To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 24 September 2007
Subject: 07-090r3 SAS-2 Transmit IDENTIFY three times

Revision history

Revision 0 (26 April 2007) First revision
Revision 1 (7 May 2007) Incorporated comments from May 2007 SAS protocol WG.
Revision 2 (13 August 2007) Changed SL_IR_IRC to notify the link layer that it is ready to go after transmitting the first IDENTIFY address frame rather than wait until it has transmitted the third.
Revision 3 (24 September 2007) Incorporated comments from September 2007 SAS protocol WG - in 4.3.2, show that the transmit MUX doesn't switch over until all 3 IDENTIFYs have been transmitted, and add cross-references between 4.3.2/4.3.3 and the transmitter/receiver subclauses. Included lists before the figures listing the mentioned state machines/transmitters/receivers, since the figures themselves cannot include cross-references.

Related documents

sas2r09a - Serial Attached SCSI - 2 (SAS-2) revision 9a
05-086r0 SAS-1.1 Link layer timeout race conditions (Rob Elliott, HP)
05-094r0 SAS Protocol WG minutes 7 March 2005 (Ralph Weber, ENDL, and John Lohmeyer, LSI)
07-334 SAS-2 Add minimum number of dwords after IDENTIFY frame (Gerry Houlder, Seagate)

Overview

In most cases, SAS recovers fairly gracefully from single-bit errors. If a single bit error occurs in an IDENTIFY address frame, however, the only recovery mechanism available is to rerun the entire link reset sequence (OOB + speed negotiation). Although the IDENTIFY address frame includes a CRC, there is no NAK to trigger retransmission on failed reception.

To better tolerate single-bit errors, the IDENTIFY address frame should be sent 3 times rather than just 1 time. Transmitting 3 times rather than 2 allows for DFE receiver error expansion of a single-bit error on the wire to a short burst of errors in the receiver and allows the frames to be sent back-to-back (SOAF ... EOAF, SOAF ... EOAF, SOAF ... EOAF).

This should not confuse a SAS-1.1 receiver, which is supposed to only honor the first SOAF it sees. As noted by Jeff Gauvin (LSI), the SL_IR state machine in SAS-1.1 (which handles IDENTIFY address frames) differs from the SL_RA state machine (which handles OPEN address frames) in its handling of unexpected address frames and SOAFs before EOAFs. SL_IR only honors the first SOAF it sees, while SL_CC restarts on any SOAF. It is likely that some SL_IR implementations don't make that distinction and restart on each SOAF, to share logic with SL_RA. If they do make the distinction, then SAS-2 to SAS-1.1 just falls back to the SAS-1.1 case of rerunning the link reset sequence if an error occurs.

This proposal:

a) Changes SL_IR_RIF to restart reception of an address frame on any SOAF
b) Changes SL_IR_RIF to keep receiving address frames until it finds a correct IDENTIFY address frame. It does not quit if the CRC is bad.
c) Allows SL_IR_TIF to transmit either one or three IDENTIFY address frames.

Two complaints were lodged against revision 1 of this proposal in the May T10 plenary (where the vote failed):

1. Frames too fast. 07-334 asks that 3 idle dwords be included after each IDENTIFY address frame to avoid triggering a bug in existing SAS-1.1 designs that would be exacerbated by sending three IDENTIFYs back-to-back.

2. Steve Finch complained that there is no maximum time between frames; the transmitter has to understand the receiver timeout and send the frames fast enough to meet it, but there is no advice in the transmitter section about how long to take. However, this is no different than the rest of the standard; on 3/7/2005, discussing proposal 05-086, the SAS protocol WG voted 9-3 to not define any transmitter time limits in SSP state machines and leave the standard vague. Unless this is changed everywhere, then just including time limits in SL_IR could cause confusion (e.g., why are there in SL_IR but not SSP?).

Suggested changes to SAS-2

4.3.2 Transmit data path

Figure 31 shows the transmit data path in a SAS phy, showing the relationship between:

a) the SP state machine (see 6.8) and the SP transmitter (see 6.8.2);
b) multiplexing (see 6.10);
c) the SL_IR state machines (see 7.9.4) and the SL_IR transmitter (see 7.9.4.2);
d) physical link rate tolerance management (see 7.3);
e) the SL state machines (see 7.14) and the SL transmitter (see 7.14.2);
f) rate matching (see 7.13); and
g) the SSP transmit data path (see figure 32), SMP transmit data path (see figure 33), and STP transmit data path (see figure 34).

Editor's Note 1: Note - there is no SP_DWS transmitter (just a receiver).

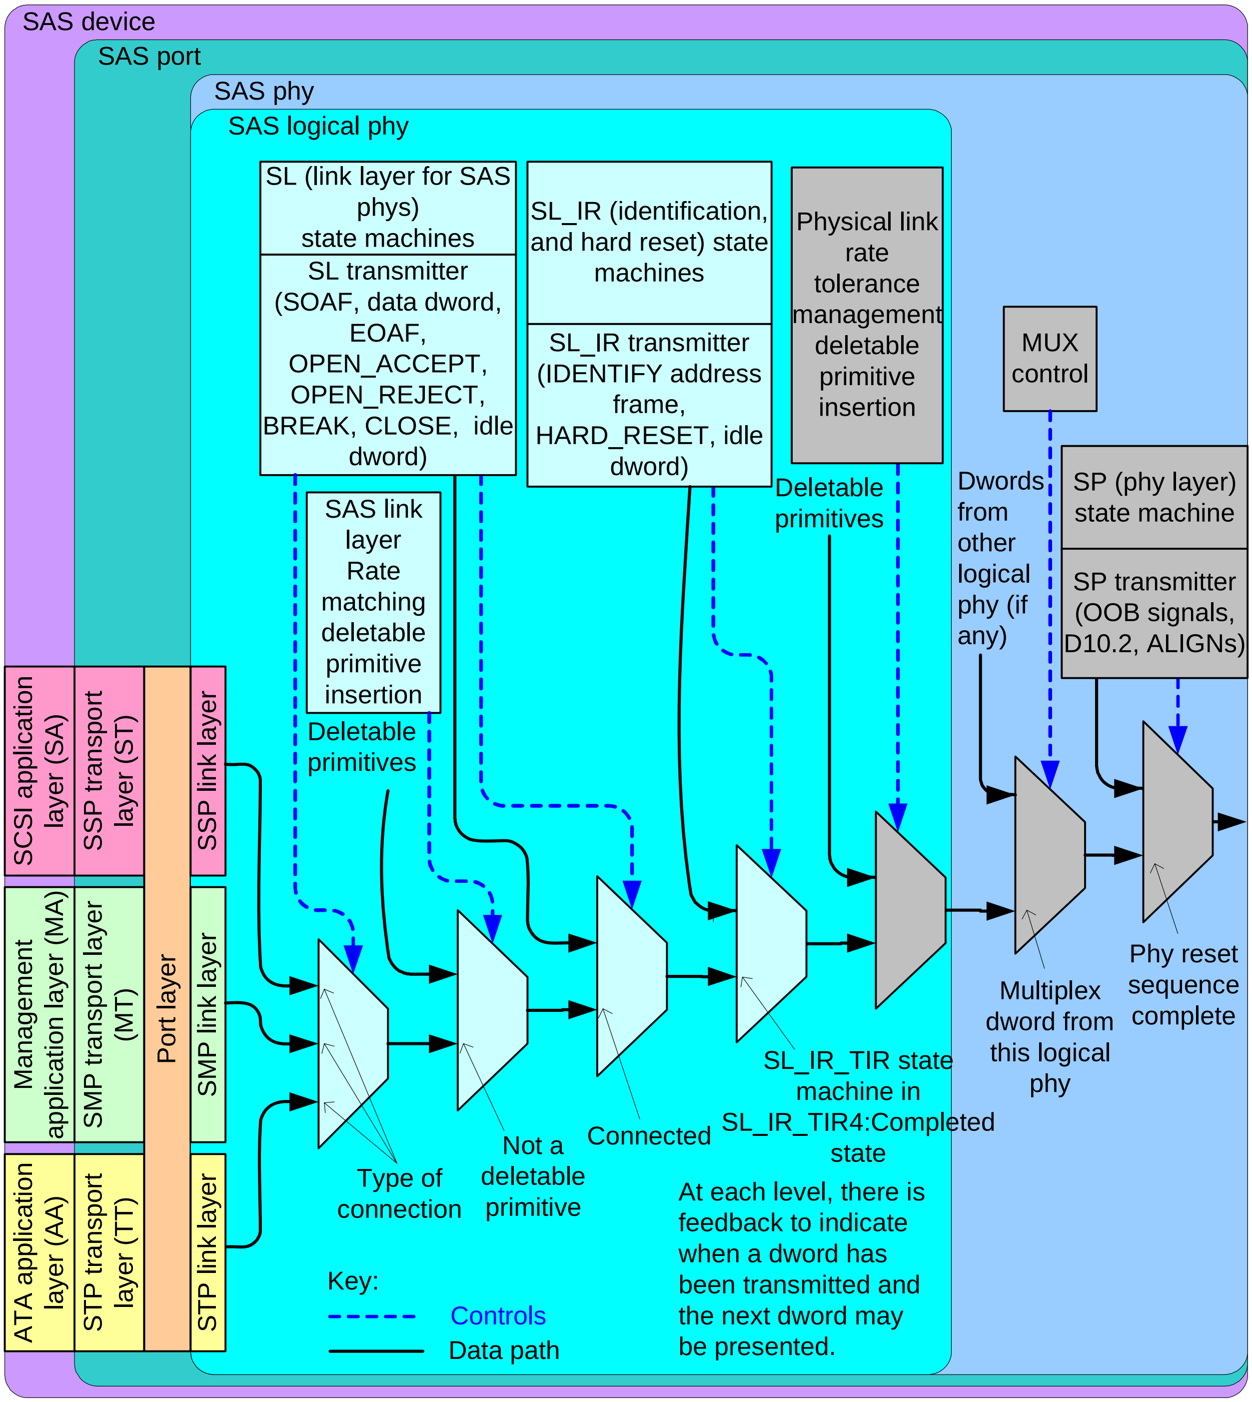

Figure 31 — Transmit data path in a SAS phy [changed]

Editor's Note 2: In the figure, changed "Link reset sequence complete" to "SL_IR_TIR state machine in SL_IR_TIR4:Completed state". Changed "Clock skew management" to "Physical link rate tolerance management". Changed "SL_IR state machine" to "SL_IR state machines." Changed "SL state machine" to "SL state machines."

Figure 32 shows the transmit data path for the SSP link layer, including the SSP state machines and the SSP transmitter (see 7.16.8 and 7.16.8.2), and the communication to the port layer, SSP transport layer, and SCSI application layer. Only the SSP link layer (i.e., not the port, transport, or application layer) transmits dwords.

Figure 32 — [not shown: SSP link, port, SSP transport, and SCSI application layer state machines]

Figure 33 shows the transmit data path for the SMP link layer including the SMP state machines and the SMP transmitter (see 7.18.5 and 7.18.5.2), and the communication to the port layer, SMP transport layer, and management application layer. Only the SMP link layer (i.e., not the port, transport, or application layer) transmits dwords.

Figure 33 — [not shown: SMP link, port, SMP transport, and management application layer state machines]

Figure 34 shows the transmit data path for the STP link layer including the STP state machines and the STP transmitter (see 7.17.8), and communication to the port layer, STP transport layer, and ATA application layer. Only the STP link layer (i.e., not the port, transport, or application layer) transmits dwords.

Figure 34 — [not shown: STP link, port, STP transport, and ATA application layer state machines]

Figure 35 shows the transmit data path in an expander phy, showing the relationship between:

a) the SP state machine (see 6.8) and the SP transmitter (see 6.8.2);
b) multiplexing (see 6.10);
c) the SL_IR state machines (see 7.9.4) and the SL_IR transmitter (see 7.9.4.2);
d) physical link rate tolerance management (see 7.3); and
e) the XL state machine (see 7.14) and the XL transmitter (see 7.15.2).

Editor's Note 3: Note - there is no SP_DWS transmitter (just a receiver).

.

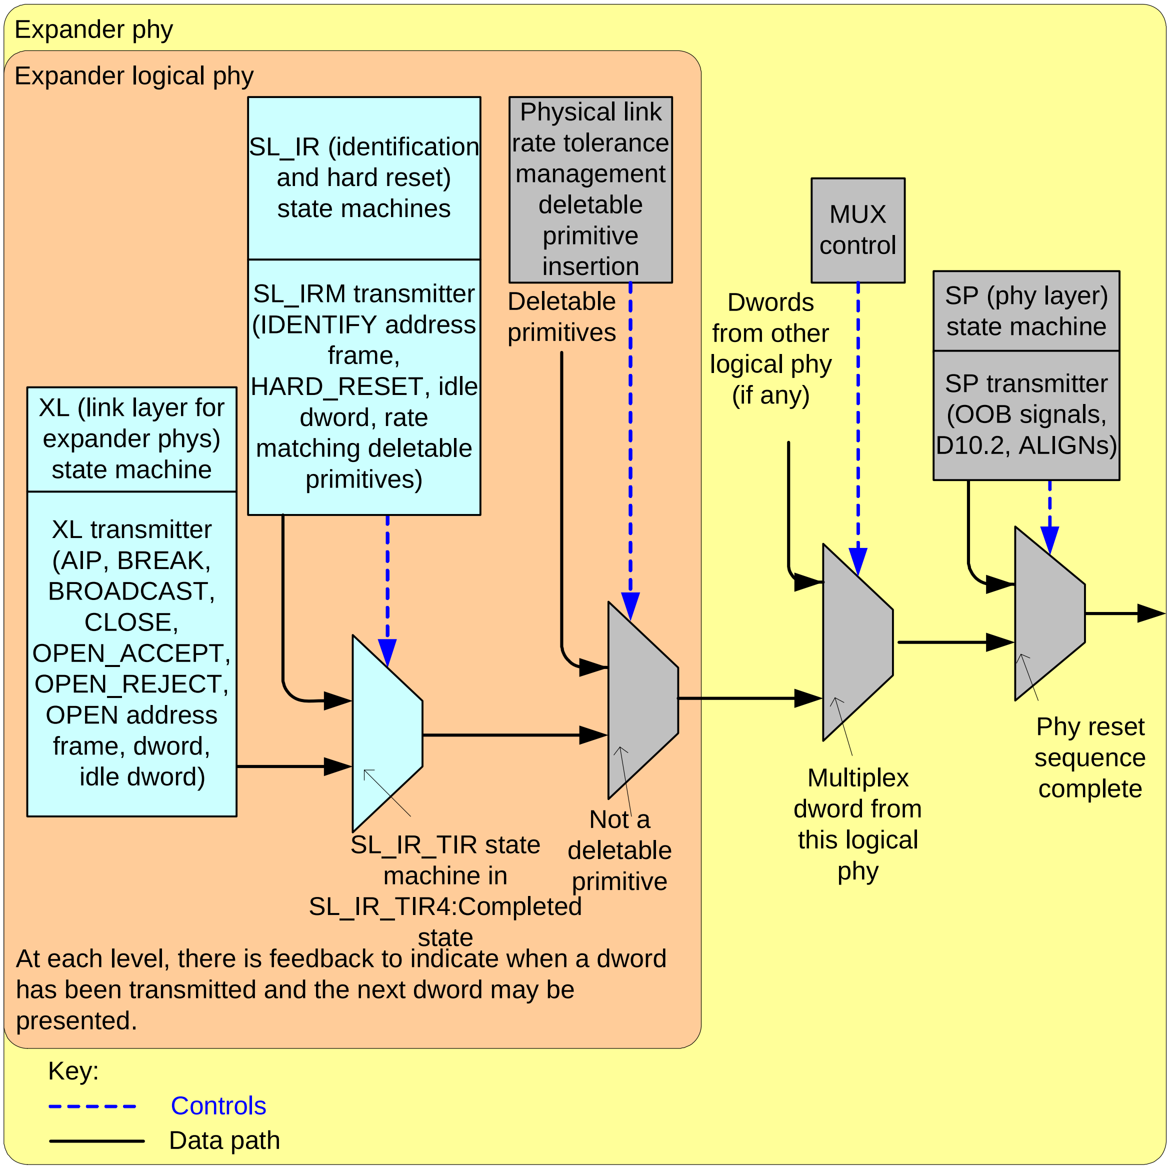

Figure 35 — Transmit data path and state machines in an expander phy [changed]

Editor's Note 4: In the figure, changed "Link reset sequence complete" to "SL_IR_TIR state machine in SL_IR_TIR4:Completed state." Changed "Clock skew management" to "Physical link rate tolerance management." Changed "SL_IR state machine" to "SL_IR state machines."

4.3.3 Receive data path

The SP_DWS receiver (see 6.9.2) establishes dword synchronization and sends dwords to the SP_DWS state machine (see 6.9) and to the link layer state machine receivers.

If multiplexing (see 6.10) is enabled (see table 90 in 6.7.4.2.3.3), the SP_DWS receiver uses incoming MUX primitives to determine the logical link numbers and route the dwords to the appropriate link layer receivers.

Figure 36 shows the receive data path in a SAS phy, showing the relationship between:

a) SP_DWS receiver (see 6.8.2);
b) SL_IR receiver (see 7.9.4.2);

c) SL receiver (see 7.14.2);
d) SSP receiver (see 7.16.8.2), STP receiver (see 7.17.8), and SMP receiver (see 7.18.5.2).

Editor's Note 5: The SP receiver is not currently shown in this figure.

.

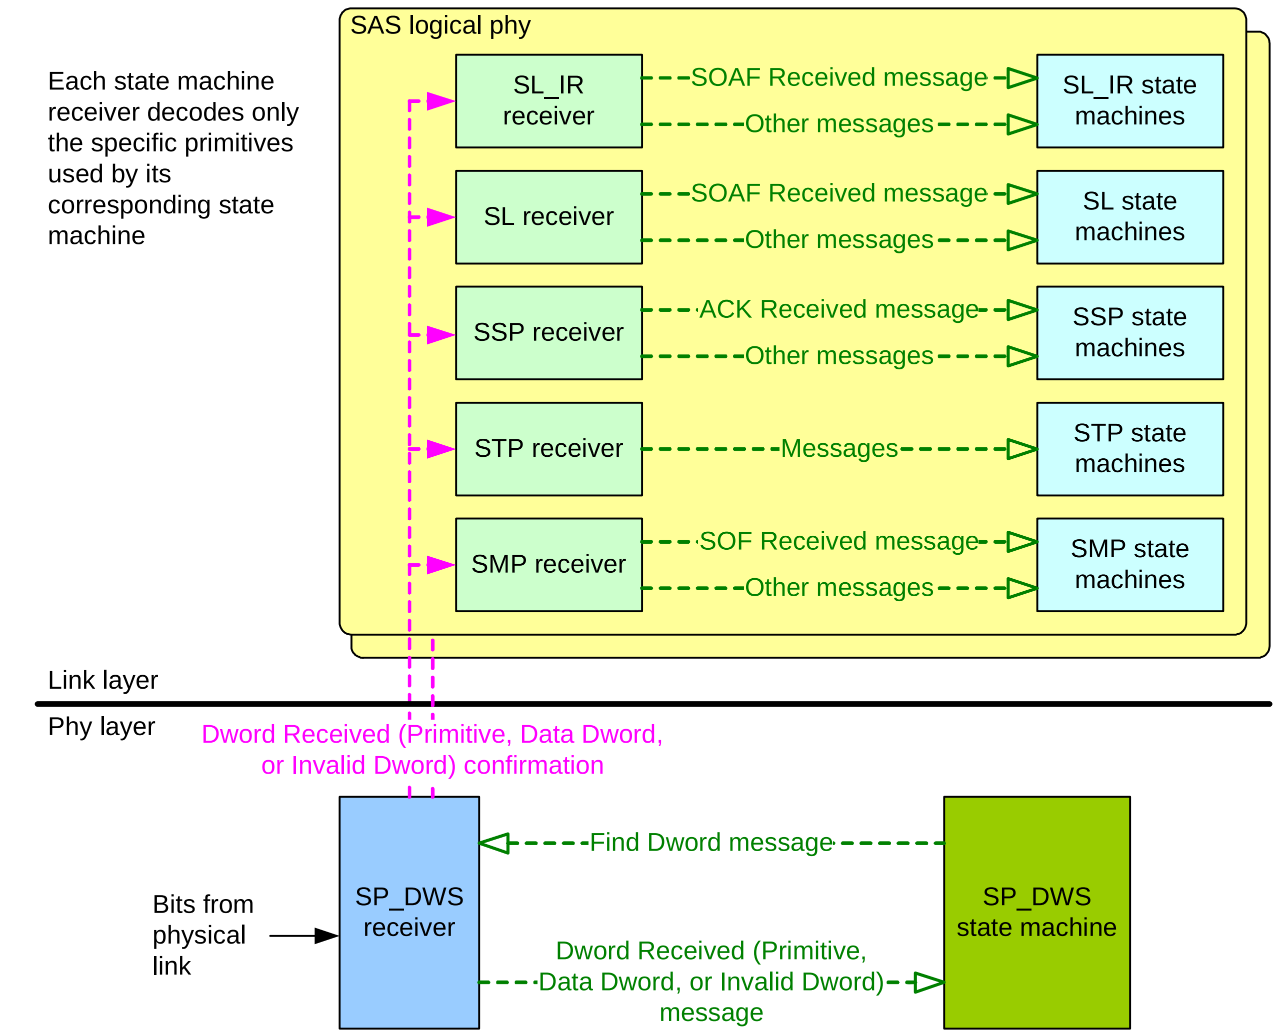

Figure 36 — Receive data path in a SAS phy [changed]

Editor's Note 6: Changed "SL_IR state machine" to "SL_IR state machines"

Figure 37 shows the receive data path in an expander phy, showing the relationship between:

a) SP_DWS receiver (see 6.9.2);
b) SL_IR state machines and SL_IR receiver (see 7.9.4.2); and
c) XL state machine and XL receiver (see 7.15.2);

Editor's Note 7: The SP receiver is not currently shown in this figure.

.

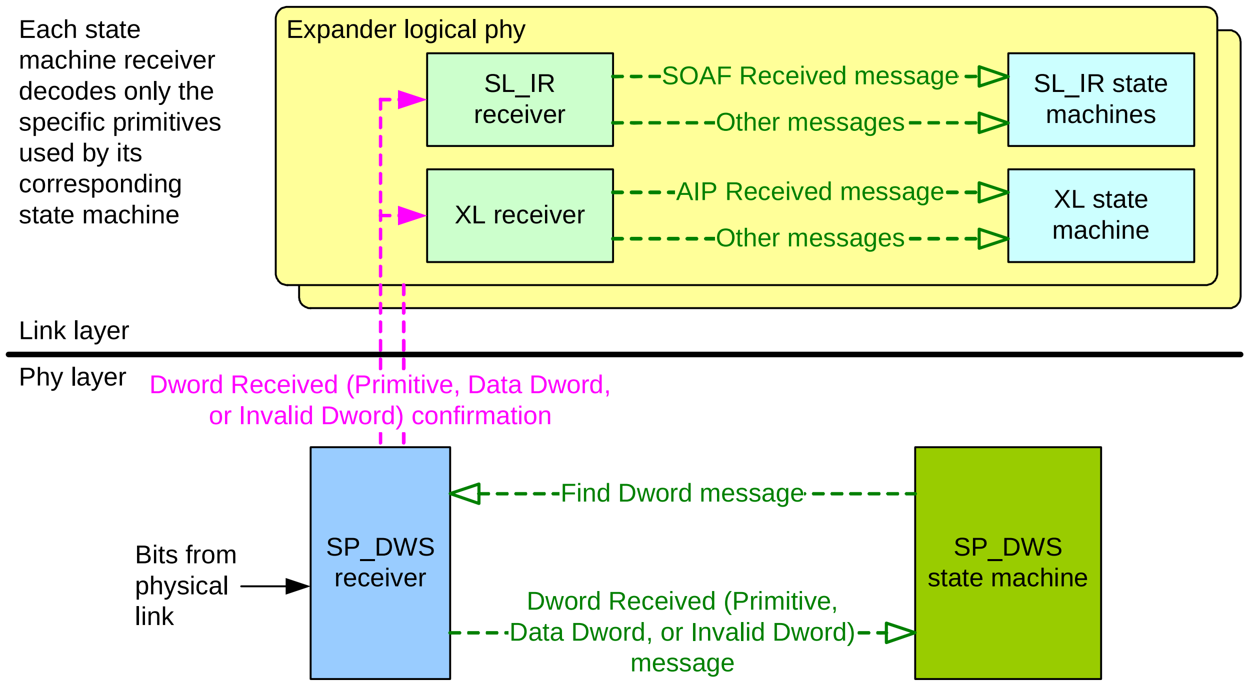

Figure 37 — Receive data path in an expander phy [changed]

Editor's Note 8: Changed "SL_IR state machine" to "SL_IR state machines"

6.8.2 SP transmitter and receiver

The SP transmitter transmits OOB signals and dwords on the physical link based on messages from the SP state machine (see 6.8).

The SP transmitter receives the following messages from the SP state machine:

a) Transmit COMINIT;
b) Transmit COMSAS;
c) Transmit COMWAKE;
d) Transmit SATA Port Selection Signal;
e) Transmit D10.2;
f) Set Rate with a Physical Link Rate argument, and an SSC On or an SSC Off argument;
g) Transmit ALIGN with an argument indicating the specific type (e.g., Transmit ALIGN (0));
h) Transmit Phy Capabilities Bits;
i) Transmit TRAIN Pattern;
j) Transmit TRAIN_DONE Pattern; and
k) Transmit MUX Sequence.

When not otherwise instructed, the SP transmitter transmits D.C. idle.

Upon receiving a Transmit MUX Sequence message, the SP transmitter transmits:

1) MUX (LOGICAL LINK 0);
2) MUX (LOGICAL LINK 1);
3) MUX (LOGICAL LINK 0);
4) MUX (LOGICAL LINK 1);
5) MUX (LOGICAL LINK 0); and
6) MUX (LOGICAL LINK 1).

The SP transmitter shall complete any physical link rate change requested with the Set Rate message within RCDT (see table 85 in 6.7.4.2).

The SP transmitter sends the following messages to the SP state machine:

a) COMINIT Transmitted;
b) COMSAS Transmitted;
c) COMWAKE Transmitted;
d) SATA Port Selection Signal Transmitted;
e) TRAIN_DONE Pattern Transmitted; and
f) Phy Capabilities Bits Transmitted.

The SP receiver receives OOB signals and dwords from the physical link and sends messages to the SP state machine indicating what it has received.

The SP receiver receives the following messages from the SP state machine:

a) Set Rate with a Physical Link Rate argument, and an SSC On argument or an SSC Off argument;
b) Start Training; and
c) Abort Training.

The SP receiver sends the following messages to the SP state machine:

a) COMINIT Detected;
b) COMSAS Detected;
c) COMWAKE Detected;
d) COMSAS Completed;
e) COMWAKE Completed;
f) ALIGN Received with an argument indicating the specific type (e.g., ALIGN Received (0));
g) Phy Capabilities Bits Received with arguments indicating the supported settings bits received;
h) Training Completed;
i) TRAIN_DONE Received; and
j) Dword Received.

The ALIGN Received, Dword Received, and TRAIN_DONE Received messages are only sent when the SP_DWS state machine has achieved dword synchronization.

For SATA speed negotiation, the ALIGN Received (0) message includes an argument containing the physical link rate at which the ALIGN (0) primitives were detected. For SAS speed negotiation, only ALIGNs at the physical link rate specified by the last Set Rate message received by the SP transmitter cause ALIGN Received messages.

The SP transmitter relationship to other transmitters is defined in 4.3.2. The SP receiver relationship to other receivers is defined in 4.3.3.

6.9.2 SP_DWS receiver

The SP_DWS receiver receives the following messages from the SP_DWS state machine:

a) Find Dword.

The SP_DWS receiver sends the following messages to the SP_DWS state machine:

a) Dword Received (Primitive);
b) Dword Received (Data Dword);
c) Dword Received (Invalid); and
d) Incorrect Mux Received.

The SP_DWS receiver also sends Dword Received confirmations to the link layer state machine receivers (e.g., SL_IR, SL, SSP, SMP, and XL). If multiplexing is enabled (see table 90 in 6.7.4.2.3.3), the SP_DWS receiver shall use the first incoming MUX primitive to determine the logical phy to which it sends each Dword Received confirmation and shall not send an Dword Received confirmations until it receives the first incoming MUX primitive.

Upon receiving a Find Dword message, the SP_DWS receiver shall monitor the input data stream and force each K28.5 character detected into the first character position of a possible dword. If the next three characters are data characters with correct disparity, it shall send the dword as a Dword Received (Primitive) message to the SP_DWS state machine. Until it receives another Find Dword message, for every four characters it receives it shall:

a) send a Dword Received (Invalid) message to the SP_DWS state machine if the dword is an invalid dword (see 3.1.112);
b) send a Dword Received (Primitive) message to the SP_DWS state machine if the dword is a primitive (i.e., the dword contains a K28.5 character in the first character position followed by three data characters); or
c) send a Dword Received (Data Dword) message to the SP_DWS state machine if the dword is a data dword (i.e., it is not an invalid dword or a primitive).

The SP_DWS receiver relationship to other receivers is defined in 4.3.3.

7.8.2 IDENTIFY address frame

Table 116 defines the IDENTIFY address frame format used for the identification sequence. The IDENTIFY address frame is sent by each logical phy after the phy reset sequence completes if the physical link is a SAS physical link. The IDENTIFY address frames sent by each logical phy in a physical phy shall be identical.

...

7.9 Identification and hard reset sequence

7.9.1 Identification and hard reset sequence overview

After the phy reset sequence (see 6.7) has been completed indicating the physical link is using SAS rather than SATA, each logical phy either:

a) performs and identification sequence by transmitting ~~an~~one or three IDENTIFY address frames (see 7.8.2); or
b) performs a hard reset sequence by transmitting a HARD_RESET primitive sequence (see 7.2.5.9).

If multiplexing (see 6.10) is enabled (see table 86 in 6.7.4.2.3.3), the identification sequence and hard reset sequence begin with a multiplexing sequence (see 6.7.4.3).

Figure 38 shows two phys with multiplexing disabled performing the identification sequence. Only one IDENTIFY address frame is shown in this example.

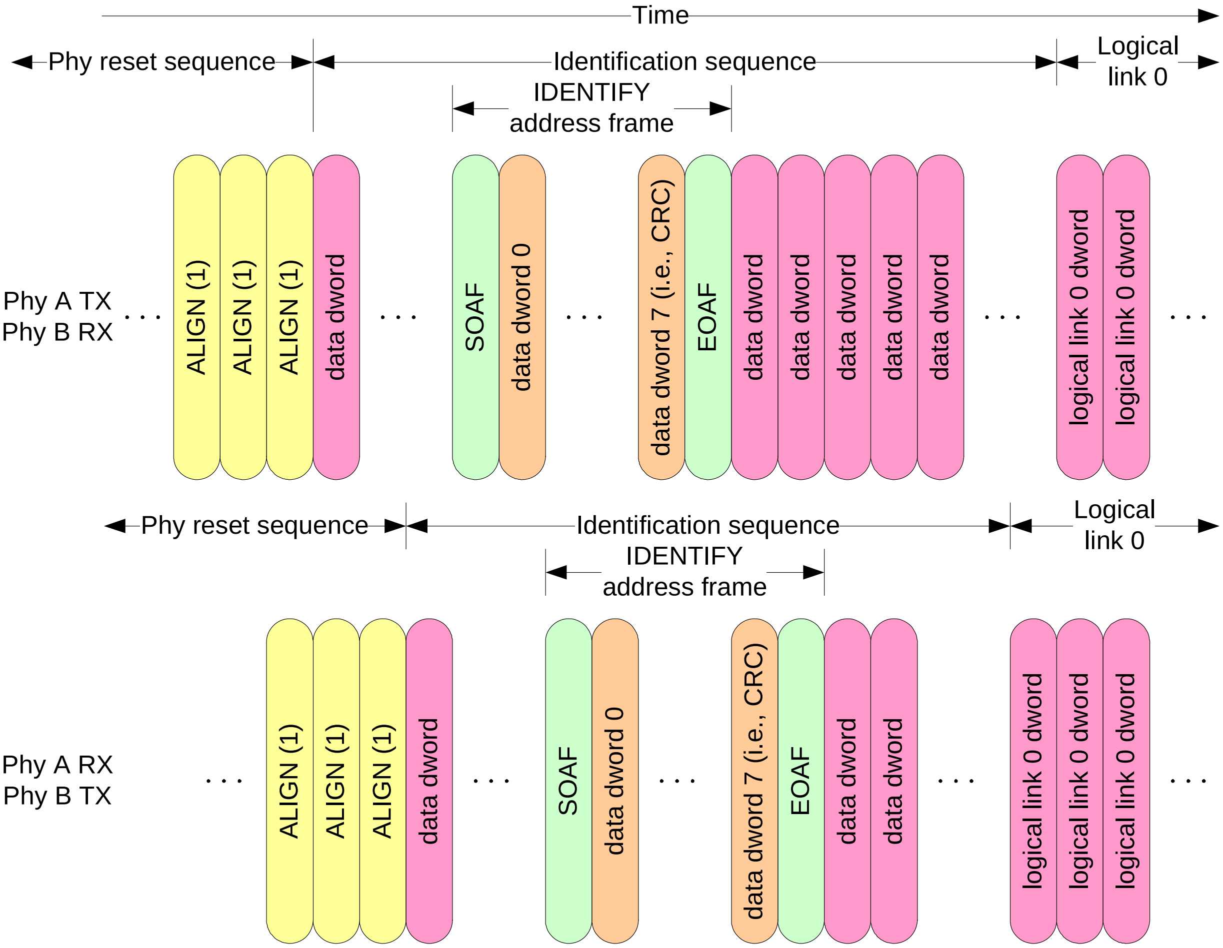

NOTE: Phys transmit deletable primitives for clock skew management after the phy reset sequence.

Figure 38 — Identification sequence

Figure 39 shows phy A performing the identification sequence and phy B performing the hard reset sequence. Multiplexing is disabled and only one IDENTIFY address frame is shown in this example.

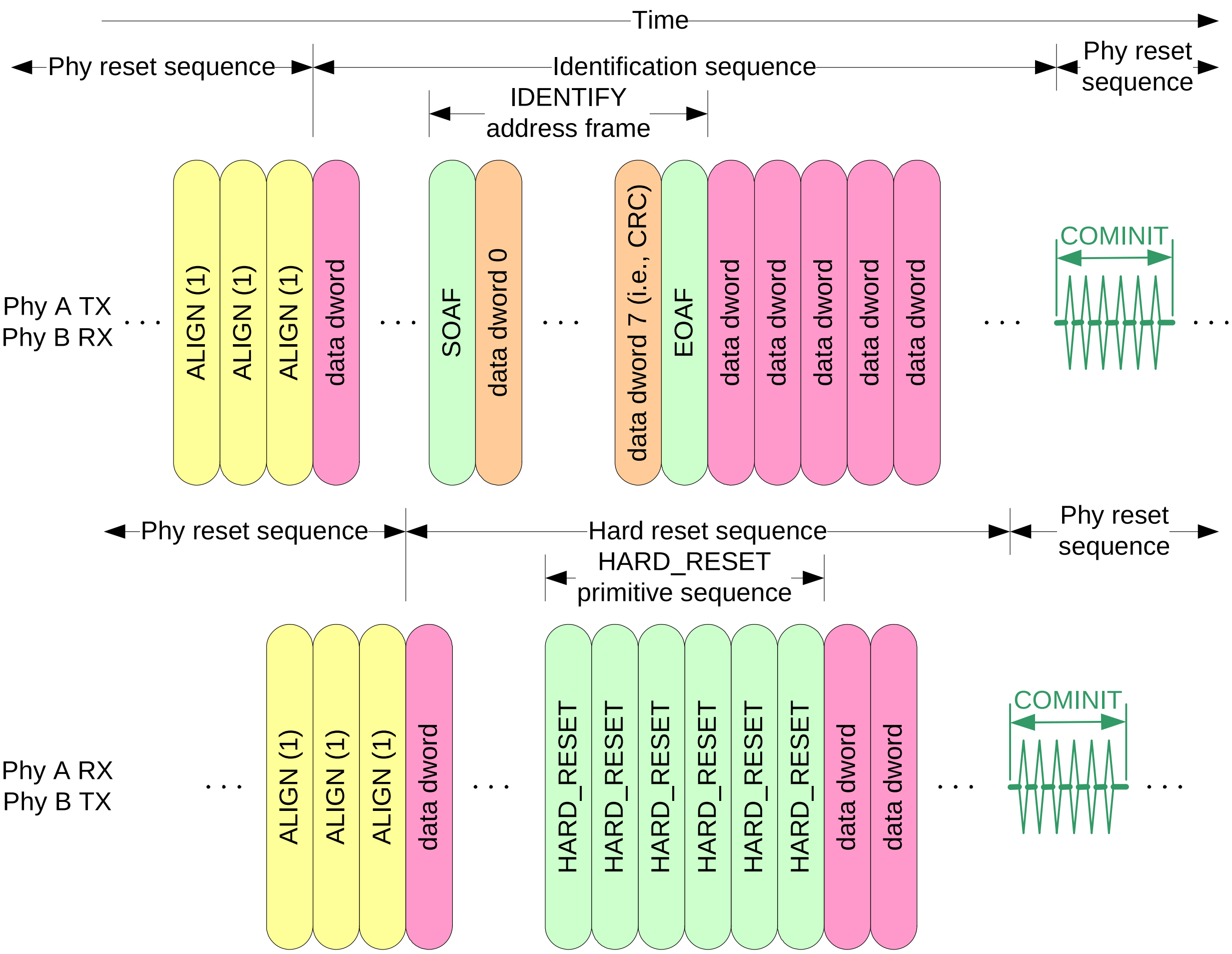

NOTE: Phys transmit deletable primitives for clock skew management after the phy reset sequence.

Figure 39 — Hard reset sequence

Each logical phy receives an IDENTIFY address frame or a HARD_RESET primitive sequence from the logical phy to which it is attached. The combination of a phy reset sequence, an optional hard reset sequence followed by another phy reset sequence, and an identification sequence is called a link reset sequence (see 4.4.1).

If a phy receives a valid IDENTIFY address frame within 1 ms of phy reset sequence completion, the SAS address in the outgoing IDENTIFY address frame(s) and the SAS address in the incoming IDENTIFY address frame determine the port to which a phy belongs (see 4.1.4). The phy ignores subsequent IDENTIFY address frames and HARD_RESET primitives until another phy reset sequence occurs.

If a phy receives a HARD_RESET primitive sequence within 1 ms of phy reset sequence completion, it shall be considered a reset event and cause a hard reset (see 4.4.2) of the port containing that phy.

If a phy does not receive a HARD_RESET primitive sequence or a valid IDENTIFY address frame within 1 ms of phy reset sequence completion, it shall restart the phy reset sequence.

7.9.2 SAS initiator device rules

After a link reset sequence, or after receiving a Broadcast (Change), a management application client behind an SMP initiator port should perform a discover process (see 4.7).

When a discover process is performed after a link reset sequence, the management application client discovers all the devices in the SAS domain. When a discover process is performed after a Broadcast (Change), the management application client determines which devices have been added to or removed from the SAS domain.

The discover information may be used to select connection rates for connection requests (see 7.8.3).

After receiving a Broadcast (Expander), a management application client behind an SMP initiator port should issue a REPORT GENERAL function (see 10.4.3.3) to all expander devices to determine:

a) the expander devices, if any, that are reducing their functionality (i.e., the REDUCED FUNCTIONALITY bit is set to one in the REPORT GENERAL response)(see 4.6.8); and
b) the amount of time remaining until the reduced functionality occurs (i.e., the contents of the TIME TO REDUCED FUNCTIONALITY field in the REPORT GENERAL response).

7.9.3 Expander device rules

After completing the link reset sequence on a phy and completing internal initialization, the ECM within an expander device shall be capable of routing connection requests through that phy. The expander device may return OPEN_REJECT (NO DESTINATION) until it is ready to process connection requests.

After a link reset sequence, or after receiving a Broadcast (Change), the management application client behind an SMP initiator port in a self-configuring expander device shall follow the SAS initiator device rules (see 7.9.2) to perform a discover process (see 4.7).

The ECM of an externally configurable expander device is dependent on the completion of the discover process (see 4.7) for routing connection requests using the table routing method.

7.9.4 SL_IR (link layer identification and hard reset) state machines

7.9.4.1 SL_IR state machines overview

The SL_IR (link layer identification and hard reset) state machines control the flow of dwords on the physical link that are associated with the identification and hard reset sequences. The state machines are as follows:

a) SL_IR_TIR (transmit IDENTIFY or HARD_RESET) state machine (see 7.9.4.3);
b) SL_IR_RIF (receive IDENTIFY address frame) state machine (see 7.9.4.4); and
c) SL_IR_IRC (identification and hard reset control) state machine (see 7.9.4.5).

The SL_IR state machines send the following messages to the SL state machines (see 7.14) in SAS devices or the XL (see 7.15) state machine in expander devices:

a) Enable Disable SAS Link (Enable); and
b) Enable Disable SAS Link (Disable).

The SL_IR_IRC state machine shall maintain the timers listed in table 1.

Table 1 — SL_IR_IRC timers

Timer	Initial value
Receive Identify Timeout timer	1 ms

Figure 40 shows the SL_IR state machines.

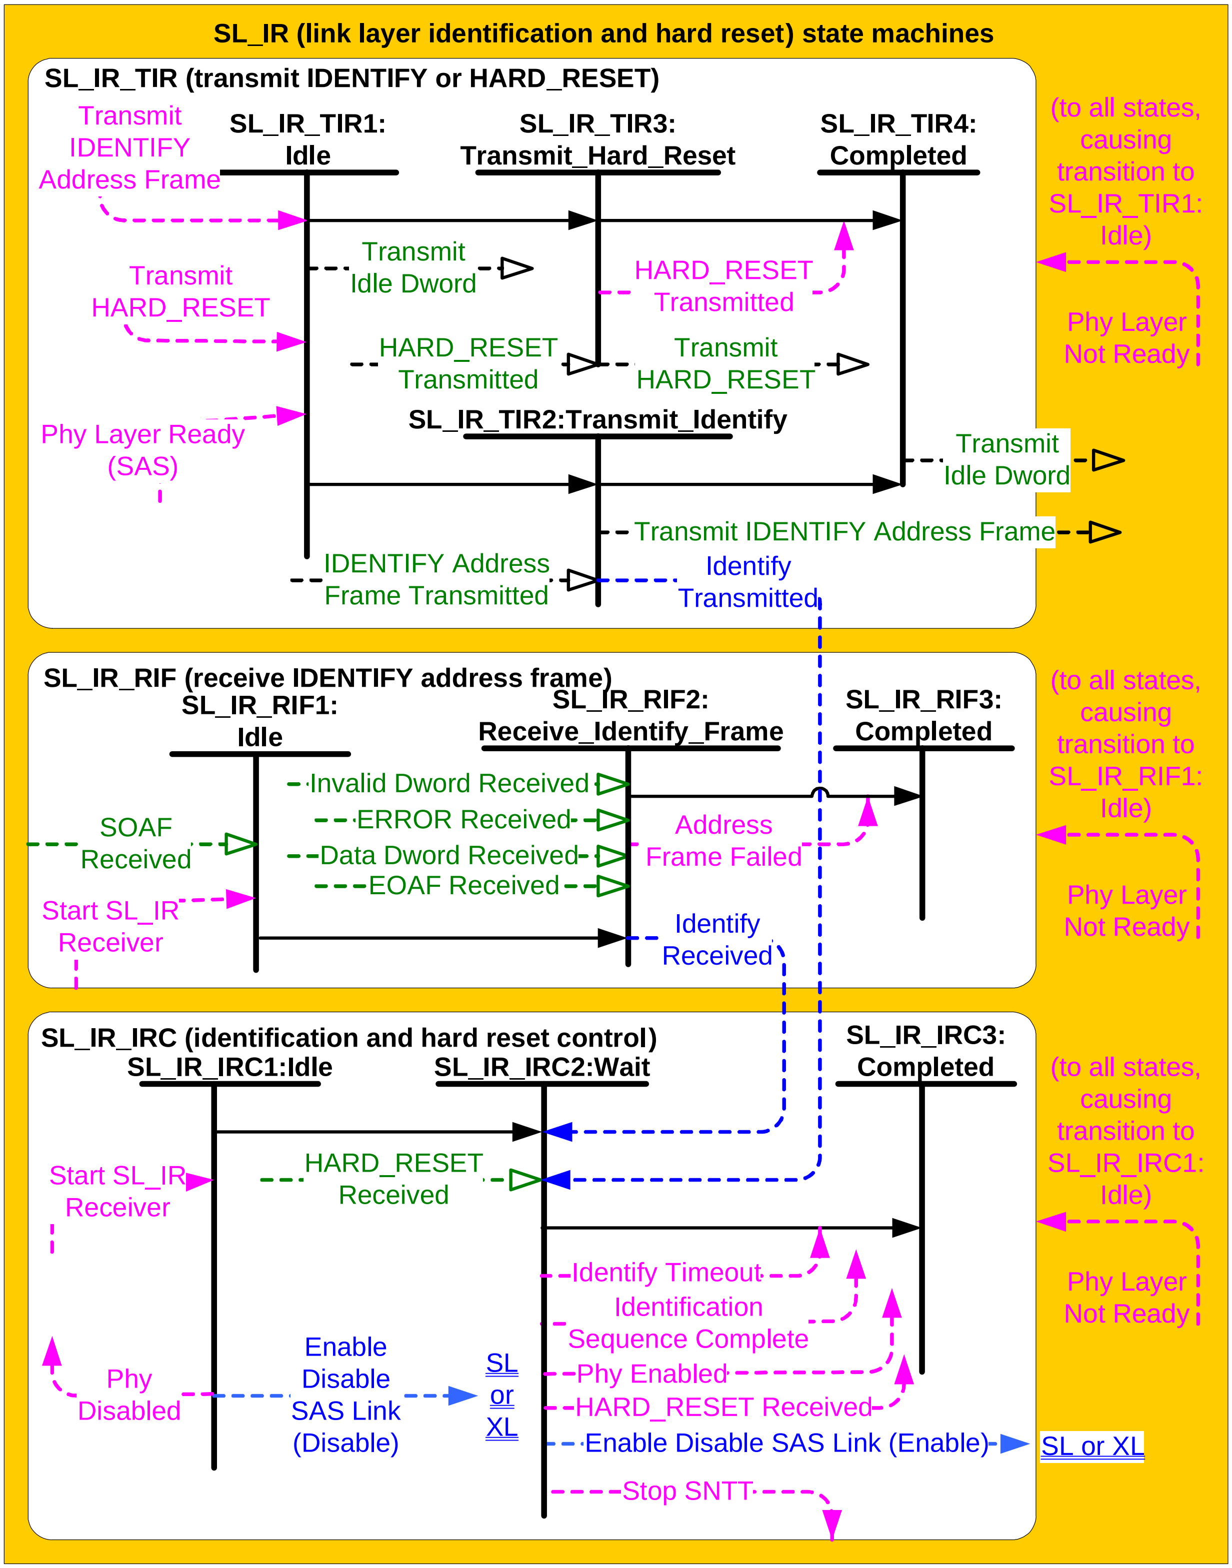

Figure 40 — SL_IR (link layer identification and hard reset) state machines

7.9.4.2 SL_IR transmitter and receiver

The SL_IR transmitter receives the following messages from the SL_IR state machines indicating primitive sequences, frames, and dwords to transmit:

a) Transmit IDENTIFY Address Frame;
b) Transmit HARD_RESET; and
c) Transmit Idle Dword.

The SL_IR transmitter sends the following messages to the SL_IR state machines:

a) HARD_RESET Transmitted; and
b) IDENTIFY Address Frame Transmitted.

The SL_IR receiver sends the following messages to the SL_IR state machines indicating primitive sequences and dwords received from the SP_DWS receiver (see 6.9.2):

a) SOAF Received;
b) Data Dword Received;
c) EOAF Received;
d) ERROR Received;
e) Invalid Dword Received; and
f) HARD_RESET Received.

The SL_IR receiver shall ignore all other dwords.

The SL_IR transmitter relationship to other transmitters is defined in 4.3.2. The SL_IR receiver relationship to other receivers is defined in 4.3.3.

7.9.4.3 SL_IR_TIR (transmit IDENTIFY or HARD_RESET) state machine

7.9.4.3.1 SL_IR_TIR state machine overview

The SL_IR_TIR state machine's function is to transmit ~~a single IDENTIFY address frame~~one or three IDENTIFY address frames or a HARD_RESET primitive after the phy layer enables the link layer. This state machine consists of the following states:

a) SL_IR_TIR1:Idle (see 7.9.4.3.2)(initial state);
b) SL_IR_TIR2:Transmit_Identify (see 7.9.4.3.3);
c) SL_IR_TIR3:Transmit_Hard_Reset (see 7.9.4.3.4); and
d) SL_IR_TIR4:Completed (see 7.9.4.3.5).

This state machine shall start in the SL_IR_TIR1:Idle state. This state machine shall transition to the SL_IR_TIR1:Idle state from any other state after receiving a Phy Layer Not Ready confirmation.

7.9.4.3.2 SL_IR_TIR1:Idle state

7.9.4.3.2.1 State description

This state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SL_IR transmitter.

7.9.4.3.2.2 Transition SL_IR_TIR1:Idle to SL_IR_TIR2:Transmit_Identify

This transition shall occur after both:

a) a Phy Layer Ready (SAS) confirmation is received; and
b) a Transmit IDENTIFY Address Frame request is received.

7.9.4.3.2.3 Transition SL_IR_TIR1:Idle to SL_IR_TIR3:Transmit_Hard_Reset

This transition shall occur after both:

a) a Phy Layer Ready (SAS) confirmation is received; and
b) a Transmit HARD_RESET request is received.

7.9.4.3.3 SL_IR_TIR2:Transmit_Identify state

7.9.4.3.3.1 State description

Upon entry into this state, this state shall send ~~a~~either one or three Transmit IDENTIFY Address Frame messages to the SL_IR transmitter.

NOTE 1 - Phys compliant with previous versions of this standard only transmitted one Transmit IDENTIFY Address Frame message.

After this state receives an IDENTIFY Address Frame Transmitted message in response to its first Transmit IDENTIFY Address Frame message, this state shall send an Identify Transmitted message to the SL_IR_IRC state machine.

7.9.4.3.3.2 Transition SL_IR_TIR2:Transmit_Identify to SL_IR_TIR4:Completed

If this state sends one Transmit IDENTIFY Address Frame message, this ~~This~~ transition shall occur after sending an Identify Transmitted message to the SL_IR_IRC state machine.

If this state sends three Transmit IDENTIFY Address Frame messages, this transition shall occur after receiving three Identify Transmitted messsages.

7.9.4.3.4 SL_IR_TIR3:Transmit_Hard_Reset state

7.9.4.3.4.1 State description

Upon entry into this state, this state shall send a Transmit HARD_RESET message to the SL_IR transmitter.

After this state receives a HARD_RESET Transmitted message, this state shall send a HARD_RESET Transmitted confirmation to the management application layer.

7.9.4.3.4.2 Transition SL_IR_TIR3:Transmit_Hard_Reset to SL_IR_TIR4:Completed

This transition shall occur after sending a HARD_RESET Transmitted confirmation to the management application layer.

7.9.4.3.5 SL_IR_TIR4:Completed state

This state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SL_IR transmitter.

7.9.4.4 SL_IR_RIF (receive IDENTIFY address frame) state machine

7.9.4.4.1 SL_IR_RIF state machine overview

The SL_IR_RIF state machine receives an IDENTIFY address frame and checks the IDENTIFY address frame to determine if the frame should be accepted or discarded by the link layer.

This state machine consists of the following states:

a) SL_IR_RIF1:Idle (see 7.9.4.4.2)(initial state);
b) SL_IR_RIF2:Receive_Identify_Frame (see 7.9.4.4.3); and
c) SL_IR_RIF3:Completed (see 7.9.4.4.4).

This state machine shall start in the SL_IR_RIF1:Idle state. This state machine shall transition to the SL_IR_RIF1:Idle state from any other state after receiving a Phy Layer Not Ready confirmation.

7.9.4.4.2 SL_IR_RIF1:Idle state

7.9.4.4.2.1 State description

This state waits for an SOAF to be received from the physical link, indicating an address frame is arriving.

7.9.4.4.2.2 Transition SL_IR_RIF1:Idle to SL_IR_RIF2:Receive_Identify_Frame

This transition shall occur after both:

a) a Start SL_IR Receiver confirmation is received; and
b) an SOAF Received message is received.

7.9.4.4.3 SL_IR_RIF2:Receive_Identify_Frame state

7.9.4.4.3.1 State description

This state receives the dwords of an address frame and the EOAF.

If this state receives an SOAF Received message, then this state shall ~~discard the address frame (i.e., the subsequent Data Dword Received and EOAF Received messages) and~~ discard any previously received dwords for the address frame, send an Address Frame Failed confirmation to the management application layer to indicate that an invalid IDENTIFY address frame was received, and start receiving the new address frame.

If this state receives more than eight Data Dword Received messages after an SOAF Received message and before an EOAF Received message, then this state shall discard the address frame and send an Address Frame Failed confirmation to the management application layer to indicate that an invalid IDENTIFY address frame was received.

If this state receives an Invalid Dword Received message or an ERROR Received message after an SOAF Received message and before an EOAF Received message, then this state shall:

a) ignore the invalid dword or ERROR; or
b) discard the address frame and send an Address Frame Failed confirmation to the management application layer to indicate that an invalid IDENTIFY address frame was received.

After receiving an EOAF Received message, this state shall check if it the received frame is a valid IDENTIFY address frame.

This state shall accept an IDENTIFY address frame and send an Identify Received message to the SL_IR_IRC state machine if:

a) the ADDRESS FRAME TYPE field is set to Identify;
b) the number of bytes between the SOAF and EOAF is 32; and
c) the CRC field contains a good CRC.

Otherwise, this state shall discard the ~~IDENTIFY~~ address frame and send an Address Frame Failed confirmation to the management application layer to indicate that an invalid ~~IDENTIFY~~ address frame was received.

7.9.4.4.3.2 Transition SL_IR_RIF2:Receive_Identify_Frame to SL_IR_RIF3:Completed

This transition shall occur after sending an Identify Received message ~~or Address Frame Failed~~ confirmation.

7.9.4.4.4 SL_IR_RIF3:Completed state

This state waits for a Phy Layer Not Ready confirmation.

7.9.4.5 SL_IR_IRC (identification and hard reset control) state machine

7.9.4.5.1 SL_IR_IRC state machine overview

The SL_IR_IRC state machine ensures that IDENTIFY address frames have been both received and transmitted before enabling the rest of the link layer, and notifies the link layer if a HARD_RESET primitive sequence is received before an IDENTIFY address frame has been received.

This state machine consists of the following states:

a) SL_IR_IRC1:Idle (see 7.9.4.5.2)(initial state);
b) SL_IR_IRC2:Wait (see 7.9.4.5.3); and

c) SL_IR_IRC3:Completed (see 7.9.4.5.4).

This state machine shall start in the SL_IR_IRC1:Idle state. This state machine shall transition to the SL_IR_IRC1:Idle state from any other state after receiving a Phy Layer Not Ready confirmation.

7.9.4.5.2 SL_IR_IRC1:Idle state

7.9.4.5.2.1 State description

This state waits for the link layer to be enabled. Upon entry into this state, this state shall:

a) send an Enable Disable SAS Link (Disable) message to SL state machines (see 7.14) or XL state machine (see 7.15) halting any link layer activity; and
b) send a Phy Disabled confirmation to the port layer and the management application layer indicating that the phy is not ready for use.

7.9.4.5.2.2 Transition SL_IR_IRC1:Idle to SL_IR_IRC2:Wait

This transition shall occur after a Start SL_IR Receiver confirmation is received.

7.9.4.5.3 SL_IR_IRC2:Wait state

7.9.4.5.3.1 State description

This state ensures that an IDENTIFY address frame has been received by the SL_IR_RIF state machine and that an IDENTIFY address frame has been transmitted by the SL_IR_TIR state machine before enabling the rest of the link layer. The IDENTIFY address frames may be transmitted and received on the physical link in any order.

After this state receives an Identify Received message, it shall send a Stop SNTT request to the phy layer.

After this state receives an Identify Transmitted message, it shall initialize and start the Receive Identify Timeout timer. If an Identify Received message is received before the Receive Identify Timeout timer expires, this state shall:

a) send an Identification Sequence Complete confirmation to the management application layer, with arguments carrying the contents of the incoming IDENTIFY address frame;
b) send an Enable Disable SAS Link (Enable) message to the SL state machines (see 7.14) in a SAS logical phy or the XL state machine (see 7.15) in an expander logical phy indicating that the rest of the link layer may start operation; and
c) send a Phy Enabled confirmation to the port layer and the management application layer indicating that the phy is ready for use.

If the Receive Identify Timeout timer expires before an Identify Received message is received, this state shall send an Identify Timeout confirmation to the management application layer to indicate that an identify timeout occurred.

If this state receives a HARD_RESET Received message before an Identify Received message is received, this state shall send a HARD_RESET Received confirmation to the port layer and the management application layer and a Stop SNTT request to the phy layer.

If this state receives a HARD_RESET Received message after an Identify Received message is received, the HARD_RESET Received message shall be ignored.

7.9.4.5.3.2 Transition SL_IR_IRC2:Wait to SL_IR_IRC3:Completed

This transition shall occur after sending a HARD_RESET Received confirmation, Identify Timeout confirmation, or an Identification Sequence Complete and an Phy Enabled confirmation.

7.9.4.5.4 SL_IR_IRC3:Completed state

This state waits for a Phy Layer Not Ready confirmation.

7.14.2 SL transmitter and receiver

The SL transmitter receives the following messages from the SL state machines specifying primitive sequences, frames, and dwords to transmit:

a) Transmit Idle Dword;
b) Transmit SOAF/Data Dwords/EOAF;
c) Transmit OPEN_ACCEPT;
d) Transmit OPEN_REJECT with an argument indicating the specific type (e.g., Transmit OPEN_REJECT (Retry));
e) Transmit BREAK;
f) Transmit BREAK_REPLY;
g) Transmit BROADCAST; and
h) Transmit CLOSE with an argument indicating the specific type (e.g., Transmit CLOSE (Normal)).

When the SL transmitter is requested to transmit a dword from any state within any of the SL state machines, it shall transmit that dword. If there are multiple requests to transmit, the following priority should be followed when selecting the dword to transmit:

1) BREAK_REPLY;
2) BREAK;
3) CLOSE;
4) OPEN_ACCEPT or OPEN_REJECT;
5) SOAF or data dword or EOAF; then
6) idle dword.

When there is no outstanding message specifying a dword to transmit, the SL transmitter shall transmit idle dwords.

The SL transmitter sends the following messages to the SL state machines based on dwords that have been transmitted:

a) SOAF/Data Dwords/EOAF Transmitted.

The SL receiver sends the following messages to the SL state machines indicating primitive sequences and dwords received from the SP_DWS receiver (see 6.9.2):

a) SOAF Received;
b) Data Dword Received;
c) EOAF Received;
d) BROADCAST Received with an argument indicating the specific type (e.g., BROADCAST Received (Change));
e) BREAK Received;
f) BREAK_REPLY Received;
g) OPEN_ACCEPT Received;
h) OPEN_REJECT Received with an argument indicating the specific type (e.g., OPEN_REJECT Received (No Destination));
i) AIP Received;
j) CLOSE Received with an argument indicating the specific type (e.g., CLOSE Received (Normal));
k) NOTIFY Received (Power Loss Expected);
l) ERROR Received; and
m) Invalid Dword Received.

The SL receiver shall ignore all other dwords.

The SL transmitter relationship to other transmitters is defined in 4.3.2. The SL receiver relationship to other receivers is defined in 4.3.3.

7.14.3 SL_RA (receive OPEN address frame) state machine

The SL_RA state machine's function is to receive address frames and determine if the received address frame is an OPEN address frame and whether or not it was received successfully. This state machine consists of one state.

This state machine receives SOAFs, dwords of an OPEN address frames, and EOAFs.

This state machine shall ignore all messages except SOAF Received, Data Dword Received, and EOAF Received.

If this state machine receives a subsequent SOAF Received message after receiving an SOAF Received message but before receiving an EOAF Received message, then this state machine shall discard the Data Dword Received messages received before the subsequent SOAF Received message.

If this state machine receives more than eight Data Dword Received messages after an SOAF Received message and before an EOAF Received message, then this state machine shall discard the address frame.

If this state machine receives an Invalid Dword Received message or an ERROR Received message after an SOAF Received message and before an EOAF Received message, then this state machine shall:

a) ignore the invalid dword or ERROR; or
b) discard the address frame.

After receiving an EOAF Received message, this state machine shall check if the address frame is a valid OPEN address frame.

This state machine shall accept an address frame if:

a) the ADDRESS FRAME TYPE field is set to Open;
b) the number of data dwords between the SOAF and EOAF is 8; and
c) the CRC field contains a good CRC.

Otherwise, this state machine shall discard the address frame. If the frame is not discarded then this state machine shall send a OPEN Address Frame Received message to the SL_CC0:Idle state and the SL_CC1:ArbSel state with an argument that contains all the data dwords received in the OPEN address frame.

7.15.2 XL transmitter and receiver

The XL transmitter receives the following messages from the XL state machine specifying primitive sequences, frames, and dwords to transmit:

a) Transmit Idle Dword;
b) Transmit AIP with an argument indicating the specific type (e.g., Transmit AIP (Normal));
c) Transmit BREAK;
d) Transmit BREAK_REPLY;
e) Transmit BROADCAST with an argument indicating the specific type (e.g., Transmit BROADCAST (Change));
f) Transmit CLOSE with an argument indicating the specific type (e.g., Transmit CLOSE (Normal));
g) Transmit OPEN_ACCEPT;
h) Transmit OPEN_REJECT, with an argument indicating the specific type (e.g., Transmit OPEN_REJECT (No Destination));
i) Transmit OPEN Address Frame; and
j) Transmit Dword.

The XL transmitter sends the following messages to the XL state machine based on dwords that have been transmitted:

a) OPEN Address Frame Transmitted.

The XL transmitter shall ensure physical link rate tolerance management requirements are met (see 7.3) while originating dwords.

The XL transmitter shall ensure physical link rate tolerance management requirements are met while forwarding dwords (i.e., during a connection) by inserting or deleting as many deletable primitives as required to match the transmit and receive connection rates (see 7.3.2).

The XL transmitter shall ensure physical link rate tolerance management requirements are met (see 7.3) during and after switching from forwarding dwords to originating dwords, including, for example:

a) when transmitting BREAK;
b) when transmitting BREAK_REPLY;
c) when transmitting CLOSE;
d) when transmitting an idle dword after closing a connection (i.e., after receiving BREAK, BREAK_REPLY, or CLOSE);
e) while transmitting a SATA frame to a SAS logical link during an STP connection, when transmitting the first SATA_HOLDA in response to detection of SATA_HOLD; and
f) while receiving dwords of a SATA frame from a SAS logical link during an STP connection, when transmitting SATA_HOLD.

NOTE 2 - The XL transmitter may always insert a deletable primitive before transmitting a BREAK, BREAK_REPLY, CLOSE, or SATA_HOLDA to meet physical link rate tolerance management requirements.

The XL transmitter shall insert a deletable primitive before switching from originating dwords to forwarding dwords, including, for example:

a) when transmitting OPEN_ACCEPT;
b) when transmitting the last idle dword before a connection is established (i.e., after receiving OPEN_ACCEPT);
c) while transmitting a SATA frame to a SAS logical link during an STP connection, when transmitting the last dword from the STP flow control buffer in response to release of SATA_HOLD;
d) while transmitting a SATA frame to a SAS logical link during an STP connection, when transmitting the last SATA_HOLDA in response to release of SATA_HOLD (e.g., if the STP flow control buffer is empty); and
e) while receiving dwords of a SATA frame from a SAS logical link during an STP connection, when transmitting the last SATA_HOLD.

NOTE 3 - This ensures that physical link rate tolerance management requirements are met, even if the forwarded dword stream does not include a deletable primitive until the last possible dword.

The XL transmitter shall ensure rate matching requirements are met during a connection (see 7.13).

When there is no outstanding message specifying a dword to transmit, the XL transmitter shall transmit idle dwords.

The XL receiver sends the following messages to the XL state machine indicating primitive sequences, frames, and dwords received from the SP_DWS receiver (see 6.9.2):

a) AIP Received with an argument indicating the specific type (e.g., AIP Received (Normal));
b) BREAK Received;
c) BREAK_REPLY Received;
d) BROADCAST Received;
e) CLOSE Received;
f) OPEN_ACCEPT Received;
g) OPEN_REJECT Received;
h) OPEN Address Frame Received;
i) Dword Received with an argument indicating the data dword or primitive received; and
j) Invalid Dword Received.

The XL receiver shall ignore all other dwords.

While receiving an address frame, if the XL receiver receives an invalid dword or ERROR, then the XL receiver shall:

a) ignore the invalid dword or ERROR; or
b) discard the address frame.

The XL transmitter relationship to other transmitters is defined in 4.3.2. The XL receiver relationship to other receivers is defined in 4.3.3.

7.16.8.2 SSP transmitter and receiver

The SSP transmitter receives the following messages from the SSP state machines specifying primitive sequences and frames to transmit:

a) Transmit RRDY with an argument indicating the specific type (e.g., Transmit RRDY (Normal));
b) Transmit CREDIT_BLOCKED;
c) Transmit ACK;
d) Transmit NAK with an argument indicating the specific type (e.g., Transmit NAK (CRC Error));
e) Transmit Frame with an argument containing the frame contents; and
f) Transmit DONE with an argument indicating the specific type (e.g., Transmit DONE (Normal)).

In response to the Transmit Frame message, the SSP transmitter transmits:

1) SOF;
2) the frame contents;
3) CRC; and
4) EOF.

The SSP transmitter sends the following messages to the SSP state machines based on dwords that have been transmitted:

a) DONE Transmitted;
b) RRDY Transmitted;
c) CREDIT_BLOCKED Transmitted;
d) ACK Transmitted;
e) NAK Transmitted; and
f) Frame Transmitted.

When there is no outstanding message specifying a dword to transmit, the SSP transmitter shall transmit idle dwords.

The SSP receiver sends the following messages to the SSP state machines indicating primitive sequences and dwords received from the SP_DWS receiver (see 6.9.2):

a) ACK Received;
b) NAK Received;
c) RRDY Received;
d) CREDIT_BLOCKED Received;
e) DONE Received with an argument indicating the specific type (e.g., DONE Received (Normal));
f) SOF Received;
g) Data Dword Received;
h) EOF Received;
i) NOTIFY Received (Power Loss Expected);
j) ERROR Received; and
k) Invalid Dword Received.

The SSP receiver shall ignore all other dwords.

The SSP transmitter relationship to other transmitters is defined in 4.3.2. The SSP receiver relationship to other receivers is defined in 4.3.3.

7.17.8 STP (link layer for STP phys) state machines

The STP link layer uses the SATA link layer state machines (see SATA-2), modified to:

a) communicate with the port layer rather than directly with the transport layer;
b) interface with the SL state machines for connection management (e.g., to select when to open and close STP connections, and to tolerate idle dwords between an OPEN address frame and the first SATA primitive);
c) communicate with an STP transmitter and receiver; and
d) support an affiliation policy (see 7.17.4).

These modifications are not described in this standard.

The STP transmitter relationship to other transmitters is defined in 4.3.2. The STP receiver relationship to other receivers is defined in 4.3.3.

Editor's Note 9: Reason: Including a few mentions of STP transmitter and receiver so cross references to 4.3 are bidirectional.

7.18.5.2 SMP transmitter and receiver

The SMP transmitter receives the following messages from the SMP state machines specifying dwords and frames to transmit:

a) Transmit Idle Dword; and
b) Transmit Frame with an argument containing the frame contents.

In response to the Transmit Frame message, the SMP transmitter transmits:

1) SOF;
2) the frame contents;
3) CRC; and
4) EOF.

The SMP transmitter sends the following messages to the SMP state machines based on dwords that have been transmitted:

a) Frame Transmitted.

When there is no outstanding message specifying a dword to transmit, the SMP transmitter shall transmit idle dwords.

The SMP receiver sends the following messages to the SMP state machines indicating primitive sequences and dwords received from the SP_DWS receiver (see 6.9.2):

a) SOF Received;
b) Data Dword Received;
c) EOF Received;
d) ERROR Received; and
e) Invalid Dword Received.

The SMP receiver shall ignore all other dwords.

The SMP transmitter relationship to other transmitters is defined in 4.3.2. The SMP receiver relationship to other receivers is defined in 4.3.3.