

A Proposal for Command Overlap and Command Queuing for ATA Devices

Pete McLean - Maxtor Corporation

1 Introduction

This document describes a method for supporting command overlap and command queuing for ATA devices. The goals for the development of this method were:

- Simplest possible protocols to implement.
- Ability to operate on the same cable with legacy devices (i.e., ATA-2 and ATAPI devices).
- Best possible future performance.

As a result, no effort was made to preserve existing hardware, firmware or software. Choices were made to provide the simplest, least overhead protocols.

A device implementing this method shall power-up in legacy mode and behave as defined in the ATA-2 specification today. The host may then determine the capabilities of the attached devices and enable the overlap and queuing functionality, if desired.

The protocols are designed to require the minimum possible overhead when both devices on the bus implement the overlap and queuing functionality and allow host DMA engines to provide high performance with minimum hardware.

2 Configurations Supported

Three configurations are supported by this proposal.

2.1 Legacy Configurations

Since the devices implementing this new functionality power up in legacy mode, they may be operated in legacy mode if desired. Therefore, a configuration may be as shown in Figure 1. In this configuration one or both of the devices may be capable of command overlap and queuing but a legacy host will operate them as shown. They operate just as ATA devices operate today where a command on one device is executed completely before a command can be issued to the other device.

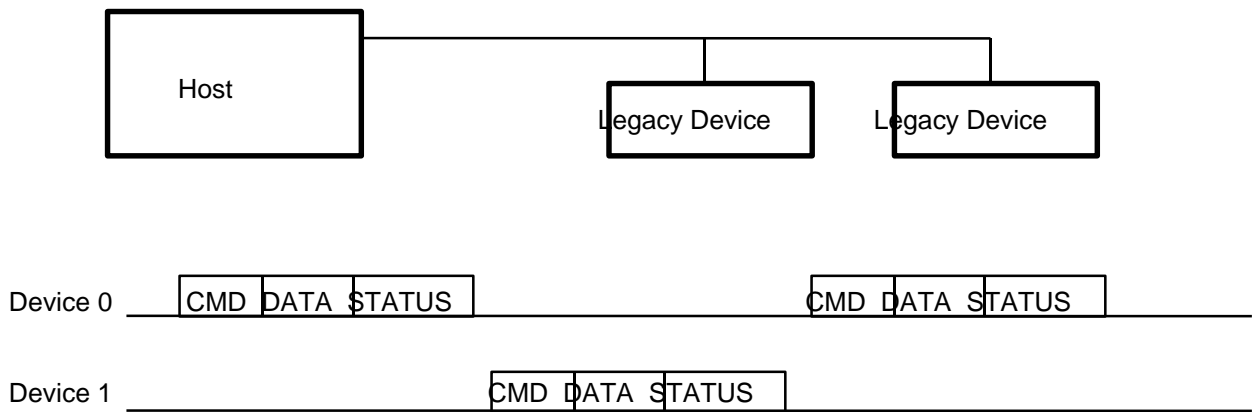


Figure 1 - Legacy Configuration

2.2 Mixed Configurations

The mixed configuration allows one device on the bus to behave as a legacy device while the other device is operating in overlapped/queuing mode as shown in Figure 2. This allows a host that is capable of supporting overlap/queuing to make use of the functionality even when a legacy device incapable of the functionality is present. Figure 2 shows how commands for the two devices may be interleaved when Device 0 is the Legacy Device and Device 1 is the Overlapped Device. The overlapped device may in fact be configured as either Device 0 or Device 1 in this configuration.

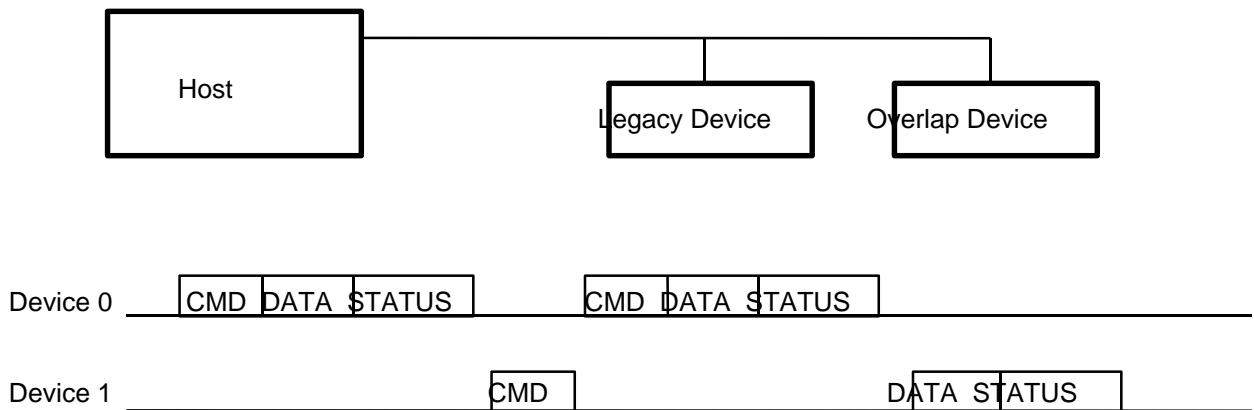


Figure 2 - Mixed Configuration

2.3 Overlapped Configuration

The overlapped configuration can occur when the host and both devices are capable of overlapped functionality as shown in Figure 3 . This proposal is designed to provide minimum overhead and maximum performance with this configuration. As shown by the first command to each device, the issuing of a command to a second device can be placed between the issuing of a command to the first device and the transfer of the first device data. If the device supports queuing as well as overlap, a subsequent commands may be issued to a device before a first command completes as shown by the second and third commands to device 0. In addition, if the data transfer is made in DMA mode, the data transfer may be completed in multiple data transfers instead of a single transfer as shown.

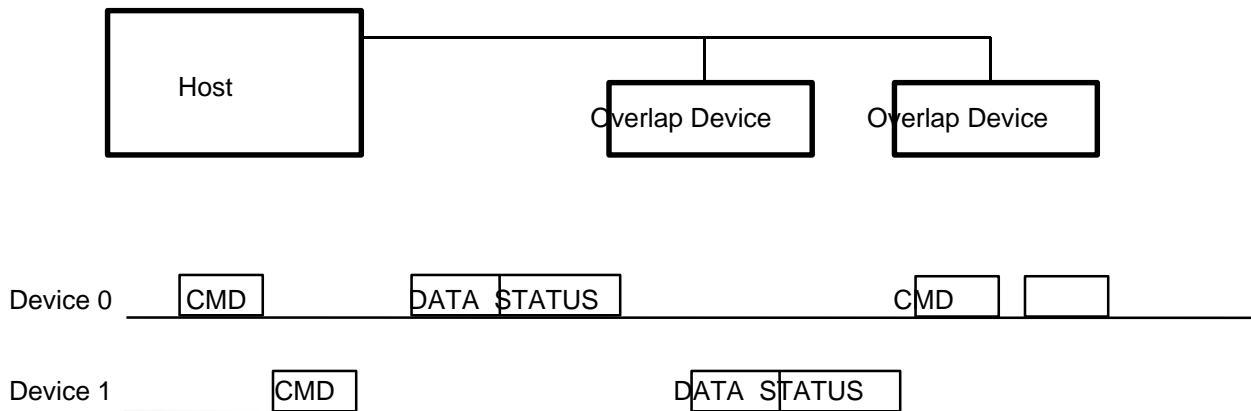


Figure 3 - Overlapped Configuration

3 Overlap/queuing Functional Issues

To implement command overlap and queuing a number of functional issues must be addressed. This clause describes each of these issues and defines solutions to them.

3.1 Register Ownership

Under the ATA specification today, register ownership is very clear. When a command is issued to a device, the device maintains a state with either BSY or DRQ asserted until the command is completed. During this time, the busy device owns the register set and the host may not write registers to either device without dire consequences. When the device has completed the command (neither BSY nor DRQ set), the host owns the registers and may write the registers for either device.

For devices implementing overlap/queuing the rules change slightly. When a command is issued to such a device BSY is asserted as in the legacy case. If the command is a non-data transfer command, BSY remains asserted until the command completes but if both devices support overlap, the host may select the other drive as soon as the command has been written. If the command is a data transfer command, the device sets BSY, inhibits writing of its command block registers, saves the contents of the applicable command block registers, and then clears BSY. If both devices support overlap, the host may select the other device as soon as the command has

been written. If the device supports queuing, the host may issue a second command as soon as BSY is cleared.

If the configuration is a Mixed Configuration, the host must obey the rules that each device expects. That is, if a command is issued to the legacy device, that legacy device own all registers until that command completes. The host may not access the other device. If the host issues a command to the overlapped device, the host may select the legacy device and issue it a command as soon as the overlapped device has cleared BSY.

3.2 Device Selection

Legacy devices today are selected by the host writing the DEV bit in the Drive/Head register. Since this register is owned by a device with BSY or DRQ set, the host is prohibited from selecting the other device when one device is in the process of executing a command.

In a legacy configuration, a device capable of command overlap/queuing is left in legacy mode and selected just as in legacy mode.

In a mixed configuration, a device capable of overlap/queuing will be placed in Overlap/DRVsel mode. In this mode, the device is selected just as in legacy mode. Since this device will clear its BSY bit as soon as the command parameters have been saved on a data transfer command, a host may then select the legacy device and execute a command while the overlap device is working on its data transfer command.

In an overlapped configuration where both devices are in overlap mode, both devices are placed in Overlap/CSsel mode. In this mode, selection is accomplished by a redefinition of the CS0- and CS1- signals. Table 1 shows both the legacy and overlapped/CSsel addressing.

Clause 4.2 describes the protocol for entering and leaving these modes.

Table 1 - Addressing for CS device selection

Address		Registers Selected	
CS0-	CS1-	Legacy Mode	Overlap/CSsel select mode
N	N	Not used	Device 0 - Control Block
N	A	Device 0/1 - Control Block	Device 1 - Control Block
A	N	Device 0/1 - Command Block	Device 0 - Command Block
A	A	Not used	Device 1 - Command Block
Key:			
A = signal asserted, N = signal negated			

3.3 Interrupt Reporting

Today, the INTRQ signal is not sharable so that a device can only assert INTRQ when it is selected by the DEV bit. When placed in Overlap/DEVsel mode, a new device implementing this function will continue to use the INTRQ signal as in legacy mode only asserting the signal when selected. When in Overlap/CSsel mode, the INTRQ signal shall be a “wire-or” signal. In this way, either device may assert an interrupt at any time. (Editor’s note: Since INTRQ is currently defined as a high asserted signal, we need to discuss how we want to do the “wire-or”.)

In addition, a new register is defined, the Interrupt register. This register is at address 5h in the Control Block and is read only. The contents of the register are described in Table 2. When this register is read, each device asserts only the two bits indicating its interrupt state. Therefore, an intelligent DMA engine can determine which device(s) is interrupting and the reason for the interrupt with the read of one register. This register should only be used when two overlapped devices exist.

Table 2 - Interrupt register

7	6	5	4	3	2	1	0
r	r	r	r	CC1	TR1	CC0	TR0

where:

TR0 indicates Device 0 has an interrupt pending to transfer data.

CC0 indicates Device 0 has an interrupt pending because a command has completed.

TR1 indicates Device 1 has an interrupt pending to transfer data.

CC1 indicates Device 1 has an interrupt pending because a command has completed.

3.4 Command Queue Tags

The Identify Device info returned by a device capable of overlap/queuing operation will contain a four bit queue depth field. If this field is 0h, the device does not support queuing. If the field contains any other value, this value plus one indicates the depth of queue supported up to a maximum queue depth of sixteen. Any tag issued by the host with a tag number greater than this value will cause the command to be aborted.

To support queue tags, all registers are defined as being sixteen bit registers just as the Data register is today. The upper byte of the Command, Status, and Alternate Status registers are as shown in Table 3. The upper byte of the remaining registers are reserved and shall be written as all zeros and ignored on read. In particular, the upper byte of the Cylinder Low and Cylinder High registers is reserved for future address space expansion.

When issuing a data transfer command to a device in overlapped/queuing mode, the upper byte of the Command register will indicate that this is an overlapped command and provide the queue tag number for this command.

When a queued device interrupts for service for a queued command, the tag for the command requiring service will be contained in the upper byte of the Alternate Status and Status register.

All queued commands shall be simple queued commands, no priority is implied.

Table 3 - Upper byte definitions

Register	15	14	13	12	11	10	9	8
Command	OVL	r	r	r	TAG(3:0)			
Status or Alternate Status	CC	r	r	r	TAG(3:0)			

Where:

OVL = one indicates an overlapped command.

TAG(3:0) = Queue tag number for command.

CC = one indicates command complete.

3.5 Host Implications

3.5.1 Hosts without DMA engines

With both devices in legacy mode, hosts function just as they have in the past.

With mixed mode devices, device selection and interrupts behave just as they have in the past. Therefore, the host will have to select the overlap device via the DRV bit to determine when an interrupt has occurred.

When both devices are placed into Overlap/CSsel mode, the host adapter must have a register that allows device selection. The writing of this register shall cause the CS1- signal to select the desired device until changed by the host software. The remaining address bits, i.e., CS0- and DA[2:0], are decoded from the host bus address just as with legacy devices. The interrupts from the two devices will occur when a device is not selected and the host will have to read the Interrupt register, or the Alternate Status or Status register of each device to determine which device interrupted and what the cause was. Note that with today's 3F6/376h host addressing only an intelligent DMA engine may be capable of reading the Interrupt register. If the host is unable to read the Interrupt register, the status of each device will have to be read to determine the cause of the interrupt.

3.5.2 Hosts with DMA engines

Command overlap and command queuing make the use of an intelligent DMA engine in the host an attractive added feature to improve performance particularly with a multithreaded operating system. However, with such an intelligent DMA engine, the accesses to the devices by the DMA engine and the host device driver must be synchronized. That is, the device driver may not access the devices while the DMA engine is in the process of transferring data to/from a device, and the DMA engine may not transfer data while the host device driver is accessing registers to issue a command or check status.

The following is an example of how this synchronization may be accomplished. The DMA engine will have a register accessible by the host. The ATAREQ bit in the register will be writable by the host. When the host device driver wishes to access a device to issue a command or read status, the host will set the ATAREQ bit to one. A register bit must exist to allow the device driver to select a device.

If the DMA engine is not currently transferring data to/from a device, it will set the ATAACK bit to one. When both the ATAREQ and ATAACK bits are one, the host device driver may access a device and the DMA engine shall not access a device.

If the DMA engine is in the process of transferring data to/from a device when the ATAREQ bit is set, it will complete the transfer and then set the ATAACK bit to one.

When the host device driver has completed the required device accesses, it will set the ATAREQ bit to zero. The DMA engine will immediately set the ATAACK bit to zero, and the DMA engine may now access a device.

When the ATAREQ and ATAACK bits are both asserted, the host adapter must pass the register value of CS1- selected by the device driver and the remainder of the address, i.e., CS0- and DA[2:0] as decoded from host bus address to the devices. When the ATAREQ and ATAACK bits are both cleared, the DMA engine must assert and latch CS1- to select the device to execute a DMA transfer. Timing for the assertion and negation for CS1- in this case must be the same as that required for DMACK today.

4 Protocols

The following clauses describe the protocols for operation of devices in overlapped/queued mode.

4.1 Power-up/ Reset Protocol

Devices supporting overlap/queuing functionality shall power-up in legacy mode following the currently specified power-up protocol. These devices shall remain in legacy mode until commanded to enter overlap mode and shall be capable of full normal legacy operation.

4.2 Entering/Leaving Overlapped Mode

To place a device into overlapped mode, the host issues a SET FEATURES command with the set overlapped mode code in the Features register. This code indicates the device should enter overlapped mode and indicates whether DEV or CS signals should be used for selection.

Once placed into overlapped mode, the device shall remain in overlapped mode until:

- a SET FEATURES command is received removing the device from overlapped mode.
- a hardware reset is received.
- the device is powered-down.

These events also cause any existing queue to be aborted.

A software reset does not remove the device from overlapped mode but does abort any existing queue.

An overlapped device in a mixed configuration is placed into Overlap/DEVsel mode. When BSY is set to one at the completion of the SET FEATURES command, the device is in Overlap/DEVsel mode and device selection is still accomplished via the DEV bit.

In a overlapped configuration both devices are placed into Overlap/CSsel mode. This is accomplished by the following protocol.

The host selects Device 1 by writing the DEV bit to one. The host then issues a SET FEATURES command to Device 1 to set Device 1 into Overlap/CSsel mode. Note that these register writes are accomplished in legacy mode as shown in the legacy column of Table 4.

When Device 1 sets the BSY bit to one at the completion of the command, the Device 1 will be in Overlap/CSsel mode and the two devices will be responding to register accesses as shown in the intermediate column in Table 4. The host shall wait 2 msec after issuing the SET FEATURES command and then check the Device 1 BSY by reading at the CS1 and CS0 signals active address.

The host will then select Device 0 by writing to its Device/Head register. Note that Device 1 no longer responds to accesses to that address. The host then issues the SET FEATURES command to Device 0 to place it into Overlap/CSsel mode.

When Device 0 has completed the command and set BSY to one, both devices are in Overlap/CSsel mode and responding to accesses as shown in the Overlap/CSsel mode column of Table 4.

Table 4 - Addressing during mode change

Address		Registers Selected		
CS0	CS1	Legacy Mode	Intermediate	Overlap/CSsel Mode
N	N	Not Used	Not Used	Device 0 Control
N	A	Device 0/1 Control	Device 0/1 Control	Device 1 Control
A	N	Device 0/1 Command	Device 0 Command	Device 0 Command
A	A	Not Used	Device 1 Command	Device 1 Command
Key: A = signal asserted, N = signal negated				

To remove the devices from the Overlap/CSsel mode the process is reversed. First, Device 0 is removed from Overlap/CSsel mode by issuing a SET FEATURES command. This returns the two devices in the state shown in the intermediate column of Table 4. When Device 0 has set BSY to one indicating the command is complete, the host issues a SET FEATURES command to Device 1 removing it from Overlap/CSsel mode. The host then waits 2 msec and selects Device 1 via the DEV bit and checks for the setting of BSY to one.

Note that the host may **not** attempt access to the Control block registers of either device when in the intermediate state.

4.3 Device Selection

When placed into overlapped mode the device will be instructed whether device selection will be accomplished by use of the DEV bit or the CS signals. In either case, when in overlapped mode a device shall only write from the bus to its command block registers when selected.

4.3.1 Selection in Overlap/DEVsel Mode

When configured in a mixed configuration the overlapped device shall be set to Overlap/DEVsel mode. In this state, device selection shall be made via the DRV bit in the Drive/Head register. In this mode, device selection is accomplished just as in legacy mode. However, when data transfer commands are received, an overlapped device in this mode shall clear BSY as soon as the required command parameters have been saved from the command block registers, and therefore, the host may select the other drive as soon as BSY is negated.

4.3.2 Selection in Overlap/CSsel Mode

When configured in an overlapped configuration, the overlapped shall be set to Overlap/CSsel Mode. In this mode, device selection is accomplished by the use of the CS signals.

The host may change the device selection via the CS signals when the currently selected device has BSY asserted. It may change the selection via the CS signals when DRQ is asserted but a PIO data transfer has not begun (i.e., no data has been written or read since the assertion of DRQ). It may change the selection via the CS signals when DMARQ is asserted if DMACK is not asserted. However, if DRQ is asserted for a PIO transfer and data has been transferred since their assertion, the data transfer shall be completed (i.e., DRQ negated) before selection shall be changed.

4.4 Register Accesses

Protocol for register accesses for a device in overlapped mode changes only slightly from that of legacy operation. When BSY or DRQ are asserted, the information returned when reading a register is undefined except for BSY and DRQ. Registers may be written any time BSY is not set.

When a non-data transfer command is issued, BSY shall remain set until the command is completed. When a data transfer command is issued BSY will remain set only until the device has saved the command and its required parameters. When not selected, a device shall ignore all register accesses on the bus.

4.5 PIO Data Transfers

When a PIO data transfer command is received, the device will set BSY, prevent all register writes, save the command and required parameters, then clear BSY and again allow register writes. If the device supports command queuing, a new command may be written; if the device does not support command queuing, the writing of a second command will cause the first command to abort.

4.5.1 PIO Data Transfers in Overlap/DEVsel Mode

When the data transfer can be accomplished, the device shall assert DRQ and, if selected, assert INTRQ. If the data transfer can be accomplished immediately upon receiving the command (e.g., a write command and a write buffer is available or read command the data is available in the device buffer), DRQ may be asserted before BSY is negated. In this case, the host may choose to execute the data transfer immediately or select the other device.

When the device has completed a command, the device shall clear DRQ, set CC, and assert INTRQ if selected.

When the host has selected the device and recognized the INTRQ, it shall read the Alternate Status register.

- If the Alternate Status register has DRQ set and BSY cleared, the device is ready to execute the data transfer. At this time, the host may either execute the data transfer or select the other device. If the host chooses to transfer the data, the data is transferred using the existing PIO data transfer protocol. The host shall not deselect the device until the data transfer has been completed.
- If the Alternate Status register has DRQ and BSY both cleared and CC set, the device has completed the command. This may be an error free completion or a completion with error depending on the state of the ERR bit just as with a legacy device.

4.5.2 PIO Data Transfers in Overlap/CSsel Mode

When the data transfer can be accomplished, the device shall assert DRQ and INTRQ. If the data transfer can be accomplished immediately upon receiving the command (e.g., a write command and a write buffer is available or read command the data is available in the device buffer), DRQ may be asserted before BSY is negated. In this case, the host may choose to execute the data transfer immediately or select the other device.

When the device has completed the requested data transfer, the device shall clear DRQ, assert ERR if applicable, and assert INTRQ.

When the host receives an interrupt from a device with a PIO data transfer command outstanding, it shall read the Interrupt register regardless of which device is selected. Note that then the host explicitly knows which device asserted the interrupt and why.

- If the Interrupt register has TRn for that device set, the device is ready to execute the data transfer. At this time, the host may either execute the data transfer or select the other

device. If the host chooses to transfer the data, the data is transferred using the existing PIO data transfer protocol. The host shall not deselect the device until the data transfer has been completed.

- If the Interrupt register has CCn for that device set, the device has completed the command. This may be an error free completion or a completion with error depending on the state of the ERR bit just as with a legacy device.

4.6 DMA Data Transfers

When a DMA data transfer command is received, the device will set BSY, prevent all register writes, save the command and required parameters, then clear BSY and again allow register writes. If the device supports command queuing, a new command may be written; if the device does not support command queuing, the writing of a second command will cause the first command to abort.

4.6.1 DMA Data Transfers in Overlap/DEVsel Mode

When the data transfer can be accomplished, the device shall assert DMARQ and INTRQ, if selected. If the data transfer can be accomplished immediately upon receiving the command (e.g., a write command and a write buffer is available or read command the data is available in the device buffer), DMARQ may be asserted before BSY is negated. In this case, the host may chose to execute the data transfer immediately or select the other device.

When the device has completed the requested data transfer, the device shall clear DMARQ, assert ERR if applicable, and if selected, assert INTRQ.

When the host has selected the device and recognized an interrupt from the device with a DMA data transfer command outstanding

- If the device has the DMARQ signal asserted, the device is ready to execute the data transfer. At this time, the host may either execute the data transfer or select the other device. If the device has a command queue outstanding the host upper byte of the Alternate Status register contains the tag of the command associated with this service request. If the host chooses to transfer the data, the data is transferred using the existing DMA data transfer protocol. The host may break up a DMA transfer in progress by deasserting DMACK with the timing protocol described in legacy mode. Having done this, the host may select the other device. The device may break up a DMA transfer in progress by deasserting DMARQ with the timing protocol described in legacy mode. The host may then wait for DMARQ to reappear or select the other device. When the device breaks a DMA transfer by deasserting DMARQ, the device shall issue its interrupt, INTRQ, when the device is selected and DMARQ is reasserted.
- If the DMARQ signal is not asserted, the host shall read the Alternate Status register. If the Alternate Status register has DRQ and BSY both cleared and CC set, the device has completed the command. This may be an error free completion or a completion with error depending on the state of the ERR bit just as with a legacy device. If the device has a command queue outstanding, the host shall check the upper byte of the Alternate Status register to determine which command this is the completion status for.

4.6.2 DMA Data Transfers in Overlap/CSsel Mode

When the data transfer can be accomplished, the device shall assert INTRQ and, if selected, assert DMARQ. If the data transfer can be accomplished immediately upon receiving the command (e.g., a write command and a write buffer is available or read command the data is

available in the device buffer), DMARQ may be asserted before BSY is negated. In this case, the host may chose to execute the data transfer immediately or select the other device.

When the device has completed the requested data transfer, the device shall clear DMARQ, assert ERR if applicable, and assert INTRQ.

When the host receives an interrupt from a device with a DMA data transfer command outstanding, it may determine the interrupting device and the reason for the interrupt by reading the Interrupt register. It shall select the interrupting device using the protocol described above.

- If the device had the TRn bit asserted, the device is ready to execute the data transfer. At this time, the host may either execute the data transfer or select the other device. If the device has a command queue outstanding the host upper byte of the Alternate Status register contains the tag of the command associated with this service request. If the host chooses to transfer the data, the data is transferred using the existing DMA data transfer protocol. The host may break up a DMA transfer in progress by deasserting DMACK with the timing protocol described in legacy mode. Having done this, the host may select the other device. The device may break up a DMA transfer in progress by deasserting DMARQ with the timing protocol described in legacy mode. The host may then wait for DMARQ to reappear or select the other device. When the device breaks a DMA transfer by deasserting DMARQ, the device shall issue INTRQ when DMARQ is reasserted.
- If the CCn bit was set, the host shall read the Alternate Status register. If the Alternate Status register has DRQ and BSY both cleared and CC set, the device has completed the command. This may be an error free completion or a completion with error depending on the state of the ERR bit just as with a legacy device. If the device has a command queue outstanding, the host shall check the upper byte of the Alternate Status register to determine which command this is the completion status for.

5 ATA Standard Modifications

This clause describes the specific modifications to the existing ATA standard to incorporate Command overlap and command queuing. Added text is shown in *italics*. Deleted text is shown with crossouts,. These modifications are presented in the order that they appear in the existing standard with clause numbers appearing in X3T10/2008D revision 0.

Clause 3.1 Definitions and Abbreviations - add:

Command overlap - The ability to issue a command to the second device while the first device is in the process of executing a command.

Command queuing - The ability to issue subsequent commands to a device while previous commands have yet to complete.

Command Tag - An identifier associated with a command that allows the host to determine which of a set of queued commands is being executed.

Clause 4.5.1 ATA Driver Types and Required Pull-ups

Table 6

Signal	Source	Driver Type(1)	Pull-up at Host (2)	Pull-up at each Device (2)	Notes
INTRQ1	Device	?			

Clause 5.2.10 INTRQ (Device interrupt) - add:

When in Overlap/CSsel mode, INTRQ is a “wire-or” signal and a device may assert INTRQ anytime the device has an interrupt pending and nIEN is cleared.

Clause 6.1 Device Addressing Considerations - add:

A device may be placed in Overlap/DEVsel mode if it shares the bus with a device incapable of supporting overlap. In this case, device select is accomplished using the DEV bit as described above. In this mode, all register accesses except Device/Head register accesses are ignored when the device is not selected.

If both devices on the bus support overlap mode, both will be placed in Overlap/CSsel mode. In this mode, the DEV bit is ignored and the CS signals are used to select the devices. In this mode all register accesses are ignored when the device is not selected except accesses to the Interrupt register. Table 9a describes selection addresses when in this mode.

When in Overlap/DEVsel or Overlap/CSsel mode, all register accesses become 16-bit accesses. The content of the upper byte for Data, Command, Status, and Alternate Status accesses are defined. The upper byte of all other register accesses is reserved and shall be written as zeros and ignored on read.

Table 9a - Overlap/CSsel mode addressing

Addresses					Functions	
CS0-	CS1-	DA2	DA1	DA0	Read (DIOR-)	Write (DIOW-)
					Device 0 Control Block registers	
N	N	0	x	x	Data bus high impeded	Not Used
N	N	1	0	0	Data bus high impeded	Not Used
N	N	1	0	1	Interrupt	Not Used
N	N	1	1	0	Dev 0 Alternate Status	Dev 0 Device Control
N	N	1	1	1	Data bus high impeded	Not Used
					Device 1 Control Block registers	
N	A	0	x	x	Data bus high impeded	Not Used
N	A	1	0	0	Data bus high impeded	Not Used
N	A	1	0	1	Interrupt	Not Used
N	A	1	1	0	Dev 1 Alternate Status	Dev 1 Device Control
N	A	1	1	1	Data bus high impeded	Not Used
					Device 0 Command Block registers	
A	N	0	0	0	Dev 0 Data	Dev 0 Data
A	N	0	0	1	Dev 0 Error	Dev 0 Features
A	N	0	1	0	Dev 0 Sector Count	Dev 0 Sector Count
A	N	0	1	1	Dev 0 Sector Number LBA bits 7-0 (2)	Dev 0 Sector Number LBA bits 7-0 (2)
A	N	1	0	0	Dev 0 Cylinder Low LBA bits 15-8 (2)	Dev 0 Cylinder Low LBA bits 15-8 (2)
A	N	1	0	1	Dev 0 Cylinder High LBA bits 23-16 (2)	Dev 0 Cylinder High LBA bits 23-16 (2)
A	N	1	1	0	Dev 0 Device/Head LBA bits 27-24 (2)	Dev 0 Device/Head LBA bits 27-24 (2)
A	N	1	1	1	Dev 0 Status	Dev 0 Command
					Device 1 Command Block registers	
A	A	0	0	0	Dev 1 Data	Dev 1 Data

A	A	0	0	1	<i>Dev 1 Error</i>	<i>Dev 1 Features</i>
A	A	0	1	0	<i>Dev 1 Sector Count</i>	<i>Dev 1 Sector Count</i>
A	A	0	1	1	<i>Dev 1 Sector Number</i> <i>LBA bits 7-0 (2)</i>	<i>Dev 1 Sector Number</i> <i>LBA bits 7-0 (2)</i>
A	A	1	0	0	<i>Dev 1 Cylinder Low</i> <i>LBA bits 15-8 (2)</i>	<i>Dev 1 Cylinder Low</i> <i>LBA bits 15-8 (2)</i>
A	A	1	0	1	<i>Dev 1 Cylinder High</i> <i>LBA bits 23-16 (2)</i>	<i>Dev 1 Cylinder High</i> <i>LBA bits 23-16 (2)</i>
A	A	1	1	0	<i>Dev 1 Device/Head</i> <i>LBA bits 27-24 (2)</i>	<i>Dev 1 Device/Head</i> <i>LBA bits 27-24 (2)</i>
A	A	1	1	1	<i>Dev 1 Status</i>	<i>Dev 1 Command</i>

Clause 6.2.1 - Alternate Status register - add:

FIELD/BIT DESCRIPTION - add:

Overlap Mode							
15	14	13	12	11	10	9	8
CC	r	r	r	Tag (3:0)			

- *Bit 15 - CC (Command Complete) When in overlapped mode, this bit is Command Complete (CC). It is set upon completion of a command when the status for that command is present in the Status and Error registers. Upon the completion of a read from the Status register, the CC bit is cleared.*
- *Bits 14:12- reserved*
- *Bits 11:8 - Queue tag*

Clause 6.2.2 Command register - add:

FUNCTIONAL DESCRIPTION - add:

When in overlapped mode, the host shall set the fact that a data transfer command shall execute in overlapped mode for each transfer command. If the device supports command queuing, the queue tag for the command shall also be provided.

FIELD/BIT DESCRIPTION - add:

Overlap Mode							
15	14	13	12	11	10	9	8
OVL	r	r	r	Tag (3:0)			

Bits 15 - OVL when set indicates the command shall be executed in overlapped mode.

Bits 14-12 - reserved

Bits 11-8 - Queue tag.

Clause 6.2.7 Device/Head register - add:

FUNCTIONAL DESCRIPTION - add:

When in Overlap/CSsel Mode the DEV bit is not used for device selection and is replaced by CS signal device selection.

Clause 6.2.12 Status register - add:

FIELD/BIT DESCRIPTION - add:

- BSY (Busy) - add:

When a command is accepted and the device is in overlapped mode, BSY shall be set as described above, the device shall save the command and applicable parameters contained in command block registers, then clear BSY. This allows the host to select the other device if desired. If the device supports command queuing, another command may be issued to the device as soon as BSY has been cleared.

15	14	13	12	11	10	9	8
CC	r	r	r	TAG (3:0)			

- CC (Command Complete) When in overlapped mode, this bit is Command Complete (CC). It is set upon completion of a command when the status for that command is present in the Status and Error registers. Upon the completion of a read from this register, the CC bit is cleared.

- TAG (3:0) When in overlapped mode with queued commands, the tag field indicates the command for which an interrupt was issued or status is being presented.

6.2.? Interrupt register -add clause

ADDRESS - CS(1:0)=1h, DA(2:1)

DIRECTION - This register is read-only to the host.

ACCESS RESTRICTIONS - The content of this register is valid any time both devices are in Overlap/CSsel mode. However, after the read of a Status register to clear a device interrupt, the host shall wait ?ns before reading this register to insure that the device has cleared the effected bit.

EFFECTIVENESS - Reading this register shall have no effect on the devices.

FUNCTIONAL DESCRIPTION - This register allows the host to determine which device(s) are interrupting and whether the interrupt is asserted to continue data transfer for a command in progress or to notify that a command is complete. The device sets the appropriate bit before asserting INTRQ and clears the bit when the INTRQ is negated due to a Status register read.

This is a unique register. It shall be read only when both attached devices are in Overlap/CSsel mode. Both devices shall respond to a read of this register address regardless of the state of CS1-. Each device shall only drive the bits associated with its assigned device number. The host shall ignore the state of all reserved bits.

FIELD/BIT DESCRIPTION -

7	6	5	4	3	2	1	0
r	r	r	r	CC1	TR1	CC0	TR0

TR0 indicates Device 0 has an interrupt pending to transfer data.

CC0 indicates Device 0 has an interrupt pending because a command has completed.

TR1 indicates Device 1 has an interrupt pending to transfer data.

CC1 indicates Device 1 has an interrupt pending because a command has completed.

r reserved

Add clause:

7.X Overlapped/Command Queuing Mode

7.X.1 Overlapped configurations

Two configurations are supported by overlapped mode.

7.X.1.1 Mixed Configurations

The mixed configuration allows one device on the bus to behave as a legacy device while the other device is operating in overlapped/queuing mode as shown in Figure X1. This allows a host that is capable of supporting overlap/queuing to make use of the functionality even when a legacy device incapable of the functionality is present. Figure X1 shows how commands for the two devices may be interleaved when Device 0 is the Legacy Device and Device 1 is the Overlapped Device. The overlapped device may in fact be configured as either Device 0 or Device 1 in this configuration.

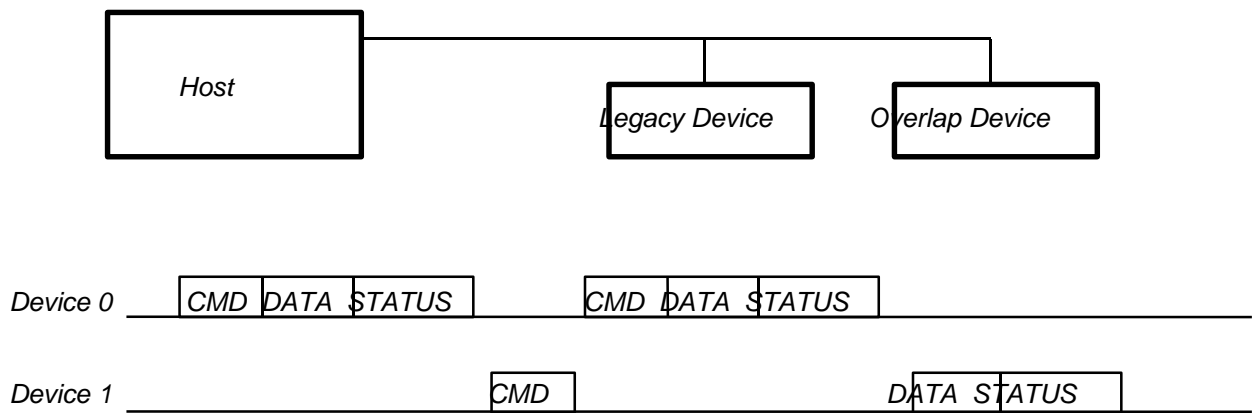


Figure X1- Mixed Configuration

7.X.1.2 Overlapped Configuration

The overlapped configuration can occur when the host and both devices are capable of overlapped functionality as shown in Figure X2. Overlapped mode is designed to provide minimum overhead and maximum performance with this configuration. As shown by the first command to each device, the issuing of a command to a second device can be placed between the issuing of a command to the first device and the transfer of the first device data. If the device supports queuing as well as overlap, a subsequent commands may be issued to a device before a first command completes as shown by the second and third commands to device 0. In addition, if the data transfer is made in DMA mode, the data transfer may be completed in multiple data transfers instead of a single transfer as shown.

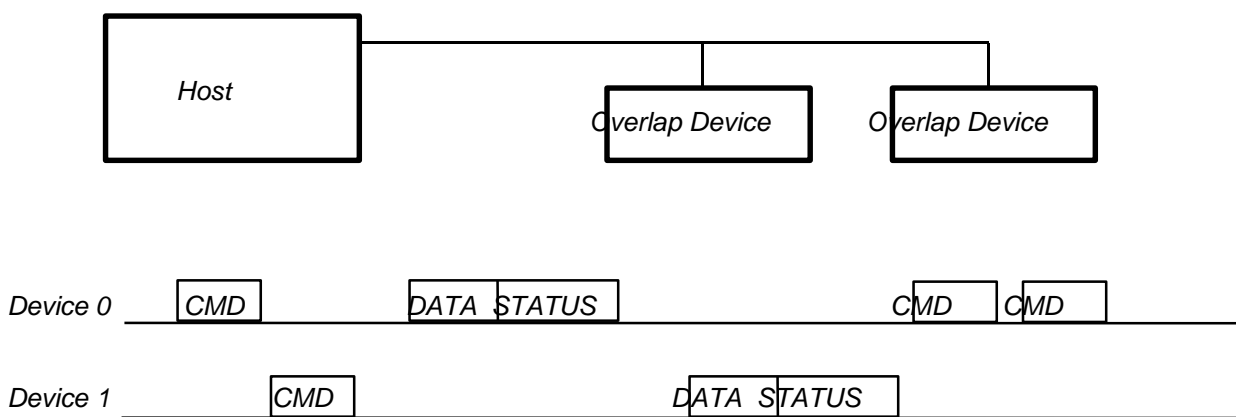


Figure X2 - Overlapped Configuration

7.X.2 Overlapped Register Ownership

When a command is issued to a device in overlapped mode, BSY is asserted. If the command is a non-data transfer command, BSY remains asserted until the command completes but if both devices support overlap, the host may select the other device as soon as the command has been written. If the command is a data transfer command, the device sets BSY, inhibits writing of its command block registers, saves the contents of the applicable command block registers, and then clears BSY. If both devices support overlap, the host may select the other device as soon as the command has been written. If the device supports queuing, the host may issue a second command as soon as BSY is cleared.

If the configuration is a Mixed Configuration, the host must obey the rules that each device expects. That is, if a command is issued to the legacy device, that legacy device own all registers until that command completes. The host may not access the other device. If the host issues a command to the overlapped device, the host may select the legacy device and issue it a command as soon as BSY has been cleared by the overlapped device.

7.X.3 Device Selection

In a mixed configuration, a device capable of command overlap/queuing is selected just as in legacy mode. Since this device will clear its BSY bit as soon as the command parameters have been saved on a data transfer command, a host may then select the legacy device and execute a command while the overlap device is working on its data transfer command.

In an overlapped configuration where both devices are in Overlap/CSsel mode, selection is accomplished with the CS signals.

In the overlapped configuration, since both devices lockout writes to their command block registers when BSY is set and only accept writes when they are selected, a host may change the device selection at any time it is not actively transferring data or accessing a register.

7.X.4 Interrupt Reporting

When placed in Overlap/CSsel mode, a device implementing this function will use the INTRQ line for reporting interrupts and this signal shall be “wire-or”. Interrupts may be asserted at any time regardless of whether or not the device is currently selected. This allows the devices to notify the host that service is required regardless of current selection. The host may read the Interrupt register with either device selected to determine which device is interrupting and the reason for the interrupt.

When in Overlap/DEVsel Mode, the device shall use INTRQ for interrupts as in legacy mode and shall assert INTRQ only when selected. In this case, the host must read the Alternate Status or Status register of each device to determine the reason for the interrupt.

7.X.5 Command Queue Tags

The Identify Device info returned by a device capable of overlap/queuing operation shall contain a four bit queue depth field. If this field is 0h, the device does not support queuing. If the field contains any other value, this value plus one indicates the depth of queue supported up to a maximum queue depth of sixteen. Any tag issued by the host with a tag number greater than this value shall cause the command to be aborted.

When issuing a data transfer command to a device in overlapped/queuing mode, the tag field in the upper byte of the Command register shall indicate that this is an overlapped command and provide the queue tag number for this command.

When a queued device interrupts for service for a queued command, the tag for the command requiring service shall be contained in the upper byte of the Status or Alternate Status register.

All queued commands are simple queued data transfer commands, no priority is implied.

7.x.6 Host implications

7.x.6.1 Hosts without DMA engines

With both devices in legacy mode, hosts function just as they have in the past.

With mixed mode devices, device selection and interrupts behave just as they have in the past. Therefore, the host will have to select the overlap device via the DRV bit to determine when an interrupt has occurred indicating that is a continuation of an overlapped command.

When both devices are placed into overlap/CSsel mode, the host adapter must have a register that allows device selection. The writing of this register shall cause the CS1- signal to select the desired device until changed by the host software. The remaining address bits, i.e., CS0- and DA[2:0], are decoded from the host bus address just as with legacy devices. The interrupts from the two devices will occur when a device is not selected and the host will have to read the Interrupt register or Alternate Status or Status register of each device to determine which device interrupted and what the cause was.

7.x.6.2 Hosts with DMA engines

Command overlap and command queuing make the use of an intelligent DMA engine in the host an attractive added feature to improve performance particularly with a multithreaded operating system. However, with such an intelligent DMA engine, the accesses to the devices by the DMA engine and the host device driver must be synchronized. That is, the device driver may not access the devices while the DMA engine is in the process of transferring data to/from a device, and the DMA engine may not transfer data while the host device driver is accessing registers to issue a command or check status.

The following is an example of how this synchronization may be accomplished. The DMA engine will have a register accessible by the host. The ATAREQ bit in the register will be writable by the host. When the host device driver wishes to access a device to issue a command or read status, the host will set the ATAREQ bit to one. A register bit must exist to allow the device driver to select a device.

If the DMA engine is not currently transferring data to/from a device, it will set the ATAACK bit to one. When both the ATAREQ and ATAACK bits are one, the host device driver may access a device and the DMA engine shall not access a device.

If the DMA engine is in the process of transferring data to/from a device when the ATAREQ bit is set, it will complete the transfer and then set the ATAACK bit to one.

When the host device driver has completed the required device accesses, it will set the ATAREQ bit to zero. The DMA engine will immediately set the ATAACK bit to zero, and the DMA engine may now access a device.

When the ATAREQ and ATAACK bits are both asserted, the host adapter must pass the register value of CS1- selected by the device driver and the remainder of the address, i.e., CS0- and DA[2:0] as decoded from host bus address to the devices. When the ATAREQ and ATAACK bits are both cleared, the DMA engine must assert and latch CS1- to select the device to execute a DMA transfer. Timing for the assertion and negation for CS1- in the case must be the same as that required for DMACK today.

Clause 8.10 IDENTIFY DEVICE - add:

Table 14 - Identify Device Information - add:

Word	F/V	
71	F	15 Overlapped Mode supported 14-11 Queue depth supported 10-0 Reserved
72-127	R	Reserved

Add **clause 8.10.36** Word 71: Overlapped Mode/Command Queuing Supported

Word 71 of the parameter information of the IDENTIFY DEVICE command is defined as the Overlapped Mode/Command Queuing supported field. If bit 15 is set to one, the device supports Overlapped Mode. This mode may be set via a SET FEATURES command.

Bits 14-11 define the queue depth supported if command queuing is supported when in overlapped mode. If 0h is present in this field, the device does not support command queuing. If this field contains a value other than 0h, that value + 1 is the maximum queue depth supported by the device.

Clause 8.32 SET FEATURES add:

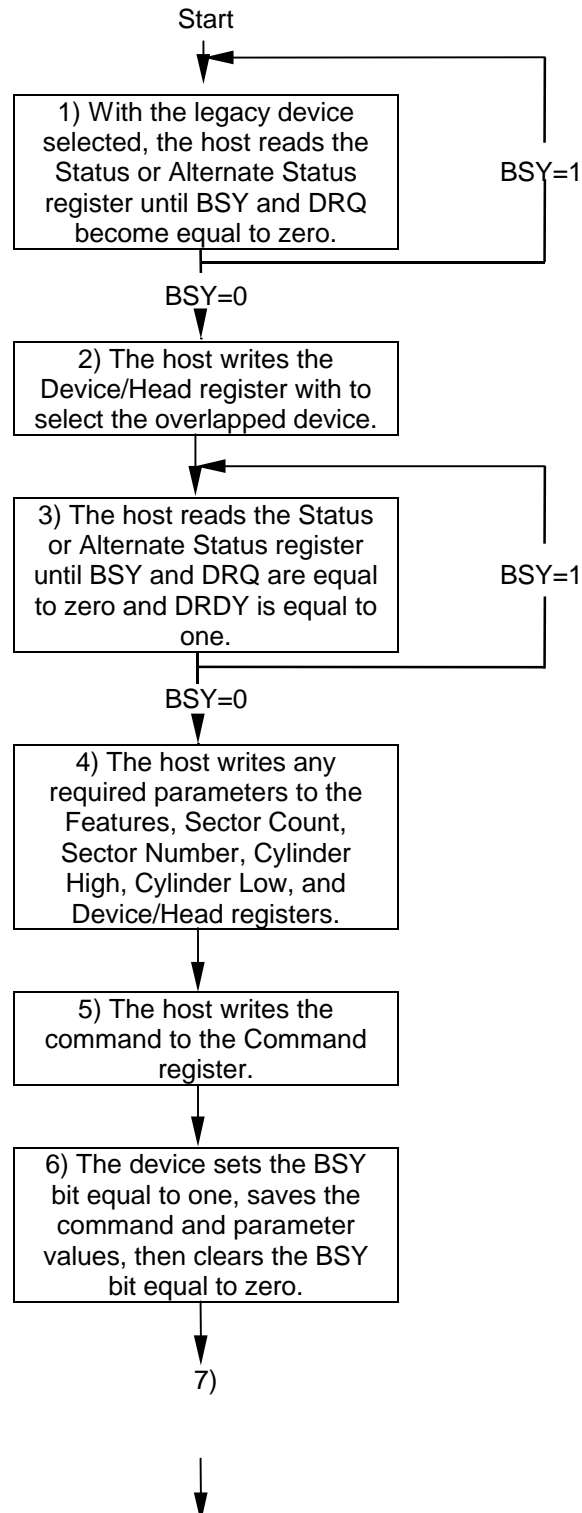
Table 17 - Set Features register Definitions add:

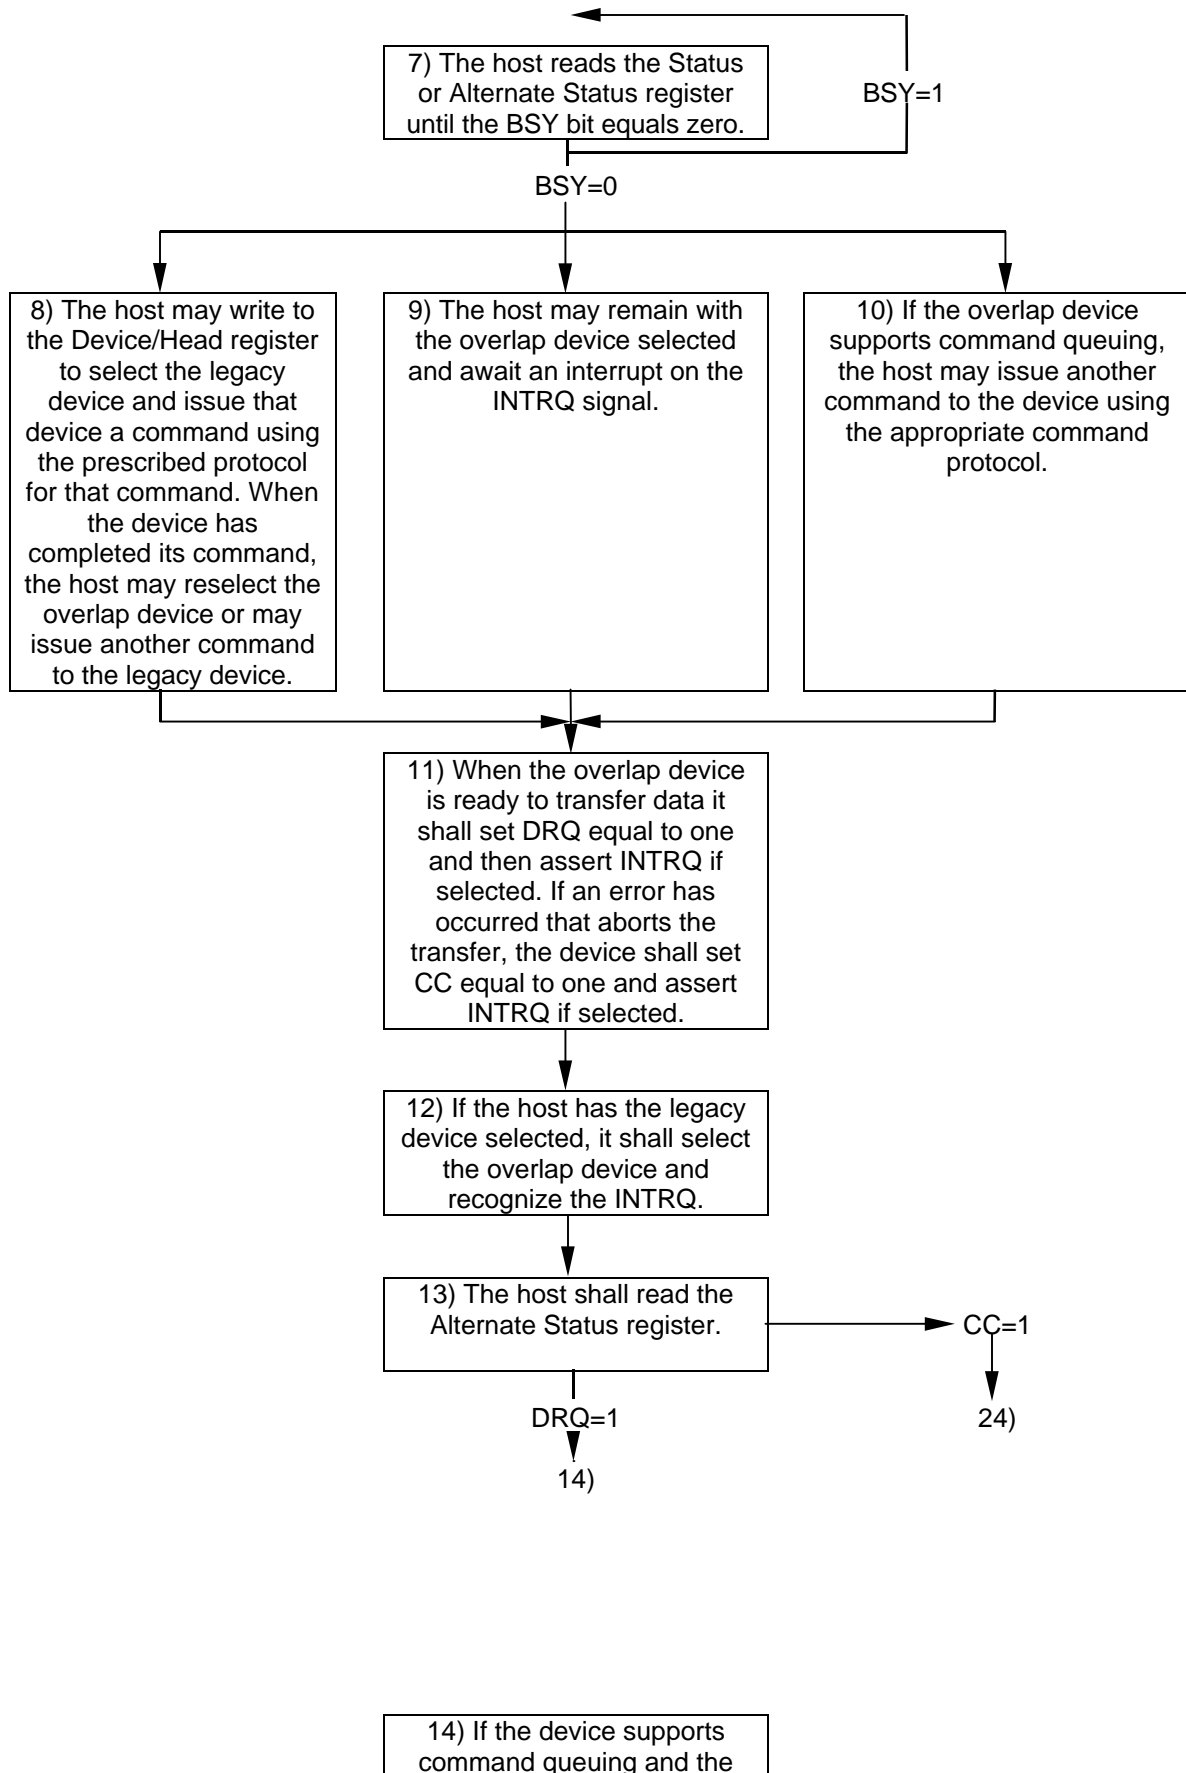
<i>DDh</i>	<i>Enable Overlap/DEVsel mode</i>
<i>DEh</i>	<i>Enable Overlap/CSsel mode</i>
<i>DFh</i>	<i>Disable Overlap mode</i>

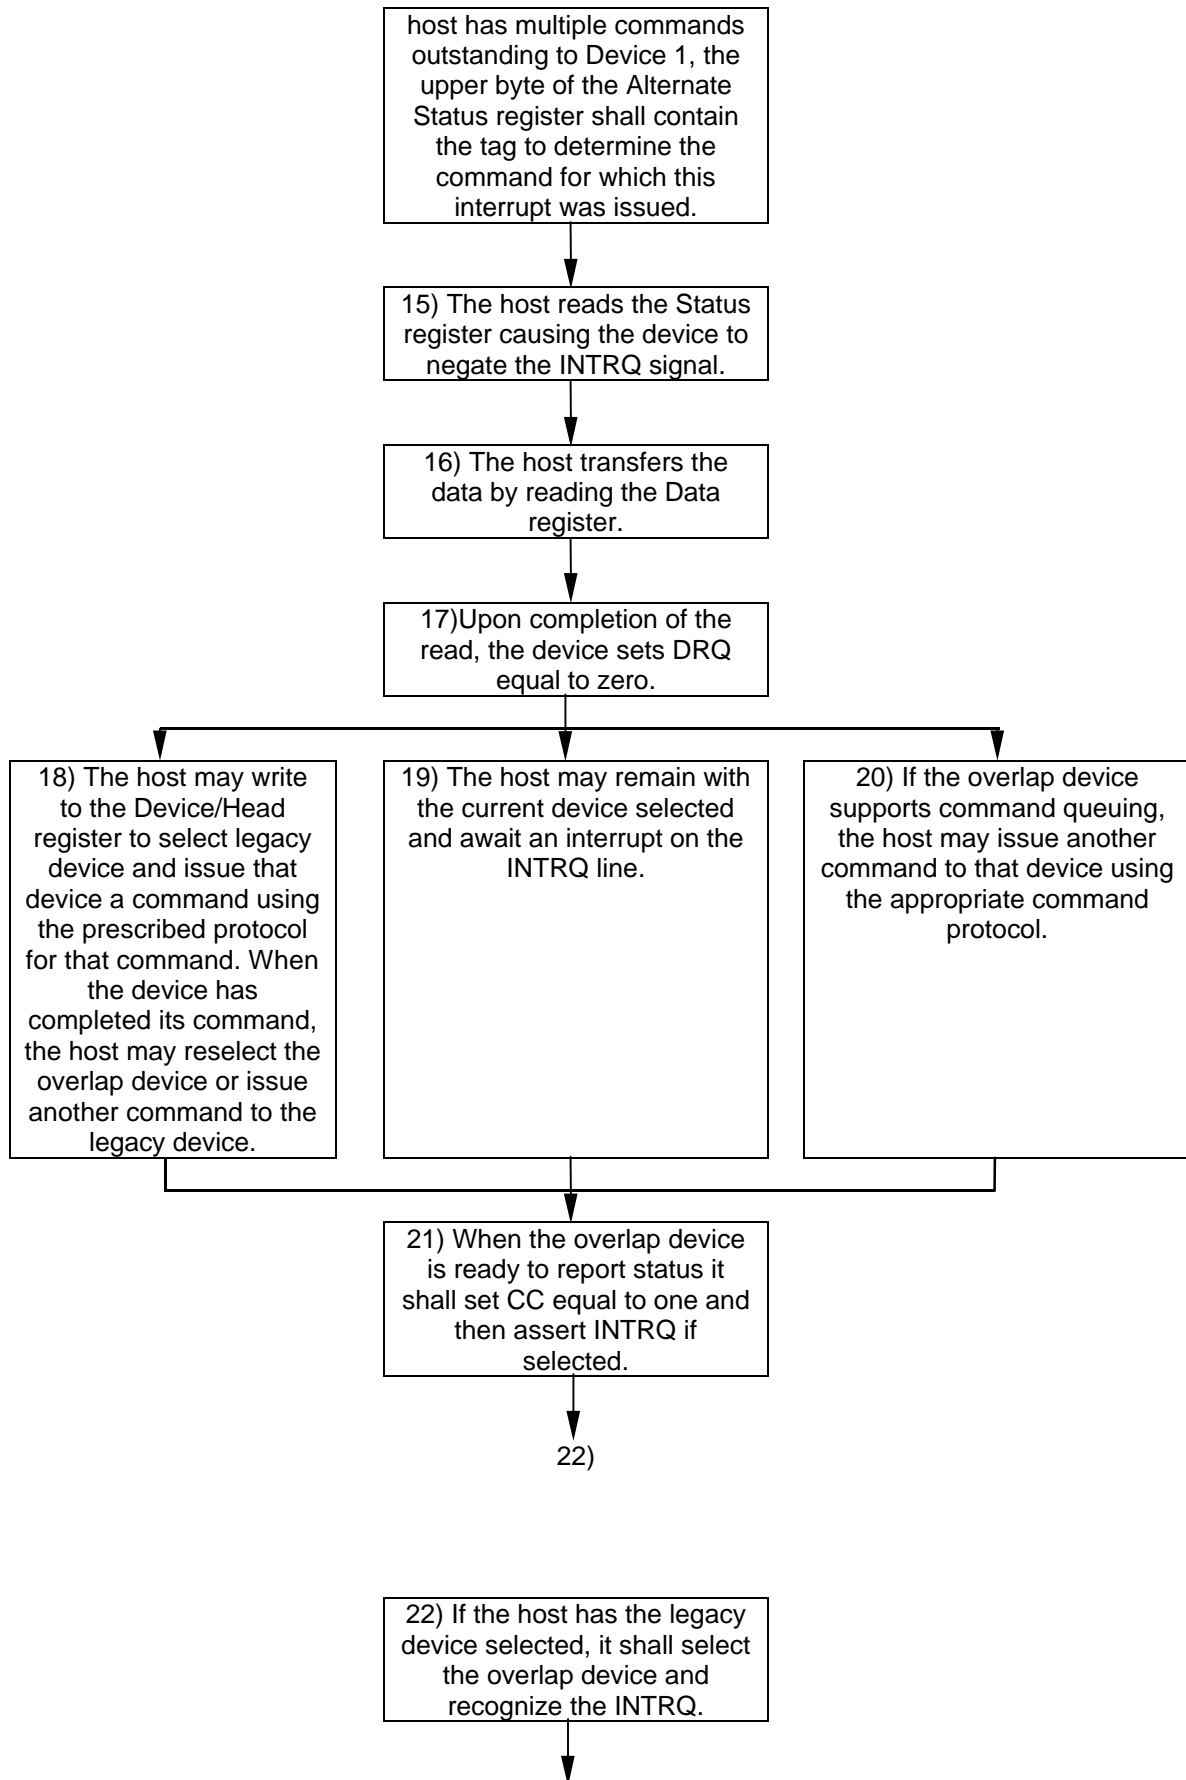
Clause 9.3 PIO data in commands - add:

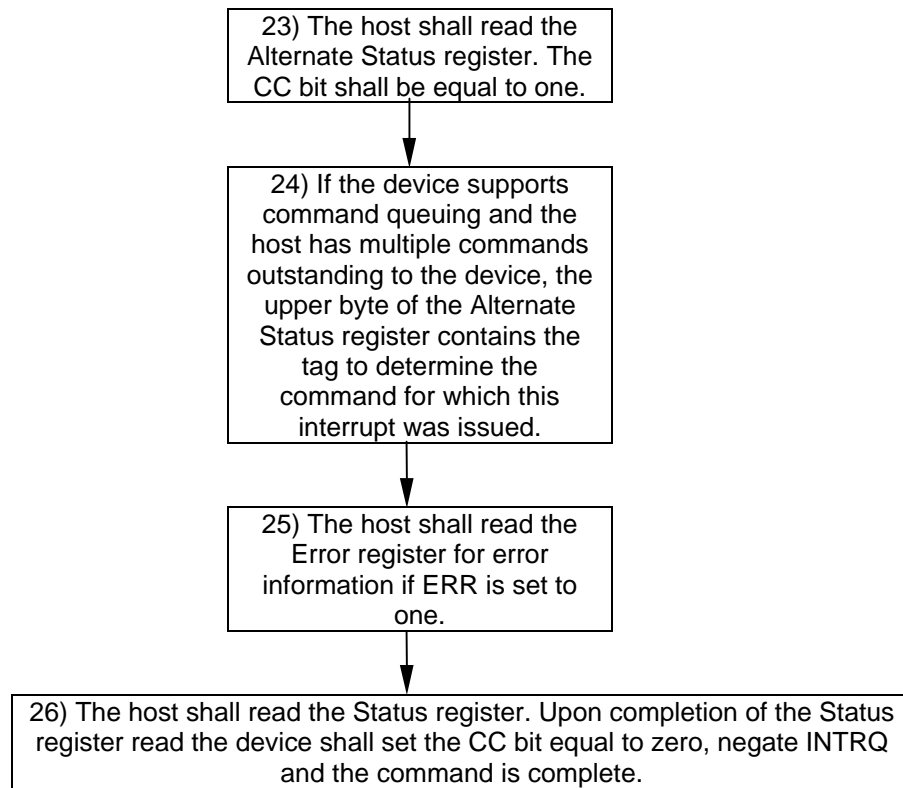
9.3.1 Overlap/DEVsel mode PIO data in commands

When a mixed configuration exists with one legacy device and one device that supports overlapped mode, the device that supports overlapped mode may be set into Overlap/DEVsel mode. The overlapped device then operates with the following protocol:



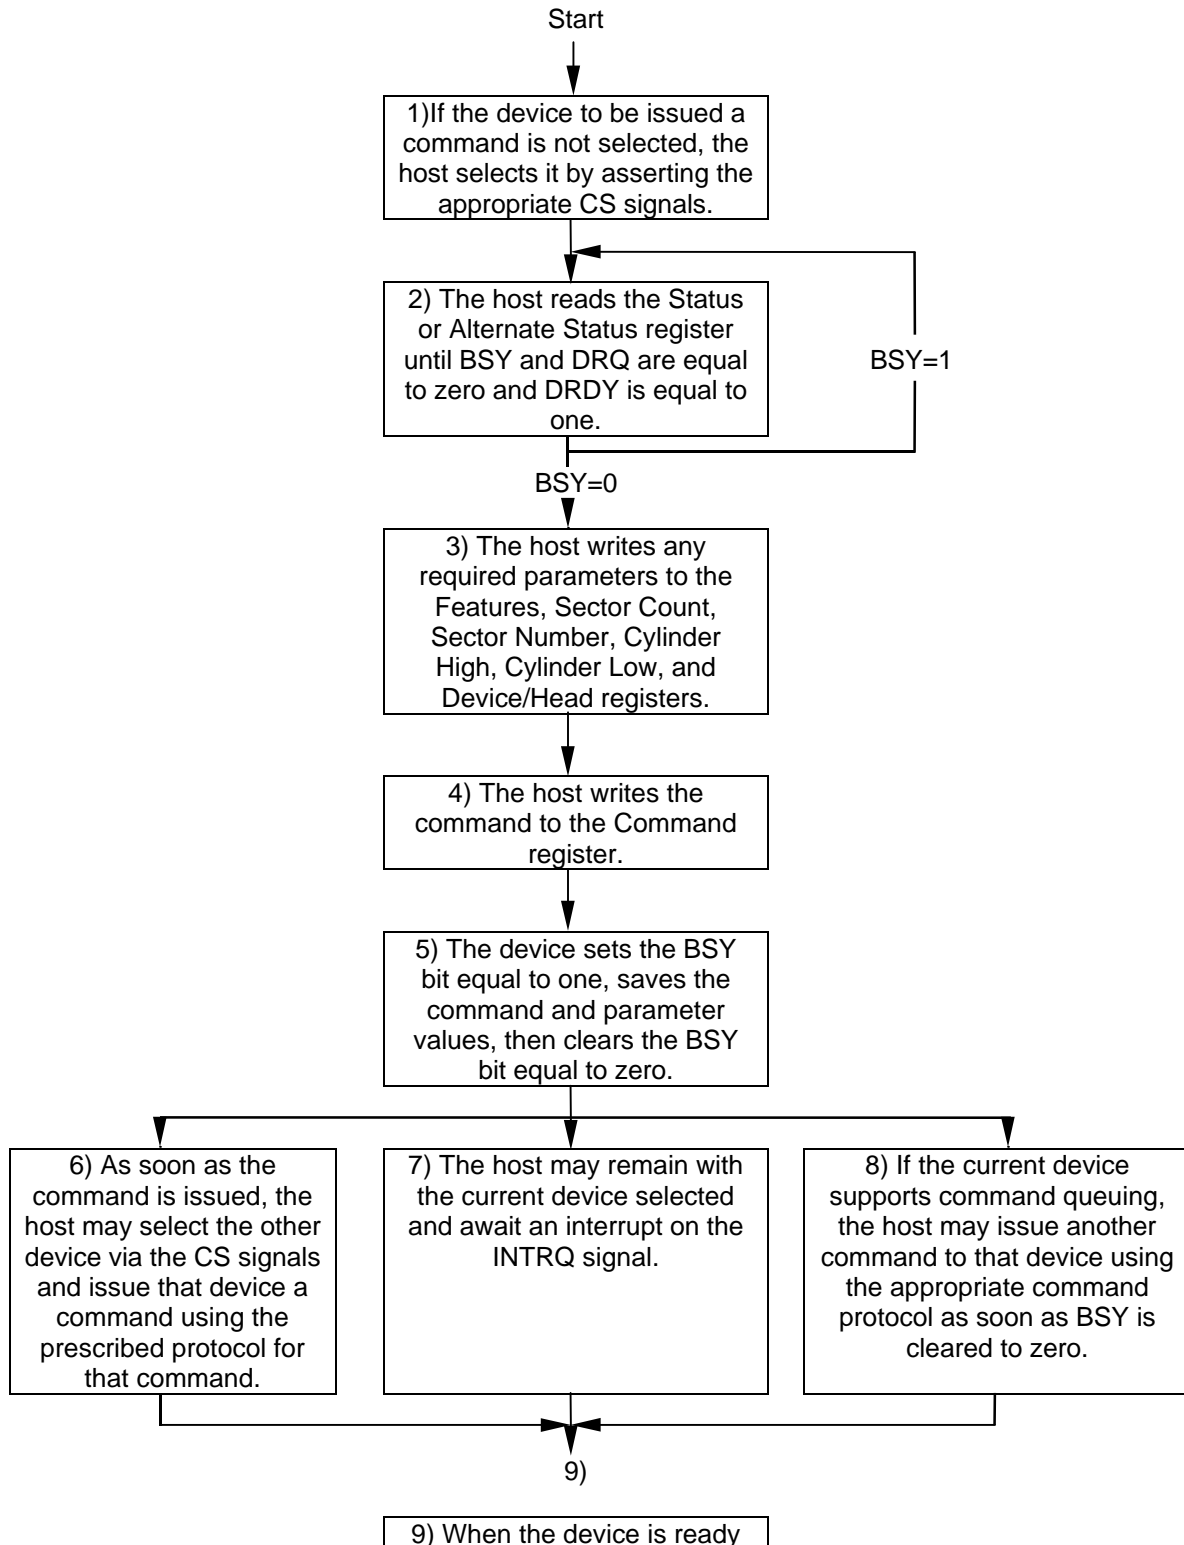


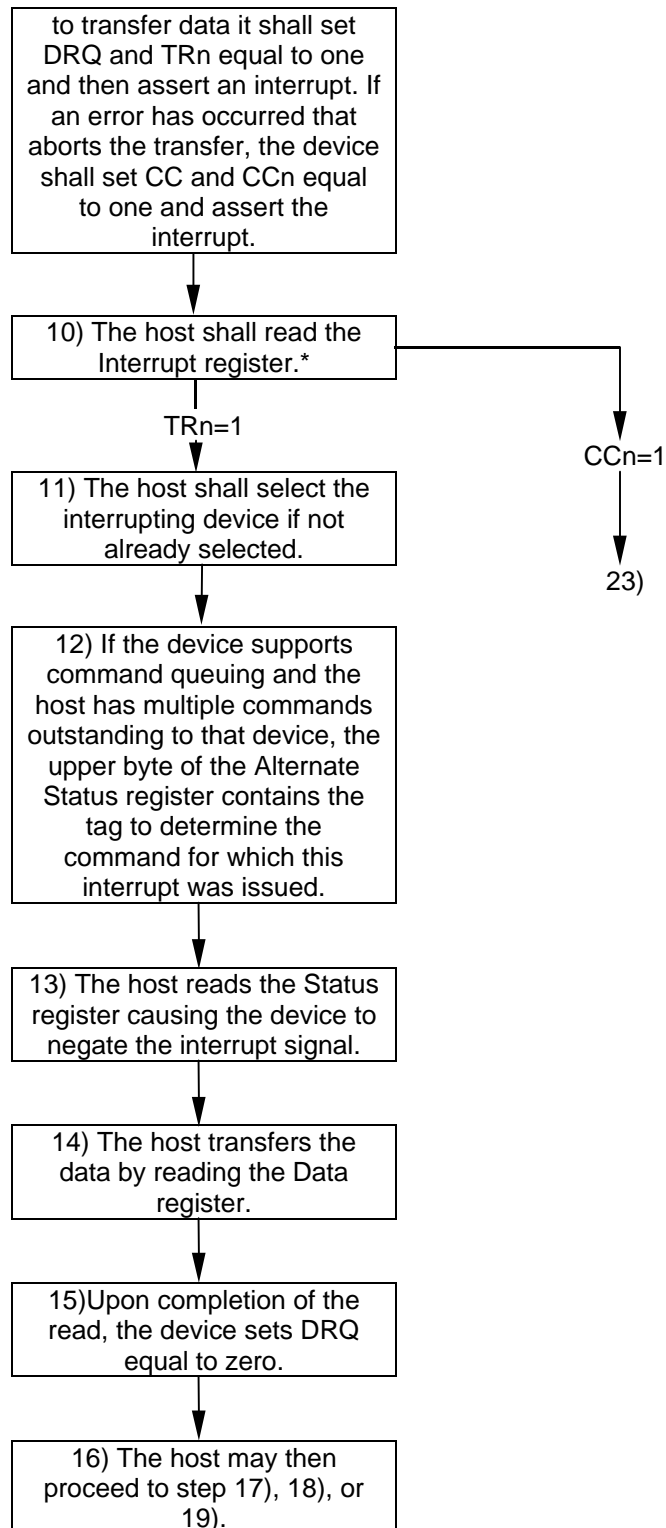




9.3.1 Overlap/CSsel mode PIO data in commands

When an overlapped configuration exists with both devices that support overlapped mode, the devices are set into Overlap/CSsel mode. In this configuration, device selection is accomplished via the CS signals. The device then operate with the following protocol:

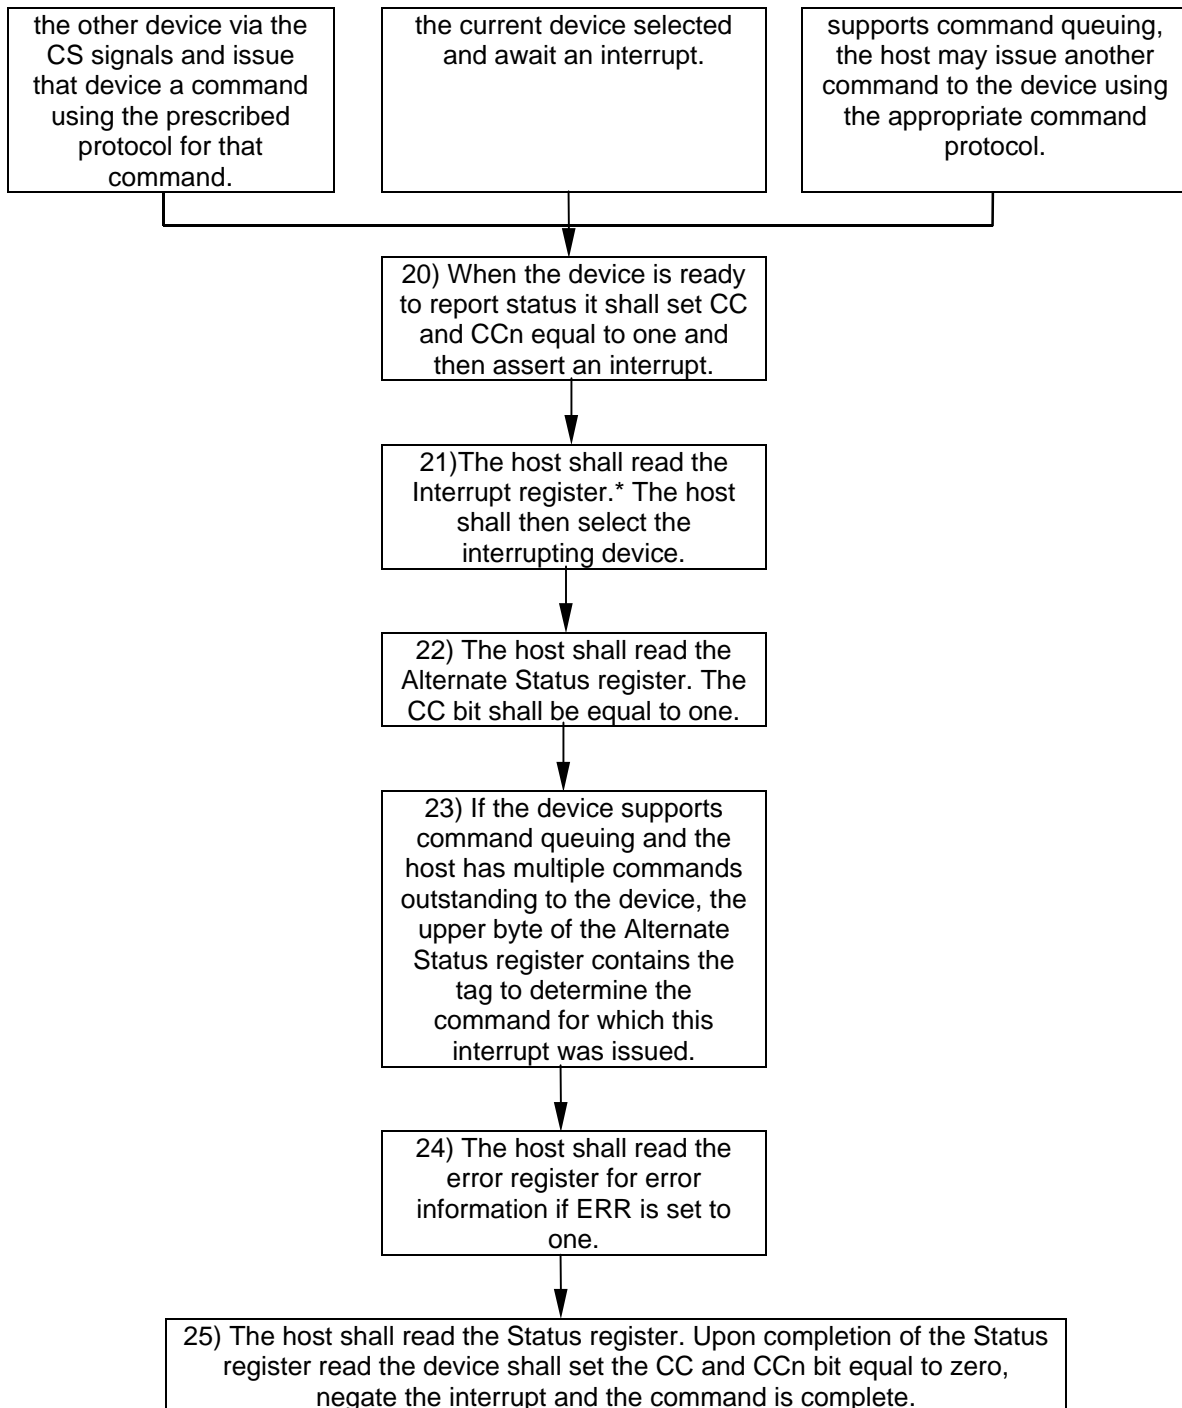




17) The host may select

18) The host may remain with

19) If the current device

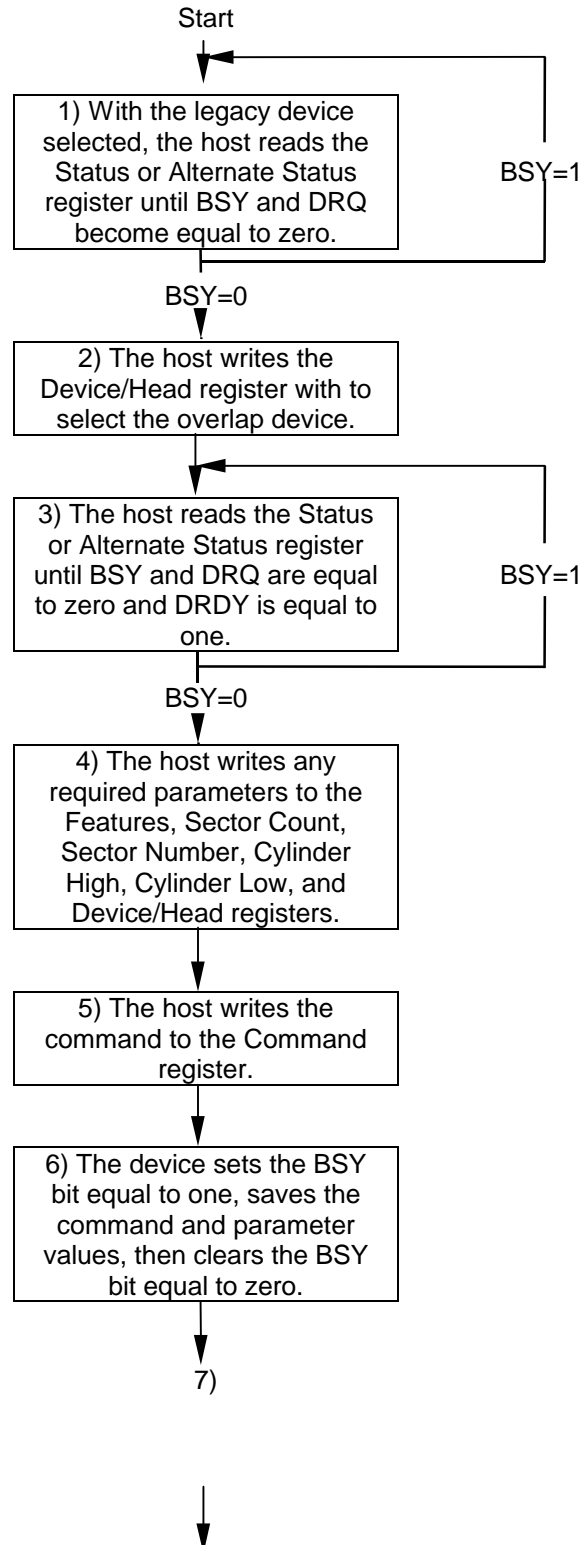


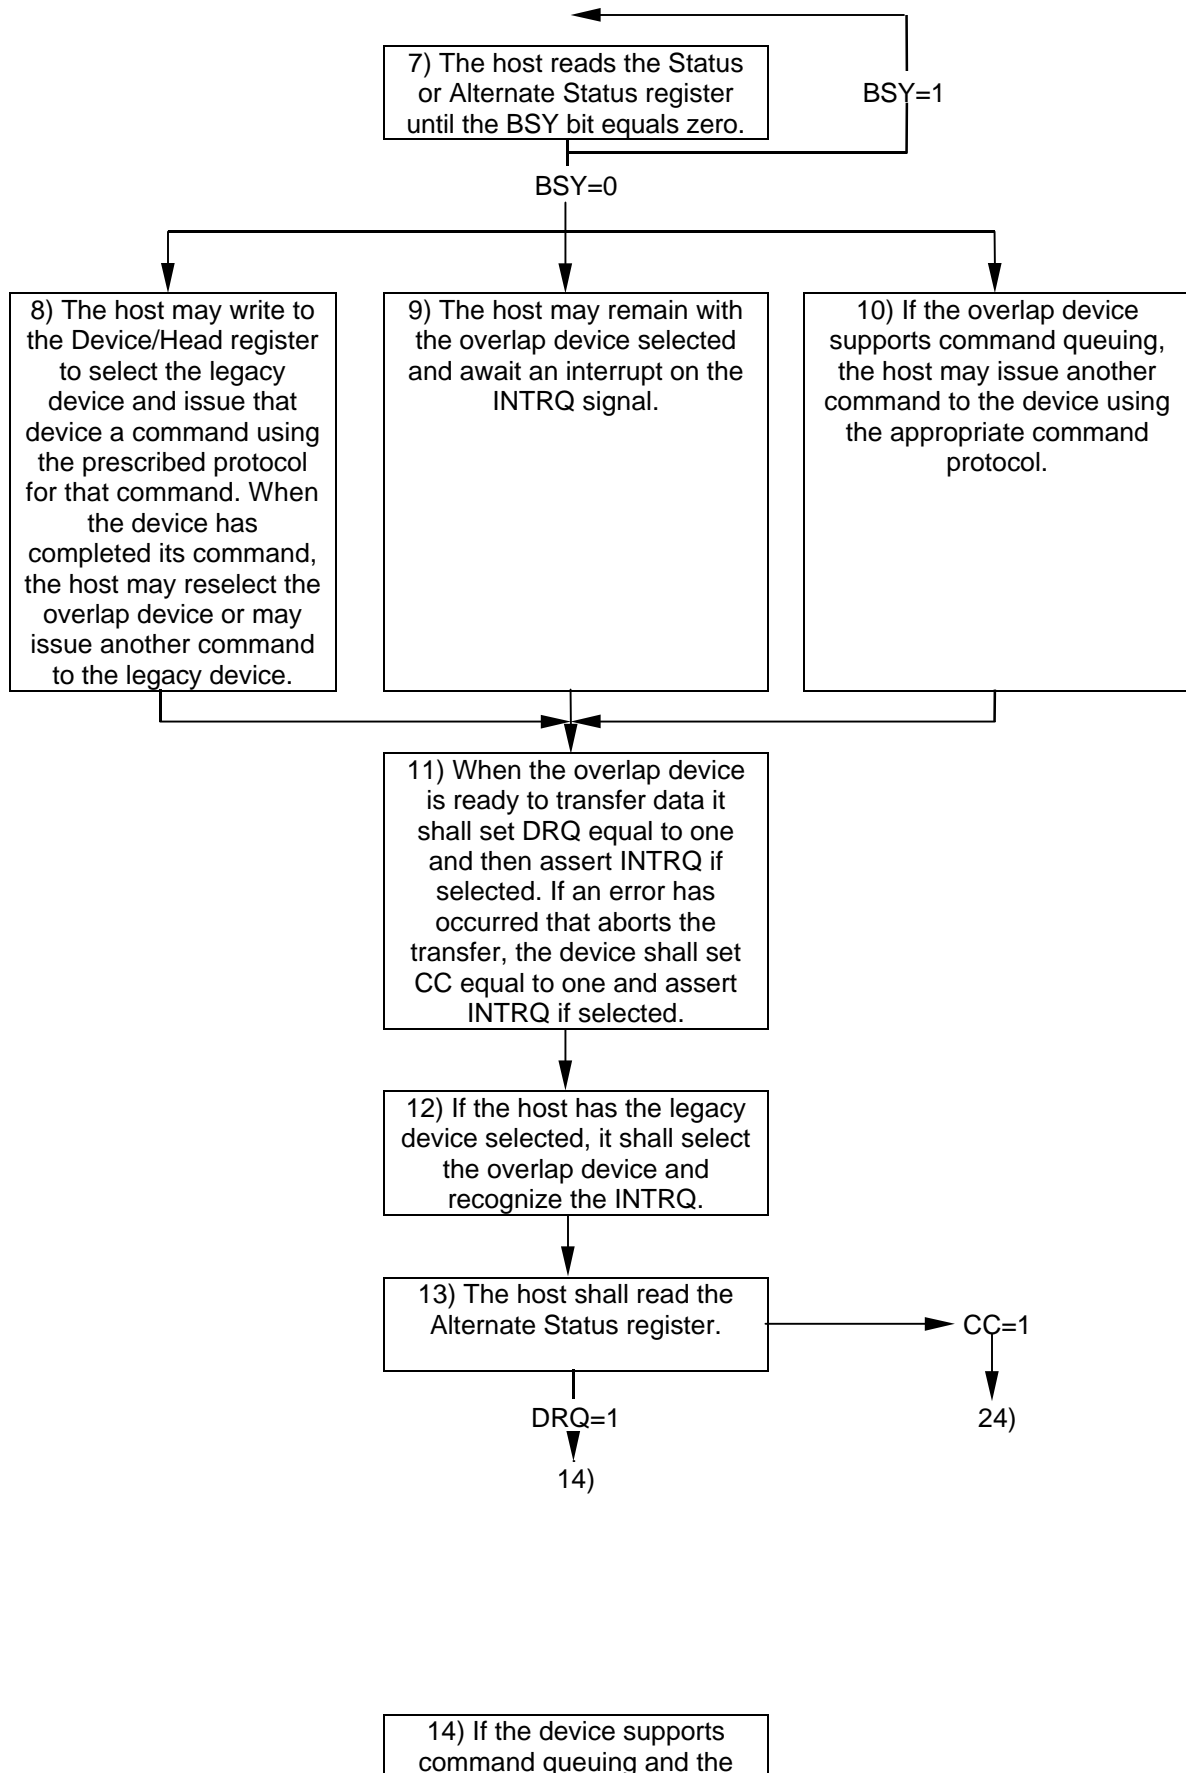
* Note: with today's 3F6/376h host addressing only an intelligent DMA engine may be capable of reading the Interrupt register. If the host is unable to read the Interrupt register, the status of each device will have to be read to determine the cause of the interrupt.

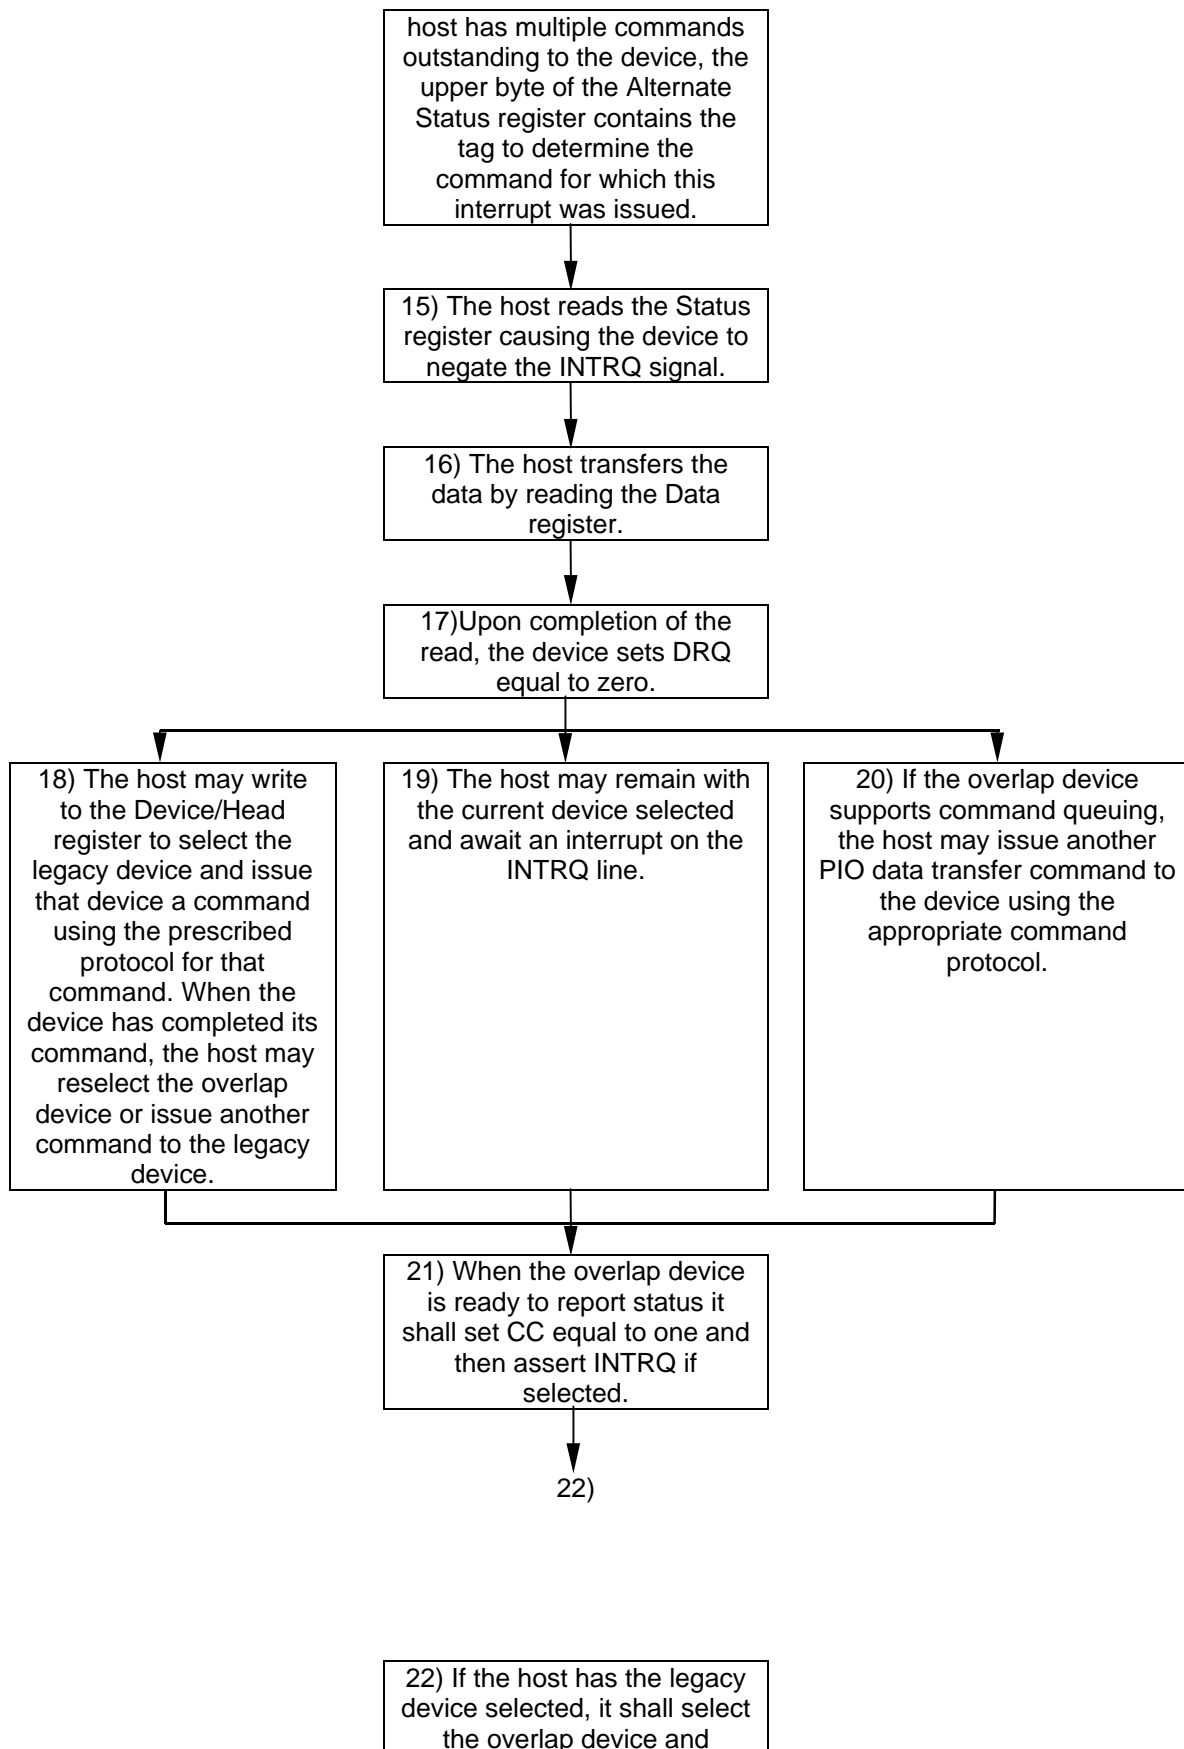
Clause 9.4 PIO data out commands - add:

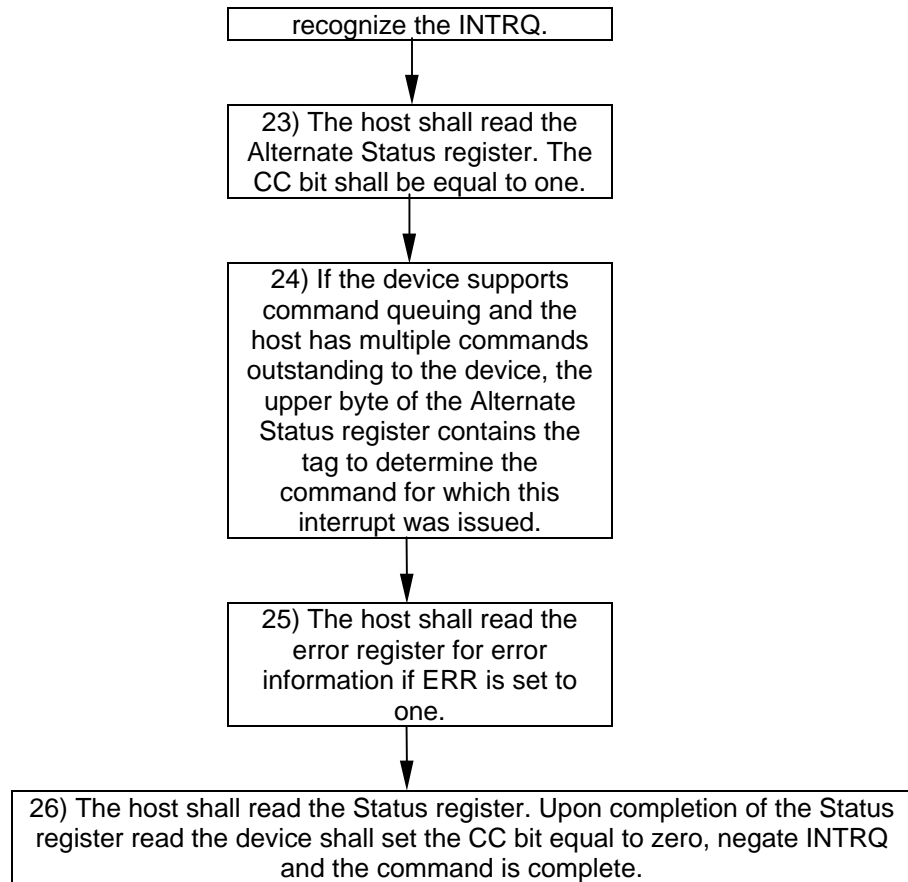
9.3.1 Overlap/DEVsel mode PIO data out commands

When a mixed configuration exists with one legacy device and one device that supports overlapped mode, the device that supports overlapped mode may be set into Overlap/DEVsel mode. The overlapped device then operates with the following protocol:



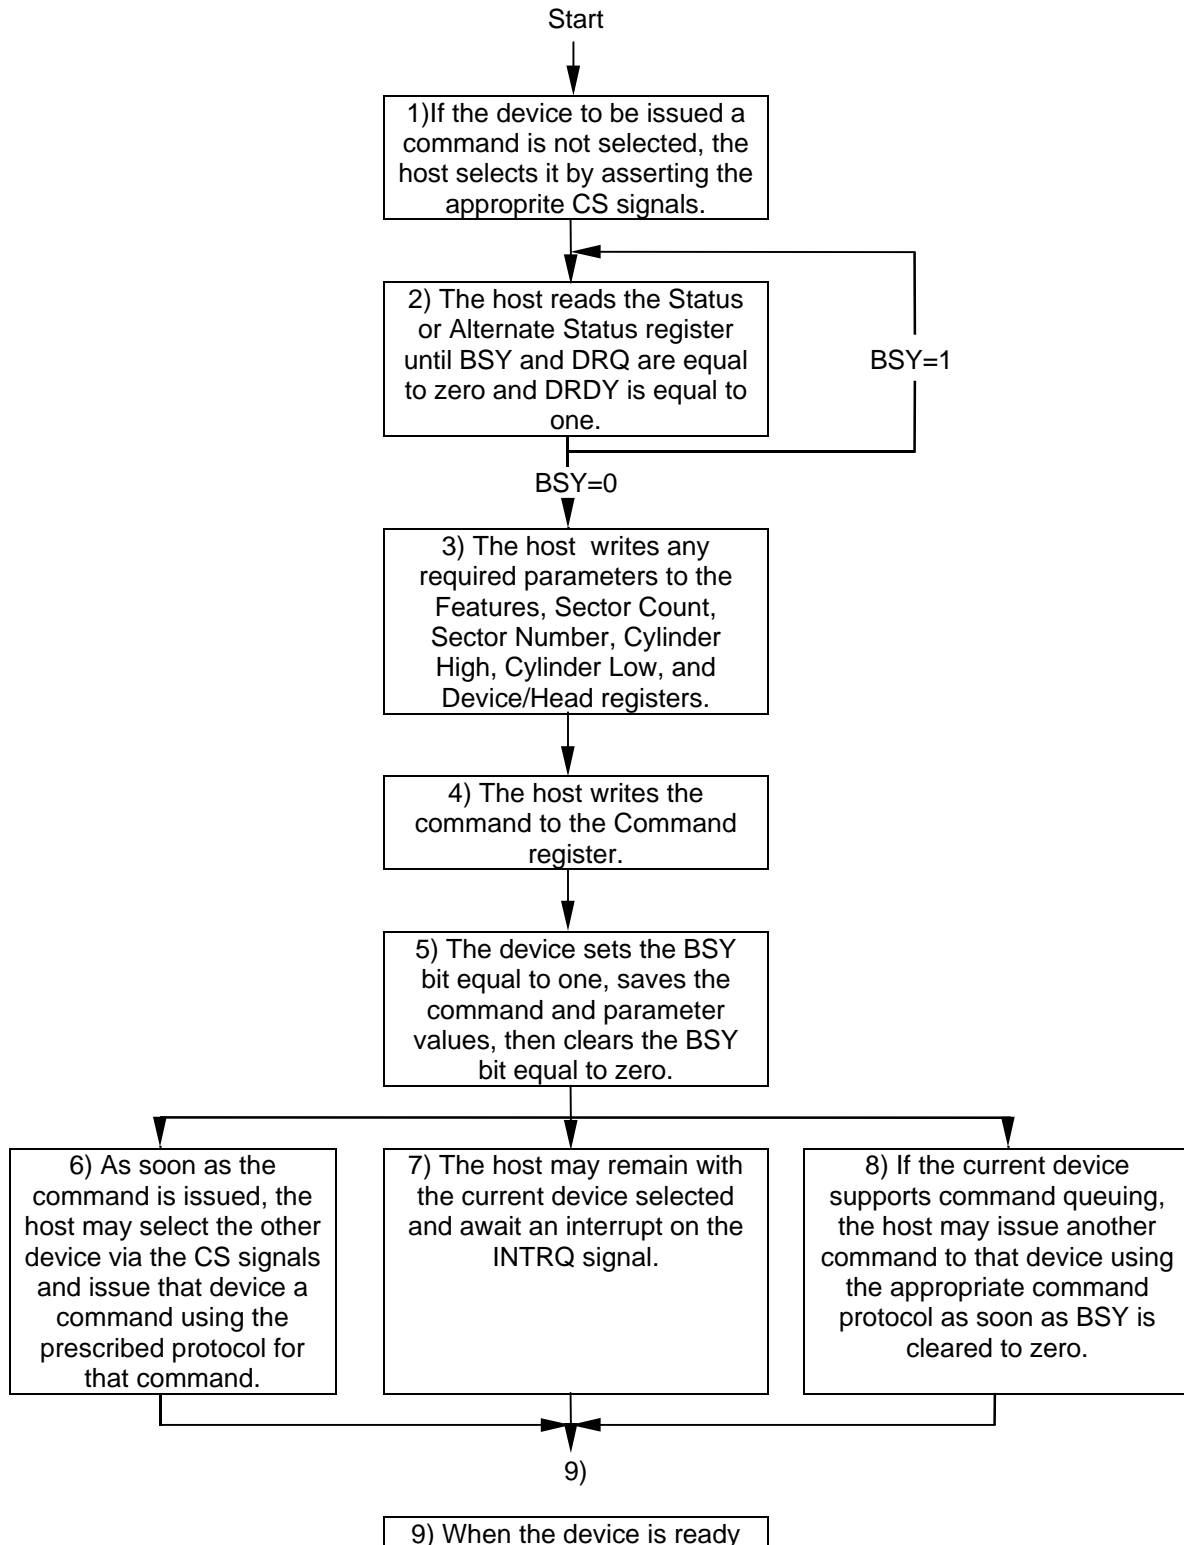


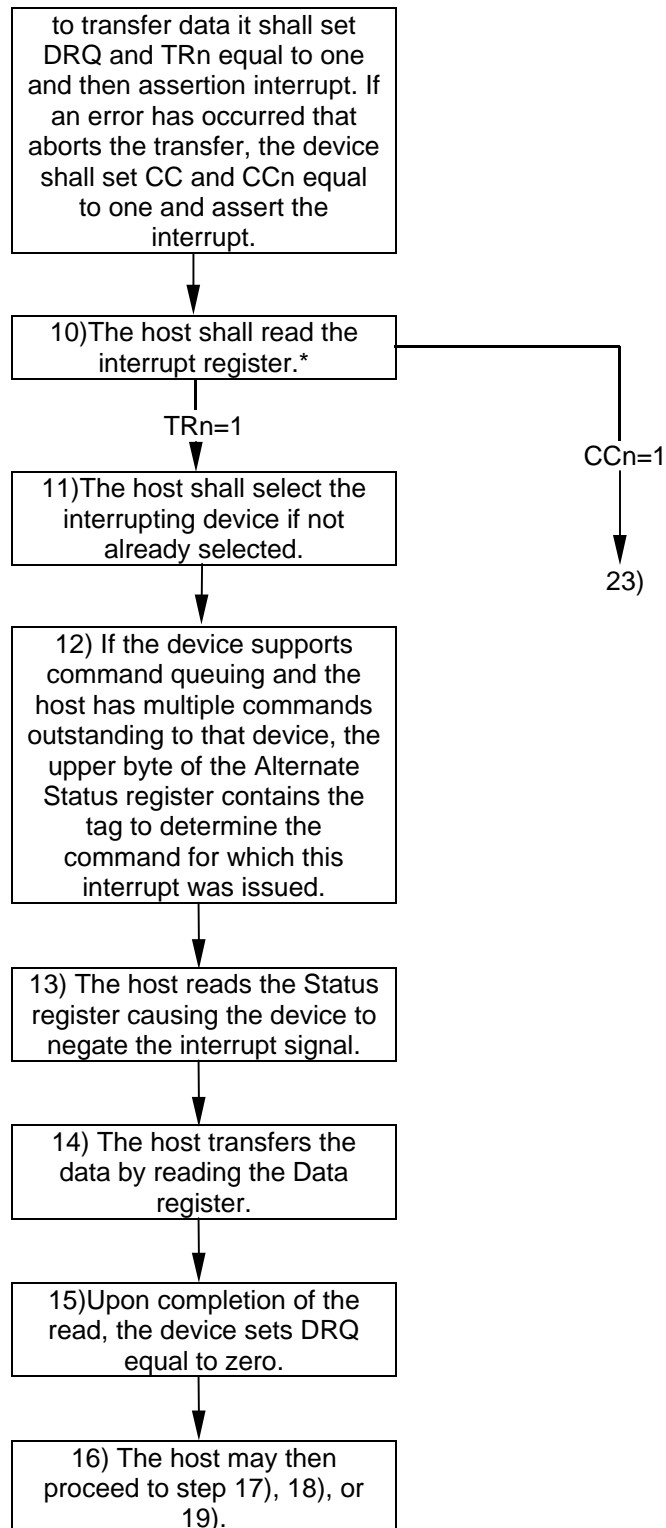




9.3.1 Overlap/CSsel mode PIO data out commands

When an overlapped configuration exists with both devices that support overlapped mode, the devices are set into Overlap/CSsel mode. In this configuration, device selection is accomplished via the CS signals. The device then operate with the following protocol:

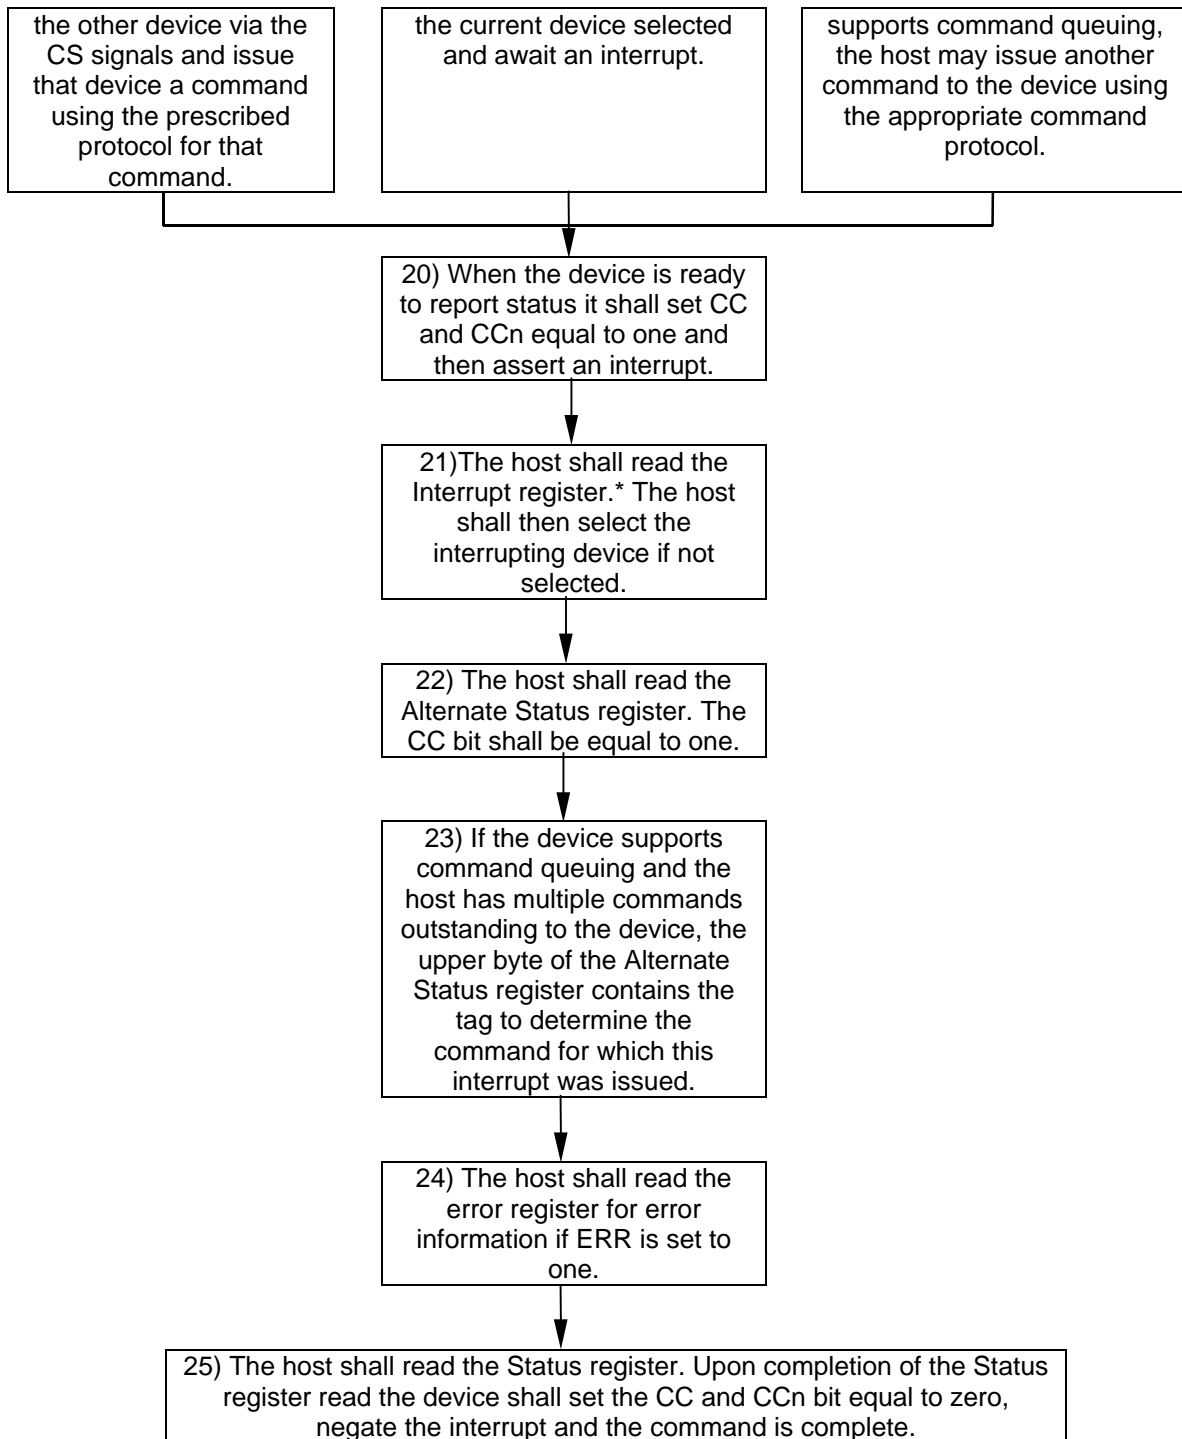




17) The host may select

18) The host may remain with

19) If the current device

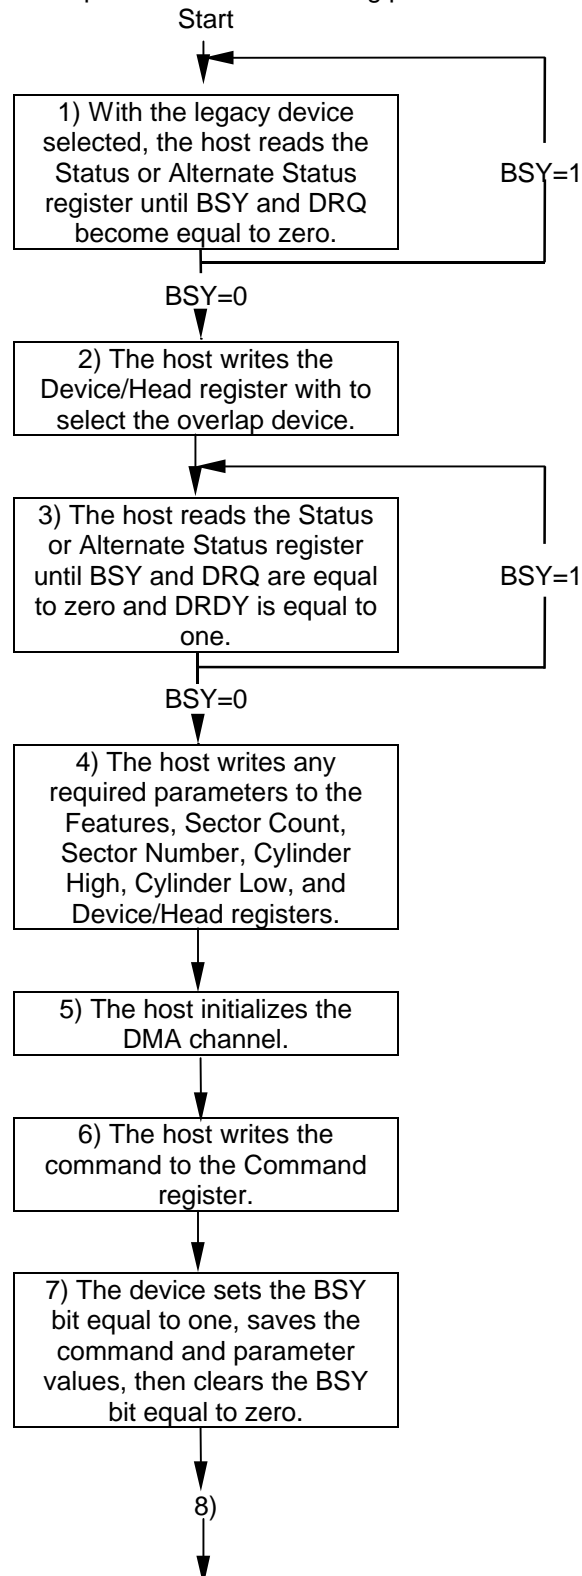


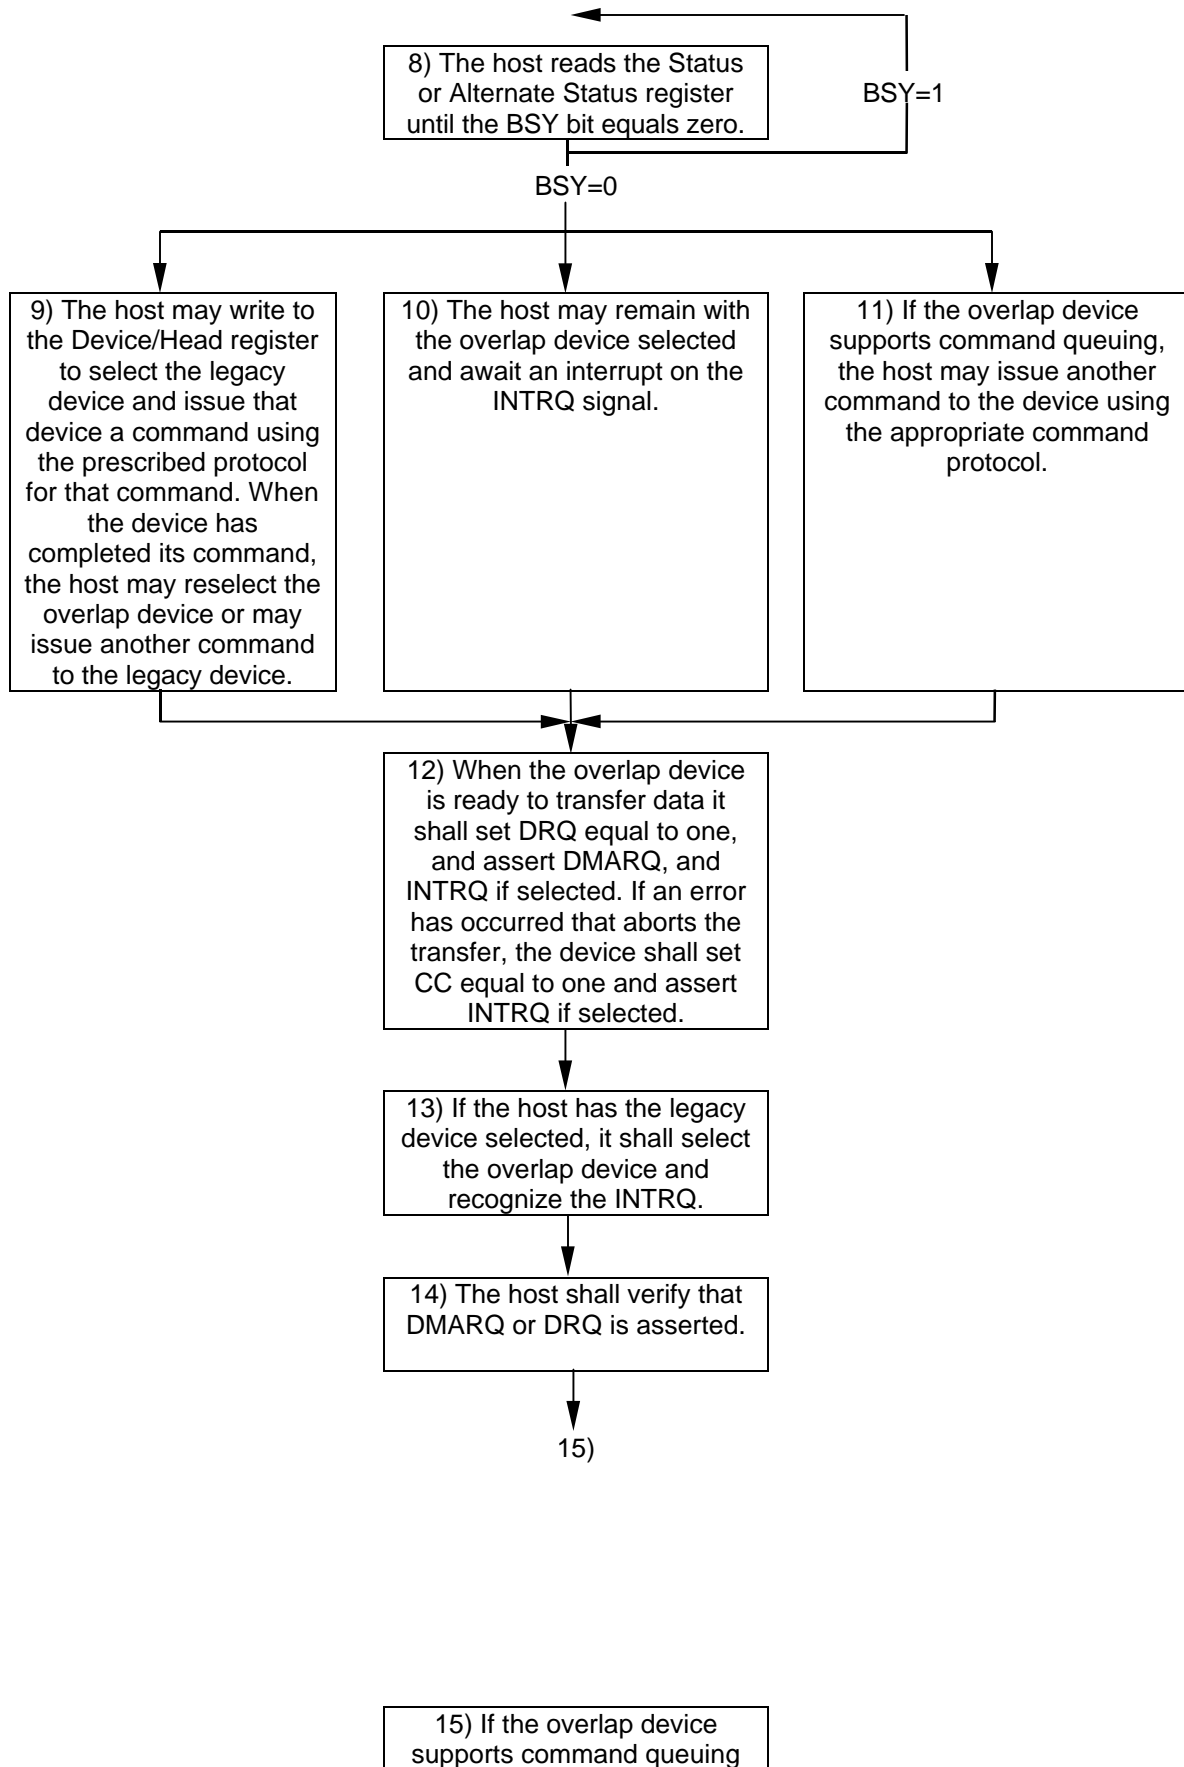
* Note: with today's 3F6/376h host addressing only an intelligent DMA engine may be capable of reading the Interrupt register. If the host is unable to read the Interrupt register, the status of each device will have to be read to determine the cause of the interrupt.

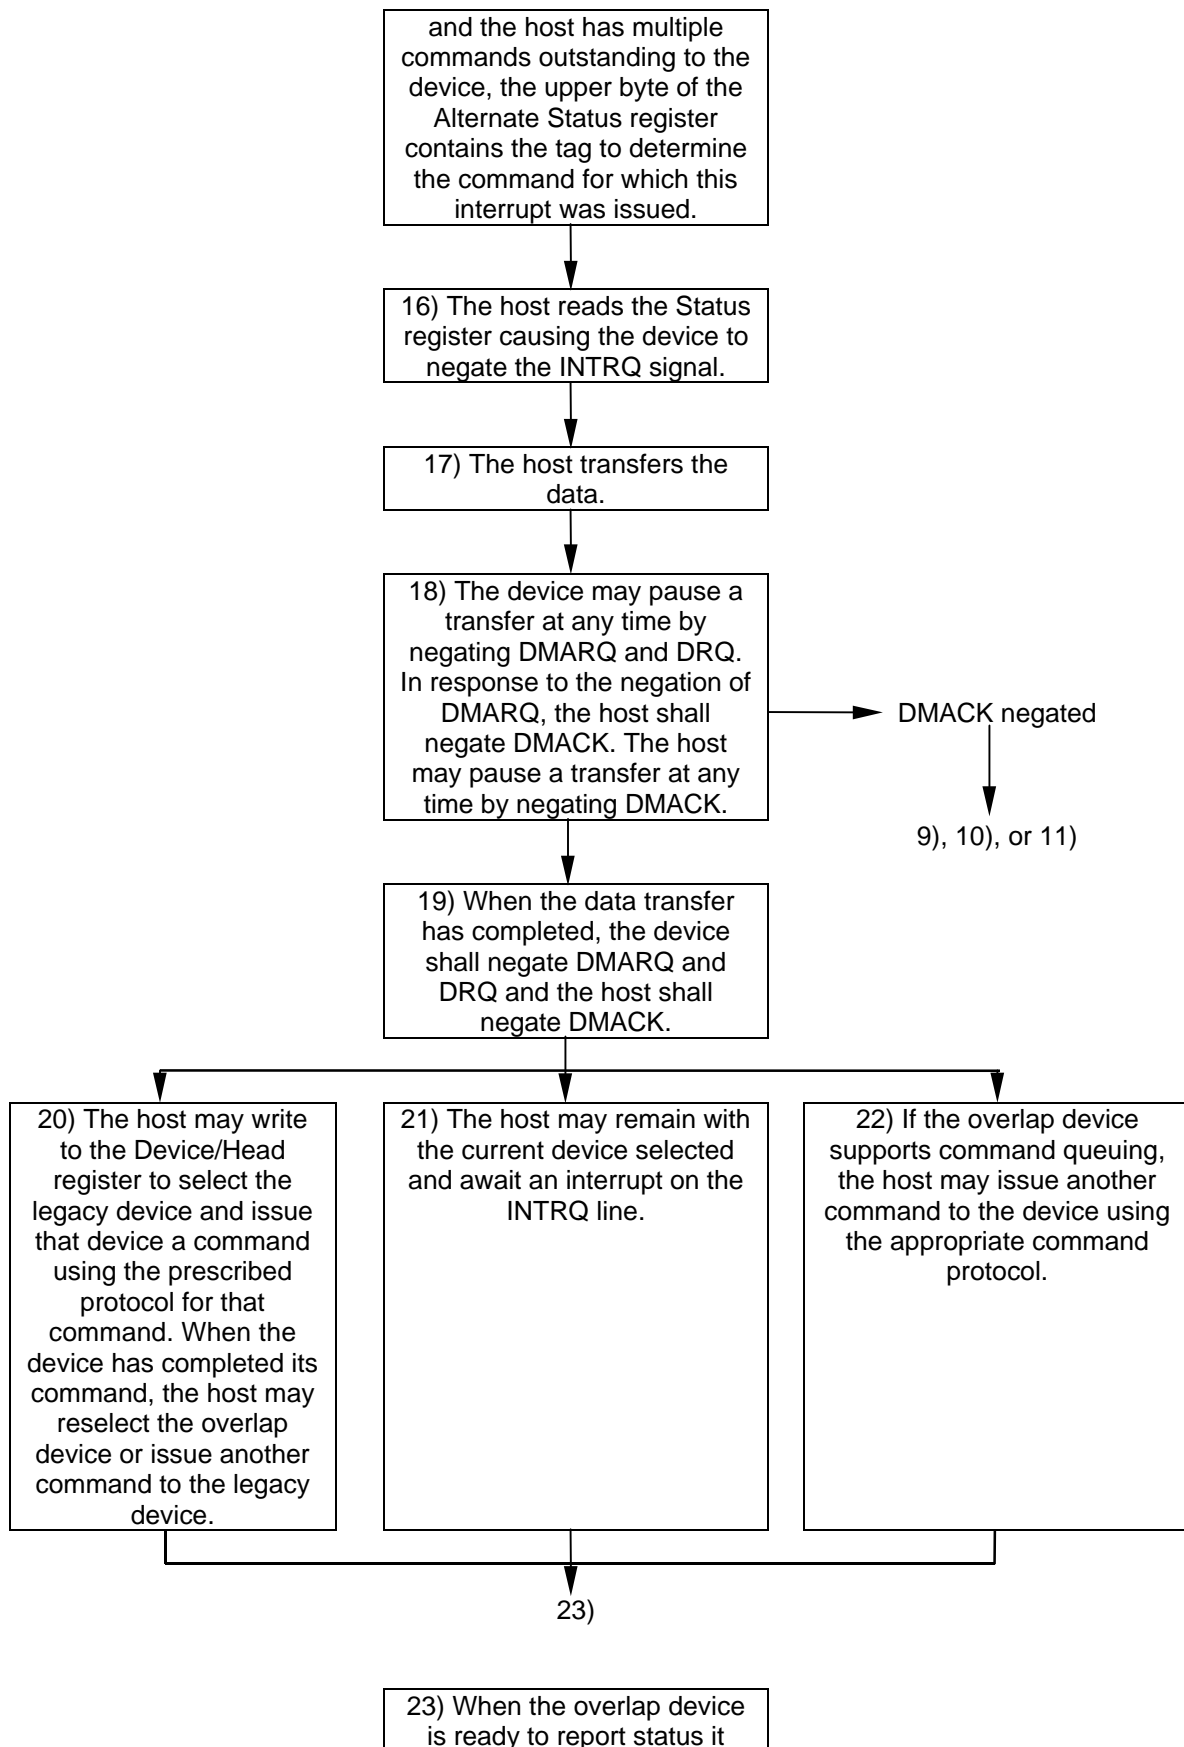
Clause 9.6 DMA data transfer commands - add:

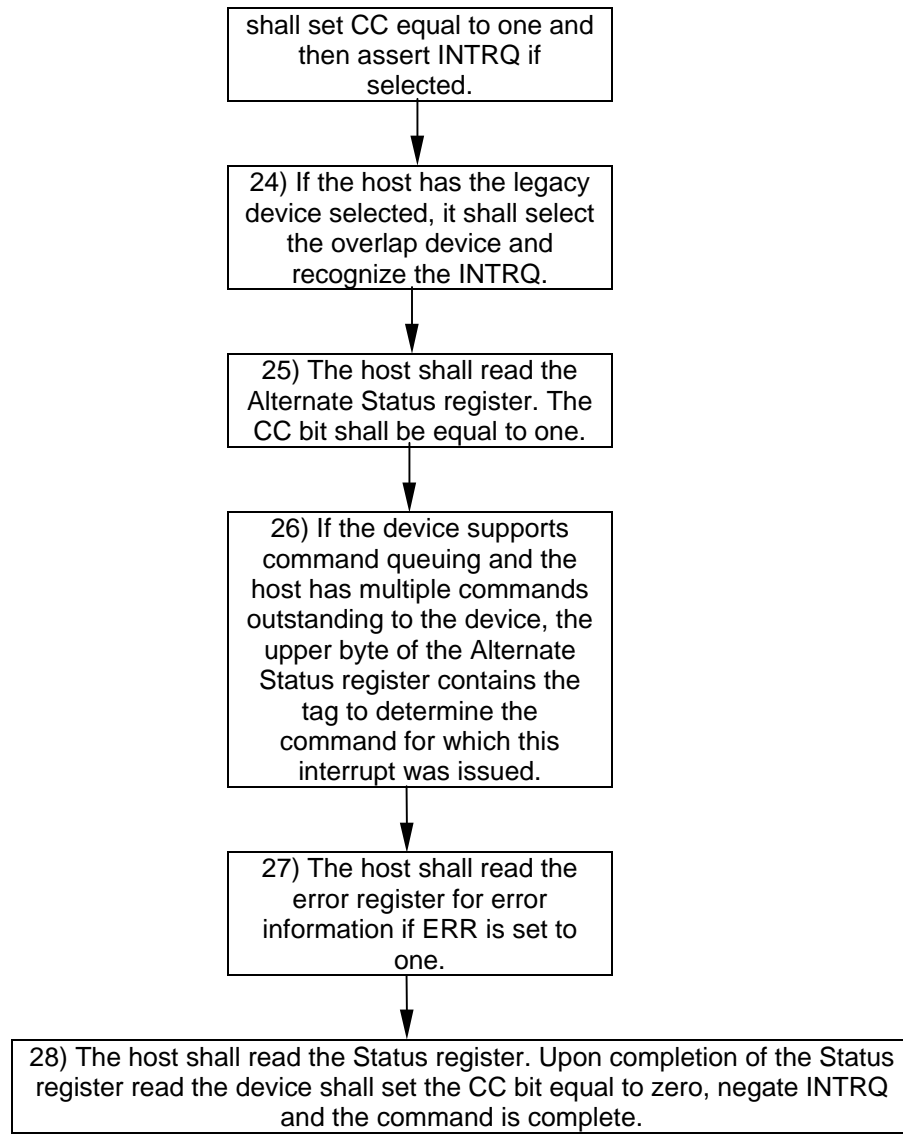
9.6.1 Overlap/DEVsel mode DMA data transfer commands

When a mixed configuration exists with one legacy device and one device that supports overlapped mode, the device that supports overlapped mode may be set into Overlap/DEVsel mode. The overlapped device then operates with the following protocol:



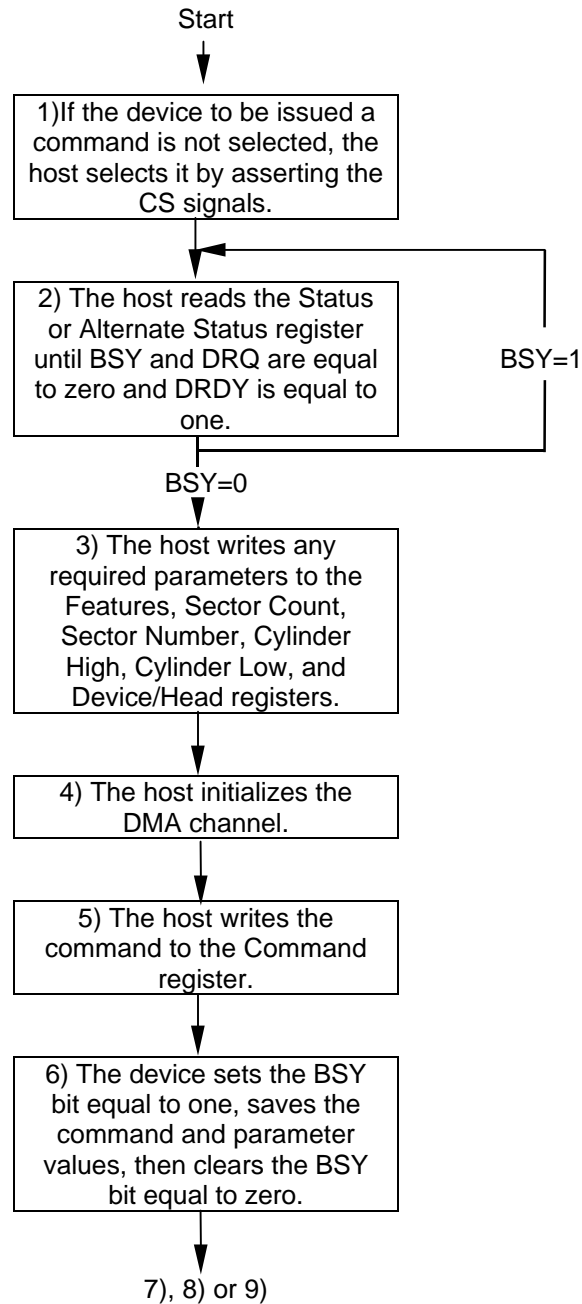


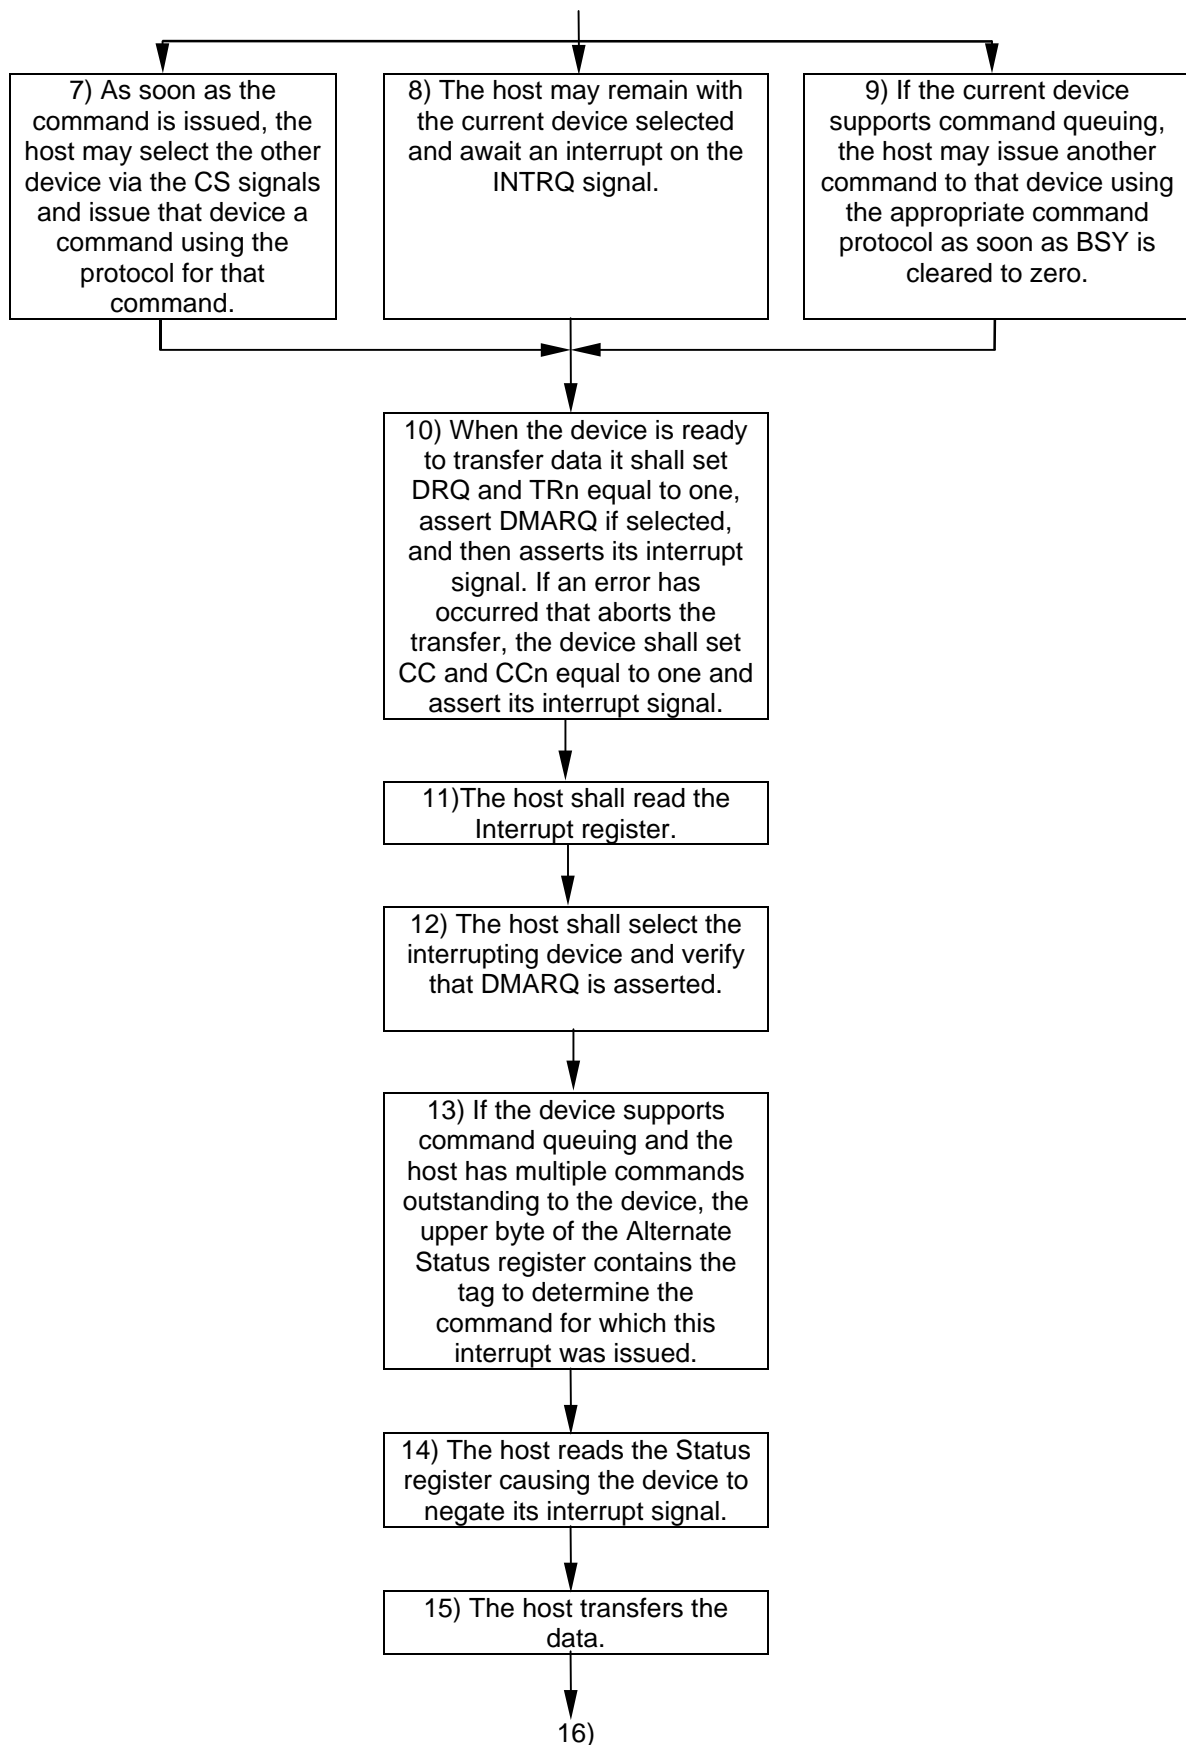


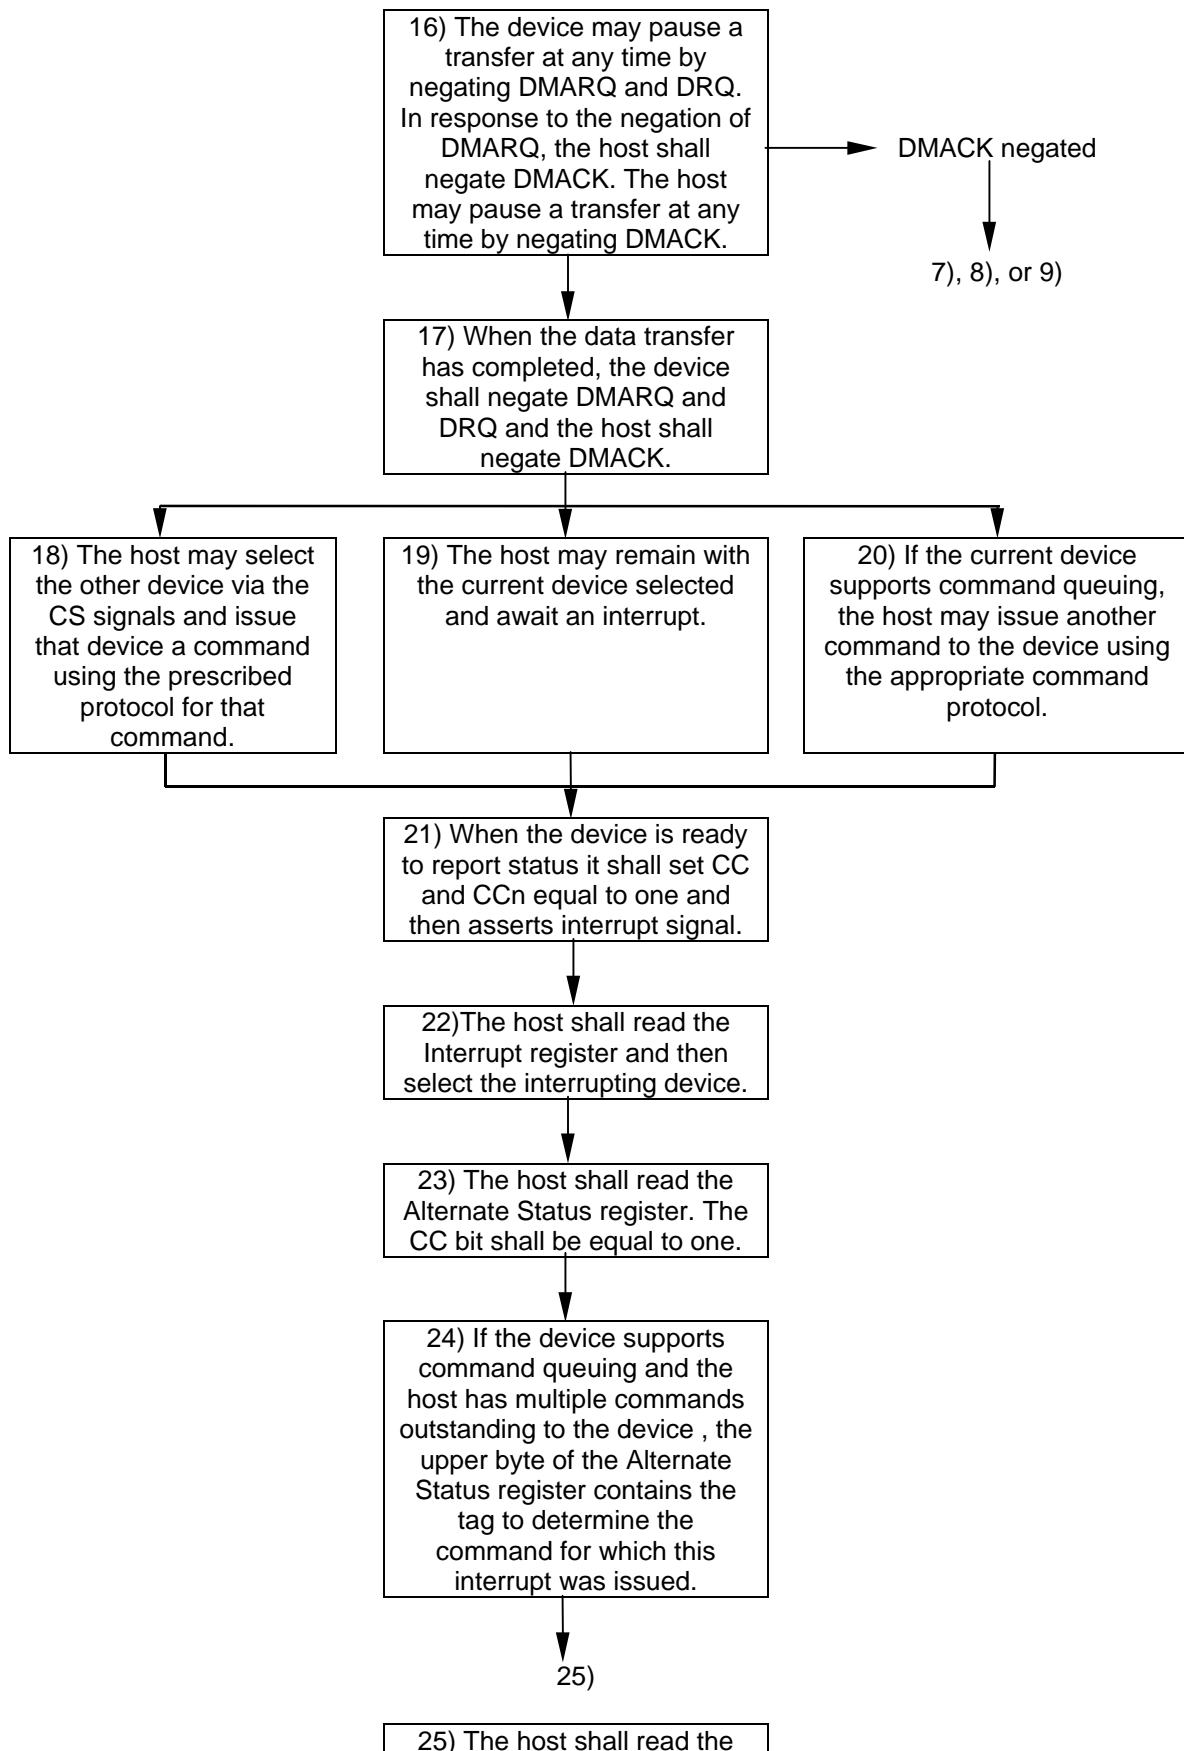


9.6.1 Overlap/CSsel mode DMA data transfer commands

When an overlapped configuration exists with both devices that support overlapped mode, the devices are set into Overlap/CSsel mode. In this configuration, device selection is accomplished via the CS signals. The device then operate with the following protocol:







error register for error
information if ERR is set to
one.



26) The host shall read the Status register. Upon completion of the Status register read the device shall set the CC and CCn bit equal to zero, negate its interrupt, and the command is complete.