

To: SPI-3 Working Group
From: Richard Moore (rmoore@corp.adaptec.com)
Subject: Quick Arbitrate & Select (QAS) Proposal
Date: Thursday, April 30, 1998

1. Background

This proposal is for Quick Arbitrate and Select (QAS). It reduces overhead associated with parallel SCSI arbitration by allowing arbitration to take place without a Bus Free phase and by decreasing timing delays associated with arbitration. The fairness algorithm defined in SPI-2 is used to provide equitable access to the bus. The proposed protocol improvement is to be compatible and interoperable with the existing parallel SCSI protocol defined in SPI-2.

[Editor's note] As of rev 7 of this document, the Initiator Preemption for QAS has been removed. This is due to the fact that preemption increases the time required for arbitration and plays havoc on the fairness algorithm.

2. Terminology

QAS - Quick Arbitrate and Select

Normal arbitration - arbitration protocol sequence described in SPI-2, used following a Bus Free state

QAS capable device - a SCSI device that supports the QAS protocol

QAS enabled device - a QAS capable device that has negotiated the use of QAS with another device. Only QAS enabled devices may send or acknowledge a QAS message and participate in the QAS protocol.

QAS participating device - a SCSI device that is attempting to gain access to the SCSI bus by participating in the current QAS protocol

QAS message code - 0x55

QAS arbitration delay - Delay required before sampling SCSI bus signals to determine win or loss in the QAS protocol; this is equal to 1000 ns.

QAS release delay - Maximum time allowed for the release of certain signals during QAS; this is equal to 200 ns.

Bus Propagation delay - 200 ns maximum

INFORMATION UNIT TRANSFER REQUEST (IUTR) message (see T10/97-230)

3. Quick Arbitrate and Select (QAS) Protocol

The Quick Arbitrate and Select (QAS) protocol allows a QAS enabled target which is currently connected to a QAS enabled initiator to transfer control of the bus to another QAS enabled device that is ready to use it without an intervening Bus Free phase. QAS capability shall be reported in the INQUIRY data.

A QAS capable initiator shall negotiate the use of QAS with each QAS capable target for each power-on cycle, using the IUTR message, in order to enable QAS. QAS enabled devices shall implement the SPI-2 fairness algorithm in all normal and QAS arbitration. Devices must have negotiated the use of QAS with a particular device before using QAS to select or reselect that device. Also, targets must have negotiated the use of QAS with a particular initiator before using a QAS message to disconnect from that initiator, and initiators must have negotiated the use of QAS with a particular target before accepting a QAS message from that target. If an initiator receives a QAS message from a target with which it has not negotiated the use of QAS, then the initiator shall assert ATN prior to sending an ACK for the QAS message, and shall report MESSAGE REJECT on the following Message Out phase.

3.1 QAS Arbitration Phase

The QAS Arbitration phase begins following a Message In phase in which a QAS message is the last message transferred. The QAS message may occur following a DISCONNECT or COMMAND COMPLETE message, or following an INFORMATION UNIT phase.

Only QAS enabled devices ready to establish a connection to a device with which they have negotiated QAS, and allowed to arbitrate under the fairness algorithm, may participate in the QAS protocol. QAS protocol begins when the QAS participating devices on the bus detect either a 0x04 0x55 message sequence, or a 0x00 0x55 message sequence, or a transition from an Information Unit In or Information Unit Out phase to a single 0x55 message. The current target shall assert REQ for a minimum of 8 ns during the MESSAGE IN phase to ensure that all QAS

participating devices see the messages. Similarly, the current initiator shall assert ACK for a minimum of 8 ns in responding to the QAS message. The current target shall hold the QAS message byte valid for at least 33 ns after the detection of ACK asserted. Each QAS participating device must detect an ACK to the QAS message, with ATN de-asserted, prior to arbitrating through QAS.

The current target, after sending the QAS message and receiving an ACK with ATN de-asserted, shall release all SCSI signals except for BSY. The maximum time from the earliest signal release (other than REQ) to the latest signal release by the current target shall be one system deskew delay.

At least two deskew delays, but within a bus propagation delay, after detecting MSG, C/D, and I/O de-asserted following a successful QAS message all QAS participating devices shall place their SCSI IDs on the data bus.

[Editor's note] During QAS arbitration, only the current target may assert BSY. Because BSY is not wired-or in QAS, and devices do not have to detect Bus Free to participate in QAS, a Bus Free delay is not required prior to asserting SCSI IDs during QAS.

After waiting at least a QAS arbitration delay (measured from its detection of MSG, C/D, and I/O de-asserted), each QAS enabled device shall examine the data bus. The non-participating QAS enabled devices shall latch their fairness register. For participating QAS enabled devices, if a higher priority SCSI ID bit is true on the data bus (see SPI-2 Table 39 for the SCSI ID arbitration priorities), then the SCSI device has lost the arbitration and the SCSI device may release its ID bit from the data bus. Any losing device must release its SCSI ID bit within a QAS release delay after detecting the assertion of SEL. If no higher priority SCSI ID bit is true on the data bus, then the participating QAS SCSI device has won the arbitration and it shall assert the SEL signal.

[Editor's note] Because the uncertainty in detecting the start of QAS is less than the uncertainty in detecting the start of normal arbitration, the longer arbitration delay required for normal arbitration is not required for QAS arbitration.

After waiting at least a QAS arbitration delay from its release of MSG, C/D, and I/O, if there are no SCSI ID bits asserted then the current target shall transition to the BUS FREE phase. The subsequent arbitration is a normal arbitration. Otherwise, the current target shall wait until it detects SEL asserted, then wait at least one QAS release delay and release BSY.

The SCSI device that wins arbitration shall wait at least a QAS release delay plus two bus settle delays after asserting SEL before changing any signals.

As in normal arbitration, devices shall not actively negate any data or parity line during the QAS arbitration phase.

3.2 QAS Selection Phase

QAS Selection is used by an initiator to select a target for the purpose of initiating an operation. QAS Selection and QAS Reselection are distinguished by the I/O line, which is de-asserted during QAS Selection and asserted during QAS Reselection.

After winning QAS arbitration, and having delayed at least a QAS release delay plus two bus settle delays after asserting SEL, an initiator begins QAS Selection by setting the data bus to a value that is the OR of its SCSI ID bit, the target's SCSI ID, and the appropriate parity bit. It shall assert the ATN signal if a non-packetized connection is desired or if packetized protocol has not been negotiated with the target; it shall not assert the ATN signal if packetized protocol has been negotiated with the target and a packetized connection is desired. The initiator shall then wait at least a bus settle delay before looking for an assertion of the BSY signal from the target.

All QAS enabled targets other than the QAS arbitration winner shall examine the data bus after detecting SEL asserted. A target shall determine that it is selected when the SEL signal and its SCSI ID bit are true and the BSY and I/O signals are false for at least a bus settle delay. The selected target may examine the data bus in order to determine the SCSI ID of the selecting initiator. The selected target shall then assert the BSY signal within a selection abort time of its most recent detection of being selected; this is required for correct operation of the selection time-out procedure.

If the device detects bad parity during the selection, or if more than two SCSI ID bits are on the data bus during QAS selection, the device shall not respond to the QAS selection.

No less than two system deskew delays after the initiator detects the BSY signal true, it shall release the SEL signal and may change the data bus. The target shall wait until the SEL signal is false before asserting the REQ signal to enter an information transfer phase.

[Editor's note] The minimum time required for a QAS sequence is the time required for the QAS message exchange (typically less than 800 ns) plus QAS arbitration phase (1415 ns) and QAS selection phase (580 ns). This compares to normal arbitration, which requires a minimum of two

bus settle delays plus a bus free delay plus an arbitration delay plus a bus clear delay (totaling 4800 ns); and normal selection, which requires a minimum of four system deskew delays plus a bus settle delay (totaling 580 ns) for selection. All timings may be increased due to cable propagation delays.

See Table at end for details.

Selection time-out for QAS is identical to SPI-2 selection time-out.

3.3 QAS Reselection Phase

QAS Reselection is used by a target to reconnect to an initiator to continue an operation already in progress. QAS Selection and QAS Reselection are distinguished by the I/O line, which is de-asserted during QAS Selection and asserted during QAS Reselection.

After winning QAS arbitration, and having delayed at least a data release delay plus a bus settle delay after asserting SEL, a target begins QAS Reselection by asserting the I/O signal and setting the data bus to a value that is the OR of its SCSI ID bit, the initiator's SCSI ID, and the appropriate parity bit. The target shall then wait at least a bus settle delay before looking for an assertion of the BSY signal from the initiator.

All QAS enabled initiators other than the QAS arbitration winner shall examine the data bus after detecting SEL and I/O asserted. An initiator shall determine that it is reselected when the SEL and I/O signals and its SCSI ID bit are true and the BSY signal is false for at least a bus settle delay. The reselected initiator may examine the data bus in order to determine the SCSI ID of the reselecting target. The reselected initiator shall then assert the BSY signal within a selection abort time of its most recent detection of being reselected; this is required for correct operation of the selection time-out procedure.

If the device detects bad parity during QAS reselection, or if the number of SCSI ID bits asserted on the data bus is not equal to two during QAS reselection, the device shall not respond to the QAS reselection.

No less than two system deskew delays after the target detects the BSY signal true, it shall also assert the BSY signal and wait at least two system deskew delays and then release the SEL signal. The target may then change the I/O signal and the data bus. After the reselected initiator detects the SEL signal false, it shall release the BSY signal. The target shall continue asserting the BSY signal until it relinquishes the SCSI bus.

After releasing the SEL signal, the target may assert the REQ signal to enter an information transfer phase.

Reselection time-out for QAS is identical to SPI-2 reselection time-out.

3.4 Mixed Bus Blocking Relief

In a mixed environment of QAS and non-QAS devices, it is possible for QAS devices to use the bus for extended time periods excluding the non-QAS devices. Implementing the following recommendations may provide a remedy.

QAS targets, when disconnecting following an initial selection, should disconnect via a BUS FREE if their fairness registers are empty. This guarantees a BUS FREE following the selection of the lowest-priority QAS target. This recommendation does not apply to targets disconnecting following reselection, or to targets whose fairness registers are not empty.

In a mixed environment non-QAS devices should disable their fairness algorithm, and be assigned higher priority SCSI IDs than QAS devices.

Current target	Current initiator	QAS enabled targets	Reselected device	QAS enabled initiators	Selected device
MSG, C/D, I/O and BSY asserted (Message In phase)	Must not be driving ATN at this point	Snoop for QAS message		Snoop for QAS message	
Assert 0x55 on data bus, then assert REQ	Detect QAS message and REQ asserted, then assert ACK				
Detect ACK asserted, then negate REQ	Detect REQ negated, then negate ACK	Detect QAS message and ACK asserted without ATN		Detect QAS message and ACK asserted without ATN	
Detect ACK negated, then release all SCSI signals except BSY	Disconnected	Detect phase lines negated, wait at least two deskew delays but no more than a bus propagation delay, and then assert ID bit		Detect phase lines negated, wait at least two deskew delays but no more than a bus propagation delay, and then assert ID bit	
One QAS arbitration delay after releasing, sample ID bits; release BSY if none asserted, else wait until SEL asserted		Wait one QAS arbitration delay after detecting phase lines negated, and then sample the data bus. Highest priority device present wins arbitration. Losers drop out.		Wait one QAS arbitration delay after detecting phase lines negated, and then sample the data bus. Highest priority device present wins arbitration. Losers drop out.	
		Winner asserts SEL, waits one QAS release delay plus two bus settle delays, then asserts its ID, the device ID to be re-selected, correct parity, and I/O.		Winner asserts SEL, waits one QAS release delay plus two bus settle delays, then asserts its ID, the device ID to be re-selected, and correct parity.	
Detect SEL asserted, wait one QAS release delay, and release BSY		Detect BSY de-asserted. Wait one bus settle delay from driving the reselected device's ID line.	Detect reselection conditions true for at least one bus settle delay, then assert BSY	Detect BSY de-asserted. Wait one bus settle delay from driving the selected device's ID line.	Detect selection conditions true for at least one bus settle delay, then assert BSY
		Detect BSY asserted, wait two system deskew delays and assert BSY, then wait two system deskew delays and release SEL	Detect SEL de-asserted and release BSY	Detect BSY asserted, then wait two system deskew delays and release SEL	Detect SEL de-asserted
		Connection established	Connection established	Connection established	Connection established