Seagate

Medalist 1080sl

SCSI Interface Drive

Product Manual

Medalist 1080sl

SCSI Interface Drive

Product Manual

MEDALIST SL

Publication Number: 36321-101, Rev. A, August 1995

Contents

Figures

Introduction

This manual describes the functional, mechanical and interface specifications for the Medalist 1080sl SCSI hard disc drive. The Medalist 1080sl SCSI is referred to throughout this manual by its model number, ST51080N.

The ST51080N is a high-capacity, high-performance, energy-efficient SCSI drive that comes in the mini 3.5-inch form-factor.

The drive uses a high-performance SCSI-2 interface that supports an asynchronous external transfer rate of up to 5 Mbytes per second and a synchronous external transfer rate of up to 10 Mbytes per second.

The interface is supported with a 128-Kbyte segmented cache and embedded servo technology. The segmented cache aids the flow of read and write data. The embedded servo allows for accurate positioning of the heads over the data and eliminates periodic thermal recalibration to assure data transfer without interruption.

The ST51080N conforms to the standard 3.5-inch footprint but have a 0.75-inch (19 mm) height profile and a 5.0-inch depth profile. The lower height and shorter depth gives the designer or integrator more room for air circulation, other peripherals or a smaller drive bay.

The following is a summary of the drive's features:

Capacity

- 1.08 Gbytes formatted

Performance

- Uses the SCSI-2 interface
- 5,376-RPM rotational speed
- 128-Kbyte segmented buffer
- 12.5-msec average seek time

Acoustics

- 30-dBA idle acoustic sound level

Mini 3.5-inch form-factor

- 19-mm height profile
- Fits standard 3.5-inch footprint

Quick specification chart

The following table serves as a quick reference for the ST51080N performance specifications. These and other specifications are discussed in the Specification summary section following the table.

Drive specification	**ST51080N**
Formatted capacity (Mbytes) ($\times 10^6$ bytes)	1,080.23
Total sectors	2,109,840
Bytes per sector	512
Sectors per track (average)	115
Physical cylinders	4,826
Physical read/write heads	4
Physical disc	2
Recording density (bits per inch, max)	73,344
Track density (tracks per inch)	4,923
Spindle speed (RPM)	5,376 ± 0.5%
Internal data-transfer rate (Mbits per second max)	33 MHz to 65 MHz
External transfer rate (Mbytes per second max)	5.0 asynchronous 10.0 synchronous
Cache buffer (Kbytes)	128
Height, inches max (mm)	0.748 (19.0)
Width, inches max (mm)	4.01 (102.8)
Depth, inches max (mm)	5.00 (127.0)
Typical weight, lb (g)	0.750 (340.2)
Track-to-track seek time read (msec typical)	3.5
Track-to-track seek time write (msec typical)	4.5
Average seek time read (msec typical)	12.5
Average seek time write (msec typical)	15.5
Full-stroke seek time read (msec typical)	25.0
Full-stroke seek time write (msec typical)	27
Average latency (msec)	5.6
Power-on to ready (sec typical)	20

continued

continued from previous page

Drive specification	**ST51080N**
Spinup current: +12V (max)	1.32A
Seek power (typical)	7.11W
Read/Write power (typical)	5.14W
Idle power (typical)	4.9W
Voltage tolerance (including noise): +5V	±5%
Voltage tolerance (including noise): +12V	±5%
Ambient temperature, operating (°C)	5° to 55°C
Temperature gradient , operating (°C per hour max)	20°C
Relative humidity, operating gradient (max)	10% per hr
Relative humidity, operating	8% to 80%
Wet bulb temperature, operating (noncondensing)	29.4°C
Altitude, operating	–1,000 to 10,000 ft.
Shock, normal operating (Gs max for 11 msec)	2 Gs
Vibration (Gs max at 22–350 Hz without nonrecoverable errors), operating	0.75 Gs 0 to Peak
Vibration (Gs max at 22–350 Hz with no physical damage incurred), nonoperating	4 Gs 0 to Peak
Drive acoustics, Idle mode (dBA)	30 dBA
Drive acoustics, seeking (dBA)	34 dBA
Nonrecoverable read errors (per bits transferred)	10^{13}
Mean time between failures (power-on hours)	300,000
Contact start-stop cycles	40,000
Service life (years)	5

1.0 Specifications summary

1.1 Formatted capacity

The capacities specified here do not include spare sectors and cylinders. The media contains one spare sector per track and two spare cylinders per drive.

Formatted capacity (Mbytes[1])	1080.23
Total sectors	2,109,840

1.2 Physical geometry

Discs	2
Read/write heads	4
Cylinders	4,826

1.3 Functional specifications

Interface	Fast SCSI-2
Zone Bit Recording method	RLL (1,7)
External data transfer rate (Mbytes per sec, max)	5.0 asynchronous 10.0 synchronous
Internal data transfer rate (Mbits per sec)	33 to 65
Spindle speed (RPM)	5,376 ± 0.5%
Bytes per sector	512
Track density (TPI)	4,923
Recording density (BPI, max)	73,344

1. One Mbyte equals 1,000,000 bytes.

1.4 Physical dimensions

Height (max)	0.748 inches (19 mm)
Width (max)	4.00 inches (102.1 mm)
Depth (max)	5.00 inches (127.0 mm)
Weight (max)	0.75 lb (340.2 g)

1.5 Seek time

All seek time measurements are taken under nominal conditions of temperature and voltage with the drive mounted horizontally. In the following table:

- *Track-to-track* seek time is the average of all possible single-track seeks in both directions.
- *Average/typical* seek time is a true statistical random average of at least 5,000 measurements of seeks in both directions between random cylinders, less overhead.
- *Full-stroke* seek time is one-half the time needed to seek from logical block address zero (LBA 0) to the maximum LBA and back to LBA 0.

Track-to-track seek time typ[2]	Average/typical seek time[3]	Full-stroke seek time typ[4]	Average latency
4.5 msec read	12.5 msec read	25.0 msec read	5.6 msec
3.5 msec write	15.5 msec write	27.0 msec write	

Note. Host overhead varies between systems and cannot be specified. Drive internal overhead is measured by issuing a no-motion seek. Drive overhead is typically less than 1.0 msec.

1.5.1 Read look-ahead and caching

The drive uses algorithms that improve seek performance by storing data in a buffer and processing it at a more convenient time. Three methods are used: read look-ahead, read caching and write caching. These are described in Appendix C.6.

---

2. All possible one track seeks are divided into the time required to perform those seeks. Only the mechanism time is used; interface overhead is excluded.

3. All possible seeks are divided into the time required to perform those seeks. Only the mechanism time is used; interface overhead is excluded.

4. The average of 1000 full stroke seeks is used in this computation. Only the mechanism time is used; interface overhead is excluded.

1.6 Start and stop time

If the motor start option is disabled, the drive becomes ready within 20 seconds after power is applied. If the motor start option is enabled, the drive becomes ready within 20 seconds after it receives the Motor Start command. If the drive receives a command to spin down or power is removed, the drive stops within 15 seconds.

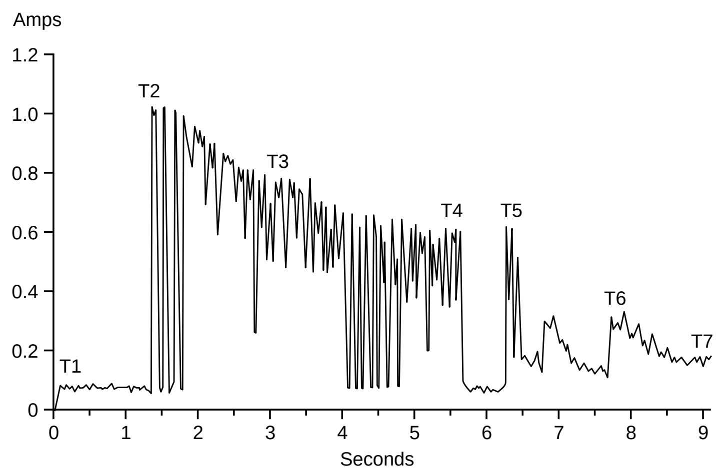

Figure 1. Typical startup current profile

1.6.1 Power-up sequence

The following typical power-up sequence is provided to assist in evaluating drive performance. This information does not constitute a specification or a performance guarantee.

1. Power is applied to the disc drive.

2. The LED comes on for about 1 second.

3. Depending on whether there is a jumper installed on pins 15 and 16 of the options jumper block (J8) shown in Figure 3 on page 17, either of the following sequences occurs:

- **a.** If a jumper is not installed, the remote start option is not enabled, and the drive begins to spin up as soon as power is applied.
- **b.** If a jumper is installed, the remote start option is enabled, and the drive begins to spin up when the host commands the motor to start.

4. Within 250 msec after power is applied, the drive responds to the Test Unit Ready, Request Sense, Mode Sense and Inquiry commands.

5. The drive begins to lock in speed-control circuits.
6. The actuator lock releases the actuator.
7. The spindle motor reaches operating speed in about 5 seconds. After 5 seconds, there are no speed variations.
8. The drive performs velocity adjustment seeks.
9. The drive seeks track 0 and becomes ready.

1.6.2 Power-down sequence

The following typical power-down sequence is provided to assist in evaluating drive performance. This information does not constitute a specification or a performance guarantee.

1. The power cable is unplugged from the drive, or the drive receives a command to spin down.
2. Within 3 seconds after the motor begins to spin down, the actuator lock engages, producing a sound.
3. The spindle stops within 15 seconds, whether the power cable is unplugged from the drive or the drive receives the power-down command.

1.6.3 Auto-park

Upon power-down, the read/write heads automatically move to the landing zone. The heads park beyond the maximum data cylinder. When power is applied, the heads recalibrate to track 0.

Caution. Do not move the drive until the spindle motor has come to a complete stop; otherwise, you may damage the drive.

1.7 Power management

The drive supports power-management modes that reduce its overall power consumption. The drive automatically changes from one mode to another in response to interface activity. You do not need to change any parameters or send any special commands to make the drive change modes. The power-management modes are described below.

- **Spinup.** Spinup is defined as the period during which the spindle is coming up to operating speed. The power consumed in this mode is equivalent to the average power during the first 10 seconds after the drive begins to spin up.

- **Seeking.** The servo electronics are active, and the heads are moving to a specific location on the disc. The read/write electronics are powered-down. The power consumed in this mode is equivalent to the average power measured while executing random seeks with a 2-revolution (26.6 msec) dwell between seeks. The drive enters this mode from the Idle mode.
- **Read/Write.** The drive is reading or writing. All electronics are active and the heads are on track.
- **Idle.** The motor is up to speed and the drive is in track follow mode.

1.7.1 Power consumption

Values in the table below were measured at the drive power connector with an RMS DC ammeter. The terminating resistors are disabled, and terminator power is supplied through the SCSI connector. All values are measured 10 minutes after the drive spins up except as noted.

	During Spinup	Seeking	Read/ Write	Idle
Current at +12V				
Amps peak	1.32	—	—	—
RMS amps typ	—	0.393	0.221	0.219
Watts typ	—	4.72	2.65	2.63
Current at +5V				
RMS amps typ	—	0.477	0.498	0.471
Watts typ	—	2.39	2.49	2.36
Power				
Total watts typ	—	7.11	5.14	4.99

1.7.2 Voltage tolerance

	+5V	+12V
Voltage tolerance (including noise)	± 5%	± 5%

1.7.3 Input noise

	+5V	**+12V**
Voltage tolerance (including noise)	± 5%	± 5%
Input noise frequency (max)	25 MHz	25 MHz
Input noise (max, peak-to-peak)	100 mV	240 mV

1.8 Environmental

This section specifies acceptable environmental conditions for the drive. The operating specifications assume that the drive is powered up. The nonoperating specifications assume that the drive is packaged as it was shipped from the factory.

1.8.1 Ambient temperature

Operating	5°C to 55°C (41°F to 131°F)
Nonoperating	–40°C to 70°C (–40°F to 158°F)

1.8.2 Temperature gradient

Operating	20°C per hour (36°F per hour)
Nonoperating	30°C per hour (54°F per hour)

1.8.3 Altitude

Operating	–1,000 ft to 10,000 ft (–305 m to 3,048 m)
Nonoperating	–1,000 ft to 40,000 ft (–305 m to 12,192 m)

1.8.4 Relative humidity

Operating	8% to 80% noncondensing Maximum wet bulb 29.4℃ (84.9℉)
Operating gradient, max	10% per hour
Nonoperating	5% to 95% noncondensing Maximum wet bulb 35℃ (95.0℉)

1.9 Shock and vibration

All shock and vibration specifications assume that the inputs are measured at the drive mounting screws. Shock measurements are based on an 11-msec, half sine wave shock pulse, not to be repeated more than twice per second.

During normal operating shock and vibration, there is no physical damage to the drive or performance degradation.

During abnormal operating shock and vibration, there is no physical damage to the drive, although performance may be degraded during the shock or vibration episode. When normal operating shock levels resume, the drive meets its performance specifications.

During nonoperating shock and vibration, the read/write heads are positioned in the shipping zone.

	Normal operating	**Abnormal operating**	**Nonoperating**
Shock	2 Gs	10 Gs	75 Gs
5–22 Hz vibration	0.020-inch displacement	0.030-inch displacement	0.160-inch displacement
22–350 Hz vibration	0.50 Gs	0.75 Gs	4.00 Gs

1.10 Acoustics

Sound pressure is measured at idle from 1 meter above the drive top cover.

	Idle	**Seek**
Sound pressure, typ	30 dBA	34 dBA
Sound pressure, max	34 dBA	38 dBA

1.11 Reliability

Read error rates are measured with automatic retries and data correction with ECC enabled and all flaws reallocated. MTBF is measured at nominal power at sea level and 40°C ambient temperature.

Nonrecoverable read errors	1 per 10^{13} bits transferred
Seek errors	1 per 10^{7} physical seeks
MTBF	300,000 power-on hours
Service life	5 years

1.12 Agency listings

This drive is listed by agencies as follows:

- Recognized in accordance with UL 478 and UL 1950
- Certified to CSA C22.2 No. 220-M1986 and CSA C22.2 No. 950-M1989
- Certified to VDE 0806/05.90 and EN 60950/1.88 as tested by VDE

1.13 FCC verification

The ST51080N drive is intended to be contained solely within a personal computer or similar enclosure (not attached to an external device). As such, a drive is considered to be a subassembly even when individually marketed to the customer. As a subassembly, no Federal Communications Commission authorization, verification or certification of the device is required.

Seagate Technology, Inc. has tested these drives in an enclosure as described above to ensure that the total assembly (enclosure, disc drive, motherboard, power supply, etc.) does comply with the limits for a Class B computing device, pursuant to Subpart J of Part 15 of the FCC rules. Operation with noncertified assemblies is likely to result in interference to radio and television reception.

Radio and television interference. This equipment generates and uses radio frequency energy and, if not installed and used in strict accordance with the manufacturer's instructions, may cause interference to radio and television reception.

This equipment is designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause interference to radio or television, which can be determined by turning the equipment on and off, you are encouraged to try one or more of the following corrective measures:

- Reorient the receiving antenna.
- Move the device to one side or the other of the radio or TV.
- Move the device farther away from the radio or TV.
- Plug the equipment into a different outlet so that the receiver and computer are on different branch outlets.

If necessary, you should consult your dealer or an experienced radio/television technician for additional suggestions. You may find helpful the following booklet prepared by the Federal Communications Commission: *How to Identify and Resolve Radio-Television Interference Problems.* This booklet is available from the Superintendent of Documents, US Government Printing Office, Washington, DC 20402. Refer to publication number 004-000-00345-4.

Note. This digital apparatus does not exceed the Class B limits for radio noise emissions from computer equipment as set out in the radio interference regulations of the Canadian Department of communications.

Le présent appareil numérique n′émet pas de bruits radioélectriques dépassant les limites applicables aux appareils numériques de Classe B prescrites dans le règlement sur le brouillage radioélectrique édicté par le Ministère des Communications du Canada.

Sicherheitsanleitung

1. Das Gerrät ist ein Einbaugerät, das für eine maximale Umgebungstemperatur von 55℃ vorgesehen ist.

2. Zur Befestigung des Laufwerks werden 4 Schrauben 6-32 UNC-2A benötigt. Bei seitlicher Befestigung darf die maximale Länge der Schrauben im Chassis nicht mehr als 5,08 mm und bei Befestigung an der Unterseite nicht mehr als 5,08 mm betragen.

3. Als Versorgungsspannugen werden benötigt:
+5V ± 5% 0,65A
+12V ± 5% 0,45A (1,9A fur ca. 10 Sek. fur ± 10%)

4. Die Versorgungsspannung muss SELV entsprechen.

5. Alle Arbeiten auf dem Festplatte dürfen nur von Ausgebildetem Servicepersonal durchgeführt werden. Bitte schaffen Sie Festplatteetiketten nicht weg.
6. Der Einbaudes Drives muss den Anforderungen gemäss DIN IEC 950V DC 0805/05.90 entsprechen.

2.0 Hardware and interface

The ST51080N drive uses a SCSI-2 interface that consists of a 9-bit bidirectional bus (8 data bits and 1 parity bit) and 9 control signals. The interface supports multiple initiators, disconnect and reconnect, self-configuring host software and logical block addressing.

The interface employs a singled-ended driver/receiver configuration that uses asynchronous or synchronous communication protocols. It supports asynchronous transfer rates of up to 5 Mbytes per second and synchronous transfer rates of up to 10.0 Mbytes per second. The bus protocol supports multiple initiators, disconnect and reconnect, additional messages and 6-byte and 10-byte command descriptor blocks. The bus cable can be up to 6 meters long for standard mode and up to 3 meters long for Fast SCSI mode.

2.1 SCSI-2 compatibility

The drive interface is described in the *Seagate SCSI-2 Interface Manual*, publication number 77738479. The drive complies with the mandatory subset of the ANSI SCSI-2 Interface. The Fast SCSI-2 interface is based on the ANSI Small Computer System Interface-2 (SCSI-2), document number ANSI X3.131-199*x* (X3T9.2/86-109 Rev. 10h).

2.2 Handling and static-discharge precautions

The ST51080N drive uses static-sensitive devices. Avoid damaging the drive and these devices by observing the following standard handling and static-discharge precautions:

- Keep the drive in its static-shielded bag until you are ready to complete the installation. Do not attach any cables to the drive while it is in its static-shielded bag.
- Before handling the drive, put on a grounded wrist strap, or ground yourself frequently by touching the metal chassis of a computer that is plugged into a grounded outlet. Wear a grounded wrist strap throughout the entire installation procedure.

  Wool and synthetic clothes, carpets, plastics and Styrofoam are contributors to electrostatic build-up. Static discharge can damage sensitive components in your drive and computer.
- Handle the drive by its edges or frame only.
- The drive is extremely fragile—handle it with care. Do not press down on the drive top cover.

- Always rest the drive on a padded, antistatic surface until you mount it in the host system.
- Do not touch the connector pins or the printed circuit board.
- Do not remove the factory-installed labels from the drive or cover them with additional labels. If you do, you void the warranty. Some factory-installed labels contain information needed to service the drive. Others are used to seal out dirt and contamination.

2.3 Electrical interface

The ST51080N drive is designed to use singled-ended interface signals. They employ singled-ended drivers and receivers and active terminator circuitry. Figure 2 shows a single-ended transmitter and receiver without the active terminator circuitry.

- **Transmitter characteristics.** The drive uses an ANSI SCSI-compatible, open-collector, single-ended driver. This driver is capable of sinking a current of 48 mA with a low-level output voltage of 0.4 volts.
- **Receiver characteristics.** The drive uses an ANSI SCSI single-ended receiver with hysteresis gate or equivalent as a line receiver.

The loss in the cable is defined as the difference between the voltages of the input and output signals, as shown below:

Logic level	Driver output (x)	Receiver input (x)
Asserted (1)	$0.0V \leq x \leq 0.4V$	$0.0V \leq x \leq 0.8V$
Negated (0)	$2.5V \leq x \leq 5.25V$	$2.0V \leq x \leq 5.25V$

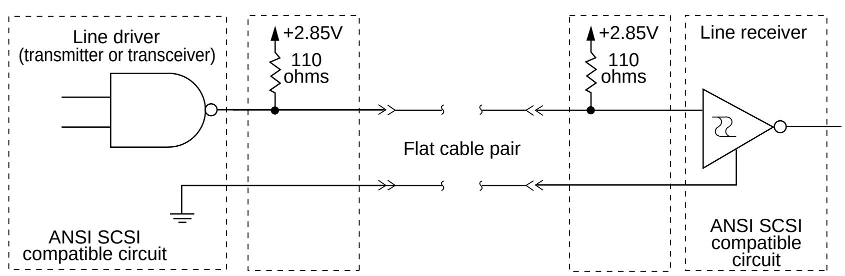

Figure 2. Single-ended transmitter and receiver

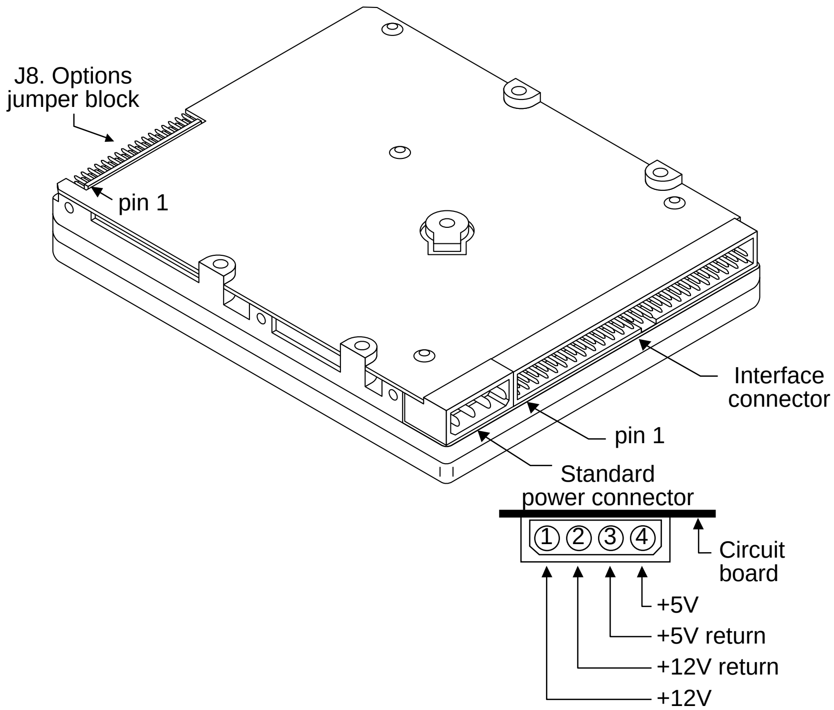

Figure 3. ST51080N connectors

2.4 SCSI interface connector

The ST51080N uses a standard 50-pin, nonshielded, keyed connector. The connector consists of two rows of 25 male contacts 0.100 inches apart. The location of pin 1 is shown in Figure 3. Recommended mating connectors are listed below with their part numbers.

Part numbers for mating 3M connectors compatible with the drive are listed below. These connectors do not have a center key and are available with or without a strain relief.

	Without strain relief Without center key	**With strain relief No center key**
Closed end (for cable ends)	3M 3425-7000	3M 3425-7050
Open end (for daisy chain)	3M 3425-6000	3M 3425-6050

Part numbers for mating Molex connectors compatible with the drive are listed below. These connectors have a center key.

Closed end (for cable ends)	Molex 39-51-2504
Open end (for daisy chain)	Molex 39-51-2501

Below are part numbers for strain reliefs that can be used with the Molex connectors.

Molex strain relief, preferred version in Europe	Molex 90170-0050
Molex strain relief, preferred version in Japan	Molex 15-25-1503

2.4.1 SCSI interface connector pin assignments

The table below shows the pin assignment for the 50-pin interface connector. A minus sign (–) indicates an active-low signal.

Signal name	Signal pin number	Ground pin number
DB(0)–	2	1
DB(1)–	4	3
DB(2)–	6	5
DB(3)–	8	7
DB(4)–	10	9
DB(5)–	12	11
DB(6)–	14	13
DB(7)–	16	15
DB(P)–	18	17
Ground	19–22	—
Reserved	23–25	—
Terminator power	26	—
Reserved	27–28	—
Ground	29–30	—

Signal name	Signal pin number	Ground pin number
ATN–	32	31
Ground	33–34	—
BSY–	36	35
ACK–	38	37
RST–	40	39
MSG–	42	41
SEL–	44	43
C/D–	46	45
REQ–	48	47
I/O–	50	49

Caution. Do not connect pin 25 to ground. If you plug in the connector upside down, the terminator power on pin 26 is shorted to ground. This may damage the drive.

2.5 Interface cable requirements

A characteristic impedance of 100 ohms +10% is recommended for the unshielded flat or twisted-pair interface cable. However, most available cables have a somewhat lower characteristic impedance. To minimize discontinuities and signal reflections, do not use cables of different impedances in the same bus. Implementation may require tradeoffs in wielding effectiveness, cable length, the number of loads and the transfer rates to achieve satisfactory system operation. If shielded and unshielded cables are mixed within the same bus, the effect of impedance mismatch must be carefully considered. This is especially important for maintaining adequate margin for Fast SCSI transfer rates.

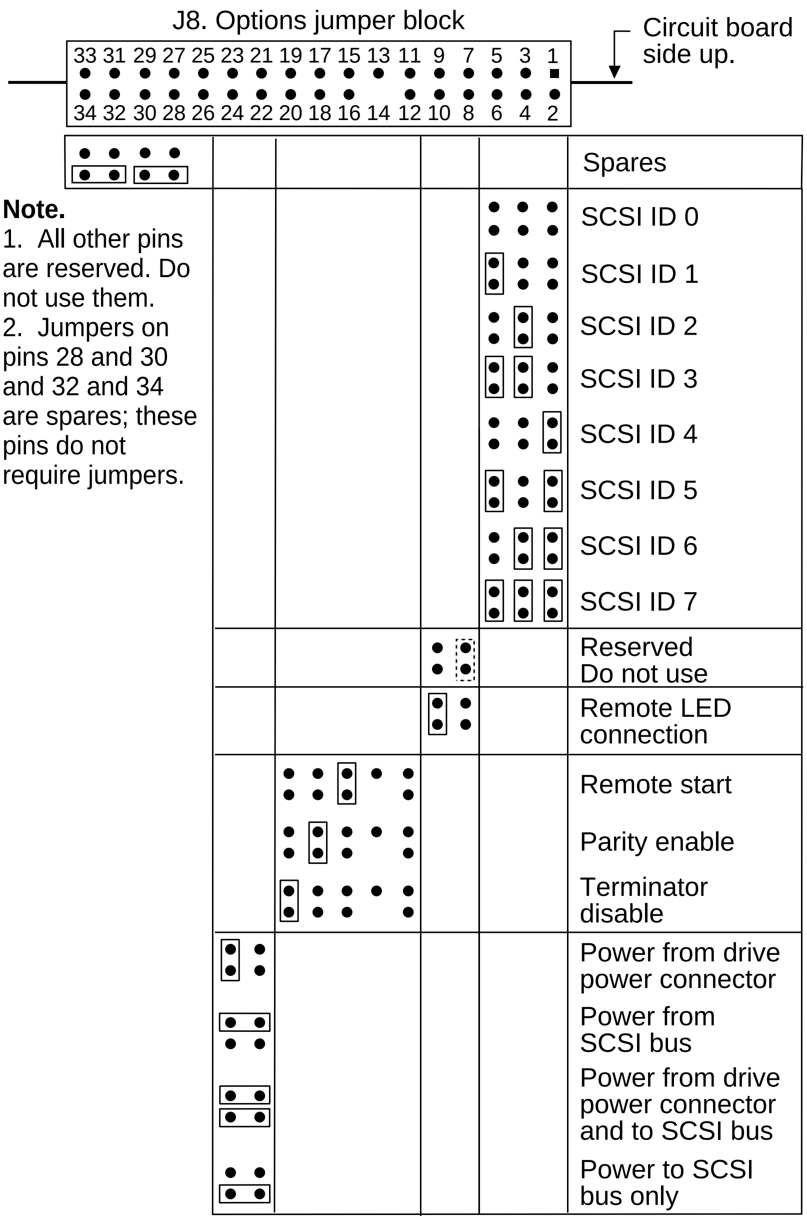

Note.
1. All other pins are reserved. Do not use them.
2. Jumpers on pins 28 and 30 and 32 and 34 are spares; these pins do not require jumpers.

Figure 4. Options jumper block settings

To minimize noise, use 28 AWG or larger 50-conductor flat cable or 25-conductor twisted-pair ribbon cable. The following are part numbers for nonshielded flat cables we recommend:

Part	**Manufacturer**
Flat Cable	3M-3365-50
Twisted Pair	Spectra Twist-N-Flat 455-248-50

2.5.1 Interface cable length for asynchronous operation

The SCSI interface cable must meet the following requirements for normal operation:

- The cable cannot be longer than 6.0 meters.
- Cable stubs cannot be more than 0.1 meter long and must be separated by at least 0.3 meter.

2.5.2 Interface cable for Fast SCSI operation

When using fast synchronous data transfer rates, the SCSI interface cable must meet the following additional requirements:

- The cable cannot be longer than 3.0 meters.
- The cable should not attenuate a 5 MHz signal more than 0.095 dB per meter.
- The DC resistance at 20℃ must not exceed 0.230 ohms per meter.
- A shielded, twisted-pair cable should not have a propagation delay delta greater than 20 nsec per meter.

2.6 Options jumper block

The ST51080N options jumper block (see Figure 4) allows you to manually:

- Enable or disable active termination
- Set the SCSI ID address
- Select the terminator power source
- Enable parity
- Activate the motor start/stop option
- Attach a remote LED

The jumper block accepts 2-mm jumpers. Spare jumpers are included with the drive attached to the pins shown in Figure 4 on page 20. If you need additional jumpers, use the jumpers listed below or equivalent.

Manufacturer	Part number
Seagate	13211-001
Du Pont	89133-001
Methode	8618-202-70

2.6.1 SCSI address

The SCSI ID address is set using pins 1 and 2, 3 and 4, and 5 and 6 on the options jumper block. The drive is shipped configured with jumpers on pins 3 and 4 and pins 5 and 6. This makes the default SCSI ID 3. To configure the drive for a different address, consult the chart in Figure 4. Refer to your host adapter reference manual for its preferred addressing scheme.

2.6.2 Active Termination

The ST51080N uses the options jumper block to configure the active termination. Active termination is enabled when no jumper is connected to pins 19 and 20 of the options jumper block. To disable active termination, place a jumper on pins 19 and 20 of the options jumper block. You can order the drive with the active termination enabled or disabled.

Note. If the active terminators are disabled, termination power must still be supplied using one of the methods described in Section 2.6.3.

2.6.3 Terminator power source selection

To select the termination power source, install jumpers as follows:

- To select the drive power connector as the termination power source for the terminator packs, install a jumper on pins 23 and 24 of the options jumper block.
- To select the SCSI bus as the termination power source for the terminator packs, install a jumper on pins 21 and 23 of the options jumper block.
- To provide terminator power to the SCSI bus and the drive terminator packs, install jumpers on pins 21 and 23 and pins 22 and 24 of the options jumper block. This is the default.
- To provide terminator power to the SCSI bus from the drive power

connector only, install a jumper on pins 22 and 24 of the options jumper block.

2.6.4 Parity enable option

Parity is enabled when a jumper is installed on pins 17 and 18 of the options jumper block. ST51080N drives are shipped with parity enabled.

2.6.5 Start/stop option

When a jumper is installed on pins 15 and 16 of the options jumper block, the drive waits for a Start/Stop Unit command from the host before starting or stopping the spindle motor.

2.6.6 Remote LED connection

Pins 9 and 10, located on the options jumper block, are reserved for a remote LED. Pin 9 is ground. The options jumper block accepts 2-mm connectors. You may need to replace the current LED cable-connector with a 2-mm connector. If you are placing the drive in an array configuration, we recommend the LiteOn (part number LTL-3231A) LED or equivalent.

2.7 Daisy chaining

The drive can be connected in a daisy chain with a maximum of eight SCSI devices (including the host) that have single-ended drivers and receivers. Each SCSI device must be set to a unique SCSI ID number. SCSI ID 7 is usually used for the host adapter.

All signals are common between all SCSI devices. The SCSI devices at both ends of the daisy chain must be terminated; the intermediate SCSI devices should not be terminated.

2.8 Hot-plugging

You can connect and disconnect the I/O and power cables for each SCSI device in a daisy chain without powering down the system. This is called hot-plugging. When hot-plugging, the following conditions must be met:

- The terminators at either end of the SCSI bus are in place.
- The drive you are disconnecting or connecting is not the device supplying terminator power or terminating resistance to the bus.

- All I/O transactions are complete before you install or remove a drive.

To avoid damage to the disc and head, the spindle must be completely stopped and the heads must be parked before you remove the drive from the system. There are two ways to stop the spindle and park the heads:

- If the drive is not configured to use the remote start/stop feature, disconnect the DC power cable from the drive DC power connector and wait 30 seconds.
- If the drive is configured to use the remote start/stop feature, issue the Start/Stop Unit command and wait 30 seconds.

2.9 Mounting the drive

The ST51080N is a 3.5-inch form-factor drive with a .75-inch height profile and a 5.0-inch depth profile. You can mount them securely in the computer using either the bottom or side mounting holes, as described below. Position the drive so that you do not strain or crimp the cables. Figure 5 shows the drive's mounting dimensions, including the side and bottom mounting holes.

Bottom mounting holes. Insert 6-32 UNC-2A mounting screws in the four available bottom mounting holes. Do not insert the screws more than 0.20 inches (6 turns) into the drive frame.

Side mounting holes. Insert 6-32 UNC-2A mounting screws in four of the six available side mounting holes. Use two mounting holes on each side of the drive. Do not insert the screws more than 0.20 inches (6 turns) into the drive frame.

Caution. To avoid damaging the drive:

- Use only mounting screws of the type specified.
- Gently tighten the mounting screws—do not apply more than 6 inch-lb of torque.

All dimensions are shown in both inches and millimeters. Millimeters are shown in parentheses.

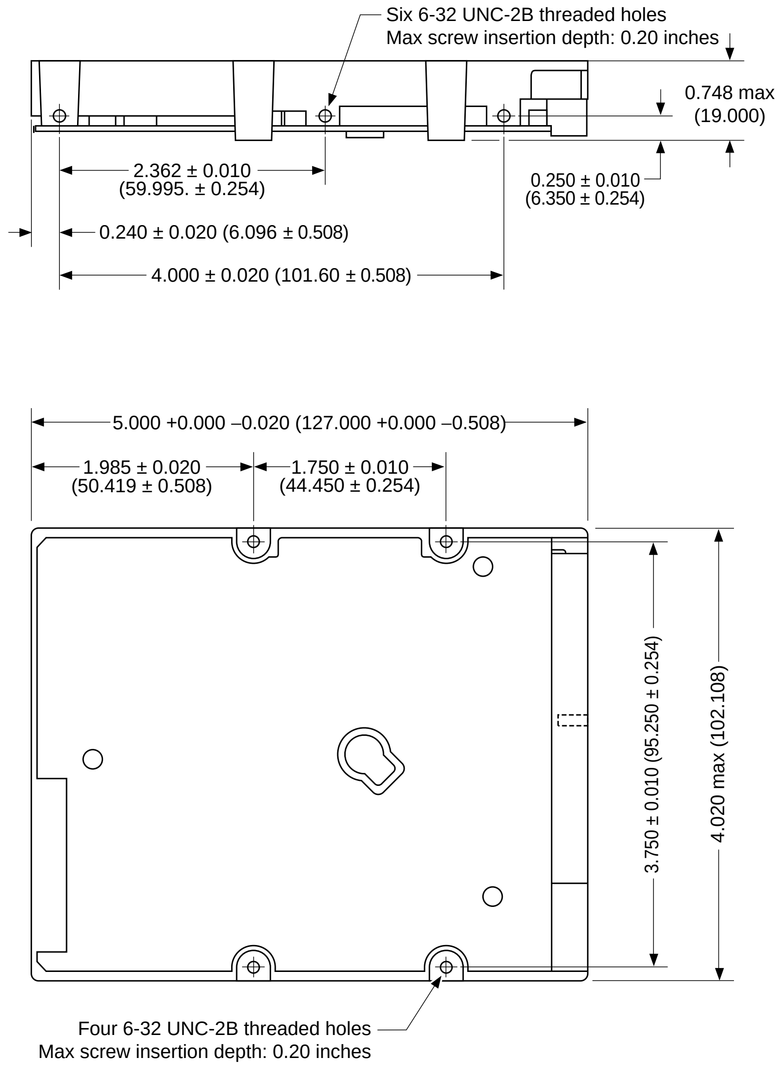

Figure 5. ST51080N mounting dimensions

3.0 Command set

The drive supports a subset of the Group 0 and Group 1 standard SCSI commands. The commands are described in this section.

3.1 Command descriptor block

The initiator makes a request to the drive by sending a command descriptor block (CDB) to the drive. Each CDB has the following common characteristics:

- Byte 0 always contains the operation code.
- The three most significant bits (bits 7–5) of byte 1 contain the logical unit number (LUN). This field is ignored if an Identify Message is sent.
- The last byte is always zero.

3.2 Status byte

The drive terminates each command by sending the status byte (shown below) to the initiator during the status phase before the command complete message.

<table>
<tr><th rowspan="2">Bytes</th><th colspan="8">Bits</th></tr>
<tr><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td rowspan="2">0</td><td colspan="2">Reserved</td><td colspan="5" rowspan="2">Status byte code</td><td>Rsrvd</td></tr>
<tr><td>0</td><td>0</td><td>0</td></tr>
</table>

The *status byte* can be any of the following:

00_H **Good status.** The drive has successfully completed a command.

02_H **Check condition status.** The drive detected an error, an exception or an abnormal condition. In response, the initiator may issue a Request Sense command to determine the nature of the condition.

08_H **Busy status.** The drive is busy and is unable to accept a command from an initiator. The initiator retries the command later. The drive returns a busy status if 1) the initiator has not sent the disconnect message and tries to queue a command or 2) the initiator rejects the disconnect message and the queue is not empty.

18_H **Reservation conflict status.** A SCSI device tried to access the drive, but was unable to because the drive was already reserved by another SCSI device.

28_H **Queue full status.** The drive received a command but rejected it because the queue was full. The drive only uses this status if tagged command queuing is implemented.

3.3 Supported commands

The drive supports the commands listed below.

Group 0 commands	Operation code
Test Unit Ready	00_H
Rezero Unit	01_H
Request Sense	03_H
Format Unit	04_H
Reassign Blocks	07_H
Read	08_H
Write	$0A_H$
Seek	$0B_H$
Inquiry	12_H
Mode Select	15_H
Reserve	16_H
Release	17_H
Mode Sense	$1A_H$
Start/Stop Unit	$1B_H$
Receive Diagnostic Results	$1C_H$
Send Diagnostic	$1D_H$

Group 1 commands	Operation code
Read Capacity	25_H
Read Extended	28_H
Write Extended	$2A_H$
Seek Extended	$2B_H$
Write and Verify	$2E_H$
Verify	$2F_H$
Read Defect Data	37_H
Write Data Buffer	$3B_H$
Read Data Buffer	$3C_H$
Read Long	$3E_H$
Write Long	$3F_H$

3.4 Group 0 commands

3.4.1 Test Unit Ready command (00_H)

The Test Unit Ready command verifies that the drive is ready; it is not a request for a self-test. If the drive can accept an appropriate media access command without encountering an error, it returns a good status.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	0	0	0	0
1	LUN			0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0

3.4.2 Rezero Unit command (01_H)

The Rezero Unit command retracts the read/write heads to the cylinder containing logical block zero.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	0	0	0	1
1	LUN			0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0

3.4.3 Request Sense command (03_H)

The Request Sense command requests the drive to transfer sense data to the initiator in the additional sense data format. The additional sense format is described in Appendix B.

The sense data applies to the previous command on which a check condition status was returned. This sense data is saved for the initiator until:

- The initiator requests the sense data using the Request Sense command, or
- Another command is received from the initiator that issued the original command that caused the check condition status.

If any of the following fatal errors occur during a Request Sense command, the drive sends a check condition status, and the sense data may be invalid.

- The drive receives a nonzero reserved bit in the CDB.
- An unrecovered parity error occurs on the data bus.
- A malfunction prevents return of sense data.

If any other error occurs during the Request Sense command, the drive returns sense data with a good status.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	0	0	1	1
1	LUN			0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	Allocation length							
5	0	0	0	0	0	0	0	0

Byte 4 The *allocation length* specifies the maximum number of bytes the initiator has allocated for returned sense data. The drive returns the number of bytes specified by the allocation length up to 22 bytes. If the allocation length is set to zero, no sense data is returned. This is not an error.

3.4.4 Format Unit command (04_H)

The Format Unit command assures that the medium is formatted so that all of the addressable data blocks can be accessed. In addition, the medium can be certified and control structures may be created for the management of the medium and defects.

If the specified logical unit is reserved, the Format Unit command is rejected with a reservation conflict status. Extent reservations are not supported. See Section 3.4.11 for more information about reservations.

The initiator can specify (or not specify) sectors to be reallocated during the formatting process.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	0	1	0	0
1	LUN			Fmt Data	Cmp lst	Defect list format		
2	0	0	0	0	0	0	0	0
3–4	Interleave							
5	0	0	0	0	0	0	0	0

Byte 1 The *format data* (Fmt Data) bit, the *complete list* (Cmp lst) bit, and the *Defect list format* field are described in Section 3.4.4.2.

Bytes 3–4 The *interleave* field is not supported. It can contain any value. However, the drive always formats the disc with an interleave of 1:1.

3.4.4.1 Defect lists

When the Format Unit command is issued, media defect information can be gathered from several sources. Four of these sources—primary defect list, certification defect list, data defect list and grown defect list—are defect lists written to the drive. They are defined below. Assignments in Byte 1 of the defect list header—described in Section 3.4.4.3—determine the use of the defect list during formatting. The Reassign Blocks and Read Defect Data commands also use these lists.

- The *primary defect list (PList)* is a list of media defects found when the drive is manufactured and written to the disc in an area that is not directly accessible by the user. These defects are considered permanent and cannot be changed.
- The *certification defect list (CList)* is a list of unrecoverable sectors that the drive reads during the certify of the Format Unit command.
- The *data defect list (DList)* is a list of sectors the initiator supplies to the drive during a data-out phase of the current Format Unit command. The drive sends the DList in the last bytes of the defect list (described in Section 3.4.4.3) and adds it to the GList.
- The *grown defect list (GList)* is a list of defects supplied by the initiator or detected by the target but does not include defects from the PList. The GList includes defects detected by the format operation during media certification, the DList, defects previously identified with a Reassign Blocks command and defects previously detected by the target and automatically reallocated.

3.4.4.2 Format Unit parameters

For each format listed below, except the default format, the initiator sends a defect list header. This header is described in Section 3.4.4.3. The bytes-from-index format is described in Section 3.4.4.4 and the physical sector format is described in Section 3.4.4.5. The block format is not supported.

Byte 1 of CDB					Description
Bit 4	**Bit 3**	**Bit 2–Bit 0**			
Fmt Data	**Cmp Lst**	**Defect List Format**			
0	0	X	X	X	*Default format*. The initiator does not send the defect list header or DList to the drive. The drive reallocates all sectors in the PList and erases the GList.
1	0	0	X	X	*Extended format.* The initiator sends a defect list header but no DList. All sectors in the PList and GList are reallocated.
1	0	1	0	0	*Format option with the GList and DList.* The initiator sends the defect list header, which may be followed by a DList in bytes-from-index format. The drive adds the DList to the existing GList. All sectors in the PList and GList are reallocated.
1	0	1	0	1	*Format option with the GList and DList.* The initiator sends the defect list header, which may be followed by a DList in physical sector format. The drive adds the DList to the existing GList. All sectors in the PList and GList are reallocated.
1	1	0	X	X	*Format option without GList or DList*. The drive erases any previous GList. The initiator sends a defect list header but no DList. All sectors in the PList are reallocated.
1	1	1	0	0	*Format option with DList only.* The drive erases any previous GList. The initiator sends the defect list header, which may be followed by a DList in bytes-from-index format. The DList becomes the new GList. All sectors in the PList and GList are reallocated.

Byte 1 of CDB					Description
Bit 4	**Bit 3**	**Bit 2–Bit 0**			
Fmt Data	**Cmp Lst**	**Defect List Format**			
1	1	1	0	1	*Format option with DList only.* The drive erases any previous GList. The initiator sends the defect list header, which may be followed by a DList in physical sector format. The DList becomes the new GList. All sectors in the PList and GList are reallocated.

3.4.4.3 Defect list header and defect list

The defect list, shown below, contains a 4-byte header, followed by one or more defect descriptors. Byte 1 of the defect list header determines whether the P and C defects are reallocated.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	0	0	0	0
1	FOV	DPRY	DCRT	STPF	0	0	0	0
2–3	Defect list length							
4–*n*	Defect descriptor (s)							

Byte 1 If the *FOV* bit is 1, the DPRY, DCRT and STPF bits are interpreted. If the FOV bit is 0, the DPRY, DCRT and STPF bits must be zeros.

The DPRY bit is always 0. The defects described in the PList are reallocated during formatting. The drive sends a check condition status if it cannot find the PList.

If the *DCRT* bit is 1, the drive does not verify the data written during the format. Therefore, no CList for this format is created or reallocated. If the DCRT is 0, the drive verifies the data written during the format, creates a CList and reallocates sectors that were unrecoverable.

If the *STPF* bit is 1, the drive stops formatting if it encounters an error while accessing either the P or G defect list. If the *STPF* bit is 0, the drive continues formatting even though it has encountered an error while accessing either the P or G defect list.

Bytes 2–3 The *defect list length* is the length, in bytes, of the defect list that follows the header. For each sector to be reallocated, the defect list contains one defect descriptor that contains 8 bytes in either the bytes-from-index format or the physical sector format. A length of zero indicates no DList follows; this is not an error.

Bytes 4–*n* The two types of defect descriptors are described in Sections 3.4.4.4 and 3.4.4.5. A length of zero indicates no DList follows; this is not an error.

3.4.4.4 Defect descriptor—bytes-from-index format

Defects are specified in the bytes-from-index format when the defect list format field is 100_{Binary}. See byte 1 of the Format Unit command in Section 3.4.4.

Each defect descriptor in the *bytes-from-index* format specifies the beginning of a single-byte defect location on the disc. Each defect descriptor is composed of the cylinder number of the defect, the head number of the defect and the number of bytes-from-index to the defect location. The defect descriptors are always listed in ascending order.

A value for defect bytes-from-index of $FFFFFFFF_H$ (which means reassign the entire track) is illegal.

The information in the following table is for each defect.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0–2	Cylinder number of defect							
3	Head number of defect							
4–7	Defect bytes-from-index							

3.4.4.5 Defect descriptor—physical sector format

Defects are specified in the physical sector format when the defect list format field is 101_{Binary}. See byte 1 of the Format Unit command in Section 3.4.4.

Each defect descriptor for the physical sector format specifies a sector-size defect location is composed of the cylinder number of the defect, the head number of the defect and the defect sector number. The defect descriptors must be in ascending order.

A defect sector number of $FFFFFFFF_H$ (which means reassign the entire track) is illegal.

Note. The initiator cannot use any previously defined C, G or D lists if the Mode Select command has changed the sector size (block length). For more information on the Mode Select command, see Section 3.4.10.

The information in the following table is for each defect.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0–2	Cylinder number of defect							
3	Head number of defect							
4–7	Defect sector number							

3.4.5 Reassign Blocks command (07_H)

When the drive receives the Reassign Blocks command, it reassigns defective logical blocks to available spare sectors.

Note. ARRE and AWRE may perform automatic reassignments independently of this command.

After sending the Reassign Blocks command, the initiator transfers a defect list containing the logical block addresses to be reassigned. The drive reassigns the logical blocks. The data contained in the logical blocks may not be preserved.

The drive can repeatedly assign a logical block to multiple physical addresses until there are no more spare locations available on the disc.

If the drive does not have enough spare sectors to reassign all of the defective logical blocks, the command terminates with a check condition status, and the sense key is set to media error. The logical block address of the first logical block not reassigned is returned in the information bytes of the sense data.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	0	1	1	1
1	LUN			0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0

3.4.5.1 Reassign Blocks defect list

The Reassign Blocks defect list contains a 4-byte header followed by one or more defect descriptors. The length of each defect descriptor is 4 bytes.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0
2–3	Defect list length							
4–*n*	Defect descriptors							

Byte 2–3 The *defect list length* specifies the total length, in bytes, of the defect descriptors that follow. The defect list length is equal to four times the number of defects.

Bytes 4–*n* The *defect descriptor* contains the 4-byte logical block address of the defect. The defect descriptors must be in ascending order.

3.4.6 Read command (08_H)

When the drive receives the Read command, it transfers data to the initiator.

The Error Recovery page (01_H) determines how the drive handles errors during a Read command. The Error Recovery page is discussed in Appendix C.1.

If there is a reservation access conflict, this command terminates with a reservation conflict status and no data is read. For more information about the reservation conflict status, see Section 3.2.

In systems that support disconnection, the drive disconnects when a valid Read command is received, unless the data is available in the cache buffer and the drive does not need to access the disc. The buffer-full ratio byte of the Disconnect/Reconnect page determines when the drive reconnects. (The Disconnect/Reconnect page is discussed in Section C.2.) The drive may disconnect, if allowed, whenever there is less than one block in the buffer.

Because the drive uses read look-ahead and caching functions, it may read more data into the buffer than specified by the transfer length in the CDB.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	1	0	0	0
1	LUN			Logical block address (MSB)				
2	Logical block address							
3	Logical block address (LSB)							
4	Transfer length							
5	0	0	0	0	0	0	0	0

Bytes 1–3 The *logical block address* specifies the logical block where the read begins.

Byte 4 The *transfer length* specifies the number of contiguous logical blocks of data to be transferred. A transfer length of 0 indicates that 256 logical blocks are to be transferred. Any other value indicates the number of logical blocks to be transferred.

3.4.7 Write command ($0A_H$)

When the drive receives the Write command, it writes the initiator's data to the disc.

The Error Recovery page (01_H) determines how the drive handles bad sectors during a Write command. The Error Recovery page is discussed in Appendix C.1.

If the system supports disconnection, the drive can disconnect and reconnect while executing this command. The drive disconnects when either an internal error recovery procedure is required or the drive's internal data buffer is full.

The buffer-empty ratio in the Disconnect/Reconnect page determines when the drive reconnects. Section C.2 documents the Disconnect/Reconnect page.

If there is a reservation access conflict, this command terminates with a reservation conflict status and no data is written. For more information about the reservation conflict status, see Section 3.2.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	1	0	1	0
1	LUN			Logical block address (MSB)				
2	Logical block address							
3	Logical block address (LSB)							
4	Transfer Length							
5	0	0	0	0	0	0	0	0

Bytes 1–3 The *logical block address* specifies the logical block where the write operation begins.

Byte 4 The *transfer length* specifies the number of contiguous logical blocks of data to be transferred. A transfer length of zero indicates that 256 logical blocks are to be transferred. Any other value indicates the number of logical blocks to be transferred.

3.4.8 Seek command ($0B_H$)

When the drive receives the Seek command, it seeks to the track of the specified logical block address. This command is seldom used because all commands that access the disc contain implied seeks. In systems that support disconnection, the drive disconnects when it receives a valid Seek command.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	1	0	1	1
1	LUN			Logical block address (MSB)				
2	Logical block address							
3	Logical block address (LSB)							
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0

Bytes 1–3 The *logical block address* specifies the logical block to which the head seeks.

3.4.9 Inquiry command (12_H)

When the drive receives the Inquiry command, it sends the inquiry data to the initiator. When the requested inquiry data cannot be returned, a check condition status is reported.

If an Inquiry command is received from an initiator with a pending unit-attention condition (before the drive reports a check condition status), the drive performs the Inquiry command and the Unit Attention condition is not cleared.

The initiator should allocate 36_H bytes for inquiry data. The inquiry data returned to the initiator is summarized in Appendix D.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	1	0	0	1	0
1	LUN			Reserved				EVPD
2	Page code							
	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	Allocation length, in bytes							
5	0	0	0	0	0	0	0	0

Byte 1 If the *enable vital product data (EVPD*) bit is zero, the drive returns the standard inquiry data. If the EVPD bit is one, the drive returns the optional vital product data specified in byte 2.

Byte 2 The *page code* field specifies which page of the vital product information the drive returns. If EVPD is zero, this field must be zero.

Byte 4 The *allocation length* specifies the number of bytes the initiator has allocated for returned inquiry data. The drive returns the number of bytes specified by the allocation length up to a maximum of 148 bytes. If the allocation length is zero, no data is returned. This is not an error. The allocation length should be at least 36_H to allow the initiator to receive all of the standard inquiry data.

3.4.10 Mode Select command (15_H)

The Mode Select command allows the initiator to change parameters stored in the mode pages. The mode pages are described in Appendix C. The drive stores four copies of each mode page:

- **Current values copy.** This copy contains the parameter values the drive uses to control its operation. After a power-on reset, hard reset or bus device reset, the current values are equal to the saved values if the saved values can be retrieved, or the default values if the saved values cannot be retrieved.
- **Changeable values copy.** This copy does not actually contain any parameters. Instead, it contains a map of each mode page indicating which parameters are changeable by the initiator. If a bit contains a 1, the corresponding value in the mode page is changeable. If a bit contains a 0, the corresponding value in the mode page is not

changeable. The changeability values for each bit of each mode page are listed in Appendix C with the default values.

- **Default values copy.** This copy contains the parameter values the drive used as its current values when it was manufactured. The drive defaults to these values after a reset condition, unless valid saved values are available. The default values are listed in Appendix C.
- **Saved values copy.** The saved values are the values the drive stores. If the parameter is changeable, these values can be set using a Mode Select command. If the parameter is not changeable, the default values are always used.

The drive has one set of mode parameters for all of the initiators on the SCSI bus. If the initiator that issued the Mode Select command changes a parameter that applies to another initiator, the drive generates a sense key of *unit attention* with an additional sense of *mode parameters changed* ($2A_H$/01) for all the other initiators. The sense keys and additional sense codes are discussed in Appendix B.

Before sending the Mode Select command, the initiator should send a Mode Sense command requesting that the drive return the changeable values for all pages. The initiator uses this information to determine which pages are supported, the proper length for those pages and which parameters in those pages can be changed for that logical unit. Also, before sending each Mode Select command, the initiator should send a Mode Sense command to request the current values.

When the drive receives the Mode Select command, it updates the saveable parameters with the current values included in the Mode Select command. After the drive saves the parameters, it reports a good status. The drive verifies all Mode Select data.

If the drive detects invalid parameter data during the Mode Select command, it sends a sense key of *illegal request* with an additional sense code of *invalid field in parameter list*, and no parameters are changed.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	1	0	1	0	1
1	LUN			PF = 1	0	0	0	SP
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	Parameter list length							
5	0	0	0	0	0	0	0	0

Byte 1 The *page format (PF)* bit is always one. This means that the data sent by the initiator after the mode select header and block descriptors complies with the page format.

When the *save pages (SP)* bit is 1, the drive saves the saveable pages in nonvolatile memory.

When the *save pages (SP)* bit is 0, the drive saves the saveable pages in RAM only, which means that the parameters are lost when the drive is powered down.

Byte 4 The *parameter list length* specifies the length, in bytes, of the header and mode page transferred to the drive. A parameter list length of 0 means that no data is transferred. To calculate the parameter list length for any given mode page, add the parameter list header (4 bytes), the block descriptor (if any, 8 bytes), the 2-byte mode page header and the length of each mode page. For the lengths of the mode pages, refer to Appendix C.

3.4.10.1 Mode Select parameter list

The Mode Select parameter list contains a 4-byte header, followed by a 1-block descriptor (if any), followed by the Mode Select parameter pages.

Each block descriptor specifies the media characteristics for all or part of a logical unit. The rest of the Mode Select parameters are grouped by function and organized into mode pages. The mode pages are described in Appendix C.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
Parameter list header								
0 (default)	Reserved (00_H)							
1 (default)	Medium type (00_H)							
2 (default)	Reserved (00_H)							
3 (default)	Block descriptor length (00_H or 08_H)							
Block descriptor data								
4 (default)	Density code (00_H)							
5–7	Number of blocks							
8 (default)	Reserved (00_H)							
9–11	Block length							
Parameter information								
12–*n*	Mode pages							

Byte 1 The *medium type* field is always 00_H, which means that the drive is a direct-access device.

Byte 3 If the *block descriptor length* is 8 bytes, a block descriptor is sent to the drive. If the *block descriptor length* is 0 bytes, no block descriptor is sent to the drive.

Byte 4 The *density code* is always 00_H and cannot be changed.

Bytes 5–7 The *number of blocks* is equal to the guaranteed sectors, which is listed in the formatted capacity section of the appropriate product manual.

Bytes 9–11 The *block length* is always 0200_H and cannot be changed.

3.4.11 Reserve command (16_H)

When the initiator issues a Reserve command, it requests that the drive be reserved for exclusive use by the initiator until the reservation is:

- Superseded by another Reserve command from the initiator that made the reservation. An initiator that has already reserved the drive can modify that reservation by issuing another Reserve command. When the drive receives the superseding Reserve command, the previous reservation is canceled.
- Released by a Release command from the same initiator. See the Release command in Section 3.4.12.
- Released by a bus device reset message from any initiator.
- Released by a hard reset.

After the drive honors the reservation from one initiator, it accepts only Request Sense and Inquiry commands from other initiators; the drive rejects all other commands with a reservation conflict status.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	1	0	1	1	0
1	LUN			3rd pty	3rd party device ID			Extent
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0

Byte 1 If the *3rd pty* bit is 0, the initiator reserves the drive for itself. If the 3rd pty bit is 1, the initiator reserves the drive for another

initiator. The SCSI ID of the third-party initiator is specified in the *3rd party device ID* field.

The *extent* bit must always be 0. The drive does not support extent reservations. If the extent bit is 1, the drive generates a check condition status.

3.4.12 Release command (17_H)

When an initiator that had reserved the drive using the Reserve command issues the Release command, it cancels the reservation. If the drive is not currently reserved and it receives a Release command, the drive returns a good status.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	1	0	1	1	1
1	LUN			3rd pty	3rd party device ID			Extent
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0

Byte 1 If the *3rd pty* bit is 0, the initiator releases its own reservation. If the 3rd pty bit is 1, the initiator releases the drive for another initiator. An initiator can only release a third-party reservation that it made. The SCSI ID of the third-party initiator is specified in the *3rd party device ID* field.

The *extent* bit must always be 0. The drive does not support extent reservations. If the extent bit is 1, the drive generates a check condition status.

3.4.13 Mode Sense command ($1A_H$)

When the initiator sends this command to the drive, it returns mode-page parameters to the initiator. This command is used in conjunction with the Mode Select command.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	1	1	0	1	0
1	LUN			0	0	0	0	0
2	PCF		Page code					
3	0	0	0	0	0	0	0	0
4	Allocation length							
5	0	0	0	0	0	0	0	0

Byte 2 The *page control field (PCF)* determines the content of Mode Parameter bytes. Regardless of the value of the PCF, the block descriptor always contains the current values.

PCF bit 7	PCF bit 6	Effect
0	0	Return current values.
0	1	Return changeable values.
1	0	Return default values.
1	1	Return saved values.

The *page code* is the designator that is unique to each page. The page codes are listed in Section 3.4.13.1.

Byte 4 The *allocation length* specifies the number of bytes that the initiator has allocated for returned Mode Sense data. An allocation length of 0 means that no Mode Sense data is to be transferred. This condition is not considered an error. Any other value represents the number of bytes to be transferred. For a description of the allocation length, see Section 3.4.13.1.

3.4.13.1 Page code and allocation length

The Mode Sense command descriptor block contains a page code (byte 2, bits 5–0) and an allocation length (byte 4). These parameters are described in the following table for SCSI-2 devices. You can transfer mode pages to the initiator either of two ways:

- Transfer all mode pages at once by using page code $3F_H$, or
- Transfer one mode page at a time by using the page codes and any number greater than or equal to the allocation length of the mode page.

Page code	Allocation length	Mode Sense data returned
01_H	18_H	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 10 bytes of Error Recovery parameters
02_H	$1C_H$	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 14 bytes of Disconnect/Reconnect parameters
03_H	24_H	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 22 bytes of Format Device parameters
04_H	24_H	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 22 bytes of Rigid Disc Geometry parameters
07H	18_H	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 10 bytes of Verify Error Recovery Page parameters
08_H	20H	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 18 bytes of Caching parameters
$0A_H$	18_H	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 10 bytes of Control Mode page parameters

Page code	Allocation length	Mode Sense data returned
$0C_H$	24_H	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 22 bytes of Notch and Partition parameters
$3C_H$	0FH	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 1 bytes of Soft ID parameters
38H	$1C_H$	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 14 bytes of Caching parameters
00_H	10_H or 11_H	4 bytes of Mode Sense header 8 bytes of block descriptor 2 bytes of mode-page header 2 or 3 bytes of Operating parameters

3.4.13.2 Mode Sense data

The Mode Sense parameter list contains a 4-byte header followed by an 8-byte block descriptor (if any), followed by the mode pages. The header and block descriptor are shown below. The mode pages are described in Appendix C.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Mode Sense data length							
1 (default)	Medium type (00_H)							
2	WP=0	Reserved						
3 (default)	Block descriptor length (08_H)							
Block descriptor								
4 (default)	Density code (00_H)							
5–7	Number of blocks							
8 (default)	Reserved (00_H)							
9–11	Block length							
Mode pages								
12–*n*	Mode pages							

Byte 0 The *Mode Sense data length* specifies the number of bytes minus 1 of the Mode Sense data to be transferred to the initiator.

Byte 1 The *medium type* is always 0.

Byte 2 The *WP (write protect)* bit is always 0, which means the media is write-enabled.

Byte 3 The *block descriptor length* is the number of bytes in the block descriptor. This value does not include the page headers and mode pages that follow the block descriptor, if any.

Byte 4 The *density code* is not supported.

Bytes 5–7 The *number of blocks* field contains the total number of blocks available to the user, which is specified on page 1.

Byte 8 Reserved

Bytes 9–11 The *block length* specifies the number of bytes contained in each logical block described by the block descriptor.

3.4.14 Start/Stop Unit command ($1B_H$)

When the drive receives the Start/Stop Unit command, the drive either spins up or spins down, depending on the setting of the start bit in byte 4.

If the host adapter supports disconnection, the drive disconnects when it receives the Start/Stop Unit command and reconnects when it is up to speed and ready.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	1	1	0	1	1
1	LUN = 0			0	0	0	0	Immed
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	Start
5	0	0	0	0	0	0	0	0

Byte 1 If the *immediate (Immed)* bit is 0, the drive returns the status after the command is completed. If the Immed bit is 1, the drive returns the status when it receives the command.

Byte 4 If the *start* bit is 1, the drive spins up. If the *start* bit is 0, the drive spins down.

3.4.15 Receive Diagnostic Results command ($1C_H$)

When the drive receives the Receive Diagnostics command after power-up or after a Send Diagnostic command with PF=0, it sends eight diagnostic data bytes to the initiator. The initiator sends the Receive Diagnostic Results command after the drive completes the Send Diagnostic command. The page length is describe in the Send Diagnostics command. The most recent Send Diagnostic command determines the data returned by the Receive Diagnostic Results command.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	1	1	1	0	0
1	LUN = 0			0	0	0	0	0
2	0	0	0	0	0	0	0	0
3–4	Allocation length							
5	0	0	0	0	0	0	0	0

Bytes 3–4 The *allocation length* specifies the number of bytes the initiator has allocated for returned diagnostic result data. An allocation length of 0 means that no diagnostic data is transferred. The drive sends the allocation length or the bytes available, whichever number is less.

3.4.15.1 Diagnostic data format

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0–1 (default)	Additional length (0006_H)							
2–5	FRU code							
6	Diagnostic error code							
7	Vendor-unique error code							

Byte 0–1 The *additional length* value indicates the number of additional bytes included in the diagnostic data list. A value of 0000_H means that there are no additional bytes. A value of 0006_H means that no product-unique bytes are available.

Bytes 2–5 If the *field replaceable unit (FRU)* code is 00_H, there is no FRU information. If the FRU code is 01_H, replace the drive. Other values are drive-unique.

Byte 6 The *diagnostic error code* is not supported.

Byte 7 The *vendor-unique error codes* are listed in Section 3.4.15.2.

3.4.15.2 Diagnostic error codes

The following diagnostic error codes are reported in byte 7 of the diagnostic data format in Section 3.4.15.1.

Error code	Description
01_H	Sequencer test error
02_H	Microprocessor RAM diagnostic error
09_H	Fatal hardware error during drive diagnostics
44_H	EEPROM test error
80_H	Buffer controller diagnostic error
81_H	Buffer RAM diagnostic error

3.4.16 Send Diagnostic command ($1D_H$)

When the drive receives this command, it performs diagnostic tests on itself. In systems that support disconnection, the drive disconnects while executing this command.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	1	1	1	0	1
1	LUN = 0			PF	0	Self Test	Dev OfL	Unit OfL
2	0	0	0	0	0	0	0	0
3–4 (default)	Parameter list length (00_H)							
5	0	0	0	0	0	0	0	0

Byte 1 If the PF (Page Format) bit is set to 0 and the *self test* bit is 1, the drive performs the buffer RAM diagnostics, which is the default self-test. If the default self- test is requested, the parameter list length is 0 and no data is transferred. If the self-test passes successfully, the command terminates with a good status. If the self-test fails, the command terminates with a check condition status and the sense key is hardware error.

If the PF bit is set to 1, *SelfTest*, *device off line* (DevOfl) and *unit off line* (UnitOfl) are ignored and a diagnostic page is sent as the parameter list. The supported pages are the *Supported Pages* (00_H) page and the *Translate Address* (40_H) page. The parameter length is 4 (04_H) bytes for page 00_H and 14 ($0E_H$) bytes for page 40_H.

The DevOfL bit is not supported and must be zero if *SelfTest* bit =1.

The UnitOfL bit is not supported and must be zero if *SelfTest* bit =1.

Bytes 3–4 The *parameter list length* must be zero if *SelfTest* bit =1.

3.5 Group 1 commands

3.5.1 Read Capacity command (25_H)

The initiator uses the Read Capacity command to determine the capacity of the drive. When the drive receives the Read Capacity command, it sends the initiator read capacity data, which is described in Section 3.5.1.1.

Bytes	Bits							
	7	6	5	4	3	2	1	0
0	0	0	1	0	0	1	0	1
1	LUN			0	0	0	0	0
2–5	Logical block address							
6	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	PMI
9	0	0	0	0	0	0	0	0

Bytes 2–5 The *logical block address* specified in the CDB cannot be greater than the logical block address reported by the drive in the read capacity data.

Byte 8 If the *partial medium indicator (PMI)* bit is zero, the logical block address in the CDB is also zero. The read capacity data returned by the drive contains the logical block address and block length of the last logical block of the drive.

If the PMI bit is one, the drive returns the read capacity data, which contains the logical block address and block length of the last logical block address, after which a substantial delay (approximately 1 msec) in data transfer occurs. This logical block address must be greater than or equal to the logical block address specified in the CDB. This reported logical block address is a cylinder boundary.

3.5.1.1 Read Capacity data

The Read Capacity data is shown below.

Bytes	Bits							
	7	6	5	4	3	2	1	0
0–3	Logical block address							
4–7	Block length (00000200_H)							

Bytes 0–3 The logical block address is determined by the PMI bit in the CDB of the Read Capacity command. The PMI bit is described in Section 3.5.1.

Bytes 4–7 The block length is always 512.

3.5.2 Read Extended command (28_H)

When the drive receives the Read Extended command, it transfers data to the initiator. This command is the same as the Read command discussed in Section 3.4.6, except that in the CDB for the Read Extended command, a 4-byte logical block address and a 2-byte transfer length can be specified.

If there is a reservation access conflict, this command terminates with a reservation conflict status and no data is read. For more information about the reservation conflict status, see Section 3.2.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	1	0	1	0	0	0
1	LUN			DPO	FUA	0	0	0
2–5	Logical block address							
6	0	0	0	0	0	0	0	0
7–8	Transfer length							
9	0	0	0	0	0	0	0	0

Byte 1 If the *disable page out (DPO)* bit is one, the cached data that the drive receives during this command has the lowest priority for being retained in the cache. If the DPO is zero, the cached data has the highest priority for being retained in the cache.

If the *forced unit access (FUA)* bit is one, the drive must access the disc to get the data requested by the initiator, even if the data is available in the cache. If the FUA bit is zero, the drive can get the data from the cache or the disc.

Bytes 2–5 The *logical block address* specifies the logical block where the read operation begins.

Bytes 7–8 The *transfer length* specifies the number of contiguous logical blocks of data to be transferred. A transfer length of 0 means that no logical blocks are to be transferred. This condition is not considered an error.

3.5.3 Write Extended command ($2A_H$)

When the drive receives the Write Extended command, the drive writes the data from the initiator to the disc. This command is like the Write command, except that the CDB for this command contains a 4-byte logical block address and a 2-byte transfer length. For more information about the Write command, see Section 3.4.7.

If there is a reservation access conflict, this command terminates with a reservation conflict status and no data is written. For more information about the reservation conflict status, see Section 3.2.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	1	0	1	0	1	0
1	LUN			DPO	FUA	0	0	0
2–5	Logical block address							
6	0	0	0	0	0	0	0	0
7–8	Transfer length							
9	0	0	0	0	0	0	0	0

Byte 1 If the *disable page out (DPO)* bit is one, the cached data that the drive receives during this command has the lowest priority for being retained in the cache. If the DPO is zero, the cached data has the highest priority for being retained in the cache.

If the *forced unit access (FUA)* bit is one, the drive must access the disc to write the data sent by the initiator, even if the data can be stored in the cache. If the FUA bit is zero, the drive can write the data to the cache or the disc.

Bytes 2–5 The *logical block address* specifies the logical block where the write operation begins.

Bytes 7–8 The *transfer length* specifies the number of contiguous logical blocks of data to be transferred. A transfer length of zero means that no logical blocks are to be transferred; this is not an error.

3.5.4 Seek Extended command ($2B_H$)

The Seek Extended command requests that the drive seek to the specified logical block address. This command is the same as the Seek command, except that the CDB includes a 4-byte logical block address. The Seek command is described in Section 3.4.8.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	1	0	1	0	1	1
1	LUN			0	0	0	0	0
2–5	Logical block address							
6	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0

3.5.5 Write and Verify command ($2E_H$)

When the drive receives the Write and Verify command, it writes the data sent by the initiator to the media and then verifies that the data is correctly written.

If the host adapter supports disconnection, the drive disconnects while it is executing this command.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	1	0	1	1	1	0
1	LUN			0	0	0	BytChk	0
2–5	Logical block address							
6	0	0	0	0	0	0	0	0
7–8	Transfer length							
9	0	0	0	0	0	0	0	0

Byte 1 If the *byte check (BytChk)* bit is zero, the drive verifies the media after a write by checking the ECC syndromes. If the BytChk bit is one, the drive verifies the media after a write by performing a byte-by-byte comparison of the data stored.

Bytes 2–5 The *logical block address* field specifies the logical block where the drive begins writing and verifying the data.

Bytes 7–8 The *transfer length* field specifies the number of contiguous logical blocks to be transferred. If the transfer length is zero, the initiator does not transfer any data and the drive does not write or verify any data. This condition is not considered an error.

3.5.6 Verify command (2F$_H$)

When the drive receives the Verify command, it verifies the data on the disc. If the host adapter supports disconnection, the drive disconnects while it is executing this command.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	1	0	1	1	1	1
1	LUN			0	0	0	Byt Chk	0
2–5	Logical block address							
6	0	0	0	0	0	0	0	0
7–8	Verification Length							
9	0	0	0	0	0	0	0	0

Byte 1 If the *byte check (BytChk)* bit is zero, the drive verifies the media by checking the ECC syndromes. If the BytChk bit is one, the drive verifies the media by performing a byte-by-byte comparison of the stored data.

Bytes 2–5 The *logical block address* field specifies the logical block where the drive begins verifying the data.

Bytes 7–8 The *verification length* field specifies the number of contiguous logical blocks to be verified. If the verification length is zero, the drive does not verify any logical blocks, although an implied seek is still performed. This condition is not considered an error.

3.5.7 Read Defect Data command (37$_H$)

When the drive receives this command, it reads the defect data from reserved cylinders or flash memory and transfers the defect data to the initiator.

This command can be used in conjunction with the Format Unit command. Read Defect Data reads the defect lists off the reserved cylinders or flash memory and resends the lists as defect data but does not change the lists.

The Read Defect Data command can be used to access two types of defect lists: the *primary defect list (PList)* and the *grown defect list (GList)*. These lists are described in Section 3.4.4.1.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	1	1	0	1	1	1
1	LUN			0	0	0	0	0
2	0	0	0	PList	GList	Defect list format		
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0
7–8	Allocation length							
9	0	0	0	0	0	0	0	0

Byte 2 If the *PList* bit is 1, the drive sends the primary defect list. If the PList bit is 0, the drive does not send the primary defect list.

If the *GList* bit is 1, the drive sends the grown defect list. If the GList bit is 0, the drive does not send the grown defect list.

If both the PList and GList bits are zero, the drive returns the defect list header only.

If the *defect list format* field contains 100_H, the drive returns the defect data in the bytes-from-index format. If the defect list format field contains 101_H, the drive returns the defect data in the physical sector format. If the defect list format field contains 000_H, the drive returns the defect data in the default format, which is the physical sector format, and generates a check condition status.

Bytes 7–8 The *allocation length* specifies the number of bytes the initiator has allocated for the returned defect data. An allocation length of 0 indicates that no defect data is transferred. The data-in phase ends when the allocation length bytes have been transferred or when all available defect data has been transferred to the initiator, whichever is less.

3.5.7.1 Defect list header

The defect data always begins with a 4-byte header, followed by a 8-byte descriptor for each defect. The defect list header format is described below.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	0	0	0	0	0	0
1	0	0	0	PList	GList	Defect list format		
2–3	Defect List Length							

Byte 1 If the *PList* bit is 1, the defect data contains the primary defect list. If the PList bit is 0, the defect data does not contain the primary defect list.

If the *GList* bit is 1, the defect data contains the grown defect list. If the GList bit is 0, the defect data does not contain the grown defect list.

The *defect list format* field is described in Section 3.5.7.

Bytes 2–3 The *defect list length* specifies the length of the defect data in bytes. If the PList and GList bits are 0, no defect descriptor bytes are sent to the initiator and the defect list length is 0. If the allocation length (in the CDB) is not large enough to accommodate all the defect descriptors, the defect list length contains the same value as the allocation length.

3.5.8 Write Data Buffer command ($3B_H$)

The Write Data Buffer command supports several different features.

The Write Data Buffer command can be used along with the Read Data Buffer command to diagnose problems in the drive's data buffer memory and to test the integrity of the SCSI bus.

You can also use the Write Data Buffer command to download microcode to the buffer and also to save it in flash memory.

Note. This command treats the buffer as a single segment, regardless of the number of segments specified in Caching page 08_H. (Caching page 08_H is described in Section C.6.3.)

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	1	1	1	0	1	1
1	LUN			0	0	Mode		
2	Buffer ID (00_H)							
3–5	Buffer offset							
6–8	Parameter list length							
9	0	0	0	0	0	0	0	0

Byte 1 If the *mode* bits contain 000_B, the initiator transfers data to the drive buffer with a 4-byte header that contains all zeros. This mode is called *write combined header and data*.

If the *mode* bits contain 010_B, the initiator transfers data to the drive buffer without the header. This mode is called *write data*.

If the *mode* bits contain 101_B, the initiator downloads microcode to the drive buffer, and the drive saves the microcode in flash memory. The drive uses the new microcode for all future operations. This mode is called *download microcode and save*.

After the microcode has been successfully downloaded, the drive generates a unit attention condition of *microcode has been downloaded* for all initiators except the one that issued the current Write Data Buffer command.

All other settings for the mode bits are reserved.

Byte 2 The *buffer ID* is not supported and must always be zero.

Byte 3–5 The *buffer offset* is added to the starting address of the buffer to determine the destination of the first data byte. The bytes that follow are placed in sequential addresses. If the sum of the buffer offset and the transfer length exceeds the buffer size reported by the Read Data Buffer command (see Section 3.5.9), the drive generates a check condition status and the initiator does not transfer any data. This field is ignored if the mode bit is 101_B.

Bytes 6–8 The *parameter list length* field specifies the maximum number of bytes the initiator transfers. If the initiator transfers the 4-byte header, the transfer length includes the header. If the transfer length is zero, no data is transferred to the drive buffer; this is not considered an error.

3.5.9 Read Data Buffer command ($3C_H$)

The Read Data Buffer command supports several different features.

The Read Data Buffer command can be used along with the Write Data Buffer command to diagnose problems in the drive's data buffer memory and to test the integrity of the SCSI bus.

Note. This command treats the buffer as a single segment, regardless of the number of segments specified in mode page (08_H). [Mode page (08_H), the Caching page, is described in Section C.5.3.]

<table>
<tr><th rowspan="2">Bytes</th><th colspan="8">Bits</th></tr>
<tr><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td>0</td><td>0</td><td>1</td><td>1</td><td>1</td><td>1</td><td>0</td><td>0</td></tr>
<tr><td>1</td><td colspan="3">LUN</td><td>0</td><td>0</td><td colspan="3">Mode</td></tr>
<tr><td>2</td><td colspan="8">Buffer ID (00_H)</td></tr>
<tr><td>3–5</td><td colspan="8">Buffer offset</td></tr>
<tr><td>6–8</td><td colspan="8">Allocation length</td></tr>
<tr><td>9</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr>
</table>

Byte 1 If the *mode* bits contain 000_B, the initiator reads data from the drive buffer. The data is preceded by a 4-byte header. This mode is called *read combined header and data*.

If the *mode* bits contain 010_B, the initiator reads data from the drive buffer without a header. This mode is called *read data*.

All other settings for the mode bits are reserved.

Byte 2 The *buffer ID* is not supported and must always be zero.

Byte 3–5 The *buffer offset* is added to the starting address of the buffer to determine the source of the first data byte. The bytes that follow are read from sequential addresses. If the sum of the buffer offset and the transfer length exceeds the available length reported in the Read Buffer header (see Section 3.5.9.1), the drive transfers all the data contained in the buffer.

Bytes 6–8 The *allocation length* field specifies the maximum number of bytes read by the initiator. If the 4-byte header is transferred, the transfer length includes the header. If the transfer length is zero, no data is read; this is not an error.

3.5.9.1 Read Buffer Header

The following table shows the structure of the 4-byte Read Buffer Header.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0							
1–3	Buffer capacity							

Bytes 1–3 The *buffer capacity* field specifies the size of the drive buffer. Byte 1 is MSB; byte 3 is LSB.

3.5.10 Read Long command ($3E_H$)

When the drive receives the Read Long command, it transfers data to the initiator.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	1	1	1	1	1	0
1	LUN			0	0	0	0	0
2–5	Logical block address							
6	0	0	0	0	0	0	0	0
7–8	Byte transfer length							
9	0	0	0	0	0	0	0	0

Bytes 2–5 The *logical block address* specifies the LBA where the drive begins reading data.

Bytes 7–8 The *byte transfer length* specifies the number of bytes transferred to the initiator. The drive transfers the logical block size plus eleven. If the byte transfer length is zero, the drive does not transfer any data to the initiator. This condition is not considered an error.

Transfer length must be 523 ($20B_H$) Bytes.

3.5.11 Write Long command ($3F_H$)

When the drive receives the Write Long command, it writes one logical block of data and eleven bytes of error correction code (ECC) to the disc. During this command, the drive does not perform any ECC verification.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	0	0	1	1	1	1	1	1
1	LUN			0	0	0	0	0
2–5	Logical block address							
6	0	0	0	0	0	0	0	0
7–8	Byte transfer length							
9	0	0	0	0	0	0	0	0

Bytes 2–5 The *logical block address* specifies the LBA where the drive begins writing data.

Bytes 7–8 The *byte transfer length* specifies the number of bytes the initiator transfers to the drive.

If the transfer length does not equal the sum of the logical block size plus eleven, the command is terminated with a check condition status.

If the byte transfer length is zero, the initiator does not transfer any data to the drive; this condition is not considered an error.

Transfer length must be 523 ($20B_H$) Bytes.

3.6 Group 2, 3 and 4 commands

Group 2, 3 and 4 commands are 10-byte commands. Group 2 commands are not implemented. Group 3 and 4 commands are reserved. If the drive receives one of these commands, it returns a check condition status.

Caution. Do not use Group 3 and 4 commands. If you do, you may destroy data on the disc.

3.7 Group 5 and 6 commands

Group 5 and 6 commands are 12-byte commands. Group 5 commands are not implemented. If the drive receives a Group 5 command, it returns a check condition status. Group 6 commands are reserved for Seagate use.

Caution. Do not use Group 6 commands. If you do, you may destroy data on the disc.

3.8 Group 7 commands

Group 7 commands are 10-byte commands. These commands are not implemented. If the drive receives one of these commands, it returns a check condition status.

Appendix A. Supported messages

A.1 Single-byte messages

The implemented single-byte messages are listed below.

Code	Message name	Direction	Must negate ATN before last ACK?
06_H	Abort	O	Yes
$0D_H$	Abort tag	O	Yes
$0C_H$	Bus device reset	O	Yes
$0E_H$	Clear queue	O	Yes
00_H	Command complete	I	—
04_H	Disconnect	I	—
80_H	Identify	I/O	No
05_H	Initiator detected error	O	Yes
09_H	Message parity error	O	Yes
07_H	Message reject	I/O	Yes
08_H	No operation	O	Yes
21_H	Head of queue tag	O	No
22_H	Ordered queue tag	O	No
20_H	Simple queue tag	O	No
03_H	Restore pointers	I	—
02_H	Save data pointer	I	—

A.2 Synchronous data transfer request message (01_H)

The synchronous data transfer message is the only extended message that the drive supports.

Depending on the value contained in the SSM bit (contained in byte 2 of the Operating page in Appendix C.11), the drive or the initiator can negotiate for synchronous data transfer after a reset. If any problem precludes the successful exchange of synchronous data transfer request messages, the initiator and drive default to asynchronous data transfers.

This exchange of messages establishes the minimum transfer period and the maximum allowed REQ/ACK offset.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Extended message (01_H)							
1	Extended message length (03_H)							
2	Identifier code (01_H)							
3	Minimum transfer period divided by 4							
4	REQ/ACK offset							

Byte 0 This byte identifies the message as an extended message.

Byte 1 This byte reports the length of the message.

Byte 2 This byte identifies the message as a synchronous data transfer request message.

Byte 3 The value contained in this byte is in nanoseconds. It is equal to the minimum time between leading edges of successive REQ and ACK pulses divided by four. In byte 3, the minimum value supported by the drive is 25, which is equivalent to a transfer period of 100 nanoseconds, or an external transfer rate of 10 Mbytes per second. A value of 50 is equivalent to a transfer period of 200 nanoseconds, or an external transfer rate of 5 Mbytes per second.

Byte 4 The *REQ/ACK offset* is the maximum number of REQ pulses that may be outstanding before its corresponding ACK pulse is received at the target. A REQ/ACK offset of zero indicates asynchronous mode. The drive supports a maximum REQ/ACK offset of $0F_H$.

Appendix B. Sense data

The appendix contains the descriptions for sense data returned by the Request Sense command. For more information on the Request Sense command, see the *Seagate SCSI-2/3 Interface Manual Volume 2; Version 2*, publication number 77738479-D.

B.1 Additional sense data

When the initiator issues a Request Sense command, the drive returns the following additional sense data.

<table>
<tr><th rowspan="2">Bytes</th><th colspan="8">Bit</th></tr>
<tr><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td rowspan="2">0</td><td rowspan="2">Valid</td><td colspan="7">Error code</td></tr>
<tr><td>1</td><td>1</td><td>1</td><td>0</td><td>0</td><td>0</td><td>X</td></tr>
<tr><td>1</td><td colspan="8">Segment number (00_H)</td></tr>
<tr><td>2</td><td>0</td><td>0</td><td>ILI</td><td>0</td><td colspan="4">Sense key</td></tr>
<tr><td>3–6</td><td colspan="8">Information bytes</td></tr>
<tr><td>7</td><td colspan="8">Additional sense length</td></tr>
<tr><td>8–11</td><td colspan="8">Command specific data</td></tr>
<tr><td>12</td><td colspan="8">Additional sense code</td></tr>
<tr><td>13</td><td colspan="8">Additional sense code qualifier</td></tr>
<tr><td>14</td><td colspan="8">FRU code</td></tr>
<tr><td>15</td><td>SKSV</td><td colspan="7">Sense key specific</td></tr>
<tr><td>16–17</td><td colspan="8">Sense key specific</td></tr>
<tr><td>18–22</td><td colspan="8">Product-unique sense data (00_H)</td></tr>
</table>

Byte 0 If the *valid* bit is one, the information bytes (bytes 3 through 6) are valid. If the valid bit is zero, the information bytes are not valid.

If *error code* contains a value of 70_H, the error occurred on the command that is currently pending. If *error code* contains a value of 71_H, the error occurred during the execution of a previous command for which a good status has already been returned.

Byte 1 The *segment number* is always zero.

Byte 2 If the *incorrect length indicator (ILI)* bit is zero, the requested block of data from the previous command did not match the logical block length of the data on the disc. If the ILI bit is one, the requested block of data from the previous command matched the logical block length of the data on the disc.

The *sense key* indicates one of nine general error categories. These error categories are listed in Appendix B.2.

Bytes 3–6 When the *valid* bit is 1, the *information bytes* contain the logical block address of the current logical block associated with the sense key. For example, if the sense key is media error, the information bytes contain the logical block address of the offending block.

Byte 7 The *additional sense length* is limited to a maximum of $0E_H$ additional bytes. If the allocation length of the command descriptor block is too small to accommodate all of the additional sense bytes, the additional sense length is not adjusted to reflect the truncation.

Bytes 8–11 These bytes contain command-specific data.

Bytes 12–13 The *additional sense code* and *additional sense code qualifier* provide additional details about errors. See Appendix B.3.

Byte 14 The *field replaceable unit (FRU) code* is used by field service personnel only.

Bytes 15–22 These bytes are not used and are always 00_H.

B.2 Sense key

The sense keys in the lower-order bits of byte 2 of the sense data returned by the Request Sense command are described in the following table. You can find a more detailed description of the error by checking the additional sense code and the additional sense code qualifier in Section B.3.

Sense key	Description
0_H	**No Sense.** In the case of a successful command, no specific sense key information needs to be reported for the drive.
1_H	**Recovered error.** The drive completed the last command successfully with some recovery action. When many recovered errors occur during one command, the drive determines which error to report.

Sense key	Description
2_H	**Not ready.** The addressed logical unit cannot be accessed. Operator intervention may be required to correct this condition.
3_H	**Medium error.** The command was terminated with a nonrecoverable error condition, probably caused by a flaw in the media or an error in the recorded data.
4_H	**Hardware error.** The drive detected a nonrecoverable hardware failure while performing the command or during a self-test. This includes, for example, SCSI interface parity errors, controller failures and device failures.
5_H	**Illegal request.** This indicates an illegal parameter in the CDB or in the additional parameters supplied as data for some commands (for example, the Format Unit command, the Mode Sense command and others). If the drive detects an invalid parameter in the CDB, it terminates the command without altering the media. If the drive detects an invalid parameter in the additional parameters supplied as data, the drive may have already altered the media.
6_H	**Unit attention.** The drive may have been reset. See the *Seagate SCSI-2 Interface Manual* for more details about the Unit Attention condition.
B_H	**Aborted command.** The drive aborted the command. The initiator may be able to recover by retrying.
E_H	**Miscompare.** The source data did not match the data read from the media.

B.3 Additional sense code and additional sense code qualifier

The additional sense code and additional sense code qualifiers returned in byte 12 and byte 13, respectively, of the Sense Data Format of the Request Sense command are listed in the following table.

Error code (hex)		
Byte 12	**Byte 13**	**Description**
00	00	No additional information is supplied.
01	00	There is no index/sector signal.
02	00	There is no seek complete signal.
03	00	A write fault occurred.
04	00	The drive is not ready.
05	00	The drive does not respond when it is selected.
06	00	Track 0 was not found.
07	00	More than one drive is selected at a time.
0C	00	A write error occurred.
10	00	An ID CRC or ECC error occurred during retries.
11	00	An unrecovered read error occurred.
12	00	The address mark was not found in the ID field.
13	00	The address mark was not found in the data field.
14	00	No record was found.
14	01	No record was found.
15	00	A seek positioning error occurred.
16	00	A data address mark was recovered.
17	01	The data was recovered with retries.
18	01	The data was recovered with ECC and retries.
18	02	The data was recovered and ARRE was performed.
19	00	There is an error in the defect list.
1A	00	A parameter overrun occurred.

Error code (hex)		
Byte 12	**Byte 13**	**Description**
1B	00	A synchronous transfer error occurred.
1C	00	The defect list could not be found.
1C	01	The primary defect list could not be found.
1D	00	A miscompare occurred during a verify operation.
1E	00	An ID error was recovered.
20	00	The drive received an invalid command operation code.
21	00	The logical block address was not within the acceptable range.
24	00	The drive received a CDB that contains an invalid bit.
25	00	The drive received an invalid LUN.
26	00	The drive received an invalid field in the Parameter List.
29	00	A power-on reset or a bus device reset occurred.
2A	01	The Mode Select parameters were changed by another initiator.
2F	00	The commands were cleared by another initiator.
31	01	The format command failed.
31	97	Seek error while updating head parameters during format.
31	98	Read error while updating head parameters during format.
31	99	Invalid head parameter while updating head parameters during format.
31	9A	Flash update error while updating head parameters during format.
31	9B	Bad format with good head parmeters.
31	9C	Inconsistent or bad head parameter detected during drive initilization.

continued

continued from previous page

Error code (hex)		
Byte 12	**Byte 13**	**Description**
32	00	No spare defect locations available during format.
32	01	GList update full or error during reassign.
37	00	A rounded parameter caused an error.
3D	00	The identify message contains invalid bits.
3F	00	The target operation command was changed.
3F	01	The firmware/microcode was changed.
40	80	Flash checksum failed during diagnostic.
40	81	Buffer ram failed during diagnostics.
40	82	Sequencer chip failed during diagnostics.
40	83	Disk write/read failed during diagnostics.
40	88	Save RB failed during the reassignment of blocks ARRE/AWRE.
40	89	Read RB failed during the reassignment of blocks ARRE/AWRE.
40	8A	Flash configuration sector error encountered.
40	8B	Servo command time-out.
40	8C	Servo command failed.
40	90	Attempted to add illegal DList entry to GList during format.
40	91	Host buffer parity error occurred during host DMA.
40	92	SCSI checksum error during download.
40	93	Error in write to flash.
40	94	Internal write failure during reassign.
40	95	Track reformat failed during reassign.
40	96	Restore RB failed during reassign.
43	00	A message error occurred.
44	00	An internal controller error occurred.

Error code (hex)		
Byte 12	**Byte 13**	**Description**
45	00	An error occurred during a selection or a reselection.
47	00	A SCSI interface bus parity error occurred.
48	00	The initiator has detected an error.
49	00	The initiator received an invalid message from the drive.
4E	00	The drive attempted to perform overlapped commands.

Appendix C. Mode pages

Mode pages are groups of parameters stored by the drive. These parameters can be read using the Mode Sense command and changed using the Mode Select command. These commands are described in Sections 3.4.10 and 3.4.13.

This appendix contains the default parameters and the changeable parameters for the mode pages. The current parameters and the saved parameters are not shown.

Note. The default values contained in this appendix may differ from the default values actually contained in your drive. To determine the default values, use the Mode Sense command.

Mode page	Page code	Bytes	Contains changeable parameters
Error Recovery page	01_H	10	Yes
Disconnect/Reconnect page	02_H	14	Yes
Format Device page	03_H	22	No
Rigid Disc Geometry page	04_H	22	No
Verify error recovery page	07_H	10	Yes
Caching page (SCSI-3)	08_H	18	Yes
Control Mode page	$0A_H$	10	Yes
Notch page	$0C_H$	22	No
Cache control page	38_H	14	No
Soft ID page (Flash memory)	$3C_H$	1	Yes
Operating page	00_H	2 or 3	Yes

For all mode pages:

- If the changeable value is 0, the initiator *cannot* change the bit directly. If the changeable value is 1, the initiator *can* change the bit directly.

  For example, in the header below, the changeable value for the page code bits is 0, which means that the page code cannot be changed; the changeable value of the PS bit is one, which means that the PS bit can be changed.

- During the Mode Sense command, the PS (parameter savable) bit is 1, which means the mode page is saved on the disc. During the Mode Select command, you must set the PS bit to 0.

- An “X” means that the value of the bit cannot be specified. For example, the default value of bit 0 of byte 1 of page 00_H (the Operating page) cannot be specified because the bit can be either 1 or 0.

All mode pages contain a 2-byte header that contains the page code and the page length for that particular page. The header is shown below.

<table>
<tr><th rowspan="2">Bytes</th><th colspan="8">Bits</th></tr>
<tr><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td>PS</td><td colspan="7">Page code</td></tr>
<tr><td>changeable</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr>
<tr><td>1</td><td colspan="8">Page length</td></tr>
<tr><td>changeable</td><td colspan="8">00_H</td></tr>
</table>

Byte 0 During the Mode Sense command, the *PS* (parameter saveable) bit is 1, which means the mode page is saved on the disc. During the Mode Select command, you must set the PS bit to 0.

The *page code* is the unique code that identifies the page.

Byte 1 The *page length* is the length, in bytes, of the page.

C.1 Error Recovery page (01_H)

The Error Recovery page is shown below. This table summarizes the function, the default value and the changeability of each bit.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	PS (1)	Page code (01_H)						
1	Page length ($0A_H$)							
2	AWRE	ARRE	TB	RC	EER	PER	DTE	DCR
default	0	0	0	0	0	0	0	0
changeable	1	1	1	1	1	1	1	1
3 (default)	Read retry count (20_H)							
changeable	FF_H							
4 (default)	Correction span (16_H)							
changeable	00_H							
5 (default)	Head offset count (00_H)							
changeable	00_H							
6 (default)	Data strobe offset count (00_H)							
changeable	00_H							
7 (default)	Reserved (00_H)							
changeable	00_H							
8 (default)	Write retry count (20_H)							
changeable	FF_H							
9 (default)	Reserved (00_H)							
changeable	00_H							
10–11 (default)	Recovery time limit ($FFFF_H$)							
changeable	0000_H							

Byte 2 When the *automatic write reallocation enabled (AWRE)* bit is 1, the drive automatically reallocates bad blocks detected while writing to the disc. When the AWRE bit is 0, the drive does not perform automatic reallocation; instead, the drive reports a check condition status with a sense key of media error.

Note. The AWRE bit does not apply during the Format command.

When the *automatic read reallocation enabled (ARRE)* bit is 1, the drive automatically reallocates bad blocks detected while reading from the disc. When the ARRE bit is 0, the drive does not automatically reallocate bad blocks. Instead, a check condition status is reported with a sense key of media error.

The *transfer block (TB)* bit is not supported.

When the *read continuous (RC)* bit is 1, the drive sends all data without doing any corrections. This function supersedes other bits in this byte. When the RC bit is 0, the correction is performed according to the other bits in this byte.

The *enable early recovery (EER)* bit is not supported.

The *post error (PER)* bit is not supported.

The *disable transfer on error (DTE)* bit is not supported.

When the *disable correction (DCR)* bit is 1, the drive does not apply offline ECC to the data even if it can correct the data.

Byte 3 The *read retry count* field is the maximum number of times the drive attempts its recovery algorithms. The read retry count field has a range of 0 through 20_H.

Byte 4 The *correction span* is the size of the largest read data error, in bits, on which ECC correction is attempted. Longer errors are reported as nonrecoverable.

Byte 5 The *head offset count* is not implemented. Head offsets are performed as part of the drive’s retry algorithms.

Byte 6 The *data strobe offset count* is not implemented.

Byte 7 Reserved

Byte 8 The *write retry count* field contains the maximum number of times the drive attempts its recovery algorithms. This byte is a reflection of byte 3 and is not directly changeable. When mode selecting a change to this byte, the drive responds with a good status and command complete message.

Byte 9 Reserved

Bytes 10–11 The *recovery time limit* field always has a value of $FFFF_H$, which means that the recovery time is unlimited.

C.2 Disconnect/Reconnect page (02_H)

The Disconnect/Reconnect page is shown below. This table summarizes the function, the default value and the changeability of each bit.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	PS (1)	Page code (02_H)						
1	Page length ($0E_H$)							
2 (default)	Buffer full ratio (99_H)							
changeable	FF_H							
3 (default)	Buffer empty ratio (99_H)							
changeable	FF_H							
4–5 (default)	Bus inactivity limit (0000_H)							
changeable	0000_H							
6–7 (default)	Disconnect time limit (0000_H)							
changeable	0000_H							
8–9 (default)	Connect time limit (0000_H)							
changeable	0000_H							
10–11 (default)	Reserved (0000_H)							
changeable	0000_H							
12–15 (default)	Reserved (00000000_H)							
changeable	00000000_H							

Byte 2 The *buffer full ratio* field indicates, on Read commands, how full the drive's buffer is before reconnecting. The drive rounds up to the nearest whole logical block. This parameter is the numerator of a fraction that has 256 as its denominator.

Byte 3 The *buffer empty ratio* field indicates, on Write commands, how empty the drive's buffer is before reconnecting to fetch more data. The drive rounds up to the nearest whole logical block. This parameter is the numerator of a fraction that has 256 as its denominator.

Bytes 4–5 The *bus inactivity limit* field indicates the time, in 100-µsec increments, that the drive can assert the Busy signal without handshakes until it disconnects. The drive can round down to its nearest capable value. If the bus inactivity limit is 0000_H, the drive maintains the BSY– signal for 1 msec without handshakes.

Bytes 6–7 The *disconnect time limit* field indicates the minimum time, in 100-µsec increments, that the drive remains disconnected until it attempts to reconnect. A value of 0 indicates that the drive is allowed to reconnect immediately.

Bytes 8–9 The *connect time limit* field indicates the maximum time, in 100-µsec increments, that the drive should remain connected until it attempts to disconnect. The drive may round to its nearest capable value. A value of 0 means that the drive can remain connected indefinitely until it tries to disconnect.

Bytes 10-11 Reserved

Bytes 12-15 Reserved

C.3 Format Device page (03_H)

The Format Device page is shown below. This table summarizes the function, the default value and the changeability of each bit.

This page is sent only before the Format Unit command is sent. The drive parameters are updated immediately, but any changes between these current parameters and the existing media format do not take effect until after the Format Unit command is completed.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	PS (1)	Page code (03_H)						
1	Page length (16_H)							
2–3 (default)	Tracks per zone (0001_H)							
changeable	0000_H							
4–5 (default)	Alternate sectors per zone (0001_H)							
changeable	0000_H							
6–7 (default)	Alternate tracks per zone (0000_H)							
changeable	0000_H							
8–9 (default)	Alternate tracks per volume (0008_H)							
changeable	0000_H							
10–11 (default)	Sectors per track (0073_H)							
changeable	0000_H							

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
12–13 (default)	Data bytes per physical sector (0200_H)							
changeable	0000_H							
14–15 (default)	Interleave (0001_H)							
changeable	0000_H							
16–17 (default)	Track skew factor (00025_H)							
changeable	0000_H							
18–19 (default)	Cylinder skew factor ($002E_H$)							
changeable	0000_H							
20	SSEC	HSEC	RMB	SURF	Reserved			
default	1	0	0	0				
changeable	0	0	0	0	0	0	0	0
21–23 (default)	Reserved (000000_H)							
changeable	000000_H							

Bytes 2–3 The *tracks per zone* field indicates the number of tracks the drive allocates to each defect-management zone. Spare sectors or tracks are placed at the end of each defect-management zone. If each zone is treated as containing one track, the valid value for tracks per zone is 1. If each zone is treated as containing one cylinder, the valid value is equal to the number of read/write heads.

Bytes 4–5 The *alternate sectors per zone* field indicates the number of spare sectors to be reserved at the end of each defect-management zone. The drive defaults to one spare sector per zone. If each zone is treated as containing one track, the valid value for alternate sectors per zone is 1. If each zone is treated as containing one cylinder, the valid values are 1 through 3.

Bytes 6–7 The *alternate tracks per zon*e field indicates the number of spare tracks the drive reserves at the end of each defect-management zone. A value of 0 indicates that no spare tracks are reserved at the end of each zone for defect management.

Bytes 8–9 The *alternate tracks per volume* field indicates the number of spare tracks to be reserved at the end of the drive for defect management. The default is equal to twice the number of read/write heads.

Bytes 10–11 The *sectors per track* field indicates the number of physical sectors the drive allocates per track. The drive reports the average number of physical sectors per track because the number of sectors per track varies between the outer and inner tracks.

Bytes 12–13 The *data bytes per physical sector* field indicates the number of data bytes allocated per physical sector.

Bytes 14–15 The *interleave* field is the interleave value sent to the drive during the last Format Unit command. This field is valid only for Mode Sense commands. The drive ignores this field during Mode Select commands. The interleave is always 1:1.

Bytes 16–17 The *track skew factor* field indicates the number of physical sectors on the media between the last logical block of one track and the first logical block of the next sequential track of the same cylinder. The actual track skew factor that the drive uses is different for every zone. The default value is 0024_H, which is the track skew factor for the first zone.

Bytes 18–19 The *cylinder skew factor* field indicates the number of physical sectors between the last logical block of one cylinder and the first logical block of the next cylinder. The actual cylinder skew factor that the drive uses depends on the zone. The default value is $002D_H$, which is the cylinder skew factor for the first zone.

Byte 20 The *drive type* field bits are defined as follows:

The *soft sectoring (SSEC)* bit is set to 1. This bit is reported as not changeable. Although it can be set to satisfy system requirements, it does not affect drive performance.

The *hard sectoring (HSEC)* bit is set to 0. This bit is reported as not changeable. Although it can be set to satisfy system requirements, it does not affect drive performance.

The *removable media (RMB)* bit is always set to 0, indicating that the drive does not support removable media. This same bit is also returned in the Inquiry parameters.

The *surface map (SURF)* bit is set to 0, indicating that the drive allocates successive logical blocks to all sectors within a cylinder before allocating logical blocks to the next cylinder.

Bytes 21–23 Reserved

C.4 Rigid Disc Geometry page (04_H)

The Rigid Disc Geometry page is shown below. This table summarizes the function, the default value and the changeability of each bit.

Bytes	Bits 7	6	5	4	3	2	1	0
0	PS (1)	Page code (04_H)						
1	Page length (16_H)							
2–4	Number of cylinders ($12DA_H$)							
changeable	000000_H							
5	Number of heads (04_H)							
changeable	00_H							
6–8	Starting cylinder for write precompensation							
default	000000_H							
changeable	000000_H							
9–11	Starting cylinder for reduced write current							
default	000000_H							
changeable	000000_H							
12–13 (default)	Drive step rate (0000_H)							
changeable	0000_H							
14–16 (default)	Loading zone cylinder (000000_H)							
changeable	000000_H							
17 (default)	Reserved (00_H)							
changeable	00_H							
18 (default)	Rotational offset (00_H)							
changeable	00_H							
19 (default)	Reserved (00_H)							
changeable	00_H							
20–21	Media rotation rate							
default	1500_H							
changeable	0000_H							
22–23 (default)	Reserved (0000_H)							
changeable	0000_H							

Bytes 2–4 The *number of cylinders* field specifies the number of user-accessible cylinders, including two spare cylinders for defects. The drive uses the additional cylinders for storing parameters and defect lists or for diagnostic purposes. The number of cylinders is specified on page 1.

Byte 5 The *number of heads* field specifies the number of read/write heads on the drive. The number of heads is specified on page 1.

Bytes 6–16 The *starting cylinder for reduced write current*, *starting cylinder for reduced read current, drive step rate* and *loading zone cylinder* bytes are not used by the drive.

Byte 17 When the rotational position locking (RPL) bits are 00_B, the rotational position locking is changeable. When the RPL bits are 01_B, the drive automatically synchronizes its spindle with the synchronized master. When the RPL bits are 10_B or 11_B, the drive is the synchronized-spindle master. RPL is not supported.

Byte 18 The rotational offset is the rotational skew the drive uses when synchronized. The rotational skew is applied in the retarded direction (lagging the sync spindle master). A value of zero means no rotational offset is used. This feature is not supported.

Byte 19 Reserved

Bytes 20–21 The medium rotation rate is the spindle speed, which is specified on page 1.

Bytes 22–23 Reserved

C.5 Verify error recovery page (07_H)

The verify recovery page specifies the error recovery parameters the target shall use during the Verify command and the verify operation of the Write and Very commands.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	PS (1)	RSVD (0)	Page code (07_H)					
1	Parameter length ($0A_H$)							
2	Reserved				ERR	PER	DTE	DCR
default	0	0	0	0	0	0	0	0
changeable	0	0	0	0	1	1	1	1
3 (default)	Verify retry count (20_H)							
changeable	FF_H							
4 (default)	Verify correction span (16_H)							
changeable	00_H							
5 (default)	Reserved (00_H)							
changeable	00_H							
6 (default)	Reserved (00_H)							
changeable	00_H							
7 (default)	Reserved (00_H)							
changeable	00_H							
8 (default)	Reserved (00_H)							
changeable	00_H							
9 (default)	Reserved (00_H)							
changeable	00_H							
10 (default)	(MSB)	Verify recovery time limit (00_H)						
changeable	00_H							
11 (default)	Verify recovery time limit (00_H)							(LSB)
changeable	00_H							

Byte 2 When the *enable early recovery* (EER) bit is set to one, the target uses the most expedient form of error recovery first. This bit only applies to data error recovery and it does not affect positioning retries and the message error recovery procedures.

When EER is set to zero, the target shall use an error recovery procedure that minimizes the risk of misdetection or miscorrection.

When the *post error* (PER) bit is set to one, the target shall not report recovered errors. Error recovery procedures shall be performed within the limits established by the error recovery parameters.

When the *disable transfer or error* (DTE) bit is set to one, the target shall terminate the data phase upon detection of a recovered error. When DTE is set to zero, the target shall not terminate the data phase upon detection of a recovered error.

When the *disable correction* (DCR) bit is set to one, error correction codes shall not be used for data error recoverry. When DCR is zero, error correction codes can be used for data recovery.

Byte 3 The *verify retry count* field specifies the number of times the target attempts its recovery algorithm during a verify operation. If the verify retry count and the verify recovery time-limit are both specified, the one that requires the least amount of time for data error recovery actions shall have priority.

Byte 4 The *verify correction span* field specifies the size, in bits, of the largest burst data-error for which data-error correction may be attempted. If the target area does not implement this field, a value of zero is returned in Mode Sense data.

Byte 5–9 Reserved

Byte 10–11 The *verify recovery time limit* field specifies in increments of one millisecond the maximum time duration that the target drive shall use error recovery procedures to recover data for an individual logical block. This value may be rounded as defined in Parameter Rounding. If the verify retry count and the verify recovery time-limit are both specified, the one that requires the least amount of time for data error recovery actions shall have priority.

Note. To disable all types of correction and retries, the initiator should set the EER bit to zero, the PER, DTE and DCR bits to one and the number of retries and recovery time limit to zero.

C.6 Caching page (08_H)

The drive uses read look-ahead, read caching and write caching to improve seek times and performance.

C.6.1 Read look-ahead and caching

The drive uses an algorithm that improves seek performance by reading the next logical sectors after the last requested sector. These unrequested sectors are read into a buffer and are ready to be transmitted to the host before they are requested. Because these sectors are read before they are requested, access read time for the sectors is virtually eliminated. This process is called either read look-ahead or read caching.

Read look-ahead and read caching are similar algorithms. Read look-ahead occurs when a Read command requests more data than can be contained in one buffer segment. Read caching occurs when a Read command requests less data than can be contained in one buffer segment.

The buffer used for read look-ahead and caching can be divided into segments as shown in the following table. To change the number of segments, use byte 13 of the Caching page, which is described in Appendix C.6.3. The default is one 128-Kbyte segment.

Number of segments	Size of segment (in Kbytes)
1	128
2	64
4	32

When the buffer is divided into multiple segments, each segment functions as an independent buffer, causing dramatically increased performance in multitasking and multiuser environments.

C.6.2 Write caching and write merging

Write caching. The drive uses the write segment to store Write commands and data. After the drive caches the commands and data, it is immediately ready to process new commands. The drive writes the data to the disc at its next convenient opportunity.

Write merging. The drive accepts contiguous Write commands and executes them sequentially as one command.

C.6.3 Caching page description

The Caching page is shown below. This table summarizes the function, the default value and the changeability of each bit.

<table>
<tr><th rowspan="2">Bytes</th><th colspan="8">Bits</th></tr>
<tr><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td>PS (1)</td><td colspan="7">Page code (08H)</td></tr>
<tr><td>1</td><td colspan="8">Page length (12H)</td></tr>
<tr><td>2</td><td>IC</td><td>ABPF</td><td>CAP</td><td>DISC</td><td>SIZE</td><td>WCE</td><td>MF</td><td>RCD</td></tr>
<tr><td>default</td><td>1</td><td>0</td><td>0</td><td>1</td><td>0</td><td>1</td><td>0</td><td>0</td></tr>
<tr><td>changeable</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>1</td><td>1</td></tr>
<tr><td>3</td><td colspan="4">Demand read retention priority</td><td colspan="4">Write retention priority</td></tr>
<tr><td>default</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr>
<tr><td>changeable</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr>
<tr><td>4–5 (default)</td><td colspan="8">Disable prefetch transfer length (FFFFH)</td></tr>
<tr><td>changeable</td><td colspan="8">0000H</td></tr>
<tr><td>6–7 (default)</td><td colspan="8">Minimum prefetch (0000H)</td></tr>
<tr><td>changeable</td><td colspan="8">FFFFH</td></tr>
<tr><td>8–9 (default)</td><td colspan="8">Maximum prefetch (FFFFH)</td></tr>
<tr><td>changeable</td><td colspan="8">FFFFH</td></tr>
<tr><td>10–11 (default)</td><td colspan="8">Maximum prefetch ceiling (FFFFH)</td></tr>
<tr><td>changeable</td><td colspan="8">FFFFH</td></tr>
<tr><td>12</td><td>FSW</td><td>Rsrvd</td><td>DRA</td><td colspan="5">Reserved</td></tr>
<tr><td>default</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr>
<tr><td>changeable</td><td>0</td><td>0</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td></tr>
<tr><td>13</td><td colspan="8">Number of cache segments</td></tr>
<tr><td>default</td><td colspan="8">01H</td></tr>
<tr><td>changeable</td><td colspan="8">FFH</td></tr>
<tr><td>14–15 (default)</td><td colspan="8">Cache segment size (0000H)</td></tr>
<tr><td>changeable</td><td colspan="8">0000H</td></tr>
<tr><td>16 (default)</td><td colspan="8">Reserved (00H)</td></tr>
<tr><td>changeable</td><td colspan="8">00H</td></tr>
<tr><td>17–19 (default)</td><td colspan="8">Noncache segment size (000000H)</td></tr>
<tr><td>changeable</td><td colspan="8">000000H</td></tr>
</table>

Byte 2 The *initiator control (IC)* bit is not supported.

When the *abort prefetch (ABPF)* bit is 0, the drive controls completion of prefetch. See the description for the DISC bit, below. This is the default value and it is not changeable.

The *caching analysis permitted (CAP)* bit is not supported.

When the *discontinuity (DISC)* bit is 1, the drive may prefetch across cylinder boundaries, where head seeks consume additional processing time. This is the default value and it is not changeable.

The *size enable (SIZE)* bit is not supported.

When the *write cache enable (WCE)* bit is 0, the drive returns a good status for a Write command after successfully writing all the data to the media. When the WCE bit is 1, the drive returns a good status for a Write command after successfully receiving the data and before writing it to the media.

When the *multiplication factor (MF)* bit is 0, the drive interprets the *minimum prefetch* and *maximum prefetch* fields as the number of logical blocks to be prefetched. When the MF bit is 1, the drive interprets the minimum prefetch and maximum prefetch fields in terms of a number which, when multiplied by the transfer length of the current command, yields the number of logical blocks to be prefetched.

When the *read cache disable (RCD)* bit is 0, the drive may return data requested by a Read command by accessing either the cache or the media. If the RCD bit is 1, the cache is not used.

Byte 3 The *demand read retention priority* field is not used. The initiator cannot assign any special retention priority to the drive.

The *write retention priority* field is not used. The initiator cannot assign any special retention priority to the drive.

Bytes 4–5 The *disable prefetch transfer length* always has a value of $FFFF_H$, which means that the drive attempts an anticipatory prefetch for all Read commands.

Bytes 6–7 The *minimum prefetch* field specifies the minimum number of blocks the drive prefetches, regardless of the delays it may cause in executing subsequent pending commands. When the minimum prefetch field contains 0, the drive terminates prefetching whenever another command is ready to be executed. If the minimum prefetch equals the maximum prefetch, the drive prefetches the same number of blocks regardless of whether there are commands pending.

Bytes 8–9 The *maximum prefetch* field specifies the maximum number of blocks the drive prefetches during a Read command if there are no other commands pending. The maximum prefetch field represents the maximum amount of data to prefetch into the cache for any single Read command.

Bytes 10–11 The *maximum prefetch ceiling* field should be equal to the maximum prefetch field. The maximum prefetch ceiling and maximum prefetch fields are the same if the MF bit is 0.

Byte 12 The *force sequential write (FSW)* bit set to one indicates that multiple block writes are to be transferred over the SCSI bus and written to the media in an ascending, sequential, logical block order.

When the *disable read-ahead (DRA)* bit is 1, the drive does not read into the buffer any logical blocks beyond the addressed logical blocks. When the DRA bit equals 0, the drive can continue reading logical blocks into the buffer beyond the addressed logical blocks.

Byte 13 The *number of cache segments* field determines the number of segments into which the cache should be divided. Valid values are 1, 2, 4, 8, 16 and 32.

Bytes 14–15 The *cache segment size* field indicates the segment size in bytes. The cache segment size field is valid only when the SIZE bit is 1.

Byte 16 Reserved

Bytes 17–19 The *noncache segment size* field always contains zeros. This means that the entire buffer is available for caching.

C.7 Control Mode page ($0A_H$)

The Control Mode page is shown below. This table summarizes the function, the default value and the changeability of each bit.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	PS (1)	Page code ($0A_H$)						
1	Page length ($0A_H$)							
2	Reserved							RLEC
default	0	0	0	0	0	0	0	0
changeable	0	0	0	0	0	0	0	0
3	Queue algorithm modifier				Reserved		QErr	DQue
default	0	0	0	1	0	0	0	0
changeable	1	1	1	1	0	0	0	1
4	EECA	Reserved				RAENP	UAAENP	EAENP
default	0	0	0	0	0	0	0	0
changeable	00_H							
5 (default)	Reserved (00_H)							
changeable	00_H							
6–7 (default)	Ready AEN hold-off period (0000_H)							
changeable	0000_H							
8–9 (default)	Busy timeout period ($FFFF_H$)							
changeable	0000_H							
10–11 (default)	Reserved (0000_H)							
changeable	0000_H							

Byte 2 The RLEC bit is not implemented.

Byte 3 The queue algorithm modifier field is only effective if the disable queuing bit is zero. When bit 4 in the queue algorithm modifier field is one, the drive may use tagged command queuing to change the order in which it executes commands. When bit 4 in the queue algorithm modifier field is zero, the drive always executes commands according to the order indicated by the simple queue tag.

When the *disable queuing (DQue)* bit is zero, tagged command queuing is enabled. When the *DQue* bit is one, tagged command queuing is disabled.

Byte 4 Not implemented

Byte 5 Reserved

Bytes 6–7 Not implemented

Bytes 8–9 The busy time out period field contains the maximum possible value, which means that the drive can remain busy an unlimited amount of time.

Bytes 10–11 Reserved

C.8 Notch page ($0C_H$)

The Notch page contains parameters that describe the notches. The table below summarizes the function, default value and the changeability of each bit.

The drive uses Zone Bit Recording, which means that the outer cylinders of the disc contain more logical blocks than the inner cylinders. The cylinders are organized into groups called zones or notches. Every logical block is part of a notch. Notches do not overlap.

Bytes	Bits							
	7	6	5	4	3	2	1	0
0	PS (1)	Page code ($0C_H$)						
1	Page length (16_H)							
2	ND	LPN	Reserved					
default	1	0	0	0	0	0	0	0
changeable	0	0	0	0	0	0	0	0
3 (default)	Reserved (00_H)							
changeable	00_H							
4–5 (default)	Maximum number of notches (0013_H)							
changeable	0000_H							
6–7 (default)	Active notch (0000_H)							
changeable	$FFFF_H$							
8–11 (default)	Starting boundary (00000000_H)							
changeable	00000000_H							
12–15 (default)	Ending boundary ($0012D903_H$)							
changeable	00000000_H							
16–23 (default)	Pages notched (0000000000000008_H)							
changeable	0000000000000000_H							

Byte 2 The *notched drive (ND)* bit is always 1.The notch recording densities are not the same size. This page defines the starting and ending boundaries for each supported, active notch value.

The *logical or physical notch (LPN)* bit is 0. The notch boundaries are based on the physical parameters of the logical unit. The cylinder is most significant; the head is least significant.

Byte 3 Reserved

Bytes 4–5 The *maximum number of notches* field indicates the maximum number of notches supported by the drive.

Bytes 6–7 The *active notch* field identifies the notch to which this and all future Mode Select and Mode Sense commands refer until the active notch is changed by a later Mode Select command. The value of the active notch field must be greater than or equal to 0 and less than or equal to the maximum number of notches. An active notch value of 0 means that current and future Mode Select and Mode Sense commands refer to the parameters that apply for all notches.

Bytes 8–11 The *starting boundary* field indicates the beginning of the active notch if the active notch is not 0, or the starting boundary of the logical unit if the active notch is 0. This field is ignored by the Mode Select command.

When the LPN bit is 0, the three most significant bytes represent the cylinder number and the least significant byte represents the head number.

Bytes 12–15 The *ending boundary* field indicates the end of the active notch if the active notch is not 0, or the end of the logical unit if the active notch is 0. The default is equal to the end of zone 1.

When the LPN bit is 0, the three most significant bytes represent the cylinder number and the least significant byte represents the head number.

Bytes 16–23 The *pages notched* field contains a bit map of the mode page codes that indicates which pages may contain different parameters for each notch. When a bit is 1, the corresponding mode page can contain different parameters for each notch. When a bit is 0, the corresponding mode page contains the same parameters for all the notches. The most significant bit of this field corresponds to page code $3F_H$ and the least significant bit corresponds to page code 00_H.

C.9 Cache Control page (38_H)

The Cache Control page is shown below. This table summarizes the function, the default value and the changeability of each bit.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	PS (1)	Page code (38_H)						
1	Page length ($0E_H$)							
2	Rsrvd	WIE	Rsrvd	CE	Cache table size			
default	0	0	0	1	0	0	0	1
changeable	(00_H)							
3 (default)	Prefetch threshold (00_H)							
changeable	00_H							
4 (default)	Maximum prefetch (FF_H)							
changeable	00_H							
5 (default)	Maximum prefetch multiplier (00_H)							
changeable	00_H							
6 (default)	Minimum prefetch (00_H)							
changeable	00_H							
7 (default)	Minimum prefetch multiplier (00_H)							
changeable	00_H							
8–15 (default)	Reserved (0000000000000000_H)							
changeable	0000000000000000_H							

Byte 2 The *cache enable (CE)* bit is always the inverse of the *RCD* bit in Mode page 08_H.

The *write index enable (WIE)* bit controls the creation of cache data on Write commands. If bit 6 is 0, the next command treats the cache area as empty.

The *cache table size* field contains the same values as Mode page 08_H, byte 13, bits 3 through 0.

Byte 3 The *prefetch threshold* is not implemented. The drive reads until the buffer is full upon receipt of a Read command.

Byte 4 The *maximum prefetch* field always contains the same value as byte 9 of the Caching page. The initiator cannot change this byte directly.

Byte 5 The *maximum prefetch multiplier* field always contains the same value as byte 9 of the Caching page, which is described in Appendix C.6.3. The initiator cannot change this byte directly.

Byte 6 The *minimum prefetch* field always contains the same value as byte 7 of the Caching page. The initiator cannot change this byte directly.

Byte 7 The *minimum prefetch multiplier* field always contains the same value as byte 7 of the Caching page. The initiator cannot change this byte directly.

Byte 8–15 Reserved

C.10 Soft ID page (Flash memory) ($3C_H$)

The Soft ID page is shown below. This table summarizes the function, the default value and the changeability of each bit. This page is saved in flash memory that has a life span of 10,000 writes.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	PS (1)	Page code ($3C_H$)						
1	Page length (01_H)							
2	Soft ID	Soft Parity	Param enable	Soft remote	Remote S/S	ID 2	ID 1	ID 0
default	0	0	0	0	0	0	0	0
change-able	1	1	1	1	1	1	1	1

Byte 2 When the *soft ID* bit is 0, the drive ignores ID0, ID1 and ID2 and uses the SCSI ID jumpers to determine the SCSI ID. When the soft ID bit is 1, the drive ignores the SCSI ID jumpers and uses ID0, ID1 and ID2 to determine the SCSI ID.

When the *soft parity* bit is 0, the drive uses the parity jumper settings to determine whether the drive uses parity. When the soft parity bit is 1, the drive ignores the parity jumper settings.

When the *soft remote* bit is 0, the drive uses the remote start jumper setting to determine whether remote start is implemented. When the soft remote bit is 1, the drive ignores the jumpers and uses the remote S/S bit to determine whether remote start is implemented.

When the *remote S/S* bit is 0, the drive spins up after a delay specified by the spinup delay field (byte 4 of the Operating page, 00_H). When the remote S/S bit is 1, the drive spins up when it receives the Start/Stop Unit command. This bit is only valid if the soft remote bit is 1.

The *ID0*, *ID1* and *ID2* bits are the SCSI ID bits. These bits are only valid when the soft ID bit is 1.

When the *param enable* bit is 0, the drive does not check parity. When the param enable bit is 1, the drive checks parity. This bit is only valid if the soft parity bit is 1.

C.11 Operating page (Flash memory) (00_H)

The Operating page is shown in the table below. This table shows the function, the default value and the changeability of each bit.

The drive accepts an Operating page of two lengths: two bytes or three bytes. If the length is two bytes, then byte 4, the *spinup delay* field, is not written and is assumed to be unchanged.

In addition to being saved on the media, this vendor-unique page is saved in flash memory that has a life span of 10,000 writes.

Bytes	**Bits**							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	PS (1)	Page code (00_H)						
1 (default)	Page length (02_H or 03_H)							
2	Usage	SSM	Rsrvd	ATOFF	IQFM	Reserved		
default	1	0	0	0	1	0	0	0
changeable	1	0	0	1	1	0	0	0
3	Rsrvd	Device type qualifier						
default	(00_H)							
changeable	00_H							
4 (default)	Spinup delay (00_H)							
changeable	(FF_H)							

Byte 2 When the *usage* bit is 1, a warning message is enabled. When the write life span of the flash memory is exceeded, a warning message is generated. See additional sense error code C2 in Appendix B.3. When the usage bit is 0, the warning message is disabled. If requested, the flash mem-

ory data and the write counter are updated even after the write life span is exceeded, but the integrity of the data cannot be assured.

When the *synchronous select mode (SSM)* bit is 0, the drive does not send a synchronous data transfer message unless the initiator has already issued a synchronous data transfer message. When the *SSM* bit is 1, the drive can send a synchronous data transfer message, even when the initiator has not sent a synchronous data transfer message.

When the *disable unit attention (ATOFF)* bit is 0, the drive generates a unit attention condition during power up. When the *disable unit attention (ATOFF)* bit is 1, the drive does not generate a unit attention condition during power up.

When the *Inquiry Queue Follow Mode (IQFM)* bit is 0, the *CmdQue* bit in byte 7, bit 1 of the Inquiry data reports that the drive supports Tagged Command Queuing. When the *IQFM* bit is 1, the *CmdQue* bit in byte 7, bit 1 of the Inquiry data follows the state of the *DQue* bit Mode page $0A_H$, byte 3, bit 0.

Byte 3 The *device type qualifier* field is not supported.

Byte 4 The Spinup delay field controls the drive when it is not in the remote start mode. When the value is 00_H, the drive spins up without delay. When the value is FF_H, the drive delays spinup to a duration whose value in seconds equals five times the drive's SCSI ID address. When the value is between 01_H and FE_H, the drive delays the spinup for the corresponding decimal duration in seconds.

Appendix D. Inquiry data

When the initiator issues an Inquiry command, the drive returns either inquiry data or vital product data, depending on the value in the EVPD bit in byte 1 of the Inquiry command descriptor block.

Both types of data are discussed in this appendix. The Inquiry command is described in Section 3.4.9.

D.1 Inquiry data

When the initiator issues an Inquiry command, and the EVPD bit in byte 1 of the Inquiry command descriptor block is 0, the drive returns the following data. If the EVPD bit in byte 1 of the Inquiry command descriptor block is 1, see Appendix D.2.

Bytes	Bits							
	7	6	5	4	3	2	1	0
0	Peripheral qualifier			Peripheral device type				
	0	0	0	0	0	0	0	0
1	RMB	Device type qualifier						
	0	0	0	0	0	0	0	0
2	ISO version		ECMA version			ANSI version		
	0	0	0	0	0	0	1	0
3	AENC	TrmIOP	Reserved		Response data format			
	0	0			0	0	1	0
4	Additional length ($8F_H$)							
5–6	Reserved (00_H)							
7	RelAdr	Wbus32	Wbus16	Sync	Linked	0	CmdQue	SftRe
8–15	Vendor identification							
16–31	Product identification							
32–35	Product revision level							
36–43	Drive serial number							
44–95	Reserved							
96–143	Copyright notice							
144–147	Servo PROM part number							

Byte 0 The *peripheral qualifier* field contains a zero, which means that the drive is currently connected to the logical unit that is issuing the Inquiry command.

The *peripheral device type* field contains a zero, which means that the drive is a direct-access device.

Byte 1 The *RMB* bit is 0, which means the discs are not removable.

The *device type qualifier is user programmable.*

Byte 2 The *ISO version* field contains a zero, which means that we do not claim compliance with ISO 9316.

The *EMCA version* field contains a zero, which means that we do not claim compliance with EMCA-111.

The *ANSI version* field contains a two, which means that the drive complies with ANSI SCSI-2 standard X3.131-199x.

Byte 3 The *asynchronous event notification (AENC)* bit is zero, which means that the drive does not support asynchronous event notification.

The *terminate I/O process (TrmIOP)* bit is zero, which means that the drive does not support the terminate I/O process message.

The *response data format* field contains a two, which means that the inquiry data is in standard SCSI-2 format.

Byte 4 The *additional length* field contains 143, which is the number of bytes contained in the inquiry data beyond byte 4. This value represents a total inquiry data length of 148 bytes. If the allocation length in the CDB of the Inquiry command is less than 148, the inquiry data is truncated, but the additional length does not change.

Bytes 5–6 Reserved

Byte 7 The *RelAdr* bit is zero, which means that the drive does not support the relative addressing mode.

The *WBUS32* bit is zero, which means that the drive does not support 32-bit data transfers.

The *WBUS16* bit is zero, which means that the drive does not support 16-bit data transfers.

The *SYNC* bit is one, which means that the drive supports synchronous data transfer.

The *Linked* bit is zero, which means that the drive does not support linked commands.

The *CmdQue* bit is one, which means that the drive supports tagged command queuing.

The *Soft Re* bit is zero, which means that the drive responds to a soft reset with a hard reset.

Bytes 8–15 The *vendor identification* field contains SEAGATE in ASCII text.

Bytes 16–31 The *product identification* field contains the model number of the drive in ASCII text.

Bytes 32–35 The *product revision level* field contains the last four digits of the firmware release number in ASCII text.

Bytes 36–43 The *drive serial number* field contains the serial number of the drive in ASCII text.

Bytes 44–95 These bytes are reserved; they contain only zeros.

Bytes 96–143 The *copyright notice* field contains the following ASCII string: “ Copyright (c) 1993 Seagate. All rights reserved.”

Bytes 144–147 This is the *servo PROM part number* field.

D.2 Vital product data pages

When the initiator issues an Inquiry command, and the EVPD bit in byte 1 of the Inquiry command descriptor block is 1, the drive returns vital product data pages. If the EVPD bit in byte 1 of the Inquiry command descriptor block is 0, see Appendix D.1.

All vital product data pages contain a 4-byte header, shown below.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Peripheral qualifier			Peripheral device type				
1	Page code							
2	Reserved (00_H)							
3	Page length							

Byte 0 The *peripheral qualifier* field contains zero, which means that the drive is currently connected to the logical unit issuing the Inquiry command.

The *peripheral device type* field contains zero, which means that the drive is a direct-access device.

Byte 1 The *page code* field contains the same value contained in the page code field in byte 2 of the Inquiry command descriptor block.

If the page code field contains any of the page codes shown in the table below, the drive returns the corresponding page. The available page codes are:

Page code	Description
00_H	Supported vital product data pages
80_H	Unit serial number page
81_H	Implemented operating definitions page
$C0_H$	Firmware numbers page (vendor-unique)
$C1_H$	Date code page (vendor-unique)
$C2_H$	Jumper settings page (vendor-unique)

Byte 2 Reserved

Byte 3 The *page length* field contains the length of the supported page list.

D.2.1 Unit Serial Number page (80_H)

The Unit Serial Number page is shown below. The table summarizes the function and the default value of each bit.

Bytes	Bits							
	7	6	5	4	3	2	1	0
0	Peripheral qualifier			Peripheral device type				
1	Page code (80_H)							
2	Reserved (00_H)							
3	Page length (08_H)							
4–11	Product serial number							

Bytes 4–11 The product serial number field contains the serial number for the drive in ASCII. If the drive does not return the serial number, it returns spaces (20_H).

D.2.2 Implemented Operating Definition page (81_H)

The Implemented Operating Definition page is shown below. The table summarizes the function and the default value of each bit.

Bytes	Bits							
	7	6	5	4	3	2	1	0
0	Peripheral qualifier			Peripheral device type				
1	Page code (81_H)							
2	Reserved (00_H)							
3	Page length (03_H)							
4	SAVIMP	Current operating definition						
	0							
5	SAVIMP	Default operating definition						
	0							
6	SAVIMP	Supported operating definition						
	0							

Byte 4 The current operating definition field contains the value of the current operating definition.

Byte 5 The SAVIMP bit is always zero; therefore, the current operating definition parameter cannot be saved. If the SAVIMP bit is one, the current operating parameter can be saved.

The default operating definition field contains the value of the default operating definition. If no operating definition is saved, the drive uses the default operating definition.

Bytes 6–8 If the SAVIMP bit is zero, the default definition parameter cannot be saved. If the SAVIMP bit is one, the default definition parameter can be saved.

The supported operating definition field contains the value of the supported operating definition. If no supported operating definition is saved, the drive uses the default operating definition.

D.2.3 Firmware Numbers page ($C0_H$)

The Firmware Numbers page is shown below. The table summarizes the function and default value of each bit.

Bytes	Bits							
	7	**6**	**5**	**4**	**3**	**2**	**1**	**0**
0	Peripheral qualifier			Peripheral device type				
1	Page code ($C0_H$)							
2	Reserved (00_H)							
3	Page length (0CH)							
4–7	Controller firmware number							
8–11	Boot firmware number							
12–15	Servo firmware number							

Bytes 4–7 The controller firmware number field contains the controller firmware number in ASCII text.

Bytes 8–11 The boot firmware number field contains the boot firmware number in ASCII text.

Bytes 12–15 The servo firmware number field contains the servo firmware in ASCII text.

D.2.4 Date Code page ($C1_H$)

The Date Code page is shown below. The table summarizes the function and the default value of each bit.

Bytes	Bits 7	6	5	4	3	2	1	0
0	Peripheral qualifier			Peripheral device type				
1	Page code ($C1_H$)							
2	Reserved (00_H)							
3	Page length (03_H)							
4	Year							
5–6	Week							

Bytes 4 The *year* field contains the year, in ASCII, that the firmware was released.

Bytes 5–6 The *week* field contains the week, in ASCII, that the firmware was released.

D.2.5 Jumper Settings page ($C2_H$)

The Jumper Settings page is shown below. The table summarizes the function and the default value of each bit.

Bytes	Bits 7	6	5	4	3	2	1	0
0	Peripheral qualifier			Peripheral device type				
1	Page code ($C2_H$)							
2	Reserved (00_H)							
3	Page length (01_H)							
4	Rsrvd	MS	Rsrvd	PE	Rsrvd	SCSI ID		

Byte 4 If the *motor start (MS)* bit is 1, the remote start enable jumper is installed on pins 15 and 16 of the options jumper block. If the MS bit is 0, the remote start enable jumper is not installed.

If the *parity enable (PE)* bit is 1, the parity enable jumper is installed on pins 17 and 18 of the options jumper block. If the PE bit is 0, the parity enable jumper is not installed.

SCSI ID is the SCSI ID of the drive.

Appendix E. Timing diagrams

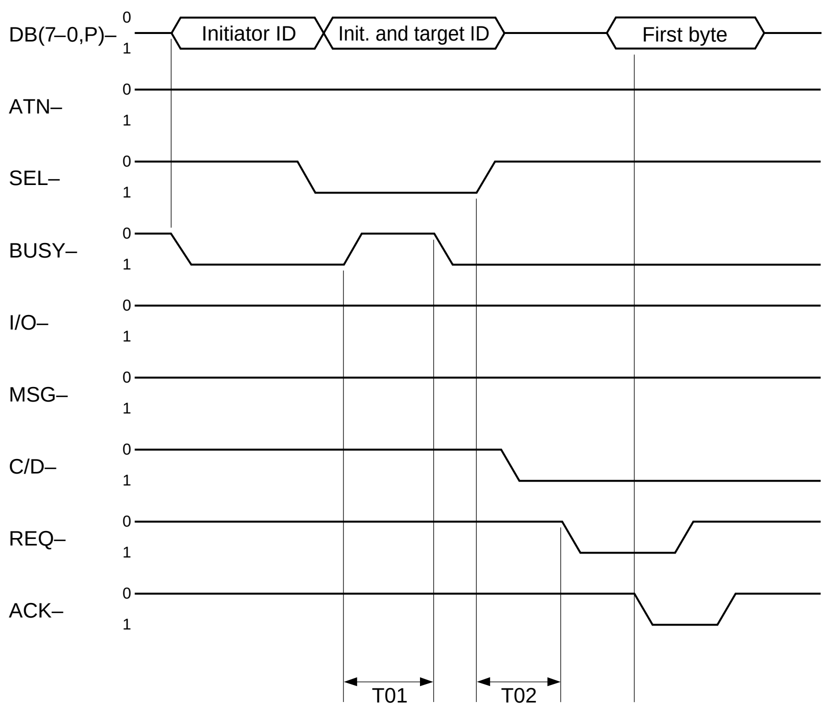

Figure 6. Arbitration, selection (without ATN) and command

Description	Symbol	Typical	Max
Target select time (without arbitration)	T00	<80 µsec	<250 msec
Target select time (with arbitration)	T01	<90 µsec	<250 msec
Target select to command	T02	<150 µsec	—

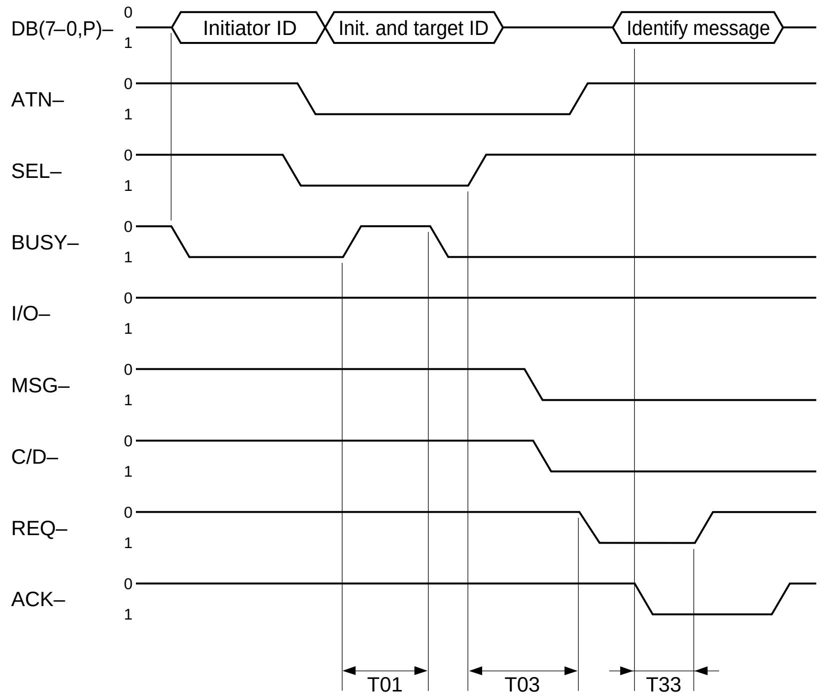

Figure 7. Arbitration, selection (with ATN) and message out

Description	Symbol	Typical	Max
Target select time (without arbitration)	T00	<1.0 µsec	<250 µsec
Target select time (with arbitration)	T01	<55 µsec	<250 µsec
Target select to message out	T03	<125 µsec	—
Message out byte transfer	T33	<0.1 µsec	0.15 µsec

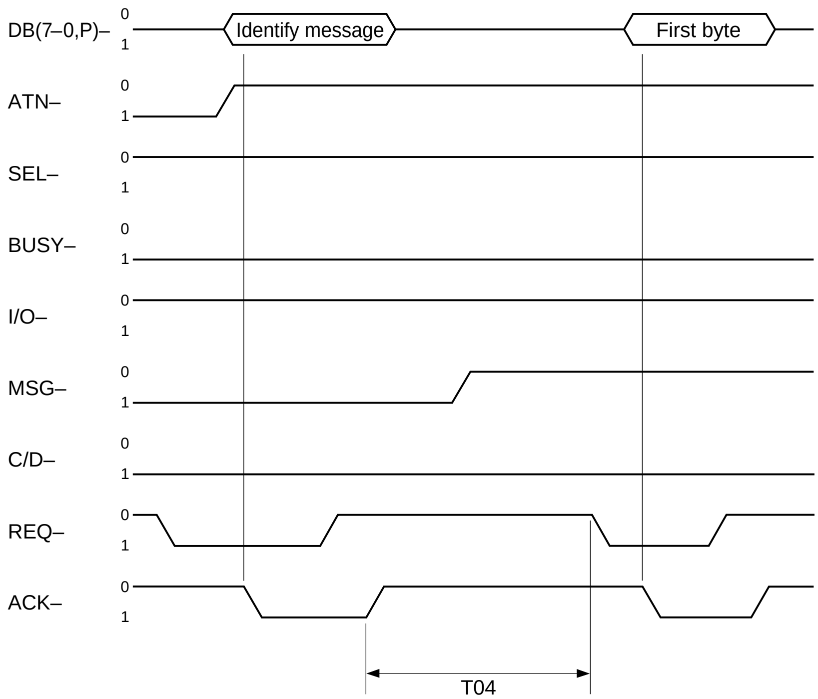

Figure 8. Identify message out to command

Description	Symbol	Typical
Identify message to command	T04	<150 μsec

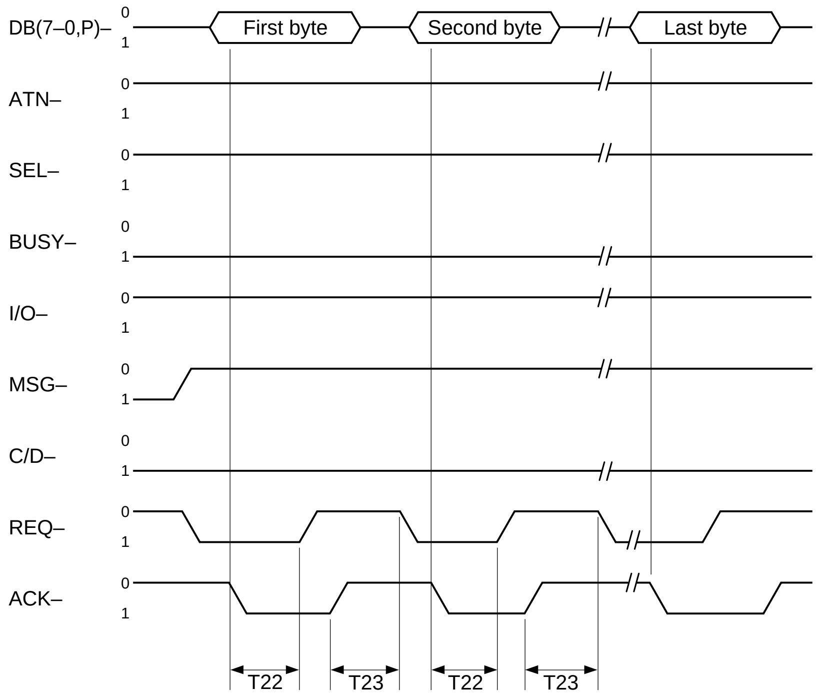

Figure 9. Command descriptor block transfer

Description	Symbol	Typical	Max
Command byte transfer	T22	<0.08 µsec	0.15 µsec
Next command byte access *	T23	<6.5 µsec	1.0 µsec

***** T23 is used, except for byte 7 of a 10-byte CDB. A 6-byte CDB requires less than 5 µsec for five T23 occurrences. A 10-byte CDB requires less than 110 µsec for nine occurrences.

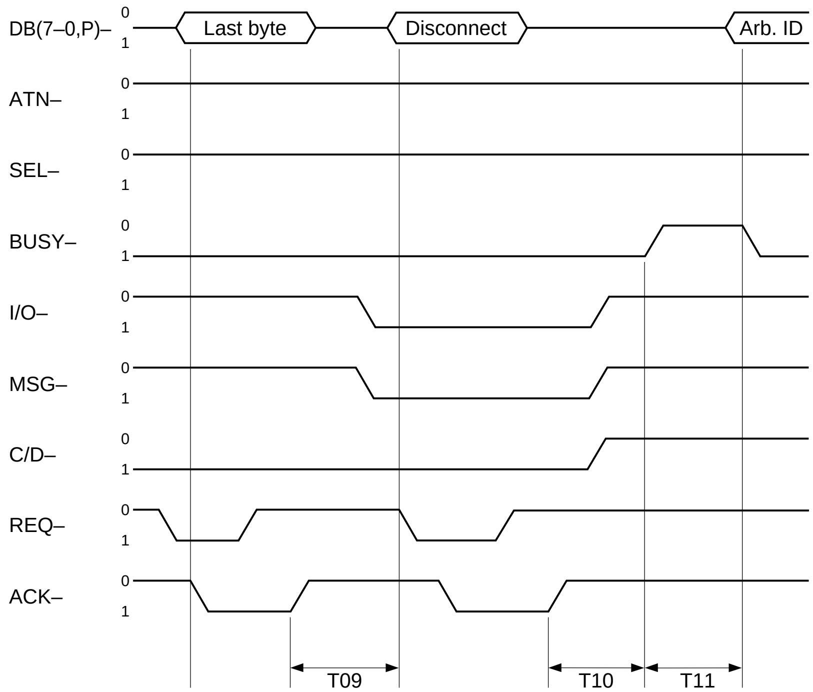

Figure 10. Command, status, command complete message and bus free

Description	Symbol	Typical	Max
Command to status	T05	Command-dependent	
Status to command complete message	T18	<150 µsec	—
Command complete message to bus free	T19	<100 µsec	—
Message in byte transfer	T32	<0.1 µsec	0.15 µsec
Status byte transfer	T34	<0.1 µsec	0.15 µsec

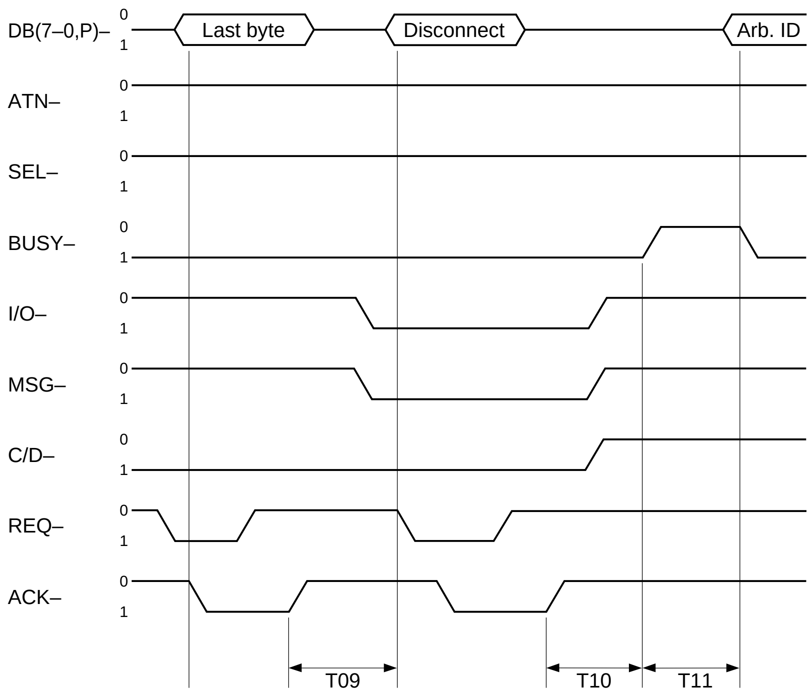

Figure 11. Last command byte, disconnect message, bus free and reselection

Description	Symbol	Typical	Max
Command to disconnect message	T09	Command-dependent	
Disconnect message to bus free	T10	<100 µsec	—
Disconnect to arbitration (for reselect). Measures disconnected command overhead.	T11 *	Command-dependent	

***** When measuring T11, no other device can be contending for the SCSI bus.

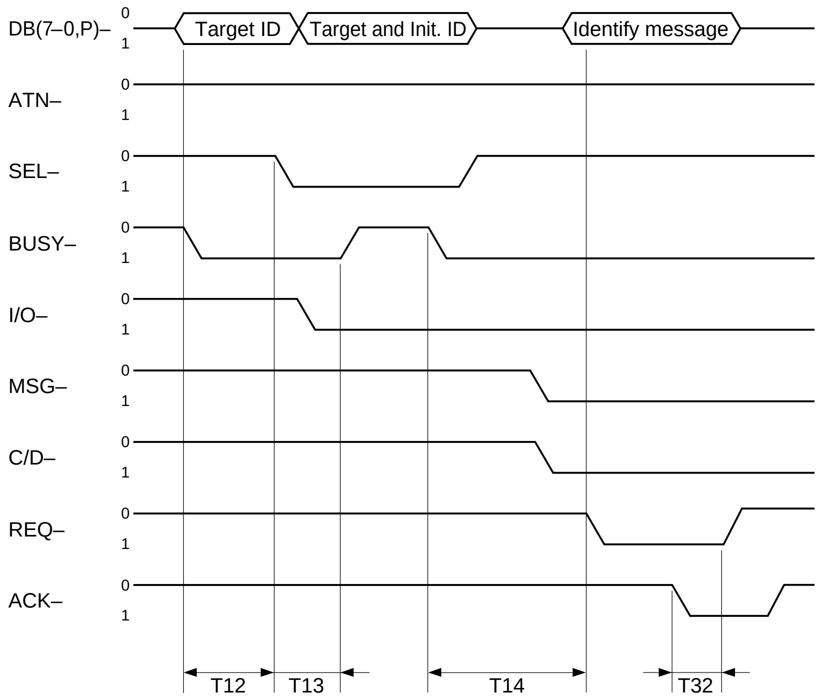

Figure 12. Arbitration, reselection and message in

Description	Symbol	Typical	Max
Target wins arbitration (for reselect)	T12	<6 µsec	—
Arbitration to reselect	T13	<5 µsec	—
Reselect to identify message in	T14	<150 µsec	—
Message in byte transfer	T32	<0.1 µsec	0.15 µsec

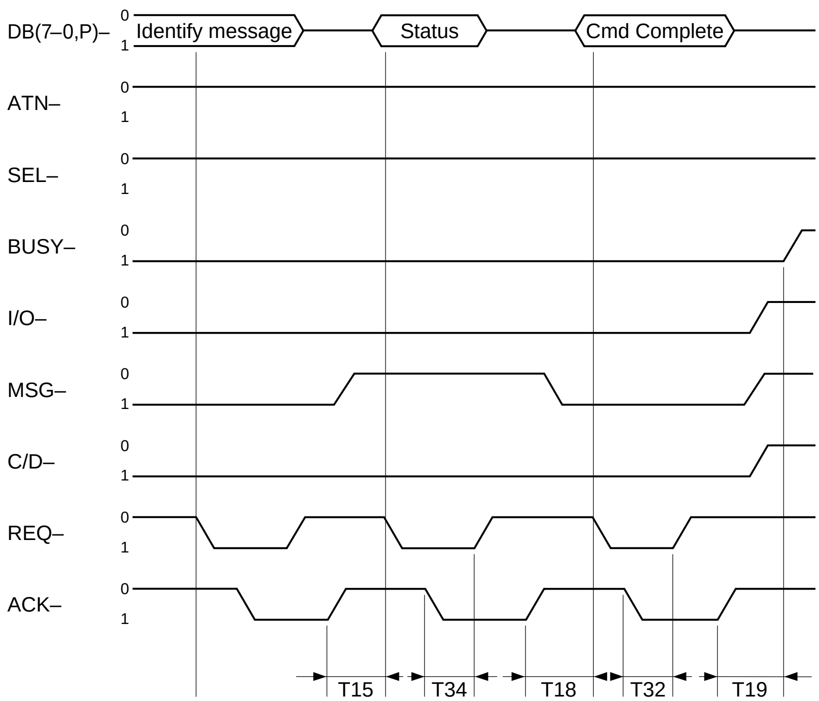

Figure 13. Reselection, status, command complete and bus free

Description	Symbol	Typical	Max
Reselect identify message to status	T15	<150 µsec	—
Status to command complete message	T18	<150 µsec	—
Command complete message to bus free	T19	<100 µsec	—
Message in byte transfer	T32	<0.1 µsec	0.15 µsec
Status byte transfer	T34	<0.1 µsec	0.15 µsec

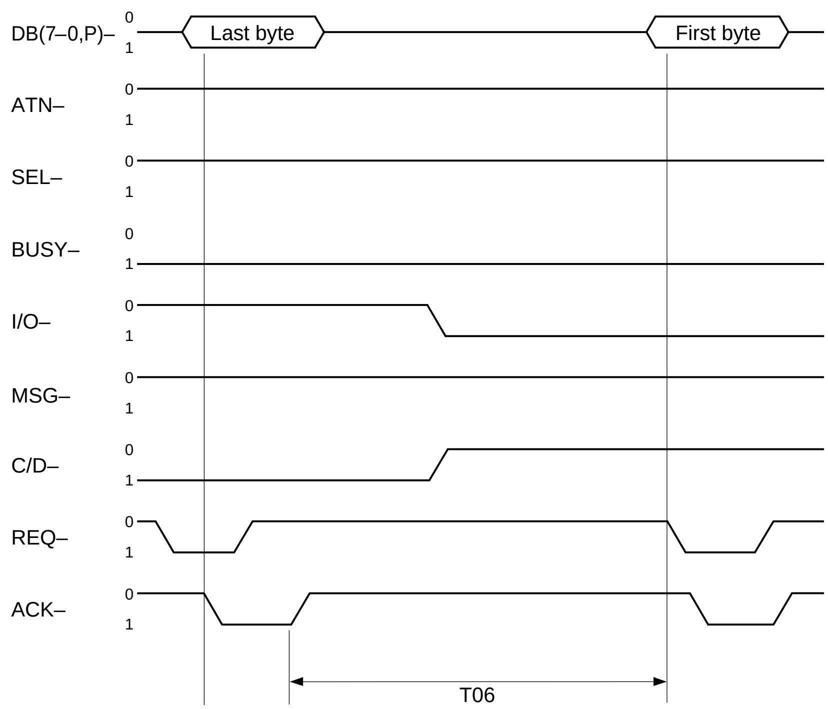

Figure 14. Last command byte to data in

Description	Symbol	Typical	Max
Command to data in or parameter in	T06	Command-dependent	

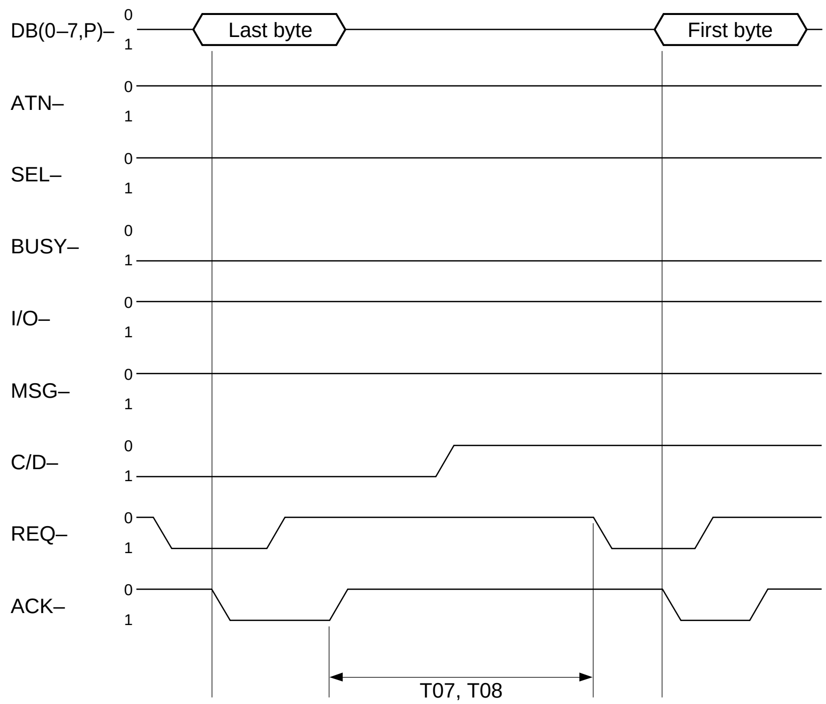

Figure 15. Last command byte to data out

Description	Symbol	Typical	Max
Command to data out or parameter out	T07	Command-dependent	
Command to data (write to data buffer)	T08	<500 µsec	1,025 µsec

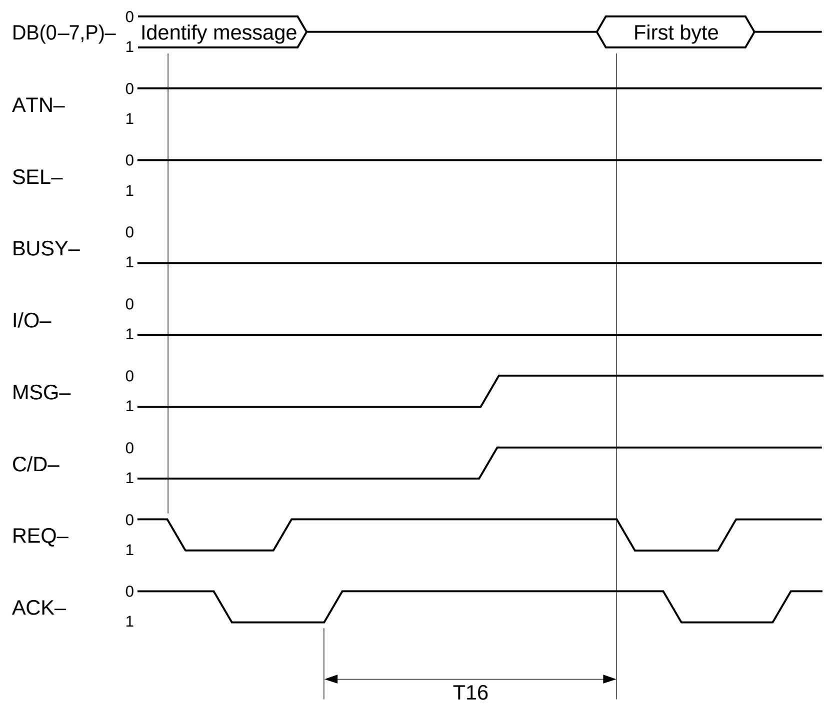

Figure 16. Reselect identify message to data in

Description	Symbol	Typical	Max
Reselect identify message to data (media)	T16	Command-dependent	

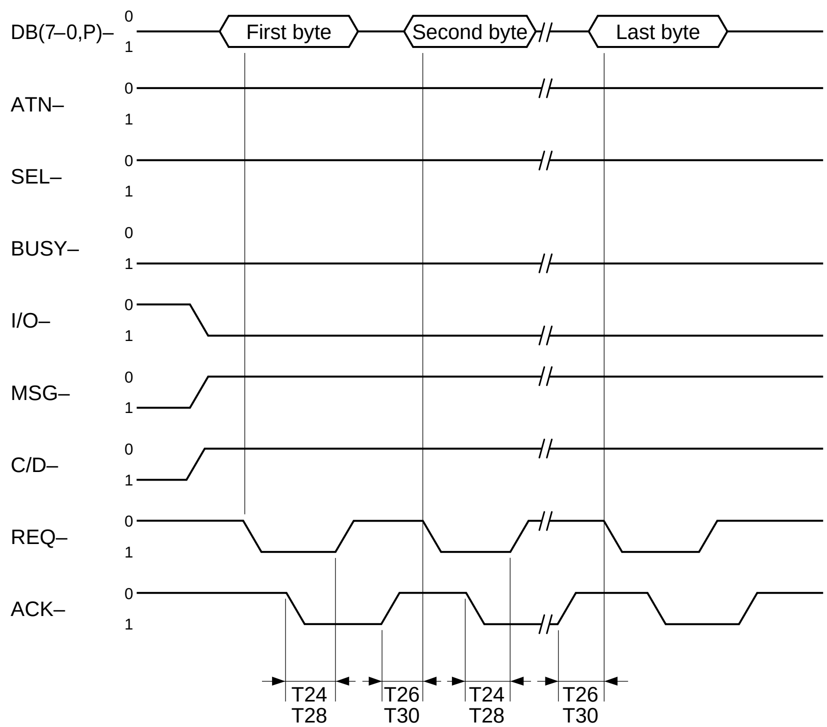

Figure 17. Data in block transfer

Description	Symbol	Typical	Max
Data in block transfer (ASYNC)	T24	<0.1 µsec	0.2 µsec
Next data in byte access (ASYNC)	T26	<0.8 µsec	1.5 µsec
Data in byte transfer (SYNC)	T28	<60 nsec	100 nsec
Next data in byte access (SYNC)	T30	<600 nsec	1.2 µsec

The maximum SCSI asynchronous interface transfer rate is 5 Mbytes per second. Therefore, the minimum time between two leading edges of a request is 200 nsec.

The maximum SCSI synchronous interface transfer rate is 10.0 Mbytes per second. Therefore, the minimum time between two leading edges of a request is 100 nsec.

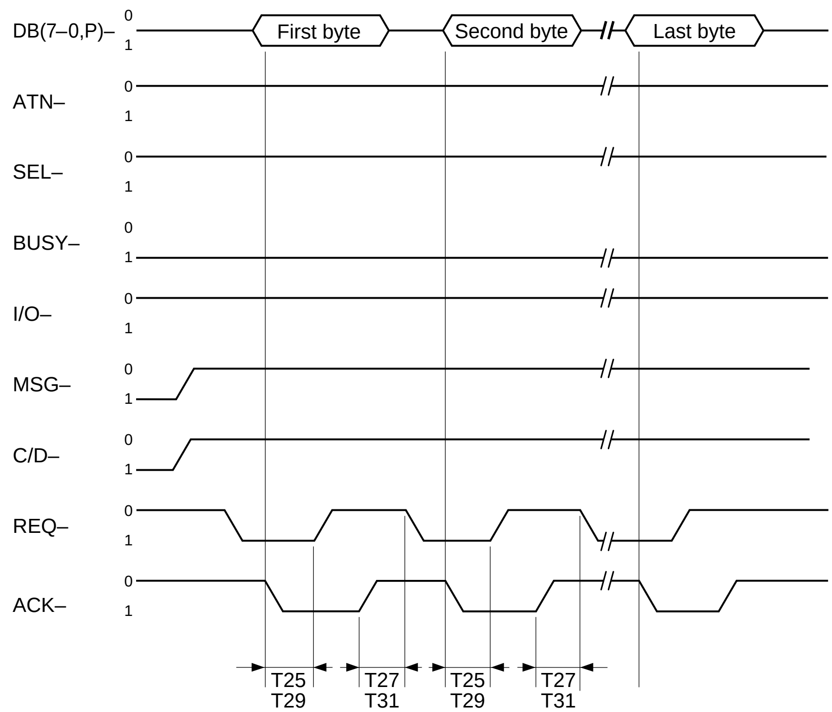

Figure 18. Data out block transfer

Description	Symbol	Typical	Max
Data out block transfer (ASYNC)	T25	<0.1 µsec	0.2 µsec
Next data out byte access (ASYNC)	T27	<0.8 µsec	1.5 µsec
Data out byte transfer (SYNC)	T29	<60 nsec	100 nsec
Next data out byte access (SYNC)	T31	<600 nsec	1.2 µsec

The maximum SCSI asynchronous interface transfer rate is 5 Mbytes per second. Therefore, the minimum time between two leading edges of a request is 200 nsec.

The maximum SCSI synchronous interface transfer rate is 10.0 Mbytes per second. Therefore, the minimum time between two leading edges of a request is 100 nsec.

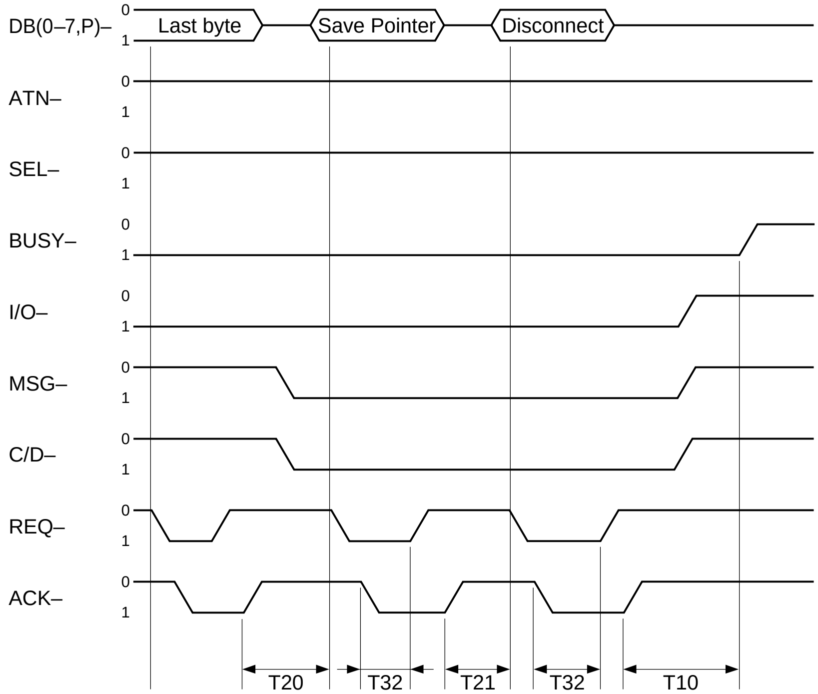

Figure 19. Last data byte, save pointer message and disconnect message

Description	Symbol	Typical	Max
Disconnect message to bus free	T10	<100 µsec	—
Data to save data pointer message	T20	<175 µsec	—
Save data pointer message to disconnect message	T21	<175 µsec	—
Message in byte transfer	T32	<0.1 µsec	0.15 µsec

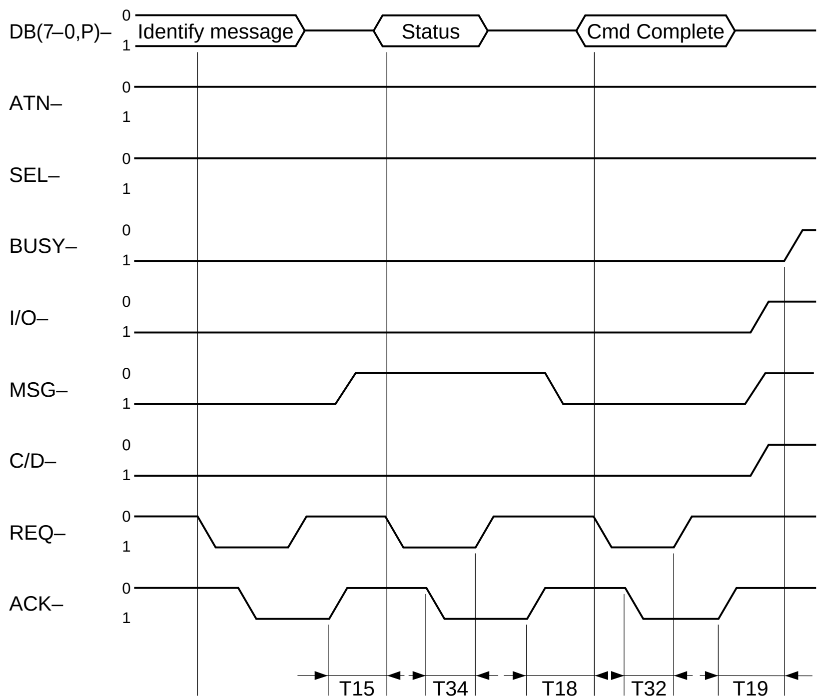

Figure 20. Data in, status, command complete message and bus free

Description	Symbol	Typical	Max
Data to status	T17	Command-dependent	
Status to command complete message	T18	<150 µsec	—
Command complete message to bus free	T19	<100 µsec	—
Message in byte transfer	T32	<0.1 µsec	0.15 µsec
Status byte transfer	T34	<0.1 µsec	0.15 µsec

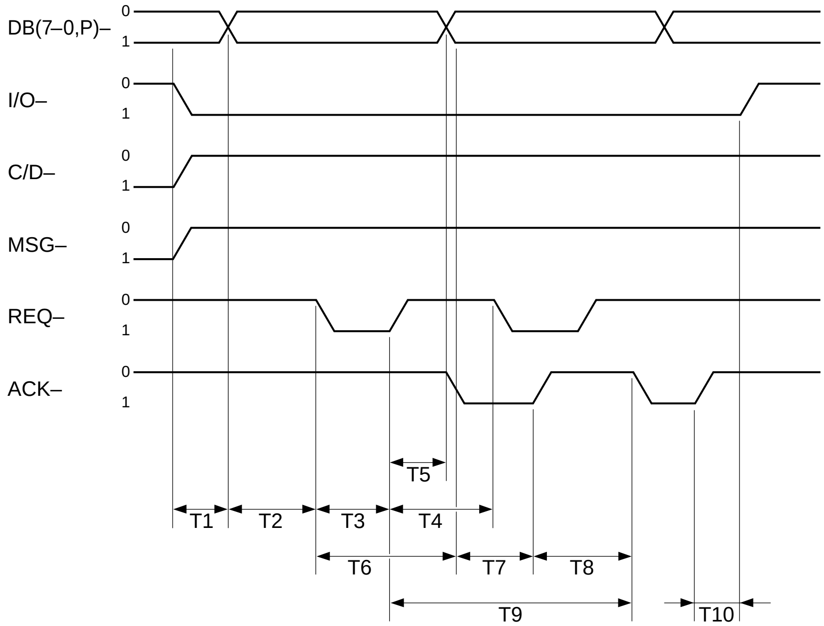

Figure 21. Synchronous timing

Description	Symbol	Min
I/O low to data bus enable	T1	400 nsec
Data bus valid to REQ– low	T2	57.5 nsec
REQ– assertion period	T3	30.0 nsec
REQ– deassertion period	T4	30.0 nsec
REQ– high to data hold	T5	—
REQ– low ACK– low	T6	10 nsec
ACK– assertion period	T7	30.0 nsec
ACK– deassertion period	T8	30.0 nsec
ACK– period	T9	100 nsec
Last ACK– pulse high to phase change	T10	125 nsec

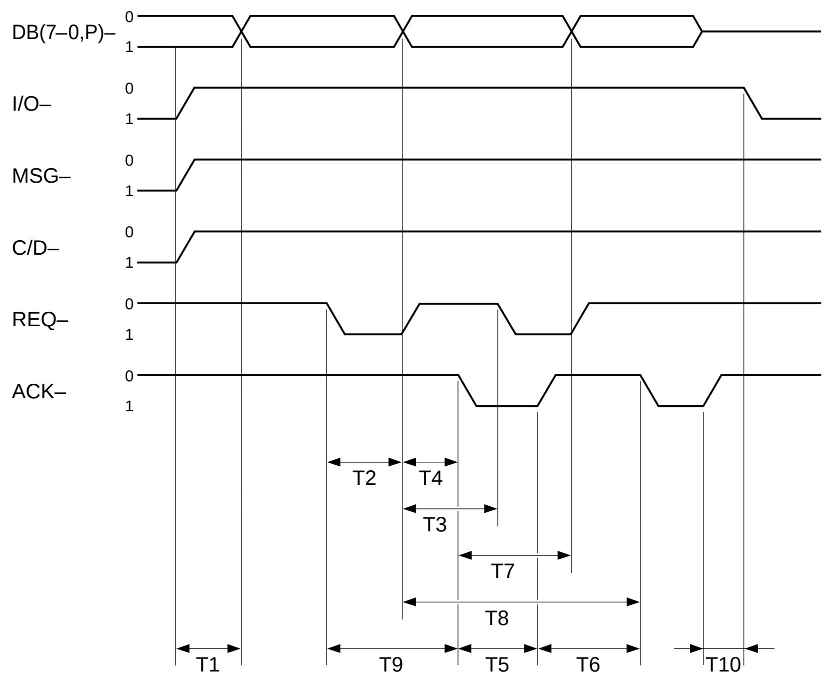

Figure 22. Synchronous write timing

Description	Symbol	Min	Max
I/O high to data bus disable	T1	—	50 nsec
REQ– assertion period	T2	30.0 nsec	—
REQ– deassertion period	T3	30.0 nsec	—
Data valid to ACK– low	T4	—	—
ACK– assertion period	T5	30.0 nsec	—
ACK– deassertion period	T6	30.0 nsec	—
ACK– low to data hold	T7	10 nsec	—
ACK– period	T8	100 nsec	—
REQ– low to ACK– low	T9	10 nsec	—
Last ACK– pulse high to phase change	T10	125 nsec	—

Seagate

Seagate Technology, Inc.
920 Disc Drive, Scotts Valley, California 95066, USA

Publication Number: 36321-101, August 1995, Printed in USA