



Additional SAS-1.1 Expander Issues

September 28, 2005

Authors:

Alexander Amezquita (alex@expertio.com) of Expert I/O

Craig Stoops (craig@expertio.com) of Expert I/O

www.expertio.com

Issue 1: SAS Expander additional qualification in Arb Won

Problem

An additional qualification is necessary for an expander link to receive an Arb Won message from the expander function. Consider the following scenario.

Expander Phys 0, 1, 2, and 3 receive open address frames (OAF) A, C, B, and D respectively, where by arbitration rules $D > C > B > A$. Oaf A received on expander phy 0 requests a connection to expander phy 2, Oaf C received on expander phy 1 requests a connection to expander phy 3. OAF B is destined for expander phy 1 and OAF D is destined for expander phy 0.

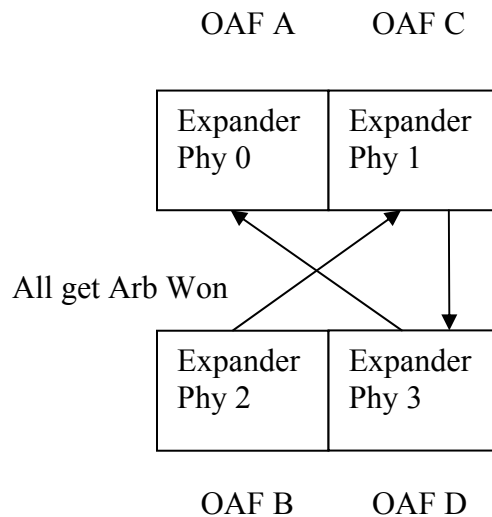
Example Summary:

OAF A received phy 0 destined for phy 2

OAF C received phy 1 destined for phy 3

OAF B received phy 2 destined for phy 1

OAF D received phy 3 destined for phy 0



The current revision of the SAS-1.1 specification (1r10) details an Arb Won is issued as follows:

The ECM shall generate the Arb Won confirmation when all of the following conditions are met:

- a) the connection request maps to an expander phy that:
 - A) supports the connection rate; and
 - B) is not reporting a Phy Status (Partial Pathway), Phy Status (Blocked Partial Pathway), or Phy Status (Connection) response, unless that expander phy is arbitrating for the expander phy making this connection request;
- b) there are sufficient routing resources to complete the connection request;

- c) no higher priority connection requests are present with this expander phy as the destination;and
- d) the connection request is chosen as the highest priority connection request in the expander device mapping to the specified destination expander phy.

Assuming there are sufficient routing resources and the connection rates are valid this scenario proves to be problematic. The expander function receives four request path messages from the respective expander phys. The value of phy status during this phase of the request path is such that Condition a) is satisfied for all ports. Condition b) is satisfied based on the assumption that the expander is designed accordingly. Condition c) rules out expander phy 0 to receive an Arb Won message since $OAF D > A$. Condition d) does not create any further restrictions since each OAF is destined to a unique expander phy.

Thus, since all conditions are met, an Arb Won message is issued to expander phy 3 with OAF D, expander phy 2 with OAF B, and expander phy 1 with OAF C. There are two error cases that need to be considered:

- 1) Expander phy 1 should not receive an Arb Won as it will cause a collision in the connection between expander phy 0 and 3
- 2) Expander phy 2 should not receive an Arb Won since OAF B sent to an end device attached to expander phy 1 will be ignored due to $C > B$. Further, expander phy 1 may have issued AIP due to receiving OAF C, which in turn could result in a hung connection attempt (i.e. an Open Timeout would not occur).

Solution

A recommendation to solve this problem is to add the following clause to the Arb Won qualification.

- e) the connection request of this expander phy is higher priority than an outstanding connection request of the destination expander phy.

Issue 2: Phy Status Clarification in XL1:Request_Path

Problem

Additional rules need to be included for the value of Phy Status while in the XL1:Request_Path state. Particularly when transitioning from XL3:Open_Confirm_Wait or XL6:Open_Response_Wait.

Consider the following example:

Expander Phy 0 receives OAF A destined for expander phy 1. Expander Phy 0 is given an Arb Won message and forwards OAF A. Shortly thereafter, Expander Phy 1 receives OAF B where $B > A$ in terms of arbitration priority. OAF B is destined for expander phy 2 which causes a backoff retry condition. The Phy Status of both Expander Phy 0 and Expander Phy 1 just prior to the BackOff Retry message was Partial Pathway due to being in XL3:Open_Confirm_Wait and XL6:Open_Response_Wait respectively. Upon transmission of the BackOff Retry message and the reception of the BackOff Retry message, each expander phy transitions to XL1:Request_Path. There is no mechanism described in XL1:Request_Path,

XL3:Open_Confirm_Wait, or XL6:Open_Response_Wait to allow the Phy Status to change from Phy Status (Partial Pathway) to Phy Status (Idle).

If we continue with the example and consider an OAF C being received on expander Phy 2 destined for expander Phy 0, we can see we have reached a deadlock. This is because an Arb Won message cannot be issued to expander phy 2 since the Phy Status of expander phy 0 is Partial Pathway. Further, if we assume OAF C > OAF B in terms of arbitration priority, an Arb Won message cannot be delivered to expander phy 1 (See first issue in this document for proof). Thus we have a standoff.

Solution

Upon entrance to XL1:Request_Path, the Phy Status is changed from Phy Status (Partial Pathway/Blocked Partial Pathway) to Phy Status (Idle). Further, the specification doesn't officially declare a Phy Status (Idle), however, by reason of deduction, one must exist. It would be helpful to define this value for Phy Status to aid the reader.

Issue 3: Specification Clarification for Phy Status (Partial Pathway)

Problem

There is text in Section 4.6.6.3 ECM Interface in Table 11 under the row for Phy Status (Partial Pathway) that is misleading. In case a) the text is as follows:

is being used for an unblocked partial pathway (i.e., the expander phy is in the XL3:Open_Confirm_Wait state or XL6:Open_Response_Wait state and the last AIP transmitted or received was not AIP (WAITING ON PARTIAL)); or

I made the assumption that the act of transmitting or receiving AIP (WAITING_ON_PARTIAL) applied to the expander link while in XL3:Open_Confirm_Wait state. Based on experience and comments in Section 7 regarding the XL Expander link definition, transmitting and receiving does not affect the value of Phy Status while in XL3:Open_Confirm_Wait state.

Solution

Reword clause after i.e. to make it clear that transmitting/receiving AIP is only applicable to XL6:Open_Response_Wait. This also applies to the wording in the Phy Status (Blocked Partial Pathway) row of the same table.

Issue 4: Expander – 7.15.4.5 Transition XL1:Request_Path to XL5:Forward_Open

This issue was originally part of our SAS 1.1 letter ballot comments but has not been resolved yet. We are including it here for completeness of all our unresolved issues with the expander architecture.

Problem

The specification indicates that if a forward open message is received after an arbitrating (NORMAL) message has been received, the forward open message is ignored. We believe there is a flaw in the statement or perhaps overall in the expander function handling. It is illustrated in the following case.

Expander Port 2 wins arbitration to open Expander Port 0. Expander Port 1 receives an Open Address Frame (OAF) to Open Expander Port 2 but has to hold off as it waits for the connection to try to open between 2 and 0. The expander function sends an arbitrating (NORMAL) message to expander Port 1 to acknowledge the receipt of the request path message. Expander Port 0 transmits an OAF at the same time an OAF resolving to port 1 is received from the device connected to Expander Port 0. The OAF received by expander Port 0 wins over the outgoing OAF by arbitration rules. Expander Port 0 sends a backoff retry message to Expander Port 2 and also a request path message to the expander function requesting a connection to port 1. The request path message from Expander Port 0 wins by arbitration rules, that is, it is more significant than the outstanding request message by Expander Port 1. As a result, Expander Port 0 issues a forward open message to Expander Port 1. However, since Expander Port 1 had already received an arbitrating (NORMAL) message while in the XL1:Request_Path state, the forward open message is ignored and a stall occurs.

Solution

Removing the restriction of transitioning to XL5:Forward_Open if an Arbitrating (Normal) message has been seen alleviates the problem. There are two cases to consider in determining that this is a valid solution. Following the example described above:

The first case is that the OAF that is forwarded to Expander Port 1 wins via arbitration rules over the OAF received by Expander Port 1. The second case being the OAF that is forwarded to Expander Port 1 loses via arbitration rules over the OAF received by Expander Port 1.

In the case that the forwarded OAF wins, with the restriction removed, the OAF will be received by the port connected and will discard the OAF it sent. This mechanism is already described in both the expander and link layer specification.

The second case can not occur because a forward open message will not be generated by Expander Port 0 destined for Expander Port 1 since Expander Port 0 did not win arbitration in the Expander Function.