To:	Distribution	**Date:**	October 27, 1999
From:	Charles Monia/Larry Lamers		
Prepared by:	Charles Monia		
Subject:	Proposal for Increasing the Addressability of Parallel SCSI		
Document Number:	T10/99-250~~XX/XX-XXX~~R2		
Reply:	ljlamers@ieee.org, cmonia@corp.adaptec.com		

1 Overview

Recent improvements in the SCSI electrical layer combined with low cost bus extenders, increases in the synchronous data rate and a more efficient protocol have increased the load carrying capacity of the bus far beyond the combined workload presented by the maximum number of connectable devices now allowed. The ability to connect more devices therefore represents an opportunity to improve systems performance at low cost by more fully utilizing latent bus capacity.

With that in mind, this paper discusses protocol extensions that increase the addressability of the parallel SCSI bus. These protocol modifications allow extended and legacy devices to interoperate with some restrictions. The intermixed operational mode is discussed in section 8.

This proposal describes new methods for arbitration, quick arbitration and device selection. Using these methods, a 16-bit parallel SCSI bus can accommodate a maximum of 64 extended and eight legacy devices. These extensions are made practical by the Ultra160m definition, which establishes a design center based on Low Voltage Differential signaling (LVD) and a 16-bit data path. They allow the superior performance and signaling capabilities of this bus type to be fully exploited.

2 Reference Documents

SCSI Parallel Interface -3 (SPI-3), Revision 7, dated 3-June-1999 (T10 project 1302D), referred to in this document as SPI-3.

3 Definitions, symbols, abbreviations and conventions

3.1 Definitions

1.1.1 Extended SCSI Device: An SCSI device connected to the bus that is using the protocol for extended addressing described in this document.

1.1.2 Fairness rotation: The set of arbitration cycles during which each device in the deferring state becomes eligible to compete for the SCSI bus according to the rules for arbitration fairness.

1.1.3 **Group arbitration cycle:** The portion of an extended arbitration or Quick Arbitrate and Select operation during which the winning group is determined.

1.1.4 Legacy SCSI Device: An SCSI device connected to the bus that is not using the protocol for extended addressing described in this document.

1.1.5 **Member arbitration cycle:** The portion of an extended arbitration or Quick Arbitrate and Select operation during which extended devices in the winning group determine whether or not they have won arbitration.

1.1.6 SCSI I/D: A device identifier including both the group and member components.

3.2 Conventions

An extended device identifier is specified as (gid, mid). Where gid is a group identifier in the range 7 - 0 and mid is a member ID component in the range 15 - 8. Since legacy devices have no member ID component, a legacy device is specified as (gid, 0).

4 Extended SCSI Addressing

An extended SCSI address consists of two components – a group I/D (GID) and group member I/D (MID). The format and associated priority of each address element is shown in Table 1 and Table 2. As described in section 5, each component forms the basis for a two-cycle arbitration scheme in which each address element is asserted on the data bus during a given cycle and tested by contending devices to determine the ultimate winner.

An extended device has a GID in the range of 0 through 7 and a MID in the range 8 through 15. The GID/MID combination must be unique. Since eight groups are allocated and each group may have up to eight members, up to 64 extended devices may be attached to the bus.

A legacy device address consists of a group address in the range of 0 through 15. Legacy devices do not have a MID component. As discussed in section 8, legacy and extended devices may not share a GID. Consequently, a legacy device with a GID in the range 0 through 7 will be the sole user of a GID that could have been shared by up to eight extended devices.

Table 1 -- Group I/D Arbitration Priority

Group I/D	DB 15	Legacy devices only							DB 8	DB 7	Legacy or extended devices						DB 0	Priority
7	-	-	-	-	-	-	-	-	1	-	-	-	-	-	-	-	1
6	-	-	-	-	-	-	-	-	-	1	-	-	-	-	-	-	2
5	-	-	-	-	-	-	-	-	-	-	1	-	-	-	-	-	3
4	-	-	-	-	-	-	-	-	-	-	-	1	-	-	-	-	4
3	-	-	-	-	-	-	-	-	-	-	-	-	1	-	-	-	5
2	-	-	-	-	-	-	-	-	-	-	-	-	-	1	-	-	6
1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	1	-	7
0	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	1	8
15	1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	9
14	-	1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	10
13	--	-	1	-	-	-	-	-	-	-	-	-	-	-	-	-	11
12	--	-	-	1	-	-	-	-	-	-	-	-	-	-	-	-	12
11	--	-	-	-	1	-	-	-	-	-	-	-	-	-	-	-	13
10	-	-	-	-	-	1	-	-	-	-	-	-	-	-	-	-	14
9	-	-	-	-	-	-	1	-	-	-	-	-	-	-	-	-	15
8	--	-	-	-	-	-	-	1	-	-	-	-	-	-	-	-	16

Table 2 – Member I/D Arbitration Priority

Member I/D	DB 15	Extended devices only						DB 8	DB 7	Unused						DB 0	Priority
15	1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	1
14	-	1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	2
13	--	-	1	-	-	-	-	-	-	-	-	-	-	-	-	-	3
12	--	-	-	1	-	-	-	-	-	-	-	-	-	-	-	-	4
11	--	-	-	-	1	-	-	-	-	-	-	-	-	-	-	-	5
10	-	-	-	-	-	1	-	-	-	-	-	-	-	-	-	-	6
9	-	-	-	-	-	-	1	-	-	-	-	-	-	-	-	-	7
8	--	-	-	-	-	-	-	1	-	-	-	-	-	-	-	-	8

5 Extended Addressing Protocol

The methods described below define a two-round elimination scheme for arbitration and quick arbitration. During the first round (the group arbitration cycle), each contesting device asserts its GID and tests to see if it is still in contention. Extended devices sharing the highest priority GID on the first round, assert their MIDs and enter a second round of competition (the member arbitration cycle). The winner is the device with the highest priority MID.

Legacy devices only compete during the group arbitration cycle. If a legacy device wins, it begins selection as described in section 8; otherwise it drops out of contention.

The following sections describe the processes in detail.

5.1 Overview of Extended Arbitration

After detecting a bus free condition (SEL and BSY false), devices wishing to arbitrate begin by asserting their group IDs. During the group arbitration cycle all arbitrating devices sample the data bus to determine whether or not they have lost~~.~~

~~Devices still in contention assert SEL and BSY and continue to assert their SCSI IDs.~~ All losing devices, including all lower priority legacy devices, are required to drop out of contention.

If a legacy device wins during the group arbitration cycle, it performs selection as described in the SPI-3 specification. If one or more extended devices win, they begin the member arbitration cycle by asserting SEL and their MIDs. After a delay, devices still in contention sample the bus once again. This time, the device with the highest priority ~~group~~ member ID wins and asserts C/D. As in the group arbitration cycle, the remaining devices must determine that they have lost and stop driving all signals.

5.2 Extended Arbitration -- Details

A device arbitrates for the bus as described below. The timing sequence is shown in Figure 1~~Figure 1~~.

a) The SCSI device shall first wait for the BUS FREE phase to occur. The BUS FREE phase is detected whenever both the BSY and SEL signals are simultaneously and continuously false for a minimum of a bus settle delay.

b) The SCSI device shall wait a minimum of a bus free delay after detection of the BUS FREE phase (i.e. after the BSY and SEL signals are both false for a bus settle delay) before driving any signal.

c) Following the bus free delay in step (b), the SCSI device may arbitrate for the SCSI bus by asserting the BSY signal and its own group I/D. However the SCSI device shall not arbitrate (i.e. assert the BSY signal and its GID) if more than a bus set delay has passed since the BUS FREE phase was last observed.

d) ~~After waiting at least an arbitration delay (measured from its assertion of the BSY signal)~~ The SCSI device shall examine the DATA BUS when one of the following occurs:

A) An arbitration delay has elapsed (measured from the device's assertion of the BSY signal) or
B) SEL is asserted by another device~~.~~

whichever occurs first. The device has lost arbitration if a higher priority group I/D is present on the bus.

e) An extended SCSI device that has not lost ~~been eliminated~~ during the group arbitration cycle shall assert the SEL signal and its member ID within a QA Assertion delay of the event in step (d). The device shall continue to assert BSY and its group ID.

f) A legacy SCSI device that has not lost during the group arbitration cycle has won arbitration. The device shall assert SEL and enter the selection phase as described in the SPI-3 specification.

g) An SCSI device that has lost during the group arbitration cycle shall release the BSY signal and its SCSI ID within a bus clear delay after the SEL signal becomes true. The device may return to step (a).

h) After waiting at least a bus clear plus twice the bus settle delay ~~plus a bus settle delay~~ from its assertion of SEL, an extended device still in contention shall examine the data bus. If no higher priority member ID is present on the DATA BUS, then the SCSI device has won the arbitration.

~~g)~~i) An extended SCSI device that has won arbitration shall assert the C/D signal. After a bus settle delay following the assertion of C/D, the device shall negate the C/D signal. The device shall release (stop driving) C/D within one QAS release delay of its being negated.~~After waiting at least a bus settle delay from the assertion of C/D, the device that has won arbitration shall release C/D.~~ After waiting at least a QAS Release Delay ~~Bus Clear Delay~~ plus twice the~~a~~ bus settle delay from its assertion of C/D, the device may begin the extended selection phase.

j) An extended SCSI device that has lost the arbitration during the member arbitration cycle shall release all signals after two deskew delays and within one QAS Release Delay ~~Bus Clear Delay~~ after the C/D signal becomes true. An SCSI device that loses arbitration may return to step (a).

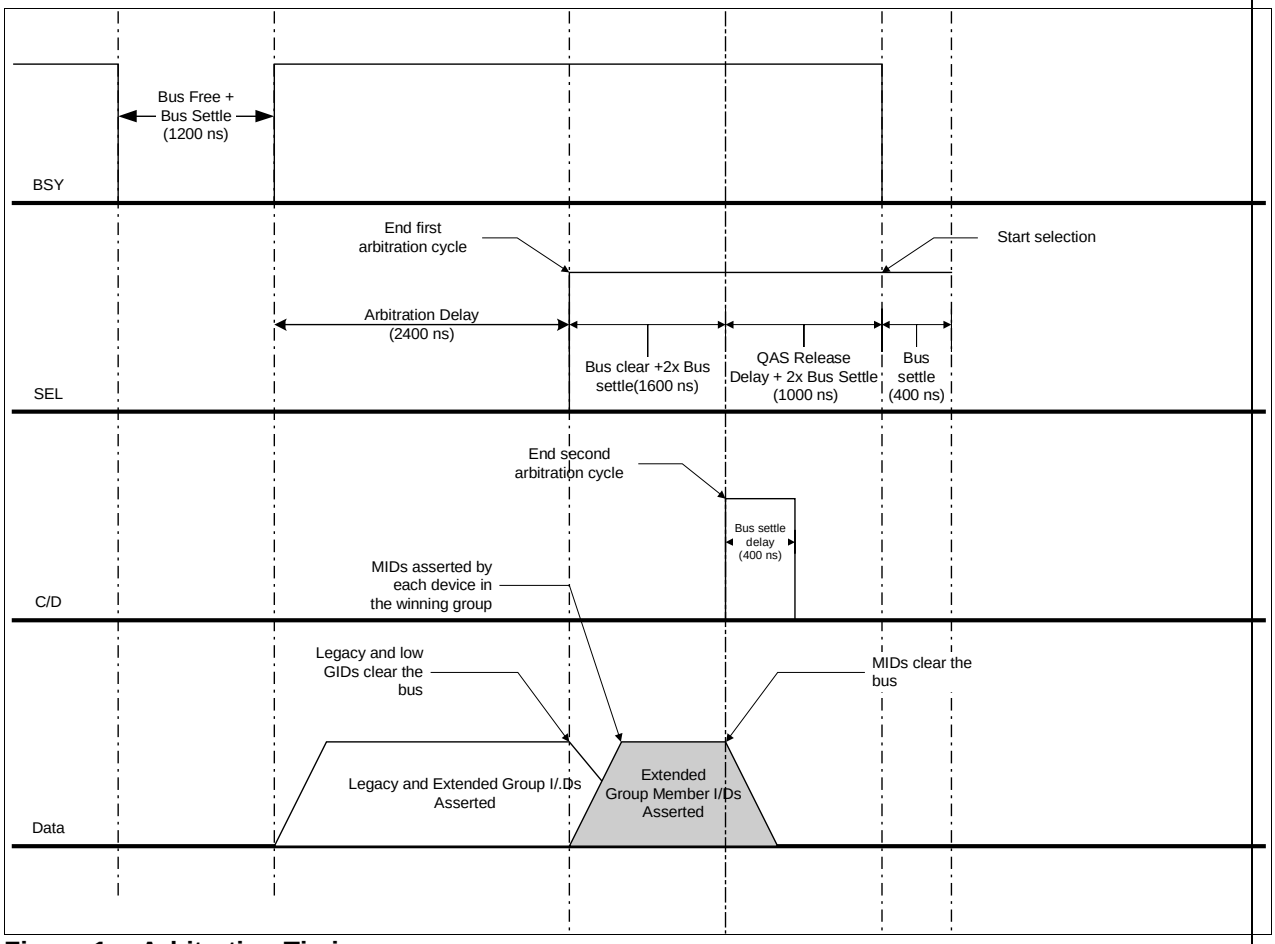

Figure 1 -- Arbitration Timing

5.3 Extended QAS

The QAS arbitration protocol reduces overhead by allowing a target to transfer control of the bus to a new device without going through the bus free phase.

A target signals the beginning of quick arbitration to all QAS-enabled devices by means of the handshake defined in SPI-3 (see SPI-3, clause 10.2.2.1). After starting QAS, the target continues to assert BSY until QAS completes as described below.

During the initial arbitration cycle, contending devices assert their group ~~and group member~~ I/Ds. After a delay, all arbitrating devices examine the bus. Those having a lower priority group I/D drop out of contention by deasserting all signals. The remaining devices assert SEL and ~~continue to drive their extended SCSI I/Ds~~their member I/Ds while continuing to assert their GIDs.

The assertion of SEL triggers the start of the second arbitration cycle and signals all devices having a lower group I/D to drop out of contention. After an appropriate delay to allow losing devices to clear the bus, devices still in contention sample the data bus again. The winning device asserts C/D to indicate the completion of the second arbitration cycle. Lower priority devices must drop out of contention within a given time of C/D becoming true. The winning device then negates C/D and releases all other signals.

On detecting the assertion of C/D, the target waits long enough to allow transients to settle and losing devices to clear the bus. It then terminates QAS arbitration by deasserting BSY. Following deassertion of BSY, the winning device performs extended selection as described in section 6.

1.1.7 Details of Extended QAS

The following detailed description is derived from the QAS protocol specified in clause 10.2.2.1 of the SPI-3 document. Procedures are given specifying the behavior of the target releasing the bus and a device wishing to gain control of the bus. QAS timing is shown in Figure 2~~Figure 2~~.

5.3.1.1 Extended QAS -- Target releasing the bus

The following steps are performed when a QAS-enabled target terminates a connection with a QAS-enabled initiator. This procedure signals the start of QAS arbitration and passes control of the bus from the target to the winning device.

a) The target shall change to a MESSAGE IN phase and issue a single QAS REQUEST (55h) message. The target shall assert REQ for a minimum of 16 ns. The current initiator shall assert the ACK signal for a minimum of 16 ns in responding. The target shall hold the message byte for a minimum of 33 ns after detection of the ACK signal being asserted.

b) After the initiator negates the ACK signal for the QAS REQUEST message and if the initiator does not create an attention condition then the initiator shall release all SCSI signals within two system deskew delays after detecting MSG, C/D, and I/O signals false.

c) After detection of the last ACK signal being false and if there is no attention condition, the target shall release all SCSI signals except the BSY, MSG, C/D and I/O signals and the target shall negate the MSG, C/D, and I/O signals within two system deskew delays.

d) If the target detects the SEL signal being true, the target shall release the MSG, C/D and I/O signals within one QAS release delay.

e) After waiting at least a QAS arbitration delay from releasing the SCSI signals in step (c), if there are no SCSI ID bits true, the target shall transition to the BUS FREE phase.

f) After waiting at least a QAS arbitration delay from releasing of the SCSI signals in step (c), if there are any SCSI ID bits asserted the target shall wait at least a second QAS arbitration delay. If the SEL signal is not asserted by the end of the second QA arbitration delay, the target shall transition to the BUS FREE phase.

g) The target shall wait for the C/D signal to be asserted.

h) After waiting at least a QAS release delay plus two times the bus settle delay from detecting the assertion of the C/D signal in step (g), the target shall release the BSY signal.

5.3.1.2 Extended QAS -- Device arbitrating for the bus

The following procedure applies to a QAS-enabled device that wishes to arbitrate for the bus.

a) The SCSI device may arbitrate for the SCSI bus by asserting its own group ID within a QAS assertion delay from detection of the MSG, C/D, and I/O signals being negated (see section 5.3.1.1, step (c).

b) ~~After waiting at least a QAS arbitration delay (measured from the detection of the MSG, C/D, and I/O signals being negated)~~ The SCSI device shall examine the DATA BUS when one of the following occurs:

A) A QAS Arbitration delay has elapsed (measured from the negation of the MSG, C/D and I/O signals) or
B) SEL is asserted by another device

whichever occurs first.

c) If no higher priority group I/D is true on the DATA BUS, then the SCSI device is still in contention and it shall assert the SEL signal and its member ID after two deskew delays and within a QA release delay of the event detected in step (b). The device shall continue to assert its group I/D.

d) If a higher priority group I/D is present on the DATA BUS (see section 4 for the SCSI ID arbitration priorities), then the SCSI device has lost the arbitration. All devices that have lost the arbitration shall release their SCSI IDs after two deskew delays and within one QA release delay after detection of the SEL signal being asserted. An SCSI device that loses arbitration may return to step (a).

e) The SCSI devices still in contention shall wait at least a QAS release delay plus twice the bus settle delay. A device still in contention shall examine the DATA BUS. If a higher priority member I/D is present on the bus, then the SCSI device has lost arbitration. If no higher priority member I/D is present on the bus, then the device has won arbitration.

f) The device that has won arbitration shall assert the C/D signal. After a bus settle delay following the assertion of C/D, the device shall negate the C/D signal. The device shall release (stop driving) C/D within one QAS release delay of its being negated.

g) The device that has lost arbitration shall release all signals after two deskew delays and within one QAS release delay after the C/D signal becomes true.

h) The SCSI device that has won arbitration shall wait at least a QAS release delay plus twice the bus settle delay after asserting C/D before changing any other signals.

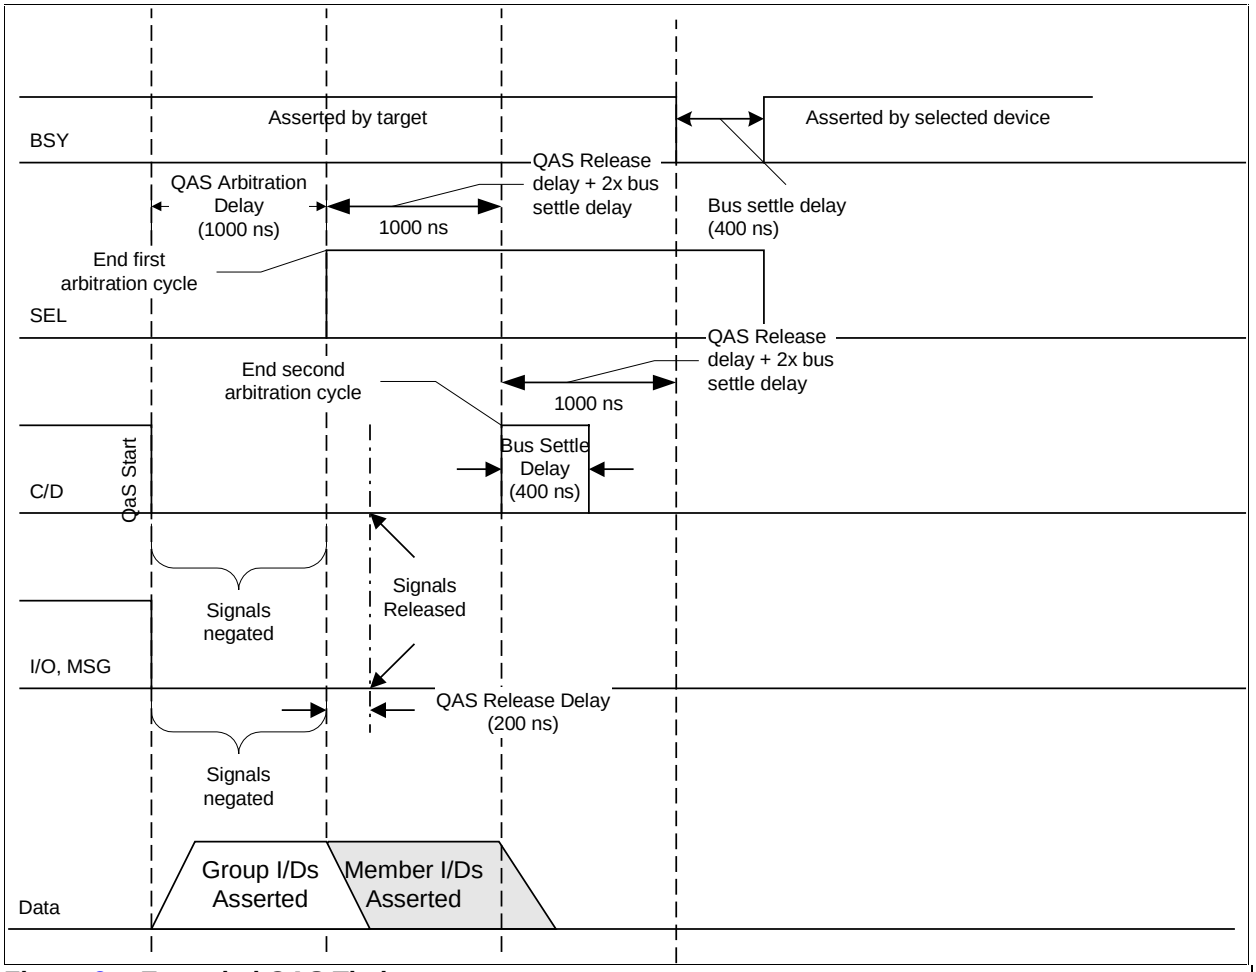

Figure 2 -- Extended QAS Timing

6 Extended Selection

To participate in extended selection, all extended devices are required to listen during each arbitration or quick arbitration cycle to determine the SCSI ID of the winning device. The inclusive OR of this value and the device's SCSI I/D are saved in a selection mask register internal to the device. A device determines that it has been selected by referencing this information during the selection phase.

Note: The ability to snoop the bus during arbitration is required for devices that implement the fair arbitration policy described in section 7.

6.1 Monitoring bus activity during arbitration

The listening device monitors the bus as follows:

a) The device shall first wait for the BUS FREE phase to occur. The BUS FREE phase is detected whenever both the BSY and SEL signals are simultaneously and continuously false for a minimum of a bus settle delay.

b) After detecting the BUS FREE condition, the listening device shall continuously monitor the BSY signal. If BSY is asserted, the listening device shall begin sampling the data bus. Sampling shall stop after C/D is asserted. The device shall save the bus I/D of the winning device.

c) The device shall set its selection mask to the inclusive OR of its SCSI ID and the SCSI ID of the winning device.

6.2 Monitoring Bus Activity during quick arbitration

Listening devices observe the bus as follows:

a) The listening device detects the start of Quick Arbitration as described in section 5.3.1.1

b) After waiting for a QAS arbitration delay, the device begins sampling the data bus continuously.

c) When C/D is asserted, the device shall stop sampling and save the bus I/D of the winning device.

d) The device shall set its selection mask to the inclusive OR of its SCSI ID and the SCSI ID of the winning device.

6.3 Extended Selection Phase

During the extended selection phase, the selecting device asserts a bit mask corresponding to its SCSI ID inclusively ORed with the SCSI ID of the device to be selected.

A device shall determine that it is selected when

a) The SEL signal is true and the BSY and I/O signals are false for at least a bus settle delay and

b) The data bus is equal to the selection mask register. Parity must be valid.

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.

An extended device ~~that supports extended addressing~~ shall not respond to selection if:

a) Parity is invalid or

b) More than four or less than three data bits are set.

6.3.1 Extended Selection Timeout

In order to facilitate timely device discovery, extended targets shall respond to extended selection within an extended selection timeout delay.

6.4 Extended Reselection

The reselection protocol is similar to the above, except that the I/O signal is true.

7 Extended Addressing Fairness

7.1 Model for Extended Devices

In this extended version of fairness, each device supporting fair arbitration implements group and ~~group~~ member fairness registers. When an SCSI device does not need to arbitrate for the SCSI bus, it monitors the arbitration attempts of the other SCSI devices and updates its fairness registers as follows.

a) During the group arbitration cycle, each listening device records in its group fairness register all asserted group I/Ds whose priority is equal to or less than the listening device.

b) If the winning group I/D in step (a) is at or below the priority of the device, it updates its ~~group~~ member fairness register with the MID of each lower priority device in the winning group (relative to the listening device) that has lost arbitration during the member arbitration cycle. (If a legacy device wins arbitration during the group arbitration cycle, the contents of the group member fairness register are set to zero.)

Whenever a requirement for arbitration arises, the device first checks to see if its group ~~both of its~~ fairness register is clear.

Note: Based on the rules for setting fairness register contents, the group fairness register will not be clear if the member fairness register is non-zero.

If so, this SCSI device may now participate in arbitration. If the group fairness register is not clear, the SCSI device postpones arbitration until all lower priority group ~~and group member~~ I/Ds have been cleared from the group fairness register as described below.

~~A device that is waiting to arbitrate shall update its fairness registers as follows:~~

~~a)Group IDs are cleared from the group fairness register as all contending devices in those groups either win or withdraw from arbitration.~~

~~b)~~a) ~~Lower priority group member I/Ds are cleared from the group member I/D register as lower priority devices in that group win arbitration.~~

To prevent low priority devices from holding off a high priority device indefinitely, new member and group I/Ds are added to the fairness registers only when the SCSI device wins or is not waiting to arbitrate. Other lower priority SCSI devices that decided to arbitrate after the waiting device will therefore not indefinitely inhibit that device from attempting to arbitrate.

7.2 Fairness Behavior for Extended Addressing

The behavior of the fairness policy for extended addressing deviates from fairness as implemented in legacy SCSI. Under the circumstances described below, a lower priority device can preempt a higher priority device for a single fairness rotation.

This preemption occurs whenever devices from several lower priority groups are in contention for the bus. Under these circumstances, higher priority devices cannot discover which group members to defer to until one of the lower priority groups wins and enters the member arbitration cycle. Consequently, a device in a lower priority group that decides to arbitrate after other devices in its peer group, may preempt devices in higher priority groups for one fairness rotation.

The effect on arbitration latency is similar to a service queue in which members closer to the head of the queue may allow a limited number of new arrivals to cut into line. In this case, the new arrivals are not detected until they reach the head of the queue. Once detected, of course, no further preemption is allowed.

7.3 Fairness Algorithm for Extended Addressing

~~In order to prevent starvation, the extended fairness algorithm obliges each participating device to defer to lower priority devices that have unsuccessfully arbitrated for the bus during a past arbitration cycle. To record the identity of such devices, an extended device implements group and member fairness registers containing device group and member I/Ds respectively. If the fairness registers are clear, a device may arbitrate otherwise the device must postpone arbitration until the lower priority devices have either won or withdrawn from arbitration.~~

A device may have one of the following arbitration states:

a) Non-participating -- The state of a device that is not participating in arbitration for itself at this time. A device in this state monitors the arbitration attempts of other devices and records the results in its fairness registers.

b) Deferring – The state of a device that wishes to arbitrate but refrains from doing so because of the fairness policy. A deferring device monitors the arbitration attempts of other devices and removes device I/Ds from its fairness registers as those devices either win or withdraw from arbitration.

c) Arbitrating – The state of a device that is actively participating in arbitration. Upon winning arbitration, the device updates its fairness registers with the I/Ds of devices that lost arbitration.

The arbitration state transitions are shown in Figure 3~~Figure 3~~.

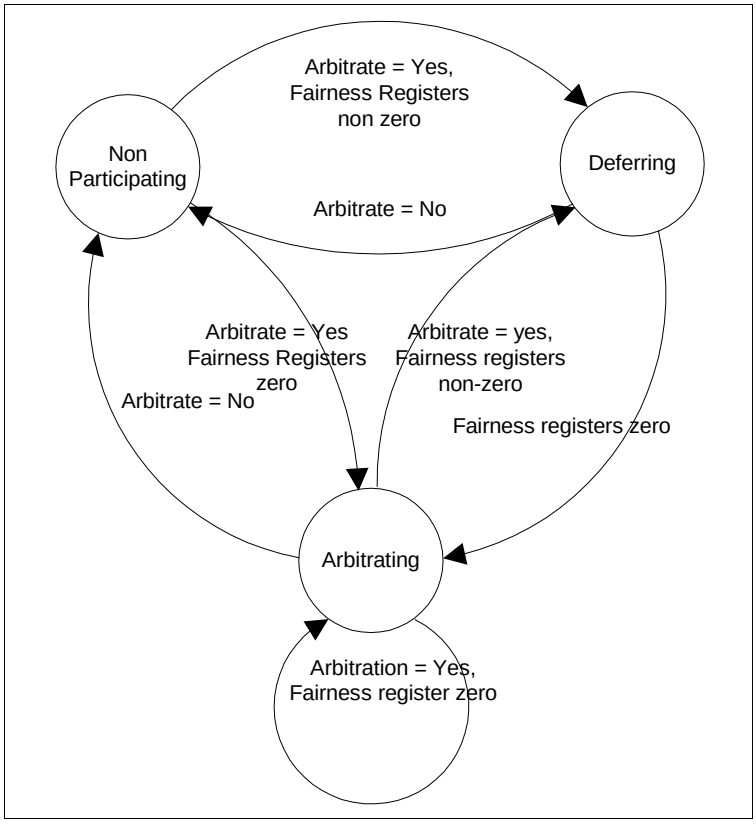

Figure 3 -- Arbitration States

A device in the non-participating state determines which lower priority devices have previously attempted and lost arbitration by updating its fairness registers as described in section 7.3.1 with the results of the last arbitration cycle.

If the device needs to arbitrate and its group fairness register is non-zero, indicating that one or more lower priority devices have lost arbitration, the device shall enter the deferring state. As described in section 7.3.2, it shall then monitor subsequent arbitration cycles until all devices represented in its fairness registers either win or withdraw from arbitration. The device may enter the arbitrating state when its group fairness register is zero.

A device remains in the arbitrating state until it either wins or withdraws from arbitration. An arbitrating device that withdraws shall return to the non-participating state. An arbitrating device that wins shall update its group and member fairness registers with the I/Ds of losing devices as described in section 7.3.3. After relinquishing the bus, the device may:

a) Return to the arbitrating state, if it wishes to arbitrate and its fairness registers are zero.

b) Enter the non-participating state if it does not wish to arbitrate.

c) Enter the deferring state if it wishes to arbitrate and its fairness registers are non-zero.

The sections below describe the rules for updating fairness registers that apply to a device in each of the arbitration states. In what follows, the “listening device” refers to a device that is monitoring the bus during arbitration in order to update its fairness registers.

7.3.1 Fairness Algorithm -- Non-participating Device

A non-participating device shall refresh its group and member fairness registers as described below each time any other SCSI devices arbitrate. The result is that the group fairness register contains the group IDs of lower priority SCSI devices (if any) that have attempted and lost arbitration. (Since several devices may reside in the same group, the group fairness register may contain the group I/D of the non-participating device.) Similarly, the member fairness register contains the member ID of any lower priority devices within the winning group that have attempted and lost arbitration.

NOTE Since non participating devices refresh their fairness registers after every arbitration, SCSI devices that have discontinued arbitration are automatically removed. Thus, the contents of the fairness registers only reflect the participants of the arbitration process that immediately precedes a subsequent arbitration in which this SCSI device may participate.

During an extended QAS or extended normal arbitration cycle, a non-participating device shall update its fairness registers as described in the following steps.

1) Latch the group I/D of all arbitration participants into the group fairness register during the group I/D sample interval.

2) Latch the member I/D of all arbitration participants in the winning group into the member fairness register during the member I/D sample interval.

3) Remove from the group fairness register, the group I/Ds of all groups whose arbitration priority is higher than the listening device.

4) If the priority of the winning group exceeds that of the listening device or a legacy device wins arbitration during the group arbitration cycle, set the member fairness register to zero.

5) If the priority of the winning group is equal to that of the listening device, then remove from the member fairness register the MID of the winning device and all device MIDs higher in priority than the listening device.

6) If the priority of the winning group is less than that of the listening device, then remove the MID of the winning device from the member fairness register.

7) If the member fairness register is zero on completion of steps 4, 5 or 6, then the corresponding group I/D shall be removed from the group fairness register.

If a non-participating device is required to arbitrate, it shall enter the deferring state if its group fairness register is non-zero, otherwise it shall enter the arbitrating state.

7.3.2 Fairness Algorithm -- Deferring SCSI Device

A deferring device postpones arbitration because a lower priority SCSI device has attempted and lost arbitration earlier. The deferring device shall monitor the arbitration attempts of other devices and update its fairness registers as described below.

1) For normal arbitration, a deferring SCSI device shall start a lockout timer of greater than 2,4 microseconds. For QAS, a deferring SCSI device shall start a lockout timer of greater than 1000 nsec.

2) If no other SCSI device participates in arbitration within the bus lockout time-out then a deferring device shall clear its group and member fairness registers and may enter the arbitrating state.

3) If another SCSI device begins arbitration within the lockout time-out then a deferring SCSI device shall remove any group I/Ds for which fairness is no longer required (i.e., SCSI devices that did not participate in the last arbitration).

If member I/Ds from a removed group were represented in the member fairness register, the member fairness register shall be set to zero.

4) If the winning group is present in the group fairness register and member IDs from the winning group are represented in the member fairness register then:

   A) Remove the member ID of the winning device from the member fairness register along with the member I/Ds of any devices in the winning group that did not participate in arbitration.

   B) If the member fairness register is zero after step A, the corresponding group I/D shall be removed from the group fairness register.

5) If the winning group is present in the group fairness register and member IDs from the winning group are not represented in the member fairness register, then:

   A) the member fairness register shall be updated with the member I/Ds asserted during the member arbitration cycle.

   B) The member I/D of the winning device shall be removed from the member fairness register.

   C) If the member fairness register is zero after step B, then the corresponding group I/D shall be removed from the group fairness register.

6) If the device's group fairness register is zero after steps 3, 4 or 5, the device may enter the arbitrating state. Otherwise, the device shall remain in the deferring state and return to step (1).

7.3.3 Fairness Algorithm – Arbitrating Device

A device in the arbitrating state shall not update its fairness registers until it wins arbitration. Upon winning, the device shall:

a) Latch in its group fairness register the group I/D of all devices that have lost arbitration, including those in the winning group.

b) Latch in its member fairness register the member I/D of all devices in the winning group that have lost arbitration.

Upon releasing the bus, the device may return to the non-participating or deferring states.

7.4 Fairness for Legacy Devices

The method described above is compatible with the fairness scheme implemented by legacy devices. Such devices will update their fairness registers during the group arbitration cycle and automatically defer to devices in lower priority groups. Conversely, extended devices will detect and defer to arbitration attempts by lower priority legacy devices.

Since legacy devices do not recognize the member arbitration cycle, however, such devices should not have a group I/D that exceeds the priority of an extended device.

7.5 Extended Fairness Examples

The following examples show the sequence of winners and the progression of fairness register values for selected devices at the start of each arbitration cycle. In the progression charts, the group and member fairness registers are labeled FR_gid and FR_mid, respectively. A check mark (✓) indicates the round in which the device won arbitration. A round in which the device is non-participating is shown with a gray background. A white background denotes a round in which the device is either in the arbitrating or deferring states.

For each example, winning devices enter the non-participating state on releasing the bus.

7.5.1 Example 1

Arbitrating devices:

Group:	7	5	3	2
Members:	12	10, 9, 8	12	10

Arbitration results:

Arb cycle:	1	2	3	4	5	6
Winner:	(7,12)	(5,10)	(5,9)	(5,8)	(3,12)	(2,10)

Fairness register progression for device (7,12)

Arb Cycle		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
✓1	FR_gid											X		X	X		
	FR_mid																
2	FR_gid											X		X	X		
	FR_mid							X	X								
3	FR_gid											X		X	X		
	FR_mid								X								
4	FR_gid													X	X		
	FR_mid																
5	FR_gid														X		
	FR_mid																
6	FR_gid																
	FR_mid																

Fairness register progression for device (5,9)

Arb Cycle		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	FR_gid																
	FR_mid																
2	FR_gid																
	FR_mid																
✓3	FR_gid											X		X	X		
	FR_mid								X								
4	FR_gid													X	X		
	FR_mid																
5	FR_gid														X		
	FR_mid																
6	FR_gid																
	FR_mid																

7.5.2 Example 2

Arbitrating devices

Group:	7	5	3	2
Members:	12	10, 9, 8	12, 9	10

Non-participating devices

Group:	3
Member:	10

Arbitration results:

Arb cycle:	1	2	3	4	5	6	7
Winner:	(7,12)	(5,10)	(5,9)	(5,8)	(3,12)	(3,9)	(2,10)

Fairness register progression for device (7,12)

Arb Cycle		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
✓1	FR_gid											X		X	X		
	FR_mid																
2	FR_gid											X		X	X		
	FR_mid							X	X								
3	FR_gid											X		X	X		
	FR_mid								X								
4	FR_gid													X	X		
	FR_mid																
5	FR_gid													X	X		
	FR_mid							X									
6	FR_gid														X		
	FR_mid																
7	FR_gid																
	FR_mid																

Fairness register progression for device (3,10) (non-participating).

Arb Cycle		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	FR_gid													X	X		
	FR_mid																
2	FR_gid													X	X		
	FR_mid																
3	FR_gid													X	X		
	FR_mid																
4	FR_gid													X	X		
	FR_mid																
5	FR_gid													X	X		
	FR_mid							X									
6	FR_gid														X		
	FR_mid																
7	FR_gid																
	FR_mid																

7.5.3 Example 3

Arb cycle 1 arbitrating devices:

Group:	7	5	3	2
Members:	12	10, 9, 8	12, 9	10

Device (5,11) changes arbitration state from non-participating to defering at the start of arbitration cycle 2:

Arbitration results:

Arb cycle:	1	2	3	4	5	6	7	8
Winner:	(7,12)	(5,10)	(5,9)	(5,8)	(3,12)	(3,9)	(2,10)	(5,11)

Fairness register progression for device (7,12)

Arb Cycle		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
✓1	FR_gid											X		X	X		
	FR_mid																
2	FR_gid											X		X	X		
	FR_mid							X	X								
3	FR_gid											X		X	X		
	FR_mid								X								
4	FR_gid													X	X		
	FR_mid																
5	FR_gid													X	X		
	FR_mid							X									
6	FR_gid														X		
	FR_mid																
7	FR_gid																
	FR_mid																

Fairness register progression for device (5,11)

Arb Cycle		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	FR_gid											X		X	X		
	FR_mid																
2	FR_gid											X		X	X		
	FR_mid							X	X								
3	FR_gid											X		X	X		
	FR_mid								X								
4	FR_gid													X	X		
	FR_mid																
5	FR_gid													X	X		
	FR_mid							X									
6	FR_gid														X		
	FR_mid																
7	FR_gid																
	FR_mid																
✓8	FR_gid																
	FR_mid																

8 Mixing legacy and extended addressing modes

The following sections describe how legacy and extended address devices can interoperate on the same bus. The only restriction is that a device in legacy mode cannot use Quick Arbitrate and Select.

This proposal assumes that a SCSI-compliant legacy device will:

a) Automatically drop out of contention during arbitration when it sees SEL+BSY set at the end of the first arbitration cycle.

b) Not respond to an extended selection phase in which more than two data bits are set.

As discussed previously, up to eight legacy devices and 64 extended address devices can be supported. Each initiator must have two addresses: a legacy address used for selection and reselection consisting of a single, bit significant value in the range of 0 - 15, and an extended address used for arbitration and extended selection.

Device I/Ds are assigned so that at least three data bits are asserted during an extended selection or reselection and two when selection or reselection is performed with a legacy device. An initiator probes for legacy devices by performing a series of selections using a single target I/D bit.

In this case, the following behaviors are required:

a) With either addressing method, extended address targets shall not respond to selection if the number of data bits asserted is less than three or parity is incorrect.

b) As required by the SCSI specification, legacy targets shall not respond to selection if more than two data bits are asserted or parity is incorrect. (The degree to which existing devices comply with this requirement is unknown.)

An initiator would select a legacy device by asserting the initiator's group I/D along with the member I/D of the target device. A legacy device would perform reselection in the same manner. During reselection, the assertion of only two data bits enables the initiator to detect and respond to a legacy device while at the same time inhibiting a response from an extended device.

An initiator would select an extended target by asserting its group and member I/Ds along with those of the target device. Since at least three data bits must be asserted, a compliant legacy device will not respond.

9 Issues

1) Devices that use the SCA-2 connector today pick up their SCSI ID from the encoded value on the connector. Investigation is needed into how such a device would automatically detect and operate in extended mode while retaining the ability to be connected to a legacy bus.
2) Bus loading and configuration rules need to be specified for systems that support extended addressing.
3) Expanders will probably be needed for bus loading considerations so they would have to support the new arbitration and selection methods.
4) Since multiple devices will be asserting SEL, this signal may require filtering to remove release glitches.

10 Bus control timings

The bus control timings used in this document are listed in Table 2~~Table 2~~. These values are extracted from SPI-3, table 29.

Table 2 -- SCSI bus control timing values

Timing description	Timing values
Arbitration delay	2.4 µs
Bus Clear Delay	800 ns
Bus Free Delay	800
Bus Set Delay	1.6 us
Bus Settle Delay	400 ns
QAS Arbitration Delay	1000 ns
QAS Assertion Delay	200 ns
QAS Release Delay	200 ns
Selection Abort Time	200 µs
System Deskew Delay	45 ns
Extended Selection Timeout Delay	200 µs

11 Change History

11.1 Revision 1:

Increased overall timing for extended arbitration by 200 ns (4%) as follows:

Increased the time delay between the group and member arbitration cycles by a Bus Settle Delay (400 ns) to allow enough time for a legacy device I/D to clear the bus. Since legacy devices do not participate in the member arbitration cycle, the QAS timing rules were used to decrease the delay between the end of the second round of arbitration and the beginning of Selection. This reduced the budgeted time delay by 200 ns (from 1200 to 1000 ns).

Tightened the timing specification for completion of the group arbitration cycle by requiring all arbitrating devices to complete the cycle when SEL is asserted or when the appropriate timeout interval elapses.

Added detailed description of extended fairness algorithm

Added extended fairness examples.

Added extended selection timeout definition