October 23, 2002

To: T10 Technical Committee

From: Timothy Hoglund (LSI Logic)

Subject: SAS Expander Annex (informative)

This proposal seeks to provide additional insight into how SAS expanders process connection requests in the form of an informative annex. Information provided hereby is meant as a supplement to relevant normative sections within the SAS draft specification, i.e. this proposed annex contains examples but not in a comprehensive manner.

X SAS Expander Annex (informative)
The following examples are intended to provide additional insight into how SAS Expanders process connection requests.

X.1 SAS connection examples

Figure 1 shows the topology used by examples X.1.1 through X.1.10 below.

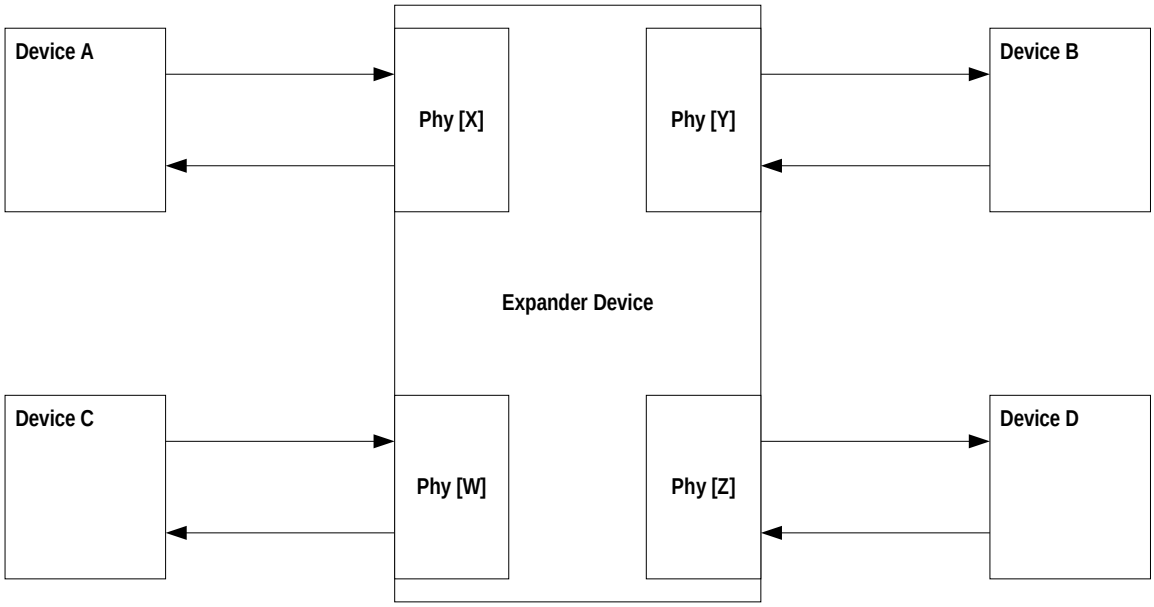

Figure 1: Example Topology #1

Table 1 defines the columns used within Figures X.1.1 through X.1.10.

Column Header	Description
PHY [X] RX	Phy [X] receive from Device A
PHY [X] TX	Phy [X] transmit to Device A
PHY [X] XL state	Phy [X] XL state
PHY [X] XL req/rsp	Phy [X] XL request and response interface signals
PHY [X] cnf/ind	Phy [X] XL confirm and indicate interface signals
PHY [Y] cnf/ind	Phy [Y] XL confirm and indicate interface signals
PHY [Y] XL req/rsp	Phy [Y] XL request and response interface signals
PHY [Y] XL state	Phy [Y] XL state
PHY [Y] TX	Phy [Y] transmit to Device B
PHY [Y] RX	Phy [Y] receive from Device B

Table 1: Column descriptions for SAS connection examples

X.1.1 Connection Request – Open Accept

Figure 2 shows the establishment of a successful connection between two end devices.

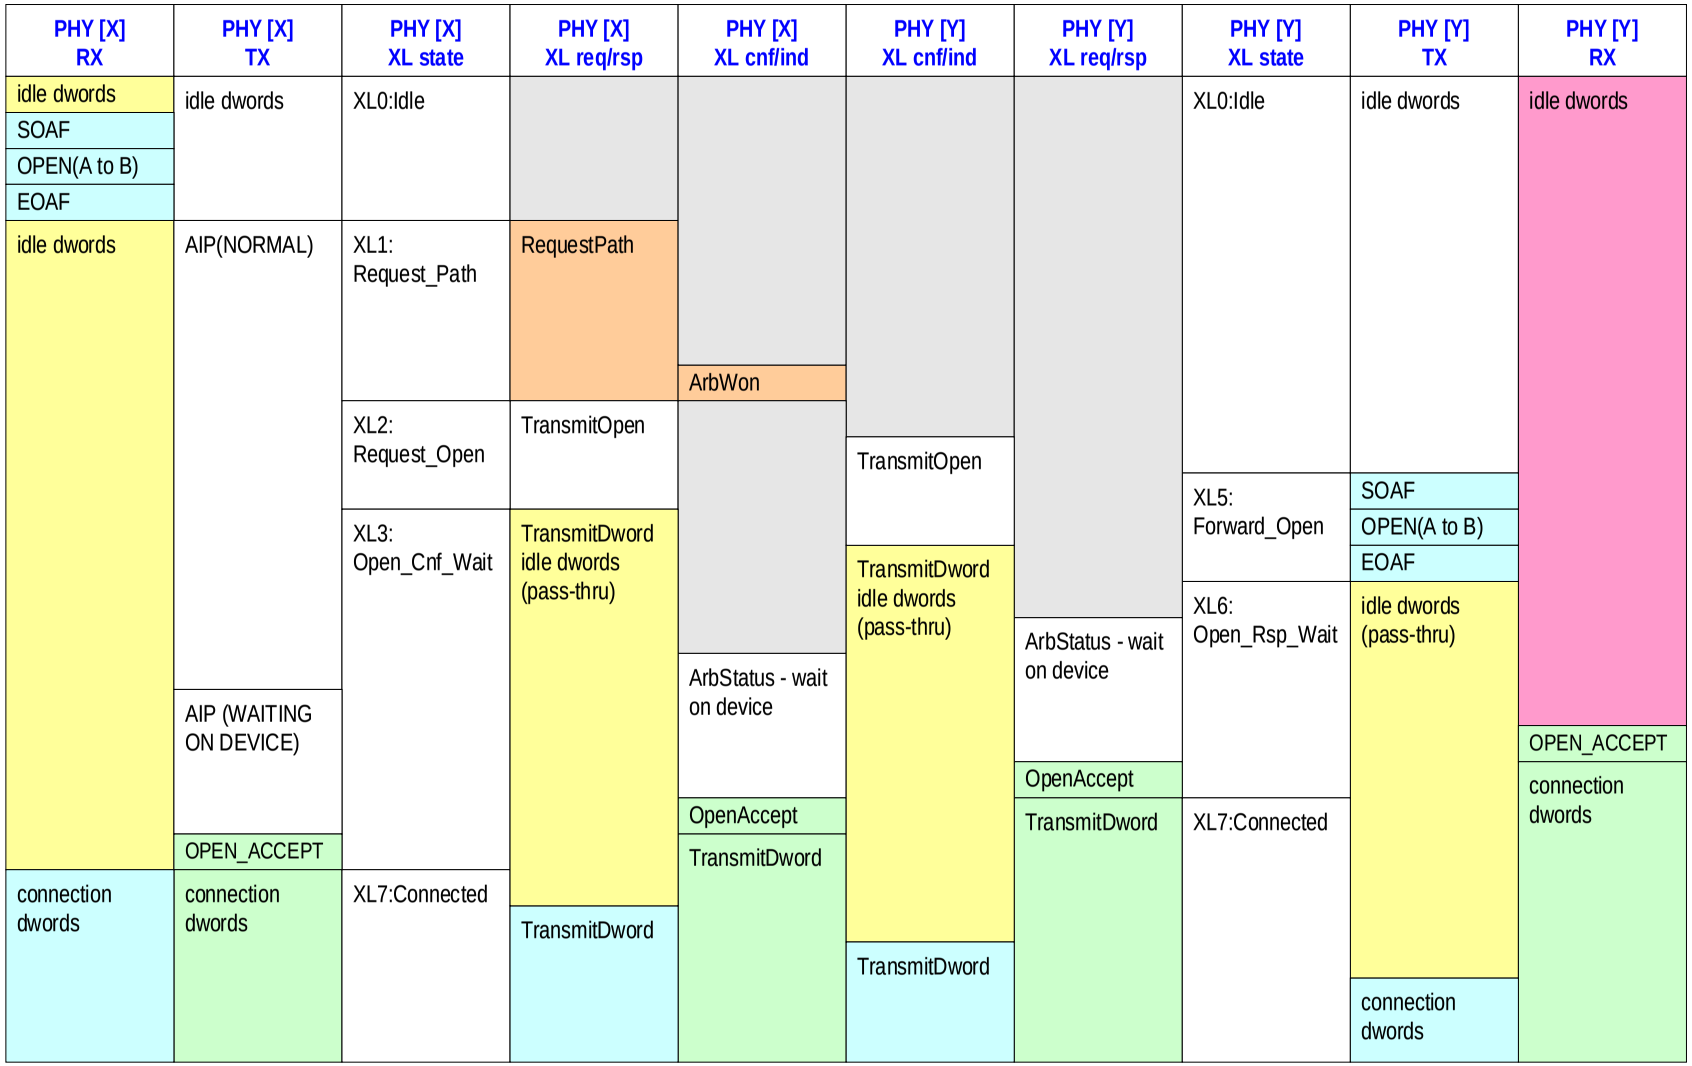

Figure 2. Open Accept

X.1.2 Connection Request – Open Reject by end device

Figure 3 shows failure to establish a connection due to rejection of the connection request by an end device.

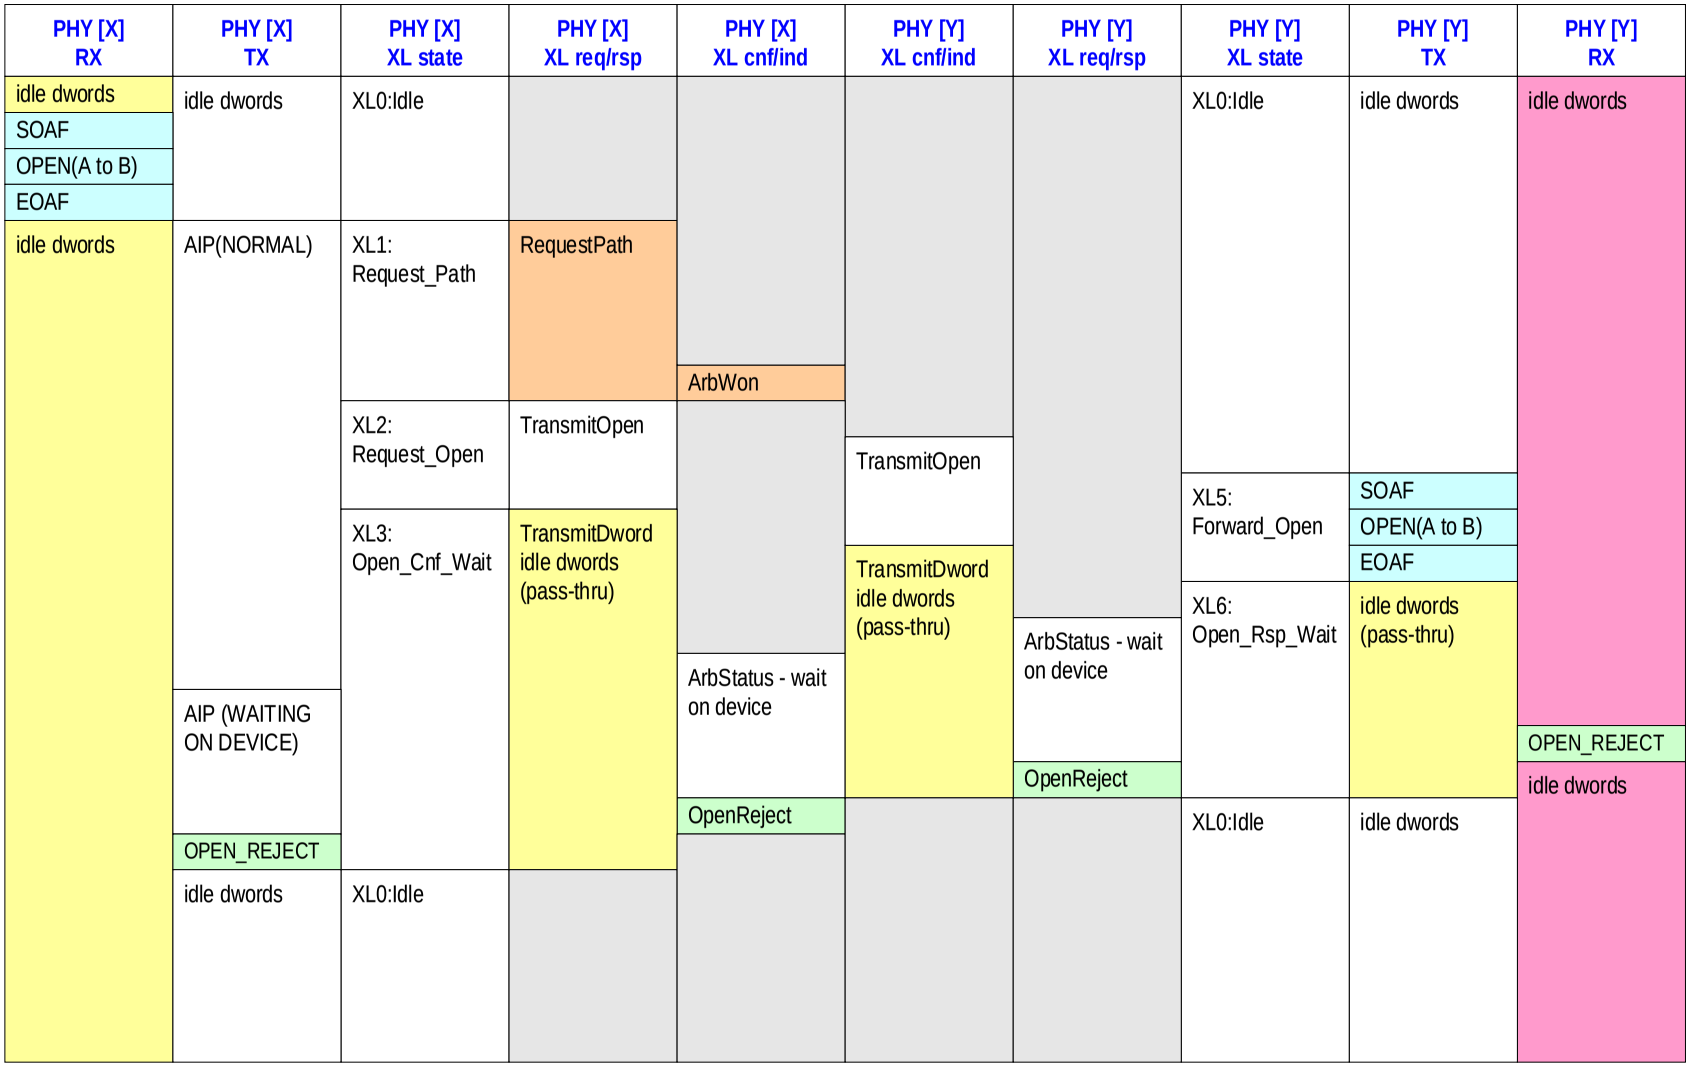

Figure 3. Open Reject by end device

X.1.3 Connection Request – Open Reject by expander device

Figure 4 shows failure to establish a connection due to rejection of the connection request by an expander.

PHY [X] RX	PHY [X] TX	PHY [X] XL state	PHY [X] XL req/rsp	PHY [X] XL cnf/ind	PHY [Y] XL cnf/ind	PHY [Y] XL req/rsp	PHY [Y] XL state	PHY [Y] TX	PHY [Y] RX
idle dwords	idle dwords	XL0:Idle					XL0:Idle	idle dwords	idle dwords
SOAF									
OPEN(A to B)									
EOAF									
idle dwords	AIP(NORMAL)	XL1: Request_Path	RequestPath						
				ArbReject					
	OPEN_REJECT	XL4:Open_Reject							
	idle dwords	XL0:Idle							

Figure 4. Open Reject by expander device

X.1.4 Connection Request – Arbitration Lost

Figure 5 shows two end devices attempting to establish a connection at the same time. This example assumes that the OPEN(A to B) address frame has higher priority than the OPEN(B to A) address frame and therefore device A wins arbitration and device B loses arbitration.

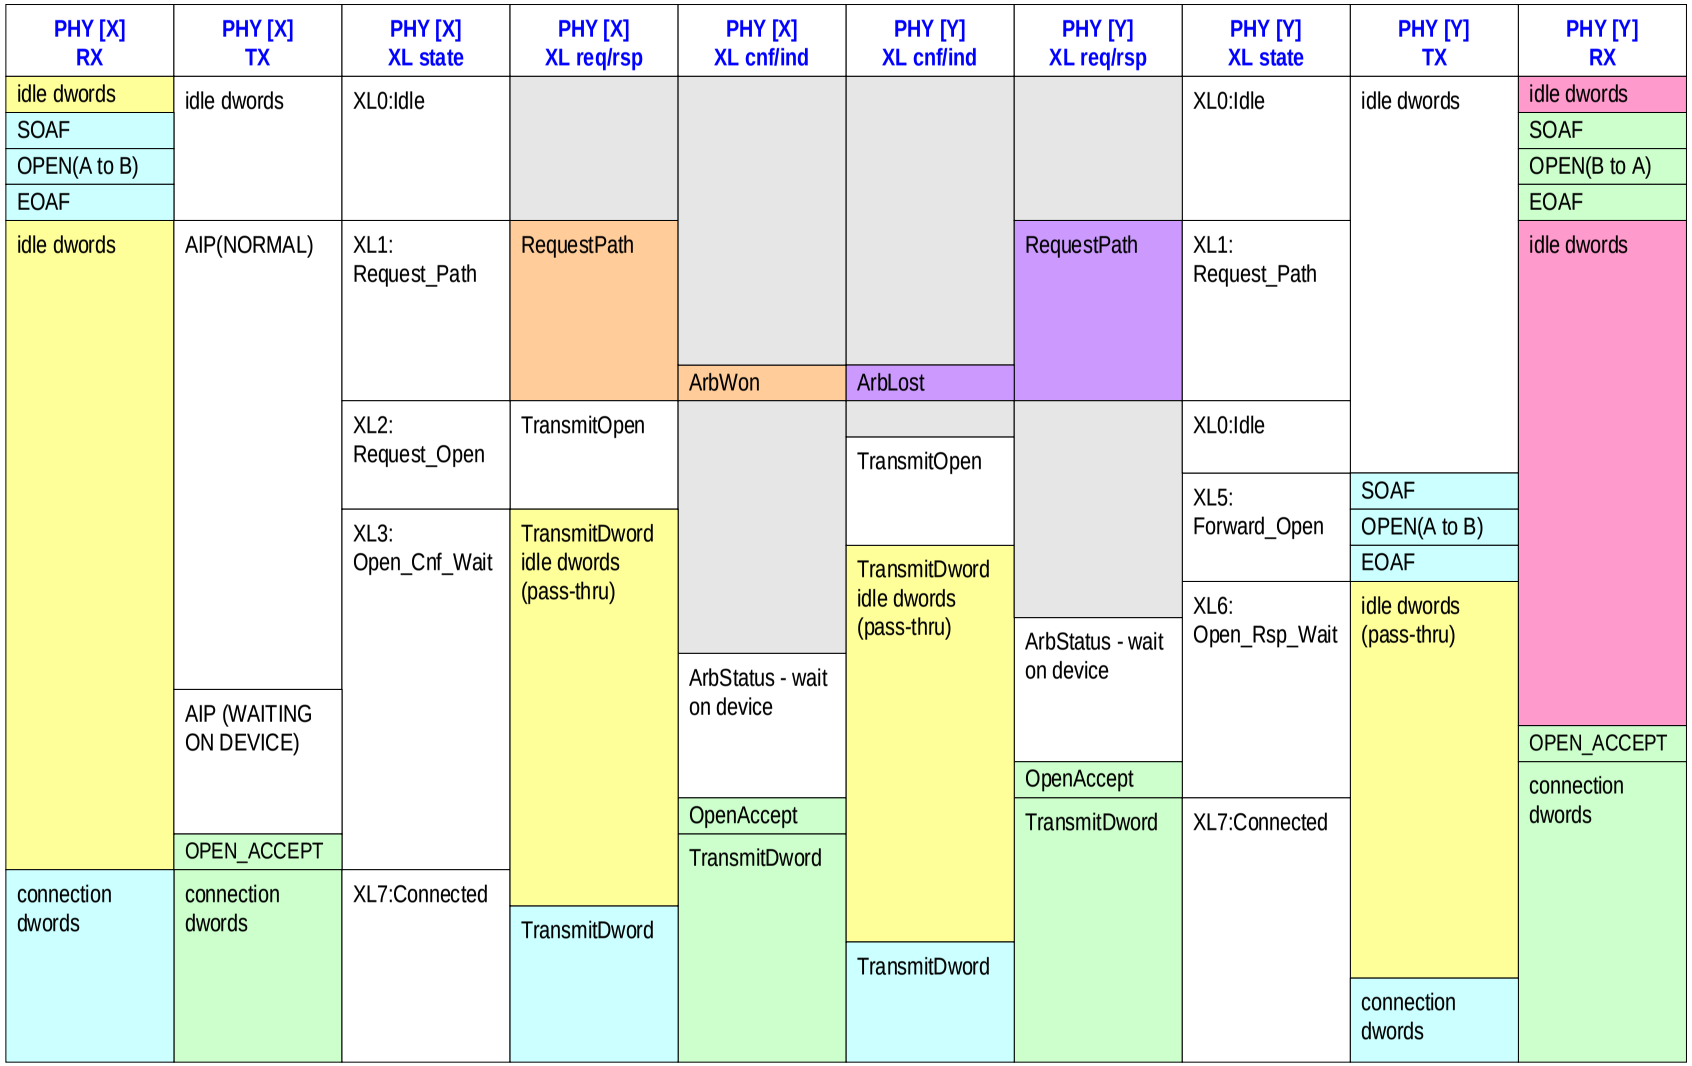

Figure 5. Arbitration Lost

X.1.5 Connection Request – Backoff and Retry

Figure 6 shows the condition which occurs when a higher priority OPEN address frame (B to C) is received by a phy which has previously forwarded an OPEN address frame to a different destination (A to B). In this case PHY[X] must back off and retry path arbitration.

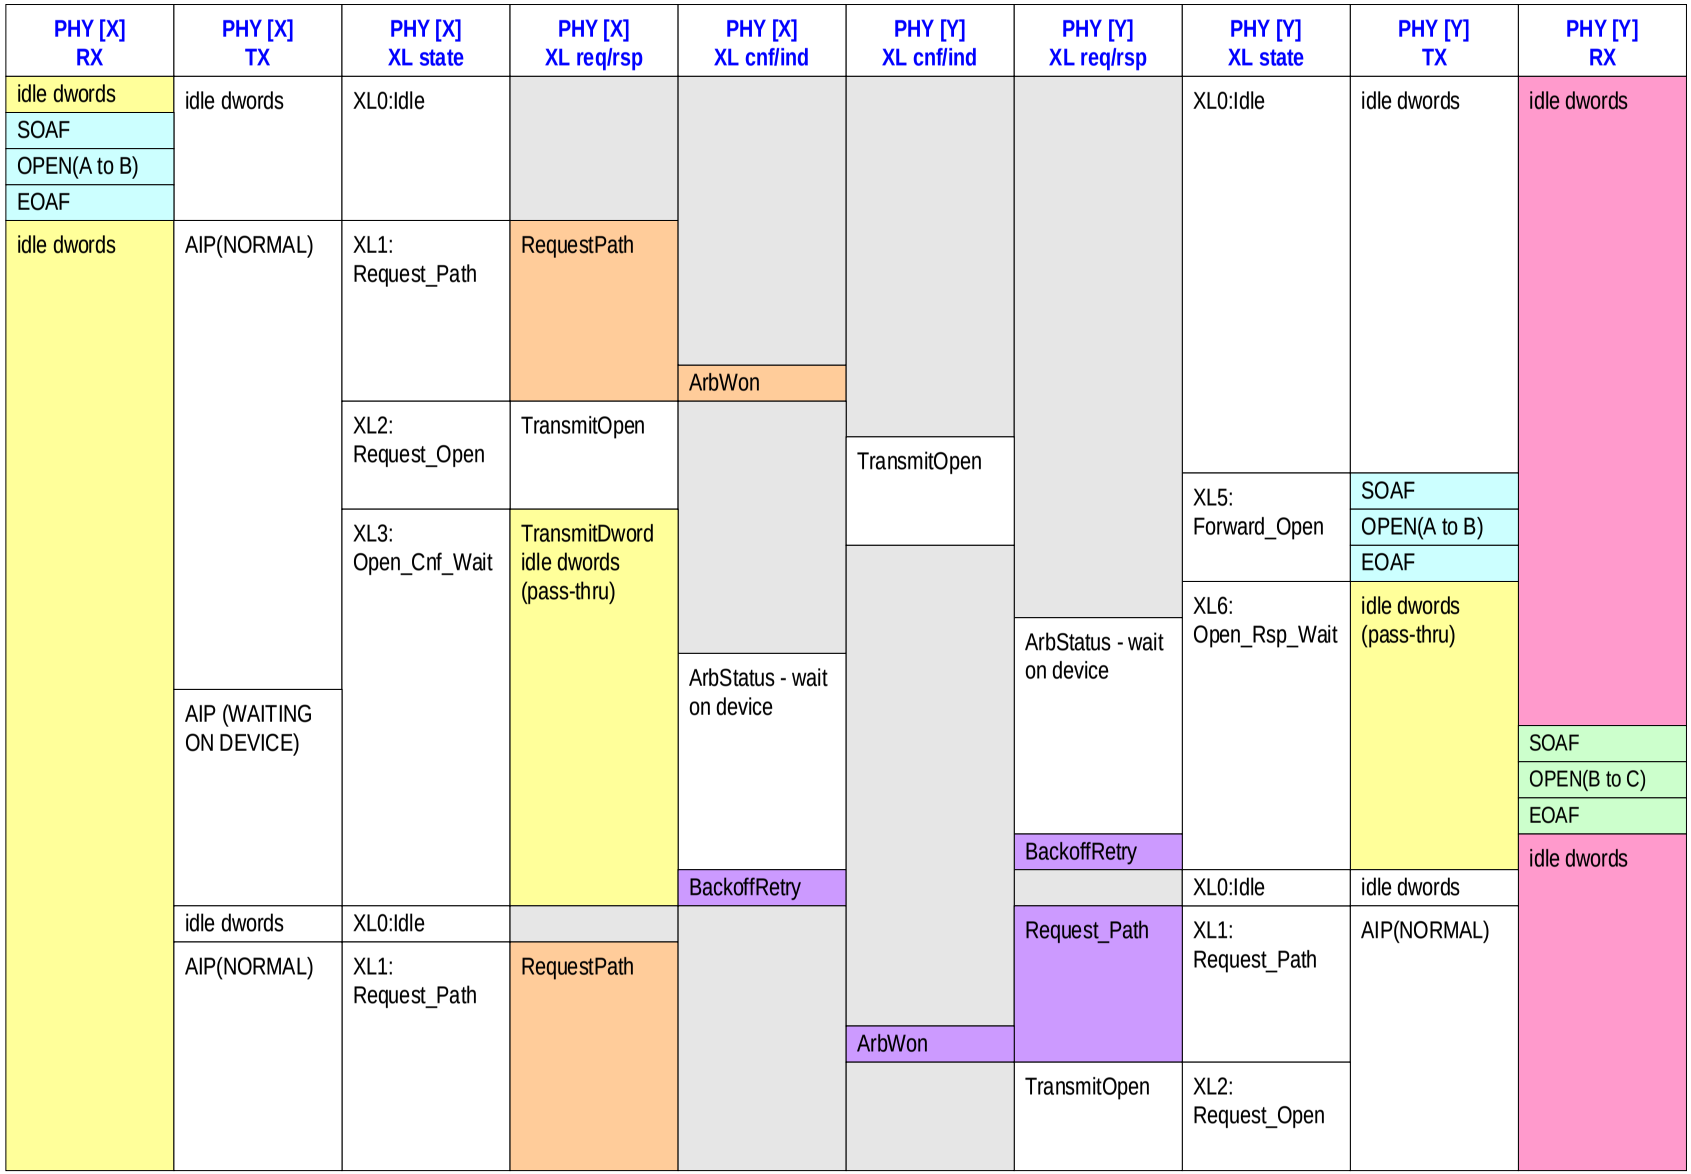

Figure 6. Backoff and Retry

X.1.6 Connection Request – Backoff and Reverse Path

Figure 7 shows the condition which occurs when a higher priority OPEN address frame (B to A) is received by a phy which has previously forwarded an OPEN address frame to the same destination (A to B). In this case PHY[Y] forwards the higher priority OPEN to PHY[X].

PHY [X] RX	PHY [X] TX	PHY [X] XL state	PHY [X] XL req/rsp	PHY [X] XL cnf/ind	PHY [Y] XL cnf/ind	PHY [Y] XL req/rsp	PHY [Y] XL state	PHY [Y] TX	PHY [Y] RX
idle dwords	idle dwords	XL0:Idle					XL0:Idle	idle dwords	idle dwords
SOAF									
OPEN(A to B)									
EOAF									
idle dwords	AIP(NORMAL)	XL1: Request_Path	RequestPath						
				ArbWon					
		XL2: Request_Open	TransmitOpen		TransmitOpen				
							XL5: Forward_Open	SOAF	
								OPEN(A to B)	
								EOAF	
		XL3: Open_Cnf_Wait	TransmitDword idle dwords (pass-thru)				XL6: Open_Rsp_Wait	idle dwords (pass-thru)	
						ArbStatus - wait on device			
				ArbStatus - wait on device					
	AIP (WAITING ON DEVICE)								SOAF
									OPEN(B to A)
									EOAF
						BackoffReverse			idle dwords
				BackoffReverse					
	idle dwords	XL0:Idle		TransmitOpen		TransmitOpen	XL2: Request_Open	AIP(NORMAL)	
	SOAF	XL5: Forward_Open							
	OPEN(B to A)						XL3: Open_Cnf_Wait		
	EOAF								
	idle dwords (pass-thru)	XL6: Open_Rsp_Wait	ArbStatus - wait on device						
					ArbStatus - wait on device			AIP (WAITING ON DEVICE)	

Figure 7. Backoff and Reverse Path

X.1.7 Connection Close – single step

Figure 8 shows an end device initiating the closing of a connection by transmitting CLOSE, followed by another end device responding with CLOSE at a later time.

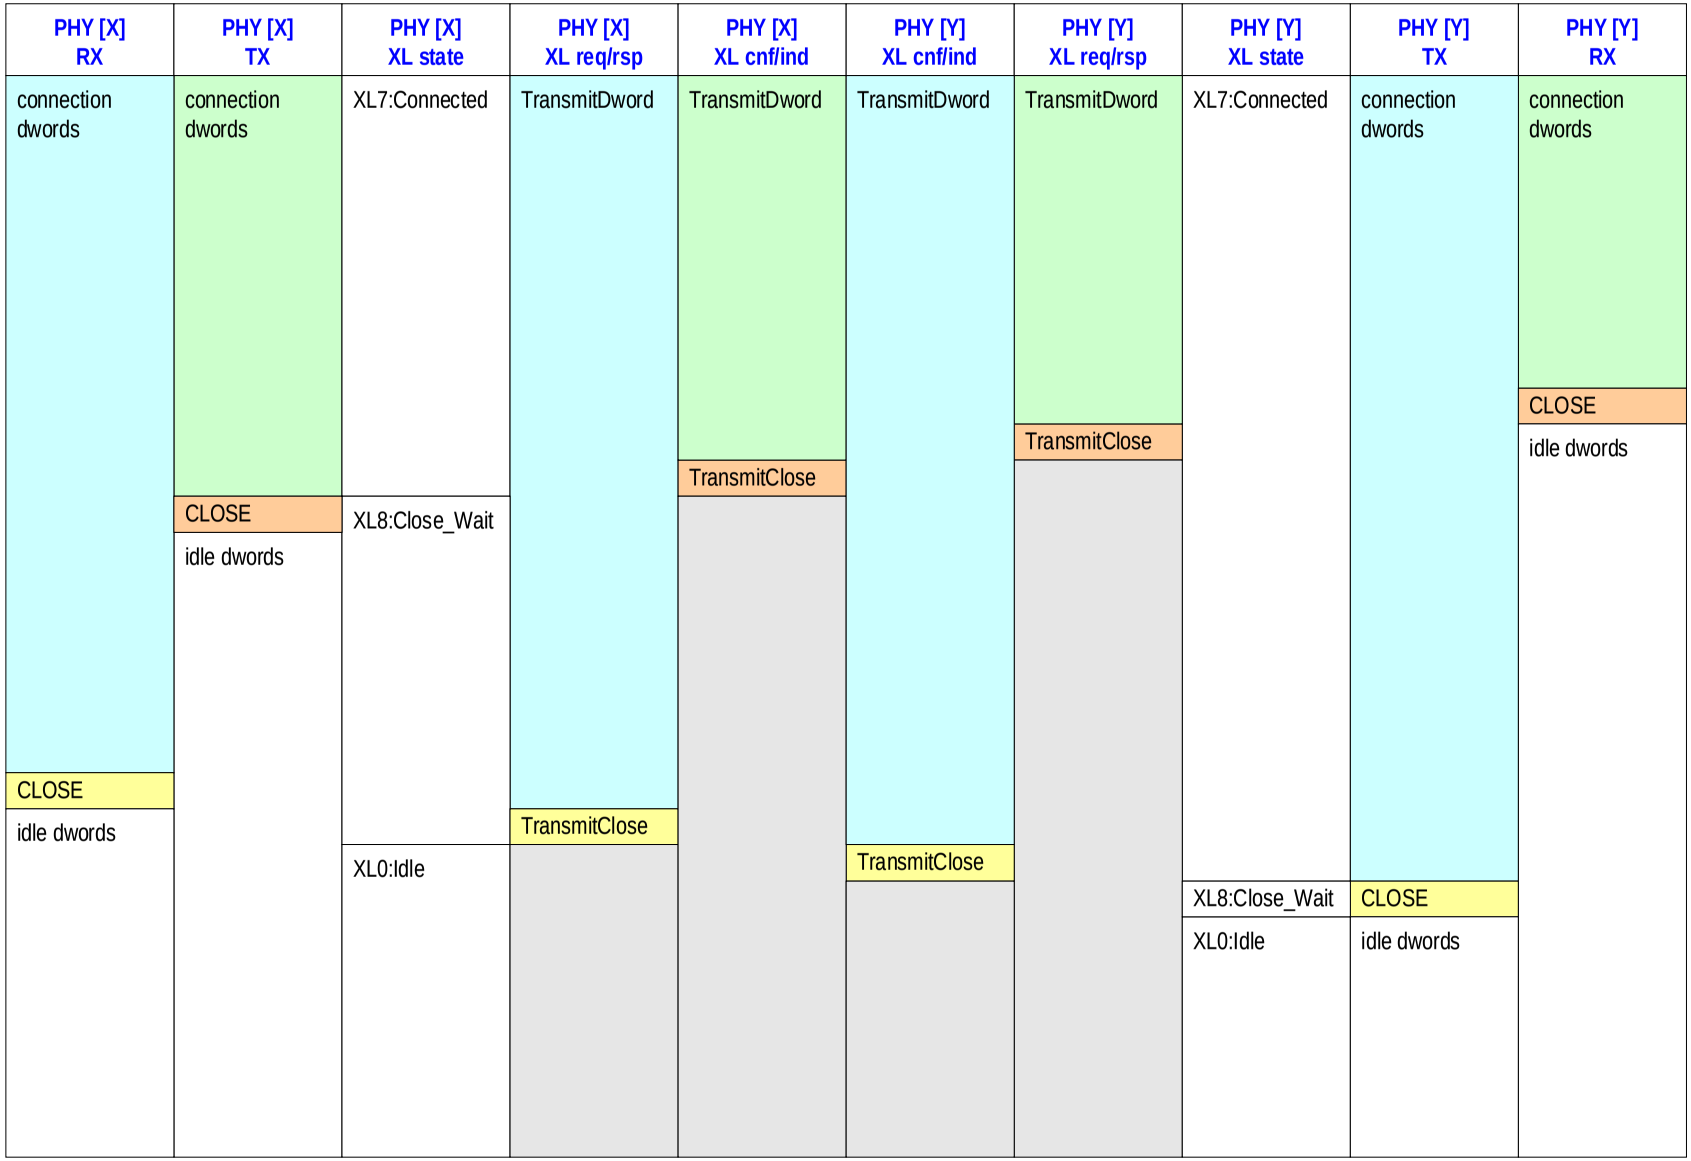

Figure 8. Connection Close – single step

X.1.8 Connection Close – simultaneous

Figure 9 shows two end devices simultaneously transmitting CLOSE to each other.

PHY [X] RX	PHY [X] TX	PHY [X] XL state	PHY [X] XL req/rsp	PHY [X] XL cnf/ind	PHY [Y] XL cnf/ind	PHY [Y] XL req/rsp	PHY [Y] XL state	PHY [Y] TX	PHY [Y] RX
connection dwords	connection dwords	XL7:Connected	TransmitDword	TransmitDword	TransmitDword	TransmitDword	XL7:Connected	connection dwords	connection dwords
CLOSE			TransmitClose			TransmitClose			CLOSE
idle dwords				TransmitClose	TransmitClose				idle dwords
	CLOSE	XL8:Close_Wait					XL8:Close_Wait	CLOSE	
	idle dwords	XL0:Idle					XL0:Idle	idle dwords	

Figure 9. Connection Close - simultaneous

X.1.9 Break handling during path arbitration

Figure 10 shows an expander device responding to the reception of a BREAK primitive sequence during path arbitration.

PHY [X] RX	PHY [X] TX	PHY [X] XL state	PHY [X] XL req/rsp	PHY [X] XL cnf/ind	PHY [Y] XL cnf/ind	PHY [Y] XL req/rsp	PHY [Y] XL state	PHY [Y] TX	PHY [Y] RX
idle dwords	idle dwords	XL0:Idle					XL0:Idle	idle dwords	idle dwords
SOAF									
OPEN(A to B)									
EOAF									
idle dwords	AIP(NORMAL)	XL1: Request_Path	RequestPath						
BREAK									
	BREAK	XL9:Break							
idle dwords	idle dwords	XL0:Idle							

Figure 10. Break handling during arbitration

X.1.10 Break handling during connection

Figure 11 shows an expander device responding to the reception of a BREAK primitive sequence during a connection.

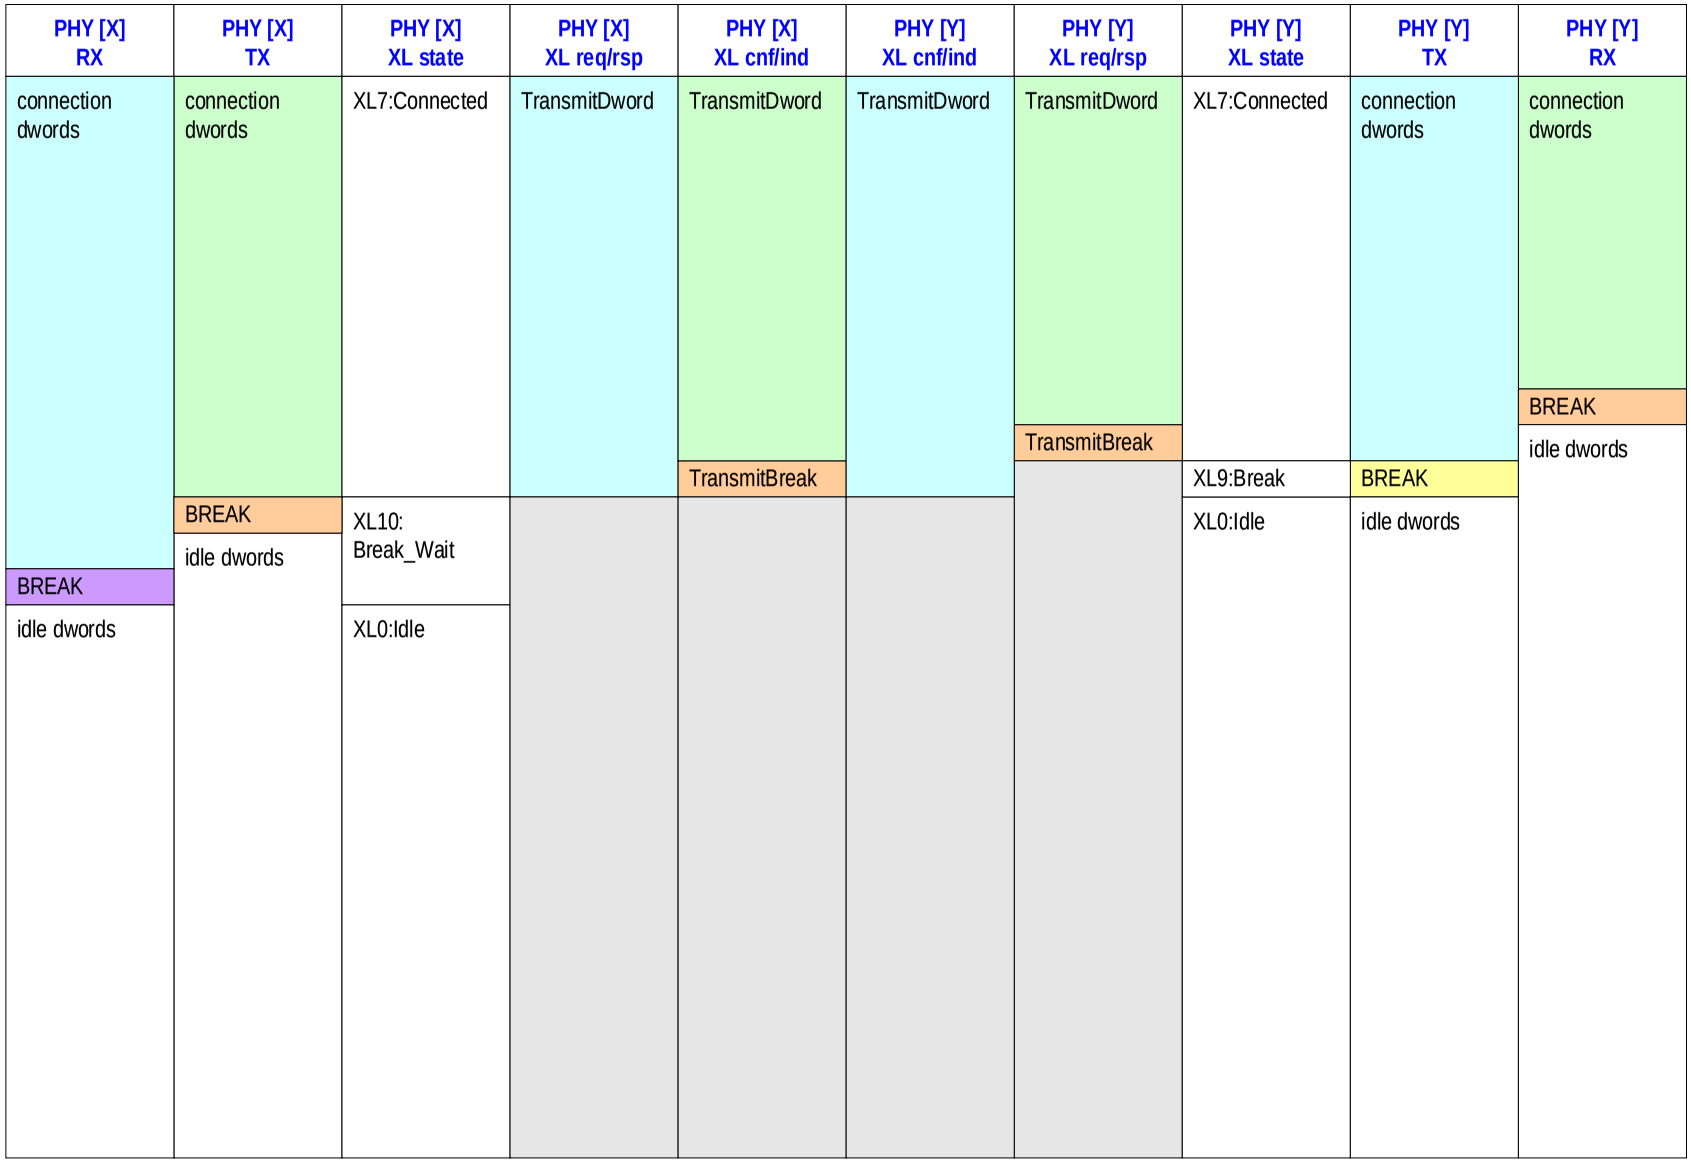

Figure 11. Break handling during connection

X.2 STP connection examples

Figure 12 shows the topology used by examples X.2.1 through X.2.3 below.

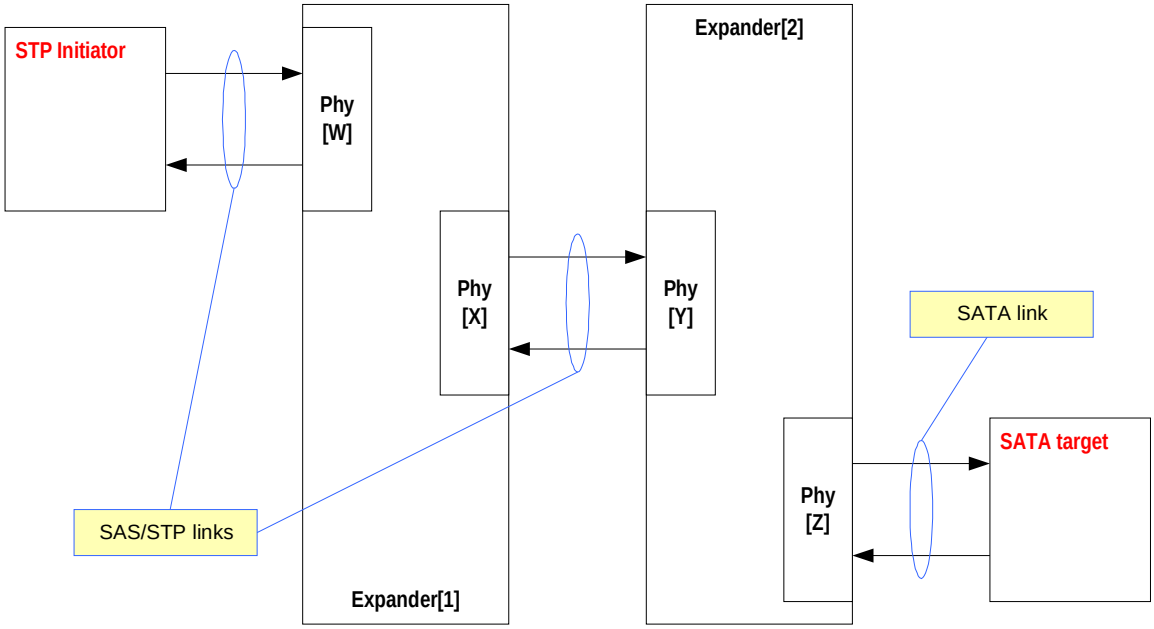

Figure 12: Example Topology #2

Table 2 defines the columns used within Figures X.1.1 through X.1.10.

Column Header	Description
PHY [W] RX	Expander [1] Phy [W] receive from STP initiator
PHY [W] TX	Expander [1] Phy [W] transmit to STP initiator
PHY [X] TX	Expander [1] Phy [X] transmit to Expander[2] Phy [Y]
PHY [X] RX	Expander [1] Phy [X] receive from Expander[2] Phy [Y]
PHY [Y] RX	Expander [2] Phy [Y] receive from Expander[1] Phy [X]
PHY [Y] TX	Expander [2] Phy [Y] transmit to Expander[1] Phy [X]
PHY [Z] TX	Expander [2] Phy [Z] transmit to SATA target
PHY [Z] RX	Expander [2] Phy [Z] receive from SATA target

Table 2: Column descriptions for STP connection examples

X.2.1 STP Connection – originated by STP initiator

Figure 13 shows an STP initiator originating a connection and frame to an idle SATA target.

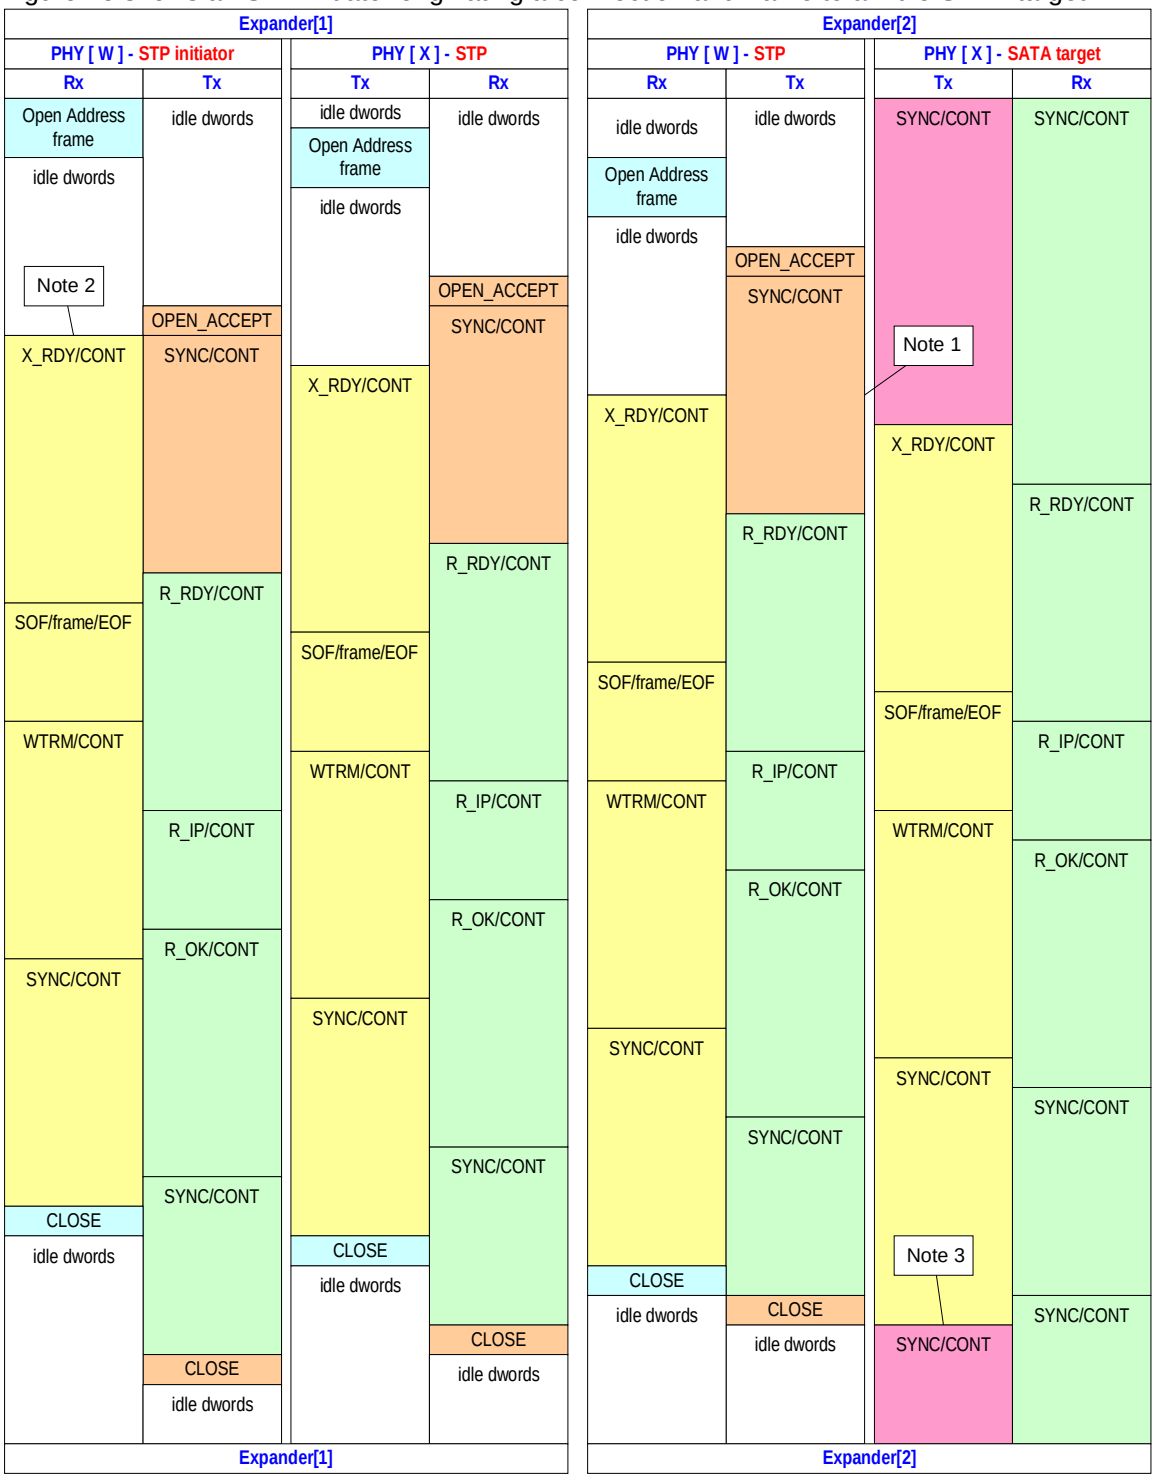

Figure 13. STP Connection – originated by STP initiator

Notes:

1. Expander must duplicate the dword stream which is being received from the SATA target.
2. SATA_X_RDY not transmitted until STP connection exists, i.e. after OPEN_ACCEPT.
3. After receiving CLOSE, expander must transmit SATA_SYNC/CONT to SATA target.

X.2.2 STP Connection – originated by STP initiator with SATA contention

Figure 14 shows an STP initiator originating a connection to a SATA target which is requesting to transmit a frame.

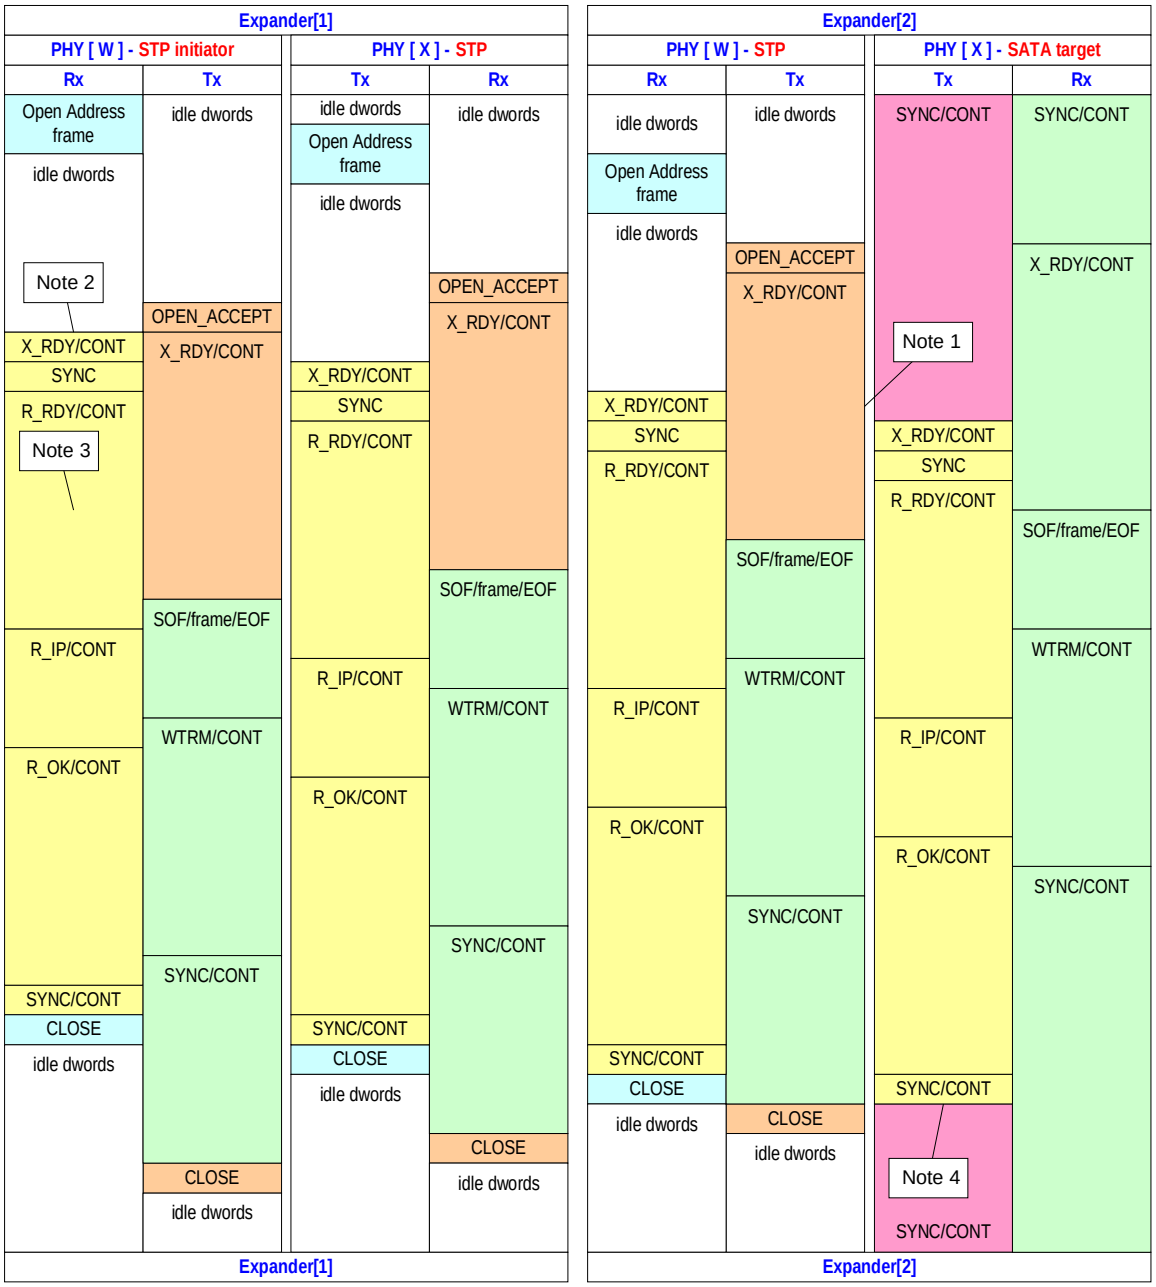

Figure 14. STP Connection – originated by STP initiator with SATA contention

Notes:

1. Expander must duplicate the dword stream which is being received from the SATA target (X_RDY/CONT).
2. SATA_X_RDY not transmitted until STP connection exists, i.e. after OPEN_ACCEPT.
3. SATA rules require STP initiator to remove request to transmit if SATA target is transmitting X_RDY.
4. After receiving CLOSE, expander must transmit SATA_SYNC/CONT to SATA target.

X.2.3 STP Connection – originated by expander on behalf of SATA target

Figure 15 shows an expander originating a connection on behalf of a SATA target which is requesting to transmit a frame.

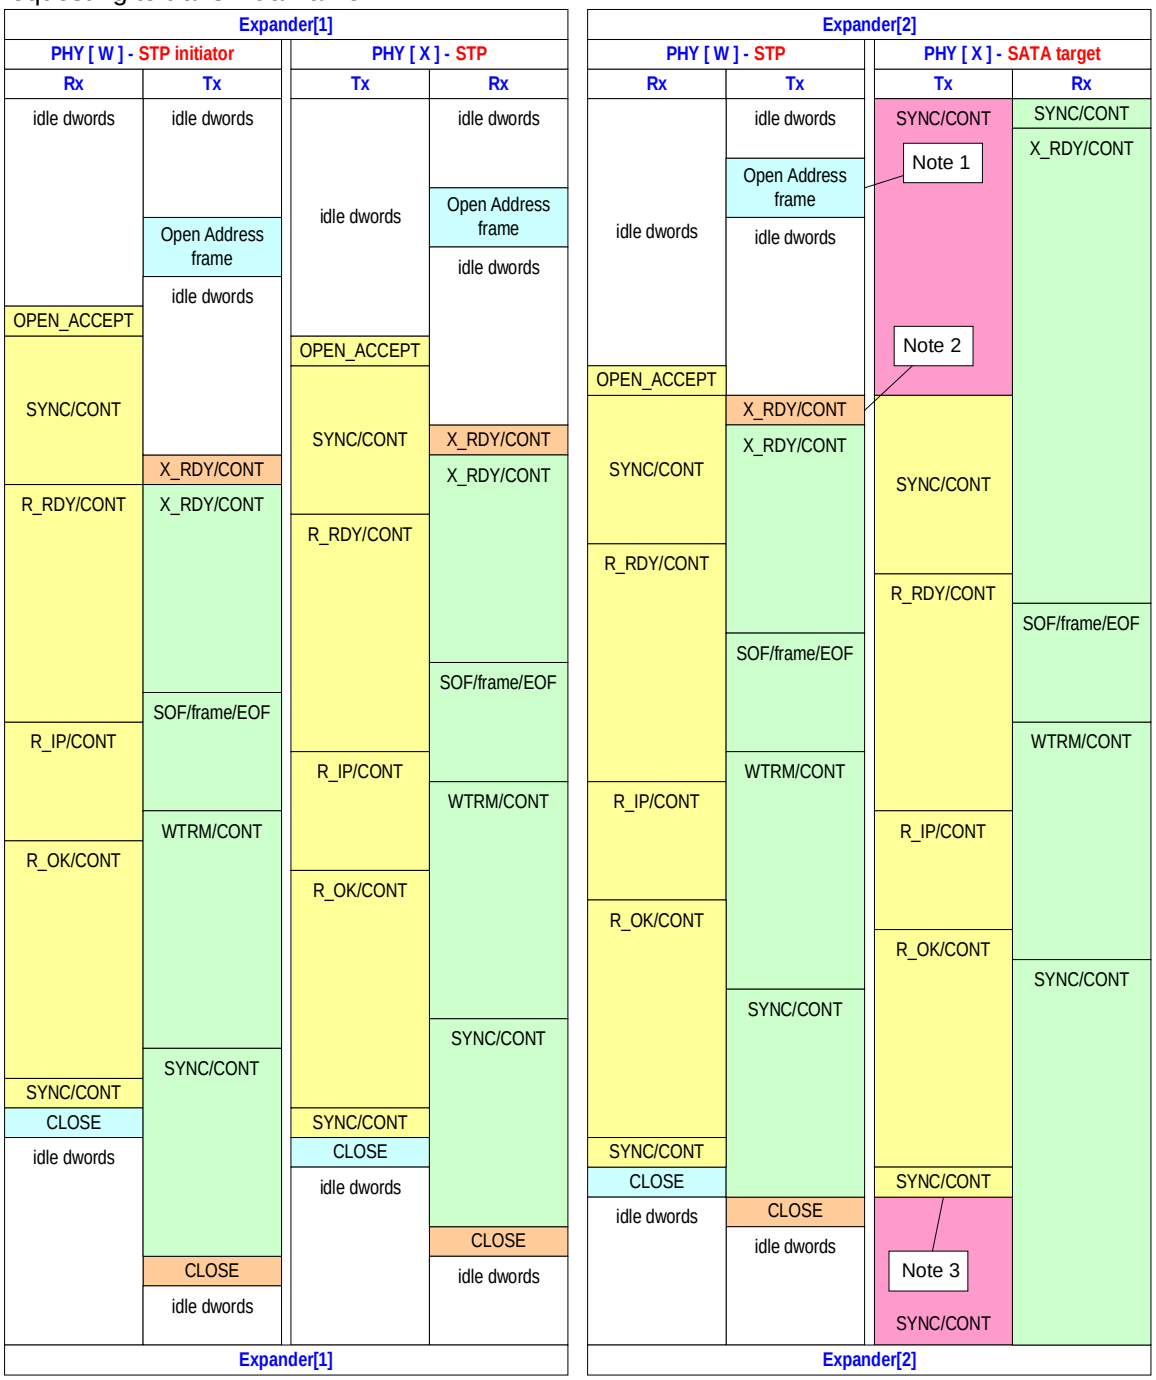

Figure 15. STP Connection – originated by expander on behalf of SATA target

Notes:

1. Expander must generate an OPEN address frame on behalf of the SATA target.
2. Expander must duplicate the dword stream which is being received from the SATA target (X_RDY/CONT).
3. After receiving CLOSE, expander must transmit SATA_SYNC/CONT to SATA target.

X.3 Pathway blocked and recovery example

Figure 16 shows the topology used to illustrate pathway recovery:

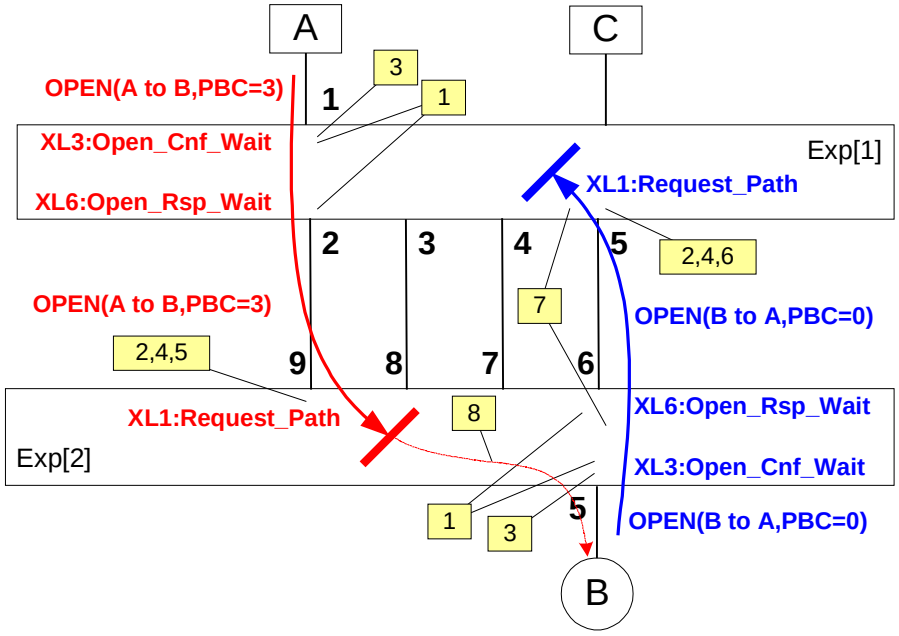

Figure 16. Partial Pathway Recovery

The sequence of events used to identify pathway blockage and to perform pathway recovery are as follows:

1. Exp[1].Phy[1,2] and Exp[2].Phy[5,6] Transmit Phy Status (Partial Pathway) to Connection Manager to indicate that they contain partial pathways.
2. Exp[1].Phy[5] and Exp[2].Phy[9] receive Arbitrating (Waiting On Partial) from Connection Manager which cause them to transmit AIP(WAITING ON PARTIAL).
3. AIP(WAITING ON PARTIAL) received by Exp[1].Phy[2] and Exp[2].Phy[6] then forwarded to Exp[1].Phy[1] and Exp[2].Phy[5] as Arb Status (Waiting On Partial). Exp[1].Phy[1] and Exp[2].Phy[5] Transmit Phy Status (Blocked On Partial) to Connection Manager as confirmation that they are blocked waiting on a partial pathway in another expander.
4. Exp[1].Phy[5] and Exp[2].Phy[9] receive Arbitrating (Blocked On Partial) from Connection Manager while all destination phys contain a Phy Status of Blocked On Partial which cause them to run their Partial Pathway Timeout timers.
5. PPT expires in Exp[2].Phy[9] – this causes a request to the Connection Manager to resolve pathway blockage. Pathway recovery priority for this phy is not lower than all phys within the destination port which are also blocked. Connection Manager does not provide Arb Reject (Pathway Blocked) to Exp[2].Phy[9] so this phy waits for pathway resolution to occur elsewhere in the topology.
6. PPT expires in Exp[1].Phy[5] – this causes a request to the Connection Manager to resolve pathway blockage. Pathway recovery priority for this phy is lower than all phys within the destination port which are also blocked. Connection Manager provides Arb Reject (Pathway Blocked) to Exp[1].Phy[5] which instructs this phy to reject the connection request using OPEN_REJECT(PATHWAY BLOCKED).
7. OPEN_REJECT(PATHWAY BLOCKED) tears down partial pathway all the way to the originating end device (Device B).
8. Partial pathway established between Exp[2].Phy[9] and Exp[2].Phy[5] and OPEN(A to B) is delivered to Device B.