432 bytes) of the drive's capacity. With DOS 3.3, you can partition the drive in multiples of 32 MB. If you use DOS 4.01 or later, or if you use a third-party partitioning program, you can create partitions that exceed 32 MB.

2. To format the drive with a high-level format and transfer the operating system to the drive, type `FORMAT C: /S`. Once this command executes, you can boot the system from the hard drive.

You do not need to perform a low-level format on the drive because it was done at the factory before shipment.

Chapter 4 SPECIFICATIONS

This chapter gives a detailed description of the physical, electrical, and environmental characteristics of the ProDrive LPS 210AT hard disk drive.

4.1 SPECIFICATION SUMMARY

Table 4-1 gives a summary of the ProDrive LPS 210AT hard disk drive.

Table 4-1 *ProDrive LPS 210AT Specifications*

DESCRIPTION	PRODRIVE LPS 210AT
Capacity (formatted)	211.0 MB
Nominal rotational speed (rpm)	3,600
Number of disks	1
Number of R/W heads	2
Data Organization	
Zones per surface	16
Tracks per surface	2,519
Total tracks	5,038
Sectors per track	
Inside zone	55
Outside zone	104
Total User Sectors	412,110
Bytes per sector	512
Number of tracks per cylinder	2
Recording	
Recording technology	Multiple Zone
Maximum linear density	41,384 fci
Maximum Recording Density	55,179 (bpi)
Encoding method	RLL 1,7
Interleave	1:1
Track density	2,670 tpi
Maximum effective areal density	147 Mbit/in ²

Table 4-1 *ProDrive LPS 210AT Specifications (Continued)*

DESCRIPTION		PRODRIVE LPS 210AT
Performance		
Typical Seek times ¹		
Average seek		15 ms
Track-to-track		5.0 ms
Full stroke		31 ms
Sequential Head Switch		4.5 ms
Rotational Latency		8.33 ms
Data Transfer rate		
Disk to Read Buffer ²		17.67 to 31.56 Mb/s ³
Read/Write Buffer to IDE-bus (PIO Mode)	Without IORDY	6.00 MB/s max.
	With IORDY	11.1 MB/s max.
Read/Write Buffer to IDE-bus (DMA Mode)		13.3 MB/s max.
Buffer size (The upper 32K is used for firmware)		128 KB
Projected MTBF		300,000 hrs
Contact Start/Stop Cycles		20,000
Auto head-park method		AirLock

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

4.2 FORMATTED CAPACITY

At the factory, the ProDrive LPS 210AT receives a low-level format that creates the actual tracks and sectors on the drive. Table 4-2 shows the capacity resulting from this process. For operation with DOS, UNIX, or other operating systems, formatting done at the user level results in less capacity than the physical capacity shown in Table 4-2.

Table 4-2 *Formatted Capacity*

	PRODRIVE LPS 210 CAPACITY
Number of bytes	211,000,320
Number of 512-byte sectors	412,110

4.3 DATA TRANSFER RATES

Data is transferred from the disk to the read buffer at a rate up to 3.945 MB/sec in bursts. Data is transferred from the read buffer to the AT bus at a rate of up to 6.0 MB/sec, using programmed I/O. If IORDY is used, then this transfer rate can be increased to 11.1 MB/sec. Refer to paragraph 6.6.15 IDENTIFY DRIVE command for additional information. Using the Multiword DMA protocol, transfer rates of 13.3 MB/sec. are achievable (burst mode only).

4.4 TIMING SPECIFICATIONS

Table 4-3 illustrates the timing specifications of the ProDrive LPS 210AT hard disk drive.

Table 4-3 *Timing Specifications*

Parameter	Typical Nominal ¹	Maximum Worst Case ²
Single Track Seek ³	5.0 ms	6.5 ms
Sequential Cylinder Switch Time ⁴	4.5 ms	6.5 ms
Sequential Head Switch Time ⁵	4.5 ms	6 ms
Random Average (Read or Seek) ⁶	15 ms	18 ms
Random Average (Write)	17 ms	18 ms
Full-Stroke Seek	31 ms	33 ms
Average Rotational Latency	8.33 ms	—
Power On ⁷ to Drive Ready ⁸	—	12 seconds
Power On to Interface Ready ⁹	—	12 seconds
Standby ¹⁰ to Drive Ready	—	12 seconds
Drive Ready to Shutdown	5 seconds	20 seconds ¹¹

¹. Nominal conditions are as follows:

- Nominal temperature (25° C)
- Nominal supply voltages (12.0V, 5.0V)
- No applied shock or vibration

². Worst-case conditions are as follows:

- Worst-case temperature extremes (4°C to 50°C)
- Worst-case supply voltages (12V ±10%, 5V±5%)

³. Seek time is defined as the time required for the actuator to seek and settle on-track. It is measured by averaging 5000 seeks of the indicated type as shown in this table. The seek times listed include head settling time, but do not include command overhead time or rotational latency delays.

⁴. Sequential Cylinder Switch Time is the time from the conclusion of the last sector of a cylinder to the first logical sector on the next cylinder.

⁵. Sequential Head Switch Time is the time from the last sector of a track to the beginning of the first logical sector of the next track of the same cylinder.

⁶. Average seek time is the average of 5000 random seeks. When a seek error occurs, recovery for that seek can take up to seven seconds.

⁷. Power On is the time from when the supply voltages reach operating range to when the drive is ready to accept any command.

⁸. 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.

⁹. Interface Ready is the condition in which the drive is ready to accept any command, before the disks are rotating at rated speed.

¹⁰. Standby is the condition in 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.

¹¹. After 20 seconds, it is safe to move the disk drive.

4.5 POWER

The ProDrive LPS 210AT hard disk drive operates from two supply voltages:

- +12V \pm 10%
- +5V \pm 5%

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

4.5.1 Power Sequencing

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

4.5.2 Power Reset Limits

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

Table 4-4 *Power Reset Limits*

DC VOLTAGE	THRESHOLD
+5 V	4.25 V to 4.70 V
+12 V	8.9 V to 10.3 V

4.5.3 Power Requirements

Table 4-5 lists the voltages and corresponding current for the various modes of operation of the ProDrive LPS 210AT hard disk drive.

Table 4-5 *Typical Power and Current Consumption*

MODE OF OPERATION	TYPICAL AVERAGE CURRENT ⁶ (mA rms)		TYPICAL AVERAGE POWER ⁷ (WATTS)
	+12 V	+5 V	
Startup ¹	1000 peak	400 peak	–
Idle ²	170	250	3.3 W
Random Read/Write ³	260	260	4.5 W
Random Seek ⁴	450	250	6.7 W
Standby/Sleep ⁵	10	175	1.0 W
Noise and Ripple	100 mv	50 mv	–

Notes: ¹ Current is rms except for startup. Startup current of peaks is greater than 10ms in duration.

² Idle mode is in effect when the drive is not reading, writing, seeking, or executing any commands. A portion of the read/write circuitry is powered down, the motor is up to speed, and the Drive Ready condition exists. The actuator resides on the last track accessed.

³ Random Read/Write mode is defined as when data is being read from or written to the disk. It is computed based on 40% seeking, 30% on-track read, and 30% on-track write.

⁴ Continuous random seek operations with minimum controller delay.

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

Sleep is when the spindle and actuator motors are off with the heads latched in the landing zone. Receipt of a reset causes the drive to transition from the sleep to the standby mode.

⁶Current is rms (except for Startup).

⁷Reflects nominal values for +12V and +5V power supplies.

4.6 ACOUSTICS

Table 4-6 and Table 4-7 specifies the acoustical characteristics of the ProDrive LPS 210AT hard disk drive.

Table 4-6 *Acoustical Characteristics – Sound Pressure*

OPERATING MODE	MEASURED NOISE	DISTANCE
Idle On Track	32 dbA (mean) 37 dbA (maximum)	39.3 in (1 m)
Random Seek ¹	36 dbA (mean) 41 dbA (maximum)	39.3 in (1 m)

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

Table 4-7 *Acoustical Characteristics—Sound Power*

Operating Mode	Measured Noise (Sound Power per ISO 7779)
Idle On Track	4.5 Bels maximum

4.7 MECHANICAL DIMENSIONS

Table 4-8 specifies the mechanical dimensions of the ProDrive LPS 210AT hard disk drive. Dimensions measurements do not include the faceplate.

Table 4-8 *Mechanical Dimensions*

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

Mechanical Weight

Pounds	Grams	Ounces
0.85	385	13.6

4.8 ENVIRONMENTAL CONDITIONS

Table 4-9 summarizes the environmental specifications for the ProDrive LPS 210AT hard disk drive.

Table 4-9 *Environmental Specifications*

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

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

4.9 SHOCK AND VIBRATION

The ProDrive LPS 210AT hard disk drive can withstand levels of shock and vibration applied to any of its three mutually perpendicular axes, or principal base axes, as specified in Table 4-10. A functioning drive can be subjected to specified *operating* levels of shock and vibration. When a drive has been subjected to specified *nonoperating* levels of shock and vibration, with power to the drive off, there will be no change in performance at power on.

When packed in either the 1-pack or 10-pack shipping container, ProDrive LPS 210AT drives can withstand a drop from 30 inches onto a concrete surface on any of its surfaces, six edges, or three corners.

Table 4-10 *Shock and Vibration Specifications*

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

4.10 RELIABILITY

Mean Time Between Failures (MTBF): 300,000 Power On Hours (POH), typical usage

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

Preventive Maintenance (PM): Not required

Mean Time To Repair (MTTR): 30 minutes

Start/Stop: 20,000 cycles (minimum)

Component Life: 5 Years

4.11 DISK ERRORS

Table 4-11 provides the error rates for the ProDrive LPS 210AT hard disk drive.

Table 4-11 *Error Rates*

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

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

Chapter 5

BASIC PRINCIPLES OF OPERATION

This chapter describes the operation of ProDrive LPS 210AT functional subsystems. It is intended as a guide to the operation of the drive, rather than a detailed theory of operation.

5.1 ProDrive LPS DRIVE MECHANISM

This section describes the drive mechanism. The ProDrive LPS 210AT hard disk drive consists of a mechanical assembly and a printed circuit board assembly (PCBA) as shown in Figure 5-1.

5.1.1 Head/Disk Assembly

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

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

The ProDrive LPS products are part of a one or two-disk product family. Both the ProDrive LPS 127AT and the 170AT have one magnetic disk and two read/write heads.

5.1.2 Base Casting Assembly

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

5.1.3 DC Motor Assembly

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




Figure 5-1. ProDrive LPS 210AT Exploded View

5.1.4 Disk Stack Assemblies

The disk stack assembly in the ProDrive LPS 210AT hard disk drive consists of one disk secured on the hub of the DC motor assembly by a disk clamp. The aluminum-alloy disks have a sputtered thin-film magnetic coating.

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

5.1.4.1 Surface Layout

The ProDrive LPS 210AT hard disk drive has 2519 tracks per recording surface. Table 5-1 details the surface layout for each platter (recording surface). The data tracks are divided into sixteen recording zones. The drive uses multiple zone recording, where each data track contains between 55 and 104 sectors (depending on the recording zone).

5.1.5 Headstack Assembly

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

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

5.1.6 Rotary Positioner Assembly

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

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

5.1.7 Automatic Actuator Lock

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

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

Table 5-1. Surface Layout

CYLINDER CONTENTS	ZONE	STARTING CYLINDER	ENDING CYLINDER	NUMBER OF CYLINDERS	SECTORS PER TRACK
Guard band					
System Test Files Diskware Firmware Diagnostics	—	-1	-10	10	89
Data	0	0	392	393	104
	1	393	537	145	104
	2	538	645	108	100
	3	646	762	117	97
	4	763	859	97	94
	5	866	1008	149	91
	6	1009	1072	64	89
	7	1073	1230	158	85
	8	1231	1353	123	82
	9	1314	1620	267	78
	10	1621	1772	152	72
	11	1773	1958	186	68
	12	1959	2107	149	65
	13	2108	2229	122	62
	14	2230	2414	185	58
	15	2415	2518	104	55
Guard band					
Total Data Tracks 5,038					

5.1.8 Air Filtration

The ProDrive LPS 210AT hard disk drives are Winchester-type drives. The heads fly very close to the media surface. Therefore, it is essential that the air circulating within the drive be kept free of particles. Quantum assembles the drive in a Class-100, purified air environment, then seals the drive with a metal cover. When the drive is in use, the rotation of the disks forces the air inside of the drive through an internal filter.

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




Figure 5-2. ProDrive LPS 210AT HDA Air Filtration

5.2 PRODRIVE LPS 210AT DRIVE ELECTRONICS

Advanced circuit design, and the use of miniature surface-mounted devices and proprietary VLSI components, enable the drive electronics, including the IDE bus interface, to reside on a single printed circuit board assembly (PCBA). The components are mounted only on one side of the PCBA. Figure 5-3 contains a simplified block diagram of the ProDrive LPS 210AT electronics.

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

5.2.1 μ Controller

The μ Controller (micro-controller) provides local processor services to the drive electronics, under program control. The μ Controller provides direction for the Disk Controller and IDE Interface ASIC (DCIIA), the Read/Write ASIC, and the Actuator Driver. In addition, it controls the head selection process.

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

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

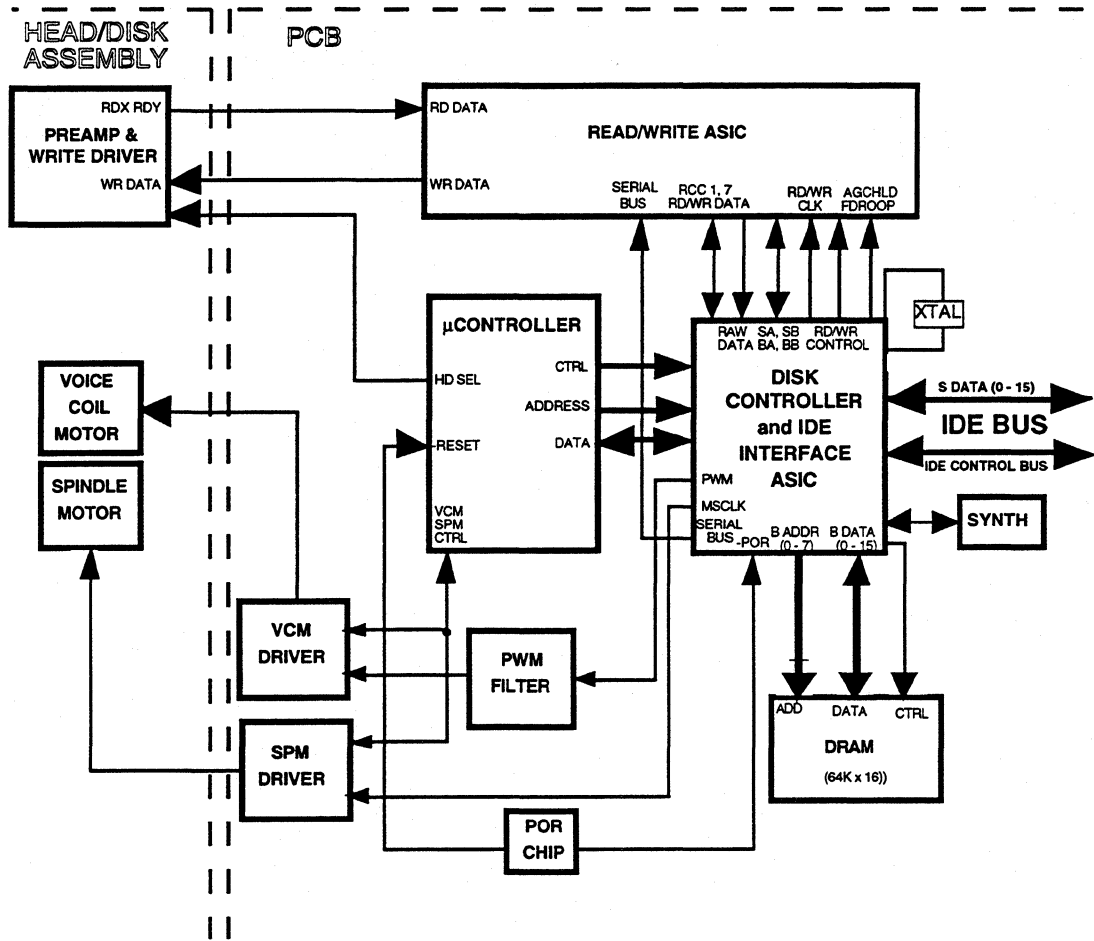


Figure 5-3. ProDrive LPS 210AT Block Diagram

5.2.2 DCIIA

The DCIIA (Disk Controller and IDE Interface ASIC) shown in Figure 5-4 provides control functions to the drive under the direction of the μ Controller.

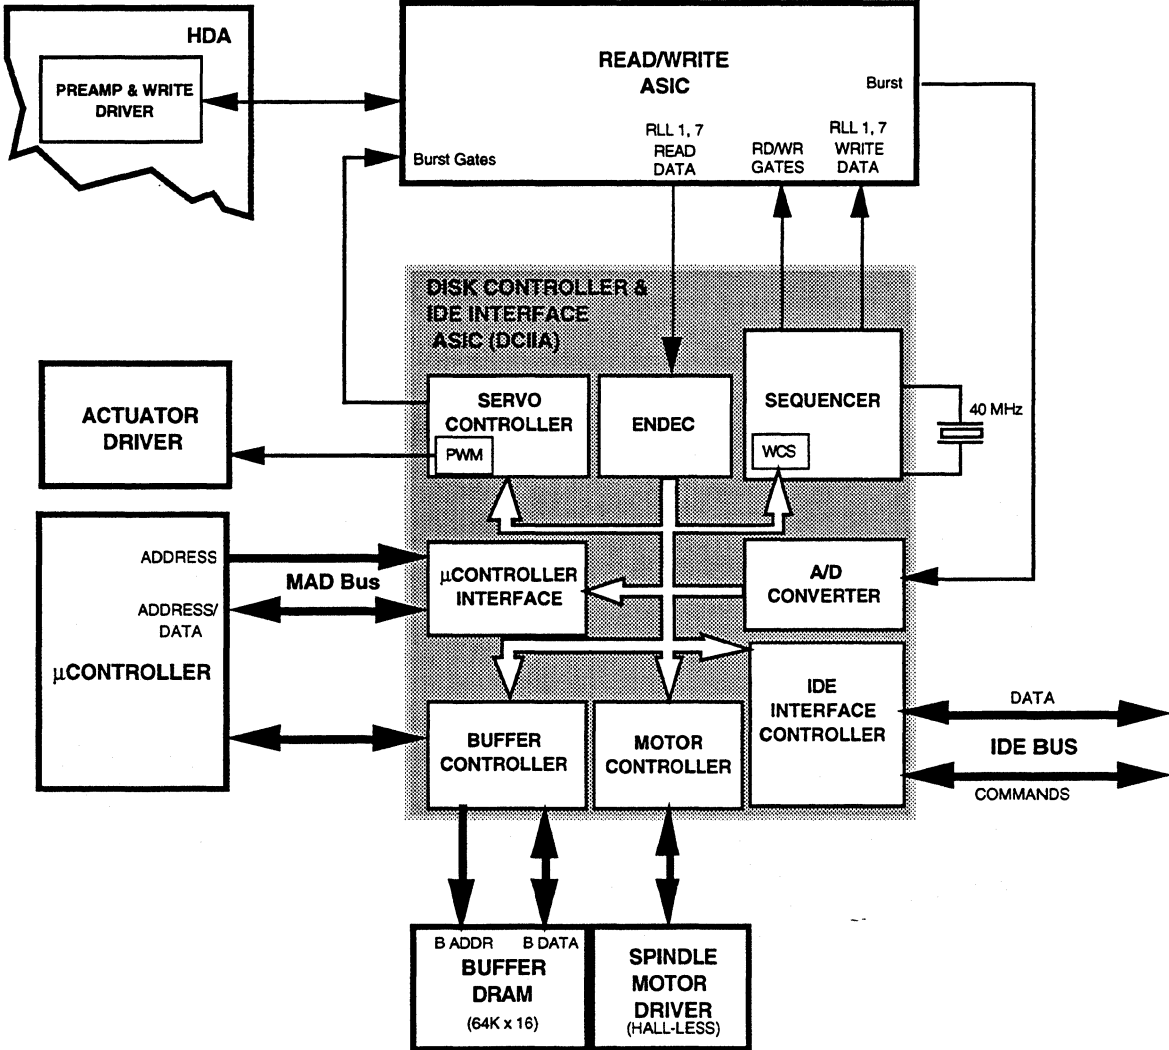


Figure 5-4 DCIIA Block Diagram

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

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

5.2.2.1 A/D Converter

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

5.2.2.2 Encoder/Decoder (ENDEC)

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

5.2.2.3 Sequencer

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

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

5.2.2.4 Buffer Controller

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

The Buffer Controller operates under the direction of the μ Controller.

5.2.2.5 μ Controller Interface

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

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

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

5.2.2.6 Servo Controller

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

5.2.2.7 Motor Controller

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

5.2.2.8 IDE Interface Controller

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

5.2.3 Read/Write ASIC

The Read/Write ASIC shown in Figure 5-5 provides write data precompensation and read channel processing functions for the drive. The Read/Write ASIC receives the RD GATE signal, reference oscillator, serial programming, and servo burst and sample gates from the DCIIA. The Read/Write ASIC sends encoded read data and the read reference clock, and receives write data from the DCIIA.

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

- Time Base Generator
- Pulse Detector
- Peak Detector and Servo Demodulator
- Data Synchronizer

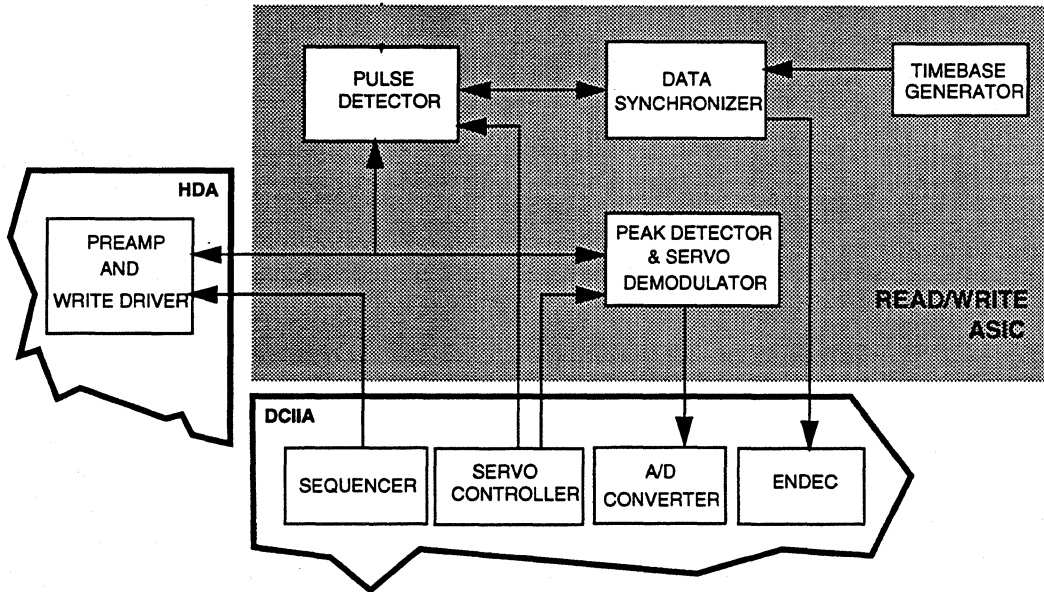


Figure 5-5 Read/Write ASIC Block Diagram

5.2.3.1 Timebase Generator

The frequency synthesizer in the timebase generator provides the reference frequency for the Data Synchronizer and the DCIIA. The μ Controller programs the Timebase Generator with the appropriate frequency for each zone on the disk by writing data over the DCIIA to R/W ASIC serial interface.

5.2.3.2 Pulse Detector

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

5.2.3.3 Peak Detector And Servo Demodulator

The peak detector extracts the servo wedge Burst A, Burst B, and Burst C amplitude information from the read data signal. The gate inputs to the peak detector are sent from the Servo Controller module in the DCIIA. The gate inputs are used to sample the burst field amplitudes in the servo wedge.

5.2.3.4 Data Synchronizer

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

5.2.4 PreAmplifier and Write Driver IC

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

5.3 SERVO SYSTEM

5.3.1 General Description

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

The ProDrive LPS 210AT uses a sectored servo system. Positioning information is radially encoded in 78 evenly-spaced servo bursts on each track. These servo burst wedges provide radial position information for each data head, 78 times per revolution. Because the drive uses multiple zone recording, where each zone has a different bit density, split data fields are necessary to optimally utilize the non-servo area of the disk. The split data fields are achieved by special processing through the DCIIA, and their presence is transparent to the host system. The servo area remains phase coherent across the surface of the disk, even though the disk is divided into various data zones. The main advantage of the sectored servo systems is that the data heads are also servo heads, which means that sectored servo systems eliminate the problems of static and dynamic offsets between heads on different surfaces.

The ProDrive LPS 210AT servo system is also classified as a digital servo because track following compensation is done in firmware. The bump detect, on-track, velocity profiles, and other "housekeeping" tasks are also done in firmware.

The servo system has three basic modes of operation (1) Velocity mode, (2) Settle mode, and (3) Track Following Mode.

1. **Velocity Mode.** The acceleration and deceleration of the actuator for seek lengths of 20 or more tracks are accomplished with a velocity loop. The velocity profile is calculated as the seek is in progress, using programmable coefficients. The bandwidth of the velocity loop is kept low, and feed forward is heavily used.
2. **Settle Mode.** The settle mode is used for all accesses—head switches, 1- to 20-track seeks, and the end of all velocity seeks. The settle mode is a position loop with velocity damping. No feed forward is used for the settle mode.
3. **Track Following Mode.** The track following mode is used when "on-track". This is also a position loop, with an integrator in the compensation.

5.3.2 Servo Burst and Track Information

Positional information is encoded on all tracks on all data surfaces. All data heads are also servo heads. The areas with servo/position information are called wedge areas. These wedge areas are evenly spaced radially around the disk, like spokes on a wheel. There are 78 wedge areas per track. Since the disk rotation is 60.0 Hz, the position information is updated at 4680 Hz (78 x 60). This is also known as the sampling frequency f_s . The sample period, T_s , is $1/f_s = 213 \mu\text{s}$. Every wedge area consists of four separate fields: (1) Automatic Gain Control (AGC)/Sync field, (2) Sector Address Mark (SAM) field, (3) Track number and (4) Burst area. Since a Phase Lock Loop (PLL) is not used in the servo wedge area, time discrimination is used. Timing for all four fields is generated from the same crystal reference. This clock has a period of 37.5 nanoseconds, and will be referred to as T_{clk} . Data is encoded in the following manner: a servo data bit 1 = 1 0 0 1 0 0 0 0 0, and a servo data bit 0 = 1 0 0 0 0 0 1 0 0, with each digit having a T_{clk} width.

1. The AGC/sync field consists of a $3T_{\text{clk}}$ pattern, and is used by the AGC to acquire the proper amplitude for the encoded track number and position bursts. It is also used by the DCIIA Servo Controller for synchronizing to the raw data pulses. The total length of the AGC/sync field is $99T_{\text{clk}}$.
2. The SAM follows the sync field. It consists of a $14T_{\text{clk}}$ pattern repeated twice, followed by a servo data bit 0. Following the servo data bit 0 is either a servo data bit 0 or 1, which is also known as the index bit. If a one is decoded, then an index pulse is generated. The total length of the SAM field is $46T_{\text{clk}}$.
3. Following the index bit is the track number. The track number is a 12-bit Gray coded number. The Gray code to binary conversion is done in the DCIIA Servo Controller. Each Gray code bit is encoded as servo data bits. The total length of the track number is $108T_{\text{clk}}$.
4. Following the track number is the burst area. There are three bursts per servo wedge time, an A burst followed by the B burst and finally the C burst. When "on-track", the A and C burst will be at half amplitude. For even tracks the B burst will be full amplitude, and for odd tracks the B burst will be zero.

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

5.3.3 Position and Velocity

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

5.4 READ AND WRITE OPERATIONS

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

5.4.1 The Read Channel

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

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

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

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

5.4.2 The Write Channel

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

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

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

5.4.3 Interface Control

The interface with the host system is through a 40-pin AT interface connector. The DCIIA AT Interface Controller implements the AT-bus protocol (interface logic). Operating under the control of the drive's microprocessor, the AT Interface Controller receives and transmits bytes of data over the bus.

The DCIIA Buffer Controller writes data to or reads data from the Cache over 16 data lines. Under μ Controller direction, the Buffer Controller controls the transfer of data and handles the addressing of the Cache for the DCIIA. The internal data transfer rate to and from the Cache is from 16 to 28 Mb/s. This high transfer rate allows the DCIIA to communicate over the IDE bus at a PIO data transfer rate of 6.0 MB/s without using IORDY, up to 11.1 MB/s with PIO using IORDY, or a DMA transfer rate of up to 13.3 MB/s while it simultaneously controls disk-to-RAM transfers and microcontroller access to control code stored in the buffer RAM.

5.5 FIRMWARE FEATURES

This section describes the following firmware features:

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

5.5.1 Disk Caching

The ProDrive LPS 210AT hard disk drive incorporates DisCache, a 96K disk cache, to enhance drive performance. This integrated feature can significantly improve system throughput. Read and write caching can be enabled or disabled using the SET CONFIGURATION command.

5.5.1.1 Adaptive Caching

The cache buffer for the LPS 210AT features adaptive segmentation for more efficient use of the buffer's RAM. With this feature, the buffer space used for read and write operations is dynamically allocated. The cache can be flexibly divided into several segments, under program control. Each segment contains one cache entry.

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

5.5.1.2 Read Cache

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

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

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

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

- RESET
- WRITE LONG
- READ LONG
- READ DEFECT LIST
- EXECUTE DRIVE DIAGNOSTIC
- FORMAT TRACK
- IDENTIFY DRIVE
- READ CONFIGURATION
- SET CONFIGURATION
- WRITE BUFFER

5.5.1.3 Write Cache

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

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

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

5.5.1.4 Performance Benefits

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 (8.33 ms on average) when requested data resides in the cache.

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

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

For read operations, the requested read data takes up a certain amount of space in the cache segment so the corresponding prefetch data can essentially occupy the rest of the space within the segment. The 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.

The requested read data takes up a certain amount of space in the cache segment so the corresponding prefetch data can essentially occupy the rest of the space within the segment. The 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.

5.5.2 Track and Cylinder Skewing

Track and cylinder skewing in the ProDrive LPS 210AT minimizes latency time, and thus, increases data throughput. See Table 5-2 for the track and cylinder data.

5.5.2.1 Track Skewing

Track skewing reduces the latency time that results when the drive must switch read/write heads to access sequential data. A track skew is employed such that the next logical sector of data to be accessed will be under the read/write head once the head switch is made and the data is ready to be accessed. Thus, when sequential data is on the same cylinder but on a different disk surface, a head switch is needed but not a seek. Since the sequential head-switch time is well defined on the ProDrive LPS 210AT (6 ms worst case), the sector addresses are optimally positioned across track boundaries to minimize the latency time during a head switch.

5.5.2.2 Cylinder Skewing

Cylinder skewing also is used to minimize the latency time associated with a single-cylinder seek. The next logical sector of data that crosses a cylinder boundary is positioned on the drive such that after a single-cylinder seek is performed, and when the drive is ready to continue accessing data, the sector to be accessed is positioned directly under the read/write head. Therefore, the cylinder skew takes place between the last sector of data on the last head of a cylinder and the first sector of data on the first head of the next cylinder. Since single-cylinder seeks are well defined on the ProDrive LPS 210AT, the sector addresses can be optimally positioned across cylinder boundaries to minimize the latency time associate with a single-cylinder seek.

Table 5-2. *Track and Cylinder Skewing*

ZONE NUMBER	TRACK SKEW WEDGE	CYLINDER SKEW WEDGE
15	28	32
14	28	32
13	28	32
12	28	32
11	28	32
10	28	32
9	28	32
8	28	32
7	28	32
6	28	32
5	28	32
4	28	32
3	28	32
2	28	32
1	28	32
0	28	32

5.5.3 Error Detection and Correction

As disk drive areal densities increase, obtaining extremely low error rates requires a new generation of sophisticated error-correction codes. Quantum ProDrive LPS 210AT series hard disk drives implement 112-bit single- and double-burst Reed-Solomon error-correction techniques to reduce the unrecoverable read error rate to less than one bit in 1×10^{14} bits read.

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

5.5.3.1 Background Information On Error Correction Code and ECC On-The-Fly

A sector for ProDrive LPS 210 AT drives is comprised of 512 bytes of user data, followed by two cross-checking (XC) bytes (16 bits), followed by twelve ECC check bytes (96 bits). The two cross-checking bytes are used to double check the main ECC correction and reduce the probability of miscorrection. Single burst errors of up to 24 bits within one sector can be corrected "on-the-fly," in real time as they occur, allowing a high degree of data integrity with no impact to the drive's performance.

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

When errors cannot be corrected on-the-fly, an automatic retry and a more rigorous double burst error correction algorithm enable the correction of any sector with two bursts of three incorrect bytes each or up to 6 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 improves the probability of miscorrection.

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




Figure 5-6 Sector Data Field with ECC Check Bytes

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

Interleave 1	↔	d1	d4	•••••	d508	d511	xc2	ecc3	ecc6	ecc9	ecc12
Interleave 2	↔	d2	d5	•••••	d509	d512	ecc1	ecc4	ecc7	ecc10	
Interleave 3	↔	d3	d6	•••••	d510	xc1	ecc2	ecc5	ecc8	ecc11	

Figure 5-7 *Byte Interleaving*

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

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

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

Each time a sector of data is read, ProDrive LPS 210AT 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 twelve *syndromes* and two cross checking syndromes, which corresponds to the number of check bytes. If all of the syndrome values equal zero, the data was read with no errors, and the sector is transferred to the host system. If any of the syndromes do not equal zero, an error has occurred. The type of correction the drive applies depends on the nature and the extent of the error.

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

Correction of Single Burst Errors On-The-Fly

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




Figure 5-8 *Single Burst Error Correctability*

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

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

5.5.3.2 Correction Of Double-Burst Errors

Through sophisticated algorithms, ProDrive LPS 210AT drives have the capability to correct double-burst errors, even though the probability of their occurrence is low. Double-burst errors can be simply viewed as two spans of errors within one sector. More specifically, correctable double-burst errors must have two or fewer erroneous bytes per interleave. This allows the drive's Reed-Solomon ECC to correct double-burst errors up to 48 bits long (provided that the error consists of two or fewer bytes residing in each of the interleaves).

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

Double-Burst Error Examples

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

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

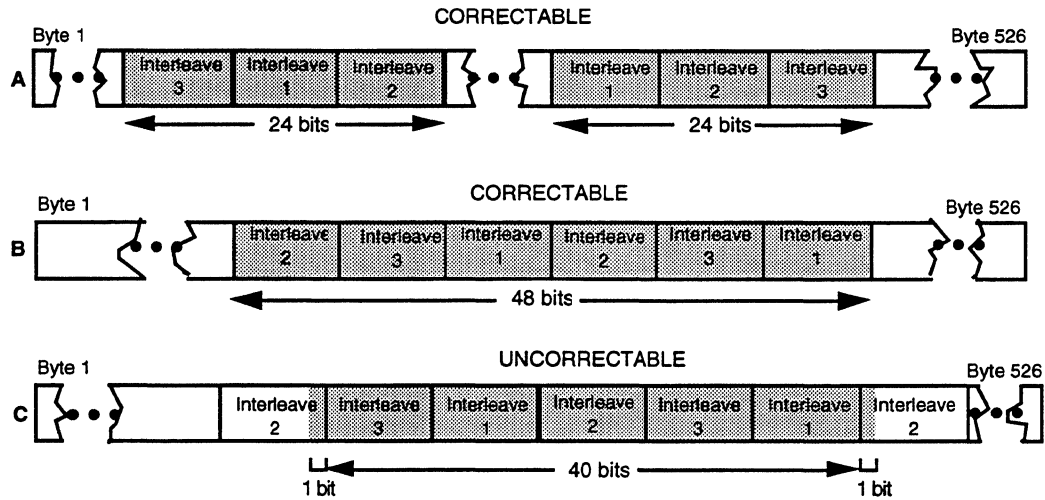


Figure 5-9. *Correctable and Uncorrectable Double-Burst Errors*

5.5.3.3 Multiple Random Burst Errors

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




Figure 5-10 Six Correctable Random Burst Errors

5.5.3.4 ECC Error Handling

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

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

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

Note: These features can be enabled or disabled through the ATA CONFIGURATION command. The EEC bit enables early ECC double-burst correction if a stable syndrome has been achieved before all of the re-reads have been exhausted. The ARRE bit enables the automatic reallocation of defective sectors.

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

If the automatic read reallocation feature is enabled, the drive will attempt up to 24 re-reads (3 times the retry count set in the AT configuration bytes) for a double-burst error. In this case, the ECC correction algorithm is divided into four parts.

Note: The ProDrive LPS 210AT is shipped from the factory with the automatic read reallocation feature enabled so that any new defective sectors can be easily and automatically reallocated.

1. When an ECC error is encountered, the drive will try to reread the correct data eight times, aborting the re-reads if the correct data was successfully read. The drive will always try to re-read the correct data before applying double-burst ECC correction and automatic sector reallocation.
2. If the first eight re-reads were unsuccessful, the drive will attempt to re-read the sector eight more times. This time, if the drive does not recover from the error with a re-read, it will look for a stable syndrome from two consecutive re-reads. If it finds a stable syndrome, it will try to correct the sector and automatically reallocate the data to a spare sector. If the sector is uncorrectable, the drive continues to re-read until the second set of eight re-reads has been exhausted.
3. In addition to the next group of eight re-reads, the drive either looks for a stable syndrome from two consecutive re-reads or for an ECC correctable error. If the drive discovers a stable syndrome from two consecutive re-reads, it will reallocate the data to a spare sector whether the data is ECC correctable or not (depending on the setting of the Reallocate Uncorrectable Error Enable or RUEE bit).

Note: The RUEE bit can be enabled or disabled through the AT CONFIGURATION command. Setting the RUEE bit enables the automatic reallocation of uncorrectable or unrecovered (e.g., ID not found) sectors.

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

Reallocation is enabled if both of the following events occur:

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

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

If the recovered data was not ECC correctable, the data is written using a special data ID field. This is done so that when data is reread, a "Read failure on uncorrectable data previously reallocated" error will be reported.

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

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

5.5.3.5 Error Types

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

Soft Errors

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

Hard Errors

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

5.5.4 Defect Management

The ProDrive LPS 210AT allocates one sector per disk or two sectors per cylinder as a spare. In the factory, the media is scanned for defects. If a sector on a cylinder is found to be defective, the address of the sector is added to the drive's defect list. Sectors located physically subsequent to the defective sector are assigned logical block addresses such that a sequential ordering of logical blocks results. This inline sparing technique is employed in an attempt to eliminate slow data transfer that would result from a single defective sector on a cylinder.

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

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

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

Chapter 6

THE IDE-BUS INTERFACE AND ATA COMMANDS

This chapter describes the interface between ProDrive LPS 210AT hard disk drives and the Integrated Drive Electronics Bus (IDE). 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

ProDrive LPS 210AT hard disk drives use the standard IBM PC IDE bus interface, and is compatible with systems that provide an IDE 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 IDE 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.

6.2 SOFTWARE INTERFACE

ProDrive LPS 210AT drives are controlled by the Basic Input/Output System (BIOS) program residing in an IBM PC AT, or 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 contains 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 IDE adapter.

If two drives are connected by a cable with two 40-pin drive connectors, the cable-select feature of the ProDrive LPS 210AT automatically configures each as drive 0 or drive 1 depending on the configuration of pin 28 on the connector. See Section 3.3.2 "Cable Select (CS) Jumper," for more information about the cable select jumper.

6.4 ELECTRICAL INTERFACE

6.4.1 IDE-Bus Interface

A 40-pin IDE 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 IDE interface contains bus drivers and receivers compatible with the standard AT bus. The AT-bus interface signals D8–D15, INTRQ, and IOCS16– require an IDE adapter board to have an extended I/O-bus connector.

The IDE interface buffers data and control 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. Neither the adapter board, motherboard interface, or drives require terminating resistors.

6.4.1.2 Drive Signals

The drive connector (J11 Section C) connects the drive to an adapter board or onboard IDE adapter in the host computer. J11-C is a 40-pin shrouded connector with two rows of 20 pins on 100-mil centers. J11 has been keyed by removing pin 20. The connecting cable is a 40-conductor flat ribbon cable, with a maximum length of 18 inches.

Table 6-1 describes the signals on the drive connector (J11-C).

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 (J11-C)

Signal	Name	Dir	Pin	Description
Reset	RESET-	IN	1	Drive reset signal from the host system, inverted on the adapter board or motherboard. Asserted for at least 300 ns during start up and deasserted thereafter, unless some event subsequently requires that the drive be reset. The RESET- signal does <i>not</i> alter the value of any drive parameters previously set, so it is not necessary to issue an INITIALIZE DRIVE PARAMETERS command after asserting RESET-.
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.
	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-; that is, the drive waits until the host asserts DMACK- before deasserting DMARQ. DMARQ is then reasserted if there is additional data to transfer.
Ground	Ground	—	22	Ground between the host system and the drive.

Table 6-1 Drive-Connector Pin Assignments (J11-C) (continued)

Signal	Name	Dir	Pin	Description
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 the host's responsibility to pull this signal up to a high level.
Cable Select (Quantum specific)		—	28	This is a Quantum specific 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).
DMA Acknowledge	DMACK1-	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	<p>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.</p> <p>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.</p> <p>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.</p>

Table 6-1 Drive-Connector Pin Assignments (J11-C) (continued)

Signal	Name	Dir	Pin	Description
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. When transferring in DMA mode, IOCS16- is not asserted and the host uses one of the 16-bit DMA channels for transferring data.
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	
Passed Diagnostics	PDIAG-	I/O	34	<p>Drive 0 (Master) monitors this Drive 1 (Slave) open-collector output signal, which indicates the result of a diagnostics command or reset. Each drive has a 10K pull-up resistor on this signal.</p> <p>Following 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.</p> <p>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-.</p> <p>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.</p>

Table 6-1 Drive-Connector Pin Assignments (J11-C)

Signal	Name	Dir	Pin	Description
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.
Drive Active/ Slave Present	DASP-	I/O	39	<p>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- within 350 ms to light the drive-activity LED.</p> <p>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.</p> <p>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.</p>
Ground	Ground	—	40	Ground between the host system and the drive.

6.4.1.3 Host System IDE Bus Signals

See Table 6-2 for the relationship between the drive signals and the IDE bus.

Table 6-2. Relationship of Drive Signals to the IDE Bus

DRIVE CONNECTOR (J11)		DIRECTION	IDE Bus	
1	RESET-	←-(INV)	B2	RESET DRV
2	GROUND	—	—	GROUND
3	DD7	↔	A2	SD7
4	DD8	↔	C11	SD8
5	DD6	↔	A3	SD6
6	DD9	↔	C12	SD9
7	DD5	↔	A4	SD5
8	DD10	↔	C13	SD10
9	DD4	↔	A5	SD4
10	DD11	↔	C14	SD11
11	DD3	↔	A6	SD3
12	DD12	↔	C15	SD12
13	DD2	↔	A7	SD2
14	DD13	↔	C16	SD13
15	DD1	↔	A8	SD1
16	DD14	↔	C17	SD14
17	DD0	↔	A9	SD0
18	DD15	↔	C18	SD15
19	GROUND	—	—	GROUND
20	KEYPIN	—	—	NO CONNECTION
21	DMARQ	→	— ¹	DRQ
22	GROUND	—	—	GROUND
23	DIOW-	←	B13	DIOW-
24	GROUND	—	—	GROUND
25	DIOR-	←	B14	IOR-
26	GROUND	—	—	GROUND
27	IORDY	→	A10	I/O CH RDY
28	CABLE SELECT	←	— ²	NO CONNECTION
29	DMACK-	←	— ³	DACK-
30	GROUND	—	—	GROUND
31	INTRQ	→	D7	INTRQ
32	IOCS16-	→	D2	I/O CS16-
33	ADDR1	←	A30	SA1
34	PDIAG-	—	—	NO CONNECTION
35	DA0	←	A31	SA0
36	ADDR2	←	A29	SA2
37	CS1FX-	←	—	CS0-
38	CS3FX-	←	—	CS1-
39	DASP-	—	—	NO CONNECTION
40	GROUND	—	—	GROUND

1. DMARQ from the drive must be jumpered to one of the DRQ lines on the motherboard or host adapter (normally connected to DRQ6).
2. Pin 28 is a vendor-specific pin that Quantum is using for a specific purpose. See Chapter 3 for details.
3. DMACK- from the drive must be jumpered to one of the DACK- lines on the motherboard or host adapter (normally connected to DACK6-).

6.4.2 Host Interface Timing

6.4.2.1 Programmed I/O (PIO) Transfer Mode

The PIO host interface timing shown in Table 6-3 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 6-3 PIO Host Interface Timing

SYMBOL	DESCRIPTION	MIN/MAX	MODE 3 ¹ (local bus)	ProDrive LPS 210AT
t0	Cycle Time	min	180	130
t1	Address Valid to DIOW-/DIOR-Setup	min	30	15
t2	DIOW-/DIOR- Pulsewidth (8- or 16-bit)	min	80	80
t2i	DIOW-/DIOR- Negated Pulsewidth	min	70	50
t3	DIOW-Data Setup	min	30	30
t4	DIOW- Data Hold	min	10	10
t5	DIOR- Data Setup	min	20	20
t5a	DIOR- to Data Valid	max	—	60
t6	DIOR- Data Hold	min	5	5
t6z	DIOR- Data Tristate	max	30	30
t7	Address Valid to IOCS16- Assertion	max	30	30
t8	Address Valid to IOCS16- Deassertion	max	30	30
t9	DIOW-/DIOR- to Address Valid Hold	min	10	10
tA	IORDY ² Setup Time	min	35	35
tB	IORDY Pulse Width	max	—	1250
tR	Read Data Valid to IORDY active (if IORDY is initially low after tA)	min	0	0

1. Timing for the proposed ATA Mode 3 specification is listed for reference only.
2. Transfer rates above 6 MB/s require the use of IORDY.




Figure 6-1 PIO Interface Timing

6.4.2.2 Multiword DMA Transfer Mode

The multiword DMA host interface timing shown in Table 6-4 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 6-4 Multiword DMA Host Interface Timing

SYMBOL	DESCRIPTION	MIN/MAX	MODE 1 ¹ (local bus)	ProDrive LPS 210 AT
t0	Cycle Time	min	150	150
tD	DIOR-/DIOW- Pulsewidth	min	80	60
tE	DIOR- to Data Valid	max	60	60
tF	DIOR- Data Hold	min	5	5
tFz	DIOR- Data Tristate ²	max	30	30
tG	DIOW- Data Setup	min	30	30
tH	DIOW- Data Hold	min	15	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	50	50
tL	DIOR-/DIOW- to DMARQ Delay	max	40	40
tZ	DMACK- Data Tristate ³	max	25	25

1. Timing for the proposed ATA Mode 1 specification is listed for reference only.
2. The ProDrive LPS 210AT tristates after each word transferred.
3. t_Z only applies on the last tristate at the end of a multiword DMA transfer cycle.

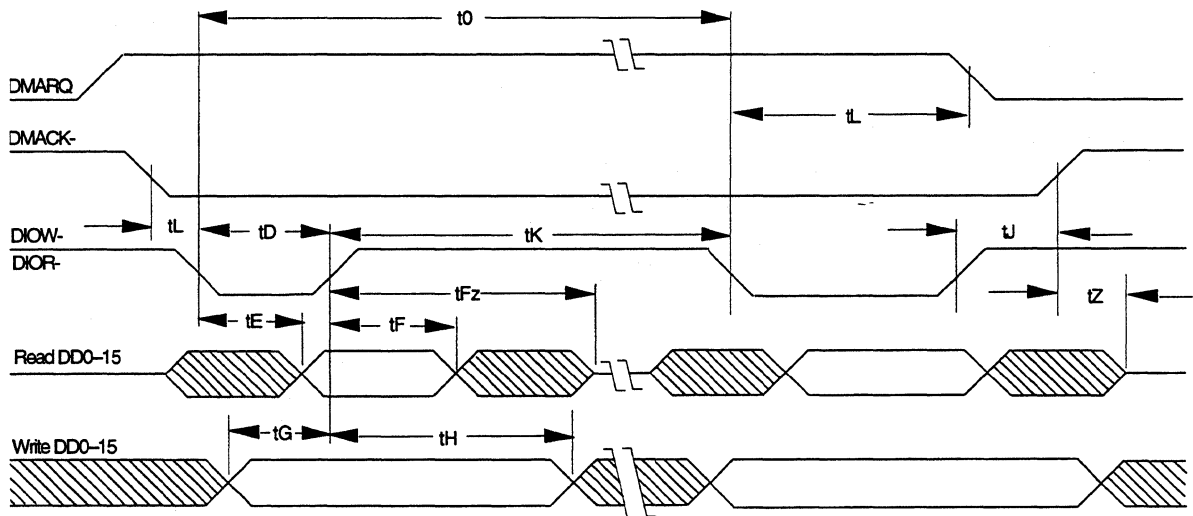


Figure 6-2 Multiword DMA Bus Interface Timing

6.4.2.3 Host Interface RESET Timing

The host interface RESET 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 6-4 provides a timing diagram.

Table 6-5 *Host Interface RESET Timing*

SYMBOL	DESCRIPTION	MINIMUM	MAXIMUM
tM	RESET- Pulse width	300	—




Figure 6-3 *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 6-6 lists the selection addresses for these registers.

Table 6-6. 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	Sector Number	A	N	0	1	1
LBA Bits 0 - 7 ⁴	LBA Bits 0 - 7 ⁴	A	N	0	1	1
Cylinder Low	Cylinder Low	A	N	1	0	0
LBA Bits 8 - 15 ⁴	LBA Bits 8 - 15 ⁴	A	N	1	0	0
Cylinder High	Cylinder High	A	N	1	0	1
LBA Bits 16 - 23 ⁴	LBA Bits 16 - 23 ⁴	A	N	1	0	1
Drive/Head	Drive/Head	A	N	1	1	0
LBA Bits 24-27 ⁴	LBA Bits 24-27 ⁴	A	N	1	1	0
Status	Command			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 LBA mode

At power on or following a reset, the drive initializes the Command Block Registers to the values shown in Table 6-7.

Table 6-7. Command Block Register Initial Values

REGISTER	VALUE (HEX)
Error Register	01
Sector Count Register	01
Sector Number Register	01
Cylinder Low Register	00
Cylinder High Register	00
Drive/Head Register	A0

6.6 REGISTER DESCRIPTIONS

The ProDrive LPS 210AT hard disk drive emulates the AT 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 (3F6)

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 Section 6.6.2.8 for definitions of bits in the Status Register.

6.6.1.2 Device Control Register (3F6)

This write-only register contains two control bits, as shown in Table 6-8.

Table 6-8. Device Control Register Bits

BIT	MNEMONIC	DESCRIPTION
7	Reserved	-
6	Reserved	-
5	Reserved	-
4	Reserved	-
3	1	Always 1
2	SRST	Host Software Reset bit
1	nIEN	Drive Interrupt Enable bit
0	0	Always 0

Note: 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.

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 (3F7)

The Drive Address Register returns the head-select addresses for the drive currently selected. Table 6-9 shows the Drive Address bits.

Table 6-9. Drive Address Register Bits

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

Note: HiZ = High Impedance bit. When the host reads the register, this bit will be in a high impedance state.

nWTG = Write Gate bit. When a write operation to the drive is in progress, nWTG equals 0.

nHS0–nHS3 = Head Address bits. These bits are equivalent to the one's complement of the binary-coded address of the head currently selected.

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. Transfers of ECC bytes during the execution of READ LONG or WRITE LONG commands are 8-bit transfers.

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—except at power on or 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 6-10.

Table 6-10. *Error Register Bits*

BIT	MNEMONIC	DESCRIPTION
7	BBK	Bad block detected in the required sector's ID field.
6	UNC	Uncorrectable data error encountered.
5	—	Not used.
4	IDNF	Requested sector's ID field not found.
3	—	Not used.
2	ABRT	Requested command aborted due to a drive status error, such as Not Ready or Write Fault, or because the command code is invalid.
1	TK0NF	Track 0 not found during execution of RECALIBRATE command.
0	AMNF	Data Address Mark not found after correct ID field format. (Not used.)

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.

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 host 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.

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 host 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.

6.6.2.7 Drive/Head Register

The Drive/Head Register contains the drive ID number and its head numbers. When the drive executes an INITIALIZE DRIVE PARAMETERS or Format Track command, the value in this register defines the number of heads per cylinder.

Table 6-11 shows the Drive/Head Register bits.

Table 6-11. *Drive/Head Register Bits*

BIT	MNEMONIC	DESCRIPTION
7	Reserved	Always 1
6	L	L = 0 CHS Mode L = 1 LBA Mode
5	Reserved	Always 1
4	DRV	0 indicates the Master drive is selected 1 indicates the Slave drive is selected
3	HS3	Most significant Head Address Bit
2	HS2	—
1	HS1	—
0	HS0	Least-significant Head Address bit

Note: In Table 6-11, the following conventions apply:

Bits 5 and 7 define the sector size set in hardware (512 bytes).

Bit 6, L = 0, only CHS mode supported.

DRV is the binary-encoded drive- select number. The Master is the primary drive; the Slave is the secondary drive.

HS0–HS3 constitute a binary encoded number that represents the address of the selected head. On command completion, the host updates these bits to reflect the address of the head currently selected.

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 is acknowledging the interrupt when the host reads the Status Register. Therefore, whenever the host reads this register, the drive clears the pending interrupt. Table 6-12 defines the Status Register bits.

Table 6-12. *Status Register Bits*

BIT	MNEMONIC	DESCRIPTION
7	BSY	<p>Busy bit. Set by the controller logic of the drive. While BSY is set, the drive has access to and the host is locked out of the Command Block Registers. BSY is set under the following conditions:</p> <ul style="list-style-type: none"> • 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. <p>While BSY=1, the host cannot write to a Command Block Register, and reading any Command Block Register returns the contents of the Status Register.</p>
6	DRDY	<p>Drive Ready bit. Indicates that the drive is ready to accept a command when set. When an error occurs, this bit remains unchanged until the host reads the Status Register. At power on, this bit is cleared, and remains cleared until the drive is up to speed and ready to accept a command.</p>
5	DWF	<p>Drive Write Fault bit. When an error occurs, this bit remains unchanged until the host reads the Status Register, then again indicates the current write fault status following the read of the Status Register.</p>
4	DSC	<p>Drive Seek Complete bit. This bit is set when a seek operation is complete and the heads have settled over a track. When an error occurs, this bit remains unchanged until the host reads the Status Register. It indicates the current seek-complete status after the read of the Status Register.</p>
3	DRQ	<p>Data Request bit. When set, this bit indicates that the drive is ready to have a word or byte of data transferred from the host to the data port.</p>
2	CORR	<p>Corrected Data bit. The drive sets this bit when it encounters and corrects a correctable data error. This condition does not terminate a multisector read operation.</p>
1	IDX	<p>Index bit. This bit is set and cleared when the drive detects the index mark, once per disk revolution.</p>
0	ERR	<p>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.</p>

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 6-13 lists the hexadecimal command codes and parameters for each executable command. The code F0_H is common to all of the extended commands. Each of these commands is distinguished by a unique subcode. For all registers except the Drive/Head register, Y indicates that the register contains a valid parameter for this command. For the Drive/Head Register, Y indicates that the drive uses both the drive and head parameters. D indicates that only the drive parameter is valid, not the head parameter. D* indicates that the host addresses the parameter to drive 0, but both drives execute the command. Where a parameter is blank, the command does not require the contents of the register.

For a detailed description of each command, see Section 6.7.

Table 6-13. ATA Command Codes and Parameters

ATA COMMANDS		PARAMETERS ¹				
NAME	CODE (HEX)	FR	SC	SN	CY	DH
RECALIBRATE	1XH					D
READ SECTORS, with retry	20H		Y	Y	Y	Y
READ SECTORS, no retry	21H		Y	Y	Y	Y
READ LONG, with retry	22H		Y	Y	Y	Y
WRITE SECTORS, with retry	30H		Y	Y	Y	Y
WRITE SECTORS, no retry	31H		Y	Y	Y	Y
WRITE LONG, with retry	32H		Y	Y	Y	Y
READ VERIFY SECTORS, with retry	40H		Y	Y	Y	Y
READ VERIFY SECTORS, no retry	41H		Y	Y	Y	Y
FORMAT TRACK	50H		Y		Y	Y
SEEK	7XH				Y	Y
EXECUTE DRIVE DIAGNOSTIC	90H					D*
INITIALIZE DRIVE PARAMETERS	91H		Y			Y
CHECK POWER MODE	98H E5H		Y			D
READ MULTIPLE	C4H		Y	Y	Y	Y
WRITE MULTIPLE	C5H		Y	Y	Y	Y
SET MULTIPLE MODE	C6H		Y			D
READ DMA, with retry	C8H		Y	Y	Y	Y
READ DMA, no retry	C9H		Y	Y	Y	Y
WRITE DMA, with retry	CAH		Y	Y	Y	Y
WRITE DMA, no retry	CBH		Y	Y	Y	Y
STANDBY MODE	E0H					
IDLE MODE	E1H 95H					
STANDBY MODE – AUTO POWER DOWN	E2H 96H					
IDLE MODE – AUTO POWER DOWN	E3H					
READ BUFFER	E4H					D
SLEEP MODE	E6H 99H					
WRITE BUFFER	E8H					D
IDENTIFY DRIVE	ECH					D
SET FEATURES	EFH	Y				D
READ DEFECT LIST (extended command)	F0H		Y	Y	Y	Y
READ CONFIGURATION (extended command)	F0H		Y	Y	Y	Y
SET CONFIGURATION (extended command)	F0H		Y	Y	Y	Y

Notes: The following information applies to Table 6-13:

- 1) SC = Sector Count Register SN = Sector Number Register CY = Cylinder Register
DH = Drive/Head Register
- 2) For all registers except the Drive/Head register, Y indicates that the register contains a valid parameter for this command. For the Drive/Head Register, Y indicates that the drive uses both the drive and head parameters.
- 3) D indicates that only the drive parameter is valid, not the head parameter.
- 4) D* indicates that the host addresses the parameter to drive 0, but both drives execute the command.
- 5) Where a parameter is blank, the command does not require the contents of the register.

6.7 COMMAND DESCRIPTIONS

The ProDrive LPS 210AT hard disk drive supports all standard AT 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. Load the required parameters into the Command Block Register.
3. Activate the Interrupt Enable (–IEN) bit.
4. Wait for the drive to set RDY (RDY=1).
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 6-13.

6.7.1 Recalibrate

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

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 Read Long

When the Long bit is set in the command code, a READ LONG command executes, returning the data and the ECC bytes contained in the data field of the requested sector. During a READ LONG operation, the drive does not check the ECC bytes to determine if a data error of any kind has occurred.

6.7.2.2 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 error has occurred, 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

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 Write Long

When the Long bit is set in the command code, a Write Long command writes the data and the ECC bytes directly from the sector buffer. The drive does not generate the ECC bytes itself.

6.7.3.2 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 has filled 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 error occurred, in what sector.

6.7.4 Read Verify Sectors

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 Format Track

The host specifies the track addresses by writing to the Cylinder and Head Registers. When the drive accepts a FORMAT TRACK command, it sets the DRQ bit, then waits for the host to fill the sector buffer. When the buffer is full, the drive clears DRQ, sets BSY, and begins command execution. The contents of the sector buffer are not written to the disk, and may be ignored or interpreted as shown in Table 6-14.

Table 6-14 Sector Buffer Contents

DD15 -- DD0		DD15 -- DD0	
First Sector Descriptor	: : : : : : : : : : : :	Last Sector Descriptor	Remainder of buffer filled with zeros

On the ProDrive LPS 210AT hard disk drive, the FORMAT TRACK command writes zeros to the data fields in the sectors on the specified logical track. The drive writes no headers at these locations. The Sector Count register contains the number of sectors per track.

One 16-bit word represents each sector (the words are contiguous from the start of the sector).

Note: Any words remaining in the buffer after the representation of the last sector must be filled with zeros.

DD15-8 contain the sector number. If an interleave is specified, the words appear in the same sequence as they appear on the track. DD7-0 contain a descriptor value defined as follows:

- 00h - Format sector as good
- 20h - Unassign the alternate location for this sector
- 40h - Assign this sector to an alternate location
- 80h - Format sector as bad

6.7.6 Seek

The SEEK command causes the actuator to seek to 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. Since 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.

6.7.7 Execute Drive Diagnostic

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 five seconds for drive 1 to assert PDIAG $\bar{}$.
- If drive 1 does not assert PDIAG $\bar{}$ to indicate a failure, drive 0 appends 80 H 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 80 H 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 00 H 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 6-14 lists the diagnostic codes.

Table 6-15. *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

6.7.8 Initialize Drive Parameters

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.

6.7.9 Read Multiple

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

- It transfers multiple-sector blocks of data to the host without intervening interrupts.
- It requires DRQ qualification of the transfer only at the start of the block—*not* at each sector.
- It 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:

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

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 Aborted 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.

6.7.10 Write Multiple

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

- It causes the controller to set BSY within 400 nsec of accepting the command.
- It causes the drive to transfer multiple-sector blocks of data to the drive without intervening interrupts.
- It requires DRQ qualification of the transfer only at the start of the block, not at each sector.
- It 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:

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

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 Aborted 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—and does not transfer subsequent blocks. The drive generates interrupts by setting DRQ at the beginning of each block or partial block.

6.7.11 Set Multiple Mode

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. Valid block sizes are 1 to 8. 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, or after a software or hardware reset, the default mode is READ MULTIPLE and WRITE MULTIPLE commands are disabled.

6.7.12 Read Buffer

The READ BUFFER command enables the host to read the current contents of the drive's buffer. When the host issues this command, the drive sets BSY, sets up the first 512 bytes of the buffer for a read operation, sets DRQ, clears 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.

6.7.13 Write Buffer

The WRITE BUFFER command allows the host to write the first 512 bytes of the drive's 64K 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 write to the buffer.

6.7.14 Read DMA

The READ DMA command is similar to the READ SECTORS command, with the following exceptions:

- Prior to issuing the command, the host must initialize a slave DMA channel
- DMARQ is used to qualify the data transfers and data transfer is performed over the slave DMA channel
- The drive issues a single interrupt for each command, which indicates that the data transfer is done and status is available

During the execution of a READ DMA command, if any unrecoverable error is found, the data transfer terminates at the sector where the error was detected. The sector containing the error is not transferred. The drive issues an interrupt to indicate that the data transfer is finished and status is available. Error posting is done in the same way as in a READ SECTORS command.

6.7.15 Write DMA

The WRITE DMA command is similar to the WRITE SECTORS command, with the following exceptions:

- Prior to issuing the command, the host must initialize a slave DMA channel
- DMARQ is used to qualify the data transfers and data transfer is performed over the slave DMA channel
- The drive issues a single interrupt for each command, which indicates that the data transfer is done and status is available

During the execution of a WRITE DMA command, if any unrecoverable error is found, the data transfer terminates. The drive issues an interrupt to indicate that the data transfer is finished and status is available in the Error Register. Error posting is done in the same way as in a WRITE SECTORS command.

6.7.16 Power Management Commands

The ProDrive LPS 210AT 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 upon which of the two commands was most recently issued:

- **Standby:**

Once a standby command is issued, the drive immediately enters the standby mode. In the standby mode, the spindle and actuator motors are off with the heads 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 exits the standby command, the drive returns to the standby mode after the APD counter reaches 0.

- **Idle:**

CAM Option — Once an idle command is issued, each time the APD counter reaches zero, the drive enters the idle mode. In the idle mode, the spindle motor and actuator power are on, and the drive is running.

Non-CAM Option (Power Save) — 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.

Note: There are two timers: the standby timer and the idle timer. If the APD is enabled using a standby command, that same command must be used to disable the APD. This also applies to the idle command.

The ProDrive LPS 210AT drives support the 30 minute incremental timeouts as per the most current CAM specifications. The values 241 (F1_H) to 251 (FB_H) inclusive are defined as 30 minutes (i.e., 241 is 30 minutes, 242 is 60 minutes, 243 is 90 minutes, up to 251, which is 5.5 hours).

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 with the heads latched in the landing zone. Receipt of a reset causes the drive to transition from the sleep to the idle 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. Each increment represents a five-second time interval.

6.7.16.1 Standby Immediate Mode (E0_H)

The Standby 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.

6.7.16.2 Idle Immediate Mode (E1_H)

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

6.7.16.3 Standby Mode (Auto Power-Down E2_H)

The Standby Mode Auto Power-Down 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. Each count represents a five second increment. The valid count range is from 12 to 255. 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 set to zero, the drive immediately enters standby mode but subsequent Auto Power-Down feature is disabled.

6.7.16.4 Idle Mode (Auto Power-Down E3_H)

The Idle Mode Auto Power-Down command immediately puts the drive in 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. Each count represents a five second increment. The valid count range is from 12 to 255. 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.

6.7.16.5 Check Power Mode (E5_H)

The CHECK POWER MODE command writes FF_H into the Sector Count Register provided that the drive is in the Idle Mode. If the drive has entered, is entering, or is recovering from the Standby Mode, the value of 00H is written to the Sector Count Register.

6.7.16.6 Sleep Mode (E6_H)

ProDrive considers the Sleep Mode to be the equivalent to the Standby Mode, except a reset is required before issuing a command that requires media access.

6.7.17 Identify Drive

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 6-15, defines the parameter words stored in the buffer. All reserved bits should be zeros. A full explanation of the parameter words is listed below:

Number of Fixed 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 (20_H).

Buffer Type:

The contents of this field are as follows:

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

Firmware Revision:

The contents of this field are left-aligned and padded with spaces (20_H).

Model Number:

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

Table 6-16. Identify Drive Parameters

WORDS ¹			PARAMETERS
WORD	BIT	BIT VALUE	DESCRIPTION (Statements below are true if the bit is set to a one)
0	15	0	Reserved for nonmagnetic drives
	14	0	Format speed tolerance gap required
	13	0	Track offset option available
	12	0	Data strobe offset option available
	11	1	Rotational speed tolerance is > 0.5%
	10	0	Disk transfer rate > 10 Mbit/s
	9	1	Disk transfer rate > 5 Mbit/s, but ≤ 10 Mbit/s
	8	0	Disk transfer rate ≤ 5 Mbit/s
	7	0	Reserved for removable-cartridge drive
	6	1	Hard disk drive
	5	0	Spindle motor control option implemented
	4	1	Head-switch time > 15 μs
	3	1	Not MFM-encoded
	2	0	Soft-sectored
	1	1	Hard-sectored
	0	0	Reserved
WORD	BIT	VALUE	DESCRIPTION
1		723	Default logical cylinders
2		00h	Reserved
3		15	Number of logical heads
4		19,456	Number of unformatted bytes per track
5		512	Number of unformatted bytes per sector
6		38	Number of logical sectors per track
7–9		5154h	Vendor-unique
10–19			Serial number (20 ASCII characters) ²
20		3	Buffer type
21		192	Buffer size in 512-byte increments

WORDS ¹			PARAMETERS
22		4	Number of ECC bytes passed on READ/WRITE LONG commands
23-26			Firmware revision (8 ASCII characters)
27-46			Model number (40 ASCII characters)
47	15-8	80h	Vendor-unique
	7-0	08h	Maximum number of sectors that can be transferred per interrupt is set to 8 for READ and WRITE MULTIPLE commands.
48		0	Cannot perform double word I/O
WORD	BIT	VALUE	DESCRIPTION
49	15-12	0	Reserved
	11	1	1 = I/O Ready is supported
	10	1	1 = I/O Ready can be disabled
	9	0	1 = LBA supported
	8	1	1 = DMA supported
	7-0	0	Vendor unique
50		0	Reserved
51	15-8	2	PIO data-transfer cycle timing mode
	7-0	0	Vendor-unique
52	15-8	2	DMA data-transfer cycle timing mode
	7-0	0	Vendor-unique
53	15-2	0	Reserved
	1	1	1 = The fields in words 64 - 70 are valid
	0	1	1 = The fields in words 54-58 are valid
54		723	Number of current cylinders
55		15	Number of current heads
56		38	Number of current sectors per track
57-58		412,110	Current capacity in sectors
59	15-9	0	Reserved
	8	1	Multiple sector setting is valid
	7-0	n ³	Current setting for number of sectors that can be transferred per interrupt on R/W Multiple commands
60-61		412,110	Total number of User Addressable Sectors (LBA Mode only)
62	15-8	4	Single-word DMA transfer mode active (Mode 2)
	7-0	7	Single-word DMA transfer modes supported (Mode 2)
63	15-8	2	Multiword DMA transfer mode active (Mode 1)
	7-0	3	Multiword DMA transfer modes supported (Mode 1)
64		1	Advanced PIO Mode is supported
65		150	Minimum multiword DMA transfer cycle time (ns) per word
66		150	Manufacturer's recommended multiword DMA cycle time (ns)
67		333	Manufacturer's PIO cycle time (ns) without flow control
68		180	Manufacturer's PIO cycle time (ns) with flow control
69-127			Reserved
128-159			Vendor-Unique
160-255			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: 00QTMTCYJJLSSSSBBB
 where: 00 = Placeholders QT = Quantum
 M = Place of manufacture C = Drive capacity (4 = 210MB)
 (1 = Japan, 2 = Ireland)
 Y = Last digit of year drive built JJJ = Julian day
 L = Manufacturing production line SSSS = Sequence of manufacture
 BBBB = Blanks (placeholders) T = Drive type family (fixed at 8)
3. n is variable from zero to eight.

6.7.18 Set Features

The SET FEATURES command is used by the host to establish parameters for the execution of certain drive features. When the drive receives this command, it sets BSY, checks the contents of the Features register, clears BSY, and generates an interrupt. The host should load the Feature register with the appropriate command. If the value in the register is not supported or is invalid, the drive aborts the command. The features function code is:

- 0AA_H Enables Read Caching
- 055_H Disables Read Caching
- 002_H Enables Write Caching
- 082_H Disables Write Caching

The settings for these features revert to their saved values after a power cycle.

6.7.19 Read Defect List

The READ DEFECT LIST command enables the host to retrieve the drive's defect list. When the host writes the bytes at addresses 1F2_H–1F6_H, followed by F0_H at address 1F7_H, the drive sets BSY, retrieves the defect list, sets DRQ, and resets BSY. The first byte, 04_H, is a subcode to the extended command code F0_H. The host can now read four sectors (512 bytes) of data. An INTRQ precedes each sector. Table 6-17 lists the bytes for this command.

Table 6-17. READ DEFECT LIST Command Bytes

ADDRESS	VALUE	DEFINITION
1F2	05H	Defect List Subcode
1F3	FFH	Password
1F4	FFH	Password
1F5	3FH	Password
1F6	A0H (Drive 0)	Drive Select
	B0H (Drive 1)	—
1F7	F0H	Extended Command Code

In Table 6-17:

- Registers 1F2_H through 1F5_H 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.
- To select the proper drive, set register 1F6_H. For execution of the command to begin, load register 1F7_H with F0_H.

The seven bytes listed in Figure 6-5 describe each defect or replacement in a physical cylinder, head, and sector.

6.7.20 Defect List Data Structure

The defect lists maintained and accessed by the defect management system consist of 7 byte defect entries. The P list contains only defect entries while the W list contains both defect and replacement cylinder information. The defect list structure is illustrated in Figure 6-5.

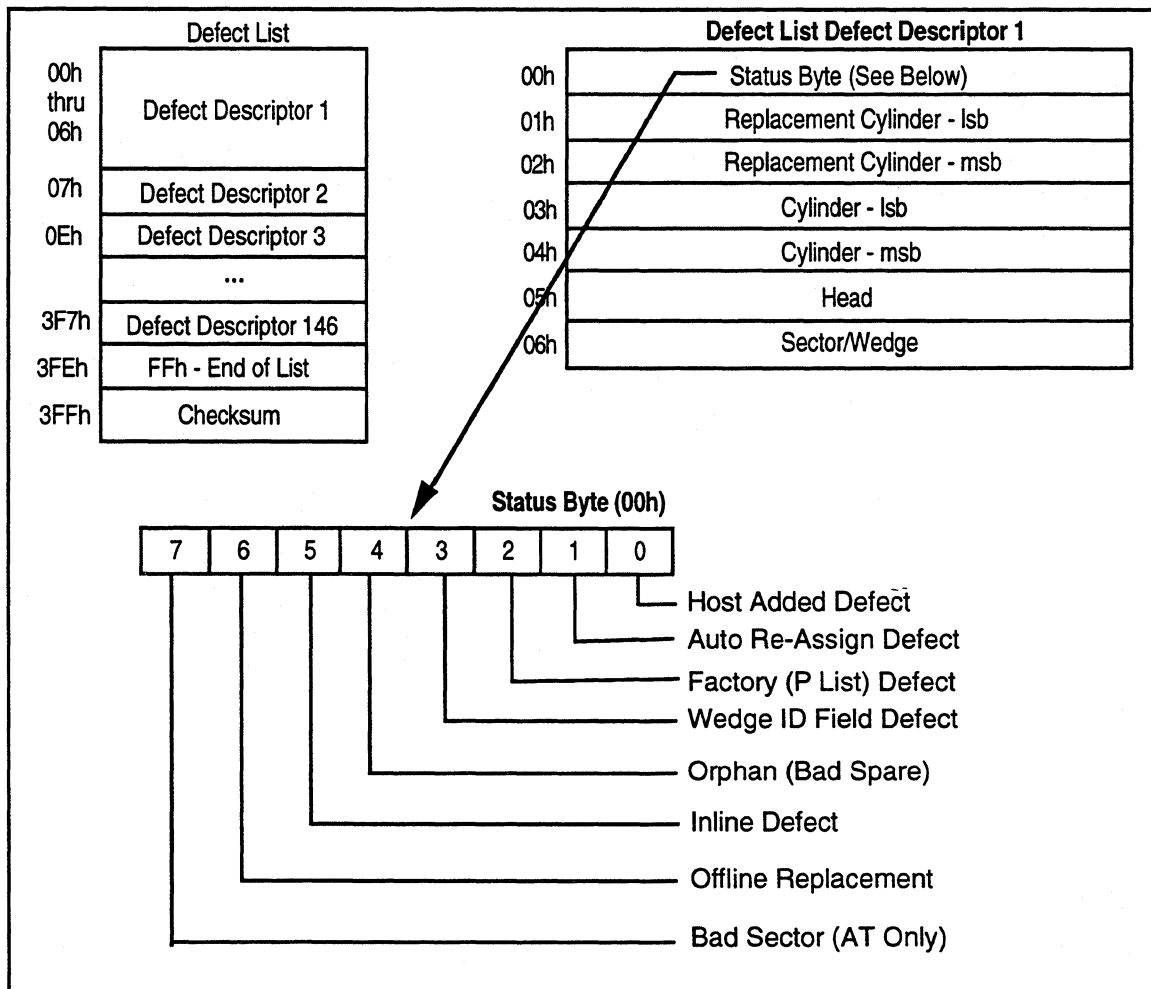


Figure 6-4. Defect List Data Structure

The end of list marker is placed after the last entry in the list (i.e., if there is only one defect, the end of list marker would be located at byte 06h).

The checksum is placed at the end of the list, and the empty area in the list is filled with zeros. When this byte is added to the rest of the bytes in the list, the **lsb** of the checksum will equal ASCII "L".

Table 6-17 shows a sample defect list.

Table 6-18. Sample Defect List

ENTRY	VALUE	DESCRIPTION
Defect Descriptor 1	24 _H	Defect found at factory, on head 3 (In-line sparing transparent to user)
	07 _H	Replacement Cylinder = 07 (LSB)
	00 _H	Replacement Cylinder = 00 (MSB)
	07 _H	Defective Cylinder = 07
	00 _H	Defective Cylinder = 00 (MSB)
	03 _H	Defective Head = 03
	02 _H	Defective Sector = 02
Defect	00 _H	Defect found in field, on head 2 (Automatically reallocated to location indicated)
	08 _H	Replacement Cylinder = 08
	00 _H	Replacement Cylinder = 00 (MSB)
	07 _H	Defective Cylinder = 07
	00 _H	Defective Cylinder = 00 (MSB)
	02 _H	Defective Head = 02
	04 _H	Defective Sector = 04
	FF _H	End of list marker

6.7.21 Configuration

In addition to the SET FEATURES command, the ProDrive LPS 210AT hard disk drive provides 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. Also see Chapter 5 for more information about error detection and defect management.

6.7.21.1 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 6-18 to the Command Block Registers. The first byte, 01H, is a subcode to the extended command code, F0H.

Table 6-19. *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 BXH – Drive 1	Drive Select
1F7H	F0H	Extended Command Code

In Table 6-18:

- 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 the table. These values function as a key. The drive issues the message ILLEGAL COMMAND if the key is not entered correctly.
- The X indicates the value can be any hex value (0–F).
- To select the drive for which the configuration is to be displayed, 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, shown in Table 6-18, is similar to that for the data field of the SET CONFIGURATION command. 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.

6.7.21.2 Set Configuration

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 6-19 to the Command Block Registers. The first byte, FFH, is a subcode to the extended command code F0H.

Table 6-20. *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 BXH – Drive 1	Drive Select
1F7H	F0H	Extended Command Code

In Table 6-19:

- Registers 1F2_H through 1F5_H 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.
- The X indicates the value can be any hex value (0–F).
- To select the drive being reconfigured, register 1F6_H should be set. For execution of the command to begin, load register 1F7_H with F0_H.

6.7.21.3 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 6-20 to the Command Block Registers. The first byte, FE_H, is a subcode to the extended command code F0_H.

Table 6-21. *Set Configuration (Without Saving) Access Pattern*

ADDRESS	VALUE	DEFINITION
1F2 _H	FE _H	Set Configuration Subcode
1F3 _H	FF _H	Password
1F4 _H	FF _H	Password
1F5 _H	3F _H	Password
1F6 _H	AX _H – Drive 0 BX _H – Drive 1	Drive Select
1F7 _H	F0 _H	Extended Command Code

In Table 6-20:

- Registers 1F2_H through 1F5_H 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.
- The X indicates the value can be any hex value (0–F).
- To select the drive being reconfigured, set register 1F6_H. For execution of the command to begin, load register 1F7_H with F0_H.

6.7.21.4 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 6-21 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 39 control the operation of DisCache and error recovery procedure. The drive does not use bytes 40 through 511, which should be set to 0.

Table 6-22. Configuration Command Format

BYTE	BIT							
	7	6	5	4	3	2	1	0
0–31	QUANTUM CONFIGURATION KEY							
32	RESERVED = 0						PE	CE
33	NUMBER OF SEGMENTS							
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	0	AR=0	ACHS	0	0	WCE	RÜEE	R = 0
40–511	RESERVED = 0							

Note: All fields marked RESERVED = 0, R = 0 (short for RESERVED = 0), or N/A should be set to zero.

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 issues the message ILLEGAL COMMAND.

DisCache Parameters (Byte 32)

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 64K RAM only as a transfer buffer. The default setting is 1.

Number of Cache Segments (Byte 33)

Byte 33 contains the number of cache segments available in the cache table. The default value for this byte is 1 for the ProDrive LPS 210AT.

Error Recovery Parameters (Byte 36)

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 ProDrive LPS 210AT 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, the same nonzero ECC syndrome occurs on two consecutive retries. 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 Early 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 size, in bits, of the largest Read Data error on which the drive can attempt error correction. The value for this byte is fixed at 16. This will correct a single burst of up to 48 bits in length, or one double burst of up to 24 bits each in length.

DRIVE PARAMETER (Byte 39)

AR (Auto Read) (Byte 39, Bit 6): When this bit is set to 1, Auto Read is enabled. When the bit is set to 0, Auto Read is disabled. The default value is 0.

ACHS (Auto CHS (Cylinder, Head, Sector)) (Byte 39, Bit 5): When ACHS is set to 1, the hardware automatically updates the CHS value in the task file register. When the bit is 0, the CHS is not automatically updated. The default value is 1.

WCE – Write Cache Enable (Byte 39, Bit 2): When this bit is set to 1, the ProDrive LPS 210AT hard disk enables the Write Cache. The host will terminate Write commands before data is completely written to the disk. When Write Caching is enabled, the WRITE command terminates when all of the data is transferred from the host to the buffer, even though all the data may not be written to the disk. Data continues to be written to the disk after the WRITE command terminates. If the following command is any command other than another WRITE command, the command is delayed until the remaining data is written to the disk. Upon completion, the new command continues normal execution.

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 ProDrive LPS 210AT hard disk drive will automatically reallocate uncorrectable hard errors, if the ARR bit (byte 36, bit 6) is set to 1. The default setting is 1.

6.8 ERROR REPORTING

At the start of a command's execution, the ProDrive LPS 210AT hard disk drive checks 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 is occurred. Table 6-22 lists the valid errors for each command.

Table 6-23. Command Errors

COMMAND	ERROR REGISTER						STATUS REGISTER				
	BBK	UNC	IDNF	ABRT	TK0	AMNF	DRDY	DWF	DSC	CORR	ERR
Check Power Mode				V			V	V	V		V
Read Defect List	V	V	V	V		V	V	V	V	V	V
Execute Drive Diagnostics											V
Format Track			V	V			V	V	V		V
Identify Drive				V			V	V	V		V
Initialize Parameters							V	V	V		
Invalid Command Codes				V			V	V	V		V
Read Buffer				V			V	V	V		V
Read DMA	V	V	V	V		V	V	V	V	V	V
Read Configuration	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 Sectors Long	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
Write Buffer				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
Write Sectors Long	V		V	V			V	V	V		V

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

Glossary

A

ACCESS – (v) Read, write, or update information stored on a disk or other medium. (n) The operation of reading, writing, or updating stored information.

ACCESS TIME – The interval between the time a request is made by the system and the time the data is available from the drive. Includes the seek time, rotational latency, and command processing overhead time. (See also *seek*, *rotational latency*, and *overhead*.)

ACTUATOR – Also known as the *positioner*. The internal mechanism that moves the read/write head to the proper track. The Quantum actuator consists of a rotor connected to head mounting arms that position the heads over the desired cylinder. Also known as rotary actuator.

AIRLOCK – A patented Quantum feature that ensures durable and reliable data storage. Upon removal of power from the drive for any reason, the read/write heads automatically park and lock in a non data area called the landing zone. AIRLOCK allows the drive to withstand high levels of non-operating shock. When power is applied to the drive, airflow created from the spinning disks causes the AIRLOCK arm to swing back and unlock the actuator, allowing the heads to move from the landing zone. Upon power down, the AIRLOCK swings back to the locked position, locking the heads in the landing zone. A park utility is not required to park the heads on drives equipped with AIRLOCK (all Quantum drives).

ALLOCATION – The process of assigning particular areas of the disk to specific data or instructions. An allocation unit is a group of sectors on the disk reserved for specified information. On hard disks for small computer systems, the allocation unit is usually in the form of a sector, block, or cluster. (See also *allocation unit*.)

ALLOCATION UNIT – An allocation unit, also known as a *cluster*, is a group of sectors on the disk that can be reserved for the use of a particular file.

ASIC – Acronym for *Application Specific Integrated Circuit*.

AT – An interface designed for IBM PCs and compatible systems. Also known as IDE. (See also *interface*.)

AVERAGE SEEK TIME – The average time it takes for the read/write head to move to a specific location. Calculated by dividing the time it takes to complete a large number of random seeks by the number of seeks performed.

B

BACKUP – A copy of a file, directory, or volume on a separate storage device from the original, for the purpose of retrieval in case the original is accidentally erased, damaged, or destroyed.

BAD BLOCK – A block (usually the size of a sector) that cannot reliably hold data due to a physical flaw or damaged format markings.

BAD TRACK TABLE – A label affixed to the casing of a hard disk drive stating which tracks are flawed and cannot hold data. This list is typed into the low-level formatting program when the drive is installed. Quantum users can ignore bad track tables since Quantum's built-in defect-management protections compensate for these flaws automatically.

BEZEL – A plastic panel that extends the face of a drive so that it covers a computer's drive bay opening. The bezel usually contains a drive-activity LED. Also known as the *faceplate*.

BIOS – Acronym for Basic Input / Output System. The firmware portion of a computer that manages the flow of signals through the system bus and to the attached cards and peripheral devices.

BIT – Abbreviation for binary digit. A binary digit may have one of two values—1 or 0. This contrasts with a decimal digit, which may have a value from 0 to 9. A bit is one of the logic 1 or logic 0 binary settings that make up a byte of data. (See also *byte*.)

BLOCK – In UNIX workstation environments, the smallest contiguous area that can be allocated for the storage of data. UNIX blocks are generally 8 Kbytes (16 sectors) in size. In DOS environments, the block is referred to as a cluster. (Note: This usage of the term block at the operating system level is different from its meaning in relation to the physical configuration of the hard drive. See *sector* for comparison.) While an operating system may have its block size set at 8K, the drive usually assigns its "logical block" size to 1 sector (1 OS block = 16 drive logical blocks).

BPI – Bits Per Inch. A measure of how densely information is packed on a storage medium. (See also *FCI*.)

BUFFER – An area of RAM reserved for temporary storage of data that is waiting to be sent to a device that is not yet ready to receive it. The data is usually on its way to or from the hard disk drive or some other peripheral device.

BUS – The part of a chip, circuit board, or interface designed to send and receive data.

BYTE – The basic unit of computer memory, large enough to hold one character of alphanumeric data. Comprised of eight bits. (See also *bit*.)

C

CACHE – High-speed RAM used as a buffer between the CPU and the hard disk. Since the CPU can get information more quickly from the cache than from main memory, the cache usually contains information that is used frequently by the system.

CAPACITY – The amount of information that can be stored on a hard drive. Also known as storage capacity. (See also *formatted capacity*.)

CLEAN ROOM – An environmentally controlled dust-free assembly or repair facility in which hard disk drives are assembled or can be opened for internal servicing.

CLUSTER – In DOS environments, the smallest contiguous area that can be allocated for the storage of data. DOS clusters are usually 2 Kbytes (4 sectors) in size.

CONTROLLER – The chip or circuit that translates computer data and commands into a form suitable for use by the hard drive. Also known as disk controller.

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

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

CRC – Cyclic Redundancy Check. An error detection procedure that identifies incomplete or faulty data in each sector.

CYLINDER – When disks are placed directly above one another along the shaft, the circular, vertical “slice” consisting of all the tracks located in a particular position.

D

DATA SEPARATOR – The circuit that extracts data from timing information on drives that store a combined data and clock signal.

DEDICATED SERVO – A positioning mechanism using a dedicated surface of the disk that contains timing and positioning information only, as compared to surfaces that are also used for data. (See also *embedded servo*.)

DEFECT MANAGEMENT – A technique ensuring long-term data integrity. Consists of scanning disk drives both at the factory and during regular use, de-allocating defective sectors before purchase and compensating for new defective sectors afterward.

DISK – In general, any circular-shaped data-storage medium that stores data on the flat surface of the platter. The most common type of disk is the magnetic disk, which stores data as magnetic patterns in a metal or metal-oxide coating. Magnetic disks come in two forms: floppy and hard. Optical recording is a newer disk technology that gives higher capacity storage but at slower access times.

DISK CONTROLLER – A plug-in board, or embedded circuitry on the drive, that passes information to and from the disk. The Quantum hard disk drives all have controllers embedded on the drive printed-circuit board. (See also *controller*.)

DMA – Direct Memory Access. A process for transferring data directly to and from main memory, without passing through the CPU. DMA improves the speed and efficiency by allowing the system to continue processing even while new data is being retrieved.

DOS – Disk Operating System. The most common operating system used in IBM PCs. Manages all access to data on the disk.

DRIVE – Short form of *disk drive*.

DRIVE GEOMETRY – The functional dimensions of a drive, including the number of heads, cylinders, and sectors per track. (See also *logical format*.)

E

ECC – Error Correction Code. The incorporation of extra parity bits in transmitted data in order to detect errors that can be corrected by the controller.

EISA – Extended Industry Standard Architecture. An enhanced AT bus architecture designed by nine manufacturers of PC compatibles and announced in September 1988. EISA promises the advantages of IBM Micro Channel Architecture with the advantage of downward compatibility to the current standard AT bus. (See also *ISA*.)

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

ENCODING – The conversion of data into a pattern of On/Off or 1/0 signals prior to being written on the disk surface. (See also *RLL* and *MFM*.)

EPROM – Erasable Programmable Read-Only Memory. An integrated circuit memory chip that can store programs and data in a non-volatile state. These devices can be erased by ultraviolet light and reprogrammed with new data.

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

F

FAT – File Allocation Table. A data table stored on the outer edge of a disk, telling the operating system which sectors are allocated to each file and in what order.

FCI – Flux Changes per Inch. The number of magnetic field patterns that can be stored on a given area of disk surface, used as a measure of data density. (See also *BPI*.)

FILE SERVER – A computer that provides network stations with controlled access to shareable resources. The network operating system is loaded on the file server, and most shareable devices (disk subsystems, printers) are attached to it. The file server controls system security and monitors station-to-station communications. A dedicated file server can be used only as a file server while it is on the network. A non-dedicated file server can be used simultaneously as a file server and a workstation.

FIRMWARE – Permanent instructions and data programmed directly into the circuitry of read-only memory for controlling the operation of the computer. Distinct from software, which can be altered by programmers.

FLUX DENSITY – The number of magnetic field patterns that can be stored in a given length of disk

surface. The number is usually stated as flux changes per inch (FCI), with typical values in the thousands. (See also *FCI*.)

FLYING HEIGHT – The distance between the read/write head and the disk surface, made up of a cushion of air that keeps the two objects from touching. Smaller flying heights permit denser data storage but require more precise mechanical designs. Also known as fly height.

FORMAT – To write a magnetic track pattern onto a disk surface, specifying the locations of the tracks and sectors. This information must exist on a disk before it can store data.

FORMATTED CAPACITY – The amount of room left to store data on a disk after writing sector headers, boundary definitions, and timing information during a format operation. The size of a Quantum drive is always expressed in formatted capacity, accurately reflecting the usable space required.

FORM FACTOR – The industry standard that defines the physical, external dimensions of a particular device. For example, most Quantum hard disk drives use a 3 1/2-inch form factor.

G

GUIDE RAILS – Plastic strips attached to the sides of a hard disk drive in an IBM PC/AT or compatible computer so that the drive easily slides into place.

H

HALF-HEIGHT – Standard drive size equivalent to half the vertical space of a 5 1/4-inch drive.

HARD DISK – A type of storage medium that retains data as magnetic patterns on a rigid disk, usually made of an iron oxide or alloy over a magnesium or aluminum platter. Because hard disks spin more rapidly than floppy disks, and the head flies closer to the disk, hard disks can transfer data faster and store more in the same volume.

HARD ERROR – A data error that persists when the disk is re-read, usually caused by defects in the physical surface.

HARD-SECTORED – The most common method of indicating the start of each sector on a disk, based on information located in the embedded servo. This method is more precise than soft-sectored techniques and results in lower overhead. (See also *soft-sectored*.)

HEAD – The tiny electromagnetic coil and metal pole used to create and read back magnetic patterns on the disk. Also known as read/write head.

HEAD CRASH – Damage to the read/write head, usually caused by sudden contact with the disk surface. Head crash can also be caused by dust and other particles.

HIGH-CAPACITY DRIVE – By industry conventions typically a drive of 100 megabytes or more.

HIGH-LEVEL FORMATTING – Formatting performed by the operating system to create the root directory, file allocation tables and other basic configurations. (See also *low-level formatting*.)

HOME – Reference track used for recalibration of the actuator. Usually the outermost track (track 0).

HOST ADAPTER – A plug-in board that acts as the interface between a computer system bus and the disk drive.

I

IDE – Integrated Device Electronics. The term used for the interface on drives with embedded controllers that can be directly connected to the AT system bus. This connection can be through a 40-pin connector on the motherboard or an adapter board that contains only the drive-select decoding logic. Quantum AT drives use the IDE interface.

INITIALIZATION – See *low-level formatting*.

INTERFACE – A hardware or software protocol, (contained in the electronics of the disk controller and disk drive) that manages the exchange of data between the drive and computer. The most common interfaces for small computer systems are AT (also known as IDE) and SCSI.

INTERLEAVE – The arrangement of sectors on a track. The Interleave Factor is the number of sectors that pass beneath the read/write heads before the next sector arrives. For example, a 3:1 interleave factor means that the heads read a sector, then let two pass by before reading another, requiring three full revolutions of the disk to access the complete data track. Quantum drives have an interleave factor of 1:1, allowing the system to access a full track of data in a single revolution.

INTERLEAVE FACTOR – The number of sectors that pass beneath the read/write heads before the next numbered sector arrives. When the interleave factor is 3:1, a sector is read, two pass by, and then the next is read. It would take three revolutions of the disk to access a full track of data. Quantum drives have an interleave of 1:1, so a full track of data can be accessed within one revolution of the disk, thus offering the highest data throughput possible.

INTERNAL DRIVE – A drive mounted inside one of a computer's drive bays, or a hard disk on a card installed in one of the computer's expansion slots.

ISA – Industry Standard Architecture. The standard 16 bit AT bus as designed by IBM for their PC/AT system. This has been the only industry standard bus for PCs until the recent release of MCA (Micro Channel Architecture) and EISA (Extended Industry Standard Architecture). (See also *EISA*.)

J

JUMPER – A tiny box that slips over two pins on a circuit board, connecting the pins electrically. Some board manufacturers use Dual In-Line Package (DIP) switches instead of jumpers.

K

KILOBYTE (K) – A unit of measure consisting of 1,024 (2^{10}) bytes.

L

LANDING ZONE – A non-data area on the disk's inner cylinder where the heads can rest when the power is off.

LATENCY – The time during which the read/write heads wait for the data to rotate into position after the controller starts looking for a particular data track. If a disk rotates at 3,600 rpm, the maximum latency time is 16.4 milliseconds, and the average latency time is 8.2 milliseconds.

LOGICAL BLOCK - See Block.

LOGICAL FORMAT – In connection with standard disk formatting, refers to low-level formatting. In relation to DOS-specific format requirements, refers to the translations accomplished by the controller in situations where the hard drive data configurations do not match DOS format limitations.

LOOK AHEAD – The process of anticipating events in order to speed up computer operations. For example, the system can buffer data into cache RAM by reading blocks in advance, preparing the system for the next data request.

LOW-LEVEL FORMATTING – The process of creating sectors on the disk surface so that the operating system can access the required areas for generating the file structure. Quantum drives are shipped with the low-level formatting already completed. Also known as *initialization*.

LOW PROFILE SERIES (LPS)— Describes drives built to the 3 1/2-inch form factor, which are only 1 inch high. The standard form factor drives are 1.625 inches high.

M

MB — See *megabyte*.

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

MEGABYTE (MB) —Quantum defines megabyte as 10^6 bytes or 1,000,000 bytes.

MEGAHERTZ — A measurement of frequency in millions of cycles per second.

MHz — See *megahertz*.

MICROPROCESSOR — The integrated circuit chip that performs the bulk of data processing and controls the operation of all of the parts of the system. A disk drive also contains a microprocessor to handle all of the internal functions of the drive and to support the embedded controller.

MICROSECOND (μ s) — One millionth of a second (.000001 sec.).

MILLISECOND (ms) — One thousandth of a second (.001 sec.).

MTBF — Mean Time Between Failure. Reliability rating indicating the failure rate expected of a product expressed in power on hours (POH). Since manufacturers differ in the ways they determine the MTBF, comparisons of products should always take into account the MTBF calculation method.

MTTR — Mean Time To Repair. The average time it takes to repair a drive that has failed for some reason. This only takes into consideration the changing of the major sub-assemblies such as circuit board or sealed housing. Component level repair is not included in this number as this type of repair is not performed in the field.

O

OVERHEAD — *Command overhead* refers to the processing time required by the controller, host adapter, or drive prior to the execution of a command. Lower command overhead yields higher drive performance. (See also *zero command overhead*.) *Disk overhead* refers to the space required for non-data information such as location and timing. Disk overhead often accounts for about ten percent of drive capacity. Lower disk overhead yields greater disk capacity.

OVERWRITE — To write data on top of existing data, erasing it.

OXIDE — A metal-oxygen compound. Most magnetic coatings are combinations of iron or other metal oxides, and the term has become a general one for the magnetic coating on tape or disk.

P

PARTITION — A portion of a hard disk dedicated to a particular operating system and application and accessed as a single logical volume.

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

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

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

PLATED MEDIA — Disks that are covered with a hard metal alloy instead of an iron-oxide compound. Plated disks can store greater amounts of data than their oxide-coated counterparts.

PLATTER — Common term referring to the hard disk.

POH — Power On Hours. The unit of measurement for Mean Time Between Failure as expressed in the number of hours that power is applied to the device regardless of the amount of actual data transfer usage. (See also *MTBF*.)

POSITIONER — See *actuator*.

PROGRAMMED I/O – In a disk drive with an AT interface, data may be transferred between the drive and host using programmed I/O (PIO). In this case, the registers in the disk drive are accessed using host address lines A0-A2, chip select signals CS1FX- and CS3FX-, and read/write signals IOR- and IOW-. The host uses PIO to write to the Command Block Registers when transmitting commands to the drive, and to the Control Block Registers when transmitting control, such as a software reset.

R

RAM – Random Access Memory. An integrated circuit memory chip that allows information to be stored and retrieved by a microprocessor or controller. The information may be stored and retrieved in any order, and all storage locations are equally accessible.

RAM DISK – A “phantom” disk drive created by setting aside a section of RAM as if it were a group of regular sectors. Access to RAM disk data is extremely fast, but is lost when the system is reset or turned off.

READ AFTER WRITE – A mode of operation requiring that the system read each sector after data is written, checking that the data read back is the same as the data recorded. This operation lowers system speed but raises data reliability.

READ VERIFY – A data accuracy check performed by having the disk read data to the controller, which then checks for errors but does not pass the data on to the system.

READ/WRITE HEAD – The tiny electromagnetic coil and metal pole piece used to create and read back the magnetic patterns (write or read information) on the disk. Each side of each platter has its own read/write head.

REMOVABLE DISK – Generally said of disk drives where the disk itself is meant to be removed, and in particular of hard disks using disks mounted in cartridges. Their advantage is that multiple disks can be used to increase the amount of stored material, and that once removed, the disk can be stored away to prevent unauthorized use.

RLL – Run Length Limited. A method of encoding data into magnetic pulses. The RLL technique permits 50% more data per disk than the MFM method, but requires additional processing.

ROM – Read-Only Memory. Integrated circuit memory chip containing programs that can be accessed and read but can not be modified.

ROTARY ACTUATOR – The rotary actuator replaces the stepper motor used in the past by many hard disk manufacturers. The rotary actuator is perfectly balanced and rotates around a single pivot point. It allows closed-loop feedback positioning of the heads, which is more accurate than stepper motors.

ROTATIONAL LATENCY – The delay between when the controller starts looking for a specific block of data on a track and when that block rotates around to where it can be read by the read/write head. On average, it is half of the time needed for a full rotation (about 8 ms.).

S

SECTOR – On a PC hard drive, the minimum segment of track length that can be assigned to store information. On Macintosh and UNIX drives, sectors are usually grouped into blocks or logical blocks that function as the smallest data unit permitted. Since these blocks are often defined as a single sector the terms block and sector are sometimes used interchangeably in this context. (Note: The usage of the term block in connection with the physical configuration of the disk is different from its meaning at the system level. See also *block* and *cluster* for comparison.)

SEEK – A movement of the disk read/write head to a specific data track.

SERVO DATA – Magnetic markings written on the media that guide the read/write heads to the proper position.

SERVO SURFACE – A separate surface containing only positioning and disk timing information but no data.

SETTLE TIME – The interval between the arrival of the read/write head at a specific track, and the lessening of the residual movement to a level sufficient for reliable reading or writing.

SHOCK RATING – A rating, expressed in “G’s”, of how much shock a disk drive can sustain without damage.

SOFT ERROR – A faulty data reading that does not recur if the same data is reread from the disk, or corrected by ECC. Usually caused by power fluctuations or noise spikes.

SOFT-SECTORED – Old time-based method of indicating the start of each sector on a disk. Soft-sectored drives require that location instructions be located in the data fields. (See also *hard-sectored*.)

SPINDLE – The drive's center shaft, on which the hard disks are mounted. A synchronized spindle is a shaft that allows two disks to spin simultaneously as a mirror image of each other, permitting redundant storage of data.

SPUTTER – A special method of coating the disk that results in a hard, smooth surface capable of storing data at a high density. Quantum disk drives use sputtered thin film disks.

STEPPER – A type of motor that moves in discrete with each electrical pulse. Stepper were originally the most common type of actuator engine, since they can be geared to advance a read/write head one track per step. However, they are not as fast, reliable, or durable as the voice coil actuators found in Quantum disk drives. (See also *voice coil*.)

SUBSTRATE – The material underneath the magnetic coating of a disk. Common substrates include aluminum or magnesium alloys for hard drives, glass, for optical disks, and mylar for floppy disks.

SURFACE – The top or bottom side of a disk, which is coated with the magnetic material for recording data. On some drives one surface may be reserved for positioning information.

T

THIN FILM – A type of coating allowing very thin layers of magnetic material, used on hard disks and read/write heads. Hard disks with thin film surfaces can store greater amounts of data.

TPI – Tracks Per Inch. The number of tracks written within each inch of disk's surface, used as a measure of how closely the tracks are packed on a disk surface. Also known as *track density*.

TRACK – One of the many concentric magnetic circle patterns written on a disk surface as a guide for storing and reading data. Also known as *channel*.

TRACK DENSITY – How closely the tracks are packed on a disk surface. The number is specified as tracks per inch (TPI).

TRACK-TO-TRACK SEEK TIME – The time required for the read/write heads to move to an adjacent track.

TRANSFER RATE – The rate at which the disk sends and receives data from the controller. The sustained transfer rate includes the time required for system processing, head switches and seeks, and accurately reflects the drive's true performance. The burst mode transfer rate is a much higher figure that refers only to the movement of data directly into RAM.

U

UNFORMATTED CAPACITY – The total number of usable bytes on a disk, including the space that will be required to later to record location, boundary definitions, and timing information. (See *formatted capacity* for comparison.)

V

VOICE COIL – A fast and reliable actuator motor that works like a loud speaker, with the force of a magnetic coil causing a proportionate movement of the head. Voice coil actuators are more durable than their stepper counterparts, since fewer parts are subject to daily stress and wear. Voice coil technology is used in all Quantum drives.

W

WEDGE SERVO – The position on every track that contains data used by the closed loop positioning control. This information is used to fine tune the position of the read/write heads exactly over the track center.

WINCHESTER DISKS – Former code name for an early IBM hard disk model, sometimes still used to refer to hard drives in general.

WRITE ONCE – An optical disk technology that allows the drive to store and read back data, but prevents the drive from erasing information once it has been written.

X

XT – On the bus level, the original 8-bit version of what is now the AT bus.

Symbols

μController 5-5

Numerics

28-bit commands 5-8

A

abbreviations 1-1
acoustical characteristics 4-6
actuator lock 5-3
adapter board 3-12, 6-1
adaptive caching 5-16
ADC (8-channel 10-bit) 5-7
AGC 5-12
air filtration 5-4
AIRLOCK 2-2
AIRLOCK® 5-3
alternate status register 6-15
altitude 4-7
Analog to Digital converter (A/D) 5-8
Automatic Gain Control (AGC) 5-12
automatic read reallocation 5-22
automatic sector reallocation 5-22
Average rotational latency 2-1
Average seek time 2-1, 4-3

B

base casting 5-1
BIOS 3-13, 6-1
BIOS requirements 3-13
block diagram 5-6
Buffer Controller 5-8
buffer controller 5-7
Burst A - C 5-11
burst correction 5-20
burst ECC correction 5-21
bus interface connector 3-11
busy bit 6-20
Byte Interleaving 5-18

C

cable select (CS) jumper 3-5
cable select feature 6-1
Canadian Standards Association 2-2
check bytes 5-19
Check Power Mode (E5H) 6-29
clearance 3-7
clock recovery 5-11
CMOS setup 3-13
command block registers 6-17
command codes and parameters 6-21
command descriptions 6-21
Command Register 6-20
command register 6-20
commands, power management 6-27
Configuration 6-34
configuration command data field 6-37
control block registers 6-15
cooling fan 3-8
correctable double burst errors 5-20
correctable random burst errors 5-21
correctable read errors 4-8
corrected data bit 6-20
cross checking (XC) 5-18
cross-check bytes 5-8
cylinder contents 5-4
cylinder high register 6-18
cylinder low register 6-18

D

daisy-chain 2-2
daisy-chained 3-5
data port register 6-17
data request bit 6-20
Data Synchronizer 5-9, 5-10, 5-11
Data transfer operations 5-9
data transfer rates 4-2
Data-transfer rate 2-1
DC motor assembly 5-1
DCAIA 5-6
DCAIA buffer controller 5-14

- DCSIA block diagram 5-7
- Defect list data structure 6-33
- defect management 5-23
- defective sectors 2-2
- device control register 6-15
- diagnostics 6-24
- Digital to Analog converter (D/A) 5-9
- dimensions 4-6
- disable correction bit 5-23
- DisCache parameters 6-38
- DisCache® 2-2
- disk capacity 3-14
- disk errors 4-8
- disk partition 3-14
- disk stack assembly 5-3
- DOS 4-2
- double burst errors 5-20
- DP8491 Read Channel Device 5-9
- drive address register 6-16
- drive clearance 3-7
- drive electronics 5-5
- drive parameters 6-39
- drive ready bit 6-20
- drive seek complete bit 6-20
- drive select (DS) jumper 3-4
- drive write fault bit 6-20
- drive/head register 6-18
- dynamic offset 5-11

E

- ECC 5-18
- ECC Check Bytes 5-18
- ECC correctable error 5-22
- ECC error handling 5-21
- ECC interleaving 5-18
- Encoder-Decoder (ENDEC) 5-7
- ENDEC 5-8
- ENDEC module 5-11
- environmental specifications 4-7
- error bit 6-20
- error rates 4-8, 5-17
- error recovery parameters 6-38
- error register 6-17
- error reporting 6-39
- error types 5-23
- errors, single burst 5-19
- European Standards 2-2
- execute drive diagnostic 6-24

F

- faceplate 3-1
- FDISK 3-14
- Federal Communications Commission 2-2
- firmware 5-11
- firmware features 5-14
- flex circuit. 5-3
- floppy drive 3-12
- format track 6-23
- formatted capacity 4-2

G

- Gray code 5-12
- Gray coded number 5-12

H

- hard errors 5-23
- head switch time 4-3
- head switches 5-11
- head/disk assembly 5-1
- high level formatting 3-14
- host software reset 6-16
- humidity 4-7

I

- IDE 3-6
- IDE interface controller 5-7, 5-9
- identify drive 6-29
- Idle Mode (E1H) 6-28
- Idle Mode (E3H) 6-29
- illustrated parts breakdown 5-2
- index bit 6-20
- initialize drive parameters 6-25
- input power connections 3-10
- integrated drive electronics 6-1
- interface control 5-13
- interleave 5-18

J

- jumper configurations 3-3
- jumper options 3-4

L

landing zone 5-3
Logical Cylinders 3-14
Logical Heads 3-14
Logical Sectors/Track 3-14
Low Profile Series (LPS) 2-1

M

maximum screw torque 3-8
mean time to repair 4-8
media-defect mapping 2-2
Microcontroller Interface 5-8
microcontroller Interface 5-7
motherboard 3-11
Motor Controller 5-7, 5-9
mounting 3-6
mounting screw maximum torque 3-6
mounting screws 3-8
MTBF 4-8
multiple random errors 5-20
multiple zone recording 5-3
multiple-sector reads 6-22
multiple-sector writes 6-23

N

noise 4-4
non-repeatable errors 5-21
number of cylinders 5-4

P

PCB power 6-28
Peak Detect 5-11
Peak Detector And Servo Demodulator 5-9
phase-lock-loop 5-11
position information 5-12
positional information 5-12
positioning information 5-11
Power 6-27
power connector 3-9
power management commands 6-27
power modes 4-4
power on hours 4-8
power requirements 4-5
power sequencing 4-4
PreAmplifier 5-11
pre-amplifier 5-11

Pre-Compensator 5-9
pre-compensator 5-10
precompensator circuit 5-8
preventive maintenance 4-8
PRODUCT OVERVIEW 2-1
Programmable Filter 5-9
programmable filter 5-10
programmed I/O 6-13
Pulse Detector 5-9, 5-10
pulse detector module 5-11
PWM 5-7

Q

Quantum configuration key 6-37

R

R/W head matrix 5-11
RD GATE 5-9
read buffer 6-27
Read Channel 5-13
read configuration 6-34
read defect list 6-32
read long 6-22
read multiple 6-25
read preamplifier 5-3
read sectors 6-22
read verify sectors 6-23
read/write ASIC 5-11
read/write operations 5-13
reallocated data errors 4-8
Recalibrate 6-22
Reed-Solomon 5-8
Reed-Solomon error correcting code 2-2
reset limits 4-4
ripple 4-4
RLL 1,7 format 5-13
rotational latency 4-3
RUEE bit 5-22

S

Sector Address Mark (SAM) 5-12
Sector Count Register 6-28, 6-29
sector count register 6-18
sector data field 5-18
sector number register 6-18
sectors per track 5-4
seek 6-24

- seek time 4-3
- Sequencer 5-8
- sequencer 5-7
- servo burst and track information 5-12
- servo bursts 5-11
- Servo Controller 5-9
- servo controller 5-7
- SERVO SYSTEM 5-11
- servo wedge 5-11, 5-12
- set configuration 6-35
- set configuration without saving to disk 6-36
- set features 6-32
- set multiple mode 6-26
- settle mode 5-11
- SETUP program 3-13
- shipping container 3-2
- shock 4-7
- skewing, track and cylinder 5-4, 5-17
- slave present (SP) jumper 3-6
- Sleep Mode (E6H) 6-29
- soft errors 5-23
- specifications 4-1
- Standby Mode (E0H) 6-28
- status register 6-19
- Status Register Bits 6-20
- storage capacity 2-1
- supply voltages 4-4
- syndrome values 5-19
- syndromes 5-21
- sys tem startup 3-13

T

- temperature 4-7
- theory of operation 5-1
- Time Base Generator 5-9
- timebase generator 5-10
- timing diagram 6-10, 6-11, 6-13
- timing specifications 4-3
- track and cylinder skewing 5-4, 5-17
- track following mode 5-11
- transferred errors 4-8
- transient voltages 4-4
- TTL logic 6-2

U

- uncorrectable double burst errors 5-20
- uncorrectable read errors 4-8
- Underwriters Laboratory 2-2

- UNIX 4-2
- UNPACKING 3-2
- unrecoverable errors 5-23

V

- velocity 5-12
- velocity mode 5-11
- ventilation requirements 3-6
- vibration 4-7
- voice-coil actuator 5-3

W

- WCS 5-8
- write buffer 6-27
- Write Channel 5-13
- Write Driver 5-11
- write driver 5-3, 5-11
- write long 6-23
- write multiple 6-26
- write sector 6-23
- WriteCache® 2-2



QUICK FAX - READER COMMENT FORM

PLEASE FAX YOUR COMMENTS AND SUGGESTIONS FOR IMPROVING THIS PUBLICATION TO:
QUANTUM TECHNICAL PUBLICATIONS (408) 894-6159

Manual Part Number 81-106898-01	Manual Title ProDrive LPS 210AT Product Manual	Current Date
How did you use this publication? <input type="checkbox"/> Reference <input type="checkbox"/> Sales <input type="checkbox"/> Installation <input type="checkbox"/> Maintenance	What is your occupation?	Do you wish a written reply? <input type="checkbox"/> Yes <input type="checkbox"/> No
What is your overall rating of this publication? <input type="checkbox"/> Very Good <input type="checkbox"/> Fair <input type="checkbox"/> Good <input type="checkbox"/> Poor	Is the material presented effectively? <input type="checkbox"/> Well organized <input type="checkbox"/> Fully covered <input type="checkbox"/> Clear <input type="checkbox"/> Correct <input type="checkbox"/> Well Illustrated Locating the information you were looking for was? <input type="checkbox"/> Easy <input type="checkbox"/> Difficult	

Please enter your comments and suggestions for improvement below. If you were in any way dissatisfied with this publication, we would like to know why. Be specific, if possible; give page, column, and paragraph number references where applicable.

<p>.....</p> <p>.....</p> <p>.....</p> <p>.....</p> <p>.....</p> <p>.....</p> <p>.....</p> <p>.....</p> <p>.....</p> <p>.....</p>	<p>Your name and FAX NUMBER:</p> <p>.....</p> <p>.....</p> <p>.....</p>
---	---

All comments and suggestions become the property of Quantum Corporation.

Thank you for your interest. We will consider all comments and will incorporate changes when appropriate.

