

July 8, 2002

To: T10 Technical Committee
 From: Timothy Hogleund (LSI Logic)
 Subject: SAS Pathway Recovery modification

The presently specified method for avoiding deadlock and reclaiming blocked pathway resources within a SAS topology requires SAS expanders to perform a two-pass algorithm to insure that deadlock is avoided. This method therefore incurs additional complexity in expander implementations and will cause additional time (and associated performance degradation) to resolve pathway conflicts. The two-pass algorithm is necessary due to the present dependency on comparing Arbitration Wait Time values which in a multiple expander topology introduces uncertainty to the comparison, i.e. a second pass is required should the first pass not resolve the conflict.

This proposal calls for a single-pass method to resolve blocked pathways through removal of Arbitration Wait Time comparisons from the Pathway Recovery algorithm and by introducing a new value to be used. This value, Pathway Blocked Count, is to be contained in all Open Address frames and specifies the number of times a port has received a OPEN_REJECT(PATHWAY BLOCKED) response to a connection request.

The port shall initialize the Pathway Blocked Count value to zero upon first connection request to a destination. The port shall increment the Pathway Blocked Count value when OPEN_REJECT(PATHWAY BLOCKED) is received in response to a connection request. The port shall not increment the Pathway Blocked count value past FFh.

Expanders use the Pathway Blocked Count value to determine which connection requests to abort as specified below in 4.8.4.4.3.

4.8.4.4.3 Pathway Recovery

Pathway recovery provides a means to abort connection requests in order to prevent deadlock. When the PPT timer for an arbitrating link expires (i.e., reaches a value of zero), the Expander Function shall determine whether to continue the connection request or to abort the connection request as follows:

The Expander Function shall instruct the arbitrating link to reject the connection request by transmitting OPEN_REJECT(PATHWAY_BLOCKED) when the PPT timer expires and one of the following conditions are met:

- a) the ~~AWT-Pathway Blocked Count~~ value of the arbitrating link (link requesting the connection) is less than the ~~AWT-Pathway Blocked Count~~ value of all links with an arbitration status of WAITING_ON_PARTIAL within the destination port ~~with an arbitration status of WAITING_ON_PARTIAL; and/or~~
- b) the ~~AWT-Pathway Blocked Count~~ value of the arbitrating link is equal to the ~~AWT-Pathway Blocked Count~~ value of ~~all-any~~ links with an arbitration status of WAITING_ON_PARTIAL within the destination port ~~with an arbitration status of WAITING_ON_PARTIAL~~ and the source device name of the arbitrating link is less than the source device name of ~~all-the~~ corresponding links within the destination ~~port with an arbitration status of~~ WAITING_ON_PARTIAL.

~~If none of the conditions above are met the Expander Function shall reinitialize the PPT timer with the PPTOV value. If the PPT timer expires a second time, the Expander Function shall instruct~~

~~the arbitrating link to reject the connection request by transmitting OPEN_REJECT(PATHWAY_BLOCKED) when the following conditions are met:~~

- ~~a) the source device name of the arbitrating link is less than the source device name of any links within the destination port with an arbitration status of WAITING_ON_PARTIAL.~~

7.5.3 OPEN address frame

~~Table 1~~~~Table 47~~ defines the OPEN address frame format used for connection requests.

Table 147. OPEN address frame format

Byte	7	6	5	4	3	2	1	0
0	INITIATOR	PROTOCOL			ADDRESS FRAME TYPE (1h)			
1	FEATURES				LINK RATE			
2	(MSB)	INITIATOR CONNECTION TAG						(LSB)
3								
4	DESTINATION DEVICE NAME							
11								
12	SOURCE DEVICE NAME							
19								
20 20	ReservedReserved							
21 21	PATHWAY BLOCKED COUNT							
22	SCALE	(MSB)	ARBITRATION WAIT TIME					(LSB)
23								
24	(MSB)	Reserved						(LSB)
27								
28	(MSB)	CRC						(LSB)
31								

The PROTOCOL field indicates the protocol for the connection being requested and is defined in ~~Table 2~~~~Table 48~~.

Table 248. Protocol

PROTOCOL	Protocol
000b	SMP
001b	SSP
010b	STP
All others	Reserved

The destination device shall reject the connection request with OPEN_REJECT(PROTOCOL NOT SUPPORTED) if the PROTOCOL field is set to a value it does not support.

An INITIATOR bit set to one indicates the source device is acting as an initiator device. An INITIATOR bit set to zero indicates the source device is acting as a target device. If a target/initiator device sets the INITIATOR bit to one, it shall operate only in its initiator role during the connection. If a target/initiator device sets the INITIATOR bit to zero, it shall operate only in its target role during the connection.

If a target/initiator device accepts an OPEN address frame with the INITIATOR bit to one, it shall operate only in its target role during the connection. If a target/initiator device accepts an OPEN address frame with the INITIATOR bit to zero, it shall operate only in its initiator role during the connection.

fixfix

An initiator-only device shall reject connection requests from an initiator device with OPEN_REJECT(PROTOCOL NOT SUPPORTED). A target-only device shall reject connection requests from a target device with OPEN_REJECT(PROTOCOL NOT SUPPORTED).

The LINK RATE field indicates the rate of the link on which the OPEN address frame originated, and the rate at which all links on the pathway must support, and is defined in [Table 349](#).

Table 349. Link rate

LINK RATE	Link rate
0h	1,5 Gbps
1h	3,0 Gbps
2h - Fh	Reserved

The FEATURES field shall be set to zero. The destination device shall reject the connection request with OPEN_REJECT(PROTOCOL NOT SUPPORTED) if the FEATURES field is set to a value it does not support.

The INITIATOR CONNECTION TAG field is used for SSP and STP connection requests to provide an initiator port an easier context lookup when the target port originates a connection request. SSP or STP initiator ports shall set the INITIATOR CONNECTION TAG field to FFFFh if they do not require the field be provided by the target port. Otherwise, an SSP or STP initiator port should set the INITIATOR CONNECTION TAG field to a unique value per target port. When requesting a connection to an initiator port, a target port shall set the INITIATOR CONNECTION TAG field to the value received in connection requests from the initiator port. An initiator port shall use the same tag for all connection requests to the same target port, and may only change the tag when it has no commands outstanding to that target port. Targets are not required to check consistency of the tags in different connection requests from the same initiator port. SMP initiator ports shall set the INITIATOR CONNECTION TAG field to FFFFh.

The destination device shall ignore the contents of reserved fields in the OPEN address frame.

The PATHWAY BLOCKED COUNT field indicates the number of times the port has retried a connection request due to receiving OPEN_REJECT (PATHWAY BLOCKED). The port shall not increment the Pathway Blocked Count value past FFh.

The SCALE bit and the ARBITRATION WAIT TIME field indicate how long the port sending the OPEN address frame has been waiting for a connection request to be accepted, and are defined in [Error! Reference source not found. Table 50](#). See [Error! Reference source not found. 7.11.3](#) for details on arbitration fairness.