X3T9.2/89-65

SCSI Bus Fairness through SELECTIVE RECONNECTION

David A Buesing
Lee D Cleveland

International Business Machines
Highway 52 & 37th Street NW
Rochester, Minnesota 55901

05/05/89

Problem: Target Bus Hogging

IBM

When two or more SCSI devices arbitrate for the SCSI bus at the same time, the highest priority device will always win. This is a part of the SCSI standard. Often times this is a desirable feature.

In some applications this architecture is undesirable as it may prevent the low priority devices from being serviced in a timely fashion. The lack of service may cause:

- data over/under runs which in turn effects overall system performance.
- the high priority devices to be favored when in fact the system application desires all devices to be treated with fairness, all getting an equal amount of bus time.

Example of Selective Reconnection

IBM

GIVEN:

1. One Initiator with SCSI ID 7 (highest priority).
2. One to seven Targets with unique SCSI IDs.
3. Physical position on the bus is not relevant.

SOLUTION:

The Initiator monitors the bus and determines when an arbitration phase is started. Point t1 in figures 1 & 2.

Note: All participants in arbitration must wait at least an Arbitration Delay (time (1)) from the assertion of their SCSI ID on the bus before determining if they should leave the bus. Further, a Device may not enter into Arbitration for a Bus Set Delay (time (5)) after the start of Arbitration (Busy asserted). One may conclude from this that all those that are participating in arbitration have placed their SCSI ID on the bus by the end of the Bus Set Delay (time (5)). This is called "Latch Arbitrators" time, point t3 in figure 1.

The Initiator enters Arbitration by placing it's SCSI ID on the bus prior to Latch Arbitrators time, point t2 in figures 1 & 2 for example. At Latch Arbitrators time the Initiator latches the SCSI ID's of the Targets participating in Arbitration, point t3 in figure 1. This establishes a Poll List of SCSI IDs. Since the Initiator has the highest priority SCSI ID, it wins the Arbitration, point t4 in figure 1.

The Initiator may then Select any one of the Targets from the Poll List. See figure 2, for example, where the Initiator selects the Target with the SCSI ID of 2. (The method used by the Initiator to determine which Target on the Poll List is selected for reconnection is beyond the scope of this discussion.) The Initiator sends a NO-OP Message to the Selected Target. See figure 3, Information Message Out Phase. The Selected Target then receives the NO-OP from the Initiator. Presuming that the Target entered arbitration for reconnection, the Target will then proceed is if it had won arbitration by sending an Identify Message to the Initiator. The Target may proceed with normal SCSI protocol.

After the Selected Target completes the activity, the bus will enter a Bus Free Phase. Another Arbitration Phase may be entered. The Initiator may or may not enter into Arbitration. It is expected that those Targets that entered into arbitration on the previous arbitration and were not selected would enter into arbitration once again. It is possible, that the Target that was just previously selected, may enter into arbitration again.

NORMAL OPERATION
FIGURE 1

DISCONNECTED DEVICES ARBITRATE; INITIATOR WINS ARBITRATION.

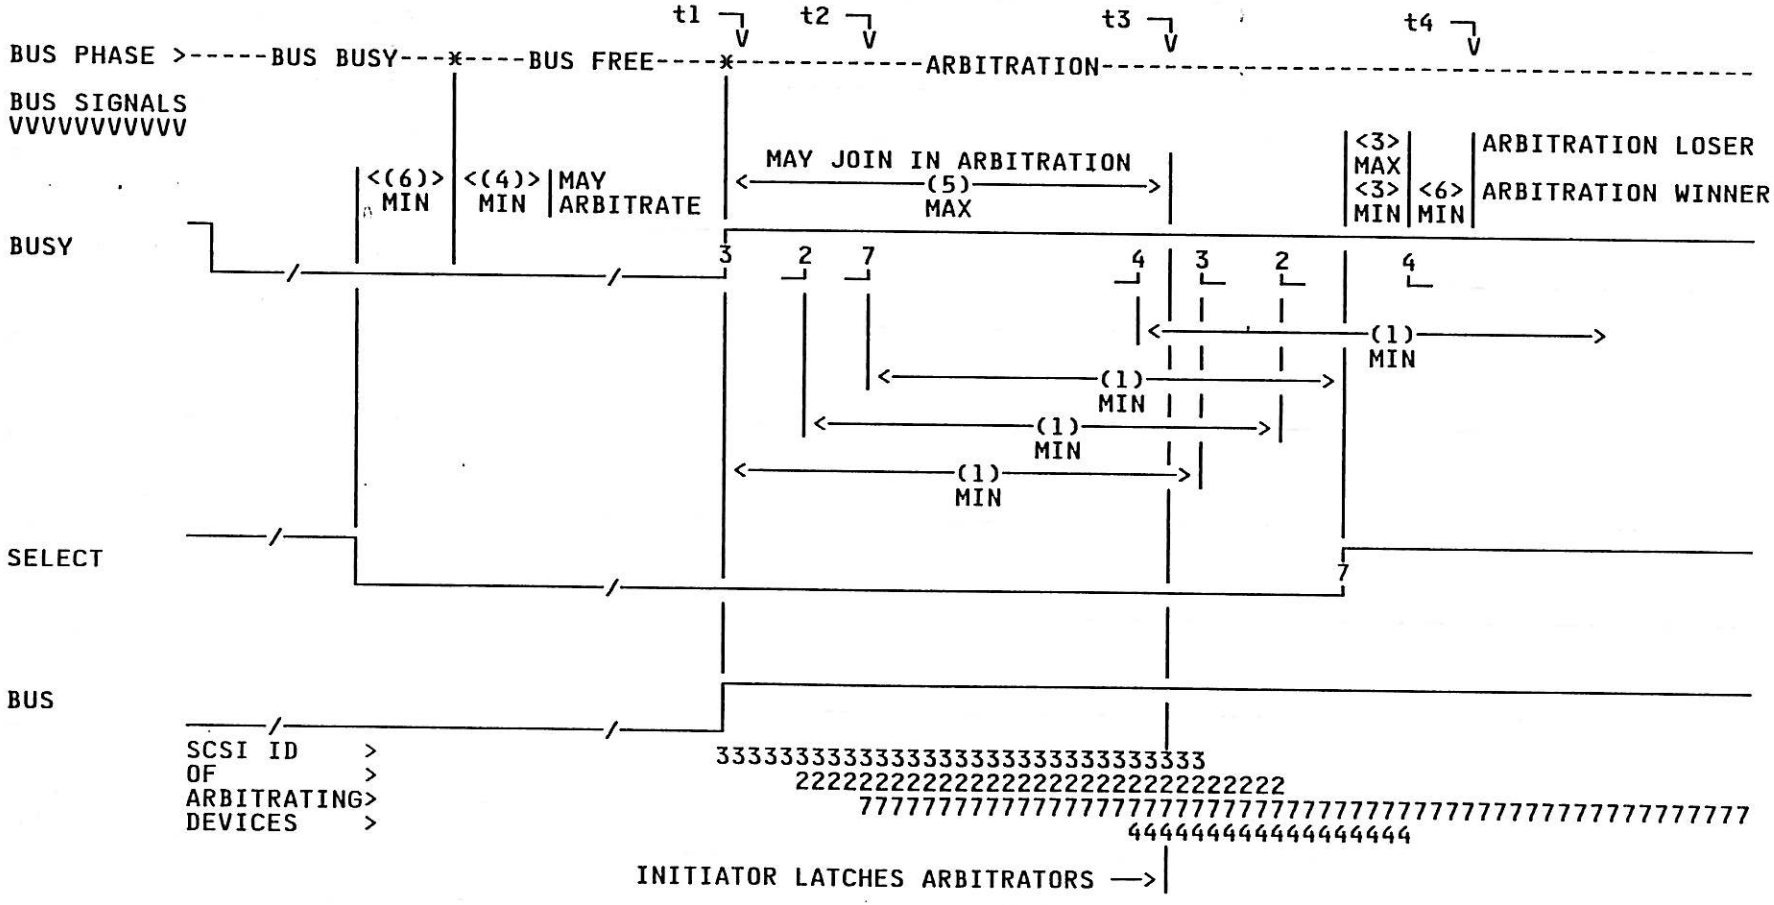

NOTE:
NUMBERS IN PARENTHESIS, IE (1), ARE "SCSI TIMING SPECIFICATIONS" WHICH ARE LISTED SEPARATELY.

dab 05/05/89
arbdet2

215

NORMAL OPERATION
FIGURE 2

DISCONNECTED TARGETS ARBITRATE WITH INITIATOR;
INITIATOR WINS ARBITRATION AND SELECTS AN ARBITRATING TARGET.

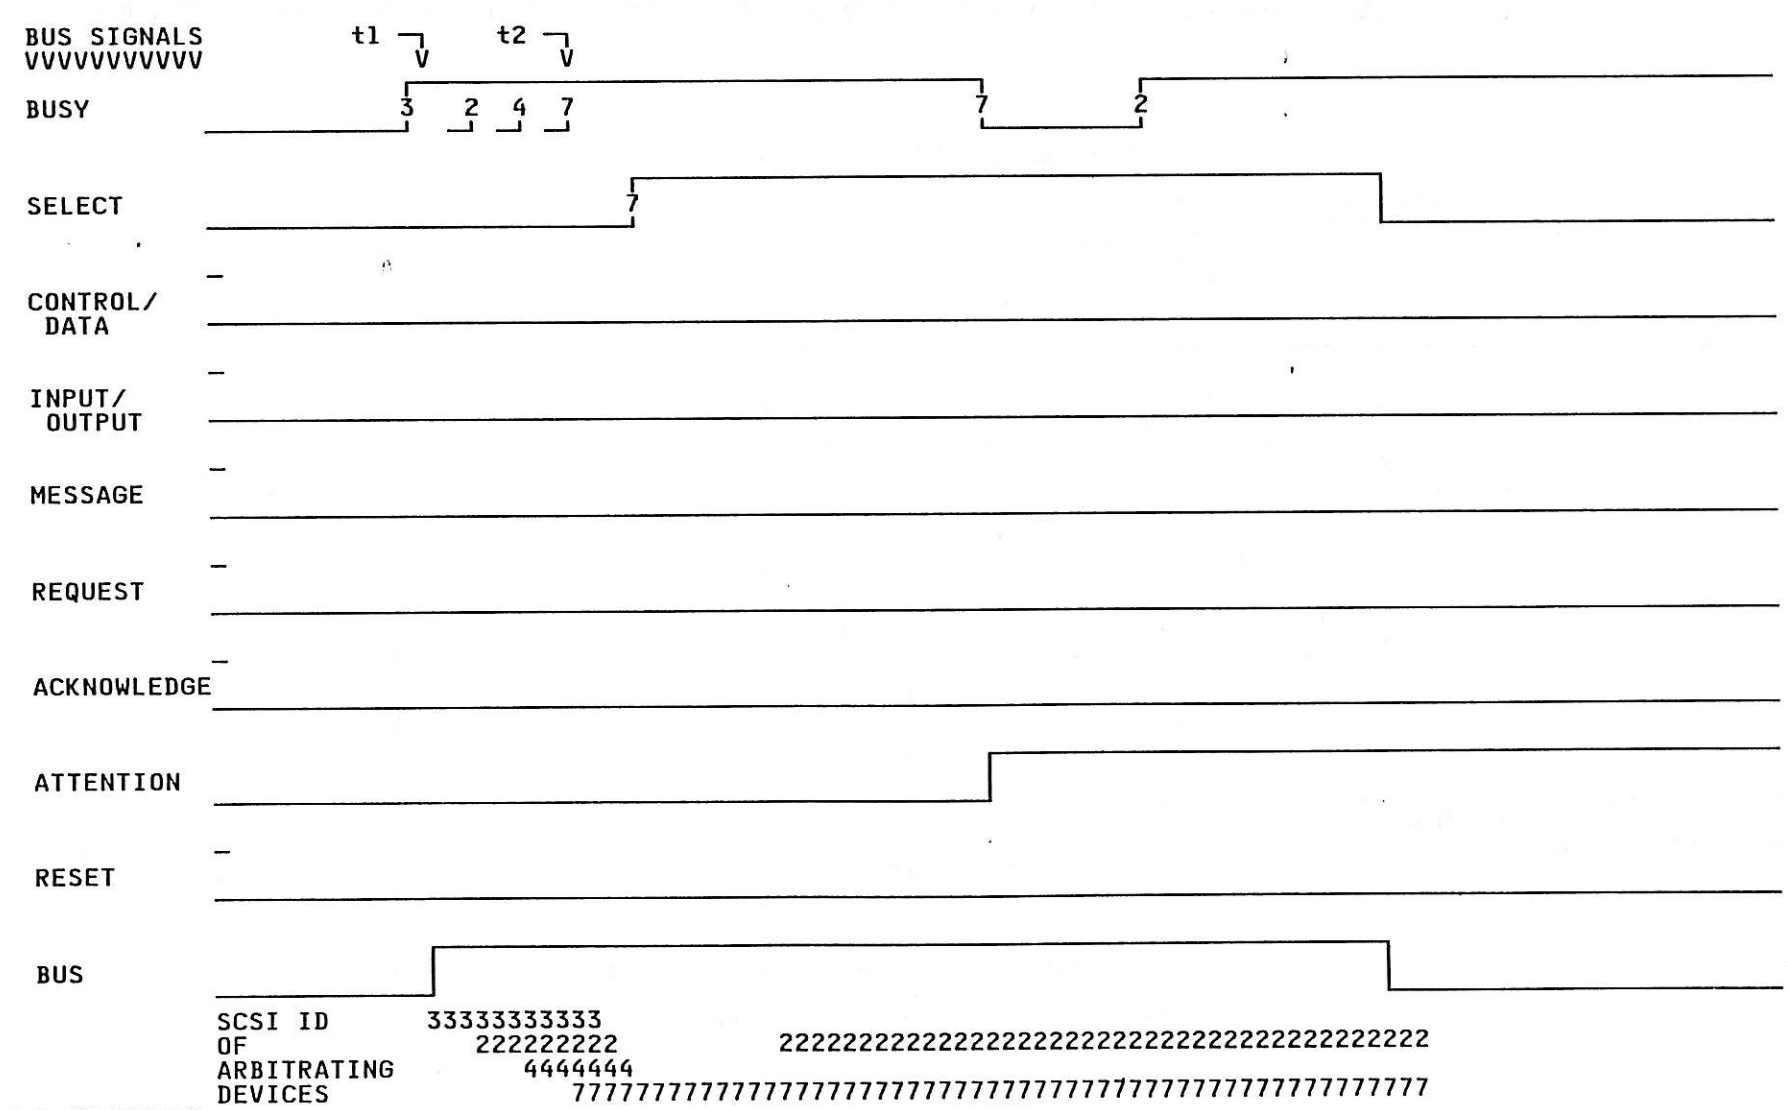

dab 05/05/89
iarbitra

NEW OPERATION
FIGURE 3

INITIATOR NO-OPS SELECTED DISCONNECTED TARGET THAT ARBITRATED;
TARGET CONTINUES SELECTION WITH DISCONNECTED INFORMATION
(THE REASON THE TARGET ARBITRATED).

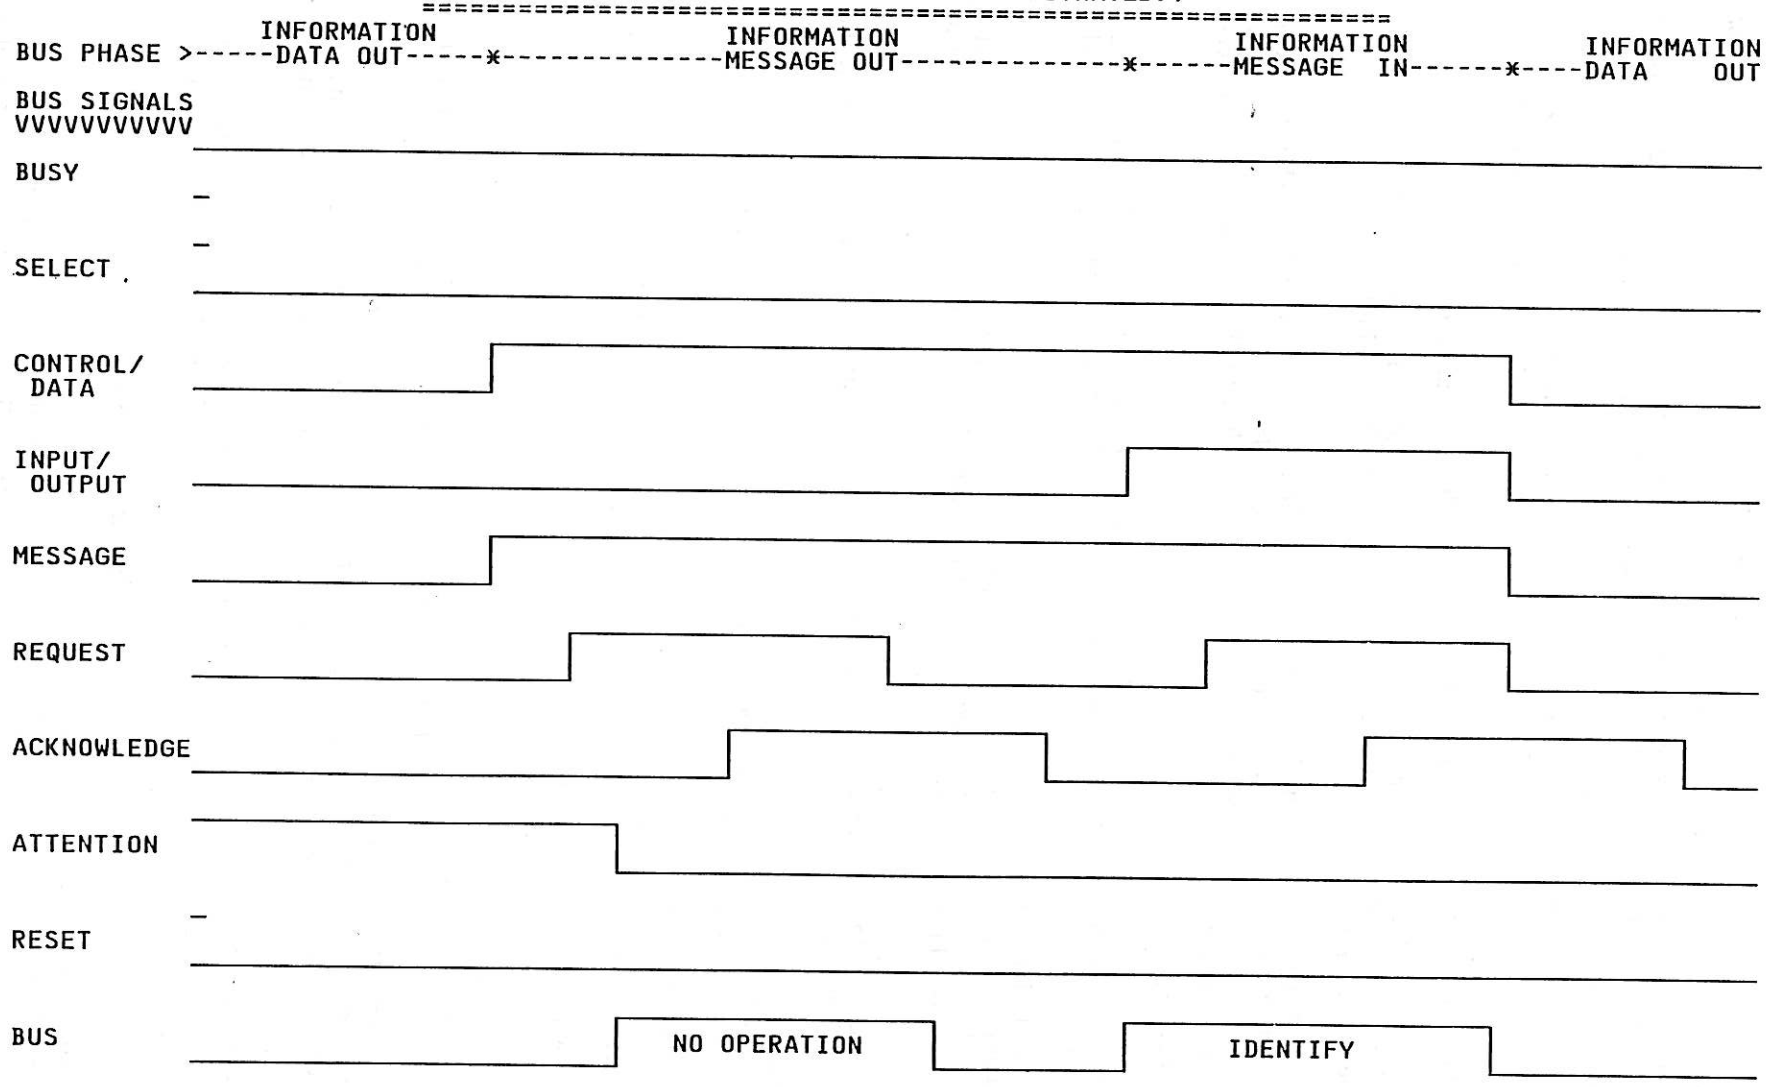

dab 05/05/89
tinforma

(1) Arbitration Delay (2.4 microseconds)
The minimum time an SCSI device shall wait from asserting BSY for arbitration until the DATA BUS can be examined to see if arbitration has been won. There is no maximum time.

(2) Assertion Period (90 nanoseconds)
The minimum time that a target shall assert REQ while using synchronous data transfers. Also, the minimum time that an initiator shall assert ACK while using synchronous data transfers.

(3) Bus Clear Delay (800 nanoseconds)
The maximum time for an SCSI device to stop driving all bus signals after:

1. The BUS FREE phase is detected (BSY and SEL both false for a bus settle delay).
2. SEL is received from another SCSI device during the ARBITRATION phase.
3. The transition of RST to true.

NOTE: For the first condition above, the maximum time for an SCSI device to clear the bus is 1200 nanoseconds from BSY and SEL first becoming both false. If an SCSI device requires more than a bus settle delay to detect BUS FREE phase, it shall clear the bus within a bus clear delay minus the excess time.

(4) Bus Free Delay (800 nanoseconds)
The minimum time that an SCSI device shall wait from its detection of the BUS FREE phase (BSY and SEL both false for a bus settle delay) until its assertion of BSY when going to the ARBITRATION phase.

(5) Bus Set Delay (1.8 microseconds)
The maximum time for an SCSI device to assert BSY and its SCSI ID bit on the DATA BUS after it detects BUS FREE phase (BSY and SEL both false for a bus settle delay) for the purpose of entering the ARBITRATION phase.

(6) Bus Settle Delay (400 nanoseconds)
The time to wait for the bus to settle after changing certain control signals as called out in the protocol definitions.

(7) Cable Skew Delay (10 nanoseconds)
The maximum difference in propagation time allowed between any two SCSI bus signals when measured between any two SCSI devices.

(8) Data Release Delay (400 nanoseconds)
The maximum time for an initiator to release the DATA BUS signals following the transition of the I/O signal from false to true.

(9) Deskew Delay (45 nanoseconds)
The minimum time required for deskew of certain signals.

(10) Disconnection Delay (200 microseconds)
The minimum time that a target shall wait after releasing BSY before participating in an ARBITRATION phase when honoring a DISCONNECT message from the initiator.

(11) Hold Time (45 nanoseconds)
The minimum time added between the assertion of REQ or ACK and the changing of the data lines to provide hold time in the initiator or target, respectively, while using synchronous data transfers.

(12) Negation Period (90 nanoseconds)
The minimum time that a target shall negate REQ while using synchronous data transfers. Also, the minimum time that an initiator shall negate ACK while using synchronous data transfers.

(13) Reset Hold Time (25 microseconds)
The minimum time for which RST is asserted. There is no maximum time.

(14) Selection Abort Time (200 microseconds)
The maximum time that a target (or initiator) shall take from its most recent detection of being selected (or reselected) until asserting a BSY response. This timeout is required to ensure that a

target (or initiator) does not assert BSY after a SELECTION (or RESELECTION) phase has been aborted. This is not the selection timeout period.

(15) Selection Timeout Delay (250 milliseconds, recommended)
The minimum time that an initiator (or target) should wait for a BSY response during the SELECTION (or RESELECTION) phase before starting the timeout procedure. Note that this is only a recommended time period. The specifications for the peripheral devices shall be consulted for the actual timing requirements.

(16) Transfer Period (set during a MESSAGE phase)
The Transfer Period specifies the minimum time allowed between the leading edges of successive REQ pulses and of successive ACK pulses while using synchronous data transfers.