

ATA Overlap Proposal

Point of contact:

Pete McLean
Maxtor Corporation
2190 Miller Drive
Longmont, CO 80501-6744
USA

Tel: 303-678-2149
Fax: 303-682-4811
Email: pete_mclean@maxtor.com

Scope

This document presents the detailed text changes to the ATA-3 document (X3T10/2008Dr2) for the Overlap proposal. It is based on the Overlap proposal submitted by Western Digital Corporation (X3T10/95-258r0) with the changes presented by Quantum Corporation (X3T10/95-257r0).

The text to be added to Clause 7 is presented first to provide an introduction to overlap and command queuing. The remainder of the document presents material in the order that it will appear in the ATA document.

Revision History

Revision 1 - Draft created from WD proposal X3T10/95-258r1

Revision 2 - Protocol diagram added

7.* Command overlap and command queuing

Command overlap is an optional feature that allows a host to issue commands or service outstanding commands with one device while a command is outstanding to the second device. For example, the host may issue a command to the first device and before the command is completed (i.e., requested data is transferred) , select the second device and issue the second device a command.

Command queuing is an optional feature that allows a host to have multiple commands outstanding to a given device. That is, before a first command has completed (i.e., requested data is transferred) additional commands may be issued to the device. A device that supports command queuing shall also support command overlap.

7.*.1 Command Overlap

Command overlap is accomplished by the use of the following mechanisms.

1. Information in the IDENTIFY DEVICE command response that notifies the host that the device supports the command overlap feature.
2. The SET FEATURES command that allows the host to enable the command overlap feature.
3. An overlap command set that identifies those commands that may be overlapped.
4. A mechanism called Release that allows the host to select the second device when the first device has an overlapped command outstanding.
5. A Proxy Interrupt mechanism that allows an unselected device to signal the host that it wishes to continue execution of an outstanding overlapped command.
6. A SERVICE command that allows the host to reselect a device that has an overlapped command outstanding so that command execution may be continued.

7.*.1.1 Overlap IDENTIFY DEVICE information

Word 49 in the IDENTIFY DEVICE response informs the host if the device supports Command Overlap and Command Queuing. Word 71 informs the host of the maximum time from the issuing of an overlap command to the negation of BSY during the Release process. Word 72 informs the host of the maximum time from the issuing of a SERVICE command to the negation of BSY during the Release process. See Clause *.* , IDENTIFY DEVICE.

7.*.1.2 Overlap SET FEATURES command

The issuing of a SET FEATURES command with the Features register content 5Fh or DFh, enables or disables the use of the Proxy interrupt mechanism and is mandatory for the use of the Overlap feature.

The issuing of a SET FEATURES command with the Features register content 5Dh or DDh, enables or disables the issuing of an interrupt when the Release process has completed (i.e., BSY negated) after the issuing of an overlapped command. If the time reported in the IDENTIFY DEVICE response word 71 is greater than 50 μ s this feature shall be enabled for overlap operation. This allows a host to determine whether it needs to or wishes to poll the BSY bit or whether it prefers an interrupt.

The issuing of a SET FEATURES command with the Features register content 5Eh or DEh, enables or disables the issuing of an interrupt when the Release process has completed (i.e., BSY negated) after the issuing of a SERVICE command. If the time reported in the IDENTIFY DEVICE response word 72 is greater than 50 μ s this feature shall be enabled for overlap operation. This allows a host to determine whether it needs to or wishes to poll the BSY bit or whether it prefers an interrupt when the BSY bit is negated.

See Clause *.* , SET FEATURES.

7.*.1.3 Overlap commands

The only commands that shall be overlapped are PACKET (see X3T10/1120D ATAPI), READ DMA OVERLAP, READ PIO OVERLAP, WRITE DMA OVERLAP and WRITE PIO OVERLAP. All other commands shall complete execution before both BSY and DRQ are negated. These five overlap commands shall be issued only when the Enable Proxy Interrupt SET FEATURE command has been executed.

The device shall execute the Release mechanism and issue Proxy interrupts only when one of these five commands is outstanding (i.e., from the issuing of the command to the Status register read at the completion of the command). The host shall only issue the SERVICE command when one of these five commands is outstanding to the device. The contents of the Features, Sector Count, Alternate Status, and Status registers are redefined for the issuing of these commands and remain redefined until the completion of these commands. See Clauses *.* , *.* , *.* , and *.* .

If an overlap command is issued while a none overlap command is id progress, both the command is progress and the newly issued command shall be aborted.

7.*.1.4 Release

Upon the receipt of an overlapped command, the device shall “save” the command and its required parameters, then Release by negating BSY. The maximum time from the receipt of the command to the negation of BSY shall be reported in the IDENTIFY DEVICE response word 71. The device shall issue an interrupt when BSY is negated if the 5Dh SET FEATURES has been issued to enable that interrupt.

If the device can transfer the data immediately (i.e., command is a write and the buffer space is available or the command is a read and the data is in the device buffer), the device may assert DRQ and DREQ before negating BSY and the transfer may be executed without an interrupt.

Once the Release process has completed, the host may select the other device if it desires.

Upon receipt of a SERVICE command while an overlap command is outstanding, the device shall set BSY, restore the register contents to reflect the outstanding overlapped command parameters, then negate BSY. The maximum time from the receipt of the command to the negation of BSY shall be reported in the IDENTIFY DEVICE response word 72. The device shall issue an interrupt when BSY is negated if the 5Eh SET FEATURES has been issued to enable that interrupt.

7.*.1.5 Proxy interrupt

The Proxy interrupt mechanism allow an unselected device to interrupt the host. The feature is enabled via the 5Fh SET FEATURES command and shall be enabled for overlapped operation in both devices.

When a device has an overlapped command outstanding, is not selected, but would like to interrupt the host to continue the processing of the overlapped command, the device shall assert the PDIAG signal.

When a device has the Proxy interrupt feature enabled via the 5Fh SET FEATURES command and the device is selected, the device shall monitor the PDIAG signal. If the PDIAG signal is asserted by the other device, the selected device shall assert the INTRQ signal on behalf of the non-selected device. If the selected device is in the process of a data transfer or other uninterruptable process, the selected device may postpone the assertion of INTRQ until the uninterruptable process has completed.

When the SERVICE command is issued, the previously selected device shall cease asserting INTRQ within ns. The newly selected device shall assert INTRQ and negate PDIAG no sooner than ns after the receipt of the SERVICE command but no later than ns.

7.*.1.6 Overlap SERVICE command

The SERVICE command is issued by the host to reselect a device that has an overlapped command outstanding. The SERVICE command shall only be issued when the device has an overlapped command outstanding. Upon receipt of the SERVICE command, the device shall assert BSY, restore its register content to reflect the overlapped command parameters, then negate BSY. If the 5Eh SET FEATURES command is enabled, the device shall assert INTRQ after BSY has been negated. This allows the host to check the status associated with the overlapped command and continue the execution of the command.

Editors note: Do we want to put one of Devon's diagrams in at this point to illustrate overlap?

7.*.2 Command queuing

Command queuing allows the host to issue multiple overlapped commands to the device. That is, additional overlapped commands may be issued to the device before previous overlapped commands are completed. To command queue, the device must have Proxy interrupts enabled via the SET FEATURES command. Only overlapped commands may be queued. If an overlapped command is outstanding when a non-overlapped command is received both commands and the queue will be aborted.

A queued command shall have a tag provided by the host in the Sector Count register to uniquely identify the command. When the device restores register parameters during the execution of the SERVICE command, this tag shall be restored so that the host can identify the command for which status is being presented. If a queued command is issued with a tag value that is identical to the tag value for a command already in the queue, the entire queue shall be aborted including the new command. If an error causes the queue to be aborted, the QA bit shall be set in the register.

*****Editors note: There is currently no way for a device to report the queue depth it supports. Must a device support the full queue depth possible with the tag field size?*****

3.1 Definitions and abbreviations (add)

Command overlap - Command overlap is an optional feature that allows a host to issue commands or service outstanding commands with one device while a command is outstanding to the second device.

Command queuing - Command queuing is an optional feature that allows a host to have multiple commands outstanding to a given device.

Proxy interrupt - The Proxy Interrupt mechanism that allows an unselected device to signal the host that it wishes to continue execution of an outstanding overlapped command.

Release - A mechanism that allows the host to select the second device when the first device has an overlapped command outstanding.

5.2.10 INTRQ (Device interrupt) (add)

When both devices attached to the interface are capable of overlapped operation and enabled for Proxy interrupts, the PDIAG- signal is used by an unselected device to request the selected device to assert INTRQ on its behalf.

When a device has an overlapped command outstanding, is not selected, but would like to interrupt the host to continue the processing of the overlapped command, the device shall assert the PDIAG- signal.

When a device has the Proxy interrupt feature enabled via the 5Fh SET FEATURES command and the device is selected, the device shall monitor the PDIAG- signal. If the PDIAG- signal is asserted by the other device, the selected device shall assert the INTRQ signal on behalf of the non-selected device. If the selected device is in the process of a data transfer or other uninterruptable process, the selected device may postpone the assertion of INTRQ until the uninterruptable process has completed.

When the SERVICE command is issued, the previously selected device shall cease asserting INTRQ within * ns. The newly selected device shall assert INTRQ and negate PDIAG- no sooner than * ns after the receipt of the SERVICE command but no later than * ns.

5.2.13 PDIAG- (Passed diagnostics) (add)

When both devices attached to the interface are capable of overlapped operation and enabled for Proxy interrupts, the PDIAG- signal is used by an unselected device to request the selected device to assert INTRQ on its behalf.

When a device has an overlapped command outstanding, is not selected, but would like to interrupt the host to continue the processing of the overlapped command, the device shall assert the PDIAG- signal.

When a device has the Proxy interrupt feature enabled via the 5Fh SET FEATURES command and the device is selected, the device shall monitor the PDIAG- signal. If the PDIAG- signal is asserted by the other device, the selected device shall assert the INTRQ signal on behalf of the non-selected device. If the selected device is in the process of a data transfer or other uninterruptable process, the selected device may postpone the assertion of INTRQ until the uninterruptable process has completed.

When the SERVICE command is issued, the previously selected device shall cease asserting INTRQ within * ns. The newly selected device shall assert INTRQ and negate PDIAG- no sooner than * ns after the receipt of the SERVICE command but no later than * ns.

6.2.8 Error register (add)

7	6	5	4	3	2	1	0
QA	UNC	MC	IDNF	MCR	ABRT	TK0NF	AMNF

Bit 7 - Queue aborted - for devices supporting command queuing, the assertion of QA indicates that an error has occurred that caused the queue to be aborted.

6.2.9 Features register (add)

FUNCTIONAL DESCRIPTION (Overlapped commands) - For overlapped commands, READ DMA OVERLAP, READ PIO OVERLAP, WRITE DMA OVERLAP and WRITE PIO OVERLAP, this register is used to provide the sector count. The operation is exactly the same as the Sector Count register for other data transfer commands.

FIELD/BIT DESCRIPTION (Overlapped commands) -

7	6	5	4	3	2	1	0
Sector Count							

6.2.10 Sector Count register (add)

FUNCTIONAL DESCRIPTION (Overlapped commands) - For overlapped commands, READ DMA OVERLAP, READ PIO OVERLAP, WRITE DMA OVERLAP and WRITE PIO OVERLAP, this register is used to provide the tag for queued commands when the overlapped commands are queued. If the device does not support command queuing, the tag value shall be set to zero.

FIELD/BIT DESCRIPTION -

7	6	5	4	3	2	1	0
Tag					REL	IO	CoD

Bits 7-3 - Tag identifier for queued command.

Bit 2 - Indicates device has Released the bus before completing the command.

Bit 1 - Indicates the direction of requested data transfer; one indicates transfer to the host, zero indicates transfer to the device.

Bit 0 - Zero indicates data to be transferred, one indicates command packet is to be transferred for PACKET command.

6.2.12 Status register (add)

FIELD/BIT DESCRIPTION (overlapped commands) -

7	6	5	4	3	2	1	0
BSY	DRDY	DMA RDY	SERV	DRQ	CORR	IDX	ERR

Bits 7 - BSY (Busy) - Upon the receipt of an overlapped command, the device shall "save" the command and its required parameters, then Release by negating BSY. The maximum time from the receipt of the command to the negation of BSY shall be reported in the IDENTIFY DEVICE response word 71. The device shall issue an interrupt when BSY is negated if the 5Dh SET FEATURES has been issued to enable that interrupt.

If the device can transfer the data immediately (i.e., command is a write and the buffer space is available or the command is a read and the data is in the device buffer), the device may assert DRQ and DREQ before negating BSY and the transfer may be executed without an interrupt.

Once the Release process has completed, the host may select the other device if it desires.

Upon receipt of a SERVICE command while an overlap command is outstanding, the device shall set BSY, restore the register contents to reflect the outstanding overlapped command parameters, then negate BSY. The maximum time from the receipt of the command to the negation of BSY

shall be reported in the IDENTIFY DEVICE response word 72. The device shall issue an interrupt when BSY is negated if the 5Eh SET FEATURES has been issued to enable that interrupt.

Bit 6 - DRQ - same as for other transfer commands.

Bit 5 - DMA RDY (DMA Ready) - This bit signals that the device is ready to start a DMA data transfer.

Bit 4 - SERV (Service) - This bit signals that the device is requesting service or interrupt. This bit is set when the interrupt is requested and not cleared until the Service command is issued.

Bits 3-0 - Same as for other commands.

8 Command descriptions (add)

In the second paragraph, add PIO data in overlapped, PIO data out overlapped, DMA data in overlapped and DMA data out overlapped to command types.

In paragraph 3, add:

DOI - DMA data in overlapped
 DOO - DMA data out overlapped
 POI - PIO data in overlapped
 POO - PIO data out overlapped

In Table 12 (add)

proto	Command	typ	Command code	Parameters used				
				FR	SC	SN	CY	DH
DOI	READ DMA OVERLAP	O	C7h	y	y	y	y	y
POI	READ PIO OVERLAP	O	24h	y	y	y	y	y
ND	SERVICE	O	A2h					D
DOO	WRITE DMA OVERLAP	O	CCh	y	y	y	y	y
POO	WRITE PIO OVERLAP	O	34h	y	y	y	y	y

8.10 IDENTIFY DEVICE (add)

Editors note: Bit 13 Word 49 conflicts with current Standby timer support bit. Words 71 and 72 conflict with the rev level added to the ATA-3 document.

Table 14 - Identify Device Information (add)

Word	F/V	
49		Capabilities
	F	15 Interleaved DMA supported
	F	14 Command queuing supported
	F	13 Overlap operation supported
71	F	Maximum time for Release after overlap command receipt
72	F	Maximum time for Release after SERVICE command receipt

8.10.* Interleaved DMA support

Bit 15 of word 49 is used to indicated that the device supports interleaved DMA data transfer for overlapped DMA commands.

8.10.* Command Queuing supported

Bit 14 of word 49 is used to indicated that the device supports command queuing for overlapped commands.

8.10.* Command overlap supported

Bit 13 of word 49 is used to indicated that the device supports command overlap operation.

8.* READ DMA OVERLAP

OPCODE - C7h

TYPE - Optional

PROTOCOL - DMA data in overlap

INPUTS - The Cylinder Low, Cylinder High, Device/Head and Sector Number registers specify the starting sector address to be read. The Features register specifies the number of sectors to be transferred. The Sector Count register contains the tag for this command if the device supports command queuing.

NORMAL OUTPUTS - The Sector Count register contains the tag for this command if the device supports command queuing.

ERROR OUTPUTS - The Sector Count register contains the tag for this command if the device supports command queuing. Aborted command if the command is not supported or if the device has not had Proxy interrupt enabled. Aborted command if the device supports command queuing and the tag is invalid. An unrecoverable error encountered during the execution of this command results in the termination of the command and the Command Block registers contain the sector where the first unrecoverable error occurred. If a queue existed, the unrecoverable error shall cause the queue to abort and the QA bit to be set.

PREREQUISITES - The host shall initialize the DMA channel.

DESCRIPTION - This command executes in a similar manner to a READ DMA except that the device shall execute the Release process upon receipt of the command and resumption of the command shall be executed upon receipt of the SERVICE command.

8.* READ PIO OVERLAP

OPCODE - 24h

TYPE - Optional

PROTOCOL - PIO data in overlap

INPUTS - The Cylinder Low, Cylinder High, Device/Head and Sector Number registers specify the starting sector address to be read. The Features register specifies the number of sectors to be transferred. The Sector Count register contains the tag for this command if the device supports command queuing.

NORMAL OUTPUTS - The Sector Count register contains the tag for this command if the device supports command queuing.

ERROR OUTPUTS - The Sector Count register contains the tag for this command if the device supports command queuing. Aborted command if the command is not supported or if the device has not had Proxy interrupt enabled. Aborted command if the device supports command queuing and the tag is invalid. An unrecoverable error encountered during the execution of this command results in the termination of the command and the Command Block registers contain the sector where the first unrecoverable error occurred. If a queue existed, the unrecoverable error shall cause the queue to abort and the QA bit to be set.

PREREQUISITES - A successful SET MULTIPLE MODE command shall precede a READ PIO OVERLAP command.

DESCRIPTION - This command executes in a similar manner to a READ MULTIPLE except that the device shall execute the Release process upon receipt of the command and resumption of the command shall be executed upon receipt of the SERVICE command.

8.* SERVICE

OPCODE - A2h

TYPE - Optional

PROTOCOL - Non-data command

INPUTS - None.

NORMAL OUTPUTS - The Sector Count register contains the interrupt reason and the tag for the command requiring service if the device supports command queuing. The Cylinder Low and Cylinder High registers contain the number of bytes remaining to be transferred for this command.

ERROR OUTPUTS - Aborted command if the device does not support the command. Aborted command if no overlap command is currently in progress.

PREREQUISITES - This command shall be issued only when there is an overlap command in progress.

DESCRIPTION - Upon receipt of this command, the device shall set BSY, place the parameters for the command requiring service in the appropriate registers, and then clear BSY.

8.32 SET FEATURES (add)

Table 19 - Set Features register definitions (add)

Value	
5Dh	Enable interrupt for Release after the receipt of an overlap command.
5Eh	Enable interrupt after the processing of a SERVICE command.
5Fh	Enable Proxy interrupt
DDh	Disable interrupt for Release after the receipt of an overlap command.
DEh	Disable interrupt after the processing of a SERVICE command.
DFh	Disable Proxy interrupt

8.* WRITE DMA OVERLAP

OPCODE - CCh

TYPE - Optional

PROTOCOL - DMA data out overlap

INPUTS - The Cylinder Low, Cylinder High, Device/Head and Sector Number registers specify the starting sector address to be read. The Features register specifies the number of sectors to be transferred. The Sector Count register contains the tag for this command if the device supports command queuing.

NORMAL OUTPUTS - The Sector Count register contains the tag for this command if the device supports command queuing.

ERROR OUTPUTS - The Sector Count register contains the tag for this command if the device supports command queuing. Aborted command if the command is not supported or if the device has not had Proxy interrupt enabled. Aborted command if the device supports command queuing and the tag is invalid. An unrecoverable error encountered during the execution of this command results in the termination of the command and the Command Block registers contain the sector where the first unrecoverable error occurred. If a queue existed, the unrecoverable error shall cause the queue to abort and the QA bit to be set.

PREREQUISITES - The host shall initialize the DMA channel.

DESCRIPTION - This command executes in a similar manner to a WRITE DMA except that the device shall execute the Release process upon receipt of the command and resumption of the command shall be executed upon receipt of the SERVICE command.

8.* WRITE PIO OVERLAP

OPCODE - 34h

TYPE - Optional

PROTOCOL - PIO data out overlap

INPUTS - The Cylinder Low, Cylinder High, Device/Head and Sector Number registers specify the starting sector address to be read. The Features register specifies the number of sectors to be transferred. The Sector Count register contains the tag for this command if the device supports command queuing.

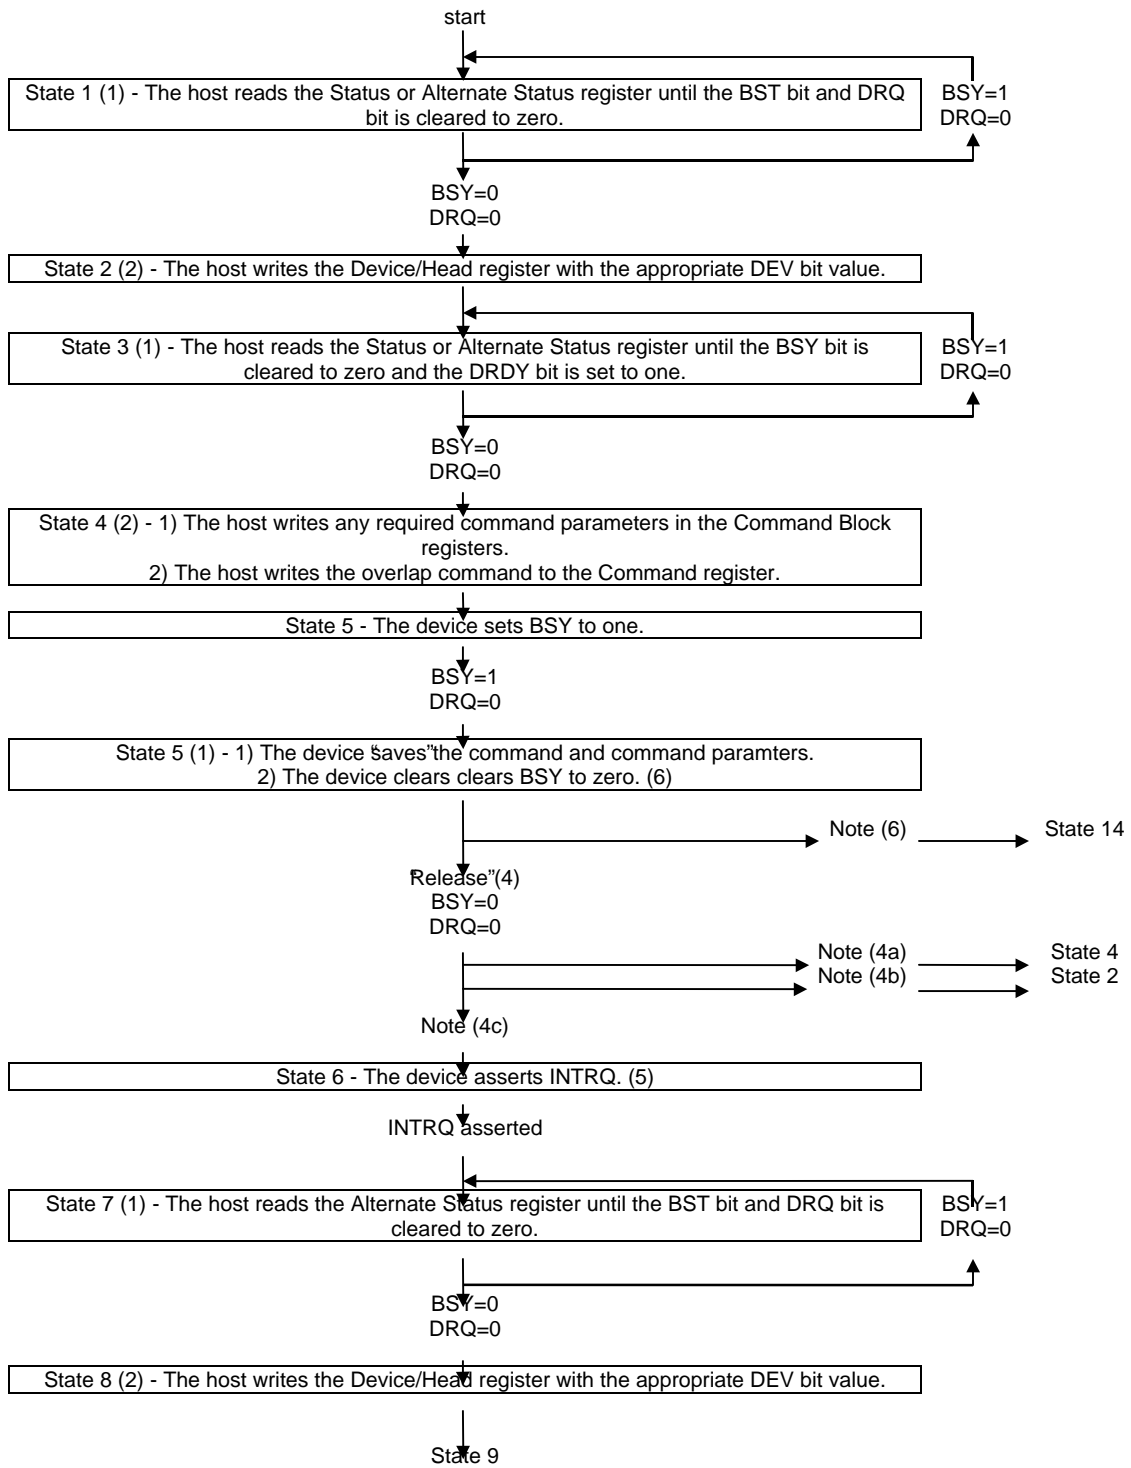
NORMAL OUTPUTS - The Sector Count register contains the tag for this command if the device supports command queuing.

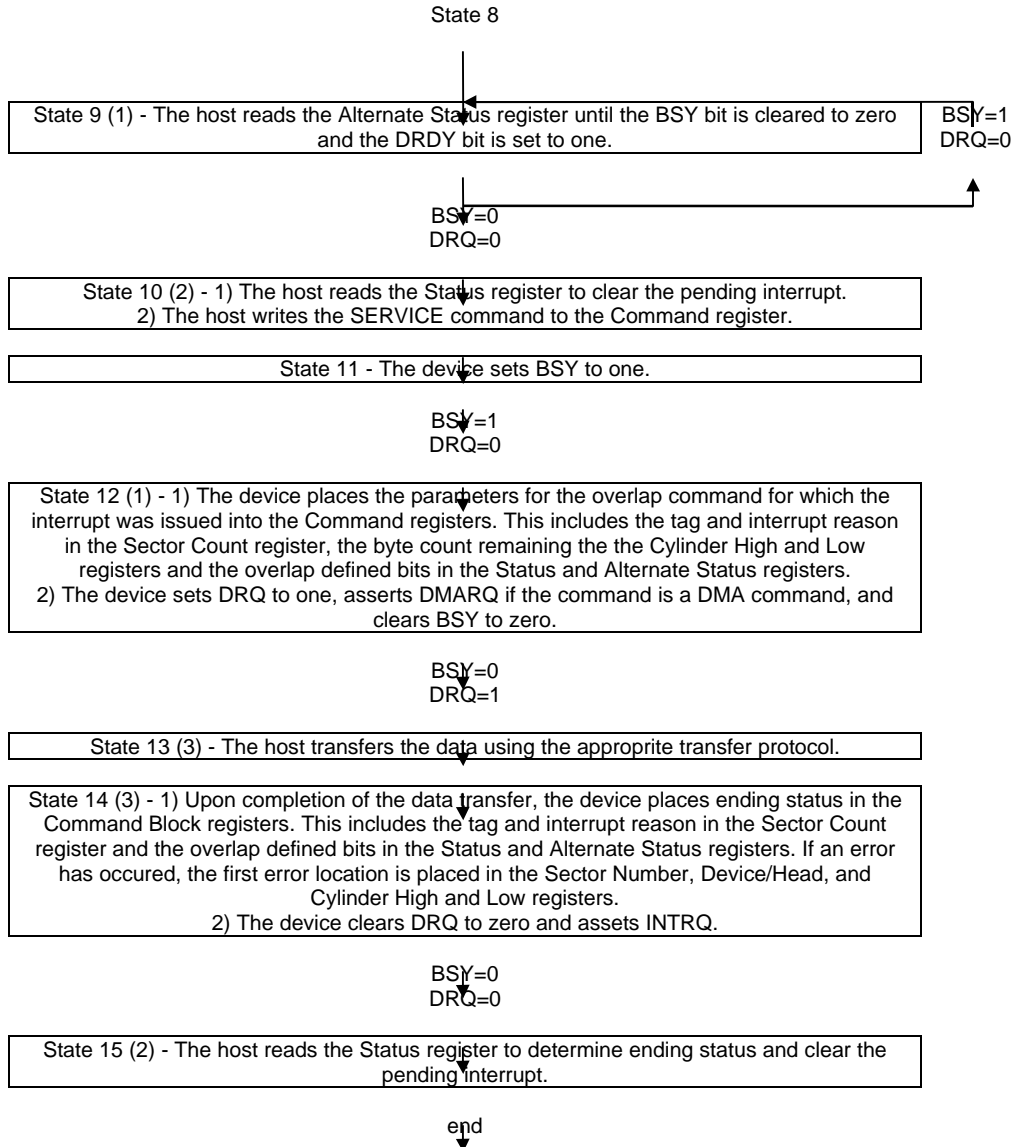
ERROR OUTPUTS - The Sector Count register contains the tag for this command if the device supports command queuing. Aborted command if the command is not supported or if the device has not had Proxy interrupt enabled. Aborted command if the device supports command queuing and the tag is invalid. An unrecoverable error encountered during the execution of this command results in the termination of the command and the Command Block registers contain the sector where the first unrecoverable error occurred. If a queue existed, the unrecoverable error shall cause the queue to abort and the QA bit to be set.

PREREQUISITES - A successful SET MULTIPLE MODE command shall precede a WRITE PIO OVERLAP command.

DESCRIPTION - This command executes in a similar manner to a WRITE MULTIPLE except that the device shall execute the Release process upon receipt of the command and resumption of the command shall be executed upon receipt of the SERVICE command.

9.* Overlap command protocol





Notes:

- (1) When BSY is set to one, the registers belong to the device. The the result of the host writing to registers is indeterminant. The state of the BSY and DRQ bit are valid when read by the host but the state of all other register bits is indeterminant.
- (2) When BSY and DRQ are cleared to zero, the registers belong to the host. They shall not be modified by the device except to set BSY when the Command register is written. The host may write the registers. When read by the host, the contents shall be the content written by the device before BSY or DRQ was last cleared if the host has not written the register, or the content last written to the register by the host since BSY or DRQ was cleared.
- (3) When DRQ is set to one, the registers belong to the device just as when BSY is set to one except that the host may access the Data register or DMA data port to transfer data.
- (4) When the device "releases" by clearing BSY, the host may:
 - a) Issue another overlap command to the device if the device supports command queuing.
 - b) Select the other device and issue it a command.
 - c) Wait for the INTRQ to transfer data.
- (5) If a device supports command queuing and has a new command issued when an interrupt is pending for an existing command, the device shall reassert INTRQ when it "releases" for the new command.
- (6) If an error makes the acceptance of the command impossible, the device goes directly to State 14.

Annex *

*****Editors Note: Should we create a new annex that includes Devon's diagrams for host behavior in overlap operation? *****