SCSI, The Nuts & Bolts

KnowledgeTek

Welcome Notes

Welcome to the KnowledgeTek SCSI Nuts and Bolts seminar.

KnowledgeTek was founded in December, 1995 to combine the technical training resources of two established companies, M5 Electronics and Zadian Technologies. KnowledgeTek offers both public and in-house courses on a large range of storage and interface technology subjects.

Consulting services are offered through our affiliate company, M5 Electronics.

For more information, please contact:

KnowledgeTek, public classes

Phone	303-465-1800
Fax	303-426-1350

KnowledgeTek, in-house classes

Phone	408-937-7200
Fax	408-937-7272

M5 Electronics

Phone	303-499-0976
Fax	303-499-0977

Copyright

Disclaimer

This material is not a specification for SCSI or any other system.

KnowledgeTek makes no warranty and assumes no liability arising out of the application or use of any product, hardware, software, system, circuit, or anything else described herein.

KnowledgeTek assumes no responsibility for errors appearing in this document.

KnowledgeTek assumes no responsibility for any claims that the concepts or details discussed in the seminar, or disclosed in the course materials, are proprietary to any person or company.

Company, brand, product, trade, and other names used in this document are trademarks of their respective holders.

Seminar participants are urged to clear any designs proposed for products with their patent and copyright council.

Conventions Used in This Book

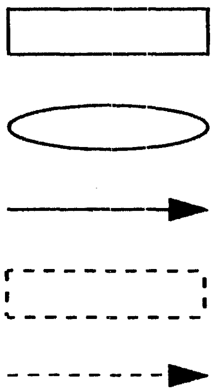

	Process
	Decision
	Flow
	Optional Process
	Optional Flow
V.U.	Vendor Unique
xxh	Hexadecimal Value
xb	Binary Value
M	Mandatory
O	Optional
iid	Initiator SCSI ID

tid	Target SCSI ID
MSB	Most Significant Bit
LSB	Least Significant Bit
H/A	Host Adapter
LU	Logical Unit
LUN	Logical Unit Number
?	Don't Care
ms	Millisecond
us	Microsecond
ns	Nanosecond

Table of Contents

Table of Contents (Continued)

Table of Contents (Continued)

Table of Contents (Continued)

Table of Contents (Continued)

INTRODUCTION

Evolution of SCSI

SCSI-1

Begun: 1982
Approved: 1986
1 standards document, approx. 200 pages
6 device types defined

SCSI-2

Begun: 1986
Approved: 1994
1 standards document, approx. 460 pages

Changes from SCSI-1
- 4 new device types
- new commands
- new features
- new requirements

SCSI-3

Begun: 1991
Approved: Not Yet
$\approx$25 standards, est. 2500 pages

Changes from SCSI-2
- 1 new device type
- new/changed commands
- new features
- new requirements
- formal architecture
- new physical interfaces and protocols

Overview of Non-SCSI System

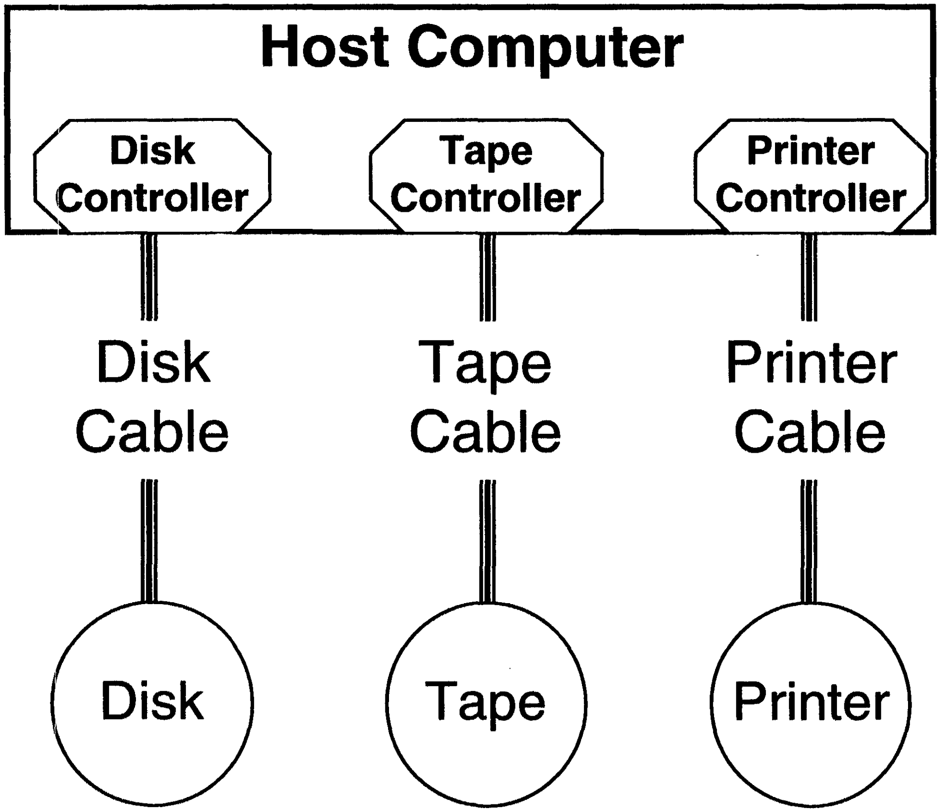

Host Computer

- → Device Dependent
- → Software

Device

- → Simple Firmware

Cable

- → Low-Level Signals
- → Low-Level Commands
- → Physical Addressing

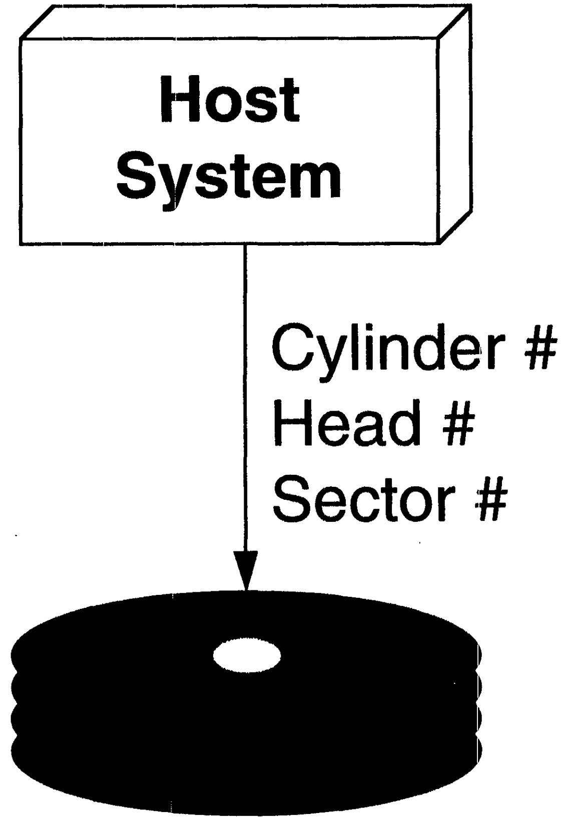

Disk consisting of
n physical sectors

Overview of a SCSI System

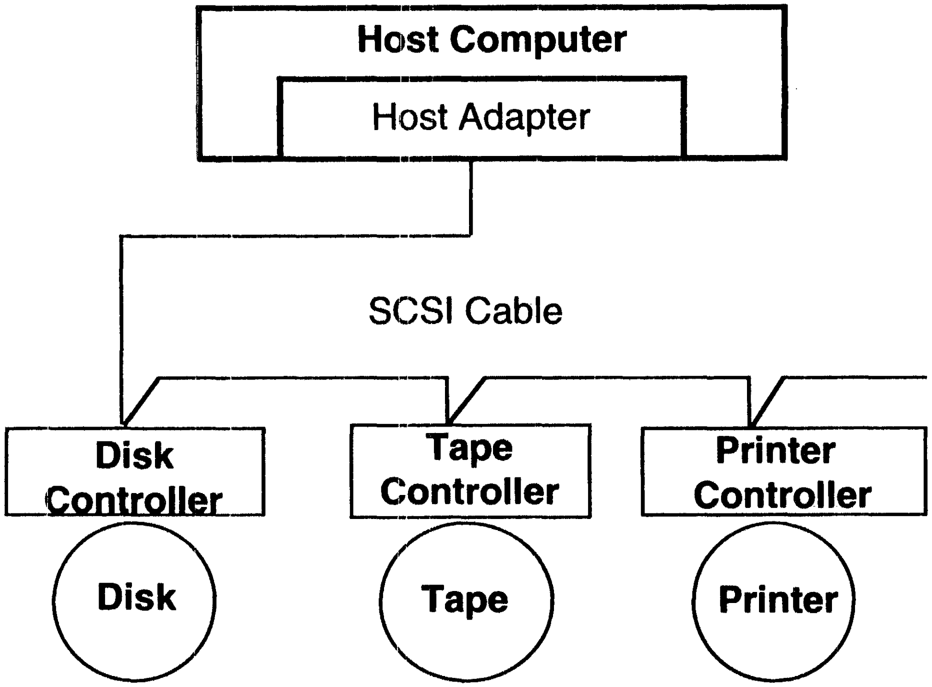

Host Computer
- generic software

Device
- intelligent firmware

Cable
- high-level signals
- high-level commands
- **logical block addressing**

Logical Block Addressing

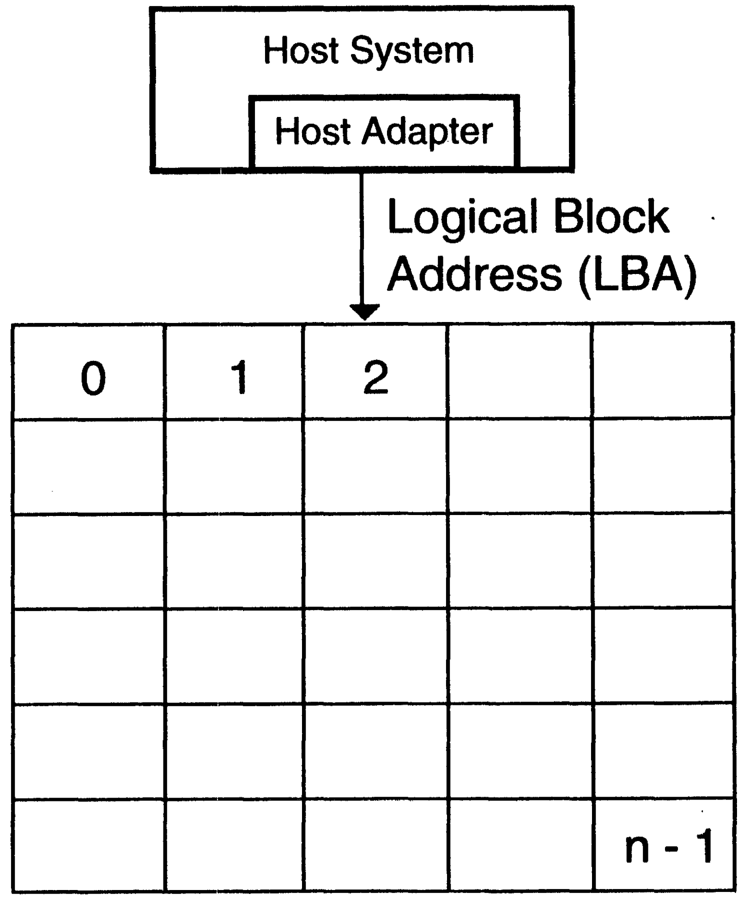

Data space consisting of n logical blocks

Advantages of SCSI

- Interface with different device types through the same cable
- Peripheral devices of the same type have similar characteristics
  - easy to replace old devices with new ones
- Peripheral devices are intelligent and independent
  - frees up the computer to do other work
- I/O is independent of system bus
  - peripheral devices can work with different computer types
  - preserves computer hardware investment
- Fast hardware
  - 20 MB/sec on 8 bit bus, 40 MB/sec on 16 bit bus
- Fast software
  - multi-threading support using disconnect/reconnect and queueing

Standards

Current Standards

SCSI-1: ANSI X3.131-1986

CCS: X3T9.2/85-52, 1986

SCSI-2: ANSI X3.131-1994

Where to Get Documents

573-7950

SCSI Bulletin Board System: (719) ~~574-0424~~

Anonymous FTP Site: ftp.symbios.com

SCSI-1, CCS, and SCSI-2:

American National Standards Institute
1430 Broadway
New York, NY 10018

Alternate source for above documents plus SCSI-3 working documents:

Global Engineering Documents
15 Inverness Way East
Englewood, CO 80112-5704
Tel: (800) 854-7179 or (303) 792-2181

SCSI Devices

SCSI Physical Interface

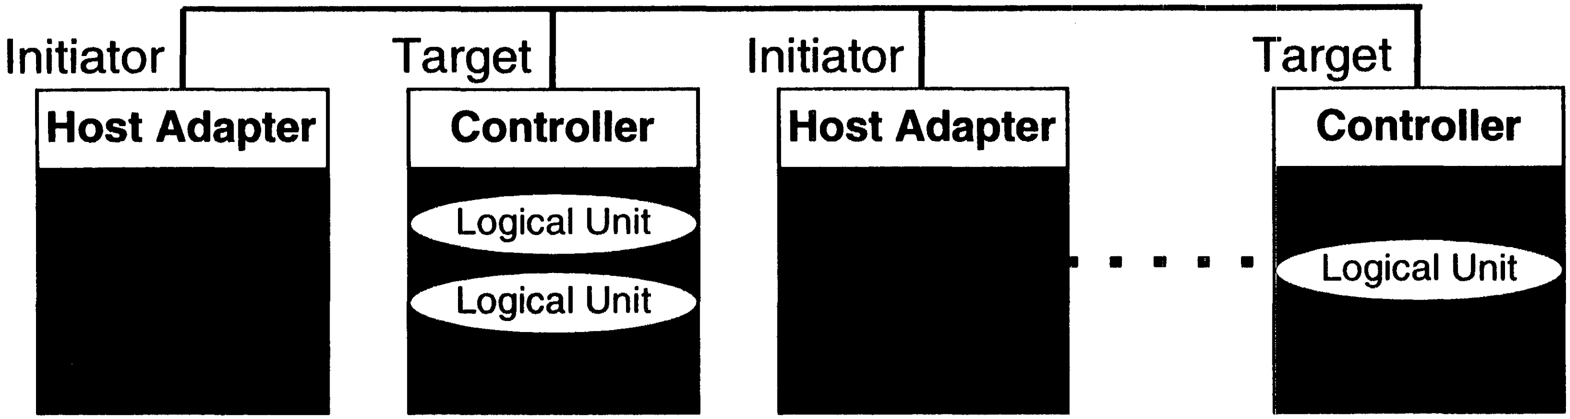

- SCSI Device is Initiator or Target
- Initiator Originates Operation (Usually Host Computer)
- Target Performs the Operation (Usually Peripheral Device)
- Each Device has Unique SCSI Address (ID)
  - 0 - 7 (Narrow)
  - 0 - 15 (Wide, 2-byte)
- Max 8 SCSI Devices (Narrow)
- Max 16 SCSI Devices (Wide, 2-byte)
- Two Devices Communicating at a Time
- Target has Controller and Logical Units
- Logical Units (LU's)
  - 8 LU's max in SCSI-2, 64 LU's max in SCSI-3, per controller. Logical units are numbered 0 - 7 or 0 - 63.

SCSI Devices
(Continued)

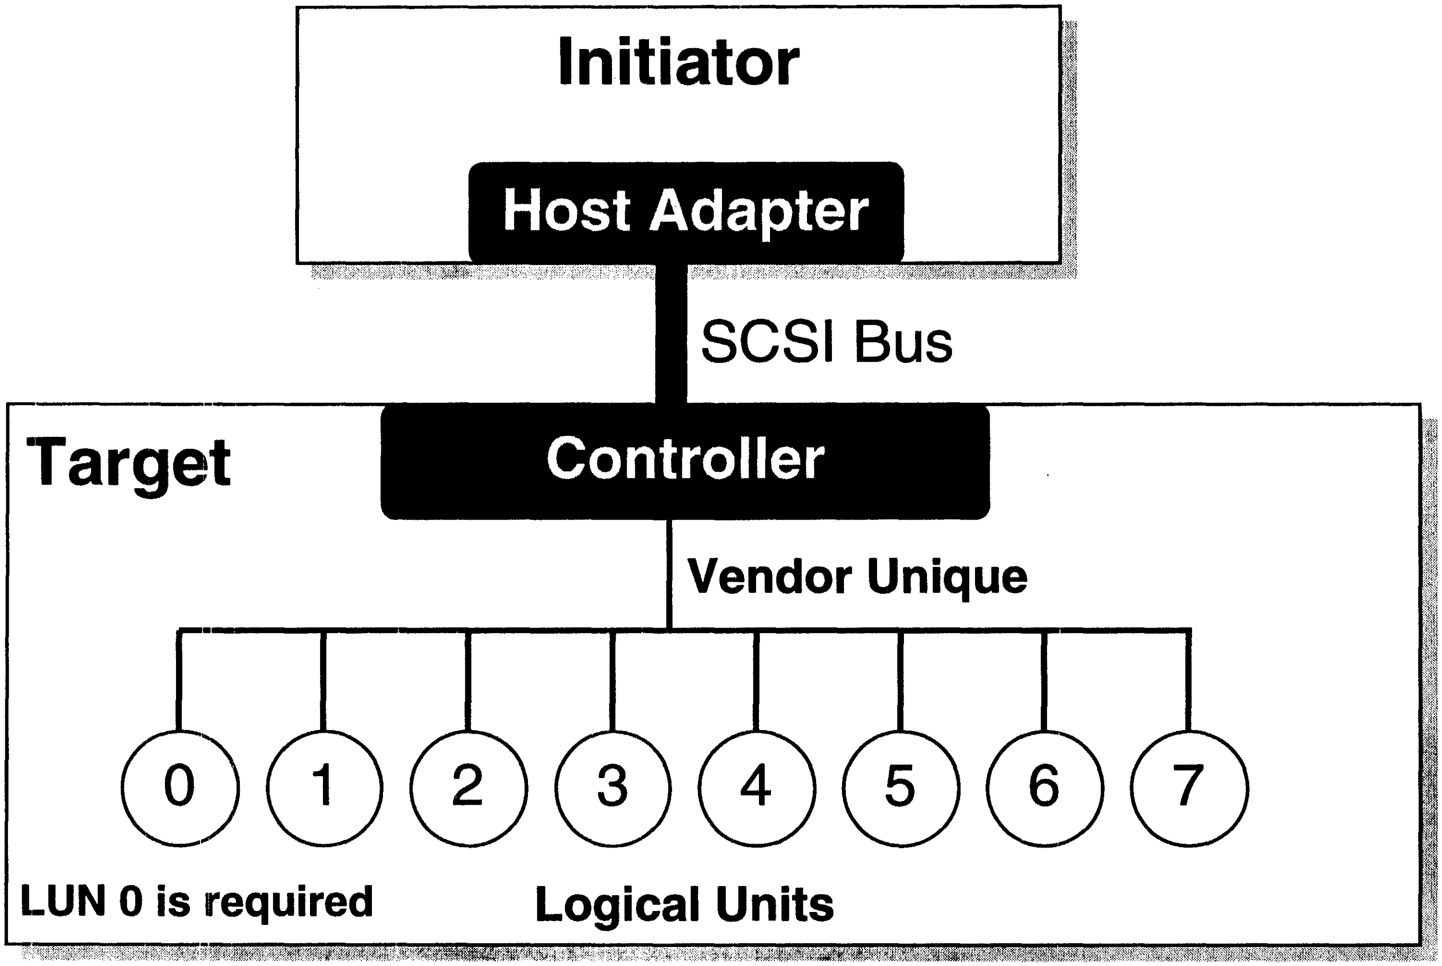

System Configurations

Single Initiator, Single Target

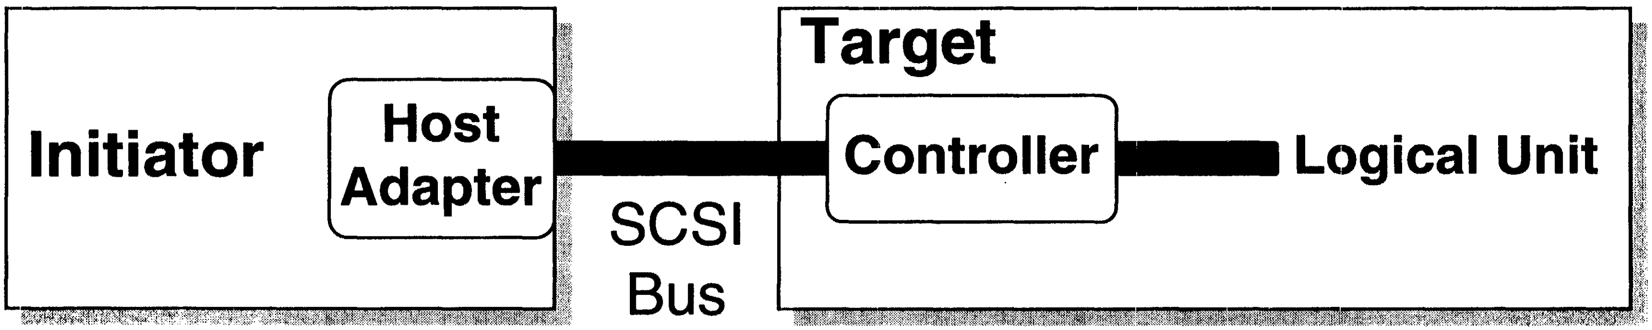

Single Initiator, Single Target, Multiple Logical Units

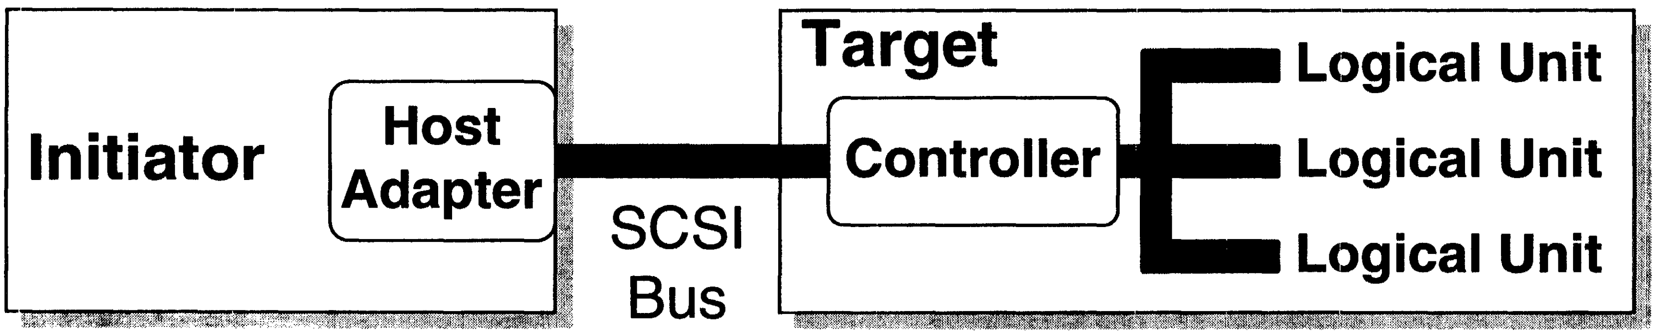

System Configurations (Continued)

Single Initiator, Multiple Targets

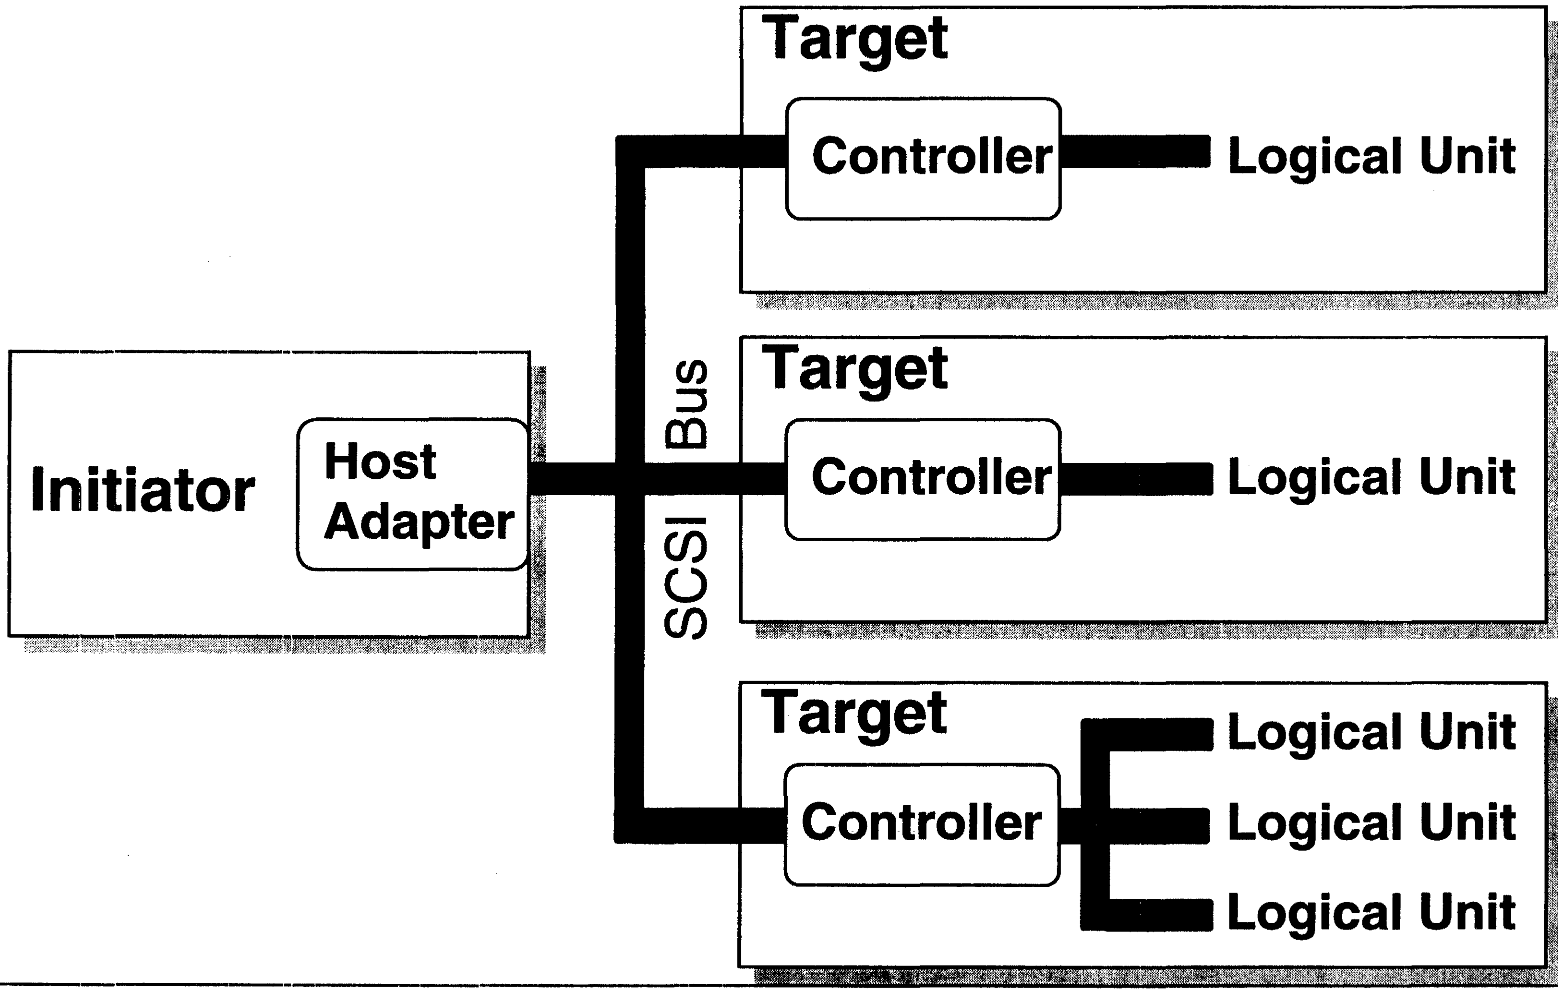

System Configurations (Continued)

Multiple Initiators, Multiple Targets

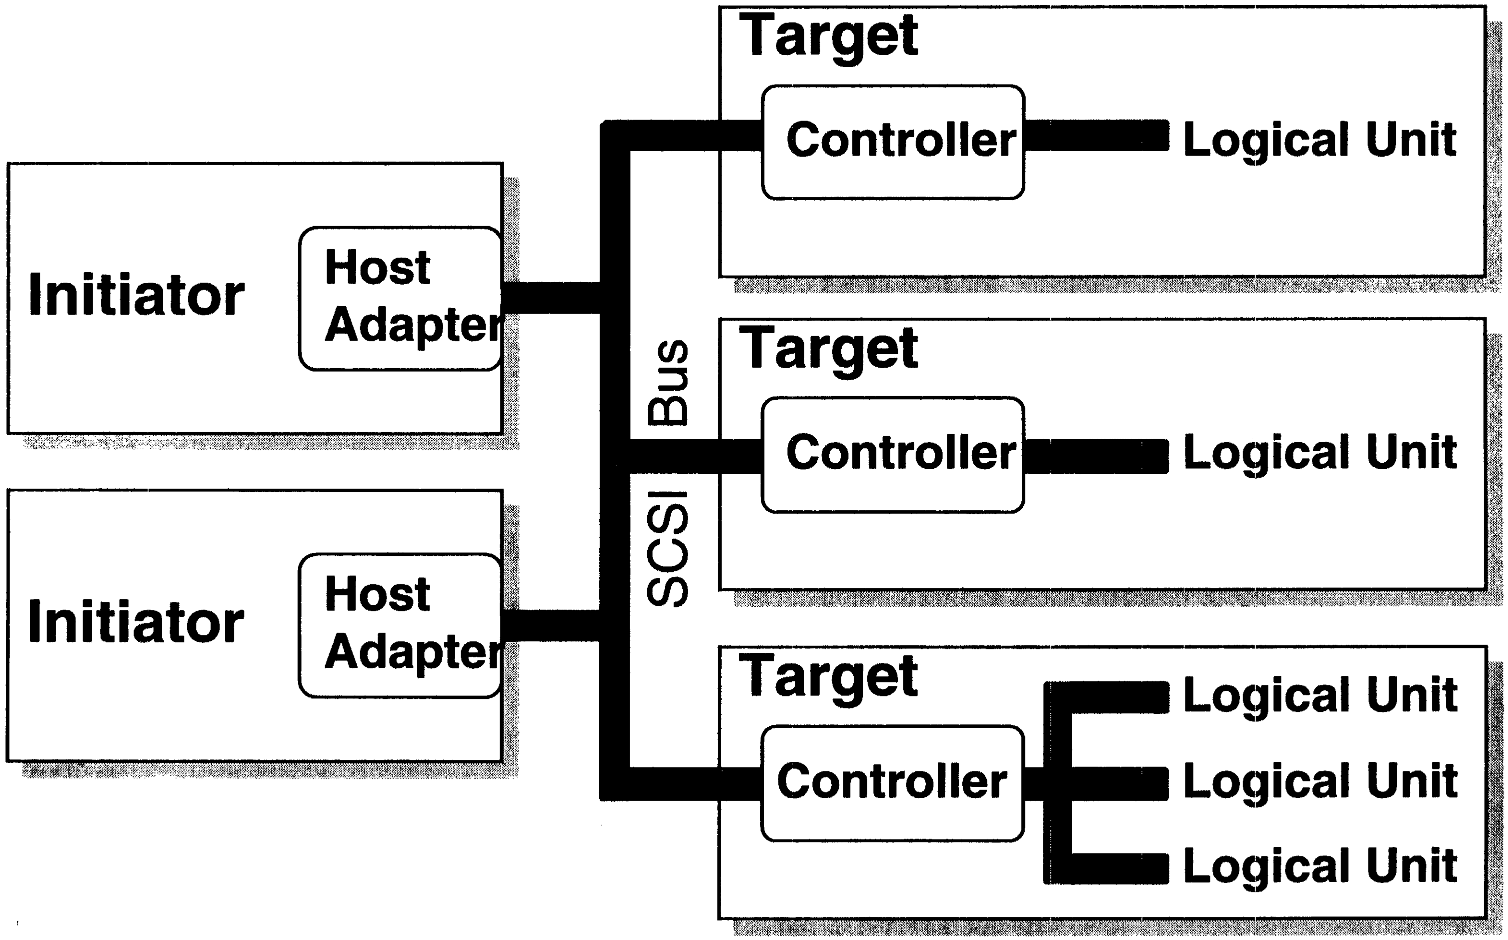

System Configurations (Continued)

Multiple Buses

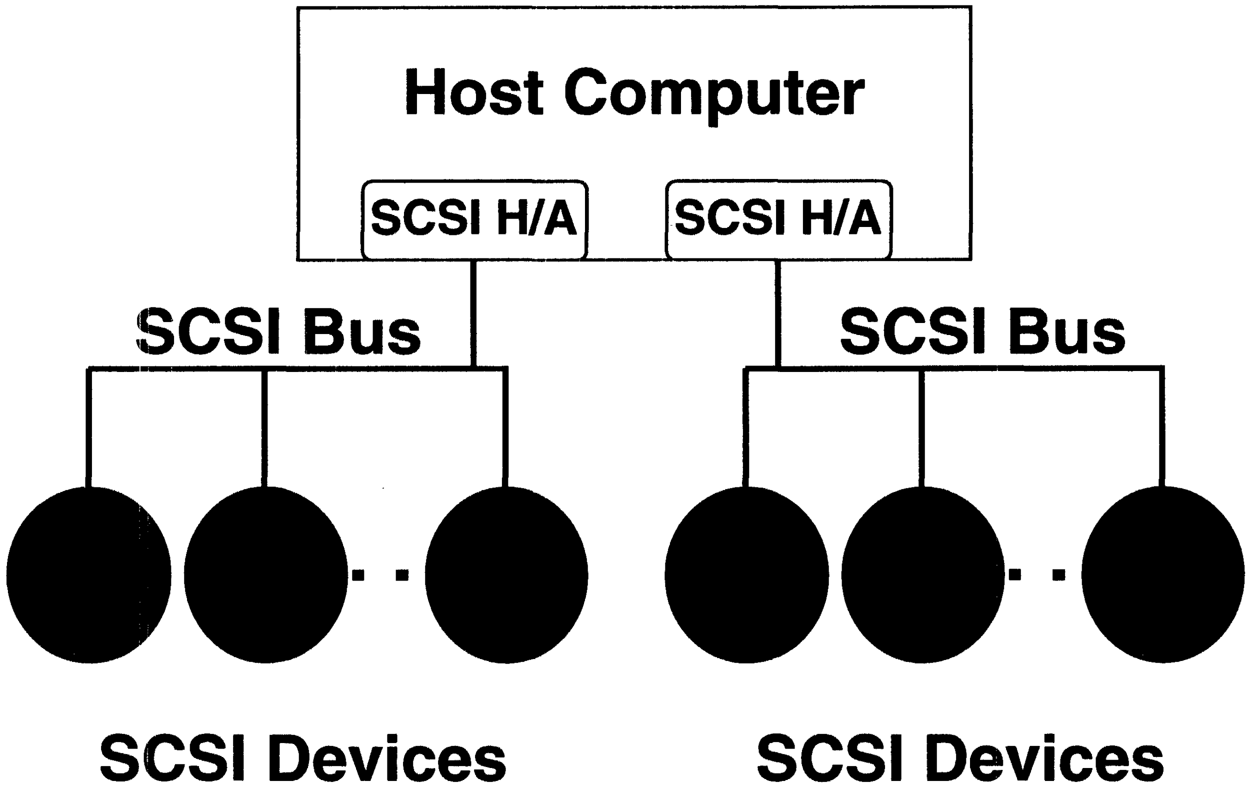

Initiator Block Diagram

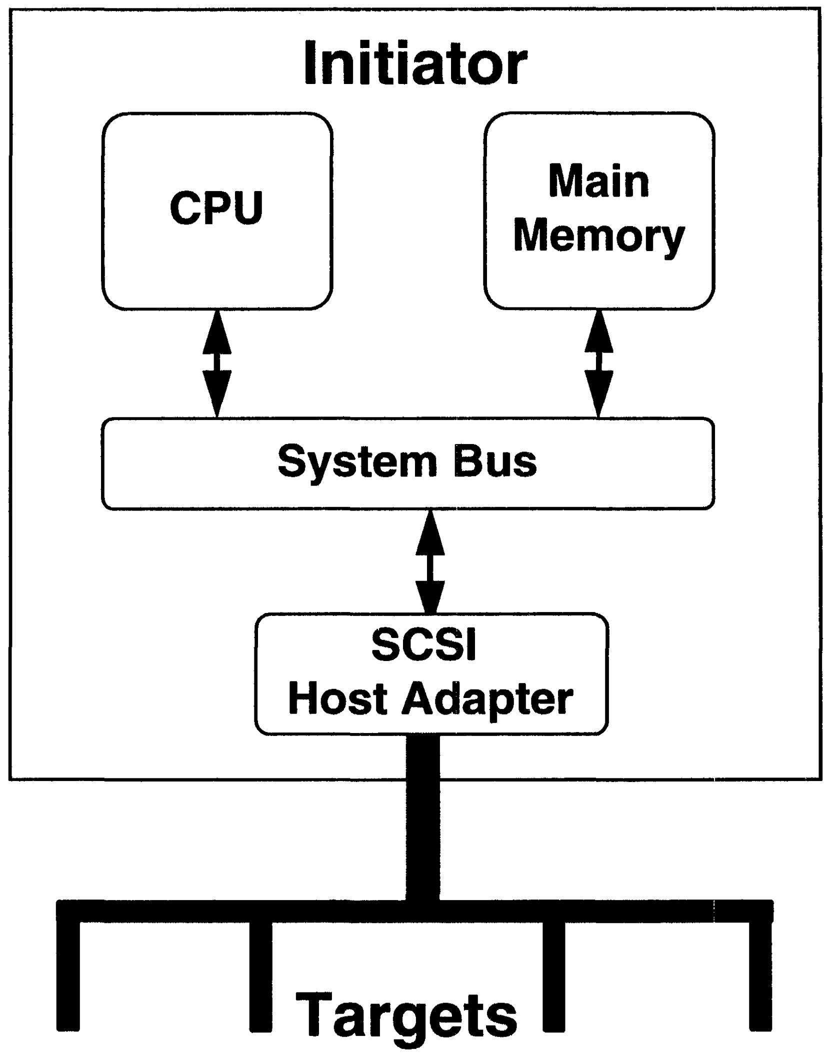

Target Block Diagram

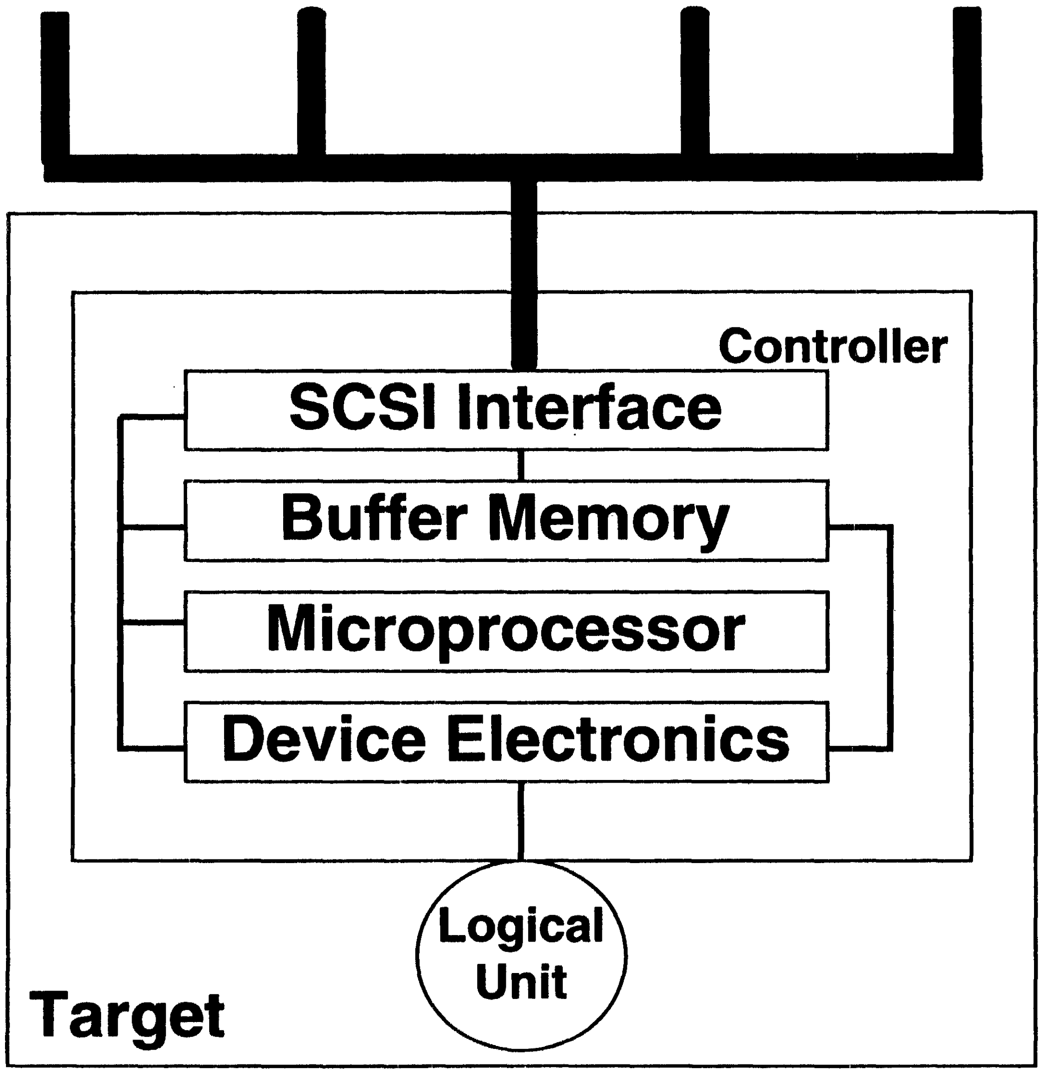

Configuration

(Not Required by SCSI Standard)

- SCSI Target Device Jumpers, Dip Switches, or software-selectable switches:
  - SCSI Address
  - Disk Wait Spin (no spin on power-up)
  - Self-Test (action is vendor-unique)
  - Parity Enable (described later)
  - Supply Internal Termination (described later)
  - Terminator Power (described later)
  - Unit Attention Disable (described later)
- SCSI Host Adapter Jumpers (or Dip Switches):
  - SCSI Address
  - Parity Enable
  - Host bus address and interrupt mapping
  - Self-Test
  - Supply Internal Termination (described later)

Disconnect/Reconnect Procedures

Connect	Initiator Selects a Target
Disconnect	Target Releases Control of the Bus
Reconnect	Target Reselects the Initiator to Resume Command Execution

Note: Initiator gives permission for disconnect, target decides if and when (details described later).

Disconnect/Reconnect Procedures

Example with No Disconnect

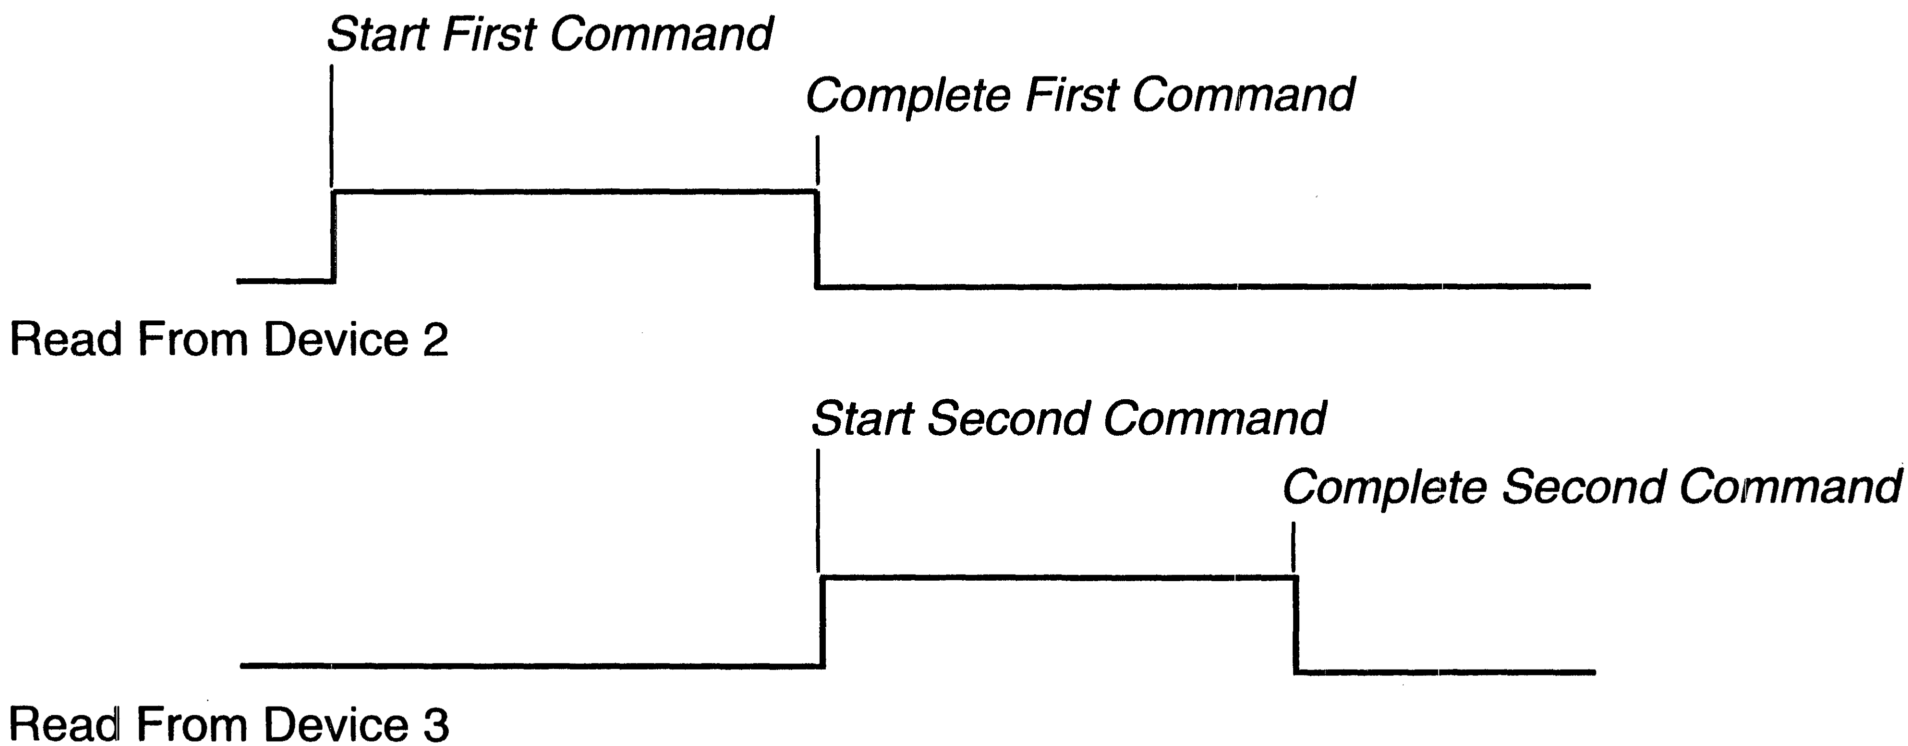

Disconnect/Reconnect Procedures

Example With Disconnect

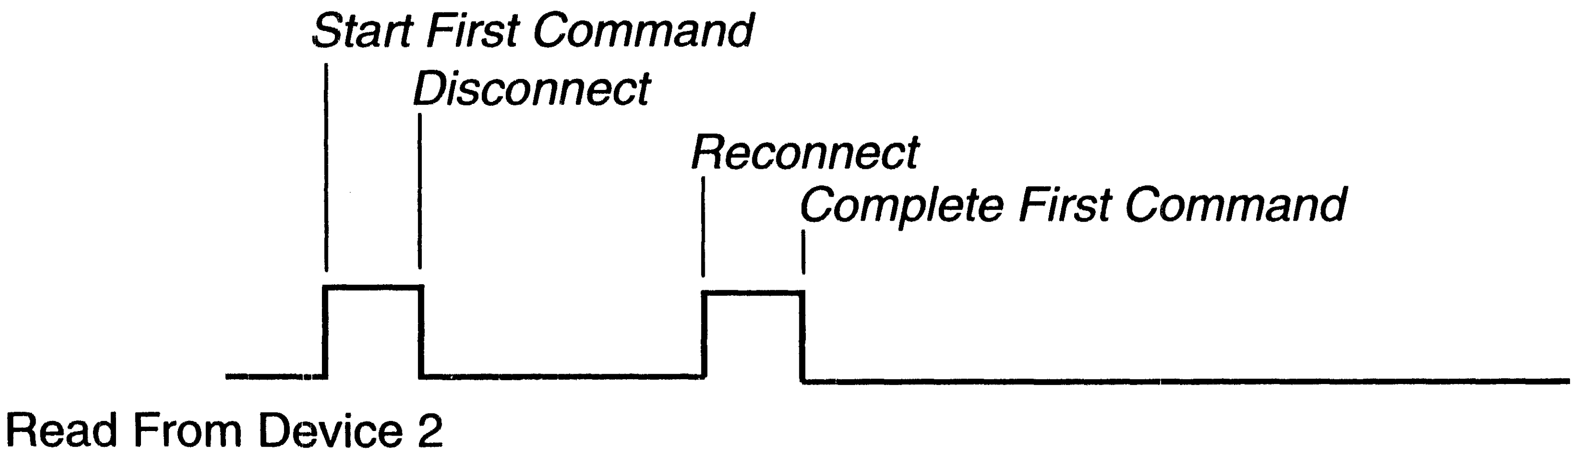

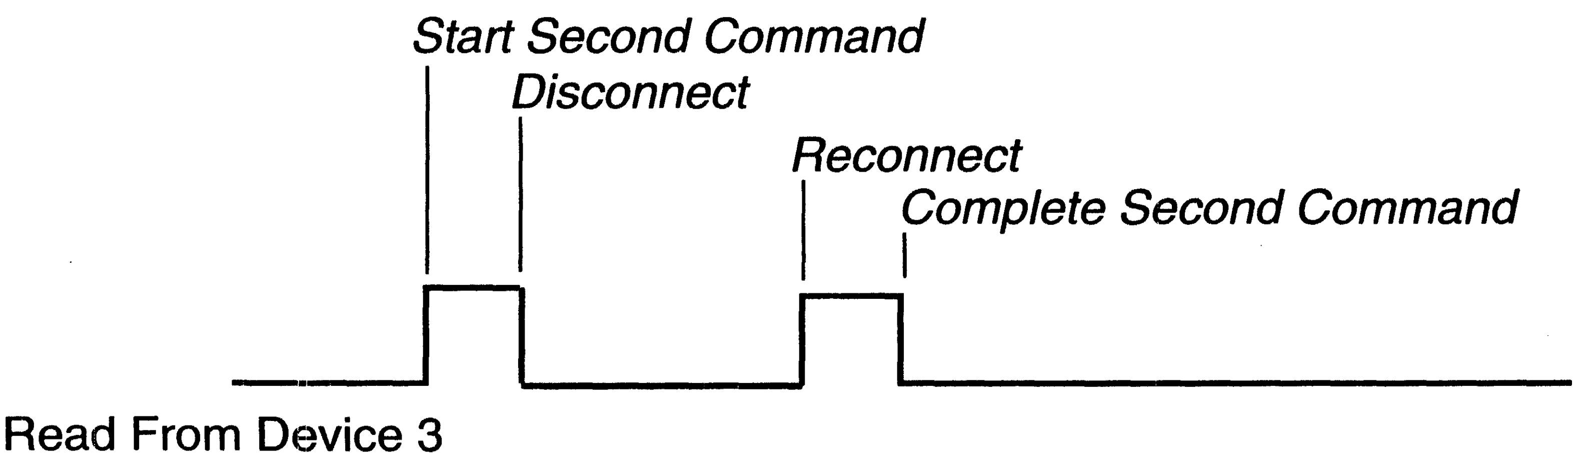

Disconnect/Reconnect Procedures
(Continued)

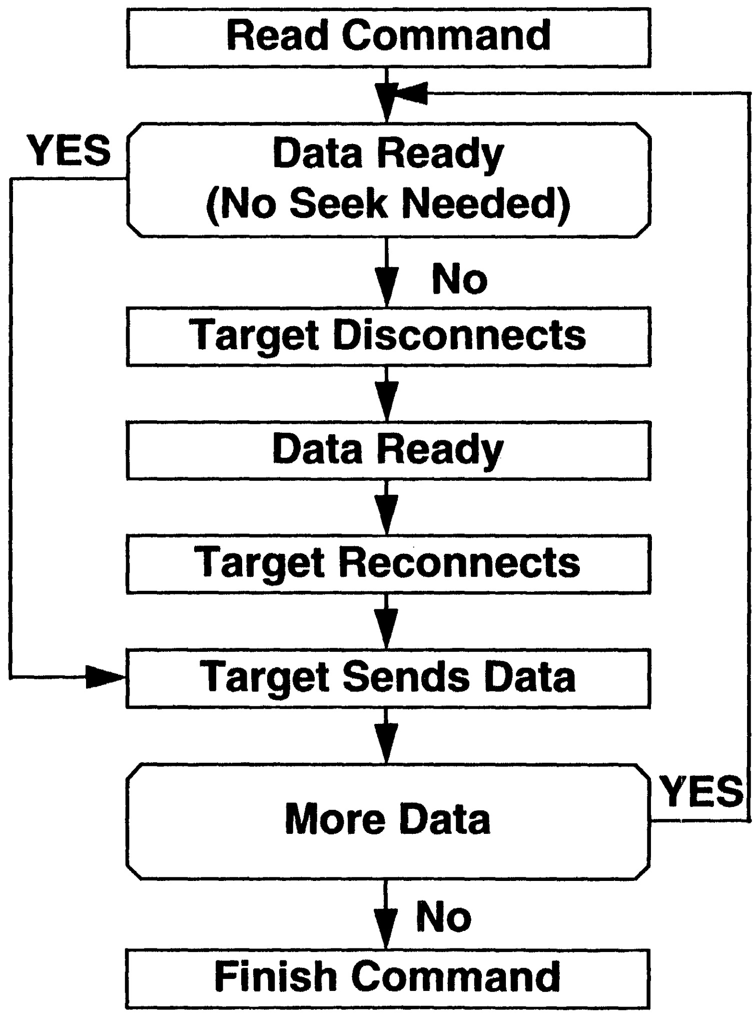

Disconnect/Reconnect Procedures
(Continued)

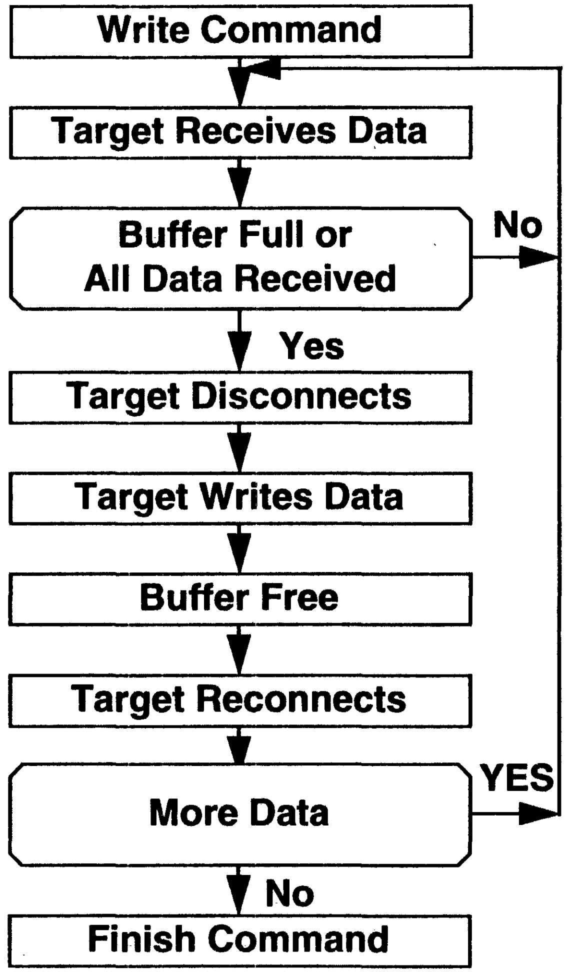

Bus Phases

Levels of SCSI Bus Activity

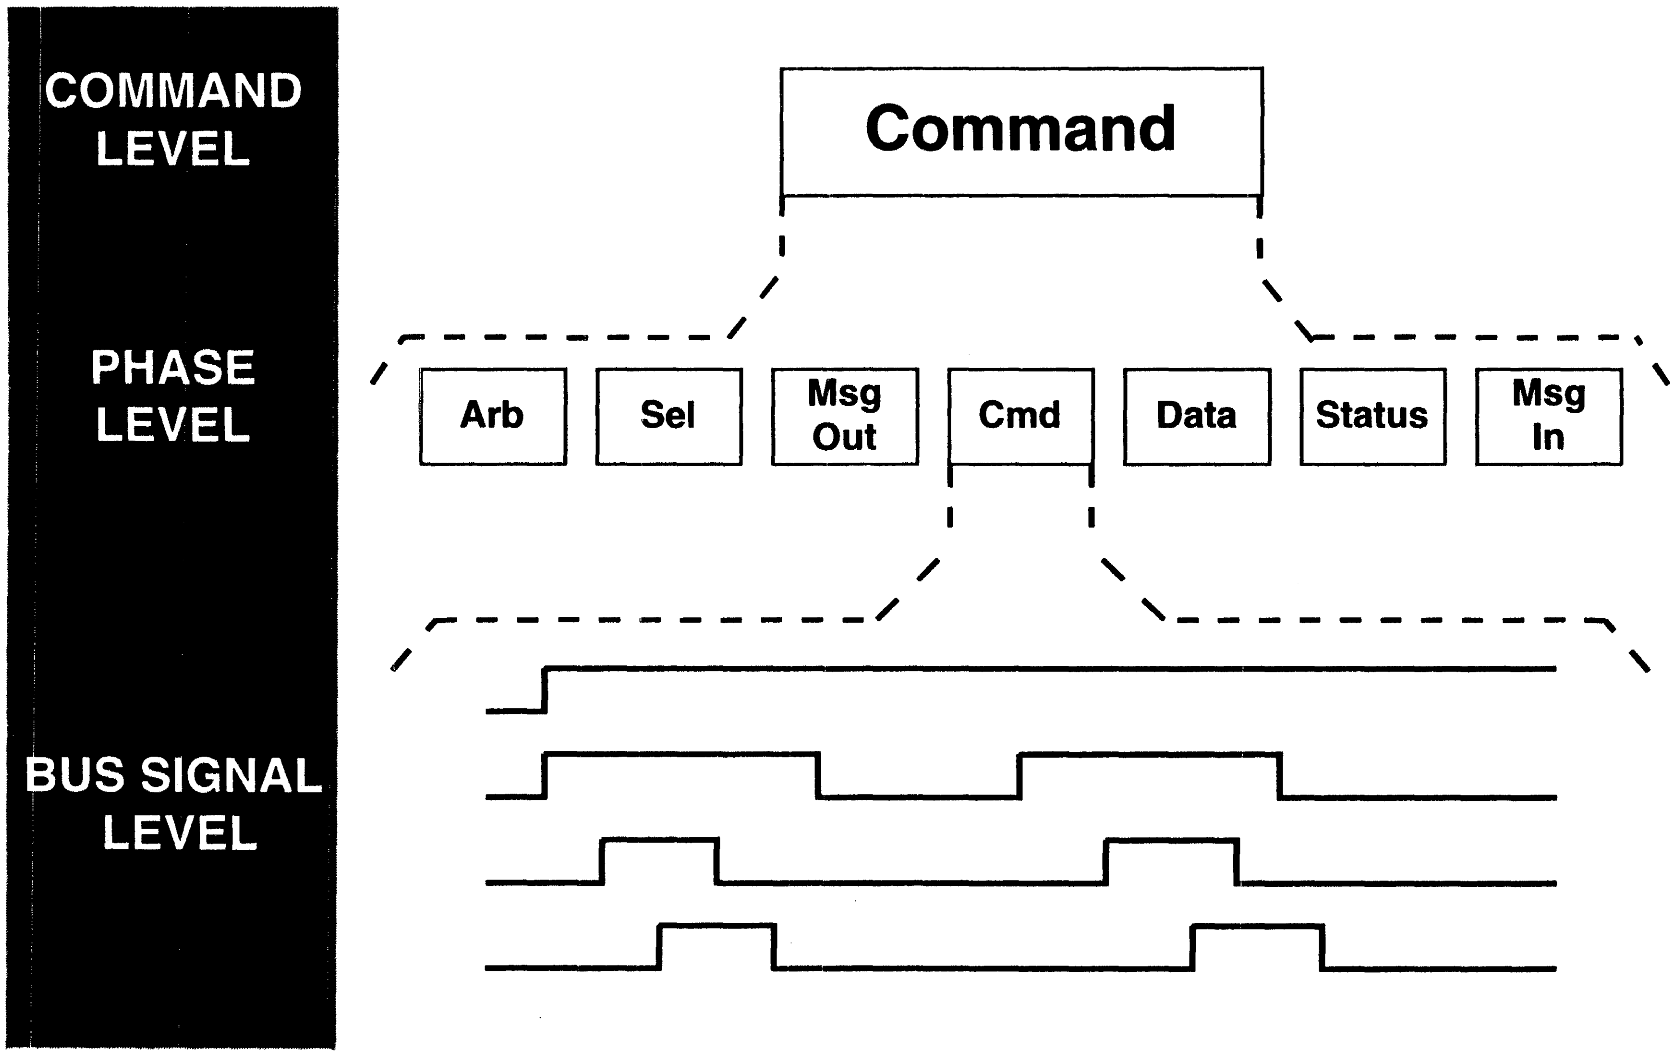

Bus Phases

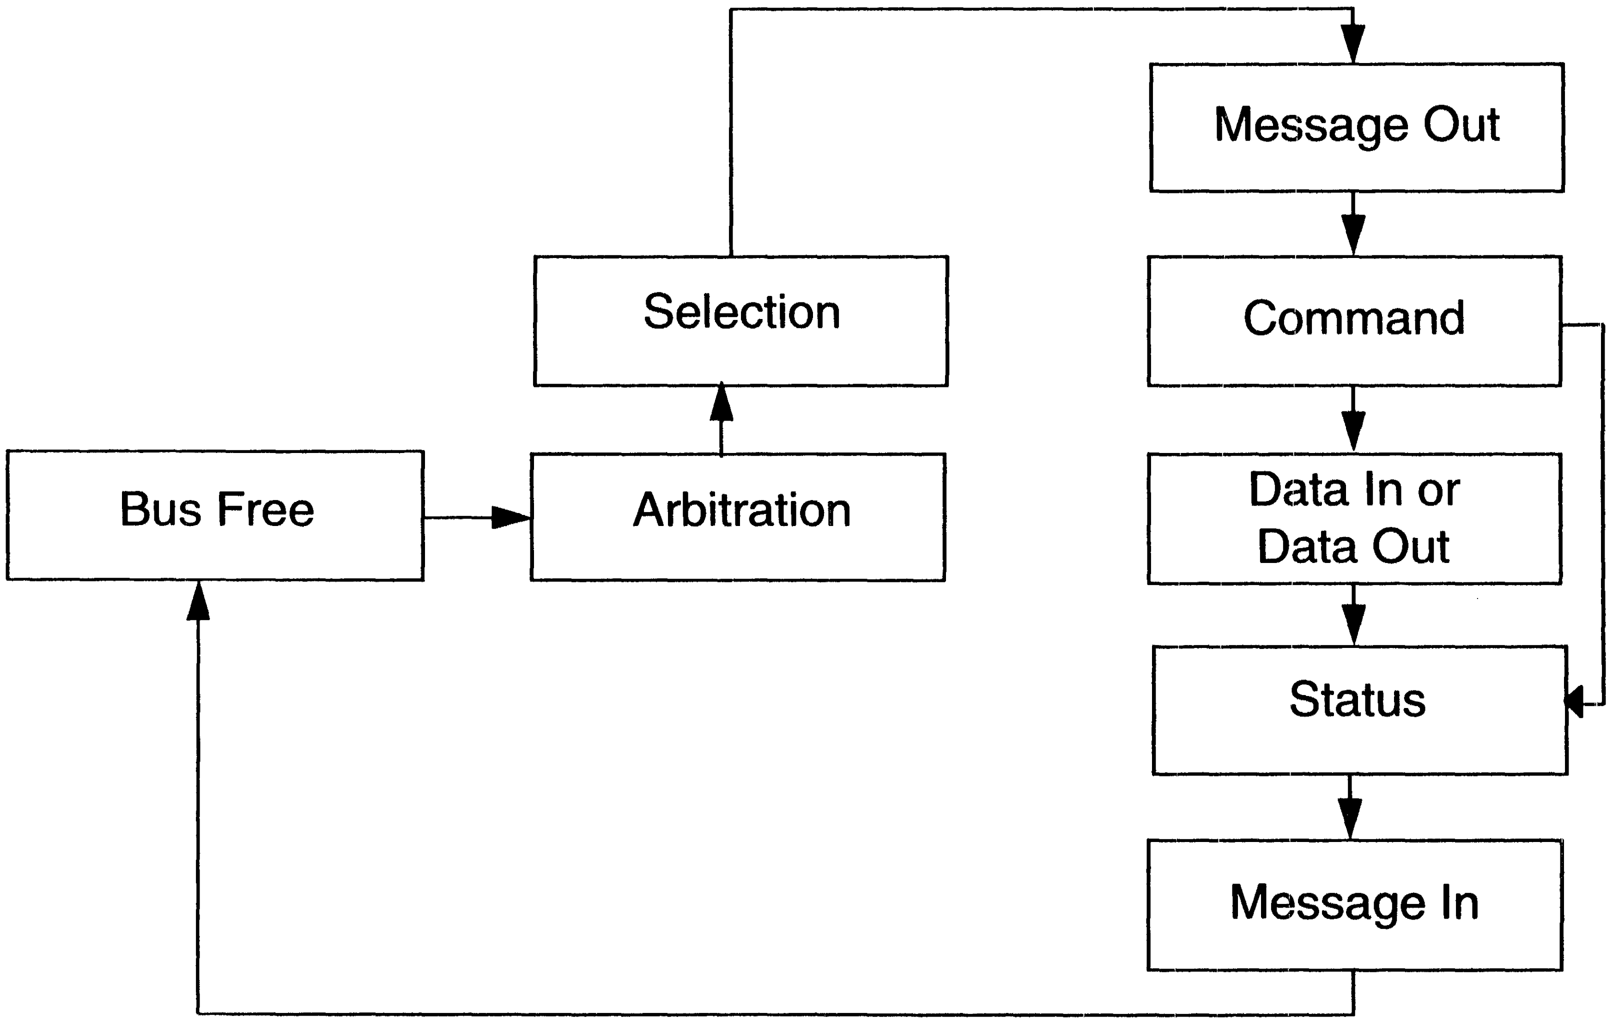
Message Out
Selection
Command
Bus Free
Arbitration
Data In or
Data Out
Status
Message In

Bus Phases (Continued)

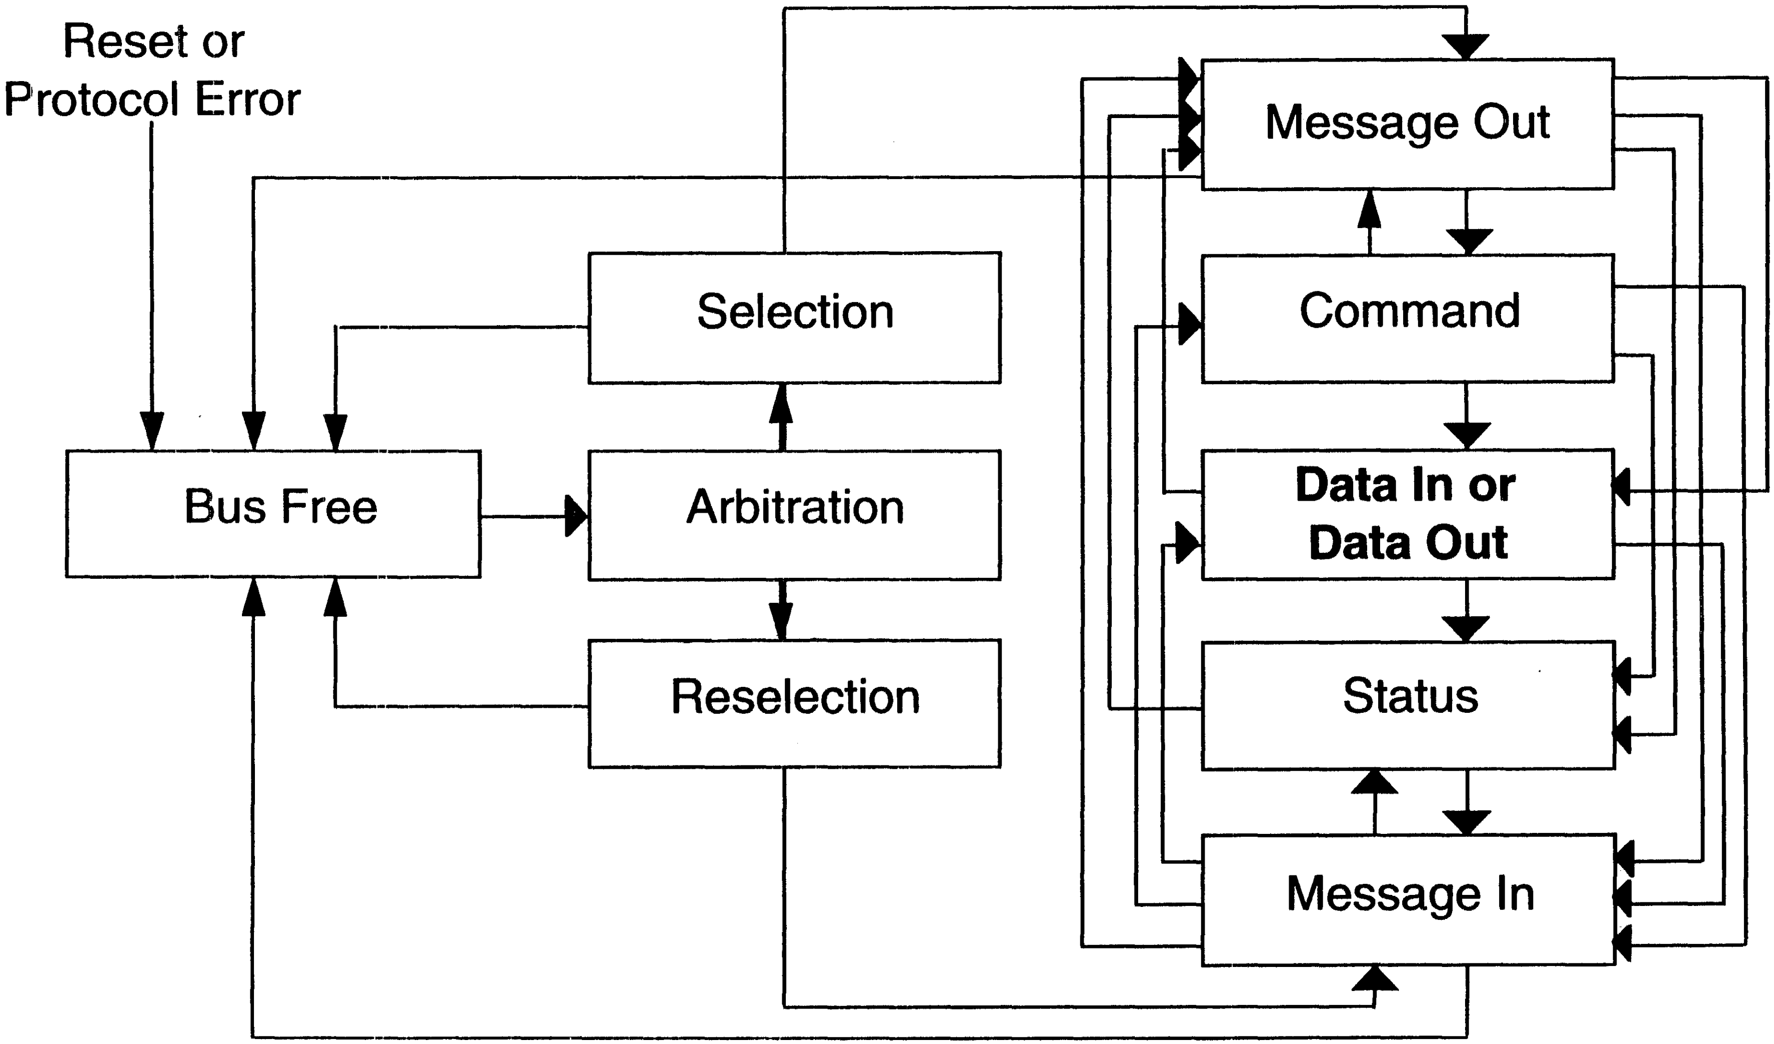

Bus Free Phase

Bus Phase	Data Bus
Bus Free	**0**

- Bus is idle
- Data Bus = 0
- Terminators control the bus signals.

Arbitration Phase

Bus Phase	Data Bus
Arbitration	Arbitrators ID's

- Used to resolve contention over the bus
- Arbitrating devices put their SCSI ID's on the Data Bus
- Distributed Arbitration (no master or slave)
- Highest SCSI ID wins and gains control over the bus
- ID priorities during arbitration:
  - Highest 7 6 5 4 3 2 1 0 15 14 13 12 11 10 9 8 Lowest
- Losers back off
- ID's typically jumper selectable

Selection Phase

Bus Phase	Data Bus
Selection	Initiator and Target ID's

- Initiator selects a target
- Initiator puts Initiator and Target SCSI ID's on Data Bus

Question: What happens if there is no target with that ID?

Reselection Phase

Bus Phase	Data Bus
Reselection	Initiator and Target ID's

- Target re-selects an initiator
- Target puts Initiator and Target ID's on Data Bus

Connect and Reconnect Procedures Revisited

- Connect
  - the __________ generates a connect sequence in order to __________, using the following sequence of bus phases:

- Reconnect
  - The ____________ generates a reconnect sequence in order to ____________, using the following sequence of bus phases:

Information Transfer Phases

Bus Phase	Data Bus
Command	CDB from Initiator (see next page)
Data In	Data from Target
Data Out	Data from Initiator
Status	Status from Target
Message In	Message from Target
Message Out	Message from Initiator

Information Transfer Phases
(Continued)

- CDB
  - Command Descriptor Block
  - Bytes describing the command to execute
  - Prepared by the initiator and sent to the target
  - Transferred during command phase
- In Versus Out (always from initiator's perspective)

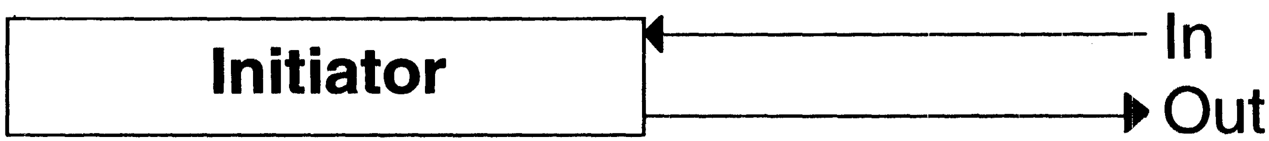

- Target Decides Information Transfer Phase - Initiator Follows

Information Transfer Phases
(Continued)

What Messages Are

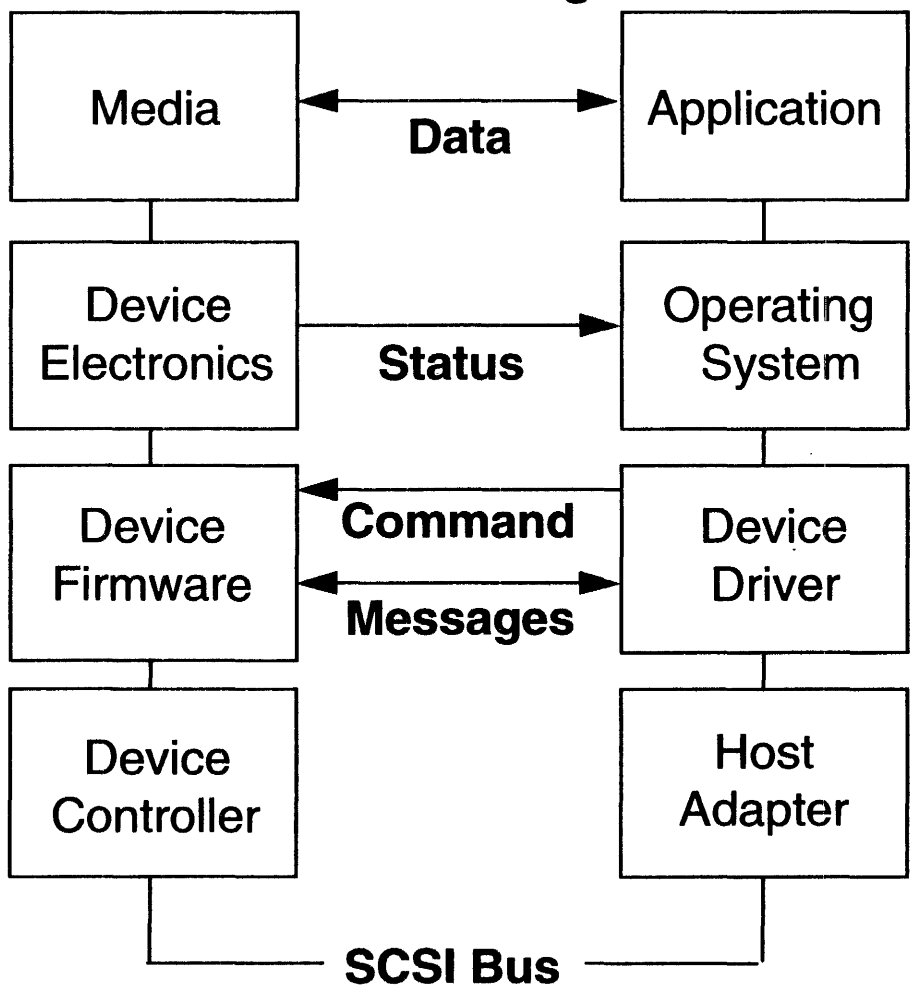

Summary of Bus Phases

Bus Phase	Who Decides Bus Phase	Data Bus Contents	Who Supplies Data Bus
Bus Free	Terminators	00h	Terminators
Arbitration	Arbitrators	Arbitrating Device ID's	Arbitrators
Selection	Initiator	Initiator and Target ID's	Initiator
Reselection	Target	Target and Initiator ID's	Target
Command	Target	CDB Byte	Initiator
Data In	Target	Data In Byte	Target
Data Out	Target	Data Out Byte	Initiator
Status	Target	Status Byte	Target
Message Out	Target	Message Out Byte	Initiator
Message In	Target	Message In Byte	Target

Nexus Levels

Nexus:

A relationship or connection between devices.

IT Nexus:

Nexus between initiator and target.
Established with Selection phase.

ITL Nexus:

Nexus between initiator, target, and logical unit.
Established with Identify message.

ITLQ Nexus:

A further level of I/O Process connection.
Will be defined further in the Command Queueing portion of this class.

Typical SCSI Phase Sequence

Without Disconnect/Reconnect

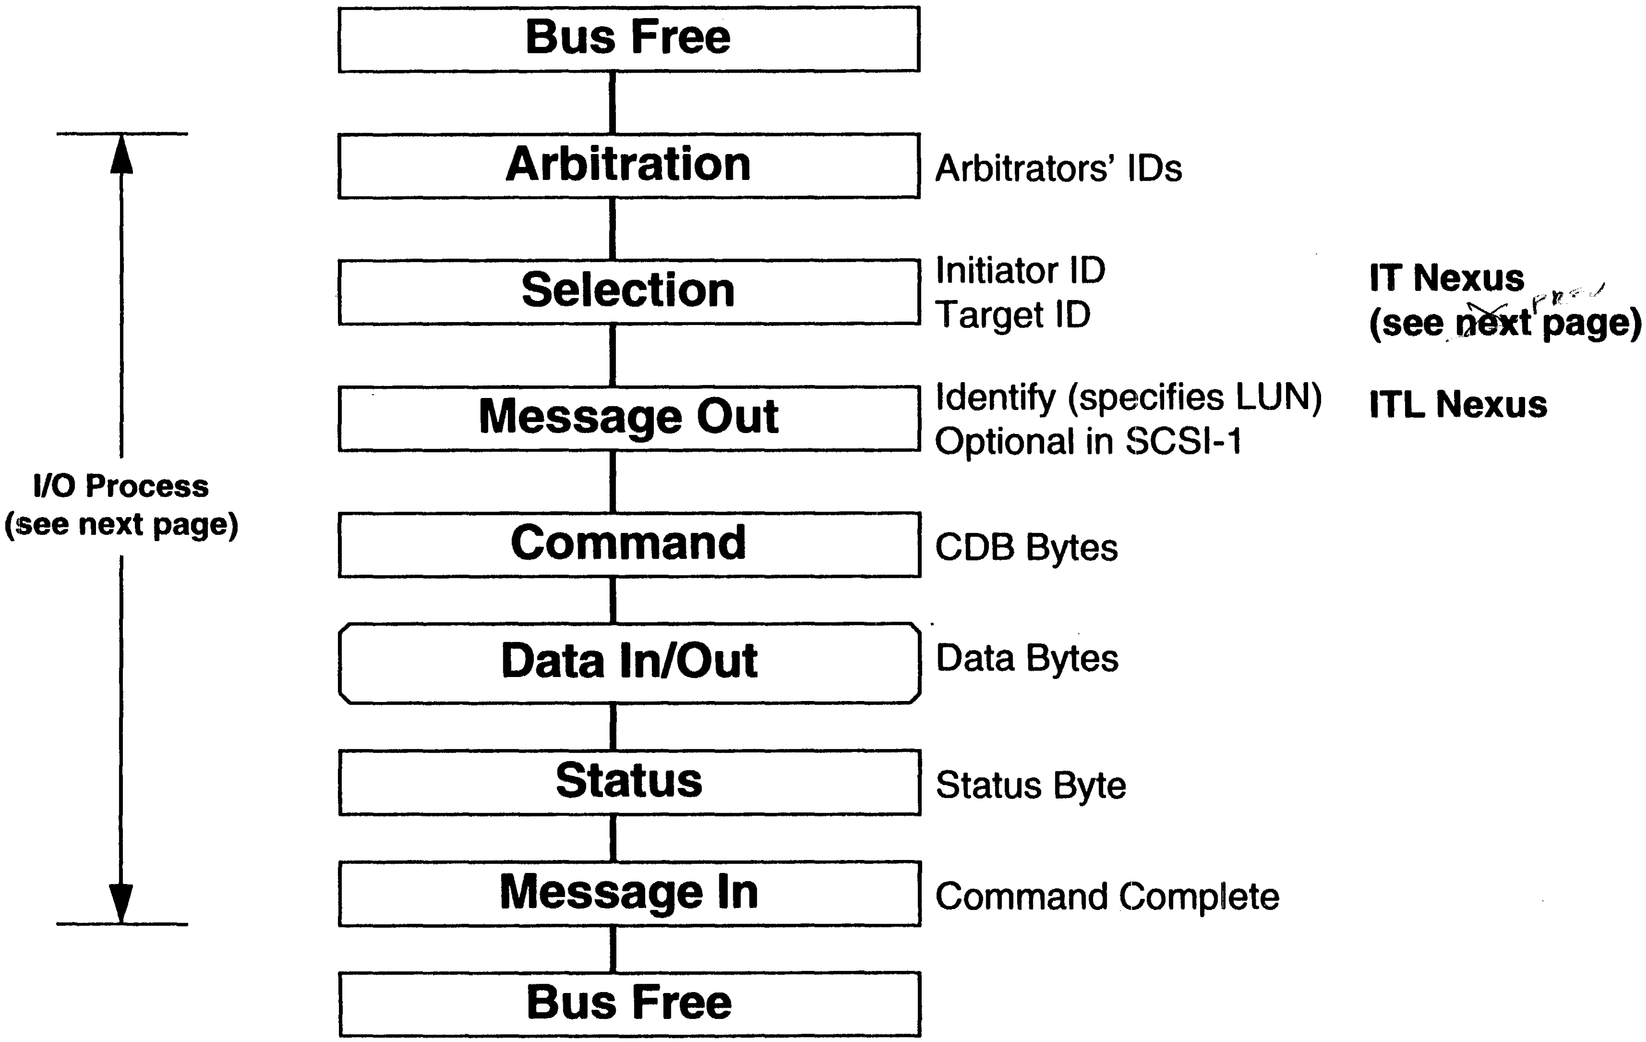

Typical SCSI Phase Sequence

(Continued)

With Disconnect/Reconnect

1. Bus Free
2. Arbitration
3. Selection
4. Message Out — Identify
5. Command
6. Message In — Disconnect
7. Bus Free

1. Bus Free
2. Arbitration
3. Reselection
4. Message In — Identify
5. Data
6. Status
7. Message In — Command Complete
8. Bus Free

I/O Process Example

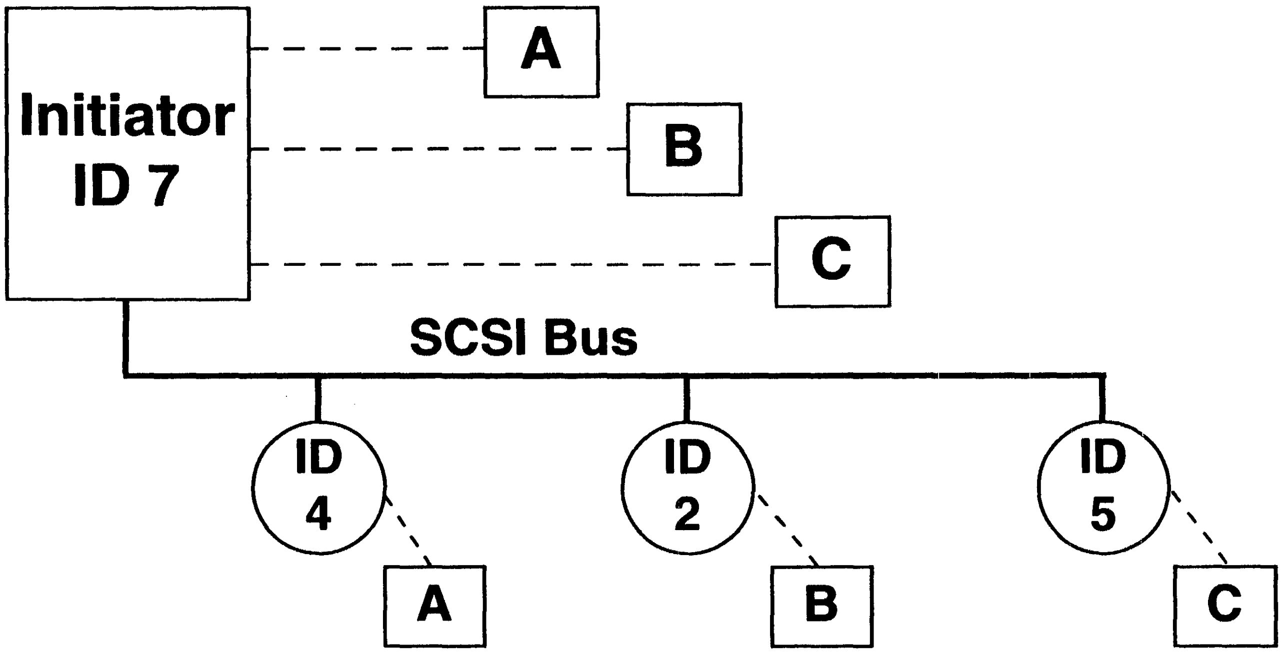

- The initiator issues command A to Device 4, command B to Device 2, and command C to Device 5. This creates 3 I/O Processes.
- In SCSI-3, the Initiator portion of an I/O process is called an Application Client and the target portion is called a Task.
- Initiators and Targets must remember the states of their I/O processes across disconnect.

Data Bus

Narrow Bus

DB-7	DB-6	DB-5	DB-4	DB-3	DB-2	DB-1	DB-0	DB-P

- 9 data signals:
  - 8 for actual data
  - 1 for parity
- Parity is odd when valid
- Data bits used for ______________ during Arbitration and Selection/ Reselection.
- Data bits used for ______________ during Information transfer phases.

Data Bus (Continued)

Wide Bus

DB-7	DB-6	DB-5	DB-4	DB-3	DB-2	DB-1	DB-0	DB-P
DB-15	DB-14	DB-13	DB-12	DB-11	DB-10	DB-9	DB-8	DB-P1

- 18 data signals:
  - 16 for actual data
  - 2 for parity (one per byte)

Parity

- Used for detecting transmission errors.
- Mandatory in SCSI-2 and SCSI-3, but in practice may be disabled for some applications.
- Valid during Selection, Reselection, and Information Transfer Phases.
- Invalid (not driven) during Bus Free and Arbitration.
- Always odd when supported (total number of 1 bits, including parity, must be odd).
- Sender generates parity and receiver checks it.
- Examples:

DB (hex)	DB (binary)	DB-P
00h	0000 0000	1
01h	0000 0001	0
15h	0001 0101	0
22h	0010 0010	1
FFh	1111 1111	1

Data Bus

Arbitration Examples

(How SCSI ID is used during Arbitration)

DB-7	DB-6	DB-5	DB-4	-	DB-3	DB-2	DB-1	DB-0	Hex	Arbitrating Device(s)	Winner
0	0	0	0		0	0	0	1	01	0	0
0	0	0	0		0	1	0	0	04	2	2
1	0	0	0		0	0	0	0	80	7	7
0	0	0	1		0	0	1	0	12	1,4	4
0	1	0	1		1	0	0	0	58	3,4,6	6
1	0	1	1		1	1	0	0	BC	2,3,4,5,7	7

Data Bus (Continued)

Selection and Reselection Examples

(How SCSI ID is used during Selection and Reselection)

DB-7	DB-6	DB-5	DB-4	DB-3	DB-2	DB-1	DB-0	Hex	Devices
0	0	0	0	0	0	1	1	03	0,1
0	0	0	1	1	0	0	0	18	3,4
0	0	1	1	0	0	0	0	30	4,5
0	1	0	1	0	0	0	0	50	4,6
0	1	1	0	0	0	0	0	60	5,6
1	1	0	0	0	0	0	0	C0	6,7

- Cannot distinguish initiator SCSI ID from target SCSI ID without examining the data bus during preceding Arbitration phase.
- During Selection, the initiator SCSI ID is the highest ID during the preceding Arbitration.
- During Reselection, the target SCSI ID is the highest ID during the preceding Arbitration

Data Bus Worksheet

0000=0h	0001=1h	0010=2h	0011=3h	0100=4h	0101=5h	0110=6h	0111=7h
1000=8h	1001=9h	1010=Ah	1011=Bh	1100=Ch	1101=Dh	1110=Eh	1111=Fh

Bus Phase		Data Bus											
phase	decided by	binary 7	6	5	4	3	2	1	0	hex		provided by	notes
Bus Free													
Arbitration										8	0		ID's=
Selection										8	4		iid= tid=
Bus Free													
Arbitration										1	0		ID's=
Reselection										9	0		iid= tid=
Bus Free													
Arbitration										C	0		ID's=
Selection										C	0		iid= tid=

This page is intentionally left blank.

Status

Status Values

00h	GOOD - (Successful I/O Process)
02h	CHECK CONDITION Error, Use REQUEST SENSE Command to Get Error Information
04h	CONDITION MET/GOOD For SEARCH DATA and PREFETCH Commands (normally unused)
08h	BUSY Target is Busy, Retry Command Later
10h	INTERMEDIATE GOOD Linked Commands (described later)
14h	INTERMEDIATE CONDITION MET/GOOD Linked Commands (described later)
18h	RESERVATION CONFLICT Device or Area Reserved by Another Initiator (described later)
22h	COMMAND TERMINATED (SCSI-2) I/O PROCESS TERMINATED (SCSI-3) After Terminate I/O Process Message (described later) Error, Use REQUEST SENSE Command to Get Error Information
28h	QUEUE FULL (SCSI-2) TASK SET FULL (SCSI-3) For Tagged Queueing (described later)
30H	ACA ACTIVE (SCSI-3)

Commands

Command Sets

Note: **In the SCSI-2 standard, there is one clause for each of the device classes. In SCSI-3, there is a separate standard for each device class.**

Device Class	SCSI-1	SCSI-2	SCSI-3	Description
General	X	X	X	Most Devices
Direct Access	X	X	X	Disk, Diskette, ...
Sequential Access	X	X	X	Tape
Printer	X	X	X	
Processor	X	X	X	Computer
Write Once	X	X	X	WORM
CD		X	X	CD-ROM, CD-R
Scanner		X	X	
Optical Memory		X	X	Erasable Optical
Medium Changer		X	X	Jukebox
Communications		X	X	Modem
Controller			X	Disk Array

Command Format

CDB = Command Descriptor Block

Bit / Byte	7	6	5	4	3	2	1	0	
0	Operation Code								What command
1	LUN			Command Dependent					LUN - see next page
2	Command Dependent								
. . .									
N - 2	Command Dependent								
N - 1	Control Byte								Used for Linking, Normally 00h

N = Command Length

Note: Reserved Fields Must Be Filled With 0

Command Format (Continued)

LUN Rules

- SCSI-1 — 3 Bit LUN, is normally supplied in the CDB
- SCSI-2 — 3 Bit LUN, is normally supplied in the Identify message
- SCSI-3 — 6 Bit LUN, must be supplied outside the CDB
LUN field in the CDB is Reserved
- In SCSI-1 and SCSI-2, if the Identify message is sent, the LUN in the CDB is ignored.
- A SCSI-2 or SCSI-3 initiator must supply the LUN in the Identify message.
- A SCSI-1 initiator may supply the LUN in the CDB or in the Identify message.
- A target may accept either method of specifying the LUN.

Command Format (Continued)

OP Code

- First CDB Byte = Operation Code

Bit	7	6	5	4	3	2	1	0
	Group Code			Command Code				

- Group Codes (Indicating CDB Length):

000	0	6-Byte Commands
001	1	10-Byte Commands
010	2	10-Byte Commands
011	3	Reserved
100	4	16-byte (SCSI-3)
101	5	12-Byte Commands
110	6	Vendor Unique
111	7	Vendor Unique (See Product Description Manual)

Command Format (Continued)

CDB Length Examples

Op Code Hex	Op Code Binary	Group Code	CDB Length
00h	**000**0 0000	0	6
12h	**000**1 0010	0	6
08h	**000**0 1000	0	6
28h	**001**0 1000	1	10
B1h	**101**1 0001	5	12
FFh	**111**1 1111	7	?

Command Format (Continued)

Control Byte

- Last CDB Byte - Control Byte

Bit	7	6	5	4	3	2	1	0
	Vendor Unique		Reserved			NACA	Flag	Link

- Typically used for Linked Commands (described later)
- Normally filled with 00h
- NACA: (SCSI-3 only, described later)

Test Unit Ready Command

CDB Format

Byte \ Bit	7	6	5	4	3	2	1	0
0	Operation Code = 00h							
1	LUN			Reserved				
2	Reserved							
3	Reserved							
4	Reserved							
5	Control Byte							

- Check if Device is Ready
- All Device Types
- Mandatory Command

Test Unit Ready Command

(Continued)

Bus Phases

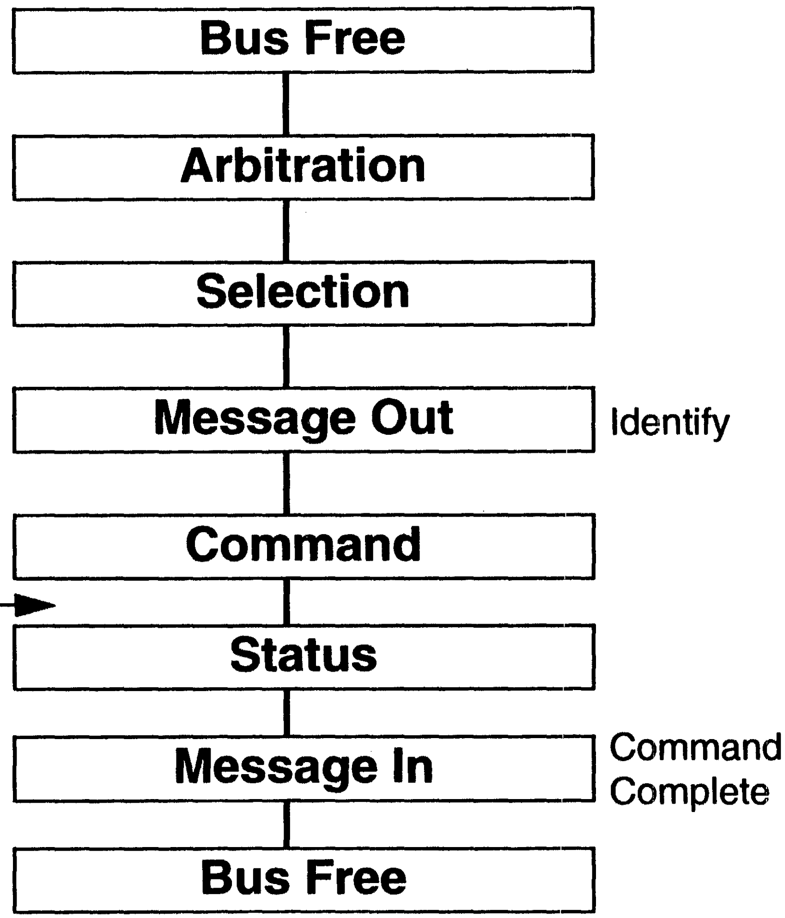

Note:
No data phase!

This page is intentionally left blank.

Inquiry

Inquiry Command

CDB Format

Byte \ Bit	7	6	5	4	3	2	1	0
0	Operation Code = 12h							
1	LUN			Reserved			CmdDt (SCSI-3)	EVPD (SCSI-2)
2	Vital Page Code/OpCode							
3	Reserved							
4	Allocation Length							
5	Control Byte							

- Mandatory Command
- EVPD = Enable Vital Product Data (normally not supported, filled with 0)
- CmdDt = Command Support data (Optional, see SPC, SCSI-3)
- Vital Page Code if EVPD = 1, or OP Code if CmdDT=1, else 0's
- Allocation Length = Maximum Number of Data In Bytes Returned (see next page)

Important Note: Data Should be Returned with good status even if Device is NOT Ready

Inquiry Command (Continued)

Allocation Length Examples

Allocation Length Bytes in the CDB	Number of Inquiry Bytes at the Target	Number of Bytes Returned
24h	24h	24h
01h	24h	01h (first)
00h	24h	00h
FFh	24h	24h

Inquiry Command (Continued)

Data In Format (SCSI-2)

Bit / Byte	7	6	5	4	3	2	1	0
0	Qualifier			Device Type				
1	RMB	Device Modifier (SCSI-1)						
2	ISO Version		ECMA Version			ANSI Version		
3	AENC	TrmIOP	Reserved		Response Data Format			
4	Additional Length							
5	Vendor-Unique							
6	Reserved							
7	Rel Adr	WB32	WB16	Sync	Link	Rsvd	CmdQue	Sft Rst
8-15	Vendor ID (ASCII)							
16-31	Product ID (ASCII)							
32-35	Revision (ASCII)							
36-55	Vendor-Unique							
56-95	Reserved							
96-	Vendor-Unique							

Note:
See next page for description of fields

Inquiry Command (Continued)

- Device Type:

00h = Direct Access	06h = Scanner
01h = Sequential Access	07h = Optical Memory
02h = Printer	08h = Medium Changer (Jukebox)
03h = Processor	09h = Communications
04h = Write Once	0A - 0Bh = Pre-Press Devices
05h = CD-ROM	0C = Controller (Disk Array)
	1Fh = Unknown or No Device

- Qualifier: 0h = Logical Unit is Connected
  1h = Logical Unit is not Connected
  3h = Logical Unit is not Supported
- Device Modifier: Usually 00h (no longer supported)
- RMB: 1h = Removable Medium, 0h = Not Removable
- ANSI Version: 1h ☛ SCSI-1, 2h ☛ SCSI-2, 3h ☛ SCSI-3
- ECMA Version: European Computer Manufacturers Association, Normally 0
- ISO Version: International Standards Organization, Normally 0
- Response Data Format: 0h ☛ SCSI-1, 1h ☛ CCS, 2h ☛ SCSI-2, 3h ~~☛ SCSI-3~~
  OR SCSI-3

Inquiry Command (Continued)

TRMIOP	1=Terminate I/O Process Supported
AENC	1=Asynchronous Event Notification Capability Supported
Additional Length	Number of bytes of the following data that is available
SftRst	1=Soft Reset Supported (SCSI-2 only)
CmdQue	1=Tagged Queueing Supported
Link	1=Linked Commands Supported
Sync	1=Synchronous Supported
WB16	1=16 Bit Wide Bus Supported
WB32	1=32 Bit Wide Bus Supported
RelAdr	1=Relative Addressing Supported

Note: The features these bits indicate support for are described later in the class.

Request Sense

SCSI Error Reporting

- SCSI devices report errors by sending a status byte of 02h (CHECK CONDITION) to the initiator during status phase.
- Initiator retrieves information about the error, called sense data, by issuing a REQUEST SENSE command.
- Sense data is normally 18 bytes long.
- The Sense Key field in sense data identifies the basic type of error, e.g. not ready, hardware error, data error.
- The ASC/ASCQ bytes define the specific error, e.g. not ready because no media present.

Request Sense Command

Bus Phases

The Command in Error		The Request Sense Command	
	Bus Free	Bus Free	
	Arbitration	Arbitration	
	Selection	Selection	
(Identify)	Message Out	Message Out	(Identify)
	Command	Command	(Request Sense)
(Error) (Build Sense Data)	Data	Data In	(Sense Data)
(Check Condition)	Status	Status	(Good)
(Command Complete)	Message In	Message In	(Command Complete)
	Bus Free	Bus Free	

Request Sense Command (Continued)

CDB Format

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code = 03h							
1	LUN			Reserved				
2	Reserved							
3	Reserved							
4	Allocation Length							
5	Control Byte							

Request Sense Command (Continued)

Data In (Sense Data) Format

Bit / Byte	7	6	5	4	3	2	1	0
0	Valid	Error Code						
1	Segment Number							
2	FM	EOM	ILI	Rsvd	**Sense Key**			
3-6	Information Bytes (MSB-LSB)							
7	Additional Sense Length							
8-11	Command Specific Data							
12	**Additional Sense Code (ASC)**							
13	**Additional Sense Code Qualifier (ASCQ)**							
14	Field replaceable Unit (Vendor-Unique)							
15-17	Sense-Key Specific							
18 -	Vendor-Unique							

Note:
See next page for description of fields

Request Sense Command (Continued)

Sense Key (Byte 2, Bits 0-3)

0h	No Sense (No error, FM, EOM, or ILI)
1h	Recovered Error
2h	Not Ready
3h	Medium Error
4h	Hardware Error
5h	Illegal Request
6h	Unit Attention
7h	Data Protect
8h	Blank Check (tape, WORM,...)
9h	Vendor Unique
Ah	Copy Aborted (error with Copy, Compare, or Copy and Verify)
Bh	Aborted Command
Ch	Equal (Search Data command)
Dh	Volume Overflow (buffered commands)
Eh	Miscompare
Fh	Reserved

Request Sense Command (Continued)

Example Additional Sense Codes - Partial List

ASC = Additional Sense Code (Byte 12)

ASCQ = Additional Sense Code Qualifier (Byte 13)

ASC	ASCQ	DESCRIPTION
17h	00h	Recovered data with no correction
17h	01h	Recovered data with retries
18h	00h	Recovered data with correction
18h	01h	Recovered data with correction and retries
20h	00h	Invalid Command OP Code
21h	00h	Invalid LBA
24h	00h	Invalid field in CDB
29h	00h	Power On or Reset or Bus Device Reset
43h	00h	Message Error
48h	00h	Initiator Detected Error Message Received
4Ah	00h	Command Phase Error
4Bh	00h	Data Phase Error

Request Sense Command (Continued)

◆ **Error Code**	Byte 0, Bits 0 - 6	
	For SCSI-1 and CCS - 70h	
	For SCSI-2 and SCSI-3 - 70h current	
	- 71h deferred error	
◆ **Valid**	Byte 0, Bit 7	
	Valid Information Bytes (Bytes 3-6)	
◆ **Information Bytes**	Bytes 3 - 6	
	Used when Valid bit is set, and contains:	
	❖ LBA of Error	(Direct Access, WORM,...)
	❖ Residue	Requested minus Actual Length (Tape, Printer,...) - described later
	❖ Residue	Requested minus Actual Blocks (Copy, Compare, Copy & Verify commands)
◆ **Segment**	Byte 1	
	For Copy and Compare commands (normally 0)	

Request Sense Command (Continued)

- **FM** Byte 2, Bit 7 (described later)
  Filemark - tape
- **EOM** Byte 2, Bit 6 (described later)
  End of Medium
  - ❖ End of tape, beginning of tape, out of paper, etc.
  - ❖ Tape and Printer use only
- **ILI** Byte 2, Bit 5 (described later)
  Incorrect Length Indicator
  - ❖ Requested block length does not match logical block length on medium
  - ❖ Normally tape only
- **Sense Key Specific** Bytes 15-17 (Optional - see SCSI standard for details)
  Depends on Sense Key value:
  - ❖ Field Pointer if Illegal Request sense key
  - ❖ Actual Retry Count if Recovered, Medium, or Hardware Error sense key
  - ❖ Format Immediate Progress Indication if Not Ready sense key

Contingent Allegiance

(A Feature for Sense Data Preservation)

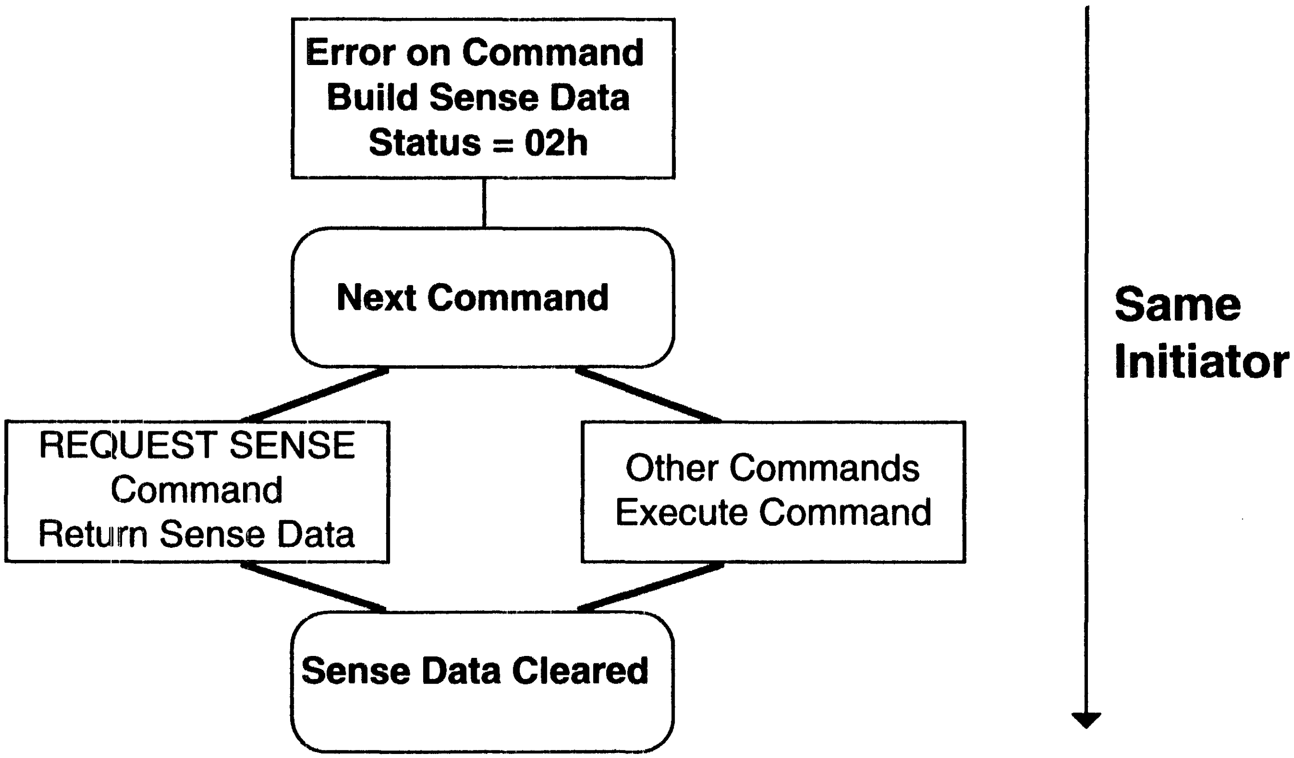

Contingent Allegiance (Continued)

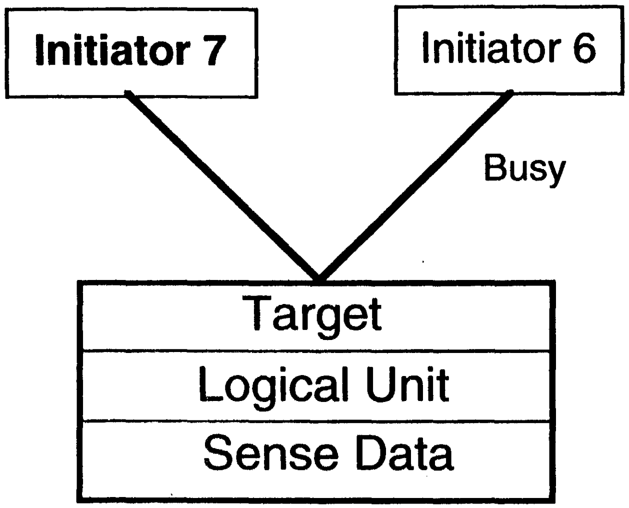

One Sense Bucket

OR

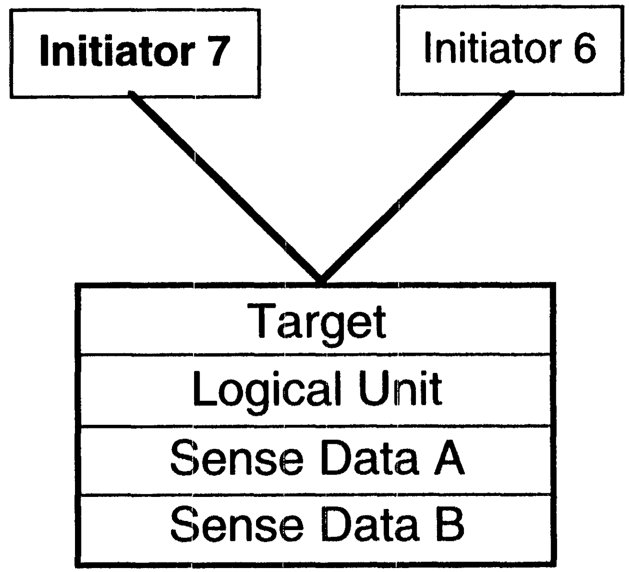

Multiple Sense Buckets

Contingent Allegiance (Continued)

- A feature provided by the target to preserve sense data for an initiator after Check Condition or Command Terminated Status, in case other initiators attempt to access the same logical unit in the mean time.
- Implemented using multiple sense buckets or Busy status.
- Cleared By:
  - Request Sense from same initiator
  - Any command other than Request Sense from same initiator
  - Power cycle
  - Hard Reset event
  - Bus Device Reset message (described later)
  - Abort message (described later)

Unit Attention

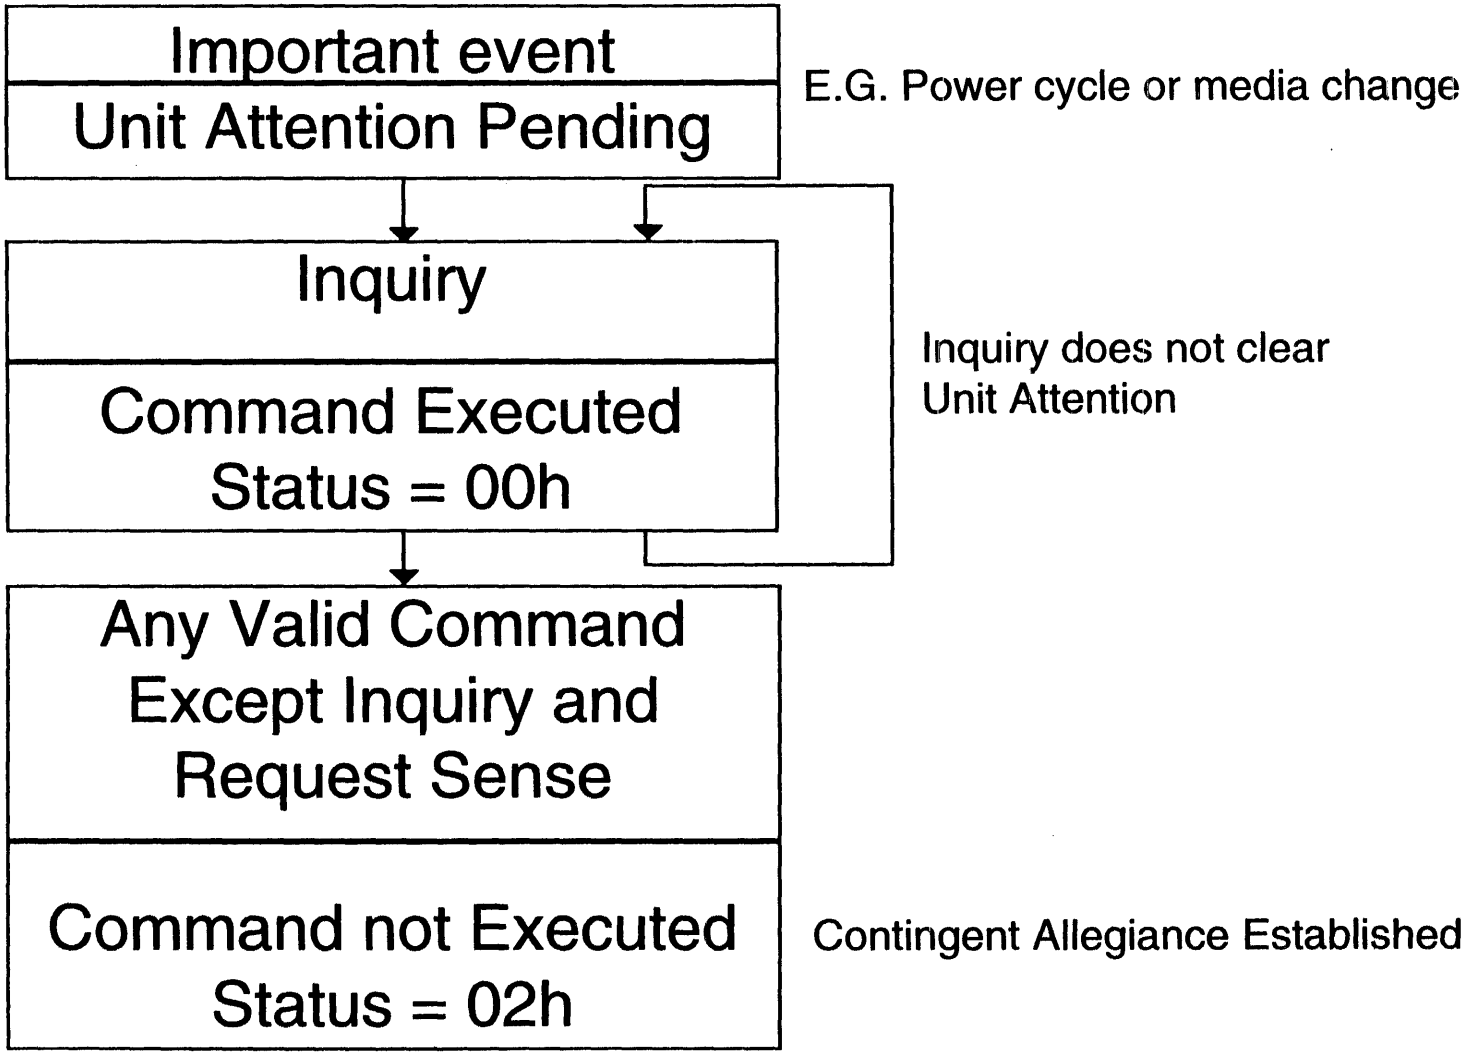

Unit Attention (Continued)

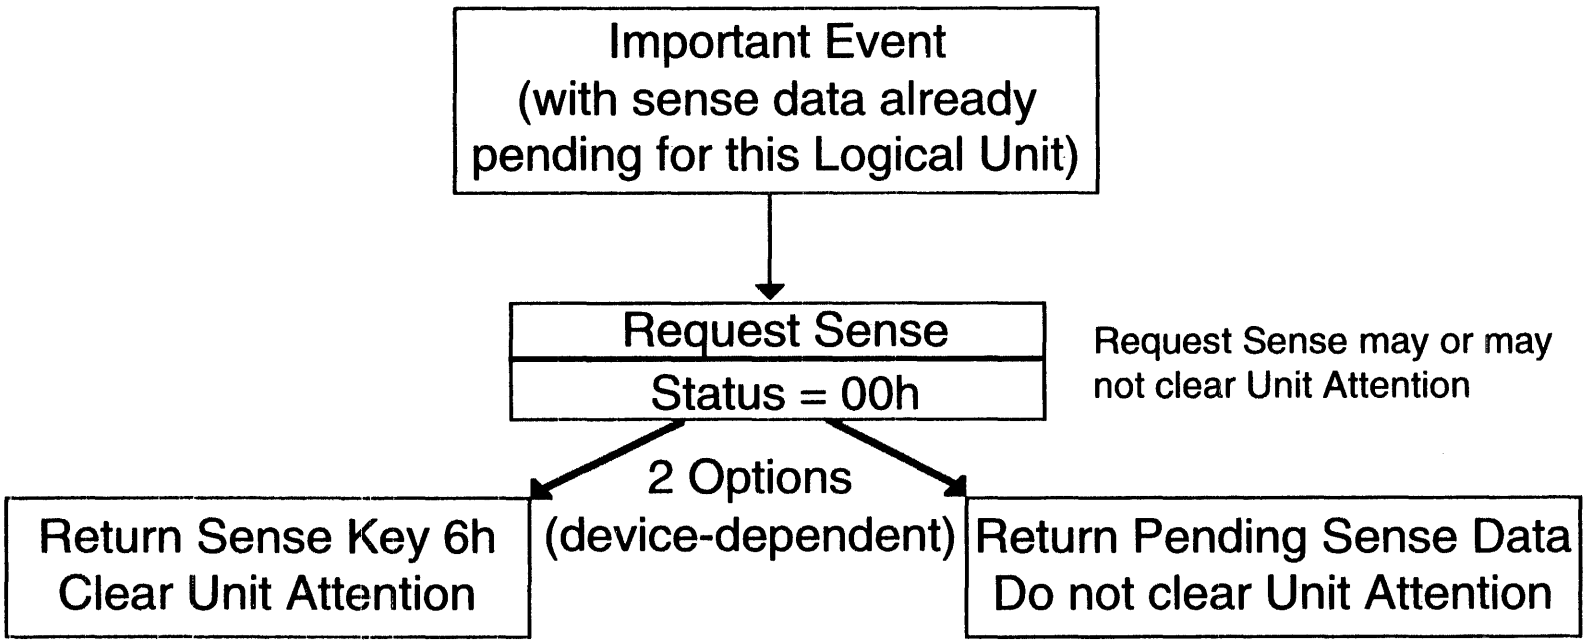

- Unit Attention can occur when another Sense Key is already pending.

Question: Which Sense Key should the device report?

Unit Attention (Continued)

Summary

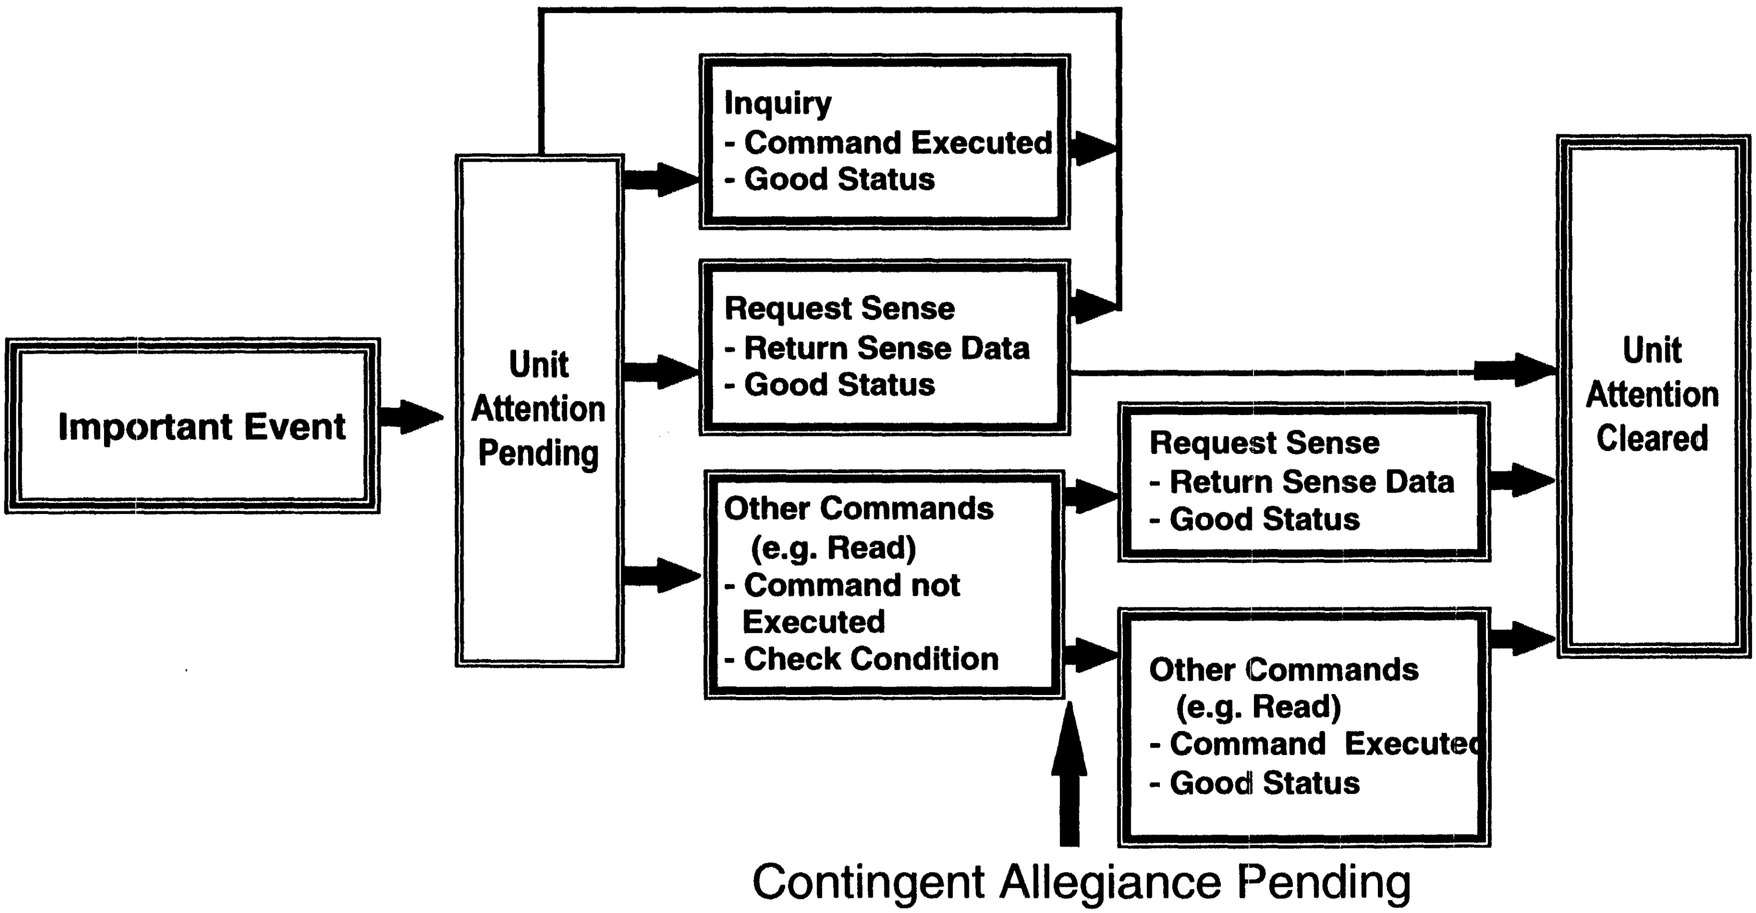

Unit Attention (Continued)

Examples

Does the command sequence shown clear Unit Attention (circle one):

Inquiry, Inquiry	yes	no	maybe
Request Sense	yes	no	maybe
Inquiry, Request Sense, Inquiry	yes	no	maybe
Test Unit Ready, Request Sense	yes	no	maybe
Read, Read	yes	no	maybe
Write, Inquiry	yes	no	maybe

Start/Stop Command

CDB Format

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code = 1Bh							
1	LUN			Reserved				IM (Imm)
2	Reserved							
3	Reserved							
4	Reserved						LoEj	Start
5	Control Byte							

- IM Immediate
- Start 1 = Start, 0 = Stop
- LoEj 1 = Load removable media (when Start = 1)
  Eject removable media (when Start = 0)
- If the drive is stopped, it responds with Check Condition Status (Not Ready Sense Key) to media access commands.

Start/Stop Command (Continued)

Non-Immediate

Start CMD

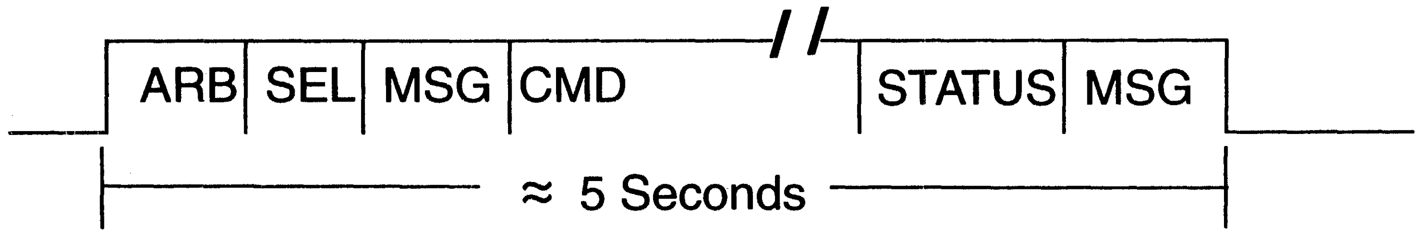

Immediate

Start CMD

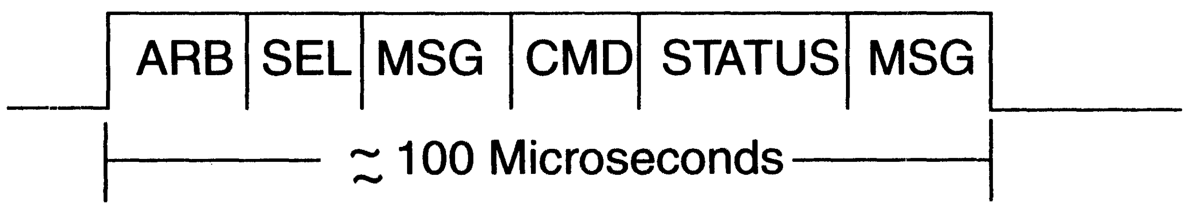

- In immediate Mode, an error detected afterwards is reported on the next command with Check Condition Status and Deferred Error Code Sense

Read Capacity Command

CDB Format

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code 25h							
1	LUN			Reserved				RA
2	LBA (MSB)							
3	LBA							
4	LBA							
5	LBA (LSB)							
6	Reserved							
7	Reserved							
8	Reserved							PMI
9	Control Byte							

- RA, LBA, and PMI fields are not commonly used and are normally filled with 0 (see SCSI Standard for details)

Read Capacity Command (Continued)

Data In Format
(When PMI = 0 in the CDB)

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	Max LBA							
4 - 7	Block Length in Bytes							

Disk Read & Write

Disk - Read Command

Bus Phases

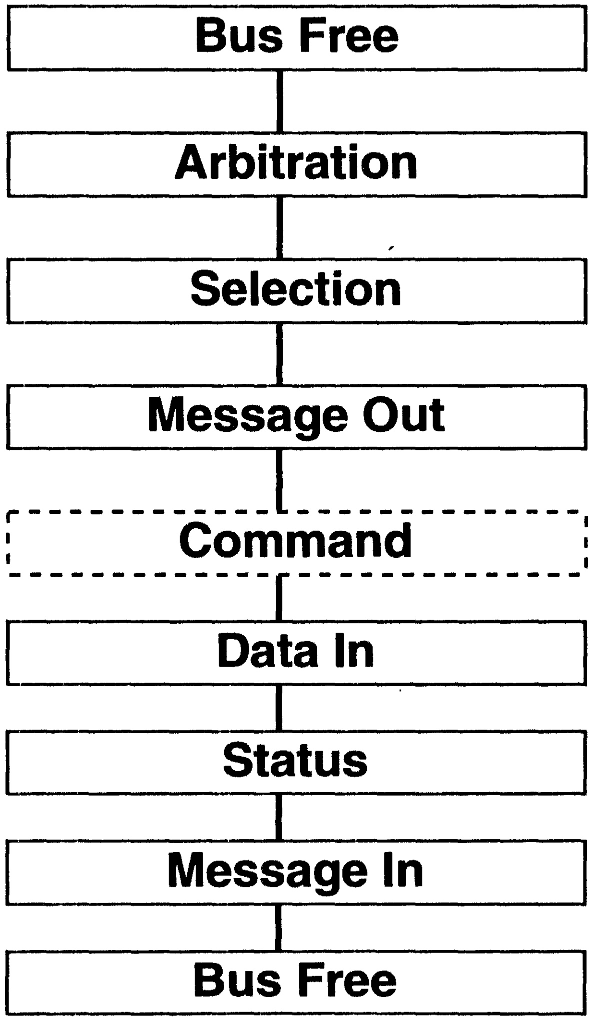

Disk - Read Command (Continued)

CDB Format
6-Byte Read

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code = 08h							
1	LUN			LBA (MSB)				
2	LBA							
3	LBA (LSB)							
4	Number of Blocks *							
5	Control Byte							

- LBA = Starting Logical Block Address
- Max accessible drive capacity is 1GB using 512 -byte length
- *Number of Blocks Value of 0 Indicates 256

Disk - Read Command (Continued)

CDB Format

10-Byte Read

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code = 28h							
1	LUN			DPO	FUA	Reserved		RA
2	LBA (MSB)							
3	LBA							
4	LBA							
5	LBA (LSB)							
6	Reserved							
7	Number of Blocks (MSB)							
8	Number of Blocks (LSB)							
9	Control Byte							

Note: See next page for description of fields.

Disk - Read Command (Continued)

- Number of Blocks value of 0 means 0 (no data transfer)
- Max accessible drive capacity 2TB using 512 -byte length
- RA = Relative Address (Linked Commands Only) - Normally 0
- **FUA** Force Unit Access - Normally 0
  1 = Read from Media
  0 = Read from Cache allowed
- **DPO** Disable Page Out - Normally 0
  1 = Replace Data (will not need block in near future)
  0 = Try to Keep Data in Cache

Disk-Write Command

Bus Phases

Bus Free

Arbitration

Selection

Message Out

Command

Data Out

Status

Message In

Bus Free

Implied Seek

Disk-Write Command (Continued)

CDB Format

6 - Byte Write

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code = 0Ah							
1	LUN			LBA (MSB)				
2	LBA							
3	LBA (LSB)							
4	Number of Blocks							
5	Control Byte							

- Number of Blocks value of 0 Indicates 256

Disk-Write Command (Continued)

CDB Format

10 - Byte Write

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code = 2Ah							
1	LUN			DPO	FUA	Reserved		RA
2	LBA (MSB)							
3	LBA							
4	LBA							
5	LBA (LSB)							
6	Reserved							
7	Number of Blocks (MSB)							
8	Number of Blocks (LSB)							
9	Control Byte							

- Number of Blocks value of 0 Indicates 0

CDB Data Addressability

512 Bytes/Logical Block

- 6 Byte CDB
  - LBA = 21 bits or 2,097,151 = 1GB addressing
  - Transfer Count = 8 bits or 256 blocks
- 10 Byte CDB
  - LBA = 32 bits or 4,294,967,296 = 2TB addressing
  - Transfer Count = 16 bits or 64k blocks
- 12 Byte CDB
  - LBA= 32 bits or 4,294,967,296 = 2TB addressing
  - Transfer Count = 32 bits or 4G blocks
- 16 Byte CDB
  - Same data addressability as a 10 byte CDB

Copy Command

- Optional Command
- Copy Manager
  - Logical Unit Receiving and Performing the Copy
- Source Device
  - Logical Unit to Copy from
- Destination Device
  - Logical Unit to Copy to
- Initiator Gives Copy Manager During Data Out:
  - Source SCSI ID and LUN
  - Destination SCSI ID and LUN
  - What to copy and how

Copy Command (Continued)

Example

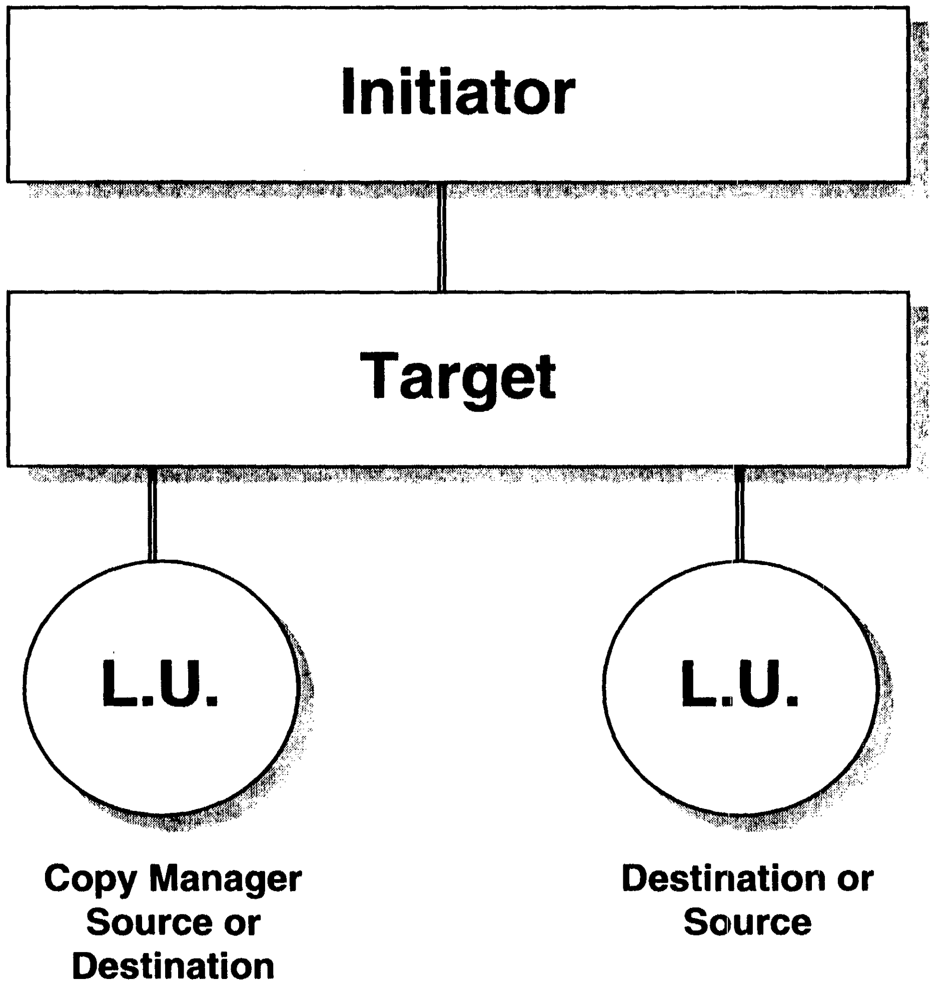

Copy Command (Continued)

Example

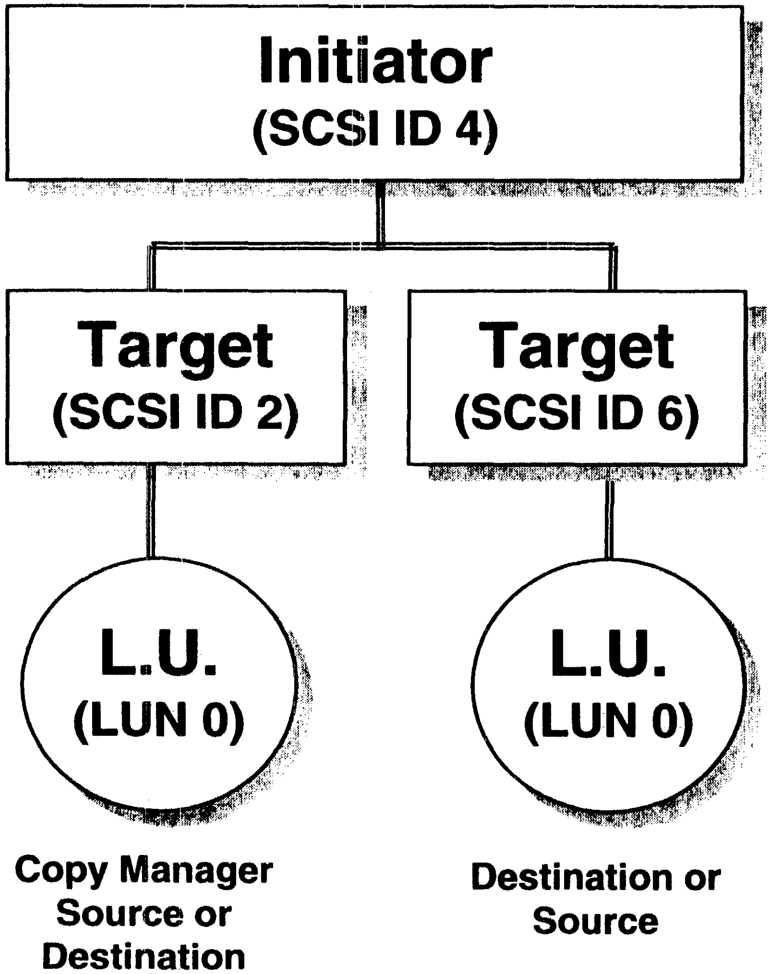

- Initiator 4 Connects to Target 2 LUN 0: Copy from Target 6 LUN 0 to Target 2 LUN 0.
- Target 2 Disconnects.
- Target 2 Becomes Initiator and Connects to Target 6 Lun 0: Execute Read Command.
- Target 6 Disconnects after Finishing Read.
- Target 2 Reconnects to Initiator 4 to Report Status.

Reserve/Release Commands

- Used for multiple command sequences that need to appear as a single operation to other initiators.
- Reserve logical unit with Reserve command and release with Release command.
- When reserved, drive responds with Reservation Conflict (18h) status to most commands from other initiators.
- A Release command to a device which is not reserved or reserved by another initiator results in Good status and no change to reservation.
- Request Sense, Inquiry, Prevent/Allow Medium Removal (Allow option only), executed even if reserved.
- Power cycle, Hard Reset, and Bus Device Reset message clear reservations.

This page is intentionally left blank.

Mode Sense / Mode Select

Mode Parameters

- SCSI permits great flexibility in configuring devices and controlling their behavior through the use of **Mode Parameters**
- Related parameters are grouped together into data structures called **Mode Pages**

Examples:

Caching Page	**Disconnect-Reconnect Page**	**R/W Error Recovery Page**
Parameters affecting read and write caching by the device	Parameters affecting when the device disconnects and reconnects	Parameters affecting retry, ECC, and error reporting

Mode Sense and Mode Select Commands

◆ Mode Sense	Query Logical Unit Parameters	
◆ Mode Select	Modify Logical Unit Parameters	
◆ Parameter Values	Default	factory settings
	Saved	settings customized and saved by user
	Current	settings currently in use

◆ For Power On or Reset, if Saved Values available, use them, else use Default

Question: Should the user be able to modify **any** parameters in the device?

Mode Sense and Mode Select Data

Data Format

Mode Parameter Header	4 bytes
Block Descriptor(s)	8 bytes normally
Page(s)	multiple bytes

Mode Sense and Mode Select Data

(Continued)

Mode Parameter Header

Bit / Byte	7	6	5	4	3	2	1	0	Typical Values: Mode Sense	Typical Values: Mode Select
0	Mode Data Length								xxh	Reserved
1	Medium Type								00h	00h
2	Device-Specific Parameter								00h	00h
3	Block Descriptor Length								08h	08h

Note: See next page for description of fields.

Mode Sense and Mode Select Data
(Continued)

- Mode Data Length:
  - For Mode Sense: number of bytes of the following data that is available
  - For Mode Select: Reserved
- Medium Type:
  - Depends on the device type (see SCSI standard)
  - For disk: Rigid = 00; floppy = sides, size, etc.
  - For tape: Reserved
- Device Specific Parameter:
  - For disk: normally 00h
  - For tape: buffered mode (described later), etc.
- Block Descriptor Length:
  - Normally 08h
  - Must be multiple of 8

Mode Sense and Mode Select Data
(Continued)

Block Descriptor

Bit / Byte	7	6	5	4	3	2	1	0	Typical Values
0	Density Code								00h
1	Number of Blocks (MSB)								00h
2	Number of Blocks								00h
3	Number of Blocks (LSB)								00h
4	Reserved								00h
5	Block Length (MSB)								00h
6	Block Length								02h
7	Block Length (LSB)								00h

Note: See next page for description of fields.

Mode Sense and Mode Select Data
(Continued)

- Density Code:
  - Depends on type of device
  - For disk = 00h (reserved)
  - For tape = type of media and recording, e.g., DDS (DAT)
- Number of Blocks:
  - Normally 0, which means that all logical blocks of the logical unit have the same medium characteristics and same size
- Block Length:
  - Number of bytes per logical block
  - Normally 512 (200h)

Mode Sense and Mode Select Data
(Continued)

Page Format

Byte	7	6	5	4	3	2	1	0
0	PS	R	Page Code					
1	Page Length (number of bytes following)							
2 -	Page Contents							

R = Reserved (0)

PS = Parameters Saveable

for Mode Sense:

1 = Page is saveable

0 = Page is not saveable

for Mode Select: 0

Mode Sense Command
CDB Format

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code = 1Ah							
1	LUN			Rsvd	DBD	Reserved		
2	PC		Page Code					
3	Reserved							
4	Allocation Length							
5	Control Byte							

- DBD — Disable Block Descriptors, not supported normally (0)
- Page Code — Desired page code, or 3F = all pages
- PC — Page Control: 0=Current, 1=Changeable, 2=Default, or 3=Saved

Mode Sense Command

(Continued)

CDB Examples

To receive Header and Block Descriptor (assuming block length 512):

hex CDB

expected hex data

To receive Header, Block Descriptor, and Current page 21h (assume page 21h is 5 bytes long and is saveable)

hex CDB

expected hex data

Binary Work Area

7	6	5	4	3	2	1	0

Mode Sense Command

(Continued)

Answers to CDB Examples

To receive Header and Block Descriptor:

hex CDB

1A	00	00	00	0C	00

expected hex data

0B	00	00	08	00	00	00	00	00	00	02	00

To receive Header, Block Descriptor, and Current page 21h (assume page 21h is 5 bytes long and is saveable)

hex CDB

1A	00	21	00	FF	00

expected hex data

10	00	00	08	00	00	00	00	00	00	02	00

A1	03	XX	XX	XX

Mode Page Example

Page 02H, Disconnect/Reconnect Parameters

Byte \ Bit	7	6	5	4	3	2	1	0	Default Values PC=2	Changeable Value PC=1
	Parameter Description									
0	PS	Rsvd	Page Code						82	82
1	Page Length								0A	0A
2	Buffer Full Ratio								80	FF
3	Buffer Empty Ratio								80	FF
4 - 5	Bus Inactivity Limit								0004	FFFF
6 - 7	Disconnect Time Limit								0000	FFFF
8 - 9	Connect Time Limit								0000	FFFF
10 - 11	Maximum Burst Size								0000	FFFF

Mode Select Command

CDB Format

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code = 15h							
1	LUN			PF	Reserved			SP
2	Reserved							
3	Reserved							
4	Parameter List Length							
5	Control Byte							

- **PF** Page Format, 1 if sending pages, else ignored
- **SP** Save Pages, 1 to save, 0 not to save
- Parameter List Length: how many data bytes the initiator wants to send

Mode Select Command (Continued)

CDB Examples
(Without saving pages)

To send Header only:

hex CDB

hex data

To send Header, Block Descriptor, page 20h (3 bytes), and page 21h (5 bytes), saving the pages:

hex CDB

hex data

Binary Work Area

7	6	5	4	3	2	1	0

Mode Select Command (Continued)

CDB Example

To send Header only:

hex CDB | 15 | 00 | 00 | 00 | 04 | 00 |

hex data | 00 | 00 | 00 | 00 |

To send Header, Block Descriptor, page 20h (3 bytes), and page 21h (5 bytes), saving the pages:

hex CDB | 15 | 11 | 00 | 00 | 14 | 00 |

hex data | 00 | 00 | 00 | 08 | | 00 | 00 | 00 | 00 | 00 | 00 | 02 | 00 |

| 20 | 01 | XX | 21 | 03 | XX | XX | XX |

Devices

This page is intentionally left blank.

Hard Disk

Hard Disk Layout

Tracks and Cylinders

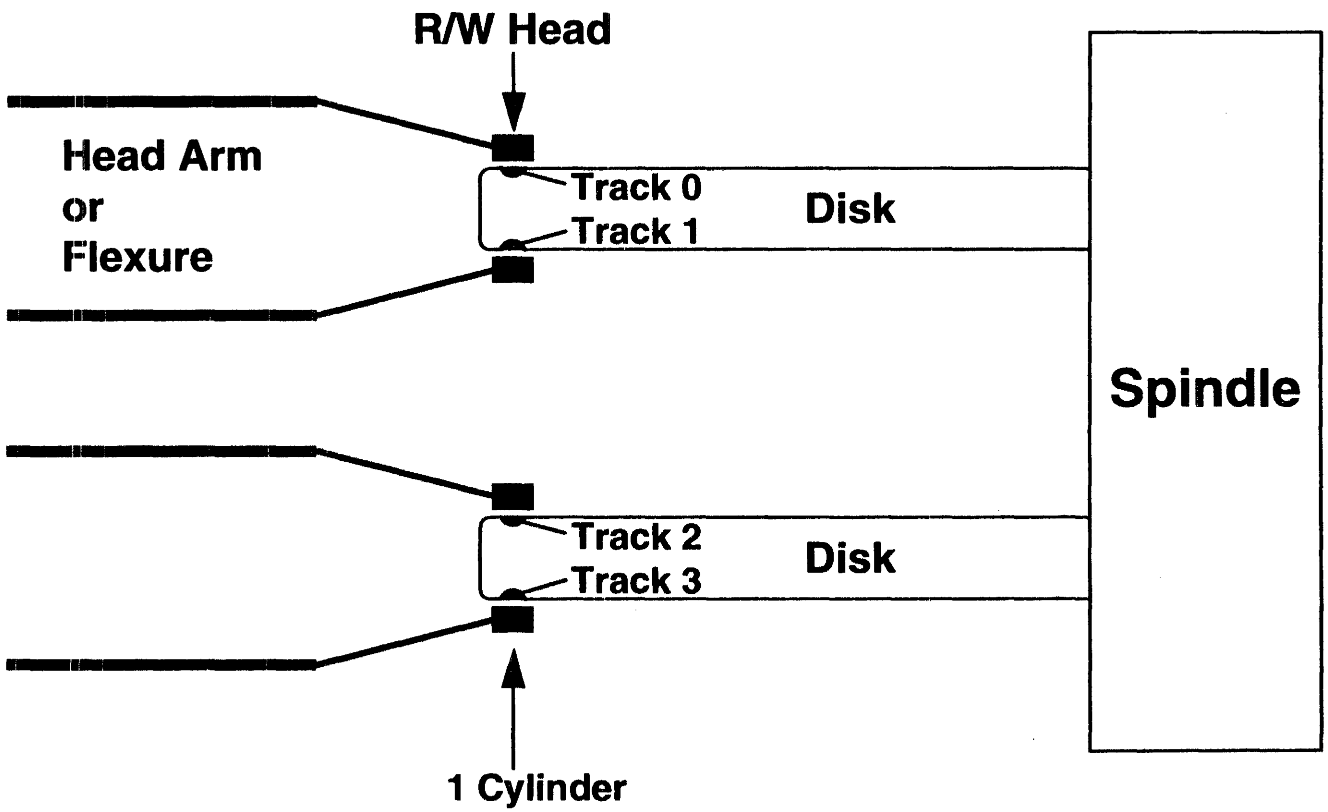

Hard Disk Layout (Continued)

Disk Surface Usage

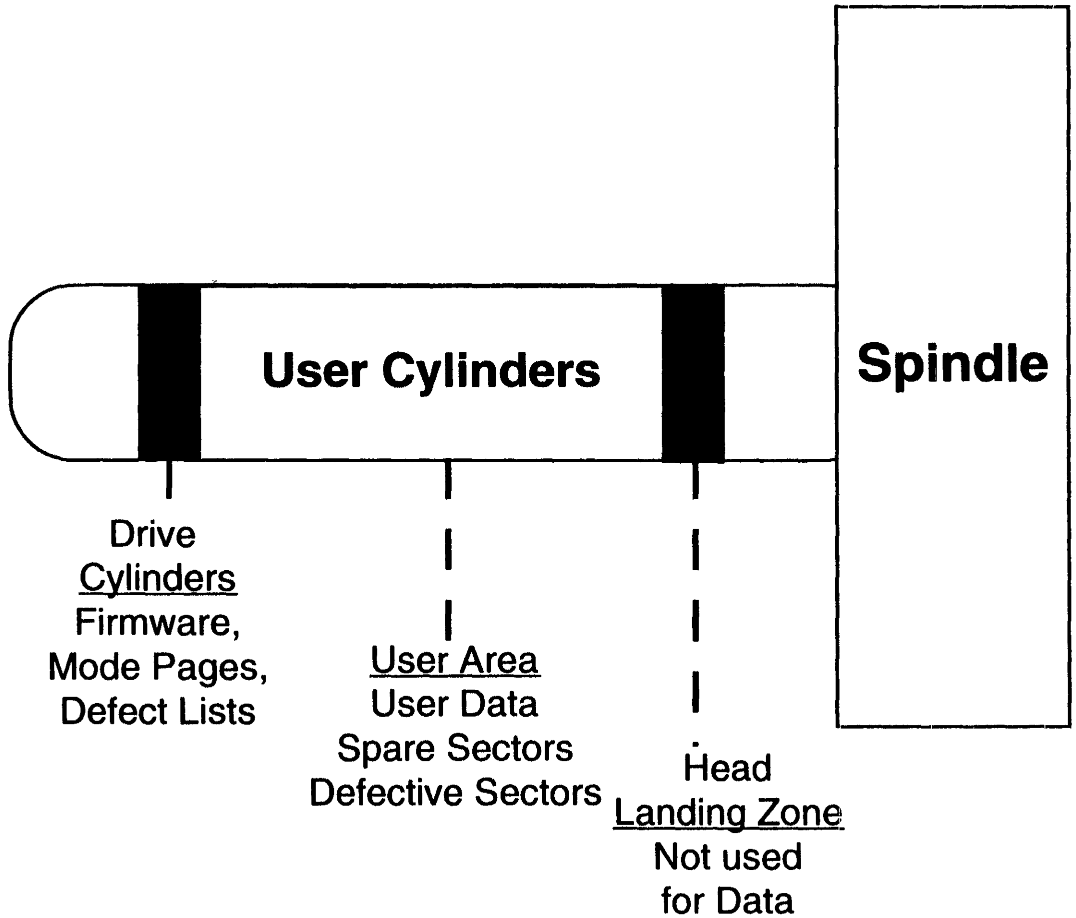

Hard Disk Layout (Continued)

Single Track

Note: A track is made up to many physical recording sectors. Normally each sector contains one Logical block.

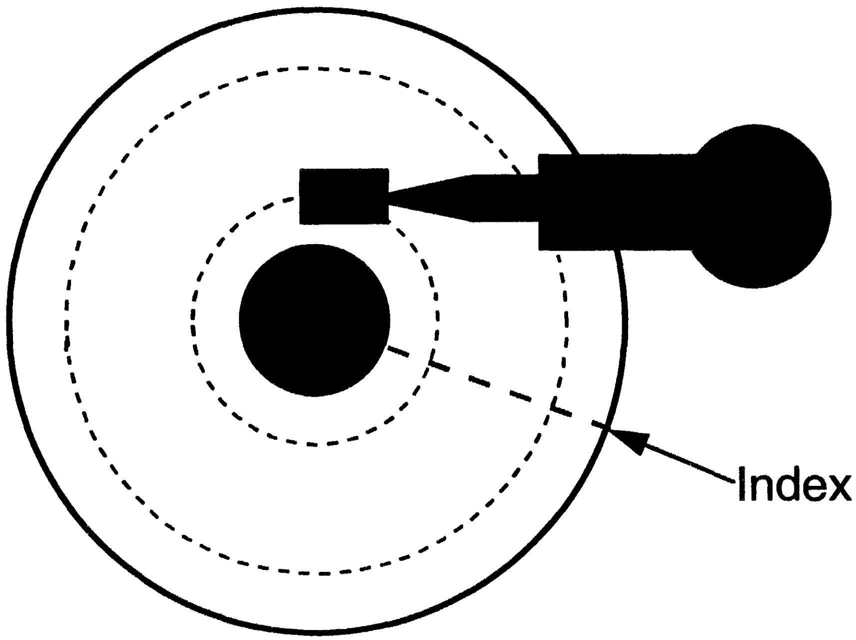

Top View

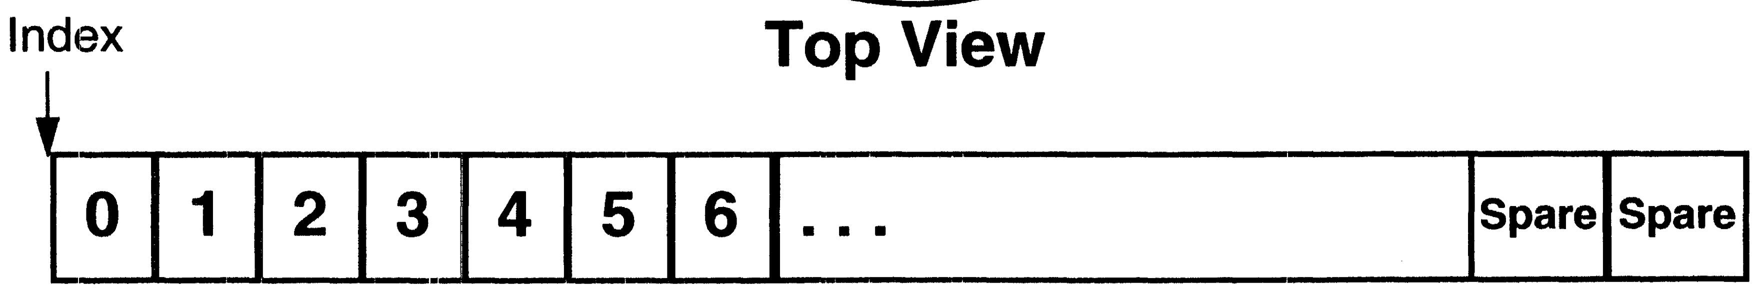

Logical View

Typical Hard Disk Sector Layout

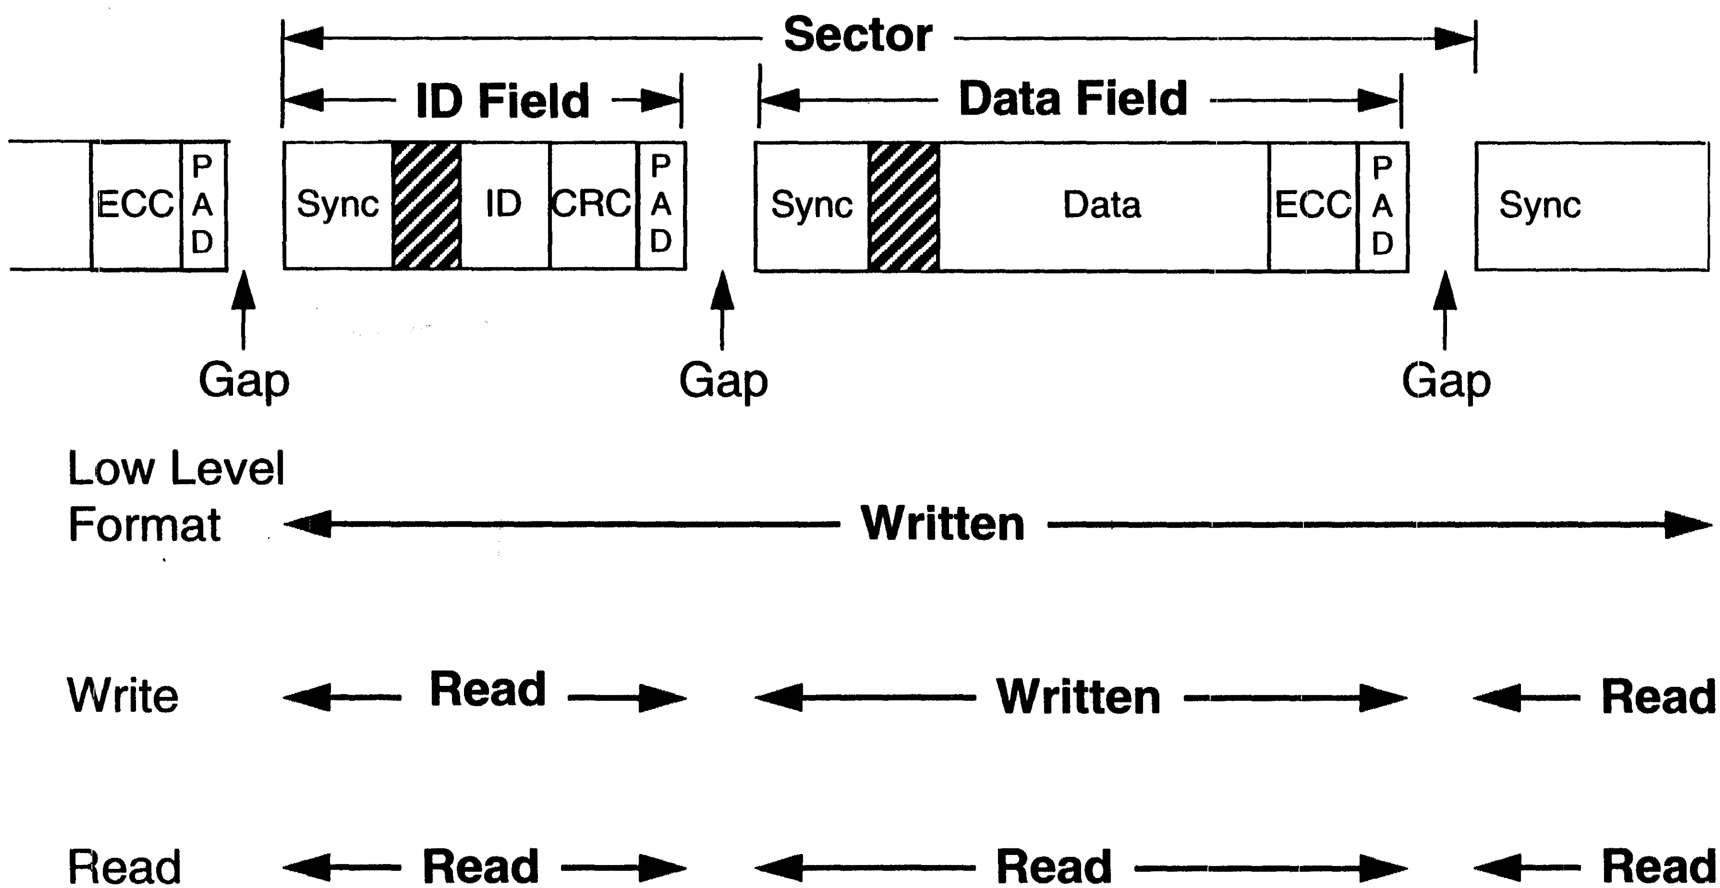

Read/Write Error Recovery

Read Error Types

- Recovered Error:
  - Block data recovered using retries or ECC correction
  - Reported with Sense Key 1
- Medium Error:
  - Block data unrecovered after retries and/or ECC correction
  - Unrecoverable error
  - Reported with Sense Key 3

Note: SCSI Provides several mode parameters to control how the drive handles R/W errors.

Read/Write Error Recovery (Continued)

Mode Page 1

Read/Write Error Recovery Parameters

Bit / Byte	7	6	5	4	3	2	1	0
0	PS	Rsvd	Page Code = 01h					
1	Page Length = 0Ah							
2	AWRE	ARRE	TB	RC	EER	PER	DTE	DCR
3	Read Retry Count							
4	Correction Span							
5	Head Offset Count							
6	Data Strobe Offset Count							
7	Reserved							
8	Write Retry Count							
9	Reserved							
10 - 11	Recovery Time Limit							

Note: See next page for description of bits.

Read/Write Error Recovery (Continued)

Mode Page 1

Bit	Name	Description
AWRE	Automatic Write Reallocation	1 = Reassign Defective Write Blocks
ARRE	Automatic Read Reallocation	1 = Reassign Recovered Read Blocks
TB	Transfer Block	1 = Transfer the bad block when a **Medium Error** is detected.
RC	Read Continuous	1 = Transfer all requested blocks, ignoring data errors
EER	Enable Early Recovery	1 = Enable correction before retries
PER	Post Error	1 = Report CHECK CONDITION status for **Recovered Error**
DTE	Disable Transfer on Error	1 = Stop after transferring a block with a **Recovered Error**
DCR	Disable Correction	1 = Do not use ECC for data error recovery

Read/Write Error Recovery (Continued)

Read Error Example

RC	DTE	TB	PER	Error	Blocks Transferred	Status	Sense Key
n/a	n/a	n/a	n/a	none		00h	0
0	0	n/a	0	Recovered Error		00h	0
0	0	n/a	1	Recovered Error		02h	1
0	1	n/a	1	Recovered Error		02h	1
0	n/a	1	n/a	Medium Error		02h	3
0	n/a	0	n/a	Medium Error		02h	3
1	n/a	n/a	n/a	Recovered or Medium		Undefined	

Mode Page 1 Read Error Handling Summary

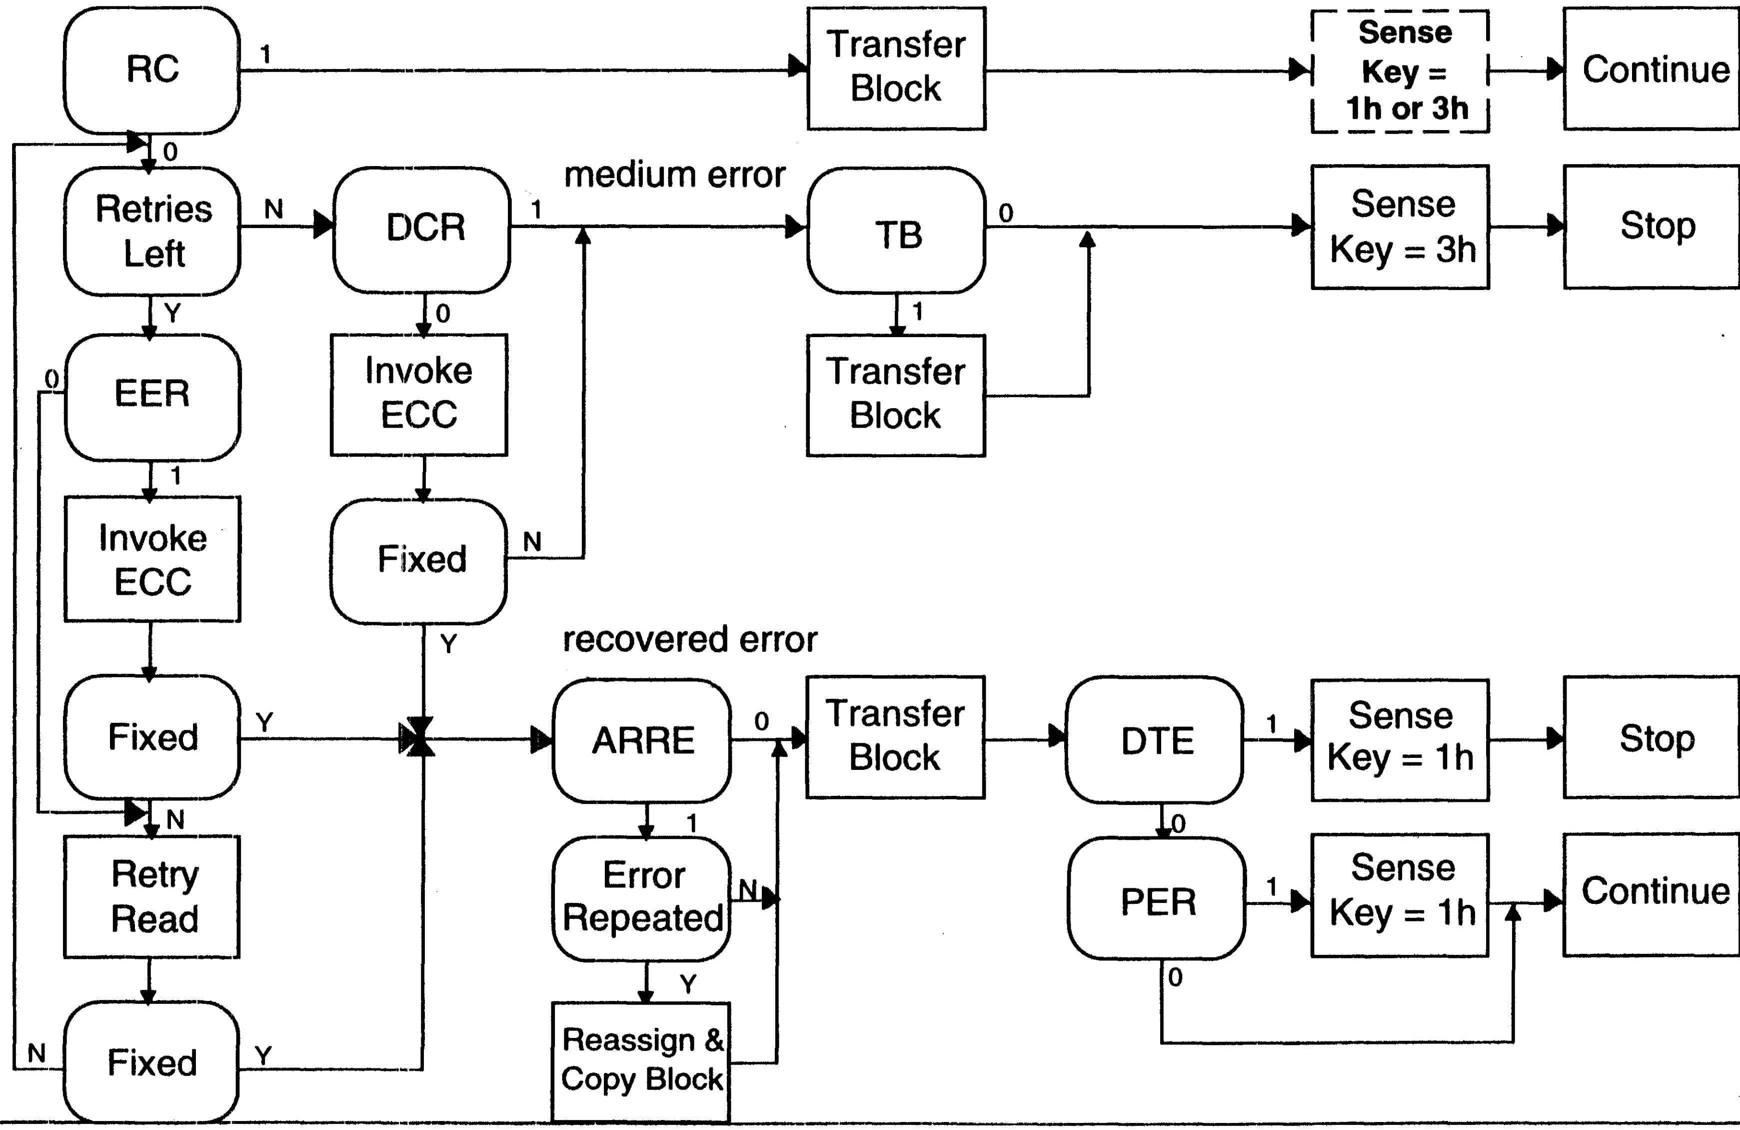

Disk Defect Management

Disk Defect Management

Modern SCSI disks are built with a small percentage of physical sectors reserved as **spares.**

When a repeatable error occurs at one location, the logical block associated with that location is remapped to one of the spare locations. The failing physical location is marked **defective**.

Defect List	When Written	How Written/Changed
P-List (Primary List)	At Factory	- Vendor Unique - User Cannot Change
G-List (Grown List)	In Field	- Automatic Reassignment - Format Unit Command - Reassign Blocks Command

Disk Defect Management (Continued)

- **Terminology** Reallocate, reassign, remap, map out, replace, spare out

Good Track

0	1	2	3	4	. . .	99	S	S

Sector Sparing

0	1	B	3	4	. . .	99	2	S

Sector Slipping

0	1	B	2	3	. . .	98	99	S

Track Sparing

0	1	2	3	4	. . .	99	S	S

(Different Physical Track)

Disk Defect Management (Continued)

- Reassigning a block does not reduce drive capacity
- Commands Affecting Defect Management:
  - Format Unit
  - Reassign Blocks - replace block(s) with alternate(s)
  - Read Defect Data:
- Defect Address Specification:
  - Logical Block Address
  - Physical - cylinder, head, and sector
  - Physical - cylinder, head, and offset from index (Byte number from beginning of Track)

Format Unit Command

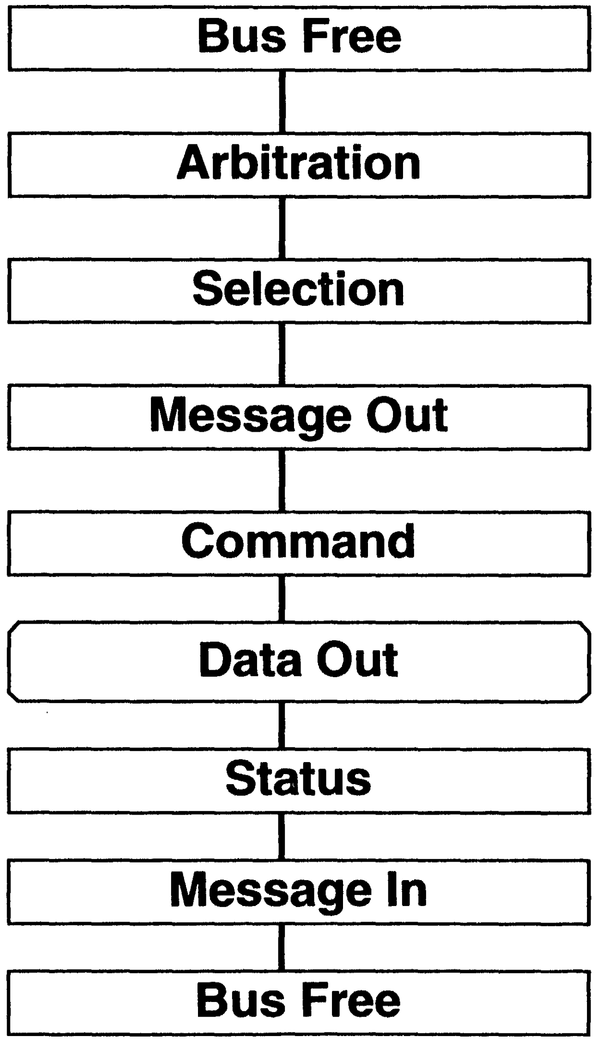

Format Unit Command (Continued)

- Before the Format Unit Command is executed, the drive has:
  - P-List (if any)
  - G-List (if any)
- Format Unit User Options:
  - P-List Existing
  - G-List Existing
  - D-List Supplied by User During Data Out
  - C-List Created During Media Certification
- Media Certification:
  - Vendor unique media scan normally using write and read
- During Format Unit:
  - If P Enabled Map Out P Defects
  - If G Enabled Map Out G Defects
  - If D Supplied Map Out D Defects
  - If C Enabled Map Out C Defects
- After the Format Unit command is executed:
  - Defect Lists contents:
    - P-List Unchanged
    - G-List Enabled G, D, and C
  - Mapped Out Defects:
    - All Enabled Lists
- Disabling the G-List causes the old G-List to be overwritten

Format Unit Command (Continued)

Format Unit Options **After the Command**

P-List existing primary	G-List existing grown	D-List supplied in data out	C-List target certification	P-List	G-List	Mapped Out Defects
No	No	No	No	unchanged	-	none
No	No	No	Yes	unchanged	C	C
No	No	Yes	No	unchanged	D	D
No	No	Yes	Yes	unchanged	D+C	D+C
No	Yes	No	No	unchanged	G	G
No	Yes	No	Yes	unchanged	G+C	G+C
No	Yes	Yes	No	unchanged	G+D	G+D
No	Yes	Yes	Yes	unchanged	G+D+C	G+D+C
Yes	**No**	**No**	**No**	**unchanged**	**-**	**P**
Yes	No	No	Yes	unchanged	C	P+C
Yes	No	Yes	No	unchanged	D	P+D
Yes	No	Yes	Yes	unchanged	D+C	P+D+C
Yes	**Yes**	**No**	**No**	**unchanged**	**G**	**P+G**
Yes	**Yes**	**No**	**Yes**	**unchanged**	**G+C**	**P+G+C**
Yes	Yes	Yes	No	unchanged	G+D	P+G+D
Yes	**Yes**	**Yes**	**Yes**	**unchanged**	**G+D+C**	**P+G+D+C**

Format Unit Command (Continued)

CDB Format

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code = 04h							
1	LUN			Fmt Data	Cmp Lst	List Format		
2	Vendor Unique							
3	Interleave (MSB)							
4	Interleave (LSB)							
5	Control Byte							

Note: See next page for description of fields

Format Unit Command (Continued)

- **Interleave** 0 = Default interleave
  (Interleave is a way of slowing down drives, and is no longer used.)
- **Fmt Data** DATA OUT supplied
  Should be set to 1
  When 0, drive may format any way it wants (i.e., vendor_unique)
- **Cmp Lst** Complete List (New complete G-list):
  - 1 = Disable old G-List
  - 0 = Enable old G-List
- **List Format** D-List format:
  - 0 = LBA
  - 5 = Cylinder, Head, Sector
  - 4 = Cylinder, Head, Offset from Index

Format Unit Command (Continued)

Data Out Format

Bit / Byte	7	6	5	4	3	2	1	0	
0	Reserved								Header
1	FOV	DPRY	DCRT	STPF	IP	DSP	IM	VU	
2	Defect List Length (MSB)								
3	Defect List Length (LSB)								
	Initialization Pattern (if any) (Normally not supplied, self-describing)								
	Defect Descriptor(s) (if any)								D-List

Note: See next page for description of fields

Format Unit Command (Continued)

- **FOV** 1 = Format Options Valid

  (Must be 1 for bits 2-6 to be valid, else 2-6 settings are vendor_unique)
- **DPRY** 1 = Disable Primary List
- **DCRT** 1 = Disable Certification
- **STPF** 1 = Stop on Defect List Error

  (When a defect list cannot be read)
- **IP** 1 = Initialization Pattern is supplied
- **DSP** 1 = Disable Saving Mode Select Parameters
- **IM** 1 = Immediate Mode
- **VU** Vendor Unique

Format Unit Command (Continued)

Hard-Versus Soft-Sectored

Hard-Sectored Drives

- Sector ID's written at factory only and not in field
- Sector size not changeable
- Logical block = 1 or more sectors
- Format Unit does not write Sector ID's
- Format Unit normally takes a short time
- Becoming more common

Soft-Sectored Drives

- Sector ID's written at factory and in field
- Sector size changeable
- Logical block = sector
- Format Unit may write Sector ID's
- Format Unit normally takes a long time
- Becoming less common

512 Byte sector - default

1024 Byte sector - soft sector

1024 Byte sector - hard sector

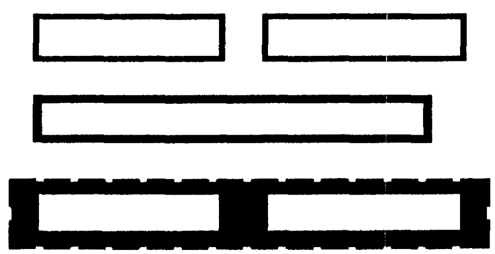

Format Unit Command (Continued)

Practical Considerations

- "Format" is used to describe several types of operations. Most common:
  - Low level - assigns sector locations and write ID fields.
  - Format Unit - used to assign LBAs to physical sectors, including reassignment of bad sectors.
  - File System Format - creates partitions and directories for organizing data on the drive.
- All user data on the disk is normally lost after SCSI FORMAT UNIT.
- Execute SCSI Format Unit if Block Length or Mode Page 3 changed.
- SCSI Format Unit in the field may cause performance improvement due to reasigned blocks (sector sparing normally changed to sector slipping).

This page is intentionally left blank.

CD-ROM

Physical Medium

Specifications

Diameter	120mm
Sides Recorded	1
Track Geometry	Single spiral track, beginning at the center of the disc.
Tracks-Per-Inch	16,000
Data Capacity	Approximately 550-630 MB ("63 minute" disc)
Physical Sector Size	2352 bytes
User Data Format	2048, 2328, or 2336 bytes/sector
Transfer Rate	150 KB/sec...1x speed, must be used to play digital audio 600 KB/sec...4x speed 1200 KB/sec...8x speed
RPM	Variable, to maintain constant linear velocity of data. Ranges from about 200 to 480 RPM for 1x speed.

Logical Sector Formats

Audio

98 sets of 6 stereo audio samples **1/75 of a second of playing time**

Mode 1 Data

<table>
<tr><td rowspan="3">Sync (12)</td><td colspan="4">Header (4)</td><td rowspan="3">User Data (2048)</td><td colspan="4">Auxiliary Data (288)</td></tr>
<tr><td colspan="3">Block Address (3)</td><td rowspan="2">Mode 1 (1)</td><td rowspan="2">EDC (4)</td><td rowspan="2">Space (8)</td><td colspan="2">ECC (276)</td></tr>
<tr><td>Min (1)</td><td>Sec (1)</td><td>Block (1)</td><td>P parity (172)</td><td>Q parity (104)</td></tr>
</table>

Mode 2 Data

<table>
<tr><td rowspan="3">Sync (12)</td><td colspan="4">Header (4)</td><td rowspan="3">All User Data (2 3 3 6)</td></tr>
<tr><td colspan="3">Block Address (3)</td><td rowspan="2">Mode 2 (1)</td></tr>
<tr><td>Min (1)</td><td>Sec (1)</td><td>Block (1)</td></tr>
</table>

CD-ROM Commands

- In addition to the commands defined for all SCSI devices, the following are defined in SCSI-3 specifically for CD-ROM devices.

Command Name	**Operation Code**	**SCSI Type**
Load/Unload CD	A6h	O
Mechanical Status	B8h	M
Pause/Resume	4Bh	O
Play Audio (10)	45h	O
Play Audio (12)	A5h	O
Play Audio MSF	47h	O
Read CD	BEh	O
Read CD MSF	B9h	M
Read CD Recorded Capacity	25h	M
Read Header	44h	M
Read Sub-channel	42h	M
Read TOC	43h	M
Scan	BAh	O
Seek	2Bh	M
Set CD Speed	B8h	O
Stop Play/Scan	4Eh	O

Key:
M = Command implementation is mandatory

O = Command implementation is optional

CD-Rom Command Notes

- Data on a CD-Rom can be addressed in 2 ways.
  - LBA - Normal SCSI logical block addressing
  - MSF - Minutes/Seconds/Frame. An MSF address defines an offset from the start of the disk or track in terms of audio playing time. 1 frame = $1/75$ second of playing time.
- PLAY AUDIO causes the drive to read audio data from the media and direct the audio stream to a speaker port. It does not transfer data across the SCSI interface.
- The standard SCSI READ command is only valid for 2048 byte data blocks. READ CD permits the host to retrieve data of any valid block size as well as header, subchannel, and ECC information.

Optical Memory Devices

- SCSI-2 and SCSI-3
- Similar to Direct Access Device Commands
- Typically for Eraseable (Reversible) Optical Drives
- Check Condition Status with Blank Check Sense Key if
  - Reading Blank (Unwritten) Block
  - Writing Previously Written Block
- Additional Commands:
  - Erase Command (10-byte and 12-byte)
  - Medium Scan Command for finding Written or Blank Blocks
  - 12-byte Commands to Read, Write,...
  - Update Block Command - see SCSI Standard
  - Read Updated Block Command - see SCSI Standard
  - Read Generation Command - see SCSI Standard
- 3.5" and 5.25" media is standard

Processor Devices

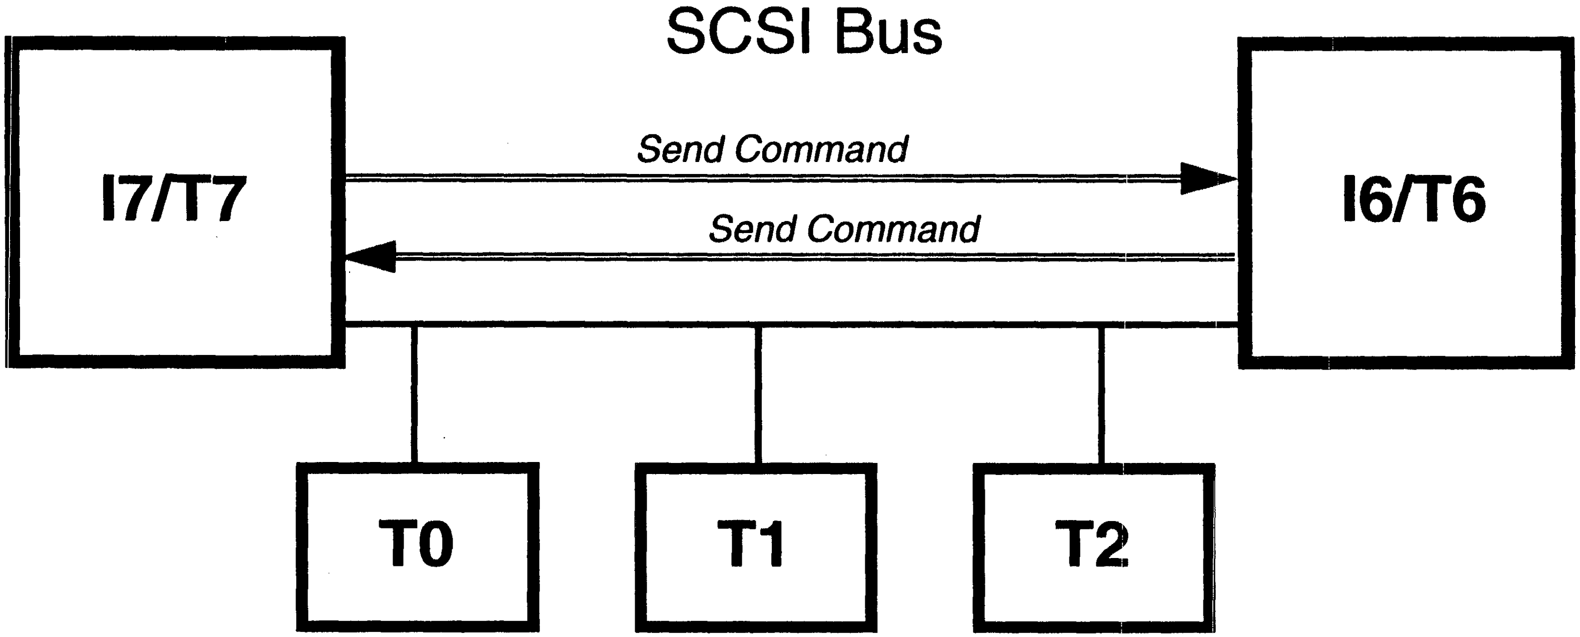

- Processors normally support Initiator Mode only.
- If Target Mode supported, should respond to INQUIRY command as device type 03h, Processor.
- Must support SEND command; may support RECEIVE.
- No standard for data length or format. Must be defined and interpreted by software in the processor.
- Used to implement processor clusters or to support AEN (described later).

Medium Changer Devices (Jukebox)

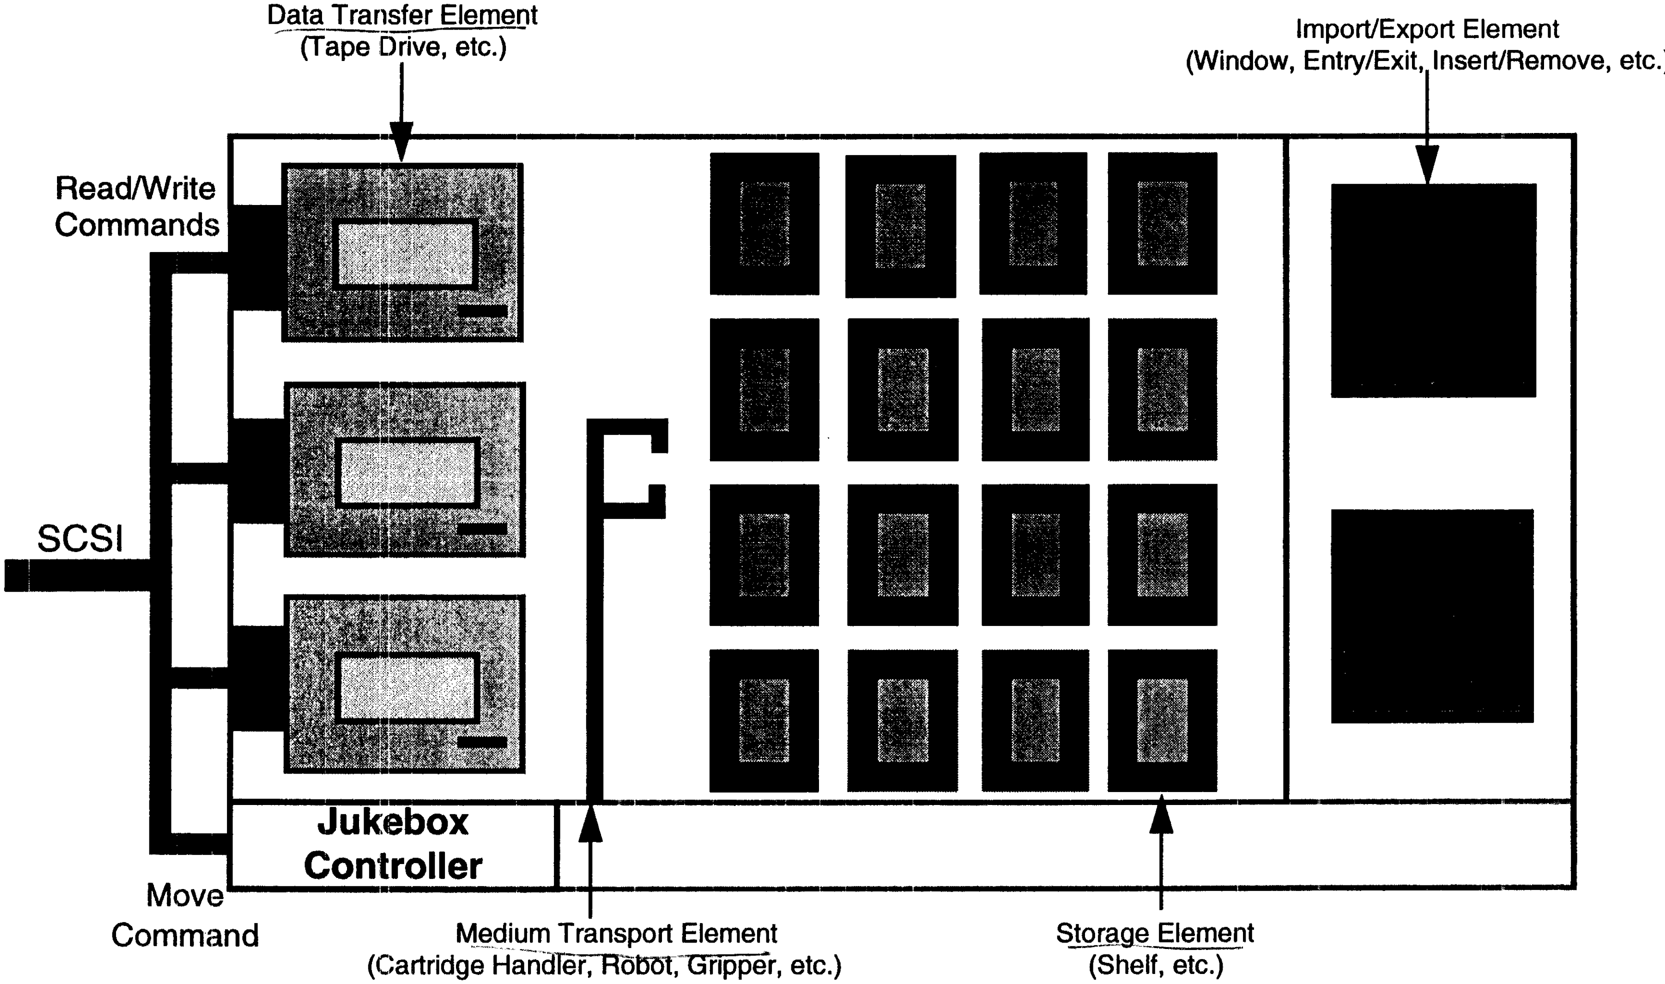

- Elements have unique element addresses
- Data transfer elements can be independent SCSI targets or logical units of the medium changer device

Typical SCSI-2 RAID Subsystem

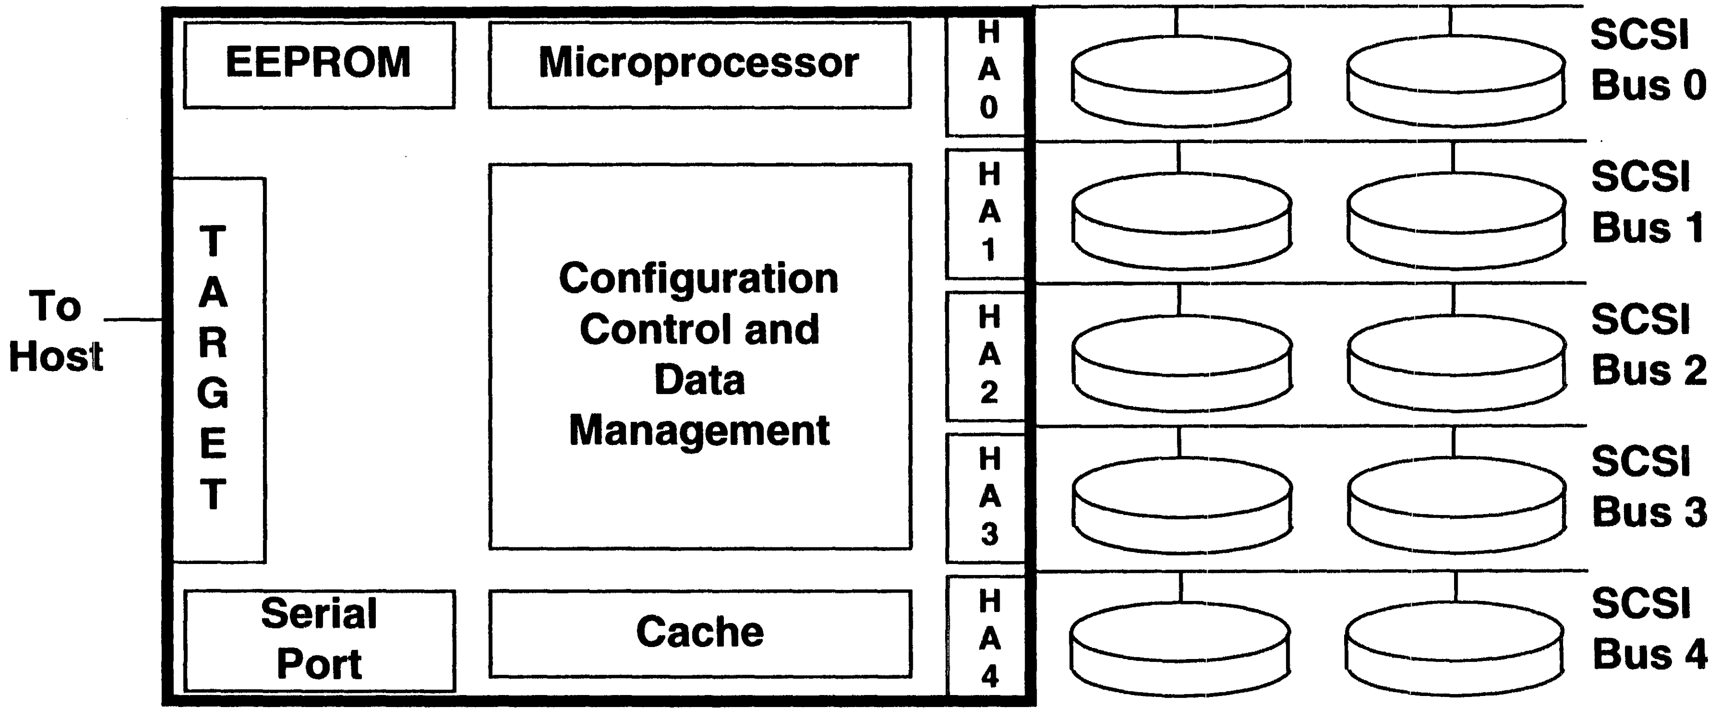

- Typically capable of RAID 0, 1, 3, and 5 configurations.
- Uses standard SCSI disk commands to access data.
- Uses vendor unique commands and mode pages to manage configuration.
- Controller may translate both LUNs and LBA's depending on the configuration used.

SCSI-3 Disk Array Example

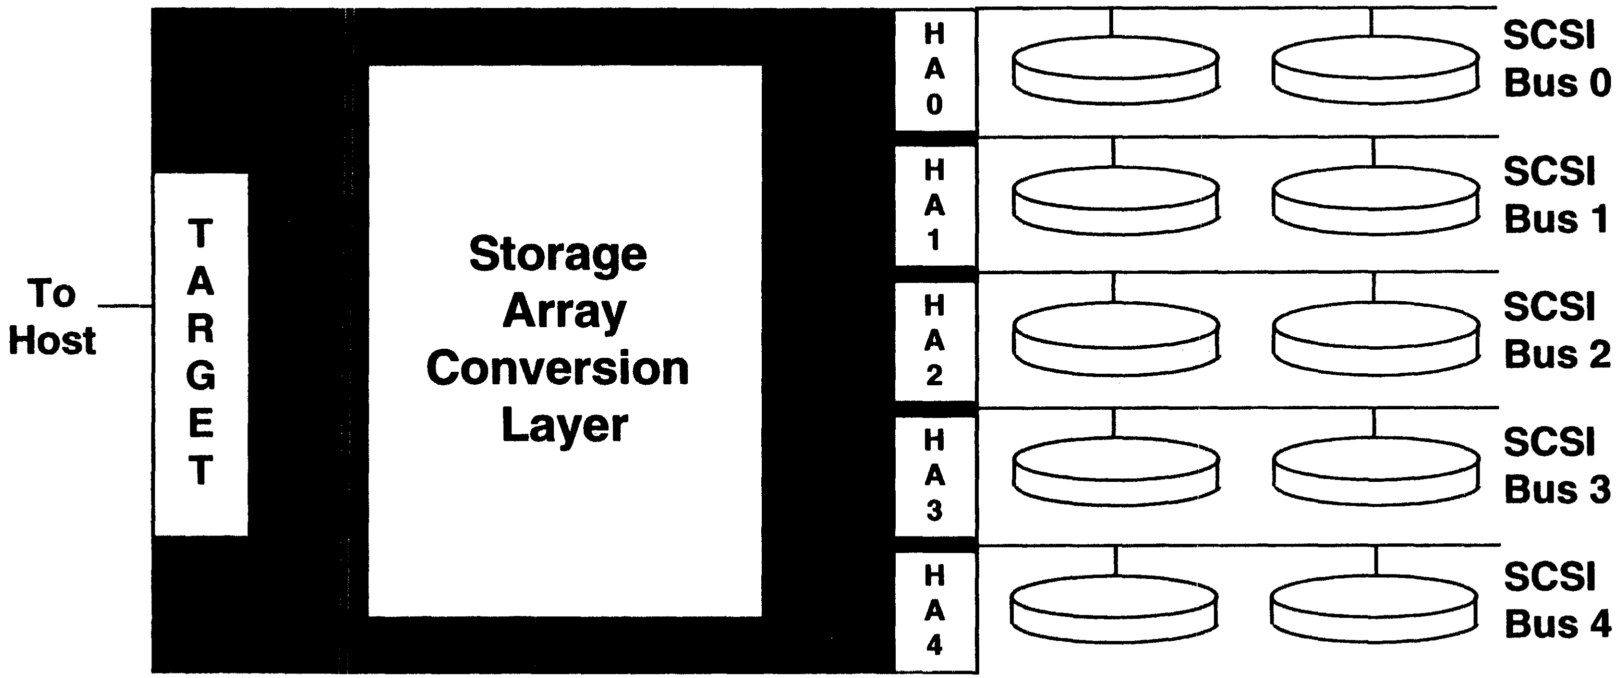

- New device type, controller, defined in SCSI-3.
- Standard command set defined to support disk array configuration and management.

Tape

SCSI Tape Device

Recording Modes

Fixed Block Mode

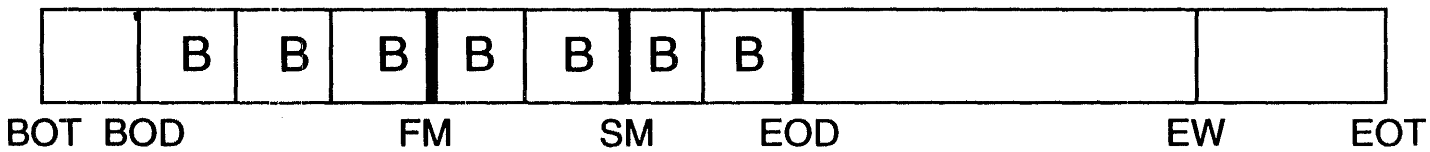

Variable Block Mode

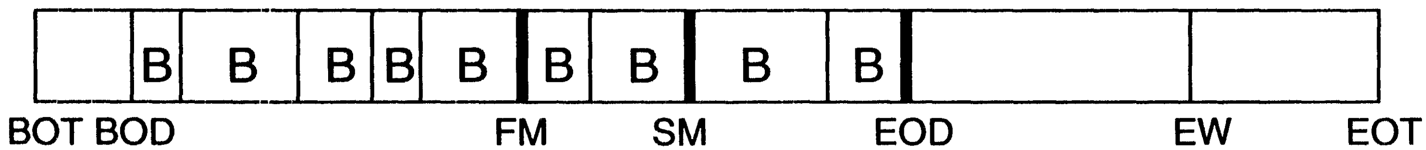

BOT	Beginning of Tape
EOT	End of Tape
BOD	Beginning of Data (also called Load Point)
EOD	End of Data
EW	Early Warning (normally for write)
FM	Filemark
SM	Setmark (typically shorter filemark)
B	Block (Fixed or Variable Length)

SCSI Tape Devices

Basic Recording Types

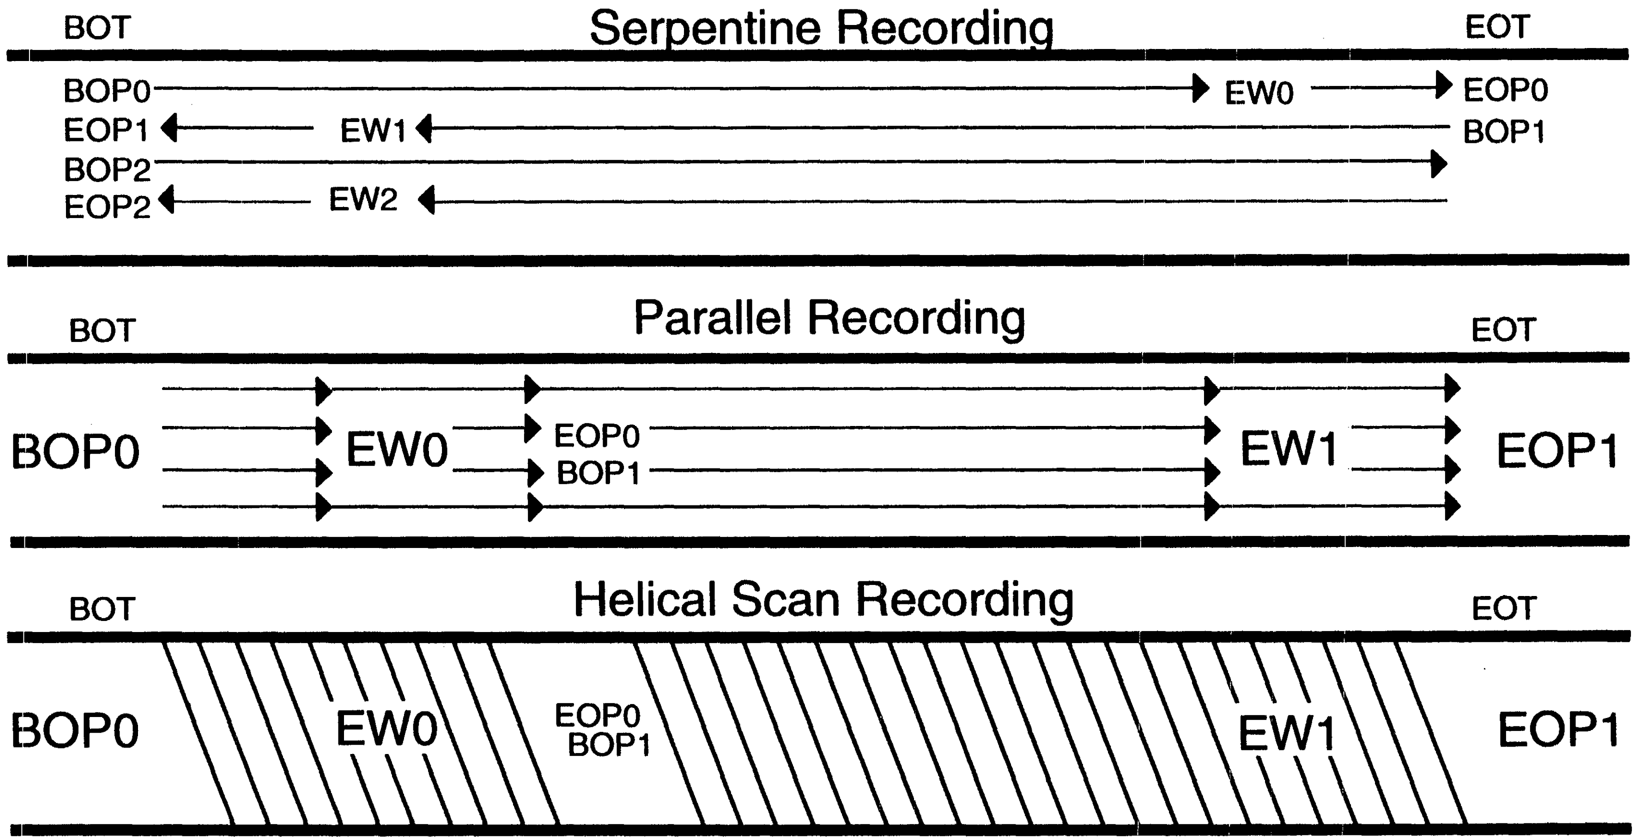

Tapes may be partitioned into independent recording areas.

BOP = beginning of partition, EOP = end of partitionm, EW = early warning

Use Mode select command or Locate command to switch partitions.

Tape - Read Command

Bus Phases

Bus Free

Arbitration

Selection

Message Out

Command

Data In

Status

Message In

Bus Free

Tape - Read Command (Continued)

CDB Format

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code = 08h							
1	LUN			Reserved			SILI	F
2	Transfer Length (MSB)							
3	Transfer Length							
4	Transfer Length (LSB)							
5	Control Byte							

Note: See next page for description of fields.

Tape - Read Command (Continued)

- F: 1 = Fixed, 0 = Variable
- Fixed Block Mode
  - Setup using Mode Select Block Length = N
  - F-Bit = 1
  - Transfer Length = Number of Blocks to Read
  - Can Read Multiple Blocks at a Time
- Variable Block Mode
  - Setup using Mode Select Block Length = 0
  - F-Bit = 0
  - Transfer Length = Number of Bytes to Read
  - Can Read One Block at a Time
- SILI: Suppress Incorrect Length Indicator
  - Variable Block Size Only

Tape - Read Command (Continued)

Error Reporting

Read CDB Type	Error Condition	Sense Key	Information Bytes (Sense data bytes 3-6)	Tape Position
V w/o SILI	Requested Length not equal actual Length	No Sense, ILI	Residue in bytes	After Block
Any	Filemark or Setmark encountered	No Sense, FM	Residue in blocks or requested length in bytes	After FM or SM
Any	EOD encountered	Blank Check	Residue in blocks or requested length in bytes	At EOD
Any	EOT Encountered	Medium Error, EOM	Residue in blocks or requested length in bytes	Undefined
F	Tape is Variable Mode	Illegal Request	None	Unchanged
F	SILI bit is set in CDB	Illegal Request	None	Unchanged

Tape - Read Command (Continued)

Fixed Block Error Handling Example

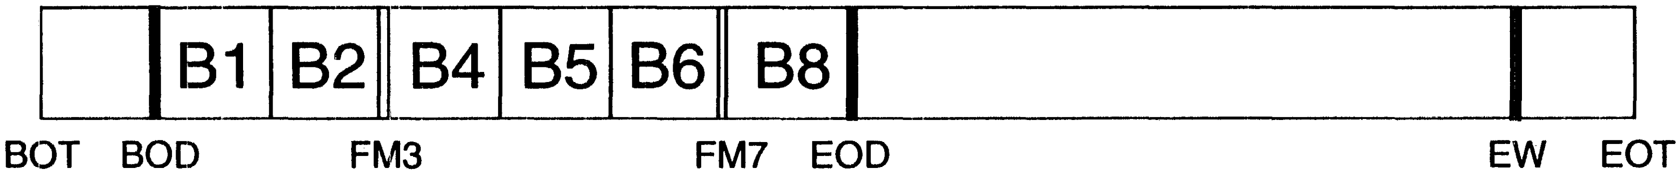

Command	Position	Status	Sense
Rewind	at BOD	00h	
Read 3 Blocks	after FM3	02h	FM
Read 3 Blocks	after B6	00h	
Read 3 Blocks	after FM7	02h	FM
Read 3 Blocks	at EOD	02h	Blank Check

Tape - Write Command

CDB Format

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code = 0Ah							
1	LUN			Reserved				F
2	Transfer Length (MSB)							
3	Transfer Length							
4	Transfer Length (LSB)							
5	Control Byte							

Tape - Write Command (Continued)

Bus Phases

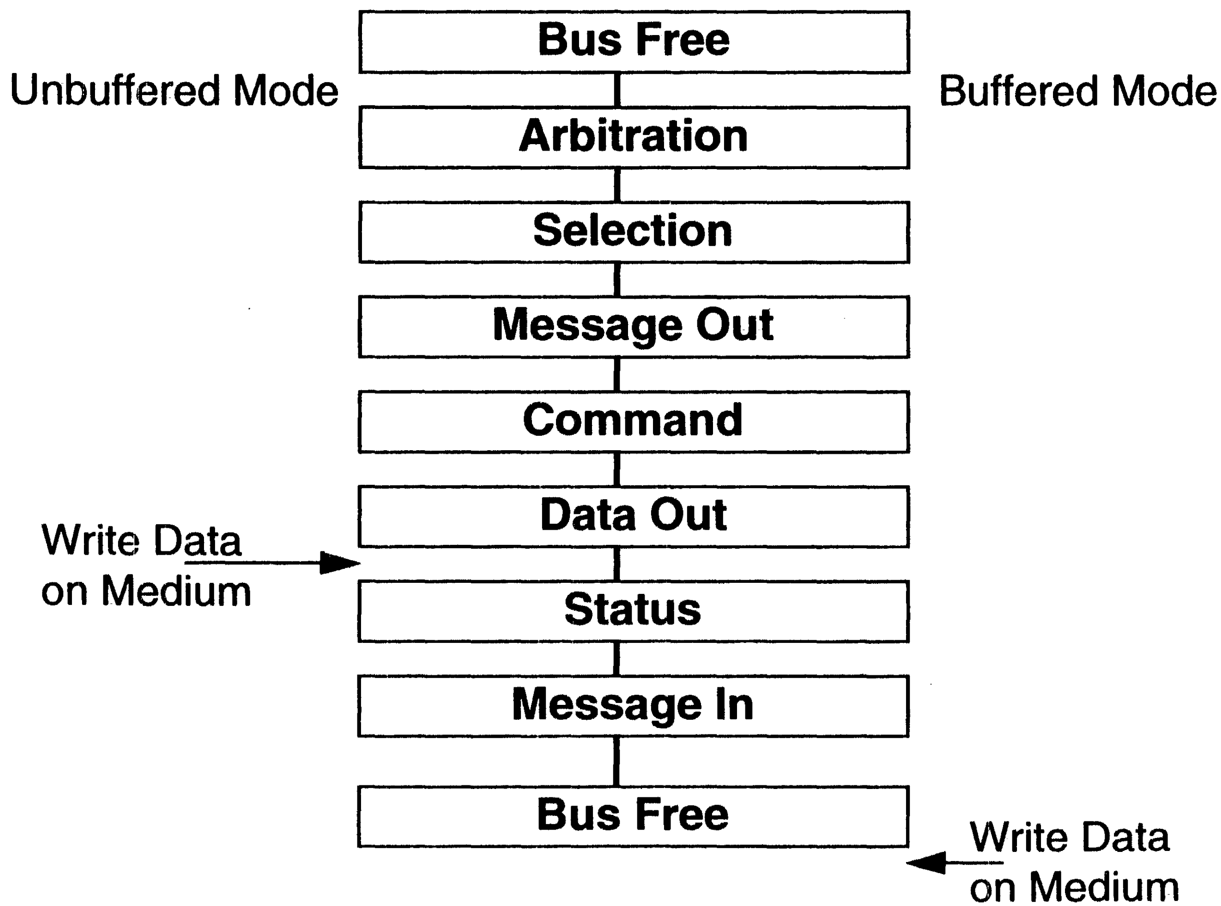

Tape - Write Command (Continued)

- Buffered/Unbuffered - Mode Selectable with Device-Specific Parameter in Mode Header
- Buffered mode error handling:
  - Errors detected after Bus Free are reported on subsequent command with Check Condition status
  - Residue in Information Bytes (Sense data bytes 3-6) = Unwritten Blocks (Bytes), Filemarks, and Setmarks
  - Deferred error code (71h) in the sense data
  - Use Recover Buffered Data Command to retrieve unwritten data
- Flushing buffered data to media:
  - Use Write Filemarks command with 0 filemarks, or any tape motion non-write-oriented command (e.g., Rewind)

Tape - Write Command (Continued)

Error Handling

Error Condition	Sense Key
Protected Medium	Data Protect
Early Warning	No Sense, EOM
End Of Tape	Volume Overflow, EOM
Write Fixed in Variable Mode	Illegal Request

Other Tape Commands

Write Filemarks	Mandatory. Specify Filemark Count in CDB. 0 Filemarks Flushes Buffered Data. Option to write Setmarks. Immediate Bit.
Space	Mandatory. Space Blocks, Filemarks (Setmarks), or End-of-Data. Space Reverse: Two's Complement Count.
Erase	Mandatory. Short/Long Erase.
Load/Unload	Optional. Load/Unload Bit. Immediate Bit. Retension Bit.
Prevent/Allow Medium Removal	Optional.
Locate	Optional. Position to LBA. Change Partition Option.
Rewind	Position tape at the beginning of the currently selected partition.
Read Reverse	Optional. Reversed Data Bytes Returned.

This page is intentionally left blank.

Physical and Electrical Characteristics

SCSI Connectors

- Different connectors are defined for shielded and unshielded cables
- The SCSI-2 standard defines two 50 pin connectors for each type of cable:
  - Low density - Large, pins widely spaced.
    Also called informally a "SCSI-1" cable.
  - High density - Small, pins closely spaced.
    Also called informally a "SCSI-2" cable.
- The SCSI-3 standard defines a 68 pin high density connector for shielded cables and a 68 pin high density connector for unshielded cables.
- Several other connectors and pinouts have been used for the SCSI bus. The SFF committiee maintaines a large set of connector specifications available from SFF FaxAccess at (408) 741-1600.

SCSI Cables

- A SCSI bus appears electrically to be a single set of conductors from one end to another. Physically it may consist of several individual cables connected together.
- A shielded cable has a conductive material, such as metal foil, surrounding the signal conductors and connected to ground through a metal connector shroud. This protects the signals from radiated electro-magnetic energy (EMI). Shielded SCSI cables are normally round bundles of twisted pair wires.
- An unshielded cable has no protection from EMI. Unshielded SCSI cables are normally flat ribbon cables.

SCSI Connectors

50 Pin A Cable (Narrow Bus)

High-Density Shielded

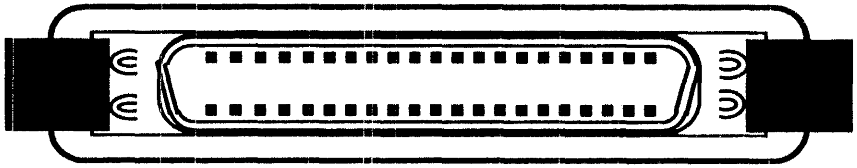

High-Density Unshielded

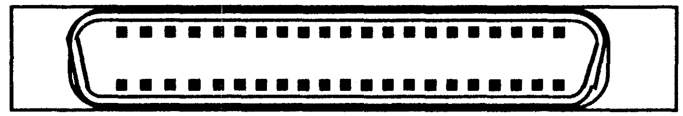

Low-Density Unshielded

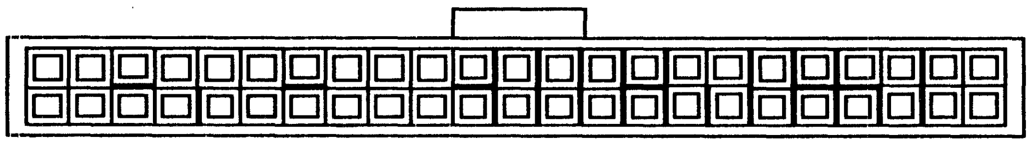

Low-Density Shielded

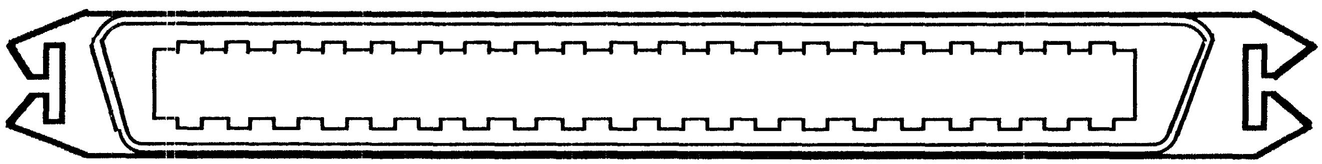

(Connectors shown approximately 2X actual size)

SCSI Connectors (Continued)

68 Pin P Cable (Wide Bus)

Shielded

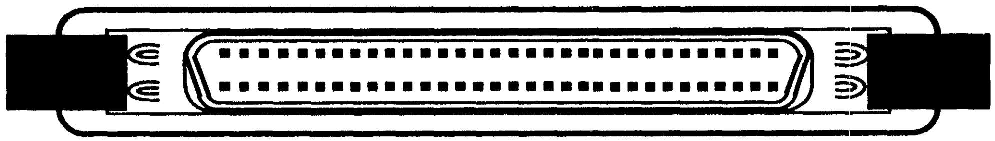

Unshielded

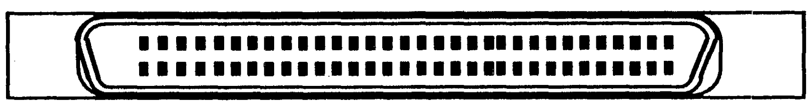

SCSI Electrical Interface Types

- Two electrical interfaces are defined.
- Single-Ended Drivers, Receivers, and Terminators
  - 6 Meters Max, for 5 MB/sec systems.. Tighter limits for Fast and Fast-20.
  - Most Common.
  - In-Cabinet Application.
- Differential Drivers, Receivers and Terminators.
  - 25 Meters Max, regardless of bus speed.
  - More Expensive.
  - Consumes More Power and Board Space.
- All devices on a bus must use the same type of electrical interface.

Termination Devices

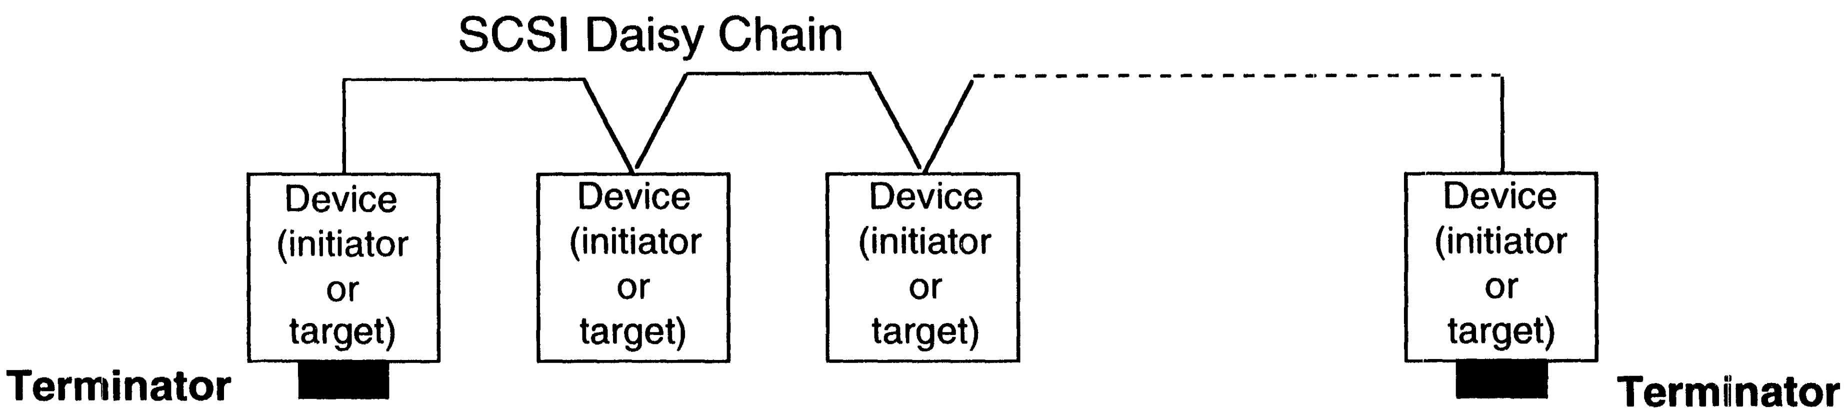

- At ends of cable
- Internal and External packages available
- Internal packages may be programmable (supply/don't supply Termination)
- SCSI standard describes two possible types of Termination, **Passive** (alternative 1) and **Active** (alternative 2)
- Other types also used, e.g., current mode and diode clamped
- Plug-N-Play requires auto termination.

Passive Termination Devices

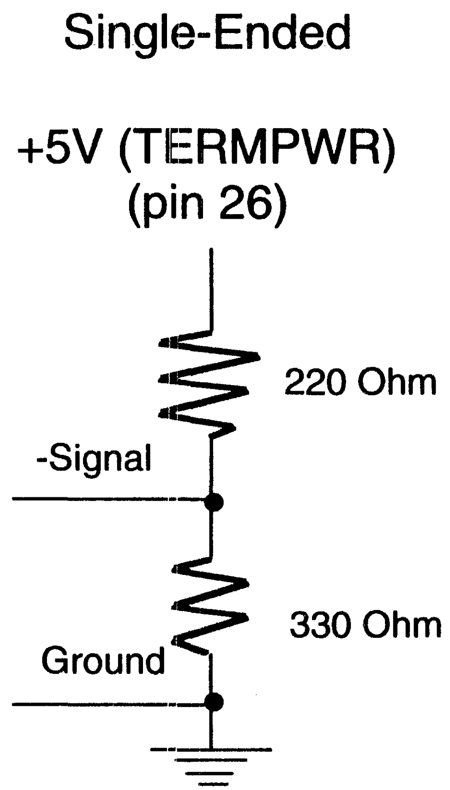

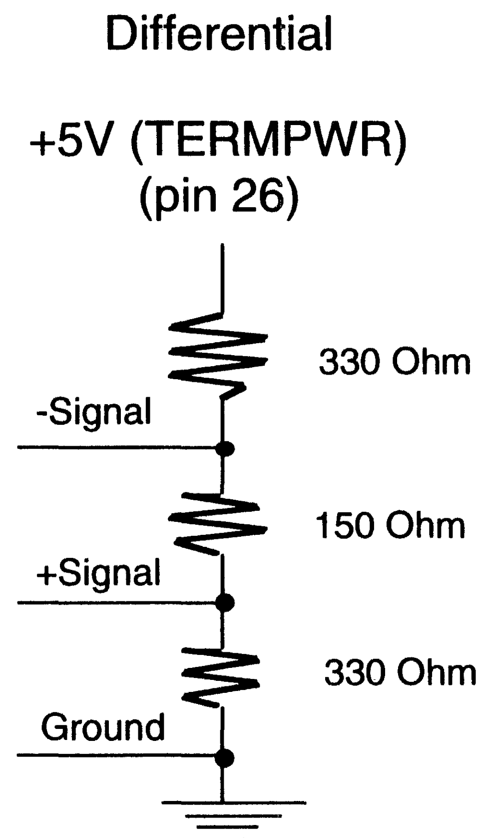

Active Termination Devices

SCSI-2 Single-Ended Alternative 2

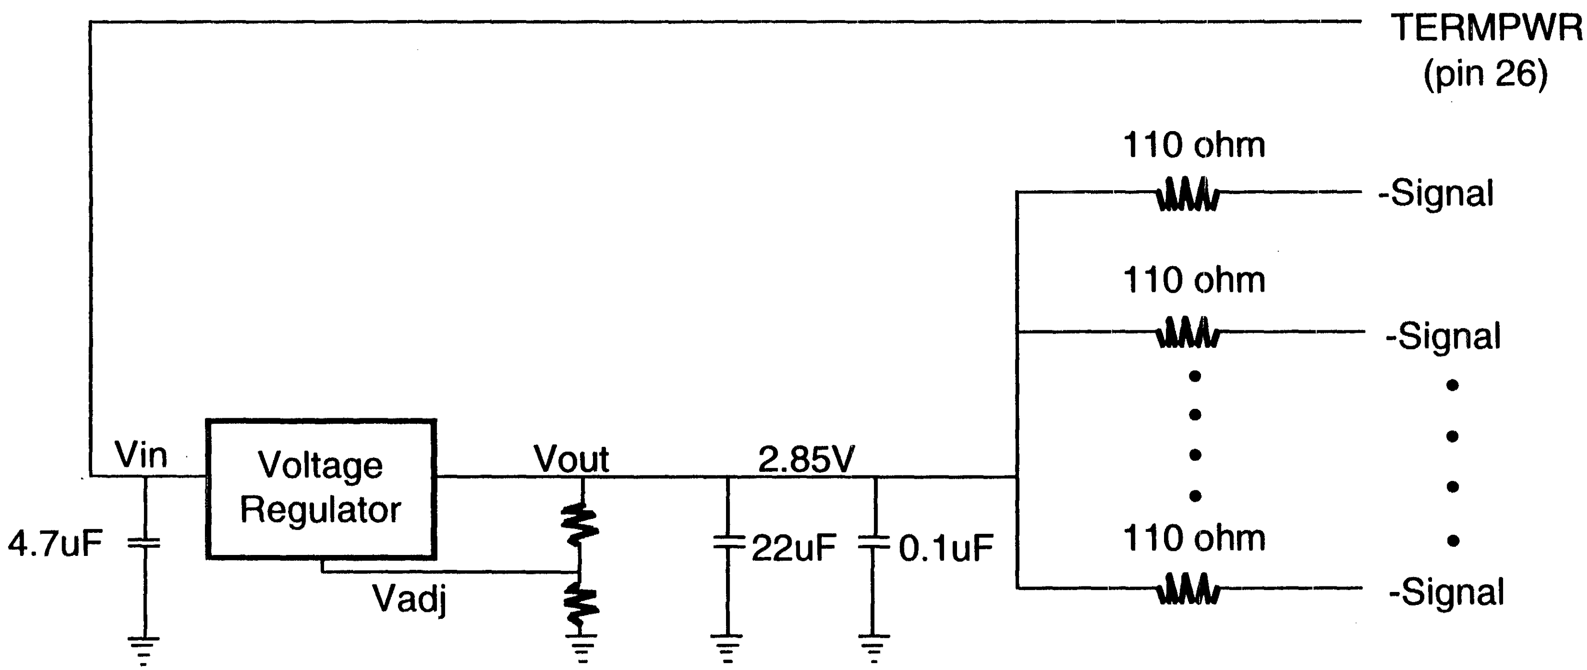

Terminator Power

- Initiators must supply terminator power to TERMPWR.
- Targets may supply terminator power.
- At least one device must supply terminator power.
- Terminator power must be supplied through a diode or similar device to prevent backflow of current.
- Fuse protects against accidental grounding.

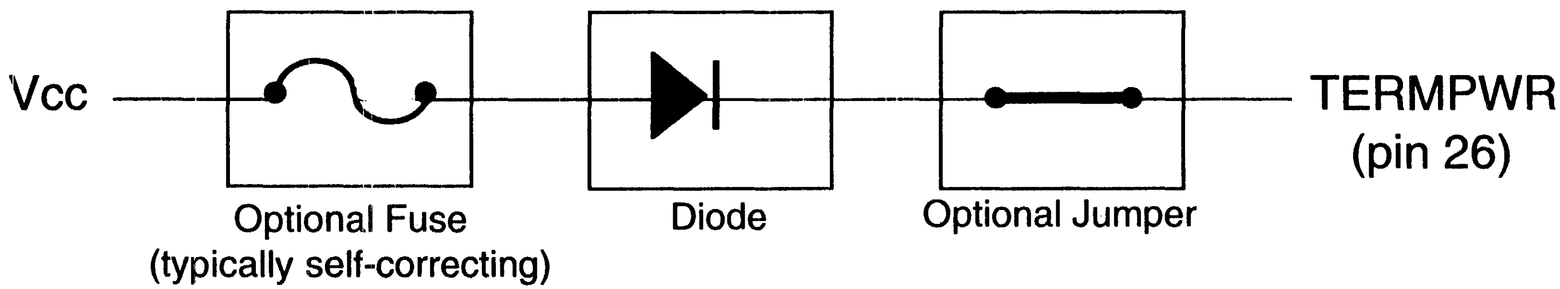

Cabling and Termination Example

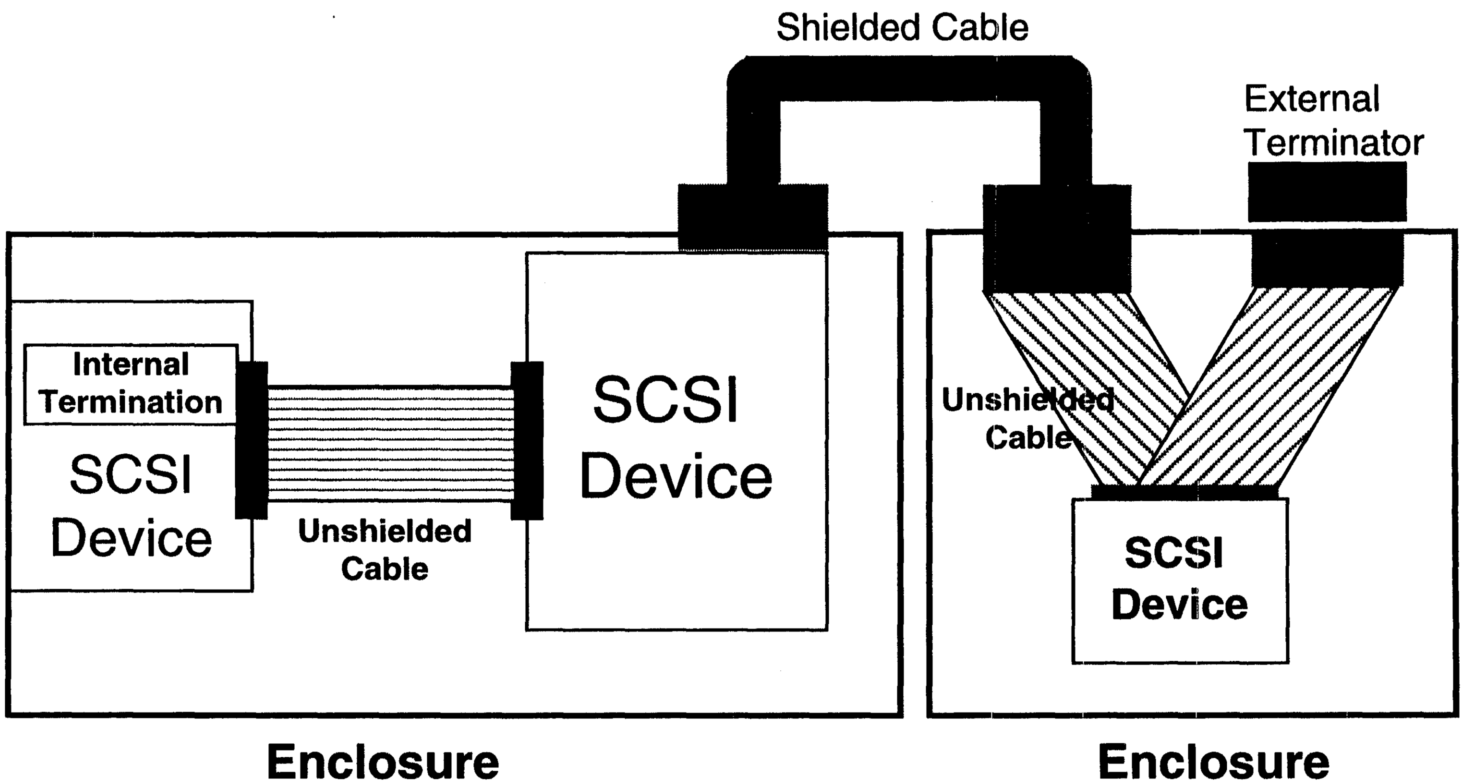

This page is intentionally left blank.

SCSI Signals

Signal Values

- Single-ended is active-low

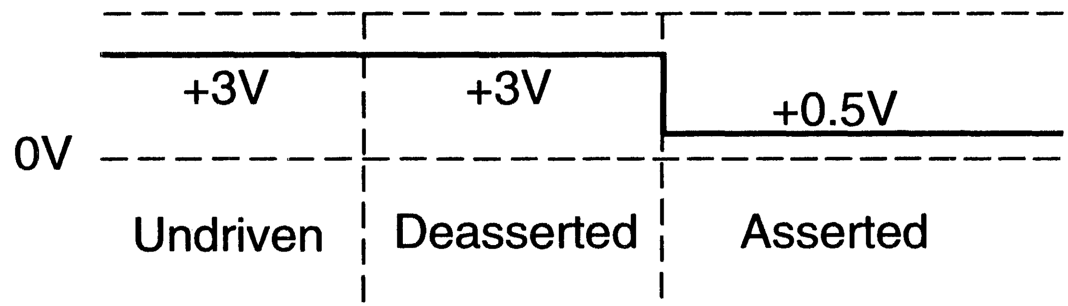

- Differential is active-high
  - i.e. + signal is more positive than - signal

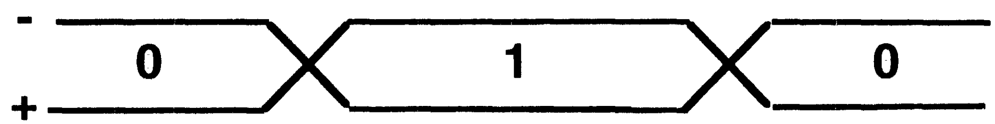

Signal Values (Continued)

Types of Signals

- OR-tied Signals:
  - Asserted (true):
    - One or more drivers are asserted.
  - Deasserted (false):
    - Terminator bias circuitry pulls the signal false.
- Non-OR-tied Signals:
  - Asserted (true):
    - One driver is asserted.
  - Deasserted (false):
    - Terminator bias circuitry pulls the signal false (open-collector).

      or
    - Signal is actively driven false, (active negation drivers).

Bus Signals

Introduction
(* = Or-Tied signal)

Signal	Name	Source	Usage
Control Signals (9 total):			
RST*	RESET	I & T	Reset Bus
BSY*	BUSY	I & T	Bus Busy or Idle
SEL*	SELECT	I & T	Select or Reselect
C/D*	CONTROL/DATA	T	Information Transfer Phase
I/O*	INPUT/OUTPUT	T	Information Transfer Phase
MSG*	MESSAGE	T	Information Transfer Phase
REQ	REQUEST	T	Transfer Byte
ACK	ACKNOWLEDGE	I	Transfer Byte
ATN	ATTENTION	I	Request Message Out Phase
Data Signals (9 or 18 total):			
DB(0-7)*	Data Bus	I & T	ID's or Bytes
DB(P)	Parity (Odd)	I & T	Transmission error detection
DB(8-15)	Upper Wide Data Bus	I & T	ID's or Bytes
DB(P1)	Upper Wide Parity (Odd)	I & T	Transmission error detection

"A" Connector Pin Assignment

Single-Ended
(- means active low)

Pin	Signal	Pin	Signal
1	GND	2	-DB(0)
3	GND	4	-DB(1)
5	GND	6	-DB(2)
7	GND	8	-DB(3)
9	GND	10	-DB(4)
11	GND	12	-DB(5)
13	GND	14	-DB(6)
15	GND	16	-DB(7)
17	GND	18	-DB(P)
19	GND	20	GND
21	GND	22	GND
23	Reserved (SCSI-1=GND)	24	Reserved (SCSI-1=GND)

"A" Connector Pin Assignment
(Continued)
Single-Ended

Pin	Signal	Pin	Signal
25	open	26	TERMPWR
27	Reserved (SCSI-1=GND)	28	Reserved (SCSI-1=GND)
29	GND	30	GND
31	GND	32	-ATN
33	GND	34	GND
35	GND	36	-BSY
37	GND	38	-ACK
39	GND	40	-RST
41	GND	42	-MSG
43	GND	44	-SEL
45	GND	46	-C/D
47	GND	48	-REQ
49	GND	50	-I/O

"A" Connector Pin Assignment
(Continued)

Differential

(TRUE = + SIGNAL more positive than - SIGNAL)

Pin	Signal	Pin	Signal
1	GND	2	GND
3	+DB(0)	4	-DB(0)
5	+DB(1)	6	-DB(1)
7	+DB(2)	8	-DB(2)
9	+DB(3)	10	-DB(3)
11	+DB(4)	12	-DB(4)
13	+DB(5)	14	-DB(5)
15	+DB(6)	16	-DB(6)
17	+DB(7)	18	-DB(7)
19	+DB(P)	20	-DB(P)
21	DIFFSENS	22	GND
23	Reserved (SCSI-1=GND)	24	Reserved (SCSI-1=GND)

"A" Connector Pin Assignment
(Continued)
Differential

Pin	Signal	Pin	Signal
25	TERMPWR	26	TERMPWR
27	Reserved (SCSI-1=GND)	28	Reserved (SCSI-1=GND)
29	+ATN	30	-ATN
31	GND	32	GND
33	+BSY	34	-BSY
35	+ACK	36	-ACK
37	+RST	38	-RST
39	+MSG	40	-MSG
41	+SEL	42	-SEL
43	+C/D	44	-C/D
45	+REQ	46	-REQ
47	+I/O	48	-I/O
49	GND	50	GND

"P" Connector Pin Assignment

Single-Ended

Signal Name	Cable Conductor Number		Signal Name
GROUND	1	2	-DB(12)
GROUND	3	4	-DB(13)
GROUND	5	6	-DB(14)
GROUND	7	8	-DB(15)
GROUND	9	10	-DB(P1)
GROUND	11	12	-DB(0)
GROUND	13	14	-DB(1)
GROUND	15	16	-DB(2)
GROUND	17	18	-DB(3)
GROUND	19	20	-DB(4)
GROUND	21	22	-DB(5)
GROUND	23	24	-DB(6)
GROUND	25	26	-DB(7)
GROUND	27	28	-DB(P)
GROUND	29	30	GROUND
GROUND	31	32	GROUND
TERMPWR	33	34	TERMPWR

"P" Connector Pin Assignment
(Continued)
Single-Ended

Signal Name	Cable Conductor Number		Signal Name
TERMPWR	35	36	TERMPWR
RESERVED	37	38	RESERVED
GROUND	39	40	GROUND
GROUND	41	42	-ATN
GROUND	43	44	GROUND
GROUND	45	46	-BSY
GROUND	47	48	-ACK
GROUND	49	50	-RST
GROUND	51	52	-MSG
GROUND	53	54	-SEL
GROUND	55	56	-C/D
GROUND	57	58	-REQ
GROUND	59	60	-I/O
GROUND	61	62	-DB(8)
GROUND	63	64	-DB(9)
GROUND	65	66	-DB(10)
GROUND	67	68	-DB(11)

SCSI-3 Interface Timing

Timing Description	Timing Values			
	Fast-20	Fast	Slow	Asynch
Arbitration Delay	2.4 us	2.4 us	2.4 us	2.4 us
Bus Clear Delay	800 ns	800 ns	800 ns	800 ns
Bus Free Delay	800 ns	800 ns	800 ns	800 ns
Bus Set Delay	1.8 us	1.8 us	1.8 us	1.8 us
Bus Settle Delay	400 ns	400 ns	400 ns	400 ns
Cable Skew Delay	3 ns	(5) 4 ns	(10) 4 ns	(10) 4 ns
Data Release Delay	400 ns	400 ns	400 ns	400 ns
Receive Assertion Period	11 ns	(30) 22 ns	(90) 70 ns	n/a
Receive Hold Time	11.5 ns	(10) 25 ns	(45) 25 ns	n/a
Receive Negation Period	11 ns	(30) 22 ns	(90)70 ns	n/a
Receive Setup Time	6.5 ns	15 ns	15 ns	n/a
Reset Hold Time	25 us	25 us	25 us	25 us
Selection Abort Time	200 us	200 us	200 us	200 us
Time-out Delay (recommended)	250 ms	250 ms	250 ms	250 ms
Deskew delay	15 ns	45 ns	45 ns	45 ns
Transmit Assertion Period	15 ns	30 ns	(90) 80 ns	n/a
Transmit Hold Time	16.5 ns	33 ns	53 ns	n/a
Transmit Negation Period	15 ns	30 ns	(90) 80 ns	n/a
Transmit Setup Time	11.5 ns	23 ns	23 ns	n/a

Note: Values in parenthesis are the SCSI-2 values for the same parameter

Bus Timings

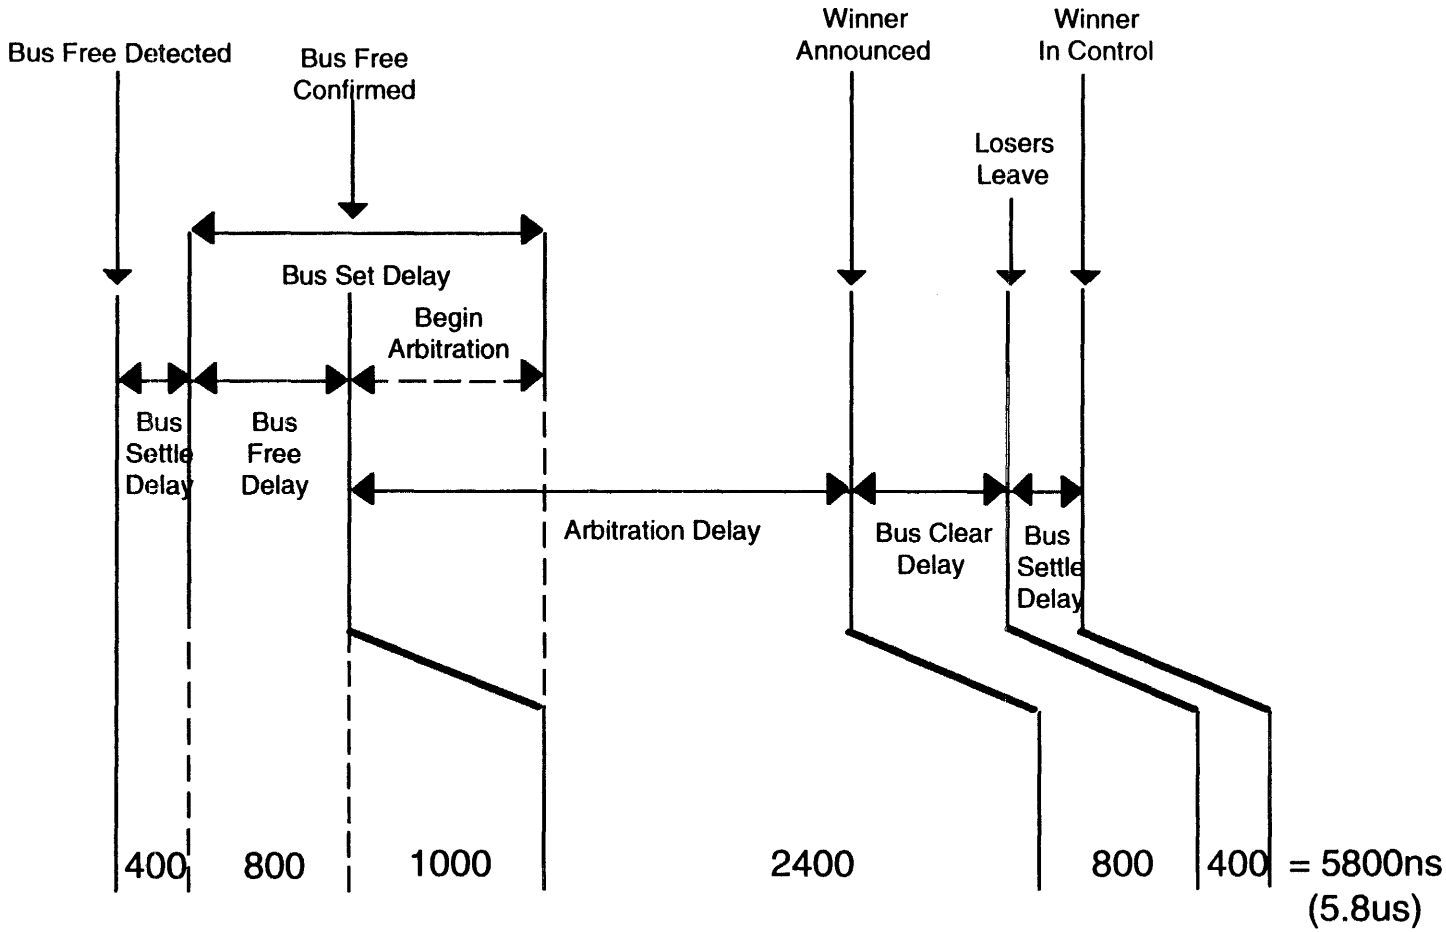

Bus Free Detected
Bus Free Confirmed
Winner Announced
Winner In Control
Losers Leave
Bus Set Delay
Begin Arbitration
Bus Settle Delay
Bus Free Delay
Arbitration Delay
Bus Clear Delay
Bus Settle Delay
400
800
1000
2400
800
400
= 5800ns
(5.8us)

SCSI Control Signals - Reset

- Any SCSI Device May Reset by Asserting the RST Line for Reset Hold Time (25us) min
- Hard Reset Alternative (widely used)
  - Release All SCSI Bus Signals
  - Clear Uncompleted Commands
  - Release Device Reservations
  - Restore Mode Select Parameters to Saved (or Default) Values
  - Post Unit Attention
- Soft Reset Alternative (not valid for SCSI-3)
  - Release All SCSI Bus Signals
  - Attempt to Complete Uncompleted Commands
  - Preserve Reservations and Parameters
- Determine which alternative is implemented using Inquiry command
- Either Hard or Soft may be used within a SCSI-~~3~~2 system, but not both

SCSI Control Signals - Attention

- Because the Target Determines the Information Transfer Phases, there must be a way for the initiator to pass control information to the target.
- How
  - Initiator Asserts the ATN signal during all phases except during Bus Free and Arbitration
  - Target Responds With Message Out Phase
  - Initiator Sends Message Out Byte(s) and Drops ATN on Last Byte

Selection Sequence Revisited

Handling Attention After Selection Phase

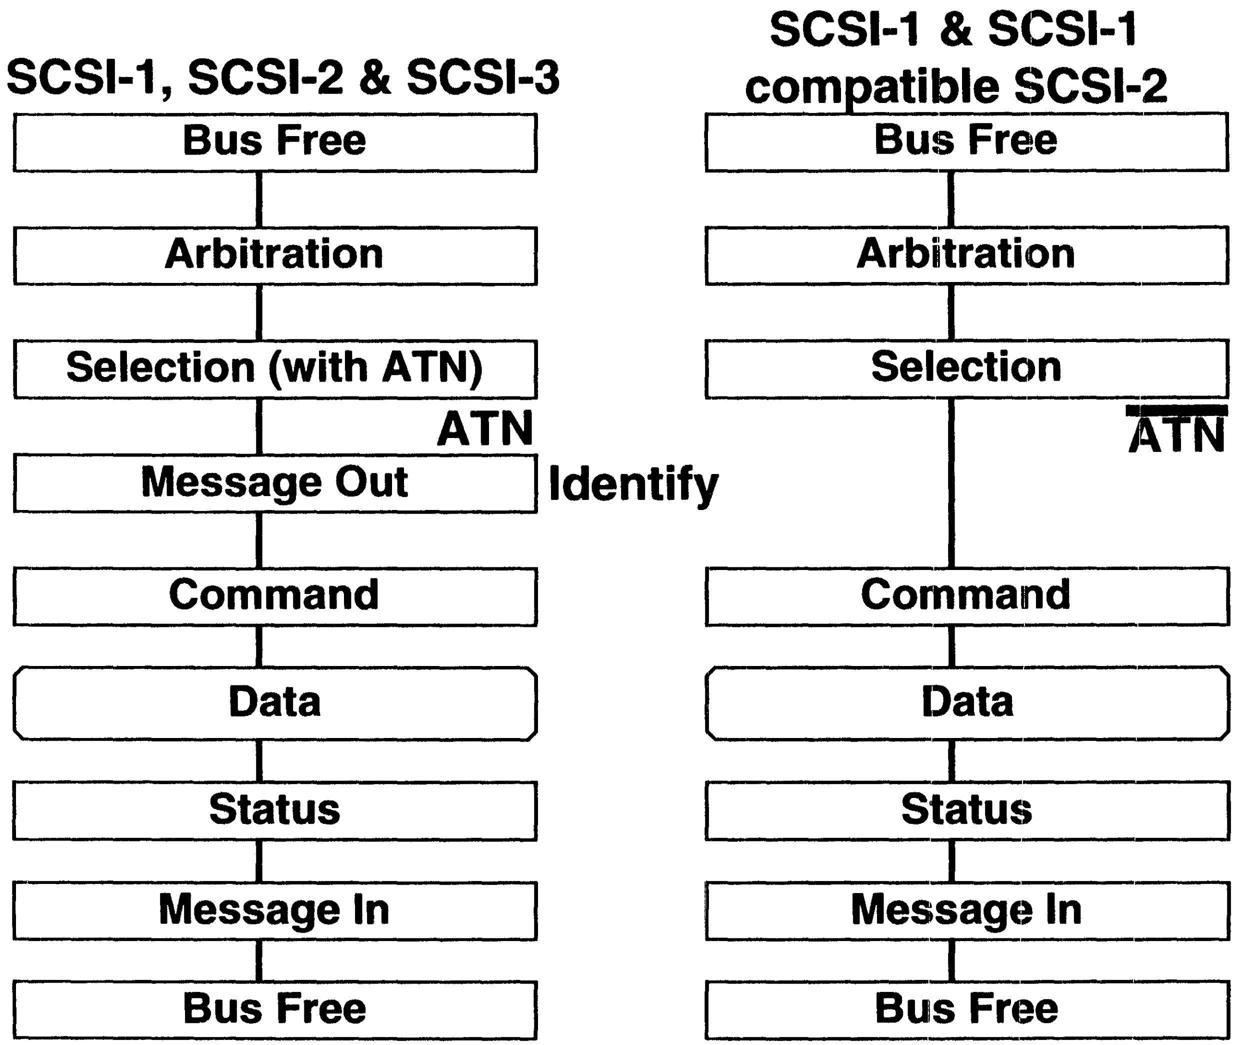

Information Transfer Phases

C/D	1	Control (Not Data)
	0	Data
I/O	1	In
	0	Out
Msg	1	Message
	0	Not Message

C/D	I/O	MSG	Phase
1	0	0	COMMAND
1	1	0	STATUS
0	0	0	DATA OUT
0	1	0	DATA IN
1	0	1	MESSAGE OUT
1	1	1	MESSAGE IN
0	0	1	Reserved
0	1	1	Reserved

Asynchronous Transfer

Asynchronous REQ/ACK Handshake

Out Direction

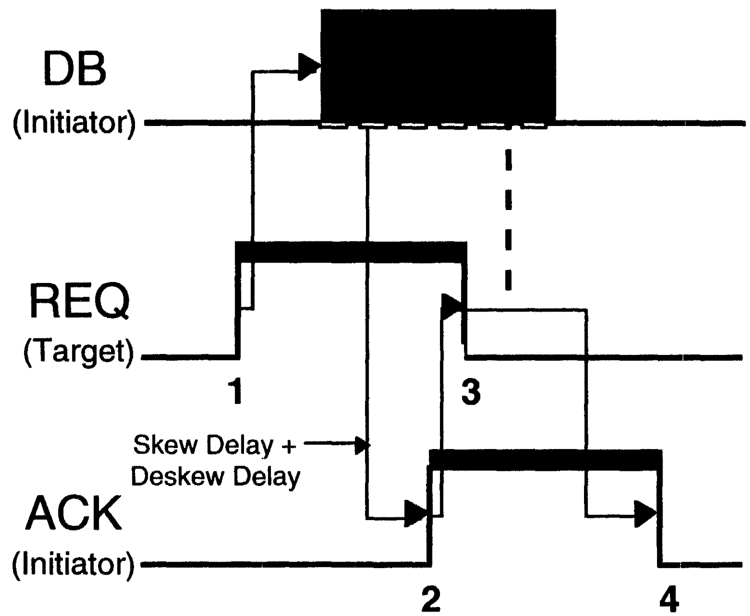

1 - Target says: Give Me

2 - Initiator says: Take It

3 - Target says: Got It

4 - Initiator says: Finished

Note: there are no timeouts in handshake transitions

Asynchronous REQ/ACK Handshake
(Continued)
In Direction

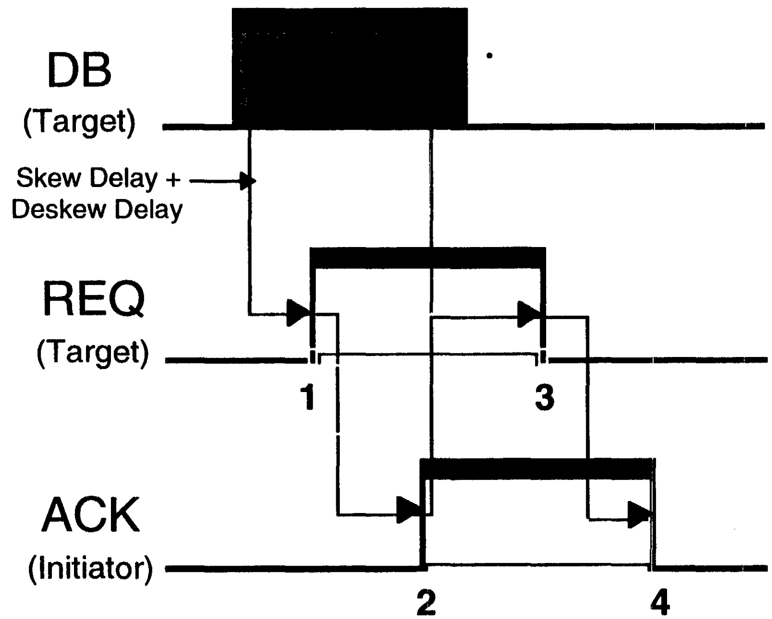

1 - Target says: Take It
2 - Initiator says: Got It
3 - Target says: Finished
4 - Initiator says: Finished

Note: there are no timeouts in handshake transitions

Bus Signals and Phase Sequences

I1 = asserted by initiator, T1 = asserted by target, T0 = deasserted by target
- = not driven (deasserted), IT1 = asserted by initiator and target

Phase	BSY	SEL	ATN	C/D	I/O	MSG	REQ	ACK	DB(7-0)	Comments
BUS FREE	-	-	-	-	-	-	-	-	-	
ARBITRATION	I1	-	-	-	-	-	-	-	I(iid)	Initiator Arbitrating
	I1	I1	-	-	-	-	-	-	I(iid)	(If Highest ID)
SELECTION	I1	I1	I1	-	-	-	-	-	I(iid,tid)	
	-	I1	I1	-	-	-	-	-	I(iid,tid)	
	T1	I1	I1	-	-	-	-	-	I(iid,tid)	
	T1	-	I1	-	-	-	-	-	-	

Bus Signals and Phase Sequences
(Continued)

Phase	BSY	SEL	ATN	C/D	I/O	MSG	REQ	ACK	DB(7-0)	Comments
MESSAGE OUT	T1	-	I1	T1	T0	T1	-	-	-	
	T1	-	I1	T1	T0	T1	T1	-	-	
	T1	-	I1	T1	T0	T1	T1	-	I(msg)	(Message Byte)
	T1	-	I1	T1	T0	T1	T1	I1	I(msg)	
	T1	-	I1	T1	T0	T1	-	I1	I(msg)	
	T1	-	I1	T1	T0	T1	-	I1	-	
	T1	-	-	T1	T0	T1	-	I1	-	Drop Attention before ACK
	T1	-	-	T1	T0	T1	-	-	-	
COMMAND	T1	-	-	T1	T0	T0	-	-	-	
	T1	-	-	T1	T0	T0	T1	-	-	
	T1	-	-	T1	T0	T0	T1	-	I(cmd)	(CDB Byte)
	T1	-	-	T1	T0	T0	T1	I1	I(cmd)	
	T1	-	-	T1	T0	T0	-	I1	I(cmd)	
	T1	-	-	T1	T0	T0	-	I1	-	
	T1	-	-	T1	T0	T0	-	-	-	

Bus Signals and Phase Sequences
(Continued)

Phase	BSY	SEL	ATN	C/D	I/O	MSG	REQ	ACK	DB(7-0)	Comments
DATA OUT	T1	-	-	T0	T0	T0	-	-	-	
	T1	-	-	T0	T0	T0	T1	-	-	
	T1	-	-	T0	T0	T0	T1	-	I(data)	(Data Byte)
	T1	-	-	T0	T0	T0	T1	I1	I(data)	
	T1	-	-	T0	T0	T0	-	I1	I(data)	
	T1	-	-	T0	T0	T0	-	I1	-	
	T1	-	-	T0	T0	T0	-	-	-	
DATA IN	T1	-	-	T0	T1	T0	-	-	-	
	T1	-	-	T0	T1	T0	-	-	T(data)	(Data Byte)
	T1	-	-	T0	T1	T0	T1	-	T(data)	
	T1	-	-	T0	T1	T0	T1	I1	T(data)	
	T1	-	-	T0	T1	T0	T1	I1	-	
	T1	-	-	T0	T1	T0	-	I1	-	
	T1	-	-	T0	T1	T0	-	-	-	

Bus Signals and Phase Sequences
(Continued)

Phase	BSY	SEL	ATN	C/D	I/O	MSG	REQ	ACK	DB(7-0)	Comments
STATUS	T1	-	-	T1	T1	T0	-	-	-	
	T1	-	-	T1	T1	T0	-	-	T(stat)	(Status Byte)
	T1	-	-	T1	T1	T0	T1	-	T(stat)	
	T1	-	-	T1	T1	T0	T1	I1	T(stat)	
	T1	-	-	T1	T1	T0	T1	I1	-	
	T1	-	-	T1	T1	T0	-	I1	-	
	T1	-	-	T1	T1	T0	-	-	-	
MESSAGE IN	T1	-	-	T1	T1	T1	-	-	-	
	T1	-	-	T1	T1	T1	-	-	T(msg)	(Message Byte)
	T1	-	-	T1	T1	T1	T1	-	T(msg)	
	T1	-	-	T1	T1	T1	T1	I1	T(msg)	
	T1	-	-	T1	T1	T1	T1	I1	-	
	T1	-	-	T1	T1	T1	-	I1	-	
	T1	-	-	T1	T1	T1	-	-	-	

Bus Signals and Phase Sequences
(Continued)

Phase	BSY	SEL	ATN	C/D	I/O	MSG	REQ	ACK	DB(7-0)	Comments
BUS FREE	-	-	-	-	-	-	-	-	-	
ARBITRATION	T1	-	-	-	-	-	-	-	T(tid)	Target Arbitrating
	T1	T1	-	-	-	-	-	-	T(tid)	(If Highest ID)
RESELECTION	T1	T1	-	-	T1	-	-	-	T(iid,tid)	(I/O Indicates Reselection)
	-	T1	-	-	T1	-	-	-	T(iid,tid)	
	I1	T1	-	-	T1	-	-	-	T(iid,tid)	
	IT1	T1	-	-	T1	-	-	-	T(iid,tid)	
	IT1	-	-	-	-	-	-	-	-	
	T1	-	-	-	-	-	-	-	-	

Bus Signals and Phase Sequences
(Continued)

Phase	BSY	SEL	ATN	C/D	I/O	MSG	REQ	ACK	DB(7-0)	Comments
SELECTION	I1	I1	I1	-	-	-	-	-	I(iid,tid)	
WITH	-	I1	I1	-	-	-	-	-	I(iid,tid)	Wait 250 ms min
TIMEOUT	-	I1	I1	-	-	-	-	-	-	Wait 200 us max
	-	-	-	-	-	-	-	-	-	End of Selection Timeout
RESELECTION	T1	T1	-	-	T1	-	-	-	T(iid,tid)	
WITH	-	T1	-	-	T1	-	-	-	T(iid,tid)	Wait 250 ms min
TIMEOUT	-	T1	-	-	T1	-	-	-	-	Wait 200 us max
	-	-	-	-	-	-	-	-	-	End of Reselection Timeout

Note: Selection Timeout Procedure avoids connecting to the wrong device.

Signal Transitions

Arbitration/Selection

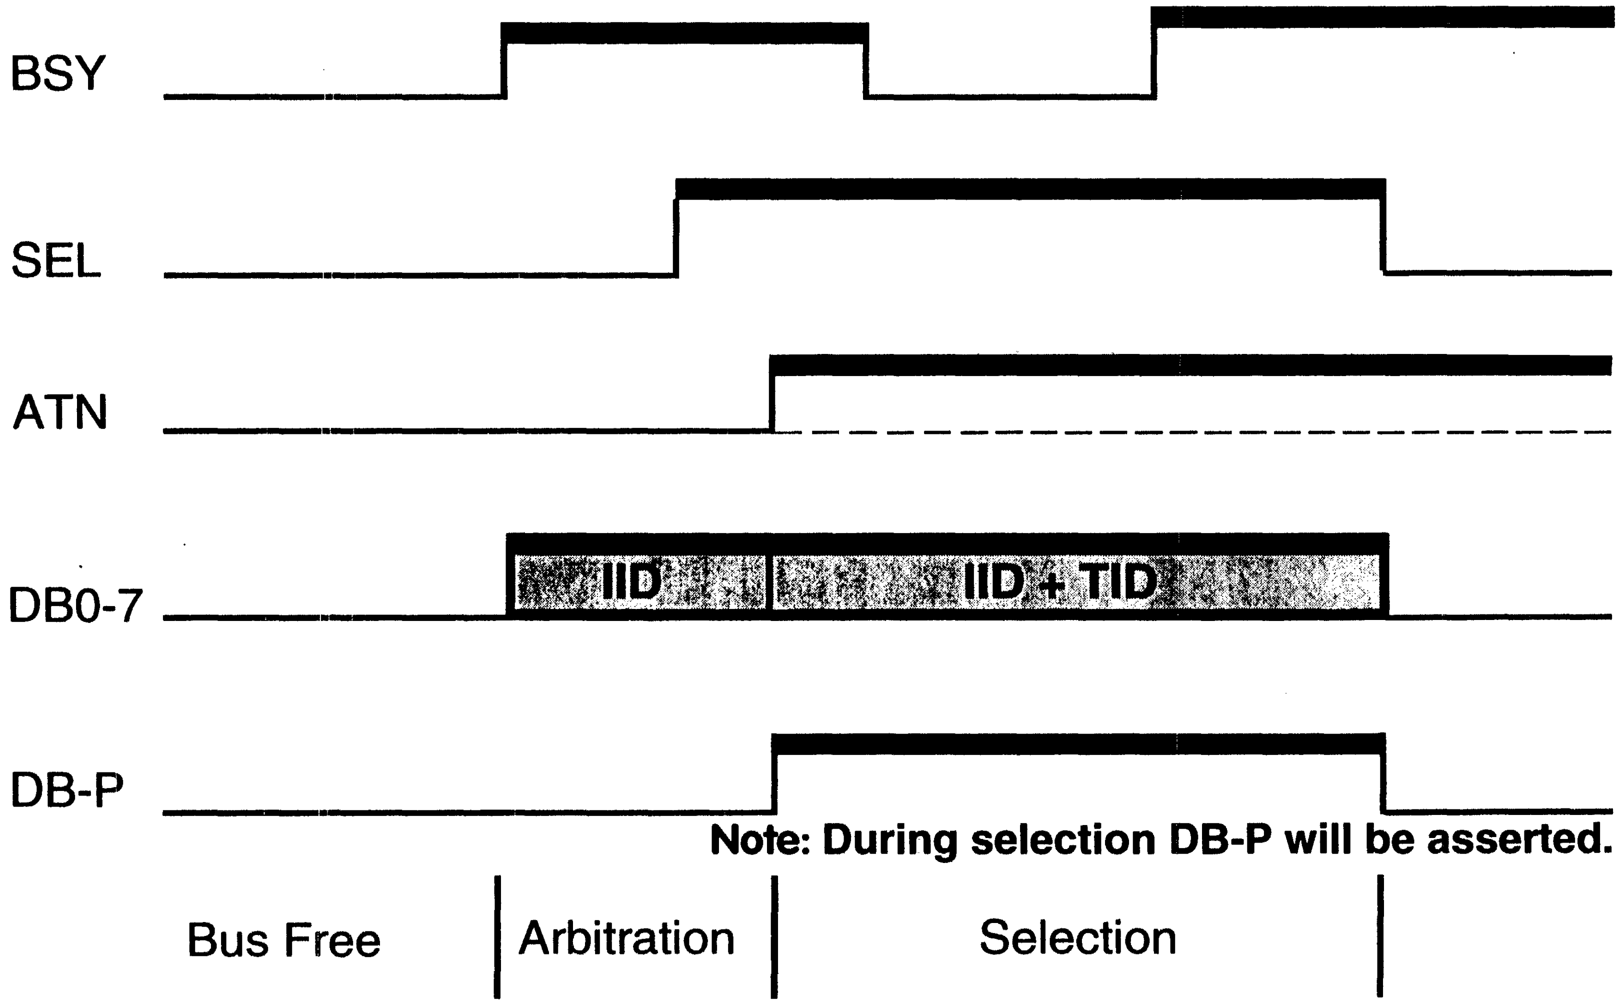

Signal Transitions (Continued)

Test Unit Ready Command

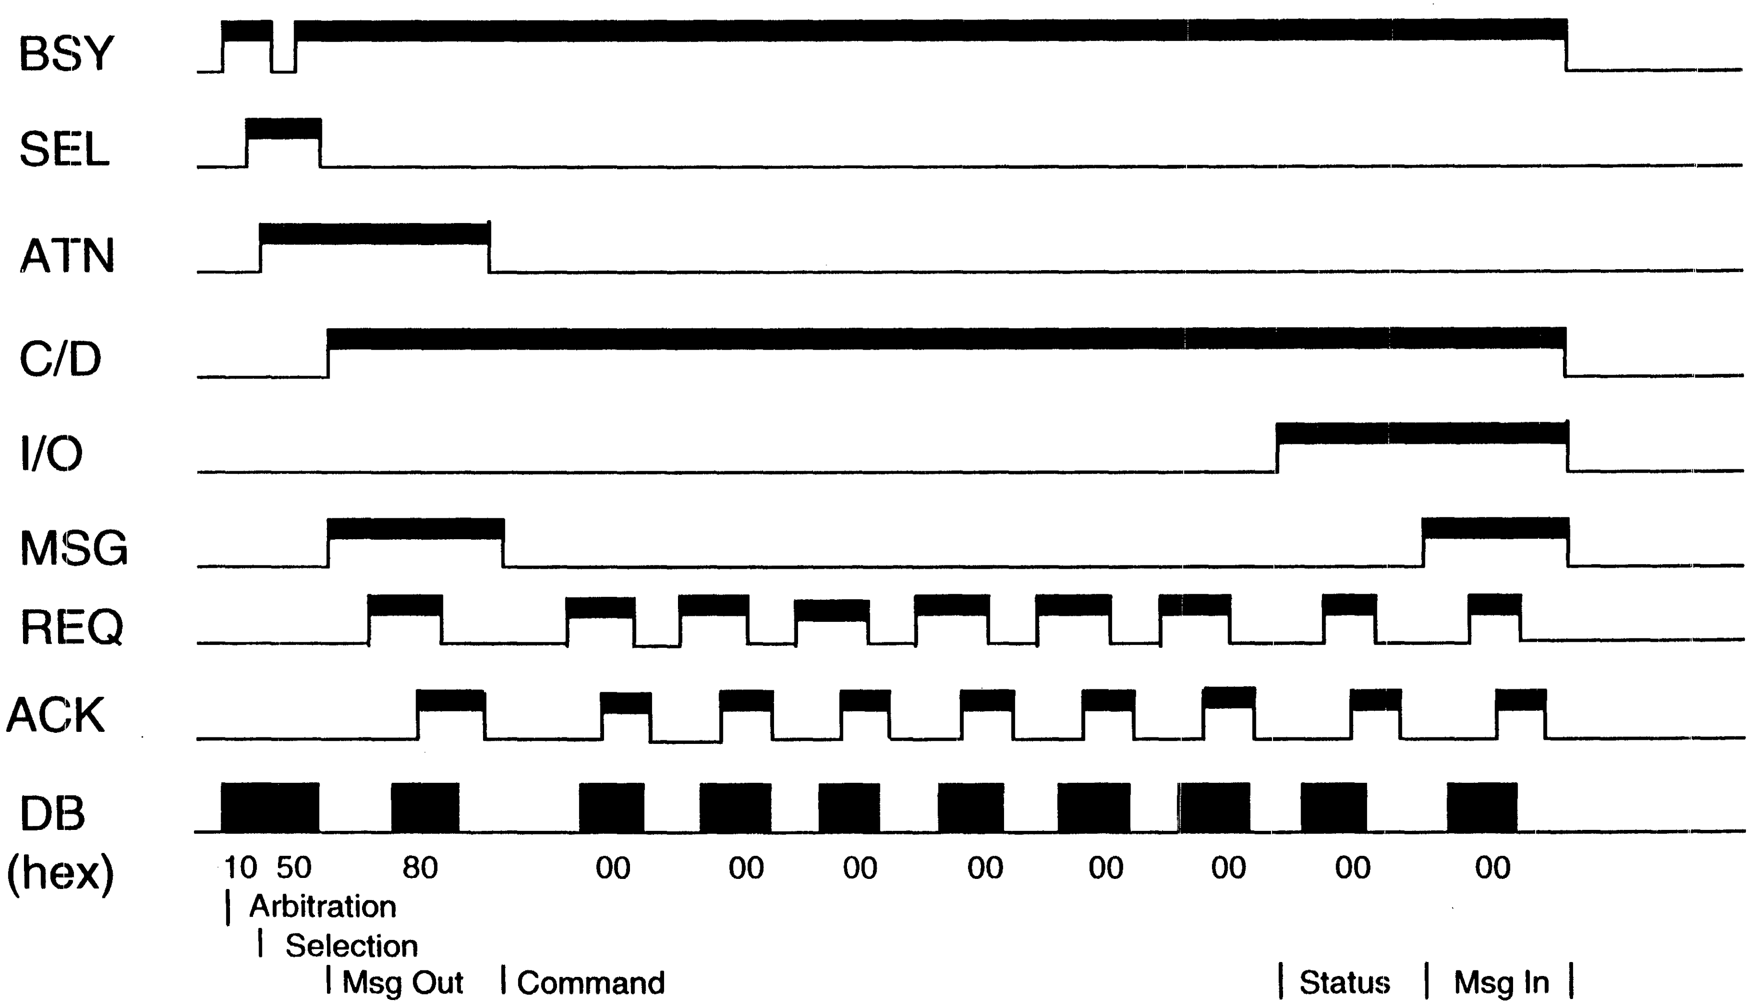

This page is intentionally left blank.

Synchronous Transfer

Synchronous SCSI

- Originally, synchronous SCSI was limted to 5 MB/sec. This is known today as slow SCSI.
- SCSI-2 added Fast SCSI, permitting up to 10 MB/sec on a narrow bus.
- SCSI-3 added Fast-20 SCSI, permitting up to 20 MB/sec on a narrow bus.
- The bus protocol for all three speed ranges is identical. Only the signal timings are different.

Synchronous Transfer Example

Data In

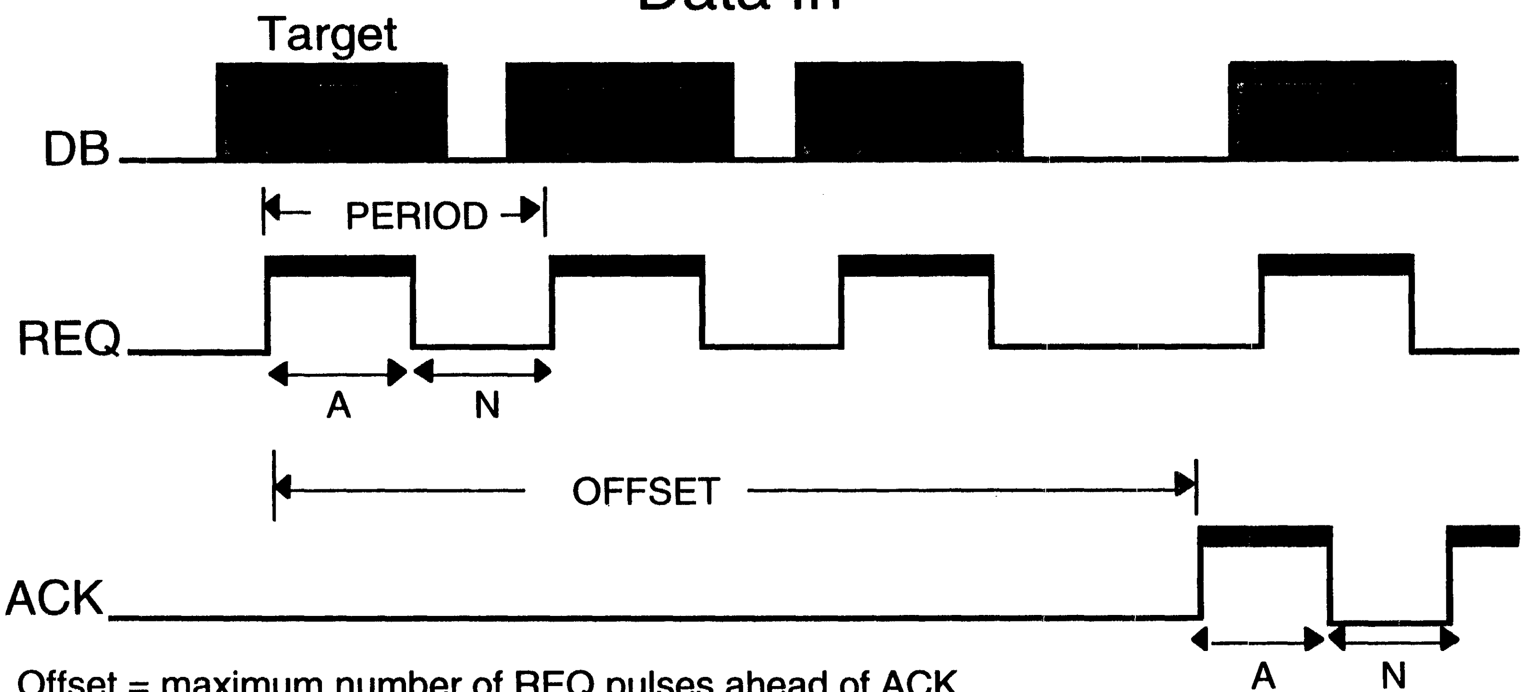

- Offset = maximum number of REQ pulses ahead of ACK
- A = Minimum Assertion Time
- N = Minimum Negation Time for REQ and ACK
- Period = Minimum interval between two REQ or ACK pulses.

Synchronous Transfer Example
(Continued)
Data Out

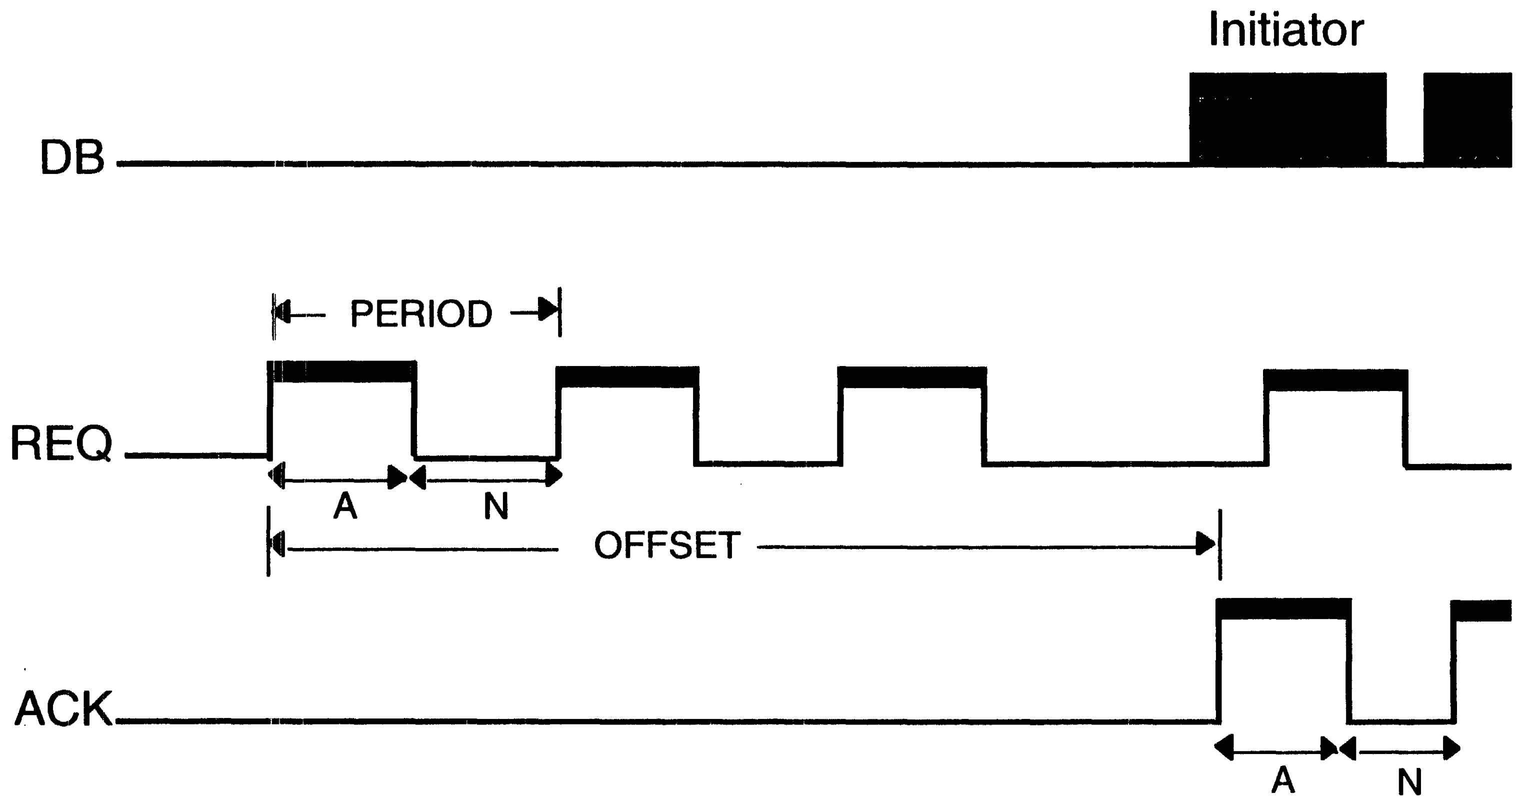

Synchronous REQ/ACK Handshake

Example with Offset 3

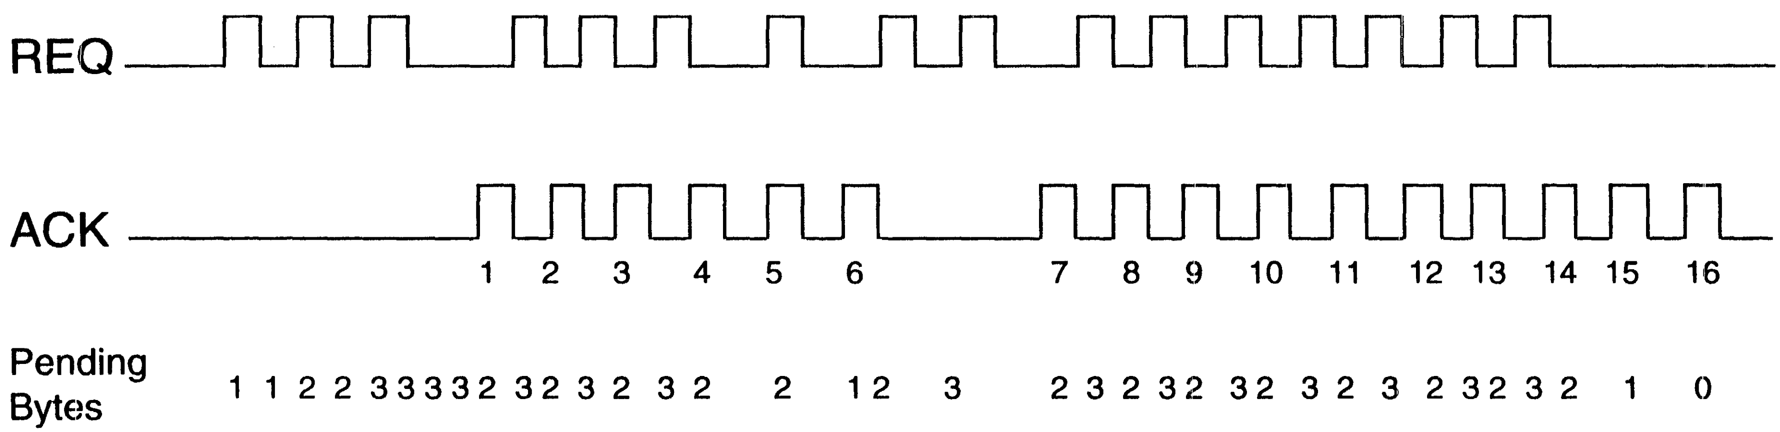

Basic Rules:

- ◆ ACK is never ahead of REQ
- ◆ REQ is never ahead of ACK by more than the offset
- ◆ Number of REQ and ACK pulses must be equal at the end of the phase
- ◆ Target doesn't change phase until all the ACK pulses are received

Slow SCSI

- Synchronous transfer with period ≥ 200 ns.
- Minimum Assertion and negation times of 90 ns (SCSI-2), 70 ns (SCSI-3)

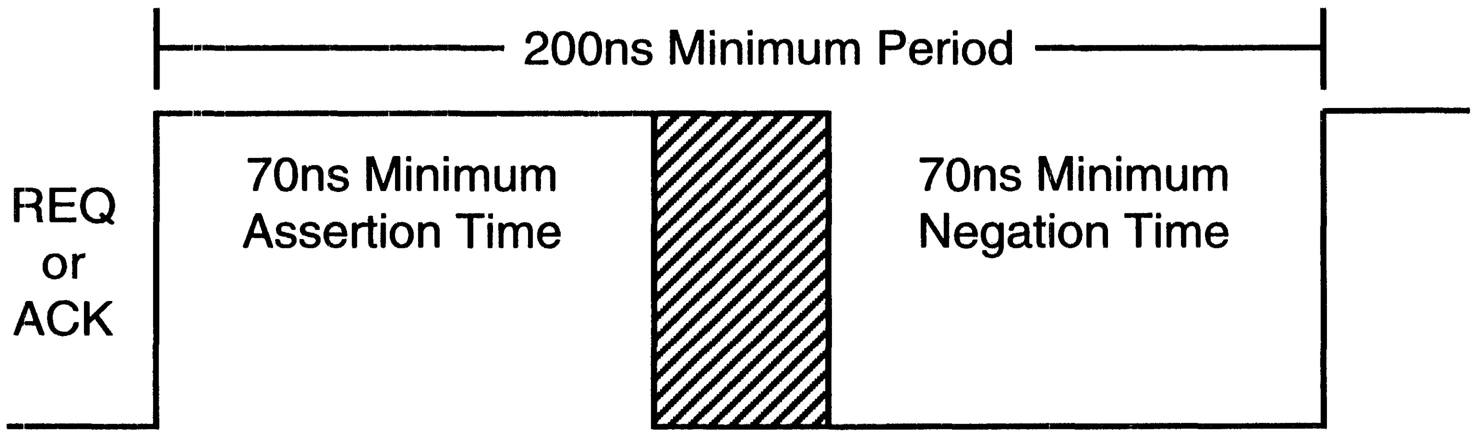

Assertion and Negation Times

Fast SCSI

- Synchronous Transfer with Period of 100-196 ns.
- Assertion and Negation Periods of 30 ns.
- Up to 10 MBPS Over 8-Bit Bus.
- Synchronous Transfer Negotiation Needed.
- Single-ended Implementation Recommendations:
  - max cable length of 3 meters.
  - active negation drivers for REQ, ACK, and DB.
  - active terminators.

Fast -20

- SCSI-3 Timing extension to parallel SCSI.
  - 20 MB/sec narrow, 40 MB/sec wide.
- Basic Rules
  - Negotiate for synchronous period of 50 - 96 ns
  - Active termination and active negation required (single-ended)
  - 25 pf node capacitance (same as SCSI-2)
- Configuration profiles (single-ended)
  - 4 devices with 3 m max cable length
  - 8 devices with 1.5 m max
- Up to 25 m for differential.

Synchronous Negotiation

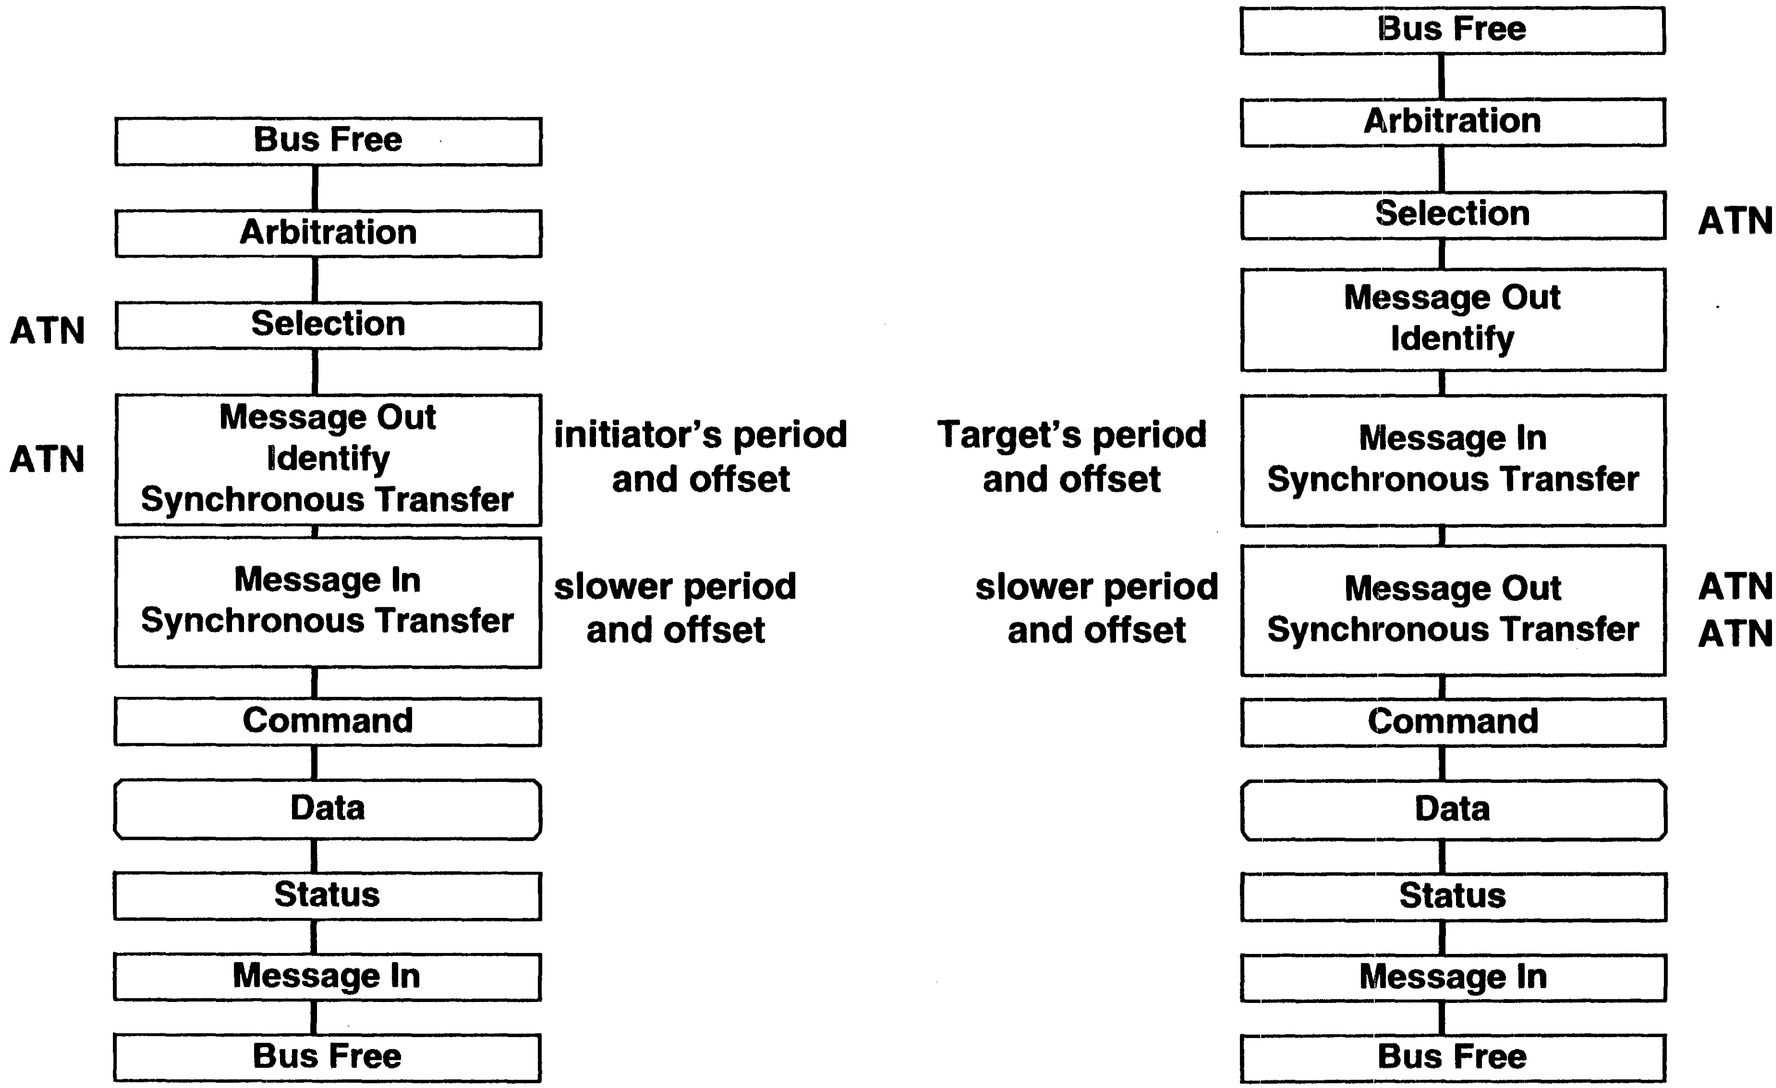

Synchronous Negotiation (Continued)

A Few Reminders

- Offset
  - 00h = Asynchronous
  - FFh = Infinite (memory is fast enough to keep up with synchronous)
- If Synchronous is not supported, respond with Synchronous message with offset 0, or Message Reject (07h) message.
- Negotiation agreements are invalidated by power cycle, hard reset, and wide negotiation
- Negotiation with every Inquiry and Request Sense Command is Recommended. This protects against cases where the target reverts to asynchronous after Reset or power cycle, while the initiator is still synchronous based on previous negotiation with that target.

Synchronous Negotiation (Continued)

Example with Offset 3

Byte	Value	Description
0	01h	Extended Message
1	03h	Extended Message Length
2	01h	Synchronous Data Transfer Request Code
3	m	Period/4 nanoseconds
4	x	Offset

- m = 12: 50 ns period
- m = 62 or 63: 250 ns period

Synchronous Negotiation(Continued)

Multi-Initiator Multi-Target Negotiation Exercise

	Initiator (offset, period) (64, 50ns)	Initiator (offset, period) (15, 125ns)
Target (offset, period) (16, 100ns)	(____, ____ ns)	(____, ____ ns)
Target Asynchronous	(____, ____ ns)	(____, ____ ns)
Target (offset, period) (8, 200ns)	(____, ____ ns)	(____, ____ ns)
Target (offset, period) (12, 100ns)	(____, ____ ns)	(____, ____ ns)

Synchronous Negotiation (Continued)

Example

	Device A	Device B
Supported Periods (ns)	50, 100, 150	60, 120, 180
Negotiate for	50	60
Agreement	60	60
Transmit at	100	60
Can Receive at up to	50	60

This page is intentionally left blank.

Wide Bus

Wide Bus

Summary

Standard	Byte Width	Bit Width	Cable Name	Pin Count	Max Transfer Rate MByte/sec	Max SCSI Devices	Description
SCSI-1	1	8	A	50	5	8	synchronous
SCSI-2	1	8	A	50	10	8	fast
SCSI-2	2	16	A+B	50+68	20	8	fast & wide with 2 cables
SCSI-2	4	32	A+B	50+68	40	8	fast & wide with 2 cables
SCSI-3	1	8	A	50	10	8	fast
SCSI-3	2	16	P	68	20	16	fast & wide with 1 cable
SCSI-3	1	8	A	50	20	8	fast-20
SCSI-3	2	16	P	68	40	16	fast-20 & wide
SCSI-3	4	32	P+Q	68+68	40	32	fast & wide with 2 cables
SCSI-3	4	32	P+Q	68+68	80	32	fast-20 & wide with 2 cables

- A is most commonly used.
- P is used for most Wide SCSI designs

Wide Bus (Continued)

Two-Byte Wide SCSI-3 Implementation

- Single Cable (P-Cable) with 68 Lines
- Signals:
  - Control = BSY SEL ATN RST REQ ACK C/D I/O MSG
  - Data = DB0-7, P
DB8-15, P1
- Narrow (1-byte) is mandatory, wide is optional, default is narrow
- Wide bus negotiation using extended messages needed to use wide
- Agreements invalidated with power cycle and hard reset
- Only data phases may use wide, other information transfer phases use narrow only
- Ignore Wide Residue message is used when data transfers do not fit on wide boundary

Wide Bus (Continued)

Examples of mixing Narrow and Wide Devices

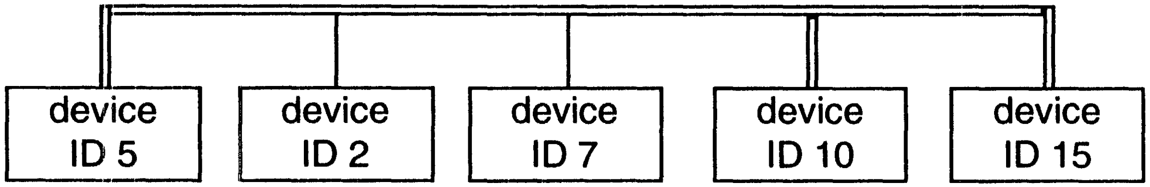

Questions:

1. If device 5 is an Initiator, which devices can it select?
2. If device 2 is an Initiator, which devices can it select?
3. If device 15 is an Intiator and selects device 7, can device 7 disconnect and reconnect?

SCSI Features

This page is intentionally left blank.

Messages

Message System

Types of Messages

First Byte of Message	Message Format
00h	One-Byte (Command Complete)
01h	Extended (Multiple-Byte)
02h-1Fh	One-Byte
20h-2Fh	Two-Byte (SCSI-1: Reserved)
30h-7Fh	Reserved
80h-FFh	One-Byte (Identify)

Byte	Value	Description	
0	01h	Extended message	
1		Extended message length	number of bytes to follow
2		Extended message code	see next page
:		Extended message arguments	

Extended Message Format

Messages

Message	Code	Dir	I	T	Notes
Command Complete	00h	In	M	M	Command Done - Bus Free
Extended Message	01h	I/O	O	O	Multiple Bytes
Modify Data Pointer	*00h*	*In*	*O*	*O*	*Pointer Management*
Synchronous Transfer	*01h*	*I/O*	*O*	*O*	*Synchronous Negotiation*
Wide Bus Transfer	*03h*	*I/O*	*O*	*O*	*Wide Bus Negotiation*
Save Data Pointer	02h	In	O	O	Pointer Management
Restore Pointers	03h	In	O	O	Pointer Management
Disconnect	04h	In	O	O	Bus Free, Reconnect Later
		Out	O	O	Request a Disconnect Normally not used
Initiator Detected Error	05h	Out	M	M	Parity Error Detected

Messages (Continued)

Message	Code	Dir	I	T	Notes
Abort	06h	Out	O	M	Clear Command, Bus Free Directly
Message Reject	07h	I/O	M	M	Inappropriate Message Received
No Operation	08h	Out	M	M	No Msg Out Available
Message Parity Error	09h	Out	M	M	Parity Err During Msg In
Linked Command Comp.	0Ah	In	O	O	Linked Commands
Linked C.C. with Flag	0Bh	In	O	O	Linked Commands
Bus Device Reset	0Ch	Out	O	M	Reset Selected Device
Abort Tag	0Dh	Out	O	O	Tagged Queueing

Messages (Continued)

Message	Code	Dir	I	T	Notes
Clear Queue	0Eh	Out	O	O	Tagged Queueing
Initiate Recovery	0Fh	I/O	O	O	Extended Contingent Allegiance (ECA)
Release Recovery	10h	Out	O	O	Clear ECA
Terminate I/O Process	11h	Out	O	O	Terminate with Status
Simple Queue Tag	20h	I/O	O	O	Tagged Queueing
Head of Queue Tag	21h	Out	O	O	Tagged Queueing
Ordered Queue Tag	22h	Out	O	O	Tagged Queueing
Ignore Wide Bus Residue	23h	In	O	O	Wide Bus Alignment
Identify	80h-FFh	Out	M	M	LUN, Disconnect
	80h-BFh	In	M	O	LUN, Pointers

Identify Messages

Rules to Remember

- The first Message Out sent by the initiator after Selection shall be:
  - Identify, Abort, or Bus Device Reset.
  - Any other message will cause the target to go to Bus Free.
- The first Message In sent by the target after Reselection is
  - Identify even if Attention is asserted.

Question: Why must the Target send an Identify message during reconnection?

- Identify Message In after Reselection implies Restore Pointers operation by initiator (described later).
- Disconnection Privilege is for current I/O process only and not subsequent ones (default is no disconnection if no Identify is sent).

Identify Message

<table>
<tr><th>Bits</th><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>SCSI-3</td><td rowspan="2">Identify Msg</td><td rowspan="2">Disconnect Privilege</td><td colspan="6">LUN</td></tr>
<tr><td>SCSI-2</td><td>LUNTAR</td><td colspan="2">Reserved</td><td colspan="3">LUN</td></tr>
</table>

7 1 indicates Identify Message

6 Disconnection Privilege

- 0 no disconnect
- 1 disconnect allowed (invalid with Message In)

0-5 LUN (SCSI-3)

5 LUNTAR (SCSI-2 only), normally 0 (not supported)

- 0 Bits 0-2 specify LUN
  (Command is for Logical Unit)
- 1 Bits 0-2 specify Target Routine
  (Command is for Target Controller)
  For use with Inquiry and Request Sense Only

3-4 RESERVED in SCSI-2 or LUN in SCSI-3

0-2 LUN or Target Routine Number (Vendor-Unique)

Linked Commands

- Last CDB Byte (Control Byte)

7	6	5	4	3	2	1	0
Vendor Unique		Reserved			NACA SCSI-3	Flag	Link

- Flag bit is normally not supported (left over from early SCSI days).

Linked Commands (Continued)

- Bus Free
- First Linked Command
  - Arbitrate
  - Select - ATN
  - Message Out (IDENTIFY)
  - Command (LINK CDB Bit Set)
  - Data
  - Status 10h (Intermediate)
  - Message In 0Ah (Linked Command Complete)

Linked Commands (Continued)

- Next Linked command(s)
  - Command (LINK CDB Bit Set)
  - Data
  - Status 10h (Intermediate)
  - Message In 0Ah (Linked Command Complete)
- Last Command
  - Command (LINK CDB Bit Clear)
  - Data
  - Status 00h (Good)
  - Message In 00h (Command Complete)
- Bus Free

Linked Commands (Continued)

- Optional Feature, not Commonly Used
- For Sending Multiple Commands Without Freeing the Bus
- If No Disconnection Privilege is Granted, Then No Bus Free, No Arbitration, and No Selection Between Linked Commands
- Link Bit in CDB is Used to Link Commands
- Intermediate (10h) Status and Linked Command Complete (0Ah) Message Between Linked Commands - Other Status Terminates I/O Process
- Normally the Initiator Does Not Give the Target the Disconnection Privilege While Linking
- If Given the Disconnection Privilege, the Target May Disconnect While Linking

This page is intentionally left blank.

Pointers

Host Adapter Pointers

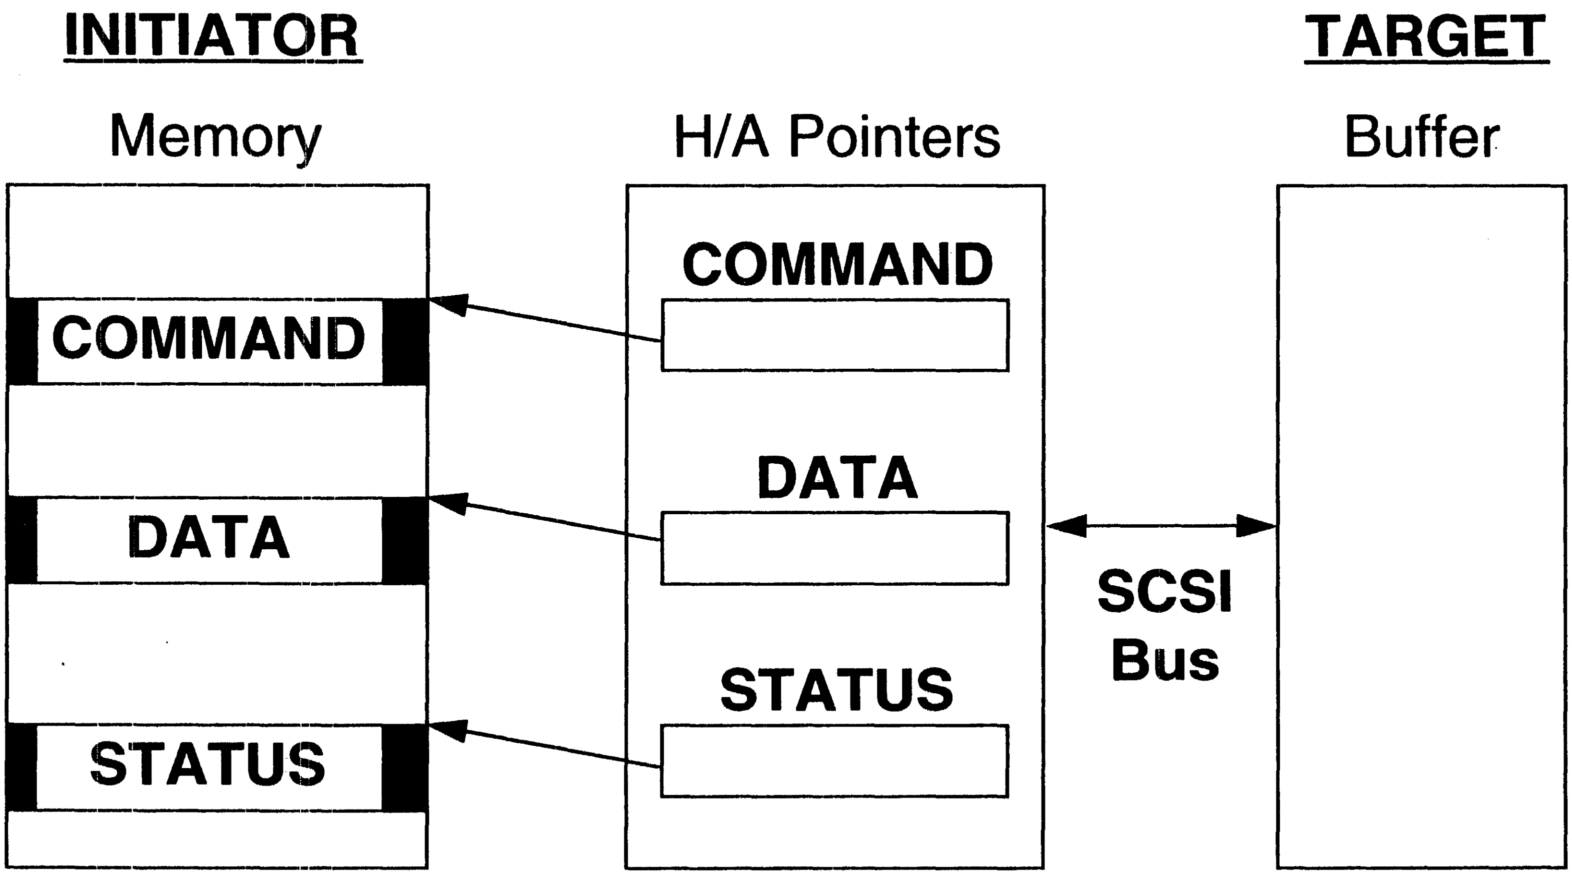

- Pointers are allocated by the host device driver and handed over to the host adapter.
- Pointers move as bytes are transferred during the I/O process.

Host Adapter Pointers (Continued)

Current (Active) Pointers

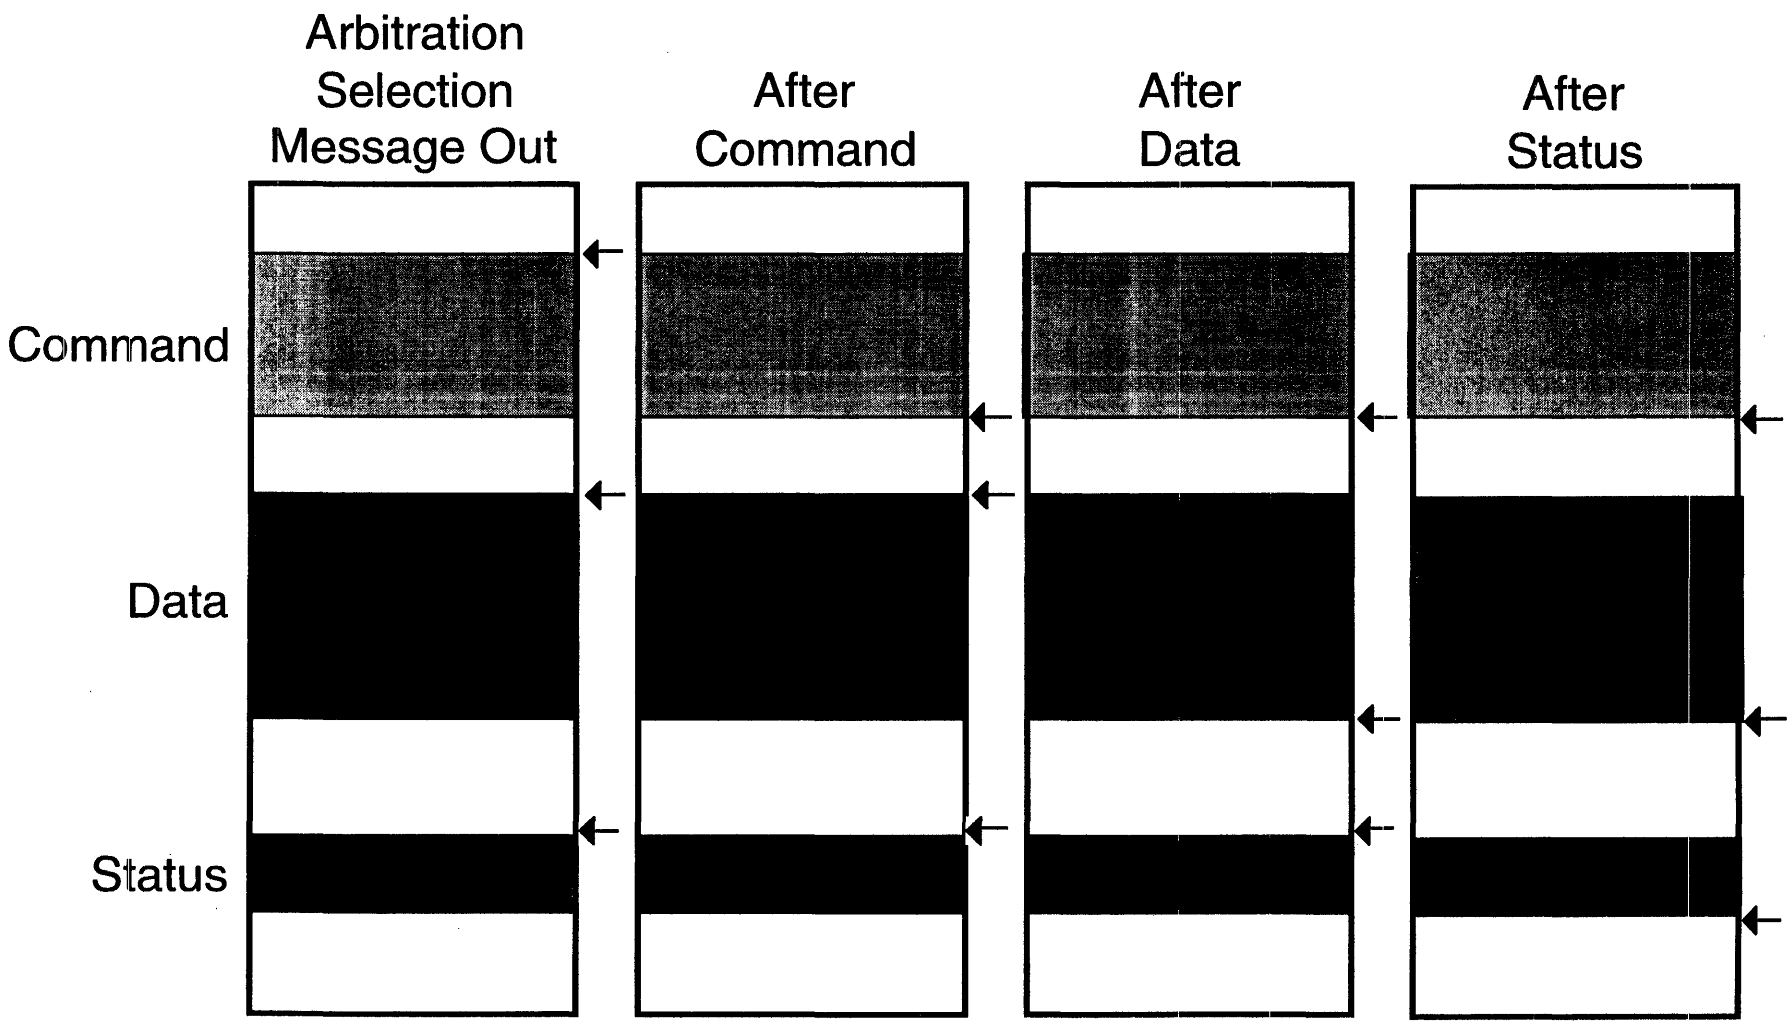

Host Adapter Pointers (Continued)

Command Pointer

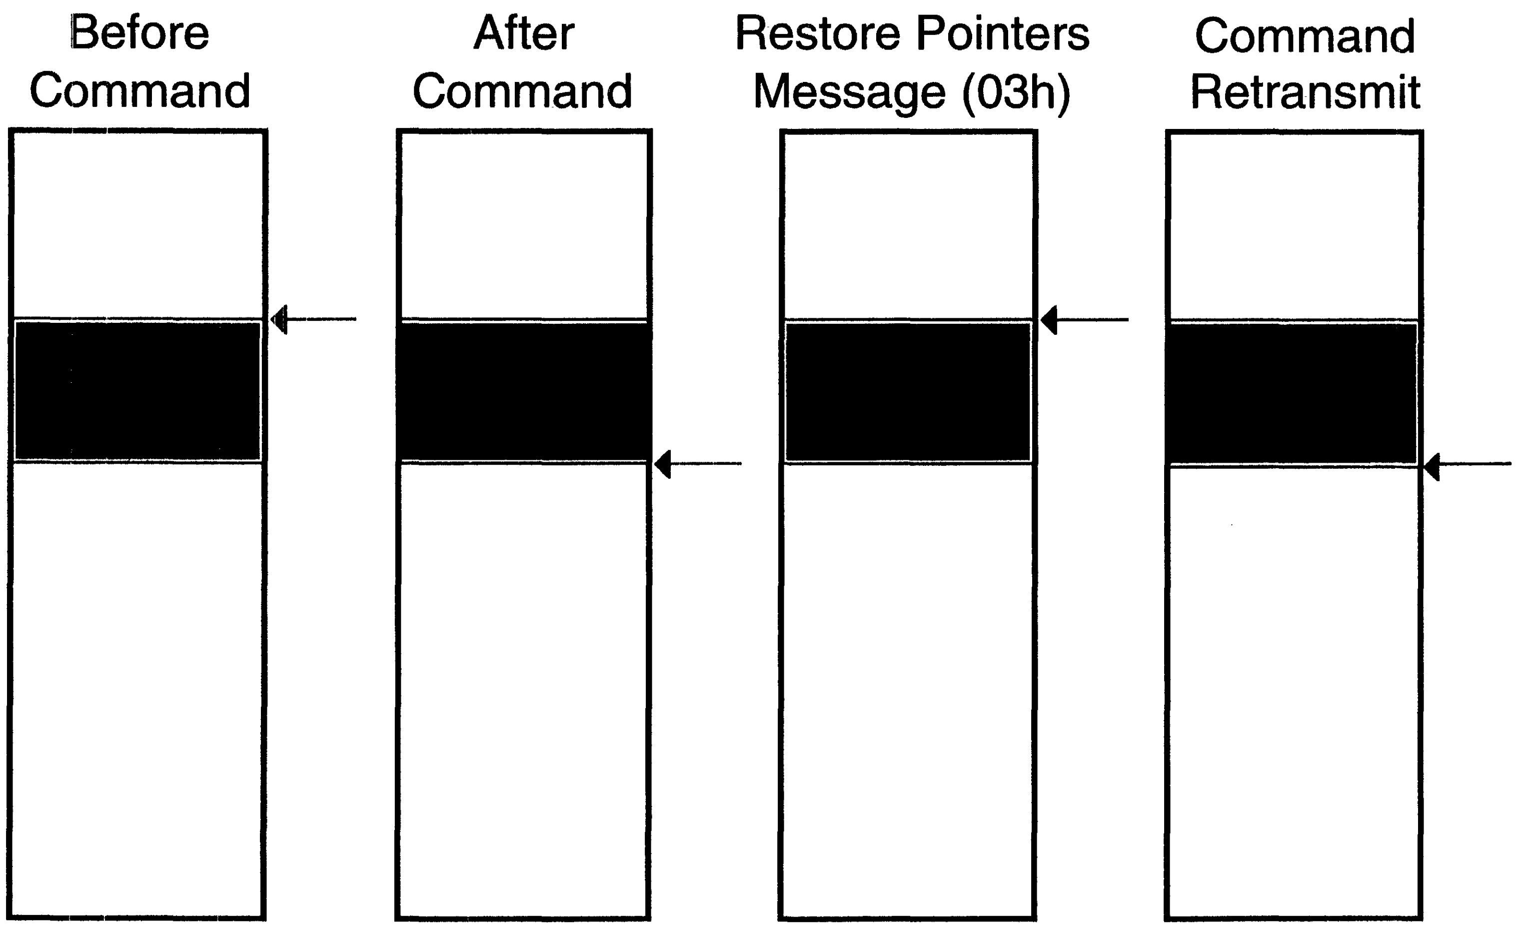

Host Adapter Pointers (Continued)

Data Pointer

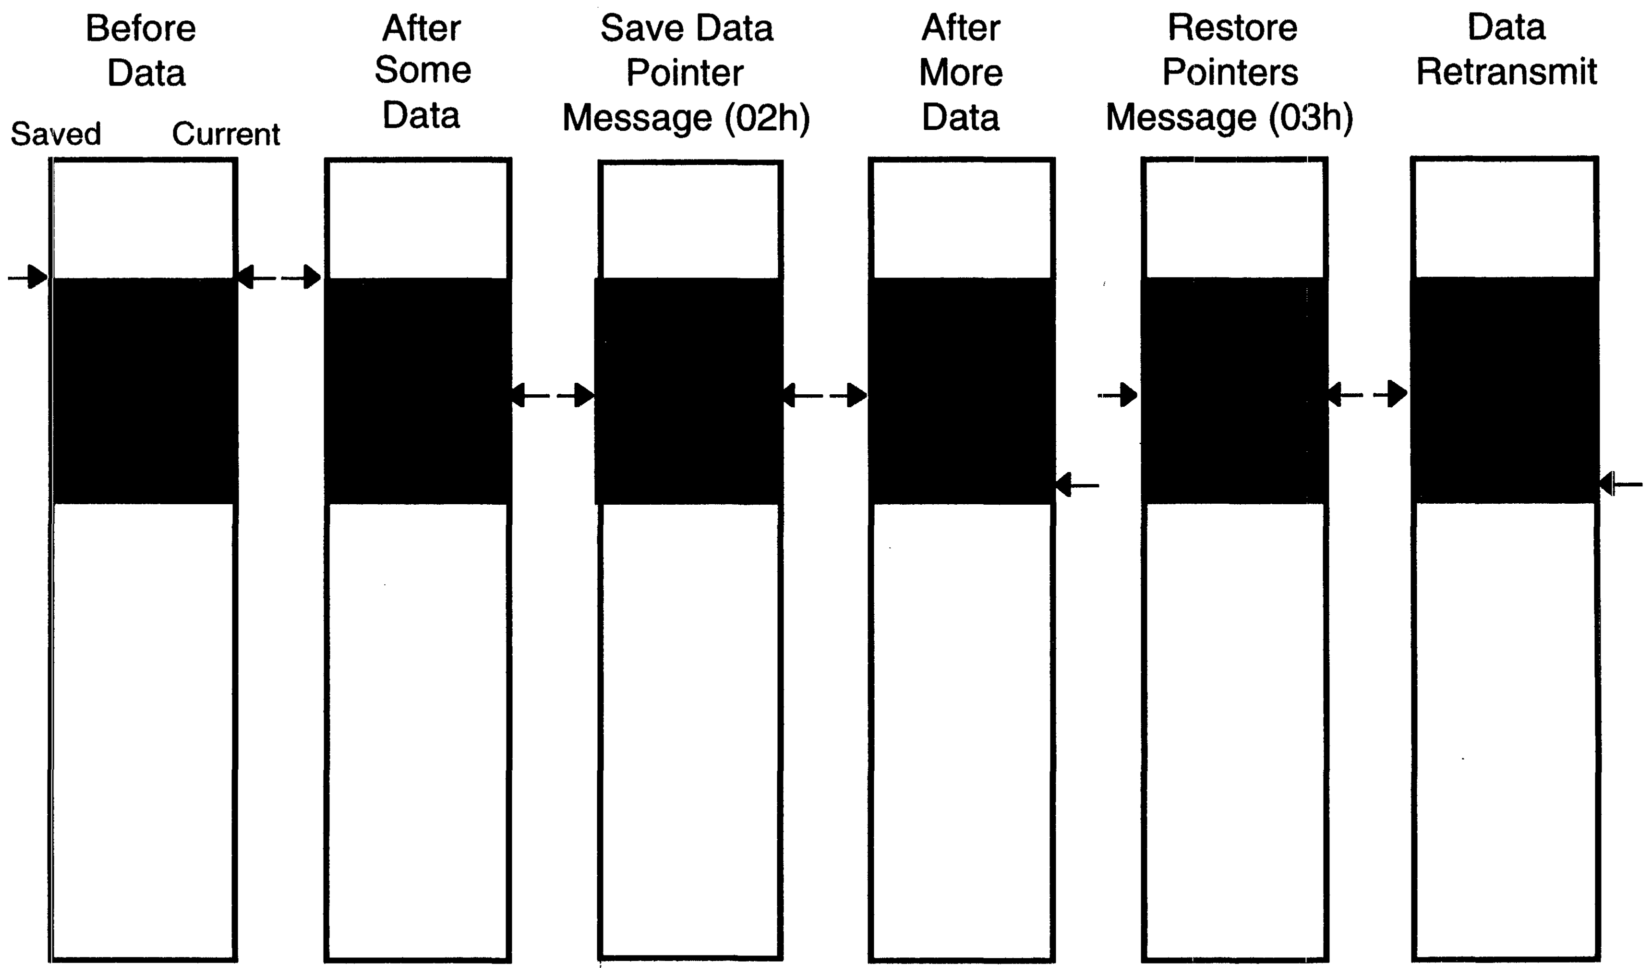

Host Adapter Pointers (Continued)

- Pointer Set:
  - Command, Status, and Data
- Save Data Pointer:
  - Saved Data Pointer = Current Data Pointer
- Restore Pointers:
  - Current Data Pointer = Saved Data Pointer
  - Current Command Pointer = Starting Command Pointer
  - Current Status Pointer = Starting Status Pointer
- One set of Pointers per Pending I/O Process. (Application Client)

Host Adapter Pointers (Continued)

Bus Free
- Arbitration
  - Selection - ATN
    - Message Out - Identify — LUN + Disconnect Allowed
    - Command
    - Data
    - Message In - Save Data Pointer
      - Disconnect

Bus Free
- Arbitration
  - Reselection
    - Message In - Identify — LUN + implied Restore Pointers
    - Data
    - Message In - Save Data Pointer
      - Disconnect

Bus Free
- Arbitration
  - Reselection
    - Message In - Identify — LUN + implied Restore Pointers
    - Data
    - Status
    - Message In - Command Complete

Bus Free

Host Adapter Pointers (Continued)

Disk Write Example

1. Why would the devices disconnect at point 1?
2. Why would the device disconnect at point 2?
3. Is this SAVE DATA POINTER message necesary?

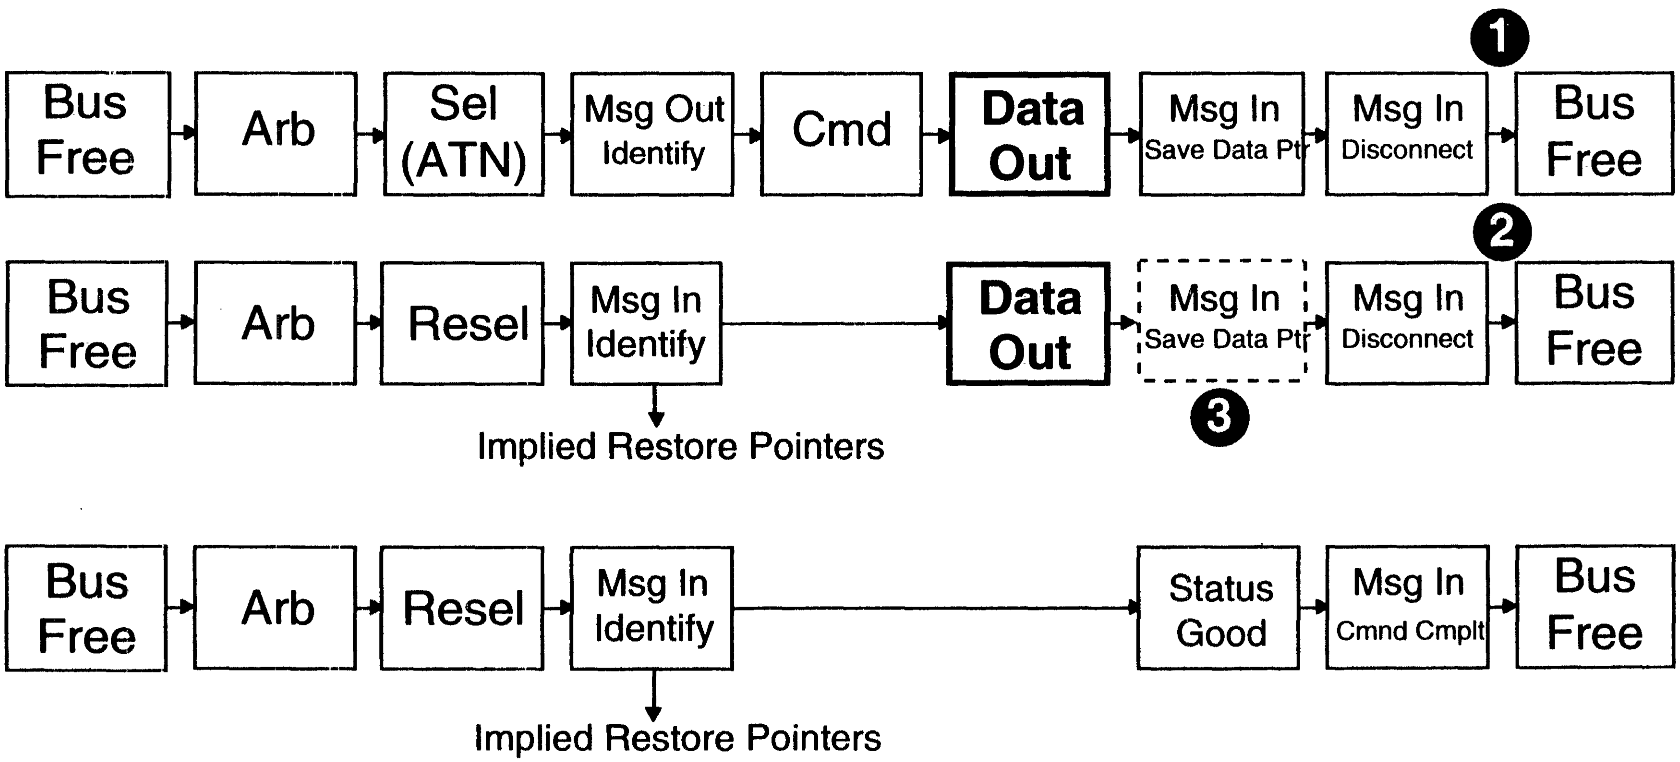

Host Adapter Pointers (Continued)

Disk Read Example

1. Is this SAVE DATA POINTER message necesary?
2. Why would the device disconnect at this point?
3. Why would the device disconnect at this point?

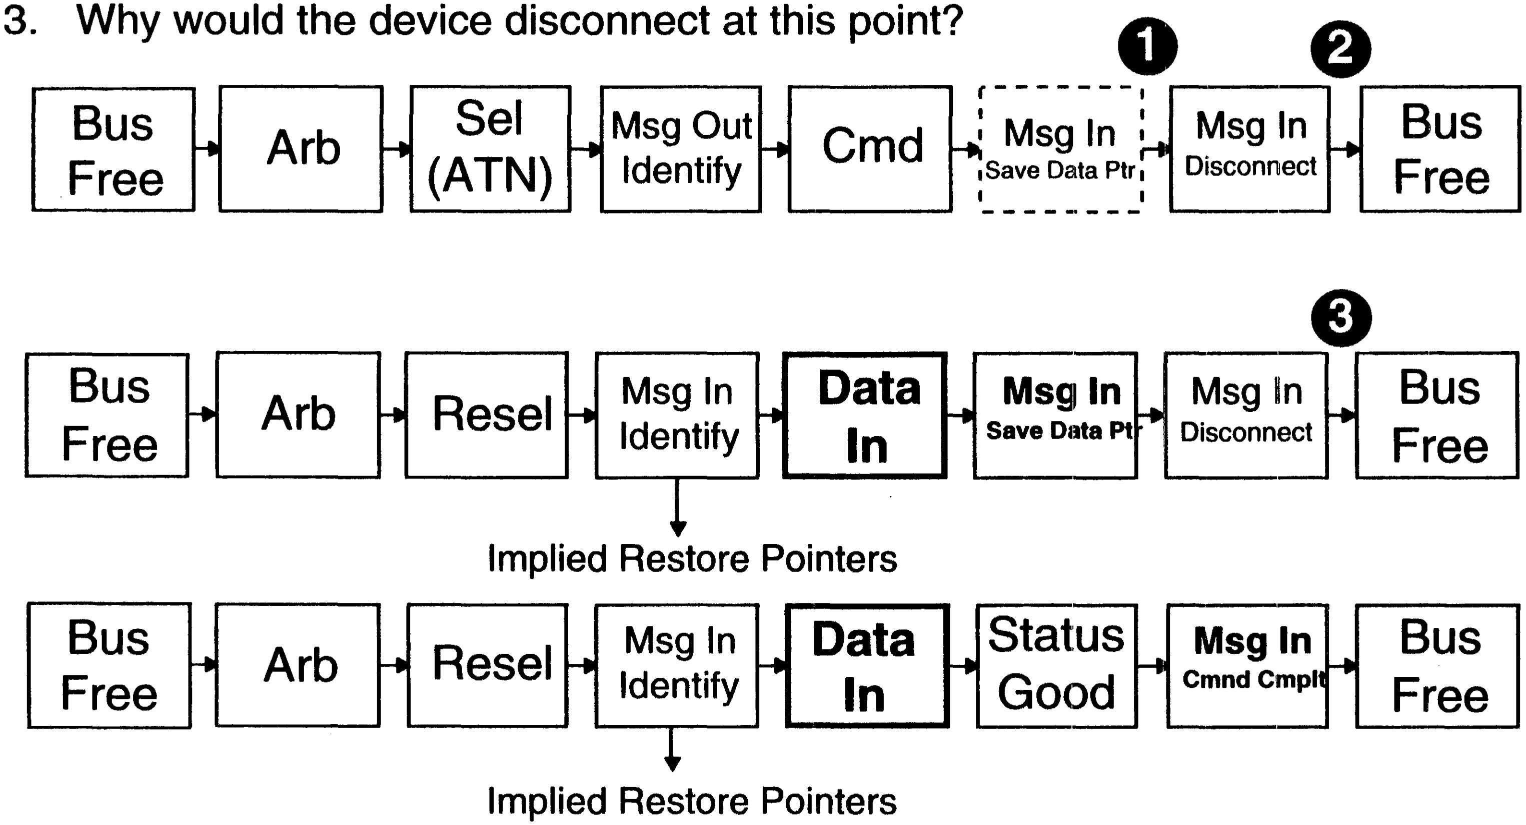

First Sector Up (Zero Latency)

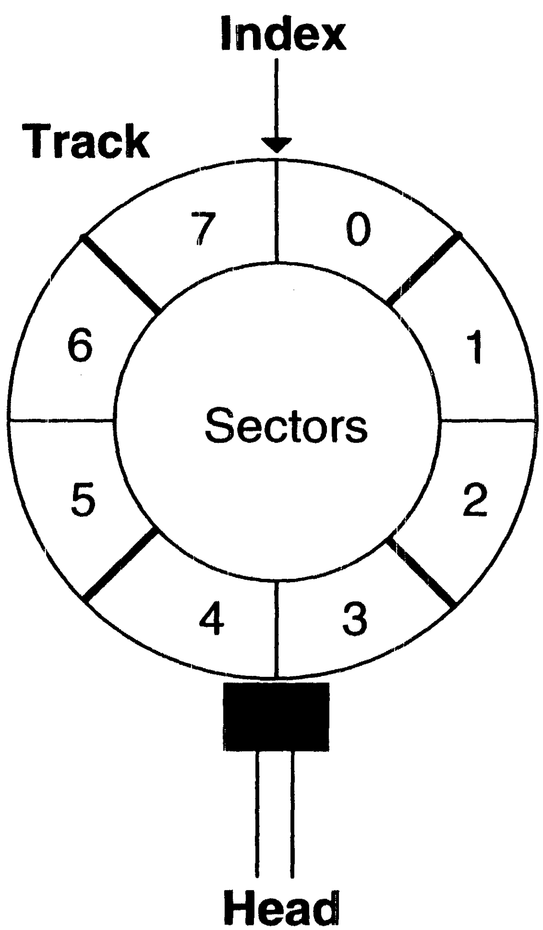

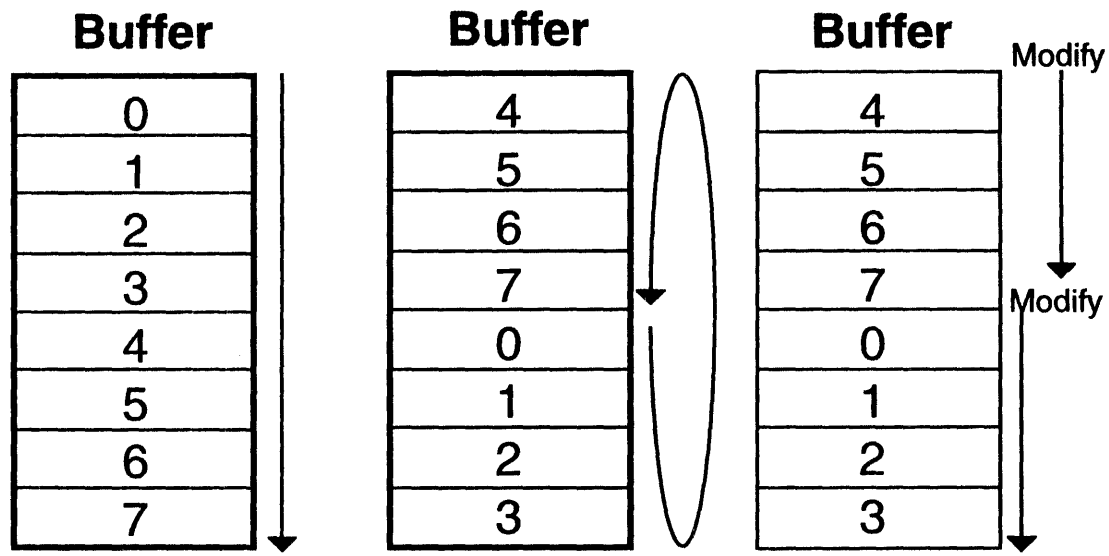

Simple firmware	Efficient firmware	Sophisticated firmware
Seek to 0 Read 0-7 Return 0-7	Read 4-7 Read 0-3 Return 0-7	Modify Data Pointer Read 4-7 Return 4-7 Modify Data Pointer Read 0-3 Return 0-3

Parity Error Handling

SCSI Tools Used

◆ **Message In**	03h	Restore Pointers
◆ **Message Out**	05h	Initiator Detected Error
	09h	Message Parity Error
◆ **Status**	02h	Check Condition
◆ **Sense Key**	4h	Hardware Error
	Bh	Aborted Command

◆ **Additional Sense Code (ASC) - in byte 12 of the sense data**

43 h	Message Error
48 h	Initiator Detected Error Message Received
4Ah	Command Phase Error
4Bh	Data Phase Error

Parity Error Handling (Continued)

General Summary

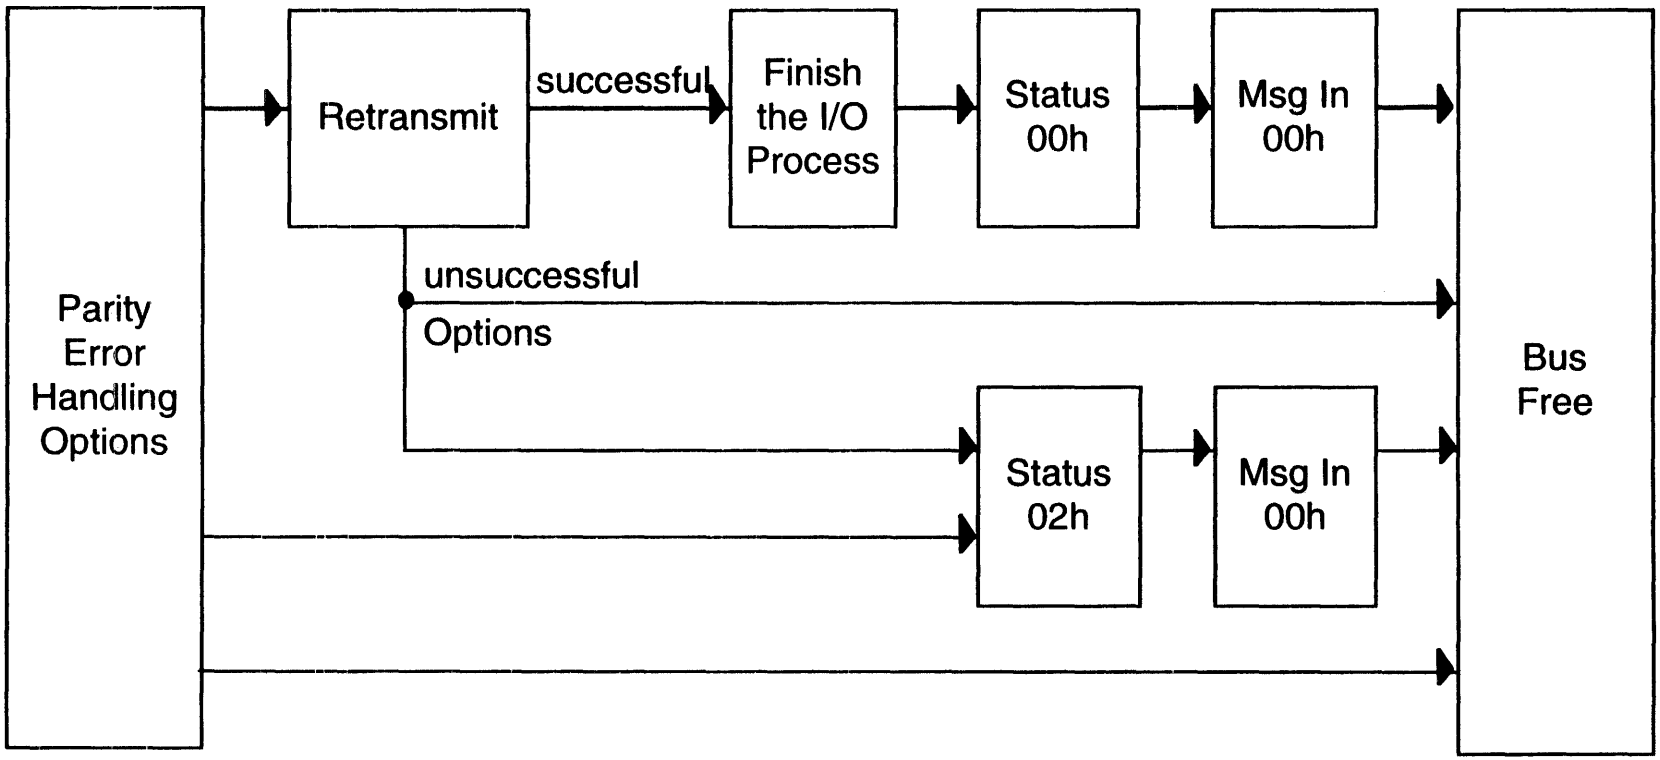

Parity Error Handling (Continued)

During Data Out Phase

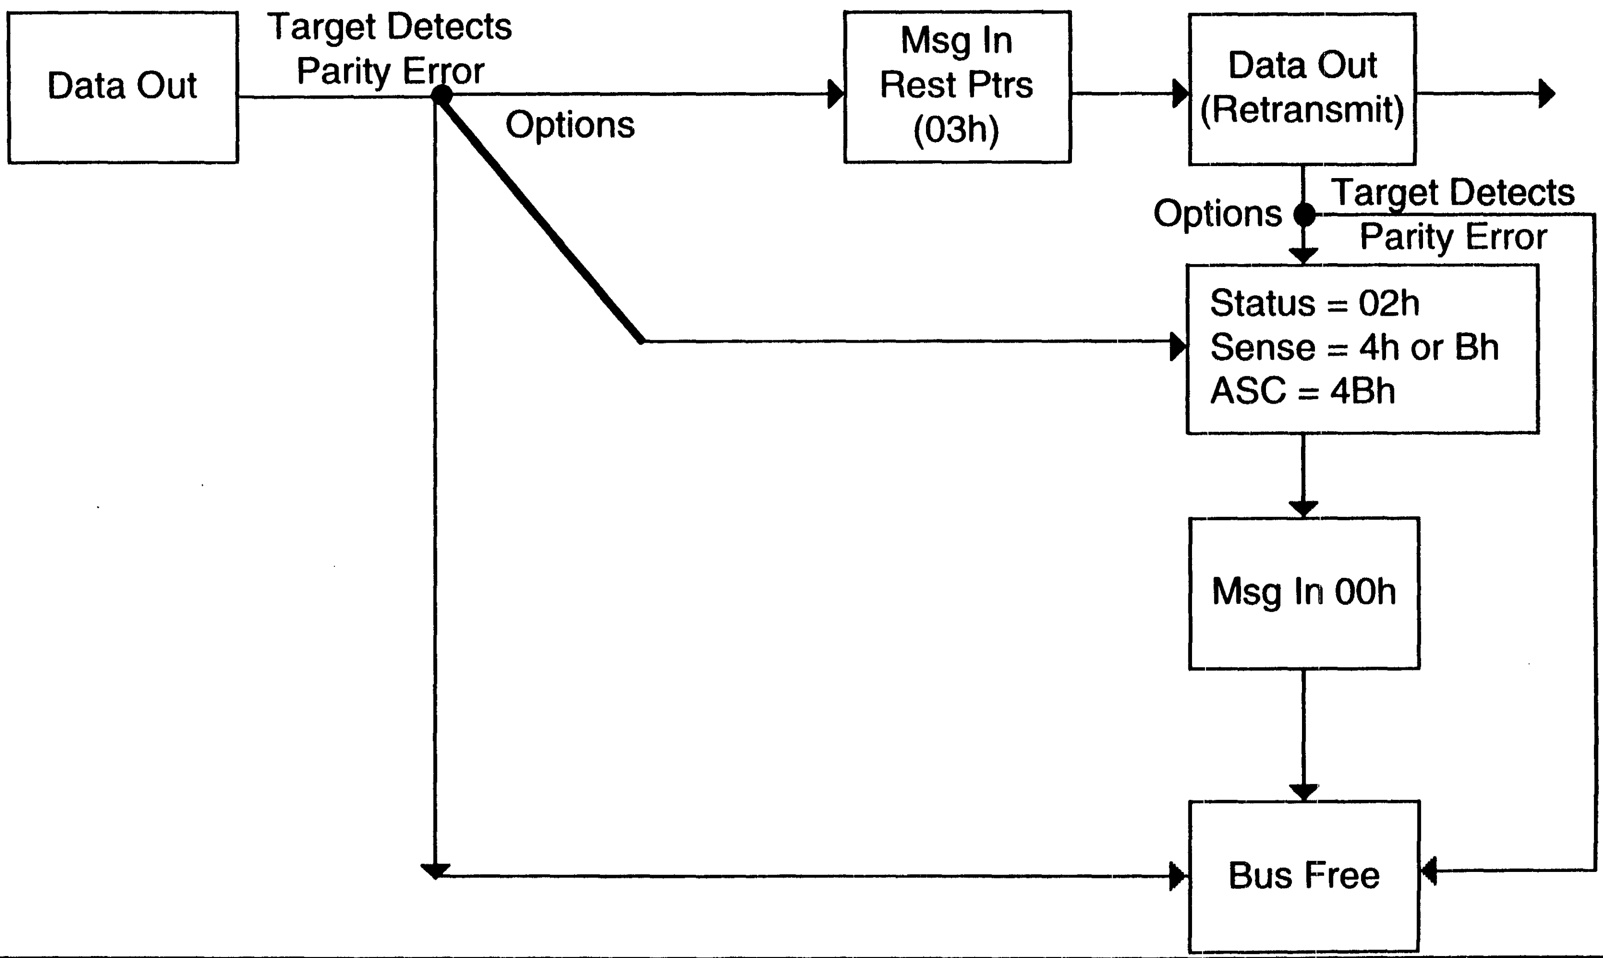

Parity Error Handling (Continued)

During Data In Phase

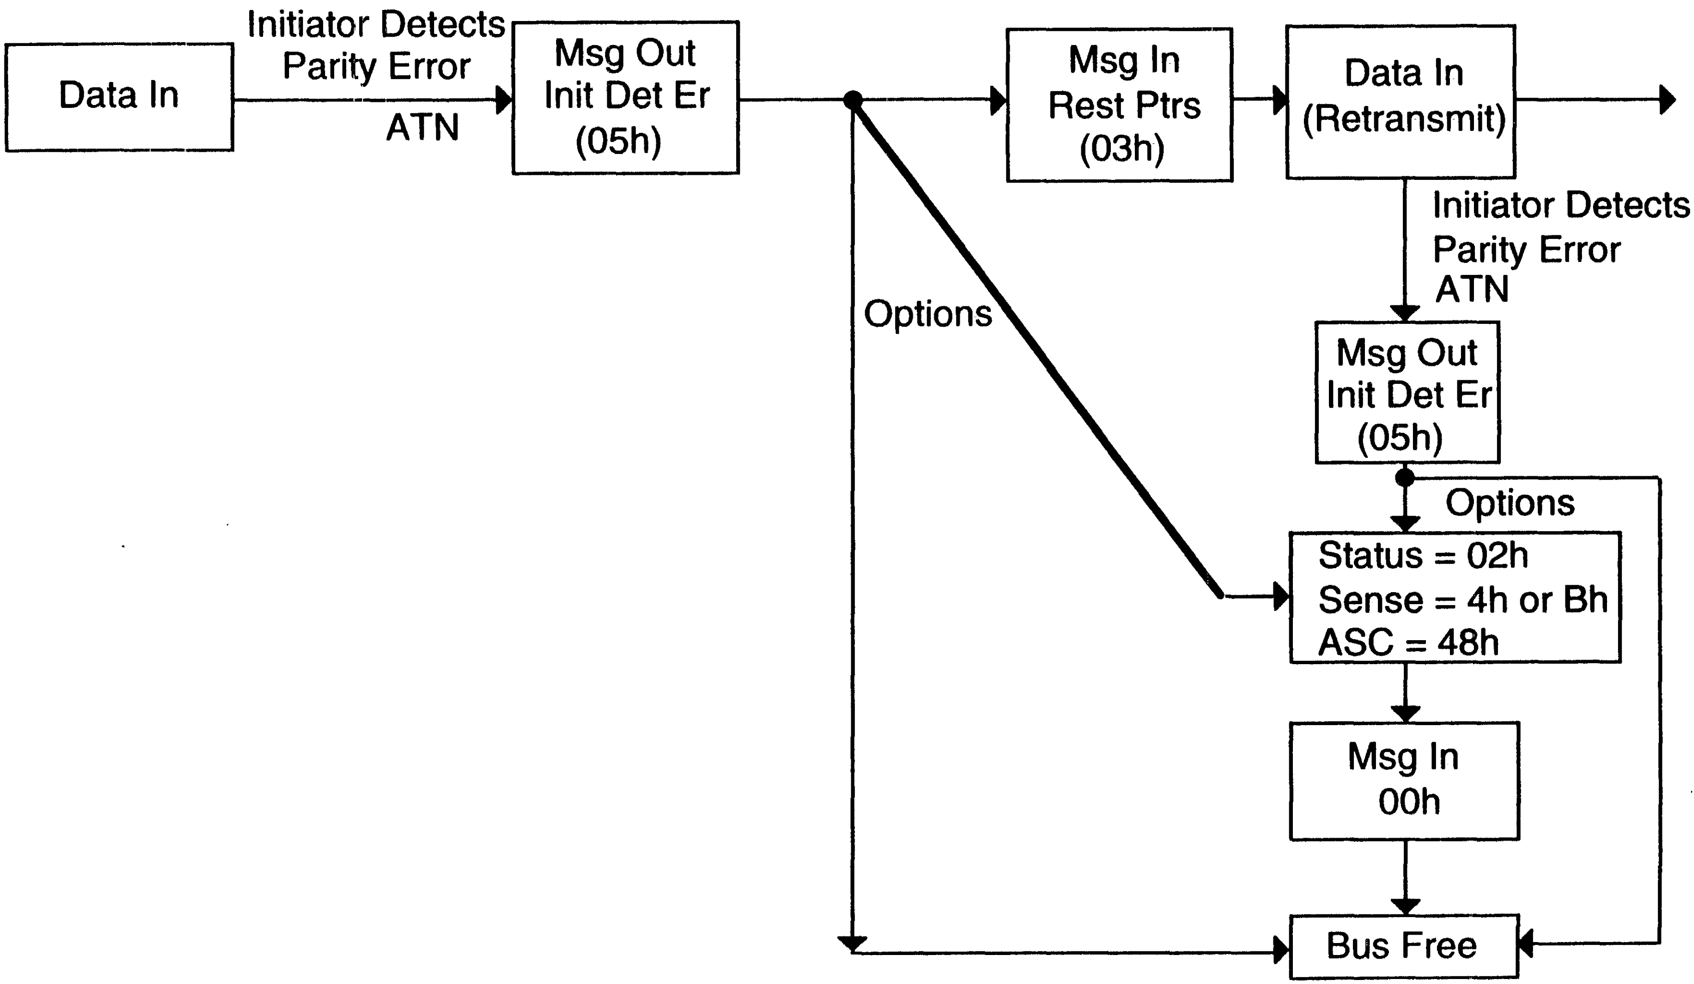

Command Queueing

Untagged Queueing

Example

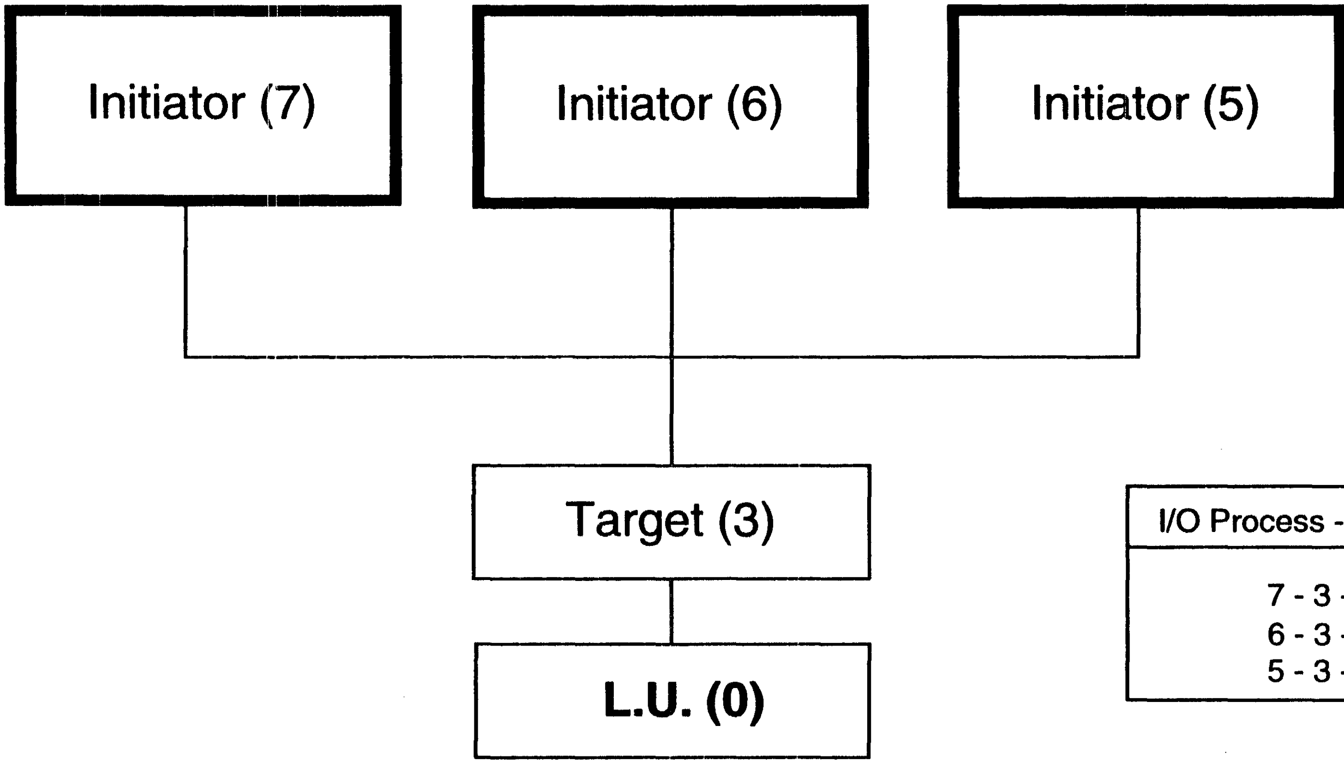

I/O Process - Nexus ID
7 - 3 - 0
6 - 3 - 0
5 - 3 - 0

- Optional feature in SCSI-1 and SCSI-2.
- Multiple initiators send one I/O process each, to the same target/L.U.
- The target queues the I/O process using disconnect/reconnect, instead of returning Busy status.

Tagged Queueing

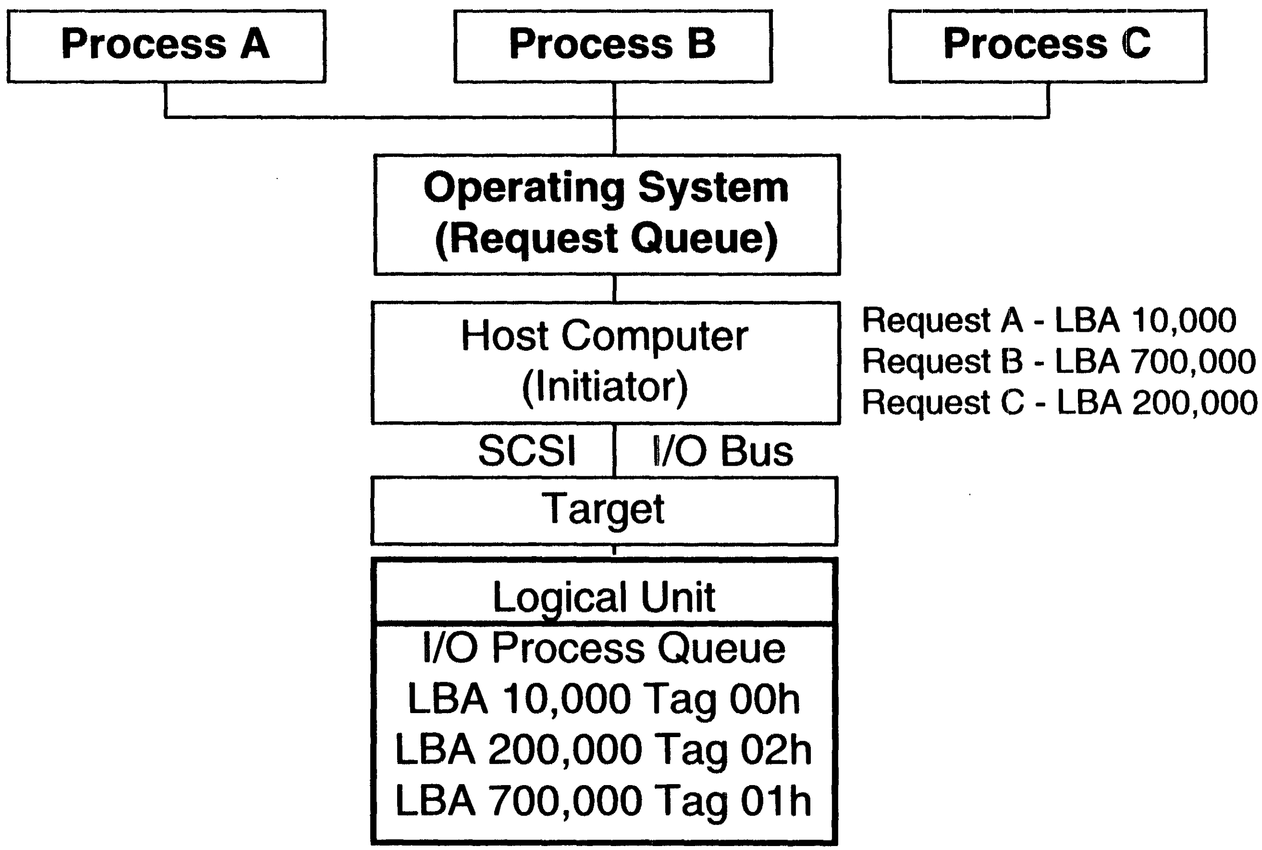

- Queueing is done at the target
- Tags are used to keep track of commands

Tagged Queueing (Continued)

- Types of Queueing:

Simple	Target decides order (e.g., elevator seek)
Head of Queue	Last In First Out
Ordered	Initiator decides order (First In First Out)

- Example: Requests for LBA 10,000 700,000 200,000 - Drive is at LBA 0

Simple	10,000	200,000	700,000
Head of Queue	200,000	700,000	10,000
Ordered	10,000	700,000	200,000

- Messages:

Code		Name	Direction
20h	Tag	Simple Queue	Out/In
21h	Tag	Head of Queue	Out
22h	Tag	Ordered Queue	Out

Tagged Queueing (Continued)

Connection

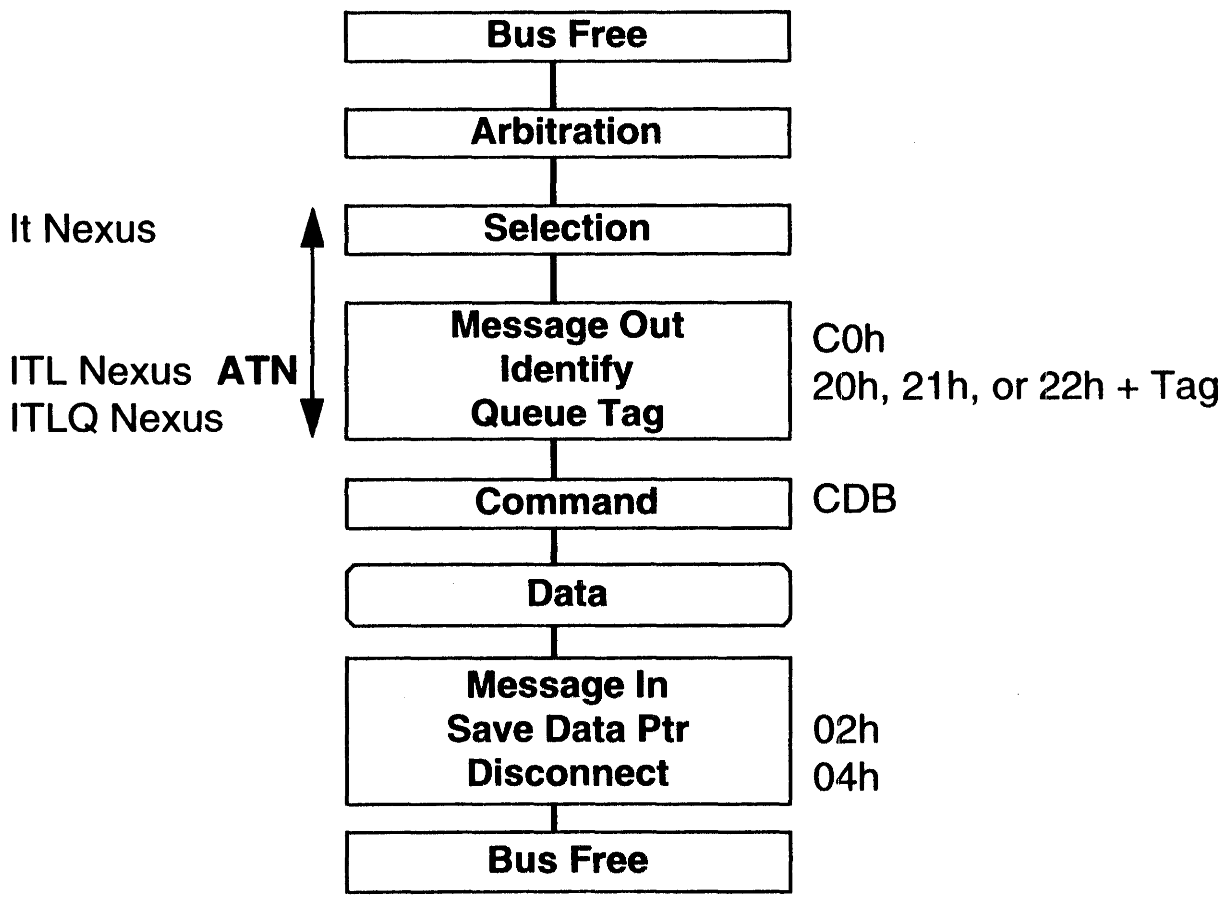

Tagged Queueing (Continued)

Reconnection

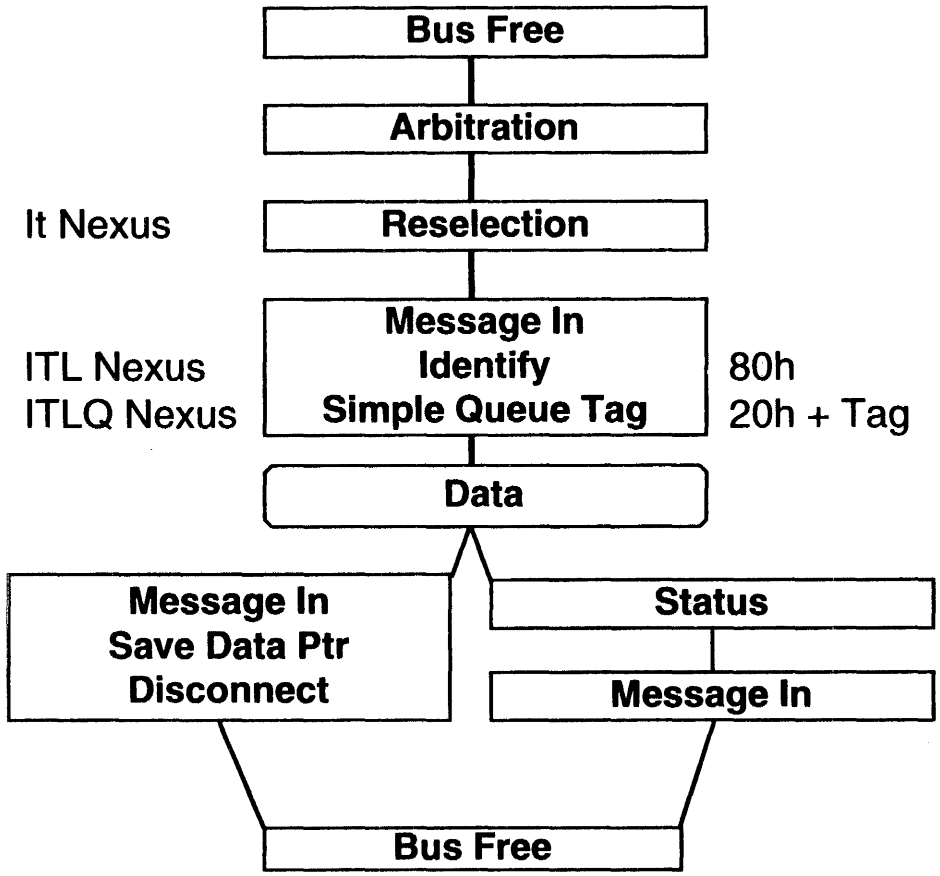

Tagged Queueing (Continued)

Mode Page 0AH, Control Parameters

Bit / Byte	7	6	5	4	3	2	1	0
0	PS	Rsvd	Page Code = 0Ah					
1	Page Length = 0Ah							
2	Reserved						GLTSD	RLEC
3	**Queue Algorithm Modifier**				Reserved		**QErr**	**DQue**
4	Rsvd	RAC	ByprtM	Bybths	Rsvd	RAERP	UAAERP	EAERP
5	Reserved							
6	(MSB) Ready AER Holdoff Period							
7	Ready AER Holdoff Period (LSB)							
8	(MSB) Busy Timeout Period							
9	Busy Timeout Period (LSB)							
10 - 11	Reserved							

Note: Highlighted fields affect Tagged Queueing behavior.
See next page for definition.

Tagged Queueing (Continued)

Control Parameter Description

DQue	0 =	Enable Tagged Queueing if supported by device.
	1 =	Disable Tagged Queueing.
QErr	0 =	Continued other queued commands after clearing (Auto) Contingent Allegiance for a command that failed.
	1 =	Abort other queued commands after clearing (Auto) Contingent Allegiance for a command that failed.
Queue Algorithm Modifier	0H =	Restricted reording required. Device must execute SIMPLE queued commands in a sequence that preserves the sequential integrity of data on the medium.
	1H =	Unrestricted reordering allowed.
	2H-7H	Reserved
	8H - FH	Vendor Specific

Question: Under what circumstances could unrestricted reordering cause a problem?

Tagged Queueing (Continued)

More Rules

- Ordered I/O Processes are Executed in the Order Received **(Initiator Order)**.
- Simple I/O Processes are Executed in the Order Determined by the Target **(Target Order)**.
- If a Set of Simple I/O Processes is Received Before a Set of Ordered Ones, the Simple I/O Processes are Executed First, and Vice Versa.
- Linked Commands Use One Queue Tag, i.e., They are Considered a Single I/O Process.

Tagged Queueing (Continued)

Rules

- Optional feature
- Maximum of 256 queued I/O processes per ITL combination, each with unique tag.
- Tags are assigned by the initiator.
- Tags must be unique for each ITL combination.
- When queue is full, QUEUE FULL status is returned to tagged I/O processes and BUSY to untagged ones.
- When queued I/O processes are pending, commands without Disconnect Privilege result in BUSY status.
- SCSI-2 did not permit mixing Tagged and untagged commands. In SCSI-3, one untagged command may be sent to a device with tagged commands pending. The untagged command is given the SIMPLE attribute.

Task Management Messages

Bus Device Reset	0Ch	Clear All I/O Processes from All Initiators Perform Reset Function Report Unit Attention
Clear Queue	0Eh	Clear All I/O Processes from All Initiators Report Unit Attention
Abort	06h	Clear All I/O Processes from this Initiator No Sense Data
Abort Tag	0Dh	Clear Current I/O Process from this Initiator No Sense Data
Terminate I/O Process	11h	Finish Command if possible Return COMMAND TERMINATED status Prepare Sense Data (Sense Key = Aborted Command)

Asynchronous Event Notification (SCSI-2)
Asynchronous Event Reporting (SCSI-3)

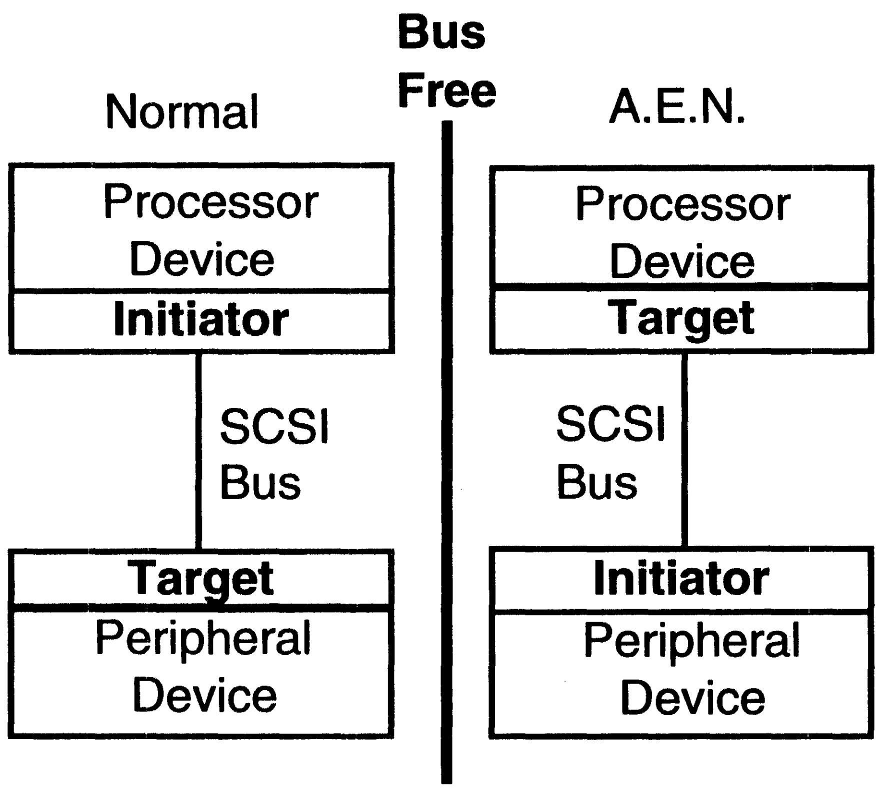

Asynchronous Event Notification
(Continued)

- Optional Feature, Currently Not widely Supported
- Used for Reporting:
  - device initialization completed
  - unit attention
  - errors
- Starts From Bus Free
- Target Becomes Initiator and Executes a "SEND" Processor Command with AEN Bit Set. Sense Data is Transferred During Data Phase.
- Reported to LUN 0 of Processor Devices that Support AEN
- After every reset, Peripheral Device must Scan the Bus, using Inquiry, to Find Processor Devices that Support AEN.

This page is intentionally left blank.

SCSI-3 Specifics

What Is SCSI-3

- It is an architecture
- It is a set of standards
- It is any of 4 different I/O interfaces
  - Parallel Bus
  - Fibre Channel (FC)
  - Serial Storage Architecture (SSA)
  - Serial Bus (also known as IEEE-1394 or Firewire)

SCSI-3 Documentation

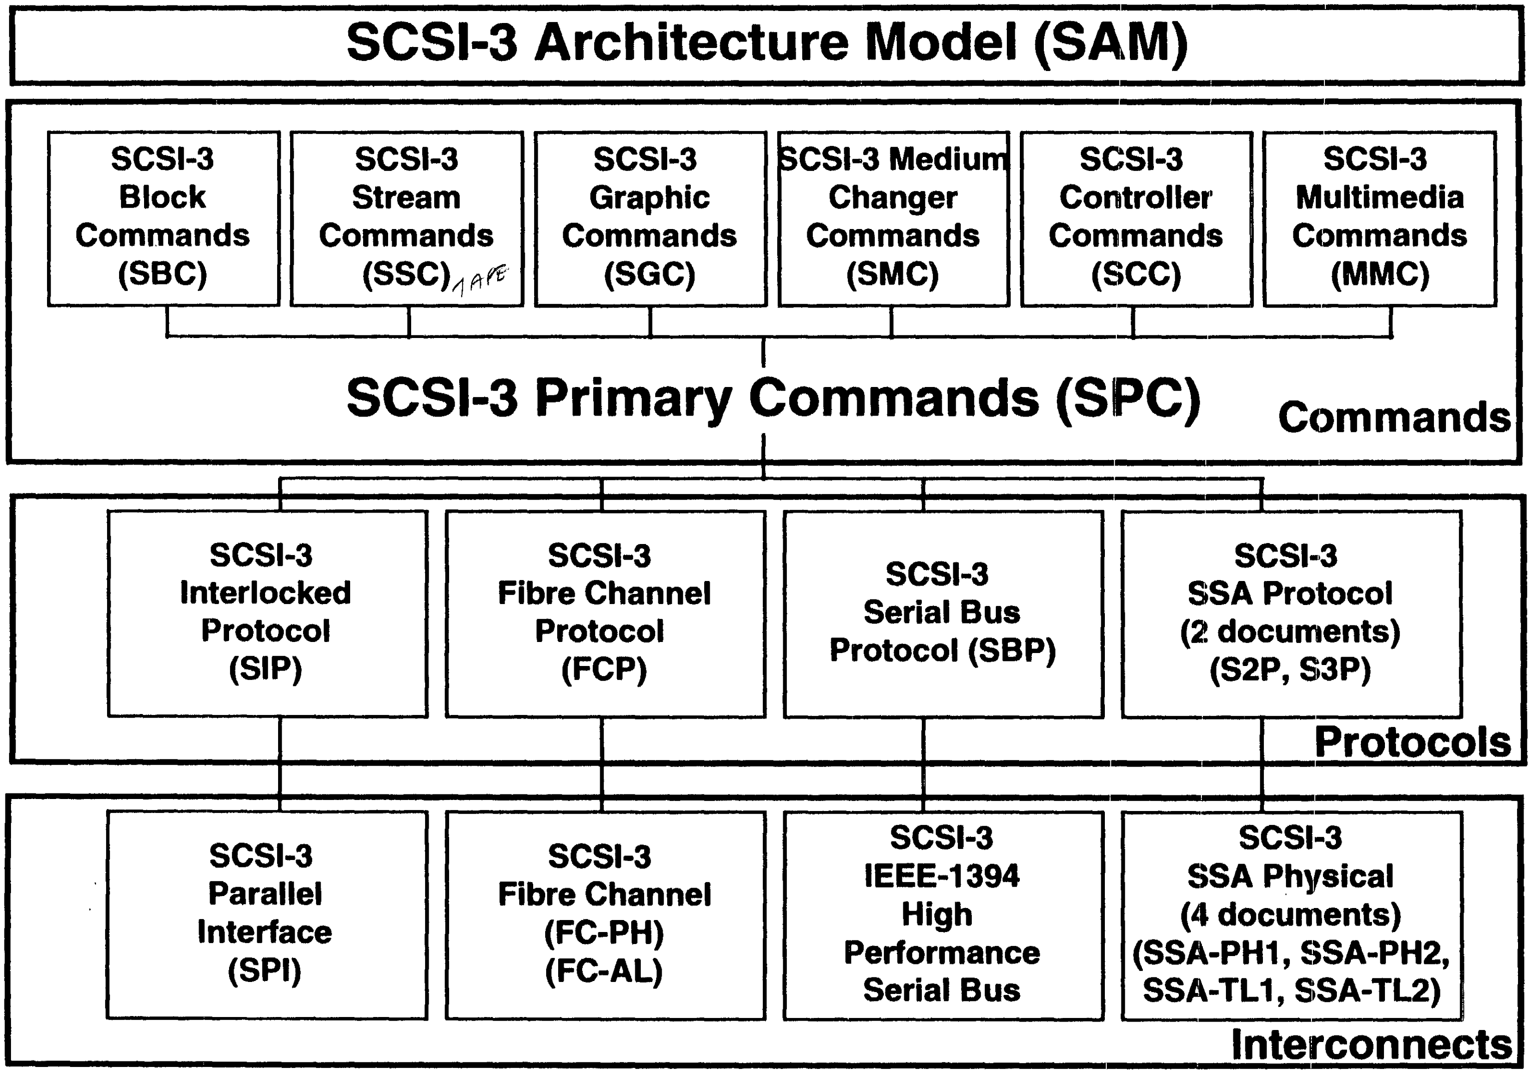

SCSI-3 Distributed Service Model

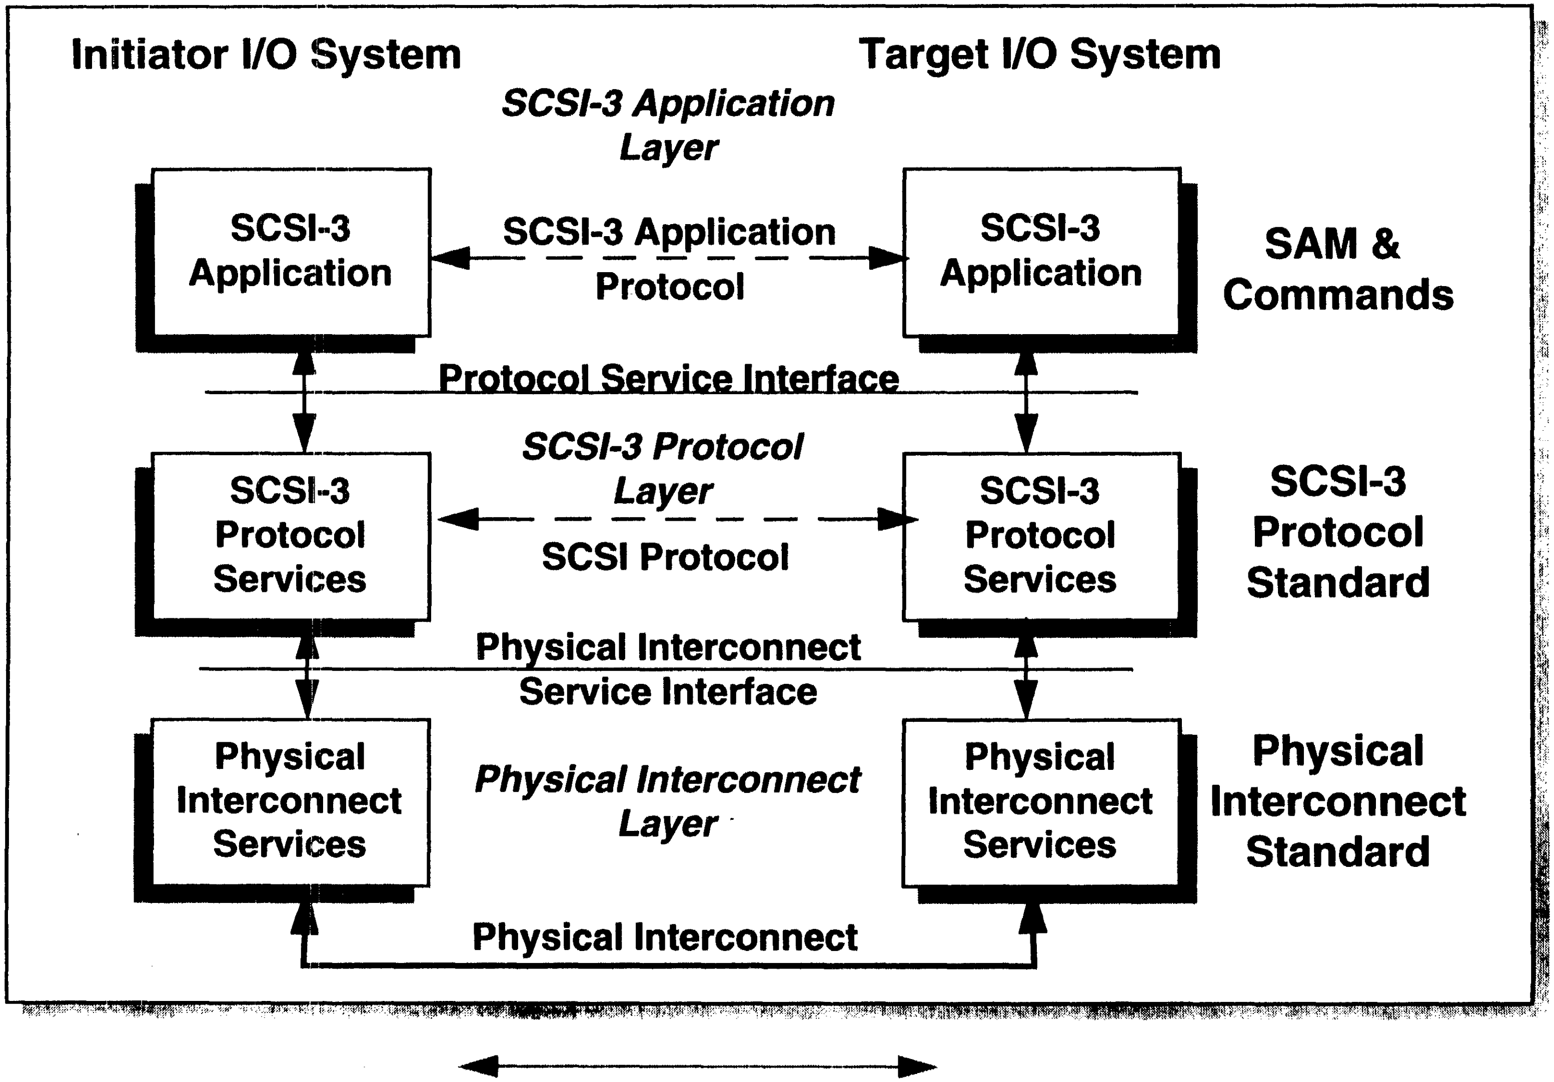

SCSI-3 Object Hierarchy

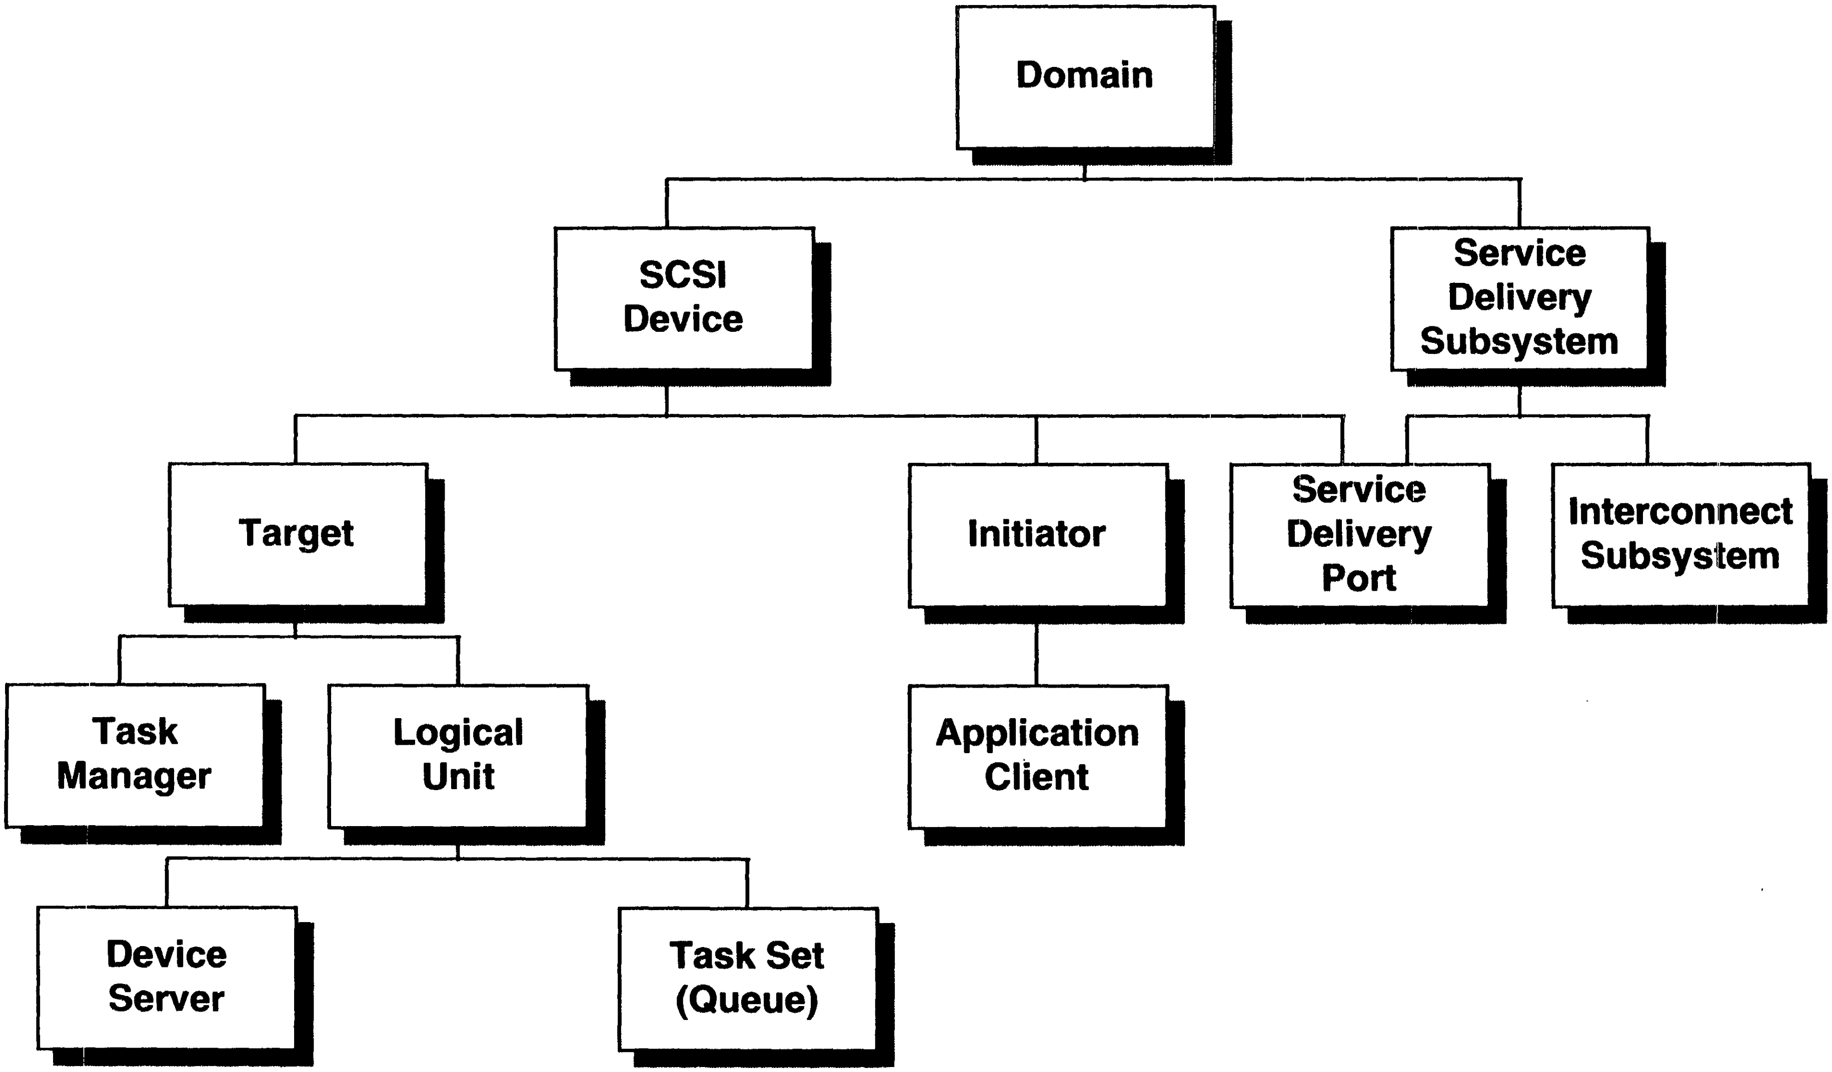

A Formal Model of Command Execution

- Service Response = Execute Command (Task Address, CDB, [Task Attribute], [Data-Out Buffer], [Command Byte Count], [Autosense Request], [Data-In Buffer], [Sense Data], Status)

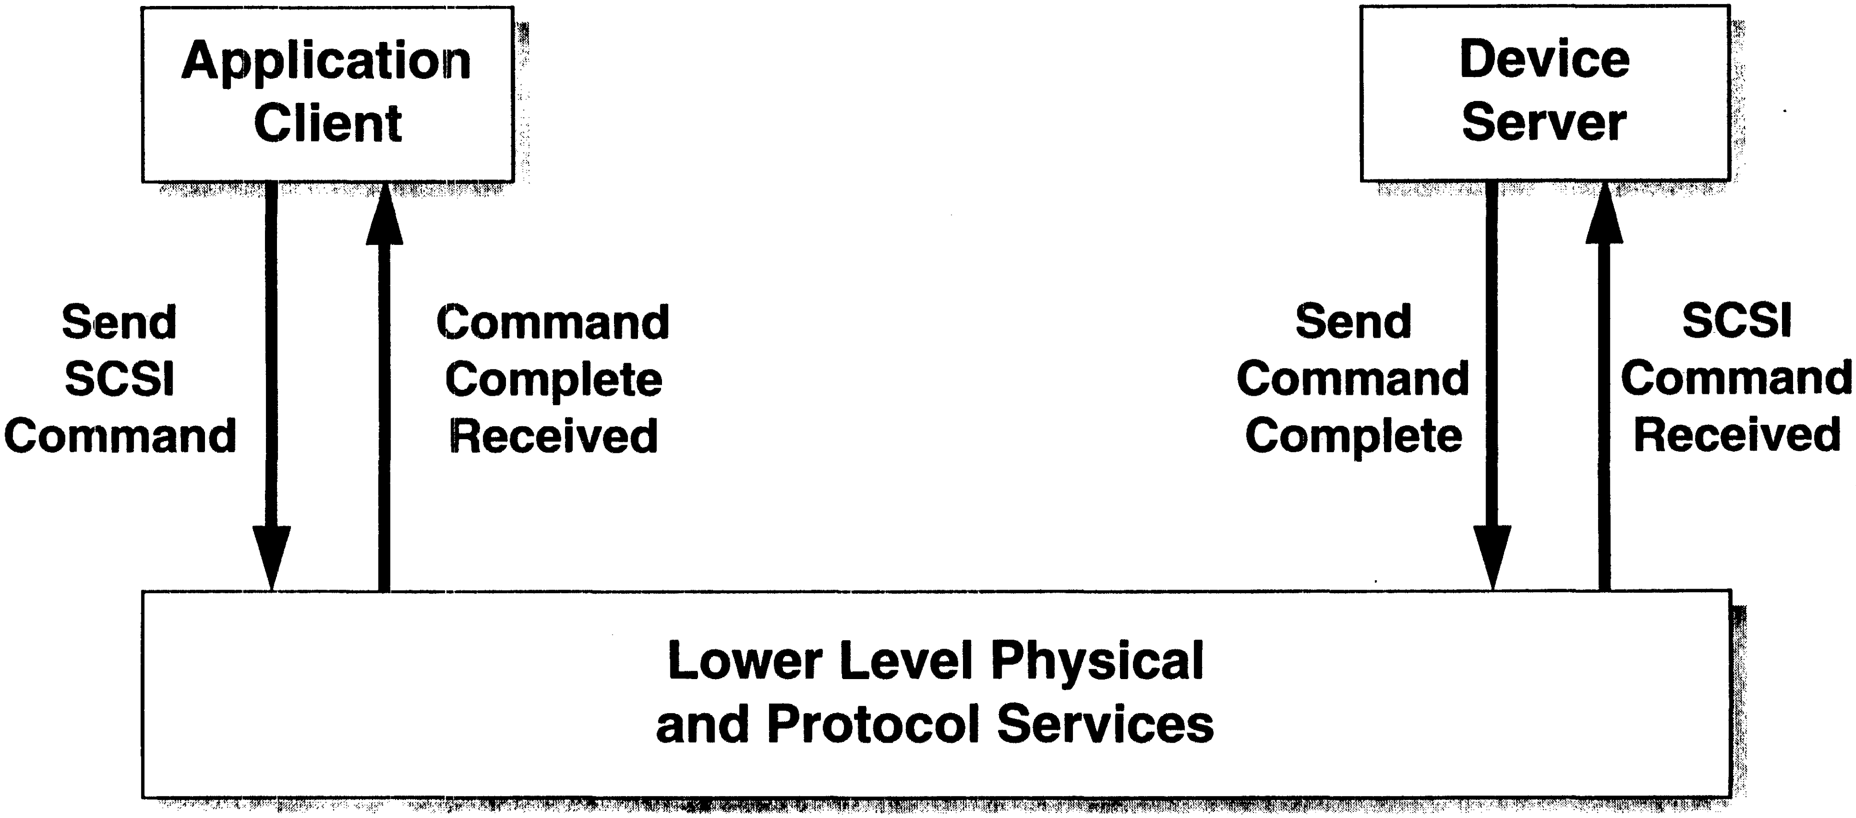

An Informal Model of Command Execution

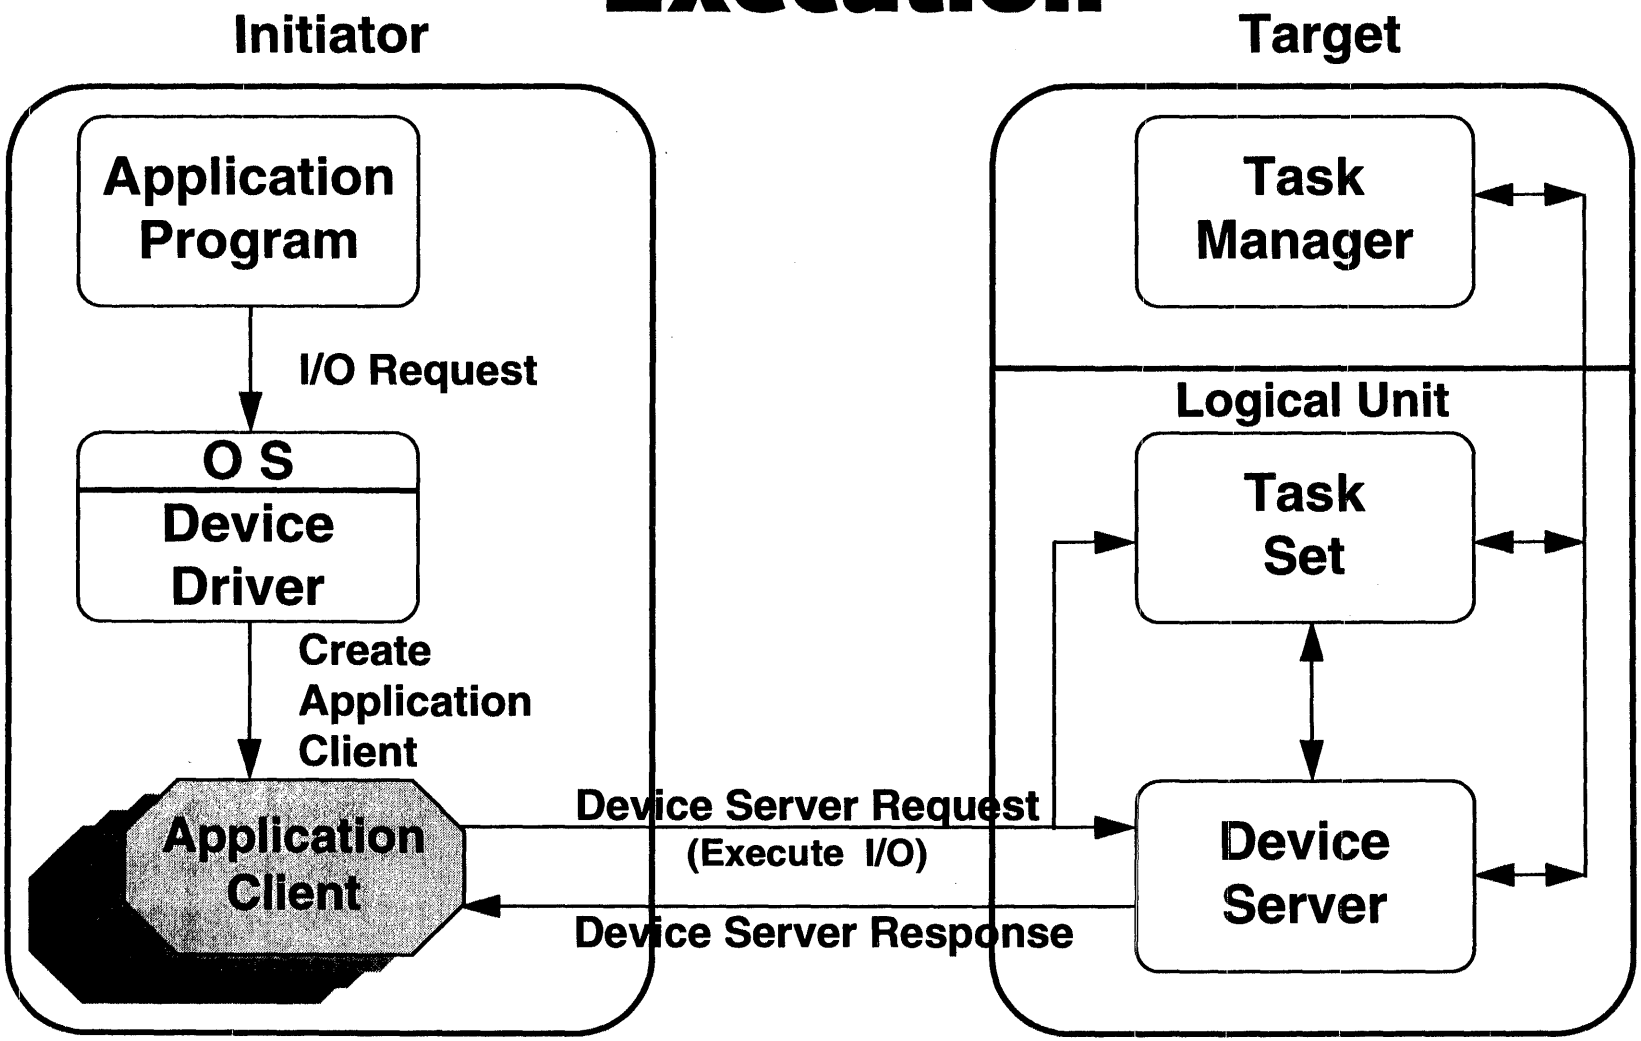

SCSI-3 Interlocked Protocol (SIP)

- Runs only on SCSI-3 Parallel Interface (SPI).
- Interlocked means each action is verified by a corresponding reaction by another device. (HANDSHAKE)
- Physically interlocked at byte level.
- Cable length sensitive.
- Half-duplex operation.
- Five new messages currently defined.
  - Logical Unit Reset
  - Target Transfer Disable
  - Continue Task
  - ACA Tag
  - Clear ACA

Changes to Standard INQUIRY Data (SCSI-3)

Bit / Byte	7	6	5	4	3	2	1	0
0	(Peripheral Qualifier)			(Peripheral Device Type)				
1	(RMB)	(Device-Type Modifier)						
2	(ISO Version)		(ECMA Version)			(ANSI-Approved Version)		
3	**AERC**	**TrmTsk**	**Norm ACA**	(Reserved)	(Response Data Format)			
4	(Additional Length)							
5	(Reserved)							
6	(Reserved)		**Multi PORT**	**Reserved**	**MChngr**	*Reserved to SIP*		
7	(RelAdr)	*Reserved to SIP*			(Linked)	*RsvSIP*	(CmdQue)	RsvSIP

Note: Fields in() above are unchanged from SCSI-2. Bytes 8-95 are unchanged

Changes to Standard INQUIRY Data (SCSI-3) (Continued)

AERC	Equivalent to SCSI-2 AEN bit
TrmTsk	Equivalent to SCSI-2 TrmIOP bit
NormACA	Device supports the NACA bit in the Control Byte of the CDB
Multiport	Identifies this target as a Multi-port device
MChngr	The addressed Logical Unit is attached to a medium changer. Only valid if RMB bit = 1.

SIP-Specific INQUIRY Data (SCSI-3)

Bit / Byte	7	6	5	4	3	2	1	0
6	Standard INQUIRY data					ACKQ REQQ	Addr32	Addr16
7	StdINQ	(Wbus32)	(Wbus16)	(Sync)	StdINQ	TranDis	StdINQ	(SftRe)

Note: Fields IN ( ) are unchanged from SCSI-2

ACKQREQQ Device supports Req/Ack in a Q cable (for 4 byte transfers)

Addr32 Device supports 32 SCSI address bits

Addr16 Device supports 16 SCSI address bits (if both are Addr32 and Addr16 = 0, the device supports only 8 address bits)

TranDis Device supports TARGET TRANSFER DISABLE and CONTINUE I/O PROCESS messages.

This page is intentionally left blank.

XOR Commands

SCSI-3 XOR Commands (Proposed) Highlights

- Allows a RAID 5 array to be implemented without a separate XOR engine in the array controller.
- Target application is a disk array with hardware RAID performance and **NO** special RAID hardware cost. Well suited to FC–AL and SSA disk configurations.
- Requires that each drive in the array have an XOR capability.
- Six new commands and two mode pages defined.

SCSI-3 XOR Commands (Proposed)

XDWRITE (16 byte CDB)	Write new data to the disk. In addition, exclusive-or the old data with the new and (optionally) transfer the XOR's data to another target.
XDWRITE (10 byte CDB)	Write new data to the disk. In addition, exclusive - or the old data with the new and save the result in the drive's buffer for retrieval by the host.
XDREAD (10 byte CDB)	Transfer XOR'd data from the target's buffer to the initiator.
XPWRITE (10 byte CDB)	Read old data from the disk. Exclusive-or it with new data from the initiator, then write the XOR'd data back to the disk.
REBUILD (16 byte CDB)	Reconstruct data on a disk by successfully reading the related data on the other drives in the array and XOR'ing it with the data on the disk being rebuilt.
REGENERATE (16 byte CDB)	Recreate lost data from a failed disk by reading and XOR'ing the related data from the other disks in the array. This is normally done to satisfy a read request from an initiator while waiting for a REBUILD to take place.

READ-MODIFY-WRITE Function
Multiple Interface Array

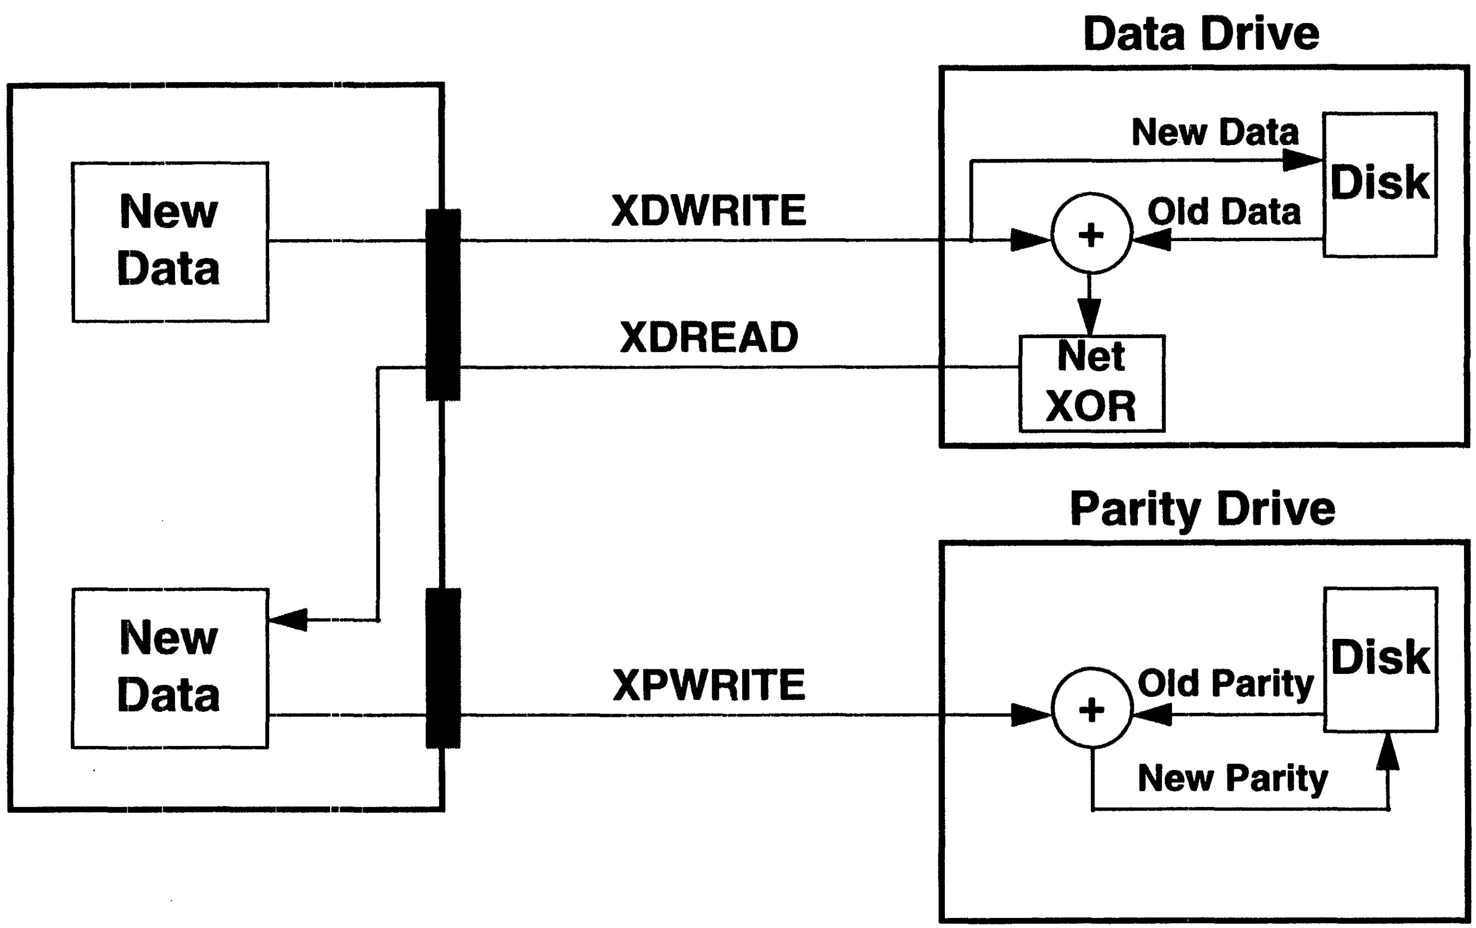

READ-MODIFY-WRITE Function

Single Interface Array

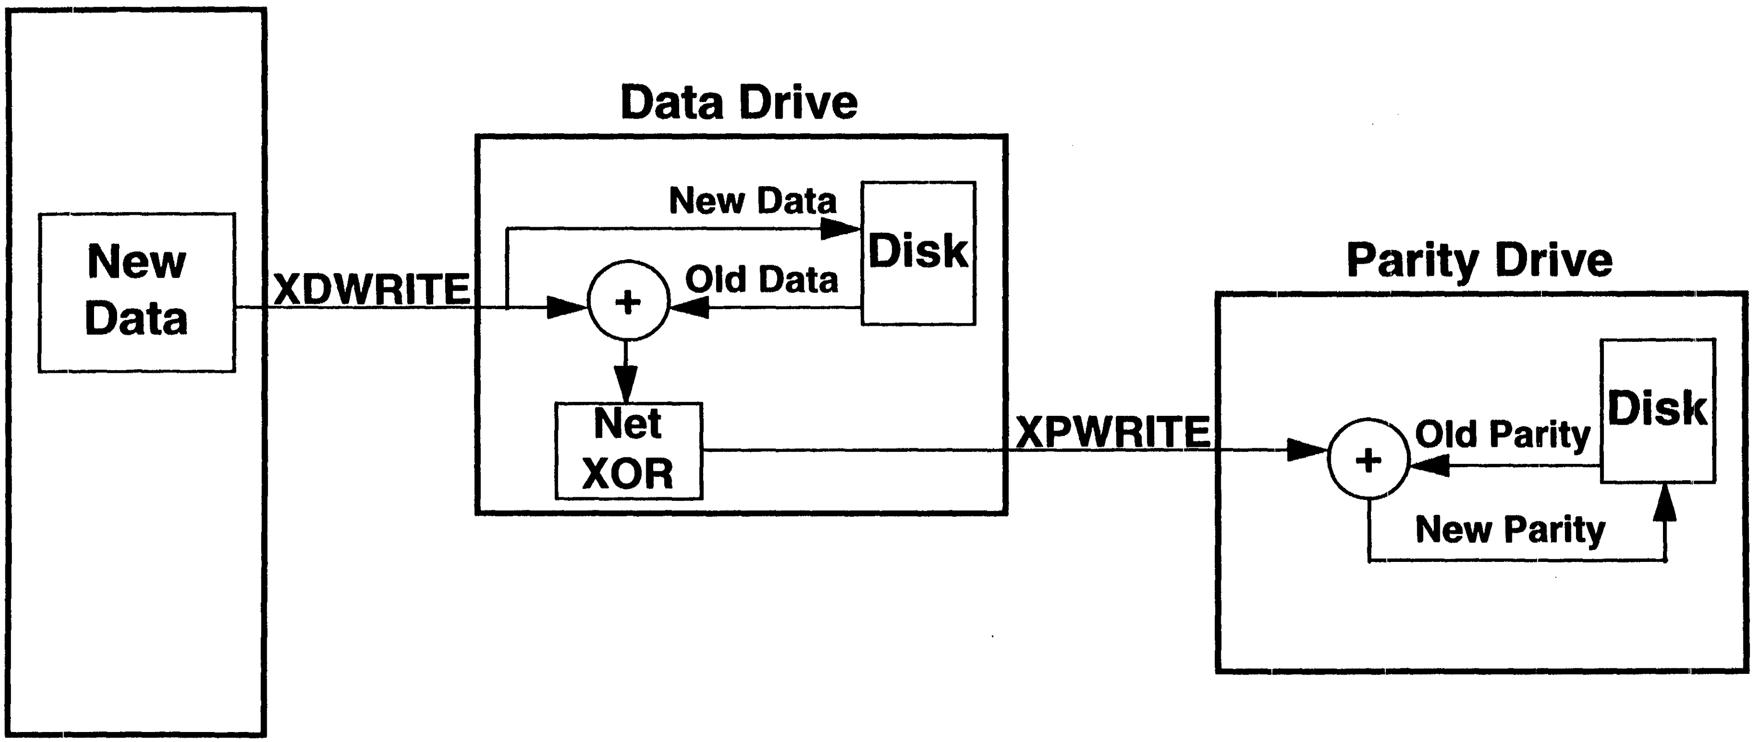

Rebuild Function
Single Interfacing

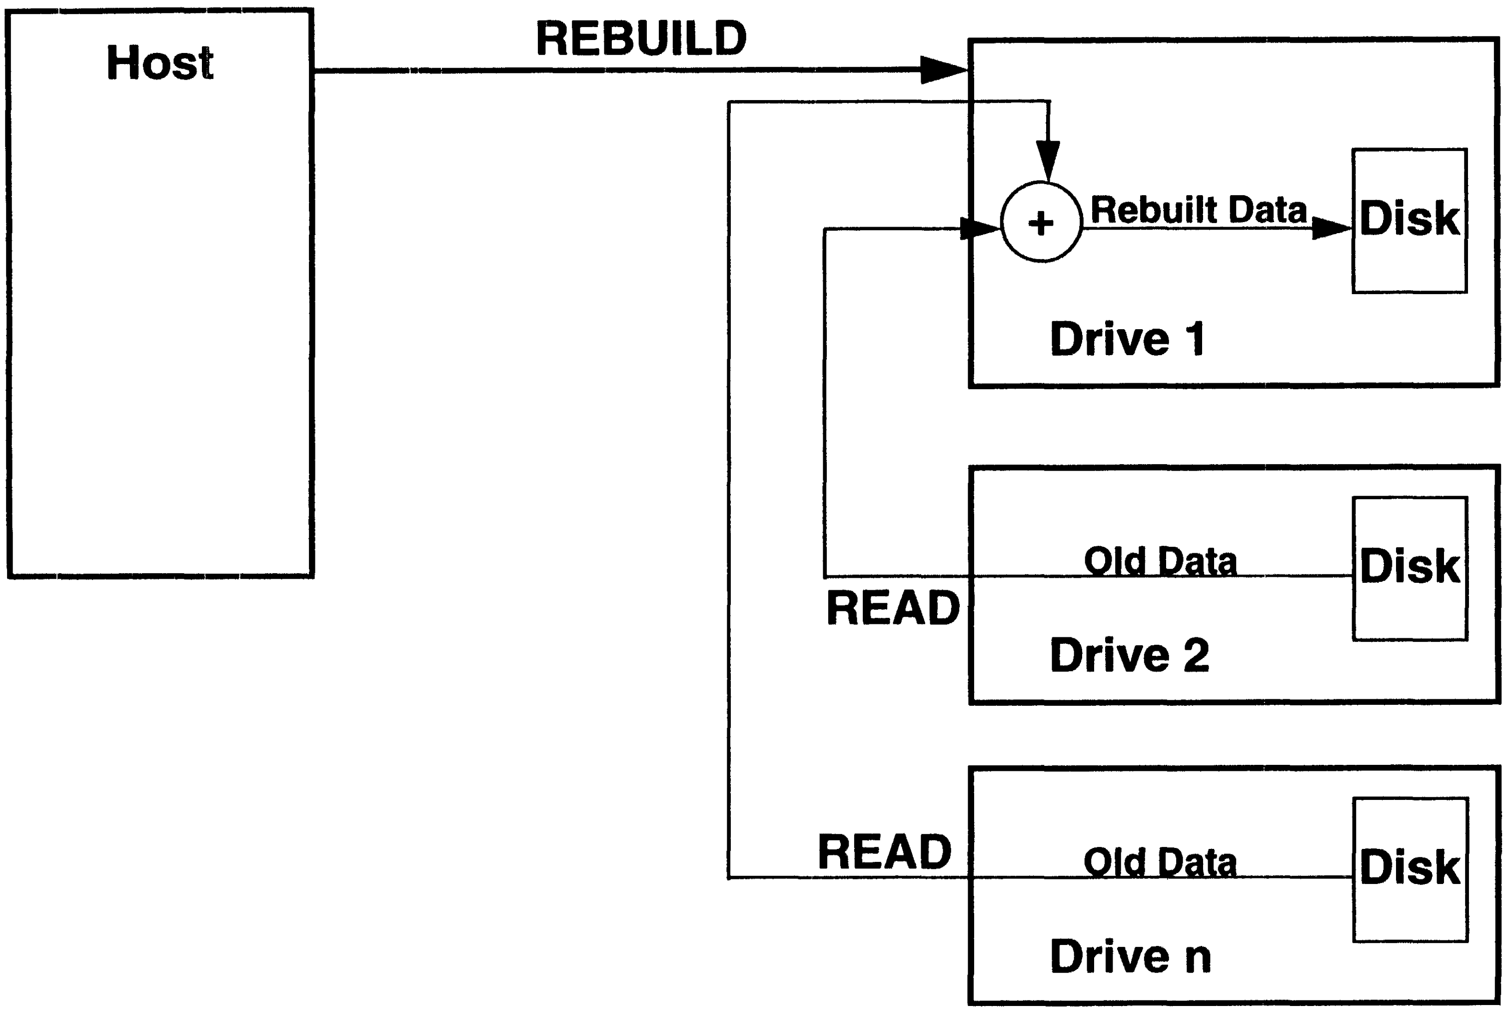

SCAM

SCSI Configuration Automatically (SCAM)

- In a SCAM system, device IDs are assigned dynamically following power-on or reset.
- A new bus phase, SCAM Selection, is defined.
- One device assumes the role of SCAM Initiator. It assigns IDs for all SCAM devices on the bus.
- New bus timings for Selection are defined.
- SCAM Selection uses a series of *transfer cycles* to isolate one device at a time for ID assignment.
- SCAM sequences are defined in the SCSI-3 Parallel Interface (SPI) document.

SCAM System Configuration Example

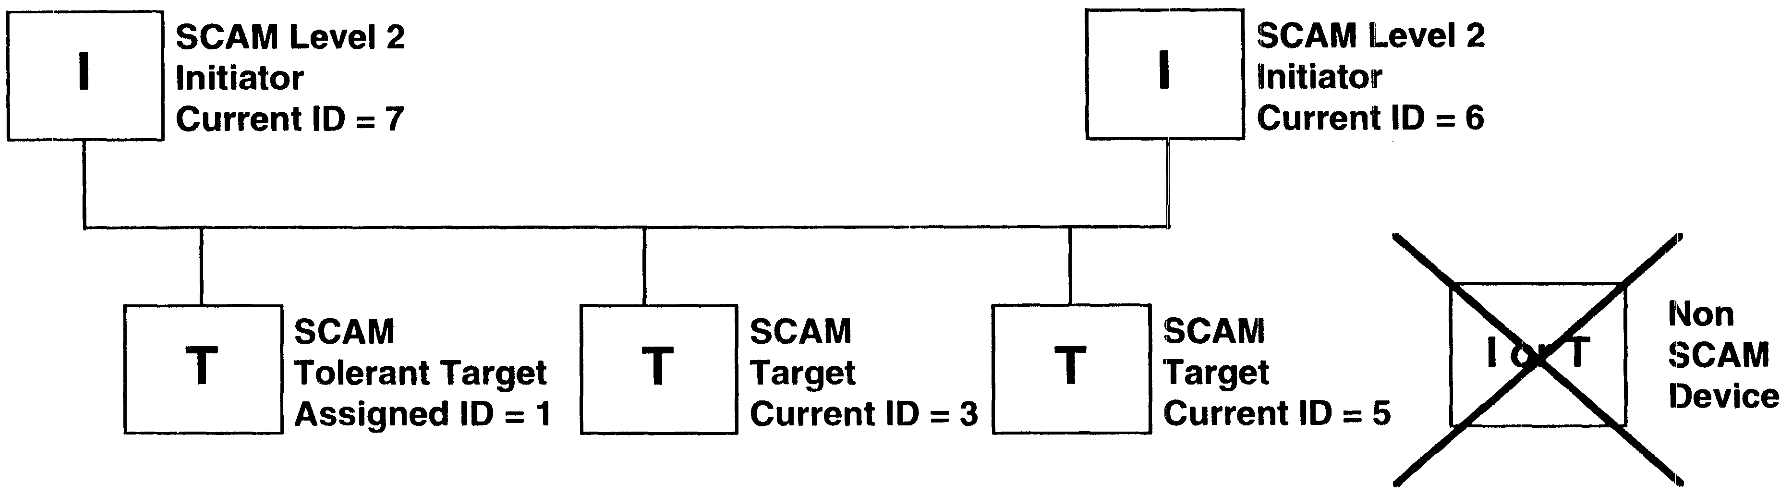

High Level Timing

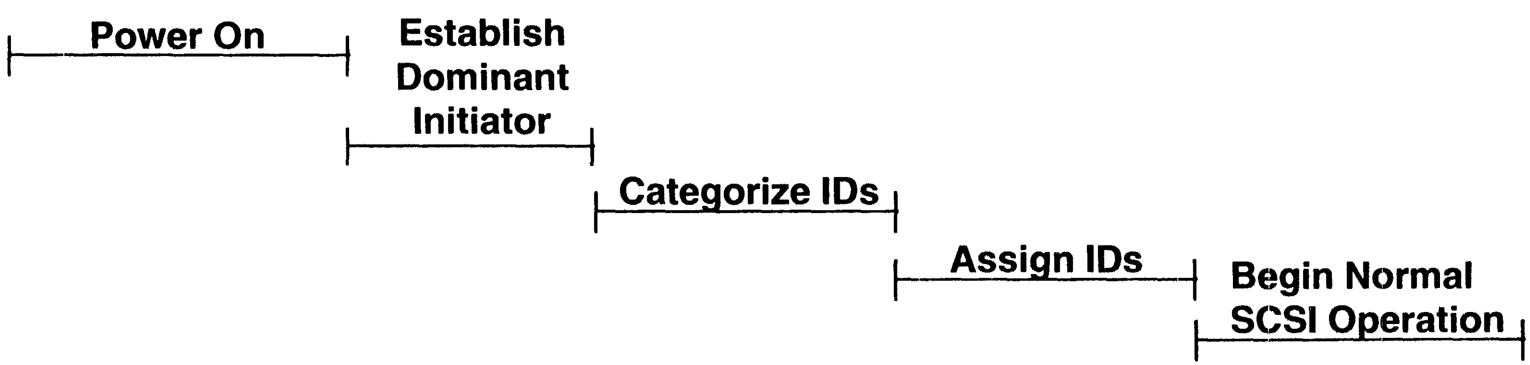

Initiating SCAM Selection

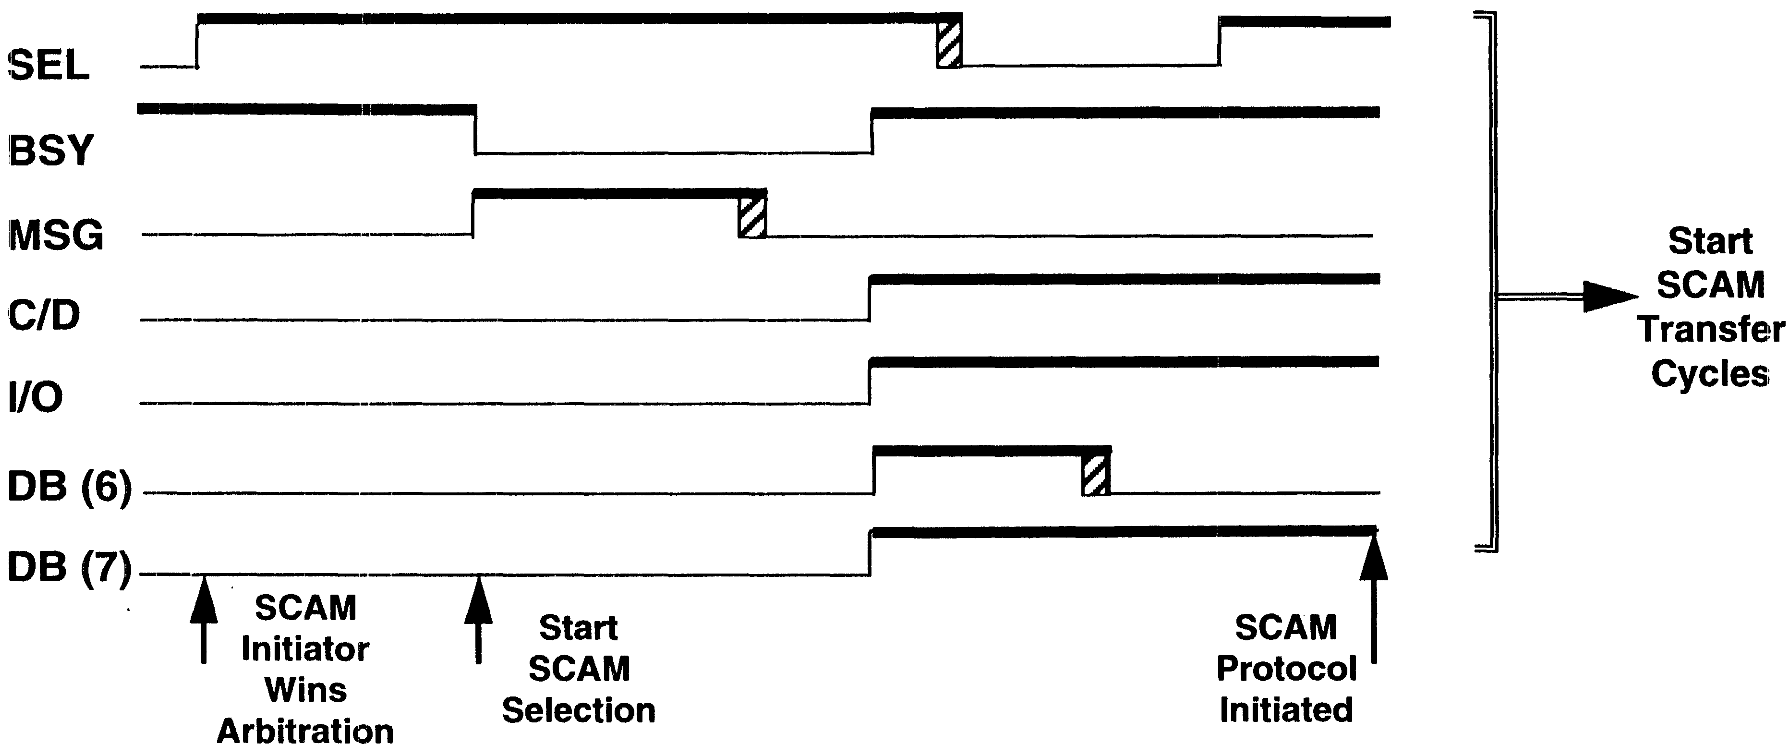

- The signals above must all be OR-tied during SCAM Selection. All participating devices drive the bus.
- SCAM Selection phase recognized **only** by devices without an assigned ID.

Typical SCAM Protocol Sequence

Assign One ID

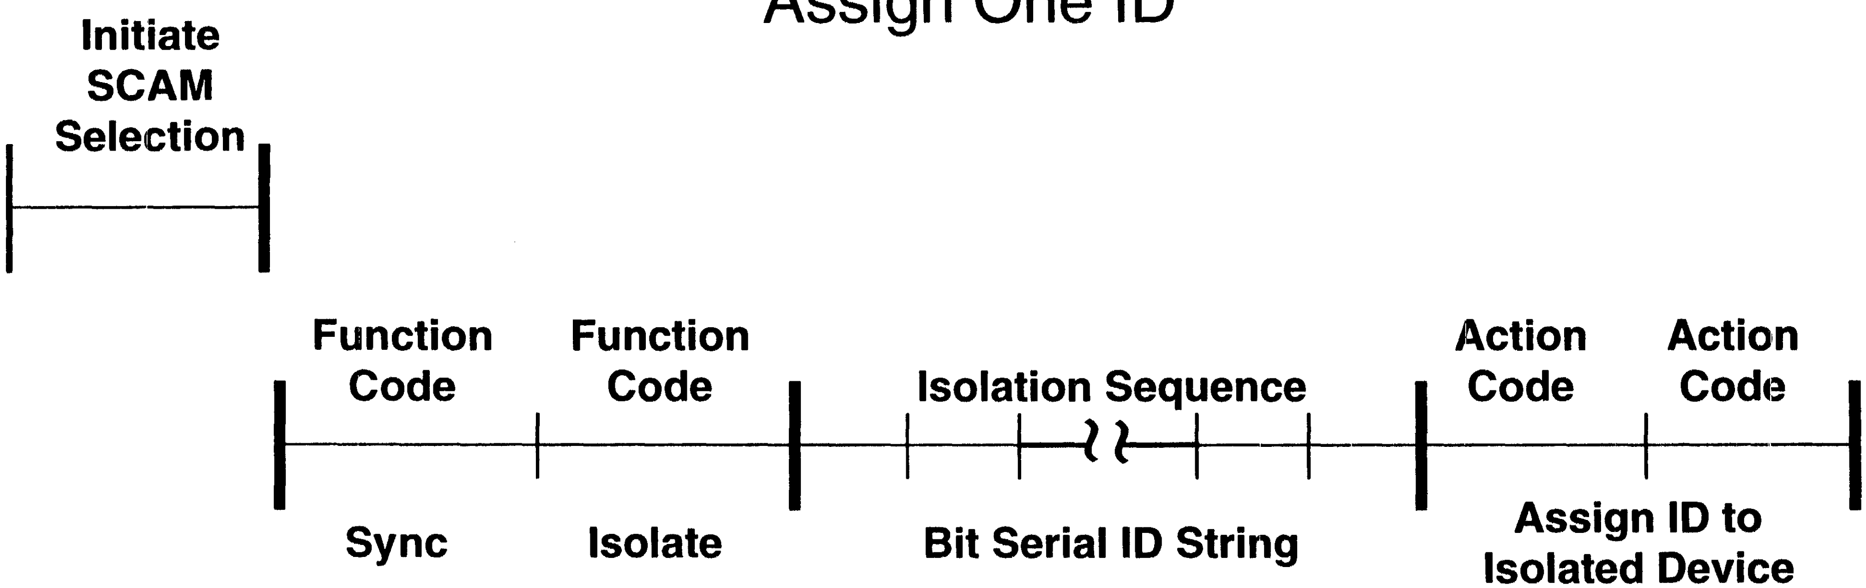

- During transfer cycles, DB (4-0) contain Function Code, Action Code, or serial data.
- Serial data encoding:
  - ID string bit = 1 if DB (1,0) = 10b
  - ID string bit = 0 if DB (1,0) = 01b

SCAM Transfer Cycles

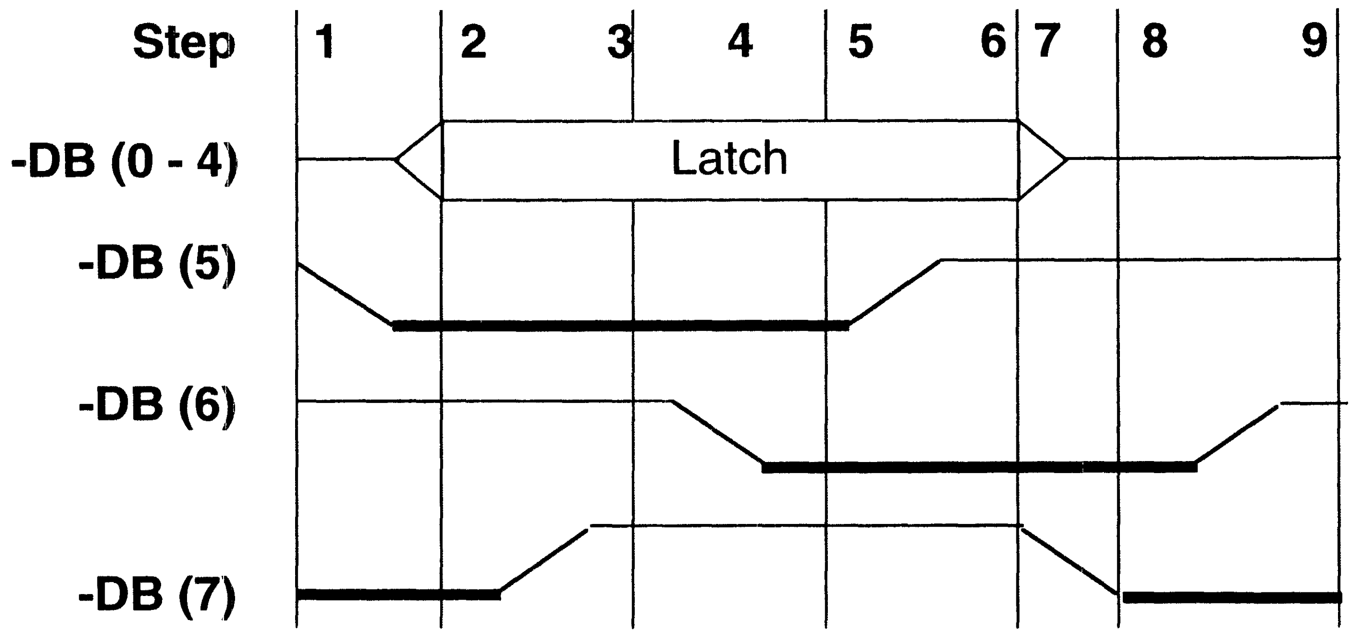

Note: Signals are shown asserted low.

SCAM Identification String

Table 1: Identification String

Byte \ Bit	7	6	5	4	3	2	1	0
0 – 1	(MSB) Type Code (LSB)							
2 – 9	(MSB) Vendor ID Code (LSB)							
10 – 31	(MSB) Vendor Unique Code (LSB)							

Table 2: Type Code

Byte \ Bit	7	6	5	4	3	2	1	0
0	Priority Code		Maximum ID Code		Reserved	ID Valid		SNA
1	Reserved			ID				

Categorizing IDs With Selection Time

Find SCAM tolerant devices (Hard IDs)

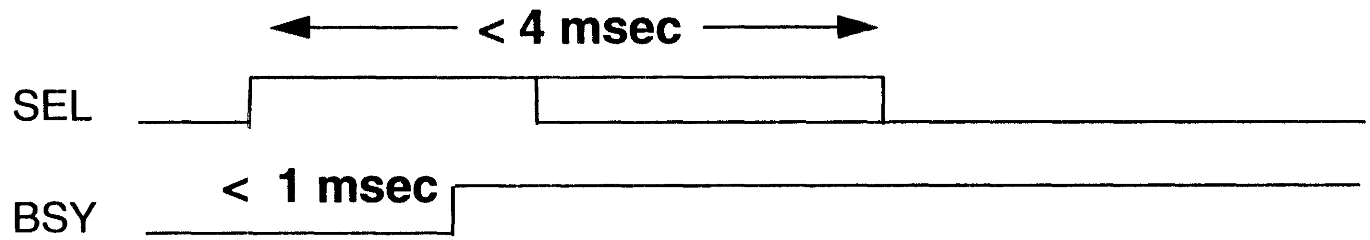

SCAM tolerant Devices must respond to selection in less than 1 msec.

Assign current IDs to SCAM device

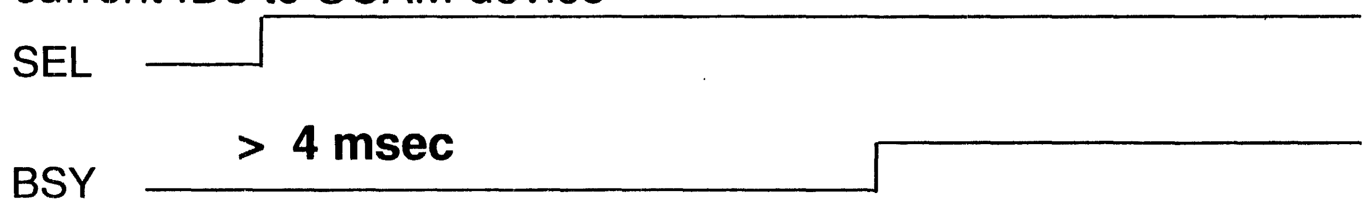

SCAM device without assigned IDS must not respond to selection in less than 4 msec.

Implementation Considerations

This page is intentionally left blank.

Host Adapters

Host Adapters

The basic function of a SCSI Host Adapter is to connect a _____________ with a _______________.

Two general classes of Host Adapter products

- Level 1: high volume, medium performance adapters.
  - Implemented in a single chips
  - Common functions handled by internal sequence
  - Complex or uncommon functions handled by host firmware
  - Typically has 16-byte FIFO for data
- Level 2: High Performance Adapters.
  - Internal RISC microcontroller replaces sequences
  - More functions handled by Host Adapter, less host firmware required
  - Larger data FIFO, 64 - 256 bytes

Host Adapters (Continued)

A Basic Host Adapter Example

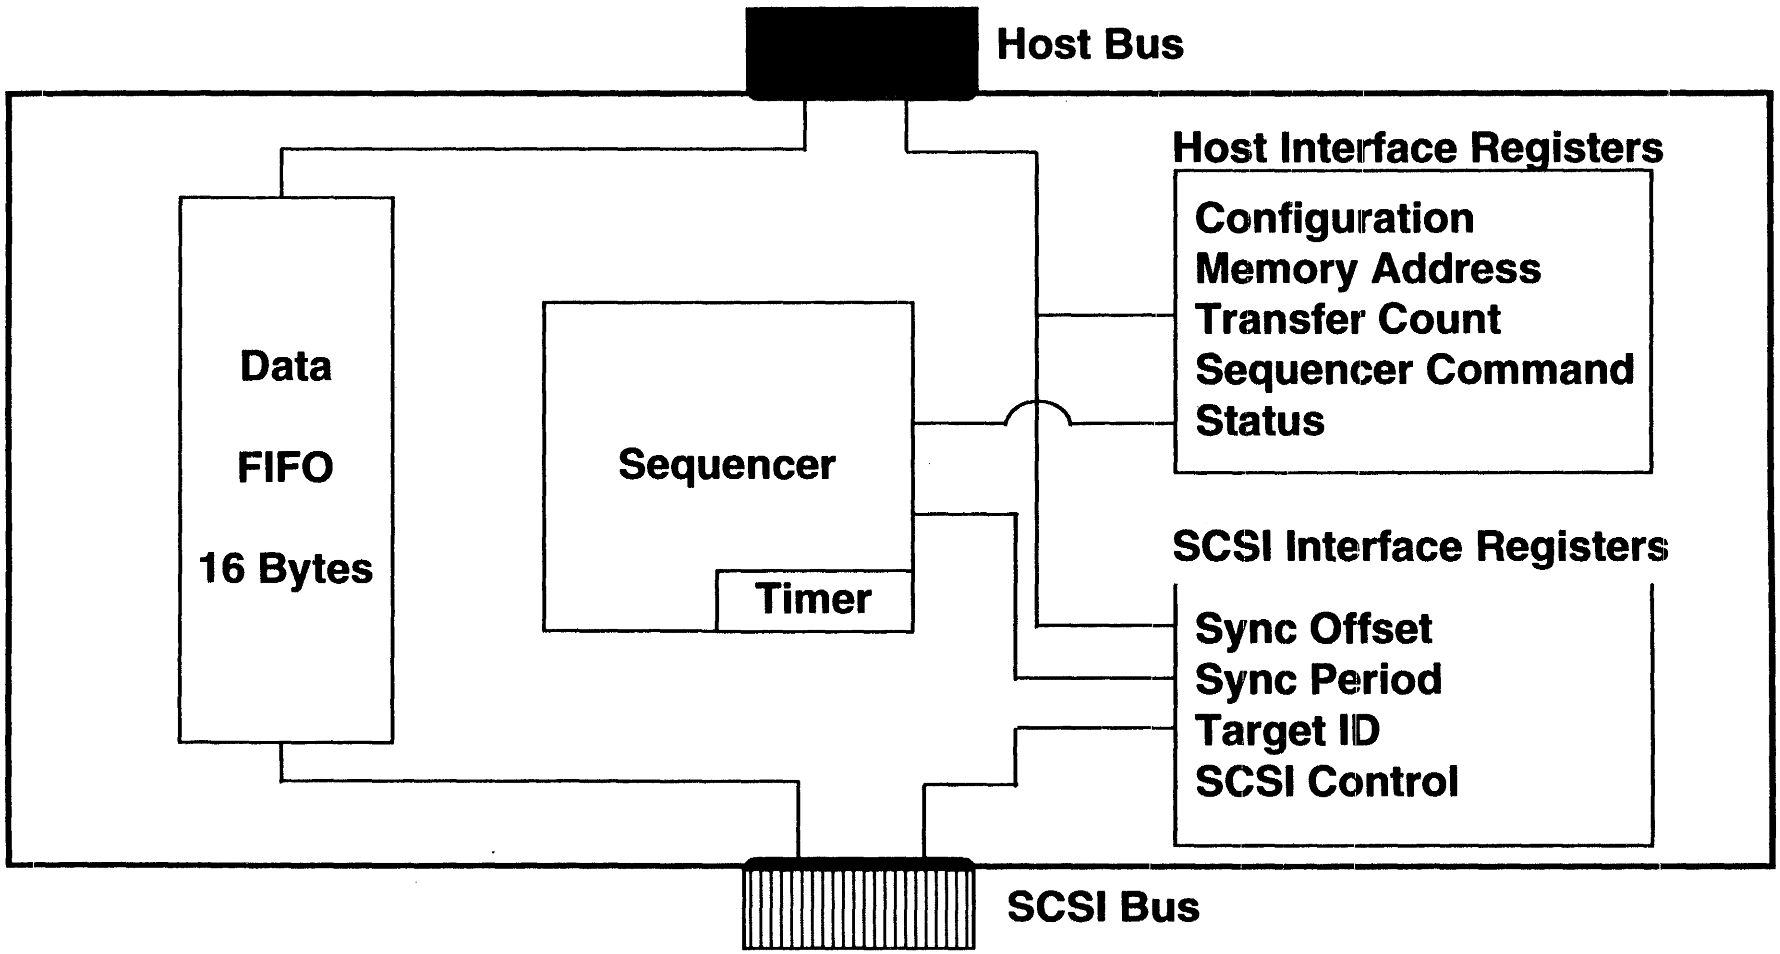

Host Adapters (Continued)

A High Performance Host Adapter Example

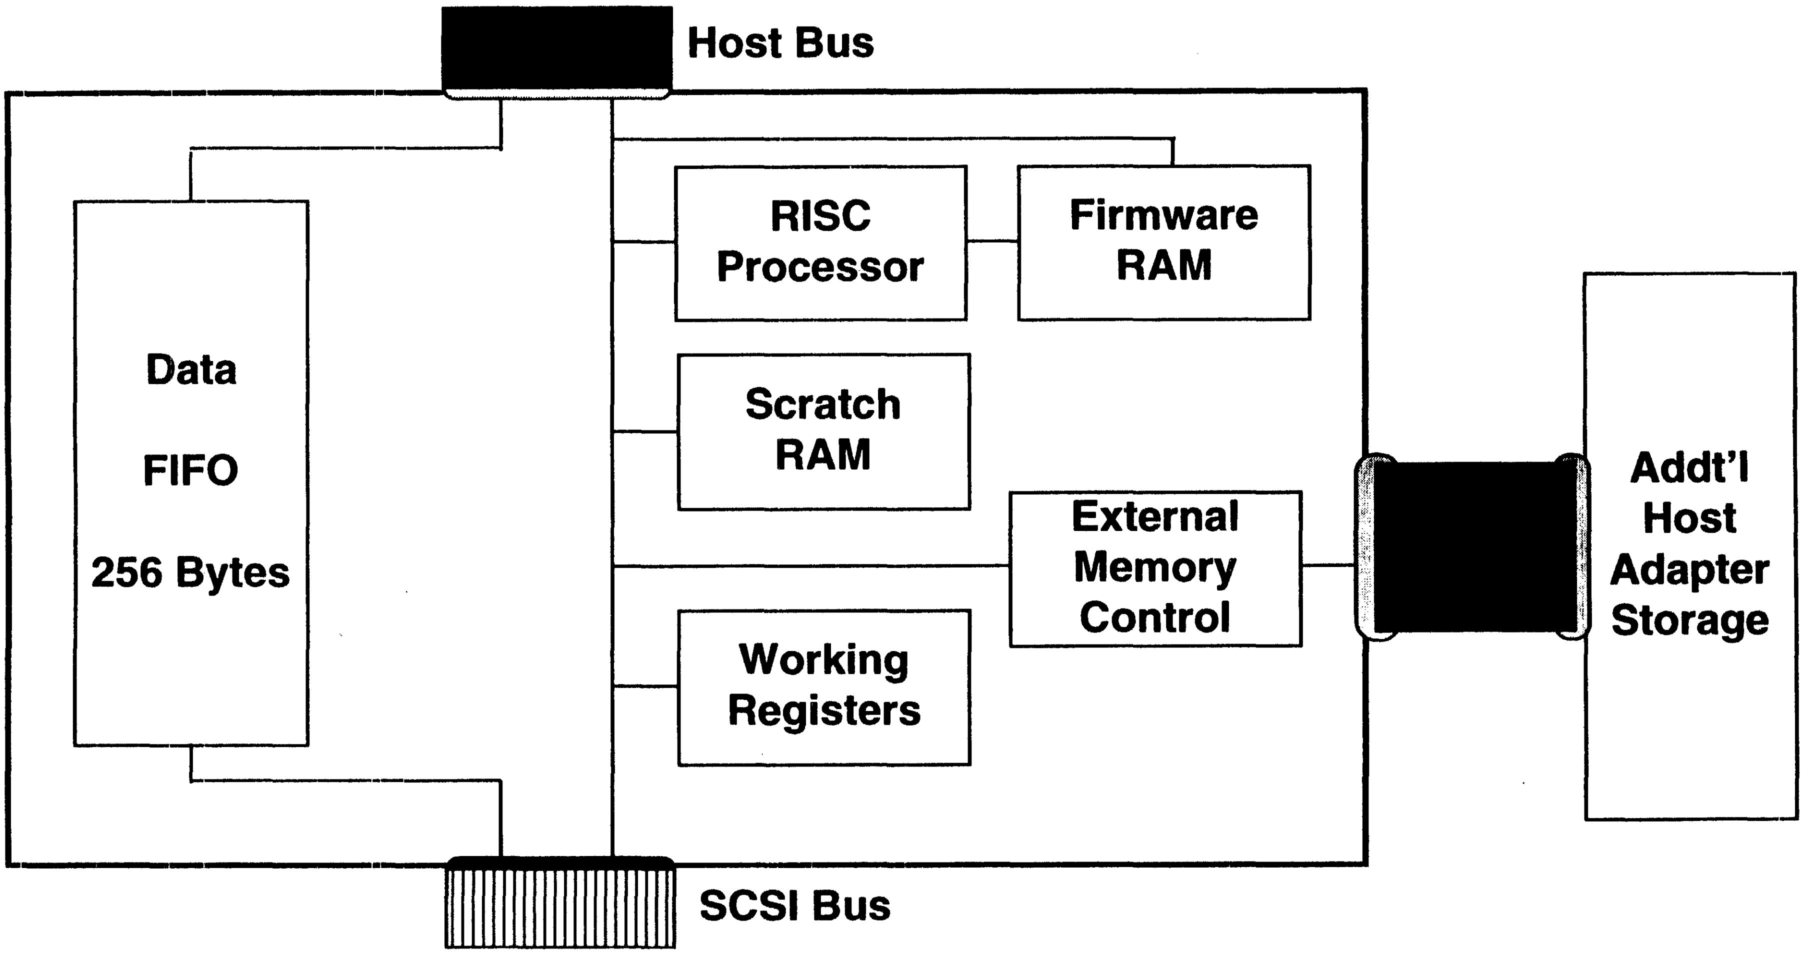

Target Controllers

Target Controller

Disk Drive Major Units

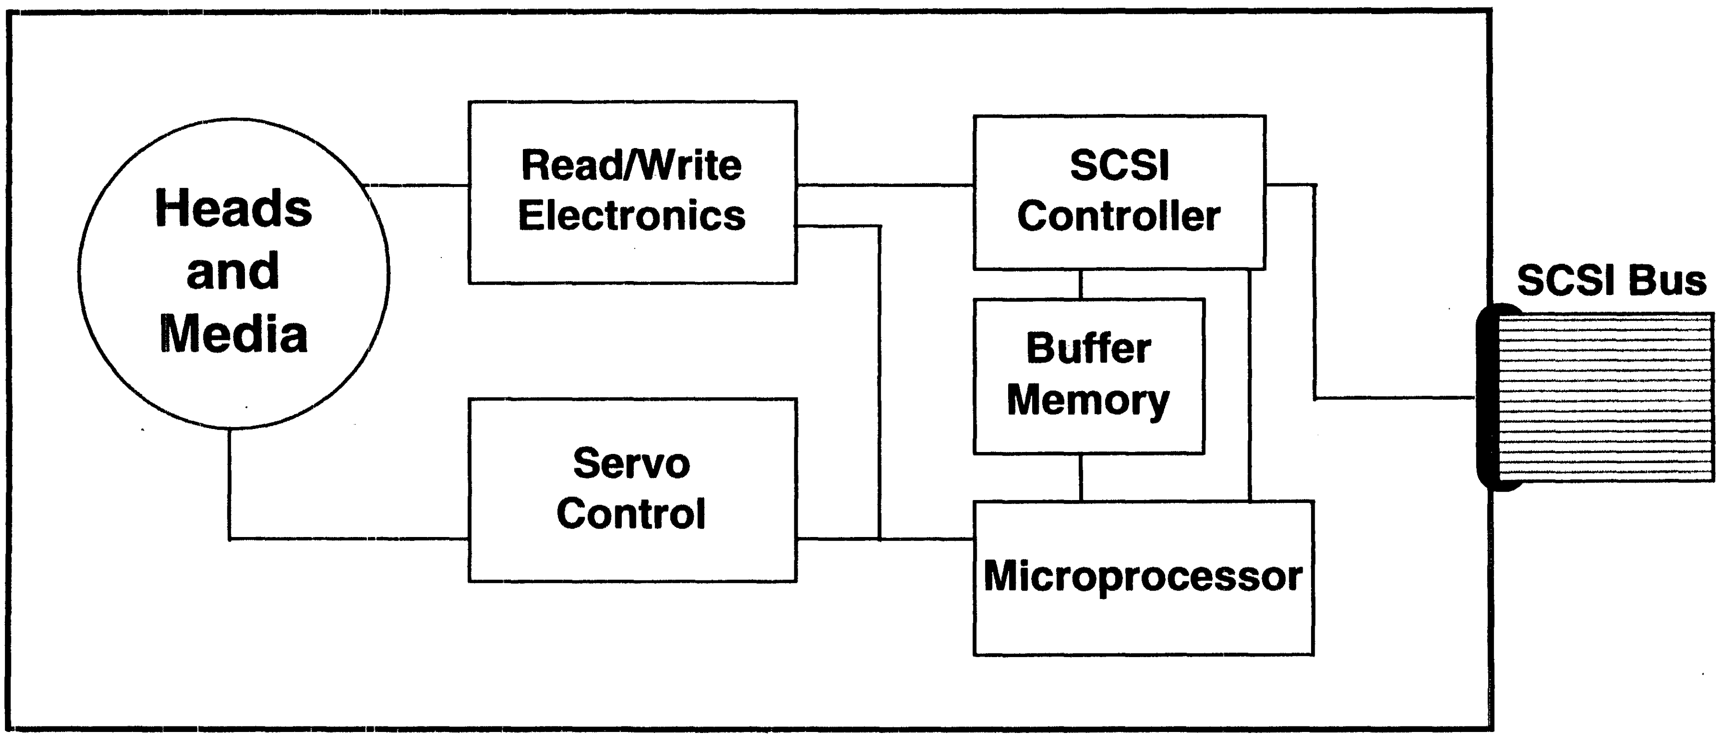

Target Controller (Continued)

Disk Drive Controller Example

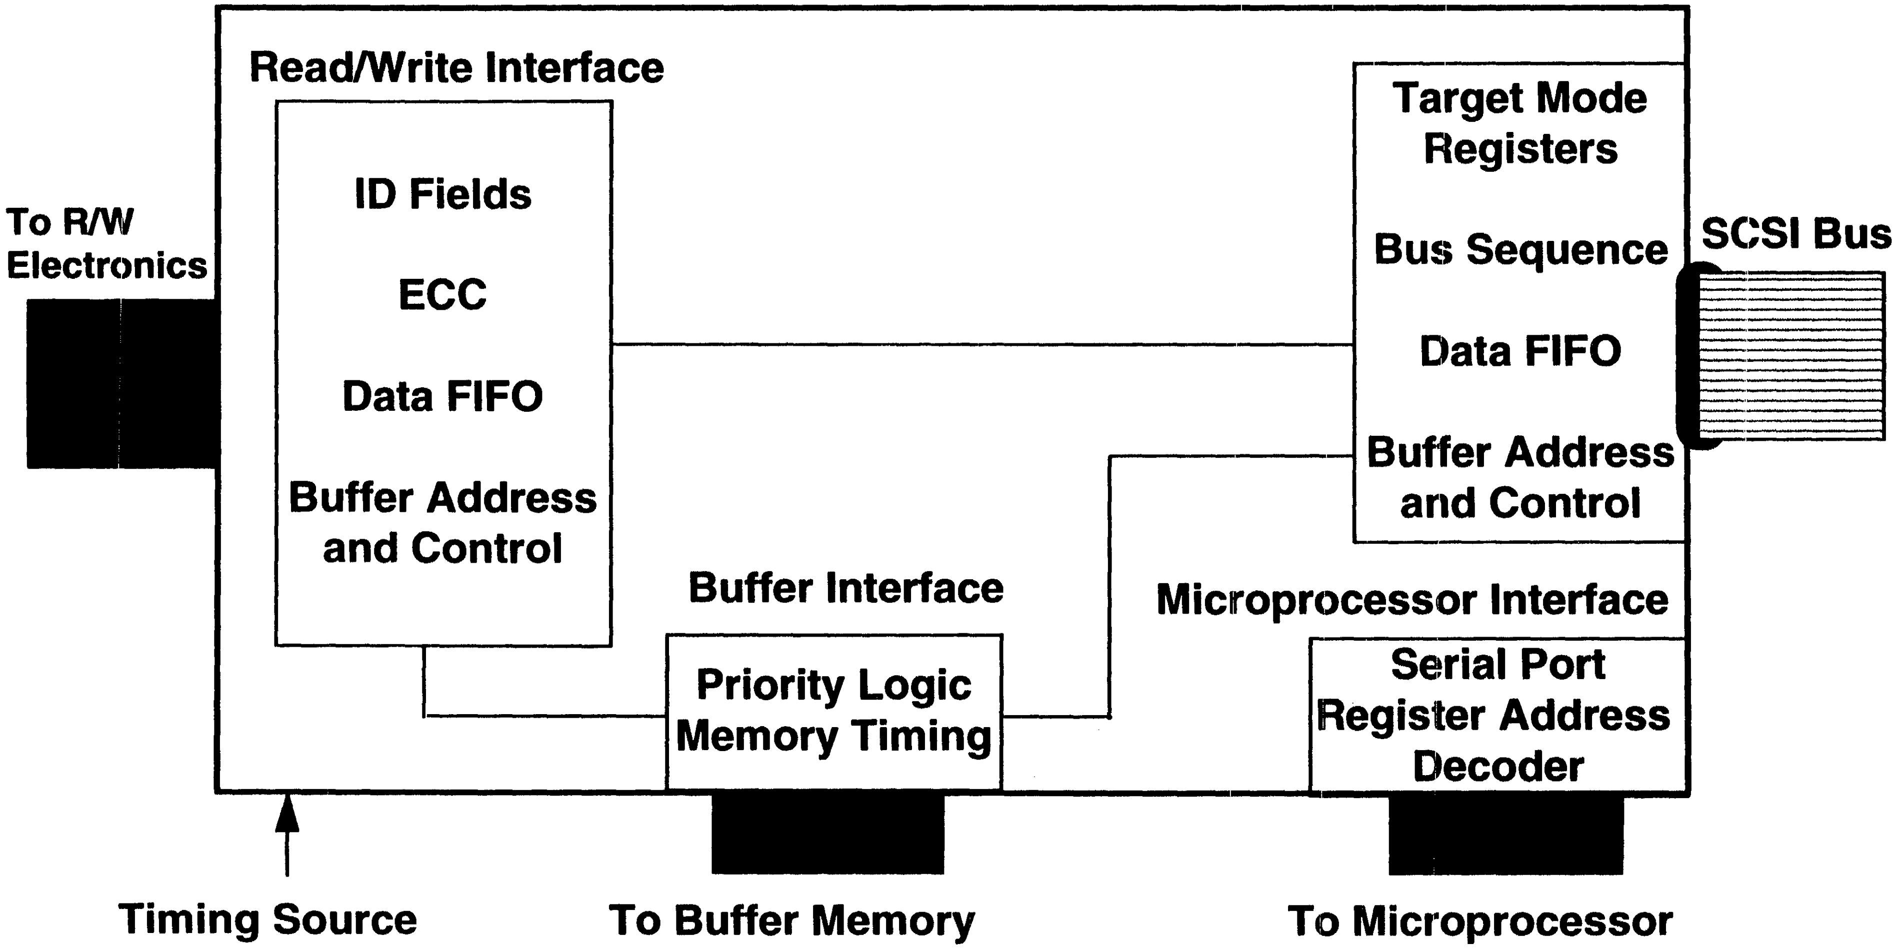

This page is intentionally left blank.

Host Software

Host Software

Layers of Function (Typical)

Application	Makes a request using file name, offset, and length
Operating System	Validates request Translates to device and LBA
Device Driver	Creates CDB Handles most errors
Host Adapter Driver	Sets up host adapter registers Handles pointers and common messages
Host Adapter Hardware	SCSI Bus Interface Transfers data to/from memory
SCSI Bus	

Host Software (Continued)

The ASPI Interface

- ASPI - Advanced SCSI Programming Interface
  - A standard for passing requests between Device Drivers (Peripheral Drivers) and Host Adapter Drivers.
  - Benefits:
    - An ASPI Device Driver can support peripherals attached to Host Adapters from different manufacturers.
    - An ASPI Host Adapter/Host Adapter Driver can work on any system with ASPI Device Drivers.

Host Software (Continued)

ASPI Layering

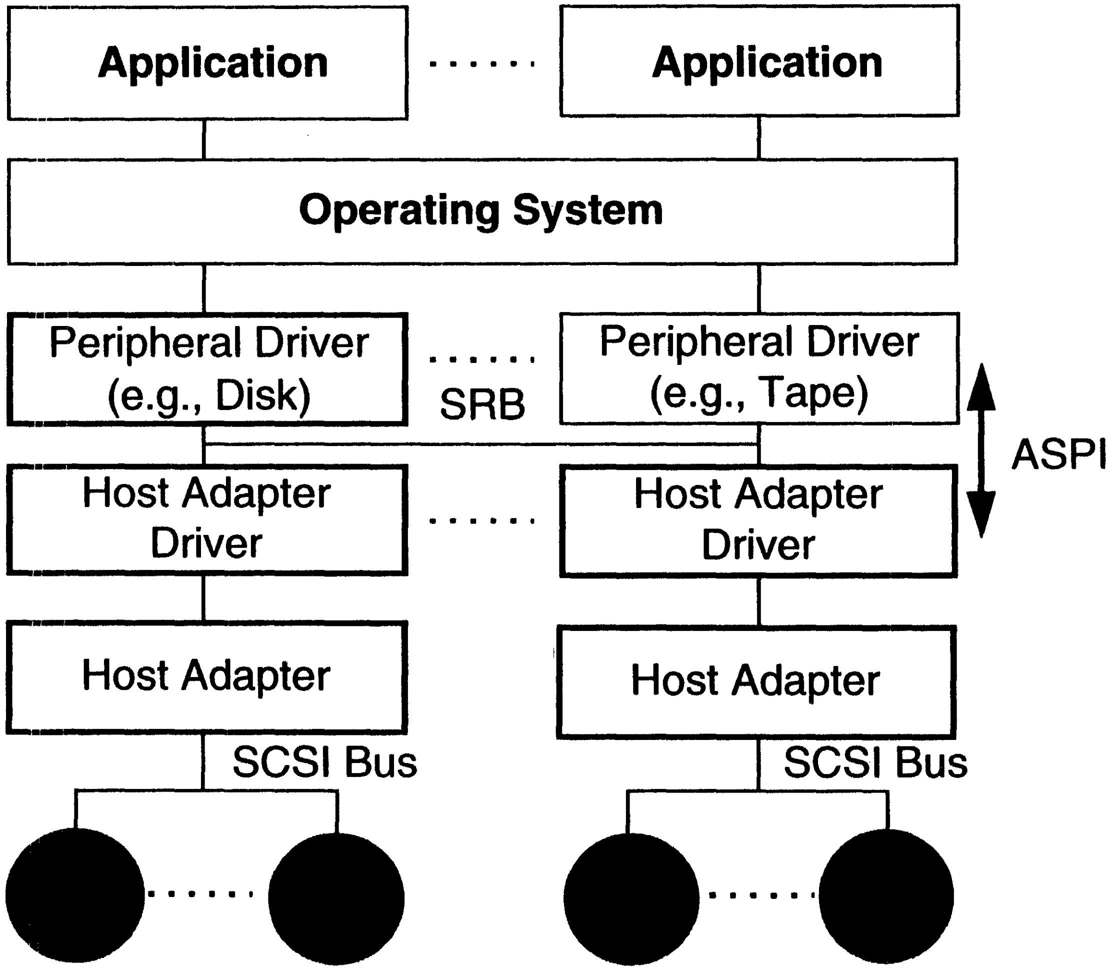

The SCSI Request Block (SRB)

The Heart of ASPI

SRB

Header

ASPI Command
Request Flags
Host Adapter Number

Status

Body

Target ID
LUN
Data Length
Data Pointer
CDB Length
Sense Length

HA Status

Target Status

Command Descriptor Block

Sense Data Buffer

The SCSI Request Block (SRB)

- ASPI Commands for Windows 3.x
  - Get ASPI Support Info
  - Host Adapter Inquiry
  - Get Device Type
  - Execute SCSI Command
  - Abort Request
  - Reset Device
- ASPI Status for Windows 3.x
  - Complete
  - Command Started
  - Aborted
  - Error
  - Invalid SRB
  - ASPI Is Busy

Note: For full details, refer to the ASPI Specifications for your operating system.