G. G. Floryance

X3T9.2/90-69

IBM Technical Disclosure Bulletin Vol. 32 No. 11 April 1990

TRANSPARENT FAIRNESS FOR DEVICES ON THE SCSI BUS

A Small Computer Systems Interface (SCSI) advertises peer-to-peer relationship bus devices. However, this is not true when multiple devices are simultaneously arbitrating for the bus since the device with the highest ID always wins and effectively dominates bus activity. A fairness technique is described which eliminates a bus dominance by a device with the highest device ID. It is a transparent feature implemented solely within a device that maintains conformance to the SCSI standard and allows coexistence with non-FAIRNESS devices on the same SCSI bus.

A device determines fairness by monitoring prior arbitration attempts by other devices. It then postpones arbitration for itself until all lower priority SCSI devices which previously lost arbitration either win a subsequent arbitration or discontinue their arbitration attempts (as in the case where the initiator aborted an outstanding command thus removing the need to re-arbitrate).

When a device does not need to arbitrate for the SCSI bus, it monitors the arbitration attempts of the other devices and updates a fairness register with the device IDs of any lower priority devices which lost arbitration.

When a requirement for arbitration arises, the device first checks to see of its fairness register is clear. If it is clear, then no lower priority device had attempted and lost the previous arbitration, and therefore, this device may now participate in arbitration. If, on the other hand, the fairness register is not clear, the device postpones arbitration until all lower priority device IDs are cleared from the fairness register. Lower device IDs are cleared as those lower level devices win arbitration. Device IDs are also cleared if a device discontinues arbitration (as a result of an internal RESET or initiator directed ABORT).

Since the fairness register is only refreshed when the device is not arbitrating for itself, the fairness register is effectively frozen by the device prior to a requirement for its own arbitration arising. Therefore, only those lower priority devices latched into the fairness register at that time arbitrate ahead of this device. Other lower priority devices which are not latched are not added to the fairness register until this device has successfully arbitrated.

During arbitration, the device ID for all arbitrating devices appear on the SCSI bus during the time prior to SEL active. If the bus

is sampled at this time, which device won and which devices lost arbitration is determined. Sampling occurs at a high enough rate to insure multiple samples within the 600 nsec window.

Fairness circuit pseudo code below summarizes the operation of a fairness algorithm. The code is described in the discussion following.

```
 1    If others arbitrating but own arbitration not required then
 2      F-REG = latched arbitration participants
 3      Mask off winner ID bit in F-REG (Most Significant bit)
 4      Mask off all ID bits >= own ID in F-REG

 5    else (* own arbitration is required *)
 6      if F-REG = 0 then (* participate in arbitration *)
 7        perform normal arbitration with own ID
 8        if arbitration won, execute lines 1-4 above
 9        else re-arbitrate at next opportunity

10      else (* F-REG <> 0 so perform fairness *)
11        Start lockout timer of > 2.4 usec
12        Wait for either lockout timeout or SEL=1

13        if SEL = 1 then (* another device started arbitration
          *)
14          latch arbitrating participants into T-REG
15          Mask off winner ID bit in T-REG
16          F-REG = F-REG and T-REG

17        else  (* lockout timeout occurred *)
18          F-REG = 0
19          go to line 5
```

Lines 1-4 describe the circuit operation if the device is not required to participate in arbitration for itself at this time. This circuitry refreshes the FAIRNESS register F-REG each time the other devices arbitrate. The result is that F-REG contains the device address bits of lower priority devices (if any) which have attempted and lost arbitration. Note that devices which have discontinued arbitration are automatically removed.

Lines 5-9 describe the circuit operation if the device is required to participate in arbitration for itself, and F-REG = 0 indicating that there is no lower priority device to be fair too.

Lines 10-18 describe the circuit operation if the device is required to participate in arbitration for itself and F-REG <> 0 indicating arbitration should be postponed because a lower priority device attempted and lost arbitration earlier.

Line 13 indicates that the fairness logic is waiting for either another device to arbitrate or for the bus lockout timeout to occur. Lines 14-16 describe the circuit operation if another device does begin arbitration within the lockout timeout.

Lines 17-18 are included to handle the case were no other device participated in arbitration within the bus lockout timeout. Note: Bus lockout occurs if all devices are waiting for someone else to start the arbitration process. This occurs when a device A disconnects to begin an operation for initiator B. After some time, A is ready to reconnect but looses arbitration to device B which selects A and aborts the operation. A third fairness device C remembers that device A has lost arbitration and continues to wait until A arbitrates, which no longer occurs since A's operation is aborted. Bus lockout is prevented through use of a timeout that clears the fairness register and returns to the start of the algorithm.