

20 September 2006

06-324r3 SAS-2 Modifications to the SAS Speed Negotiation

To: T10 Technical Committee
From: Amr Wassal, PMC-Sierra (amr_wassal@pmc-sierra.com)
Stephen Finch, STMicroelectronics (steve.finch@st.com)
Date: 20 September 2006
Subject: 06-324r3 SAS-2 Modifications to the SAS Speed Negotiation

Revision history

Revision 0 (7 July 2006) First revision

Revision 1 (24 August 2006)

Revision 2 (7 September 2006) Review updates including: training pattern definition, changes to several diagrams where training pattern is depicted and state machine.

Revision 3 (20 September 2006) Updates based on 12 September 2006 Phy Working Group meeting. A significant re-organization, plus many changes to make consistent use of new terms like Final-SNW and Train-SNW. Updates to diagrams.

Related documents

sas2r04a - Serial Attached SCSI - 2 (SAS-2) revision 4a

05-397 SAS-2 Start-up training sequence (Harvey Newman, Infineon)

06-263 SAS-2 Spread-spectrum clocking (Rob Elliott, HP)

06-295 SAS-2 Speed negotiation (Amr Wassal and Robert Watson, PMC-Sierra)

06-354 SAS-2 Startup training sequence proposal (Steve Finch, ST Microelectronics)

06-355 SAS-2 SNW-3 definition (Amr Wassal and Robert Watson, PMC-Sierra)

06-363 SAS-2 SNW-3 bit definitions (Rob Elliot, HP)

Editor's Note:

Some of the text from proposal 06-363 is included in this document for reference to allow the reader to understand the relationship of this document to 06-363. The content of this material should be reviewed during discussion of 06-363 and not during discussions of this document. Editor's notes in blue text indicate the start and end of each such inclusion.

Overview

The SAS speed negotiation defined in SAS-1 did not allow for training of the phy nor allow for exchange of parameters that are proving useful at G3 and higher rates. New proposals for SAS-2 proposed a training pattern in 05-397 and speed negotiation modifications in 06-295 that has provision for training and for exchanging parameters during SNW-3 as defined in 06-363.

Suggested changes to chapter 6 (Phy Layer)

6.7.4.2 SAS speed negotiation sequence

6.7.4.2.1 SAS speed negotiation sequence overview

The speed negotiation sequence establishes communications between the two phys of a physical link at the highest possible transmission rate.

The SAS speed negotiation sequence is a peer-to-peer negotiation technique that does not assume initiator and target (i.e., host and device) roles. The rules for speed negotiation are the same for both phys of a physical connection.

The speed negotiation sequence consists of a set of speed negotiation windows. Each speed negotiation window is identified by a name (e.g., Speed Negotiation Window 1 or SNW-1) that is used to identify the format, the timing of any transmissions within the window, and the transmission rate.

Speed negotiation windows conform to one of three defined formats:

- a) Speed negotiation without training. SNW-1, SNW-2 and Final-SNW are of this type.
- b) Speed negotiation with training. Training-SNW is of this type.
- c) Link Capabilities Exchange. SNW-3 is of this type.

The speed negotiation window formats are defined in 6.7.4.2.3.

Many of the timing parameters used for defining the speed negotiation windows are common to multiple window types. All of the timing specifications for all speed negotiation window types are defined in section 6.7.4.2.2.

A SAS speed negotiation sequence may or may not include all three types of speed negotiation windows. Phys may implement a subset of speed negotiation windows provided that the subset implements a valid speed negotiation sequence. Speed negotiation sequences are defined in section [6.7.4.2.46-7.4.2.3](#).

6.7.4.2.2 Speed Negotiation Timing Specifications

Table 74 defines the timing specifications for the SAS speed negotiation windows. The definitions of the speed negotiation window types specify the timing parameters that are applicable for that type.

Table 74 — SAS speed negotiation sequence timing specifications

Acronym	Name	Time ^a	Comments
RCDT	Rate change delay time	750 000 OOB ^b	The time the transmitter device shall transmit D.C. idle between rates during speed negotiation.
SNTT	Speed negotiation transmit time	163 840 OOB ^c	The time during which ALIGN (0) or ALIGN (1) is transmitted at each physical link rate during the speed negotiation sequence.
SNLT	Speed negotiation lock time	153 600 OOB ^d	The maximum time during SNW-1, SNW-2 or Final-SNW for a transmitter device to reply with ALIGN (1).
ALT	Actual lock time		The time at which actual lock to the received ALIGN (0) or ALIGN (1) occurs and the transmitter transitions from sending ALIGN (0) to sending ALIGN (1). This must be within SNLT.
MTT	Maximum training time	29 998 080 OOB ^e	The maximum time for training to complete during a Training-SNW.
SNWT	Speed negotiation window time	913 840 OOB ^f	The duration of SNW-1, SNW-2, SNW-3, or a Final-SNW.
MTWT	Maximum Train-SNW window time	30 748 080 OOB ^g	The maximum duration of the Train-SNW window. The duration of the Train-SNW may be less than RCDT + MTT.

^a OOB is defined in table 67.
^b 750 000 OOB is nominally 500 μ s. Derived from: OOB x 18 750 x 40.
^c 163 840 OOB is nominally 109,226 μ s. Derived from: OOB x 4 096 x 40.
^d 153 600 OOB is nominally 102,4 μ s. Derived from: OOB x 3 840 x 40.
^e 29 998 080 OOB is nominally 19,99872 ms (~ 20 ms, an integer number of training patterns)
^f 913 840 OOB is nominally 609,226 μ s. Derived from: RCDT + SNTT.
^g 30 748 080 OOB is nominally 20,5 ms. Derived from: RCDT + MTT.

Speed negotiation window

During each speed negotiation window, a phy will either:

- a) Transmit and receive as defined for the current speed negotiation window, or
- b) Transmit D.C. Idle and not attempt to receive.

If both phys transmit and successfully receive the expected transmission, then that speed negotiation window is said to have been valid.

If either phy fails to receive the expected transmission from the far end phy, then that speed negotiation window is said to have been invalid.

Note x: If a phy transmits D.C. Idle during a speed negotiation window, then the far end phy will not receive the expected transmission and that speed negotiation window is invalid.

6.7.4.2.3 Speed Negotiation Window Definitions

6.7.4.2.3.1 SNW-1, SNW-2 and Final-SNW Definition

Figure 122 defines the SAS speed negotiation window format for SNW-1, SNW-2 and Final-SNW, including:

- a) speed negotiation window time;
- b) speed negotiation window rate;
- c) rate change delay time (RCDT);
- d) speed negotiation transmit time (SNTT);
- e) speed negotiation lock time (SNLT); and
- f) actual lock time (ALT)




Figure 122 — SAS SNW-1, SNW-2 and Final-SNW

The speed negotiation window shall consist of the following transmission sequence:

- 1) transmission of D.C. idle for an RCDT; and
- 2) if the phy supports the physical link rate or if this is the Final-SNW, transmission of ALIGNs at the physical link rate specified below, for the remainder of the entire speed negotiation window time. If the phy does not support the physical link rate, transmission of D.C. idle for the remainder of the entire speed negotiation window time.

The physical link rate for SNW-1 shall be 1,5Gbps.

The physical link rate for SNW-2 shall be 3,0 Gbps.

The physical link rate for Final-SNW shall be:

- a) 1,5 Gbps if SNW-1 was valid and SNW-2 was invalid, or
- b) 3,0 Gbps if SNW-2 was valid and SNW-3 was invalid

If the phy supports the speed negotiation window rate, it shall initially transmit ALIGN (0) and it shall attempt to synchronize on an incoming series of dwords at that rate for the SNLT. The received dwords may be ALIGN (0) or ALIGN (1) primitives. If the phy achieves dword synchronization within the SNLT, it shall change from transmitting ALIGN (0) primitives to transmitting ALIGN (1) primitives for the remainder of the SNTT (i.e., the remainder of the speed negotiation window time). The point at which the phy achieves dword synchronization is the Actual lock time (ALT). The change from transmitting ALIGN (0) primitives to transmitting ALIGN (1) primitives shall occur after ALT. If the phy does not achieve dword synchronization within the SNLT, it shall continue transmitting ALIGN (0) primitives for the remainder of the SNTT (i.e., the remainder of the speed negotiation window time).

At the end of the SNTT, if a phy is both transmitting and receiving ALIGN (1) primitives, it shall consider that physical link rate valid.

SNW-1, SNW-2 and Final-SNW are always performed with spread spectrum clocking disabled.

6.7.4.2.3.2 SNW-3 Definition

SNW-3 provides a means for the phys to exchange information about their capabilities (e.g., physical link rates supported information, SSC support information, and multiplexing support).

During SNW-3 each phy transmits 32 bits of information that describes its capabilities and receives 32 bits of information that describes the capabilities of the far end phy.

Each of the 32 bits of information can be a "1" or a "0". The transmitter transmits 32 bit cells at the beginning of the SNTT portion and then transmits D.C. Idle for the remainder of SNTT.

Each bit cell consists of the transmission of a COMWAKE (see 6.6.2) or the transmission of D.C. Idle for a period of time equivalent to the length of a COMWAKE. The transmission of a COMWAKE indicates a value of "1". The transmission of D.C Idle for a period equivalent in length to a COMWAKE indicates a value of "0".

COMWAKE shall be transmitted at 1,5 Gbps.

Table 75 defines the information bits.

Figure 123 defines the SAS SNW-3, including:

- a) speed negotiation window time;
- b) speed negotiation window rate of 1,5 Gbps;
- c) rate change delay time (RCDT); and
- d) speed negotiation transmit time (SNTT)..




Figure 123 — SAS SNW-3

<<Start of text from 06-363>>

Table 75 defines the SNW-3 information.

Table 75 — SNW-3 settings bits

Bit(s)	Description
0 (first bit)	START bit
1	TX SSC TYPE bit
2 to 7	Reserved
Supported Settings	
8	G1 WITH SSC SUPPORTED bit
9	G1 WITHOUT SSC SUPPORTED bit
10	G2 WITH SSC SUPPORTED bit
11	G2 WITHOUT SSC SUPPORTED bit
12	G3 WITH SSC SUPPORTED bit
13	G3 WITHOUT SSC SUPPORTED bit
14 to 30	Reserved
Trailer	
31 (last bit)	PARITY field

The START bit shall be set to one. The phy's receiver shall use this bit to establish the timing for the subsequent bits.

A TX SSC TYPE bit set to one indicates that the phy's transmitter uses center-spreading SSC when SSC is enabled. A TX SSC TYPE bit set to zero indicates that the phy's transmitter uses down-spreading SSC when SSC is enabled, or that the phy does not support SSC.

NOTE 1 - The phy receiver may use the TX SSC TYPE bit to optimize its CDR circuitry. This bit indicates the type of SSC used when attached to a SAS phy or an expander phy; if a phy supports center-spreading when attached to a SAS phy or an expander phy and down-spreading when attached to a SATA phy, it sets the TX SSC TYPE bit to one.

The G<GENERATION NUMBER> WITH SSC SUPPORTED bits and G<GENERATION NUMBER> WITHOUT SSC SUPPORTED bits indicate the physical link rates and SSC options the phy is attempting to negotiate.

A G1 WITH SSC SUPPORTED bit set to one indicates that the phy supports G1 (i.e., 1.5 Gbps) with SSC. A G1 WITH SSC SUPPORTED bit set to zero indicates that the phy does not support G1 with SSC.

A G1 WITHOUT SSC SUPPORTED bit set to one indicates that the phy supports G1 without SSC. A G1 WITHOUT SSC SUPPORTED bit set to zero indicates that the phy does not support G1 without SSC.

A G2 WITH SSC SUPPORTED bit set to one indicates that the phy supports G2 (i.e., 3 Gbps) with SSC. A G2 WITH SSC SUPPORTED bit set to zero indicates that the phy does not support G2 with SSC.

A G2 WITHOUT SSC SUPPORTED bit set to one indicates that the phy supports G2 without SSC. A G2 WITHOUT SSC SUPPORTED bit set to zero indicates that the phy does not support G2 without SSC.

A G3 WITH SSC SUPPORTED bit set to one indicates that the phy supports G3 (i.e., 6 Gbps) with SSC. A G3 WITH SSC SUPPORTED bit set to zero indicates that the phy does not support G3 with SSC.

A G3 WITHOUT SSC SUPPORTED bit set to one indicates that the phy supports G3 without SSC. A G3 WITHOUT SSC SUPPORTED bit set to zero indicates that the phy does not support G3 without SSC.

The PARITY bit provides for error detection of all the SNW-3 information bits. The PARITY bit shall be set to one or zero such that the total number of SNW-3 information bits that are set to one is even, including the START bit and the PARITY bit. If the PARITY bit received is incorrect, the phy shall consider it a phy reset problem.

<<End of text from 06-363>>

For bits defined as reserved, the phy shall transmit zeros (D.C. Idle) in the outgoing SNW-3 information and shall ignore the bits in the incoming SNW-3 information.

<<Start of text from 06-363>>

Table 76 lists some example SNW-3 settings bits.

Table 76 — Example SNW-3 settings bits

Code ^a	Description
80A80000h	Down-spreading SSC G1, G2, and G3 with SSC supported
80FC0001h	Down-spreading SSC G1, G2, and G3 with and without SSC supported
C0540001h	Center-spreading SSC G1, G2, and G3 without SSC supported
C0FC0000h	Center-spreading SSC G1, G2, and G3 with and without SSC supported
^a Expressed as a 32-bit value with bit 0 (i.e., the START bit) as the MSB and bit 31 as the LSB (i.e., the PARITY bit).	

<<End of text from 06-363>>

If the START bit is incorrect, the phy shall consider SNW-3 to be invalid.

If the START bit is correct and the PARITY bit received is incorrect, the phy shall consider it a phy reset problem.

If SNW-3 exchanges no compatible supported features (e.g., one phy supports only 1,5 and 3 Gbps and the other phy supports only 6 Gbps), the phy shall consider SNW-3 to be invalid.

<<Start of text from 06-363>>

The phy shall negotiate to the highest commonly supported settings based on the outgoing and incoming SNW-3 information supported settings bits. For bits defined as reserved, the phy shall transmit zeros in the outgoing SNW-3 information and shall ignore the bits in the incoming SNW-3 information.

Table 77 defines the priority of the SNW-3 supported settings.

Table 77 — Negotiation priority of SNW-3 information supported settings

Priority	Supported Setting
Highest	G3 with SSC
...	G3 without SSC
...	G2 with SSC
...	G2 without SSC
...	G1 with SSC
Lowest	G1 without SSC

<<End of text from 06-363>>

The transmitter device shall use SAS signal output levels during the SAS speed negotiation sequence as described in 5.3.6.5.

6.7.4.2.3.3 Training-SNW Definition

Figure 124 defines the Training-SNW, including:

- a) final speed negotiation window time;
- b) rate change delay time (RCDT);
- c) actual training time (ATT); and
- d) maximum training time (MTT).

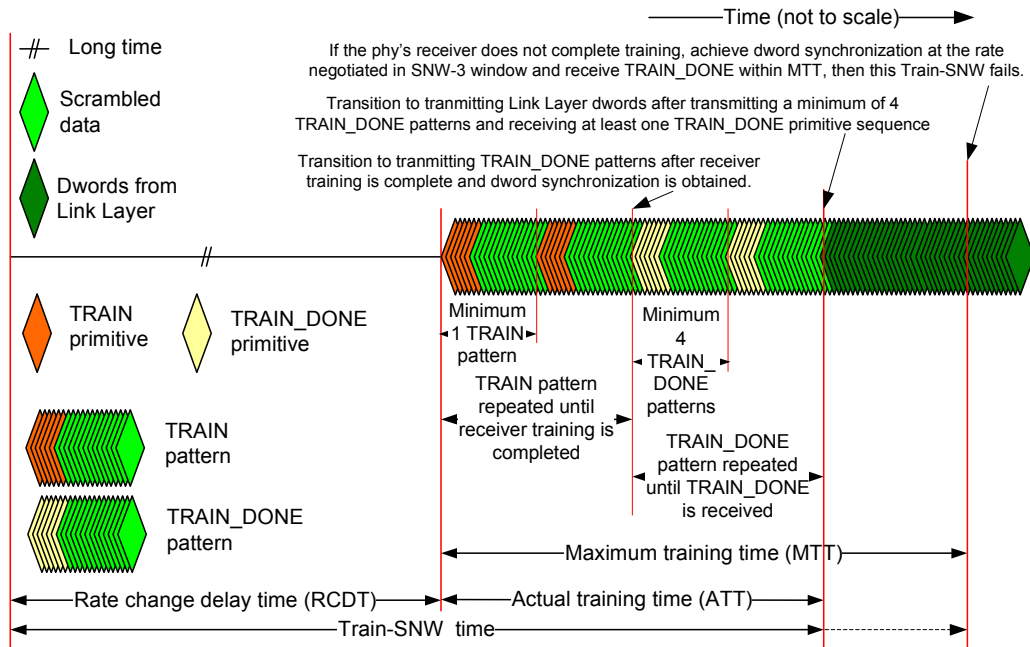


Figure 124 — SAS Train-SNW

The Train-SNW utilizes TRAIN and TRAIN_DONE primitive sequences as defined in x.x.x.x. These primitive sequences are used within training patterns.

There are two training patterns, the TRAIN pattern and the TRAIN_DONE pattern. Each training pattern consists of:

- 1) either a TRAIN primitive sequence if the pattern is a TRAIN pattern, or a TRAIN_DONE primitive sequence if the pattern is a TRAIN_DONE pattern; and
- 2) 58 dwords of value 0000_0000h that are transmitted scrambled and 8b10b encoded.

The scrambler is as defined for the link layer and shall be re-initialized at the beginning of the MTT portion of the Train-SNW window.

The transmitter shall start transmitting TRAIN patterns at the beginning of the MTT portion of the Train-SNW window. The number of TRAIN patterns transmitted is determined by the time required for the phy's receiver to complete training.

When the phy's receiver is trained and has acquired dword synchronization, the phy shall stop transmitting the TRAIN patterns and start transmitting with TRAIN_DONE patterns. A minimum of four TRAIN_DONE patterns shall be transmitted.

If the phy has transmitted four or more TRAIN_DONE patterns and has received a minimum of one TRAIN_DONE primitive sequence before the maximum training time (MTT) has been reached, the phy shall stop transmitting TRAIN_DONE patterns and shall start transmitting dwords from the link layer. In this case the Train-SNW was successful.

If either phy has not transmitted a minimum of four TRAIN_DONE patterns before the maximum training time (MTT) has been reached, the SNW-3 was unsuccessful.

6.7.4.2.4 SAS speed negotiation sequence

The SAS speed negotiation sequence consists of a set of speed negotiation windows (see) for each physical link rate in this order:

- 1) SNW-1 (used to negotiate 1,5 Gbps without training capabilities);
- 2) SNW-2 (used to negotiate 3 Gbps without training capabilities); and
- 3) One of the following sequences:
 - a) A Final-SNW negotiating 1,5 Gbps if SNW-1 was valid and SNW-2 was invalid;
 - b) An SNW-3 followed by one of the following
 - A) A Final-SNW negotiating 3,0 Gbps if SNW-2 was valid and SNW-3 was invalid;
 - B) One or more Training-SNWs if SNW-3 was valid

If SNW-1, SNW-2 and SNW-3 were invalid, a Phy Reset Problem has occurred.

If a Final-SNW window is performed and is unsuccessful, then a Phy Reset Problem has occurred.

<<Start of text from 06-363>>

~~If the phy does not successfully negotiate the highest commonly supported setting in the final SNW, it shall perform another final SNW with the next highest commonly supported setting (e.g., if the phys support both G3 without SSC and G3 with SSC but fail to negotiate the final SNW successfully using G3 with SSC, they shall perform another final SNW using G3 without SSC).~~

<<End of text from 06-363>>

If a phy participated in an SNW-3 window by transmitting its supported settings and successfully received the supported settings transmitted by the other phy, the phy shall participate in a Training-SNW utilizing the highest commonly supported settings based on the outgoing and incoming SNW-3 settings bits.

If a Training-SNW is performed and is unsuccessful and there are additional, untried commonly supported settings exchanged during SNW-3, then a new Training-SNW is performed at the next highest, untried commonly supported capability. If a Training-SNW is performed and is unsuccessful and there were no additional, untried commonly supported settings exchanged during SNW-3, then a Phy Reset Problem has occurred.

Figure 125 shows speed negotiation between a phy A that supports G1 through G3 and a phy B that only supports G2. Both phys run:

- 1) the SNW-1, supported by phy A but not by phy B;
- 2) the SNW-2, supported by both phys; and
- 3) the SNW-3, supported by phy A but not by phy B.
- 4) the Final-SNW at 3,0 Gbps

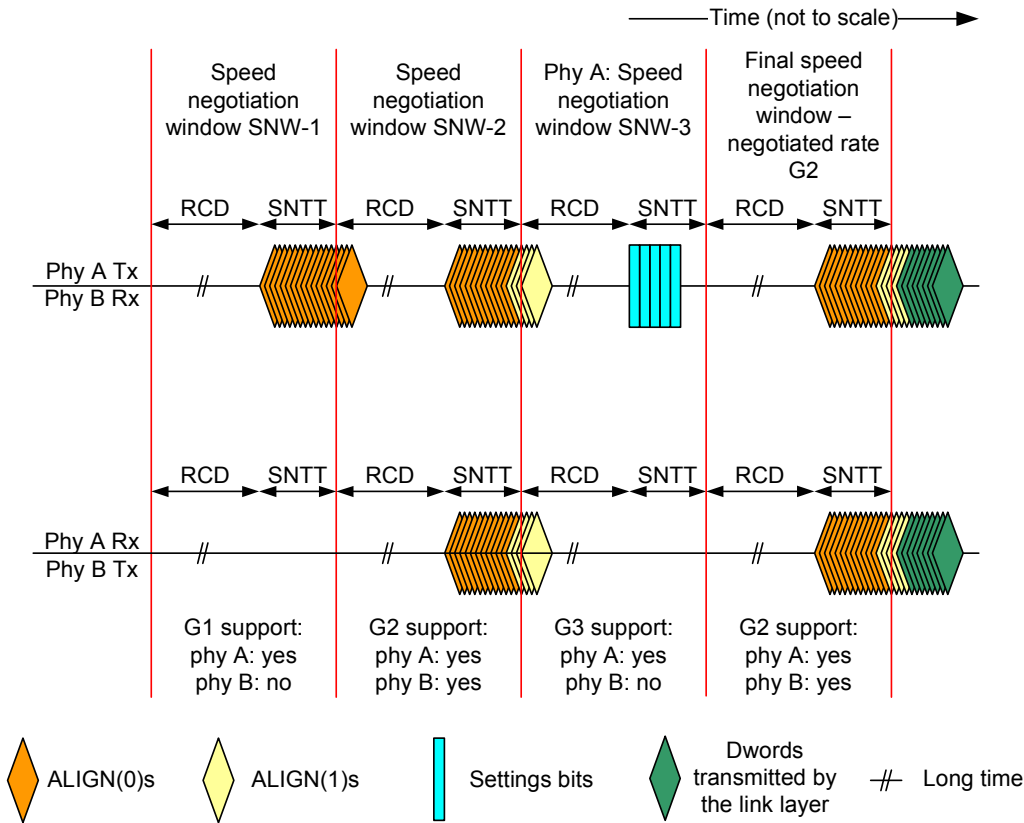


Figure 125 — SAS speed negotiation sequence (phy A: G1, G2, G3, phy B: G2 only)

Phy A determines the correct speed negotiation sequence by detecting a valid SNW-2 window followed by an invalid SNW-3 window and correctly proceeds to a Final-SNW window at 3,0 Gbps.

Phy B could determine the correct speed negotiation sequence based solely upon its capabilities, but should determine the correct speed negotiation using the same method as used by Phy A.

If the phy does not achieve dword synchronization during the Final-SNW, the SAS speed negotiation sequence fails. This is defined as a phy reset problem and may be counted and reported in the PHY RESET PROBLEM COUNT field in the SMP REPORT PHY ERROR LOG page (see 10.4.3.6) and the Protocol-Specific Port log page (see 10.2.8.1).

Figure 126 shows a speed negotiation sequence where phy B does not achieve dword synchronization during the Final-SNW. If this occurs, the handshake is not complete and the phy reset sequence is retried.

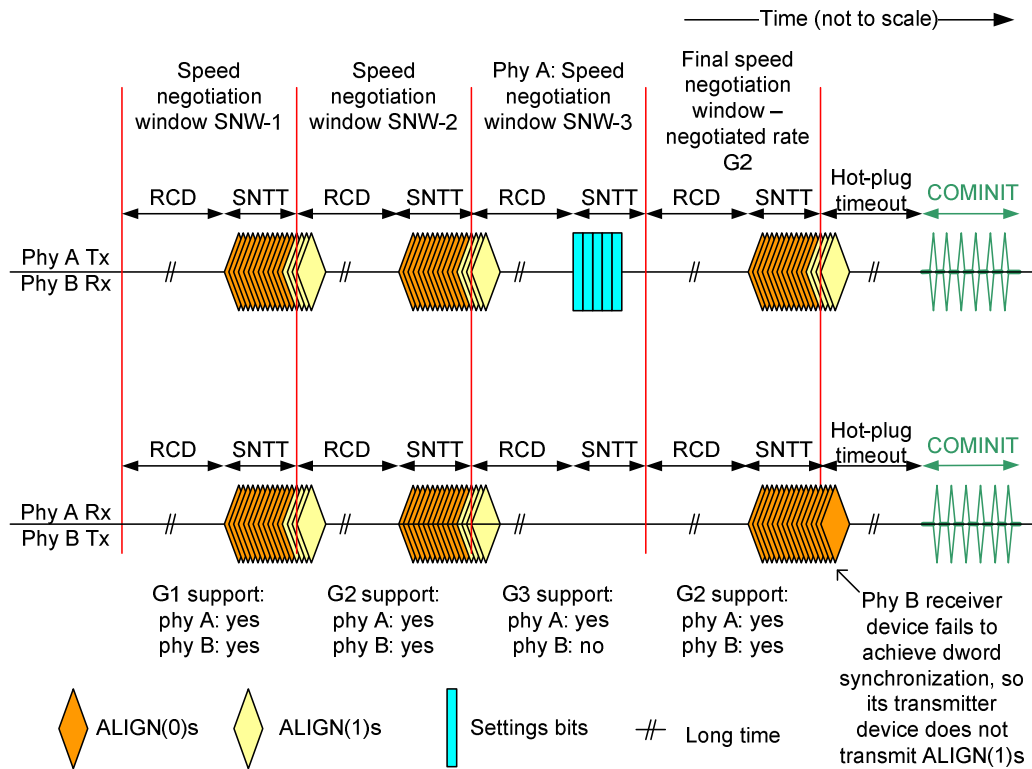


Figure 126 — SAS Final-SNW phy reset problem

Figure 127 show speed negotiation between a phy A and a phy B that both support G1 and G2 rates without requiring training and also support additional capabilities. Both phys run:

- 1) the SNW-1, supported by both phys;
- 2) the SNW-2, supported by both phys; and
- 3) the SNW-3, supported by both phys where the settings bits are exchanged and the highest common rate supported by both phys is negotiated; and
- 4) the Training-SNW, supported by both phys where training is done.

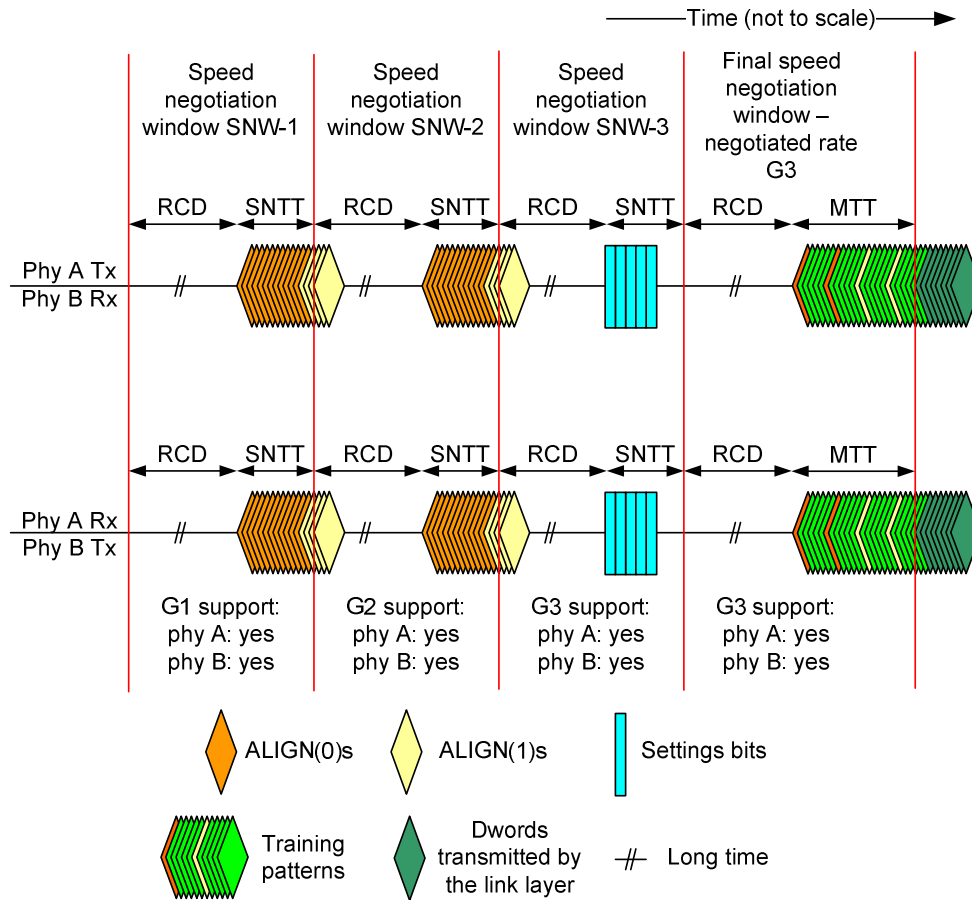


Figure 127 — SAS speed negotiation sequence (phy A: G1, G2, G3 or higher, phy B: G1, G2, G3 or higher)

A phy reset problem may be encountered if:

- 1) the phys do not exchange the settings bits properly (e.g., due to a parity error) in SNW-3 window (see Figure 128); or
- 2) either phy does not complete training within the MTT interval of several training windows exhausting all common settings exchanged in SNW-3 (see Figure 129)..




Figure 128 — SAS speed negotiation sequence - SNW-3 (maximum SNW) phy reset problem

These phy reset problems may be counted and reported in the PHY RESET PROBLEM COUNT field in the SMP REPORT PHY ERROR LOG page (see 10.4.3.6) and the Protocol-Specific Port log page (see 10.2.8.1).




Figure 129 — SAS speed negotiation sequence - training (SNW-Final) phy reset problem

For more examples of speed negotiations between phys that support various speeds, see Annex C.

6.8 SP (phy layer) state machine

6.8.1 SP state machine overview

The SP state machine controls the phy reset sequence. This state machine consists of three sets of states:

- OOB sequence (OOB) states;
- SAS and SAS-2 speed negotiation (SAS) states; and
- SATA host emulation (SATA) states.

This state machine consists of the following states:

- SP0:OOB_COMINIT (see 6.8.3.2)(initial state);
- SP1:OOB_AwaitCOMX (see 6.8.3.3);
- SP2:OOB_NoCOMSASTimeout(see 6.8.3.4);
- SP3:OOB_AwaitCOMINIT_Sent (see 6.8.3.4);
- SP4:OOB_COMSAS (see 6.8.3.4);
- SP5:OOB_AwaitCOMSAS_Sent (see 6.8.3.4);
- SP6:OOB_AwaitNoCOMSAS (see 6.8.3.4);
- SP7:OOB_AwaitCOMSAS (see 6.8.3.4);
- SP8:SAS_Start (see 6.8.3.4);
- SP9:SAS_RateNotSupported (see 6.8.3.4);
- SP10:SAS_AwaitALIGN (see 6.8.3.4);

- l) SP11:SAS_AwaitALIGN1 (see 6.8.3.4);
- m) SP12:SAS_AwaitSNW (see 6.8.3.4);
- n) SP13:SAS_Pass (see 6.8.3.4);
- o) SP14:SAS_Fail (see 6.8.3.4);
- p) SP15:SAS_PHY_Ready (see 6.8.3.4);
- q) SP16:SATA_COMWAKE (see 6.8.3.4);
- r) SP17:SATA_AwaitCOMWAKE (see 6.8.3.4);
- s) SP18:SATA_AwaitNoCOMWAKE (see 6.8.3.4);
- t) SP19:SATA_AwaitALIGN (see 6.8.3.4);
- u) SP20:SATA_AdjustSpeed (see 6.8.3.4);
- v) SP21:SATA_Transmit_ALIGN (see 6.8.3.4);
- w) SP22:SATA_PHY_Ready (see 6.8.3.4);
- x) SP23:SATA_PM_Partial (see 6.8.3.4);
- y) SP24:SATA_PM_Slumber (see 6.8.3.4);
- z) SP25:SATA_PortSel (see 6.8.3.4); and
- aa) SP26:SATA_SpinupHold (see 6.8.3.4).
- bb) SP27:SAS2_Settings (see 6.8.3.4);
- cc) SP28:SAS2_FinalSNW (see 6.8.3.4);
- dd) SP29:SAS2_Training (see 6.8.3.4); and
- ee) SP30:SAS2_TrainingDone (see 6.8.3.4).

The SP state machine shall start in the SP0:OOB_COMINIT state after:

- a) a power on;
- b) a hard reset;
- c) receiving a Management Reset request from the management layer (e.g., from the SMP PHY CONTROL function requesting a phy operation of LINK RESET or HARD RESET in an expander device); or
- d) receiving a Disable Phy request from the management layer (e.g., from the SMP PHY CONTROL function requesting a phy operation of DISABLE in an expander device).

The SP state machine shall maintain a MgmtReset state machine variable to determine whether a Management Reset request has been received. Any SP state that receives a Management Reset request shall set the MgmtReset state machine variable to one before making a transition to the SP0:OOB_COMINIT state (see 6.8.3.2). Any SP state that receives a power on, or a hard reset shall set the MgmtReset state machine variable to zero before making a transition to the SP0:OOB_COMINIT state.

If the phy supports SATA port selectors, the SP state machine shall maintain a COMWAKE_Received state machine variable to indicate whether a COMWAKE detected message was received in the SP0:OOB_COMINIT state or the SP1:OOB_AwaitCOMX state since the last time the SP0:OOB_COMINIT state was entered, and the SP state machine shall transition to the SP25:SATA_PortSel state whenever it receives a Transmit SATA Port Selection Signal request.

The SP state machine sends the following messages to the SP_DWS state machine (see 6.9):

- a) Start DWS; and
- b) Stop DWS.

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

- a) DWS Lost; and
- b) DWS Reset.

The SP state machine shall maintain the timers listed in Table 78.

Table 78 — SP state machine timers

Timer	Initial value
COMSAS Detect Timeout timer	COMSAS detect timeout (see table 66 in 6.6.1)
Await ALIGN Timeout timer	Await ALIGN timeout (see table 73 in 6.7.2.2)
Hot-Plug Timeout timer	Hot plug timeout (see table 72 in 6.7.1)
RCDT timer	RCDT (see table 77 in 6.7.4.2)
SNLT timer	SNLT (see table 77 in 6.7.4.2)
SNTT timer	SNTT (see table 77 in 6.7.4.2)
MTT timer	MTT (see table 77 in 6.7.4.2)

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 (Physical Link Rate);
- g) Transmit ALIGN with an argument indicating the specific type (e.g., Transmit ALIGN (0));
- h) Transmit Settings Bits;
- i) Transmit TRAIN sequence;
- j) Transmit TRAIN_DONE sequence; and
- k) Set SSC.

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

The SP transmitter shall complete any physical link rate change requested with the Set Rate message within RCDT (see table 77 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; and
- d) SATA Port Selection Signal 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 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) Dword Received;
- h) Settings Bits Detected;
- i) TRAIN Completed; and
- j) TRAIN_DONE Detected.

The ALIGN Received and Dword 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.

<break in paragraph sequence>

6.8.4 SAS speed negotiation states

6.8.4.1 SAS speed negotiation states overview

Figure 131 shows the SAS speed negotiation states, in which the phy has detected that it is attached to a SAS phy or expander phy rather than a SATA phy, and performs the SAS speed negotiation sequence. These states are indicated by state names with a prefix of SAS.

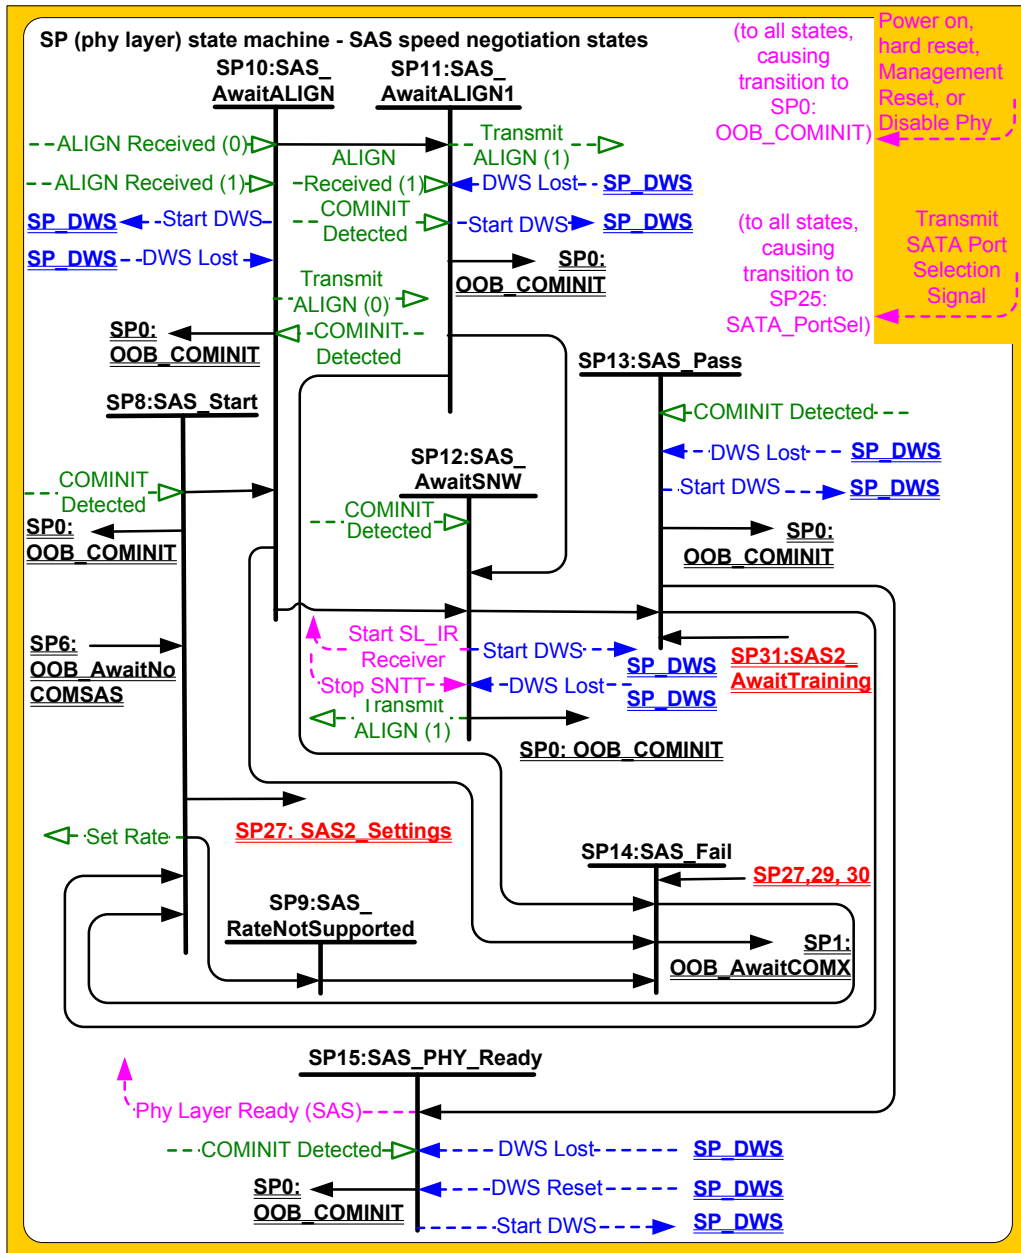


Figure 131— SP (phy layer) state machine - SAS speed negotiation states

Figure 132 shows the rest of the SAS speed negotiation states related to SNW-3 and the SNW-Final speed negotiation windows. These states are indicated by state names with a prefix of SAS2.

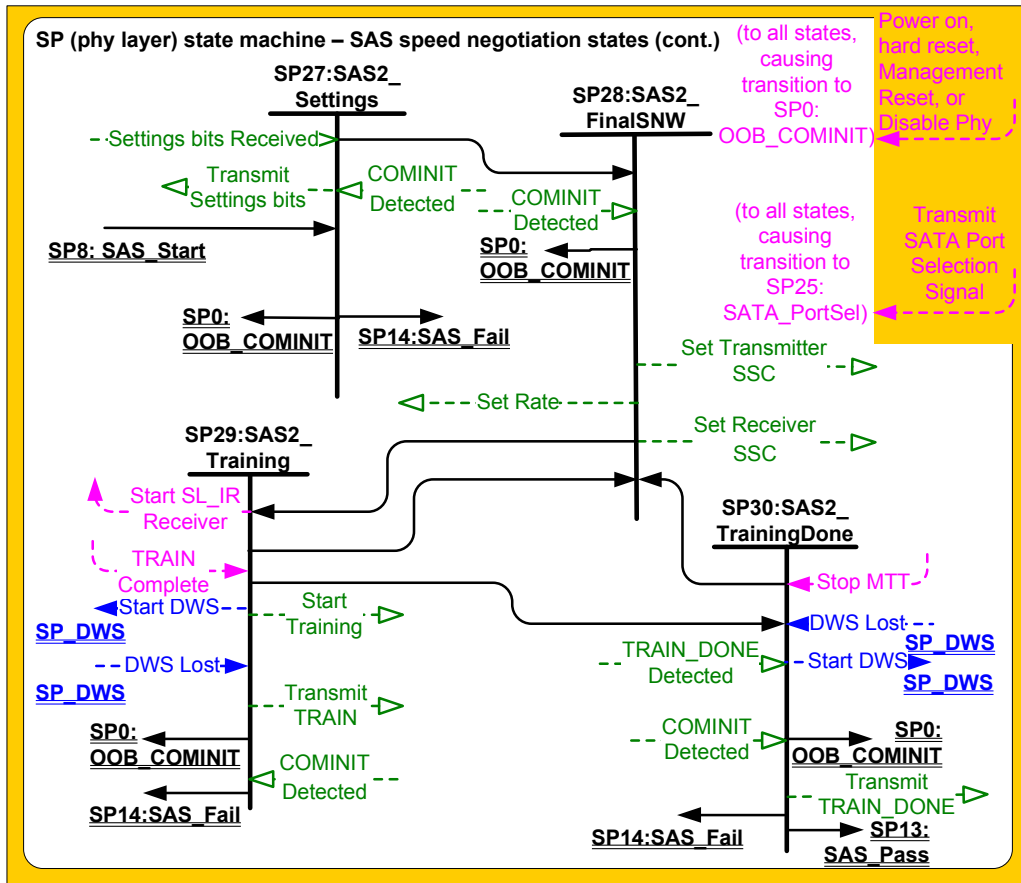


Figure 132 — SP (phy layer) state machine - SAS speed negotiation states (continued)

6.8.4.2 SP8:SAS_Start state

6.8.4.2.1 State description

This is the initial state for the SAS speed negotiation sequence.

Upon entry into this state, this state shall:

- a) initialize and start the RCDT timer;
- b) send a Set Rate message to the SP transmitter with the argument set to:
 - A) 1,5 Gbps, if the transition into this state was from the SP6:OOB_AwaitNoCOMSAS state (i.e., if this is the first speed negotiation window); or
 - B) 1,5 Gbps, if the SAS Speed Negotiaion Window is SNW-3; or
 - C) the value of the SAS Speed Negotiation Window Rate argument for G1 and G2 in SNW-1 and SNW-2.

During this state D.C. idle shall be transmitted.

6.8.4.2.2 Transition SP8:SAS_Start to SP0:OOB_COMINIT

This transition shall occur after receiving a COMINIT Detected message.

6.8.4.2.3 Transition SP8:SAS_Start to SP9:SAS_RateNotSupported

This transition shall occur after the RCDT timer expires if the current speed negotiation window rate is not supported.

6.8.4.2.4 Transition SP8:SAS_Start to SP10:SAS_AwaitALIGN

This transition shall occur after the RCDT timer expires if the current speed negotiation window is not SNW-3 and its rate is supported.

6.8.4.2.5 Transition SP8:SAS_Start to SP27:SAS2_Settings

This transition shall occur after the RCDT timer expires if the current speed negotiation window is SNW-3.

6.8.4.3 SP9:SAS_RateNotSupported state

6.8.4.3.1 State description

Upon entry into this state the SNTT timer shall be initialized and started. During this state D.C. idle shall be transmitted.

6.8.4.3.2 Transition SP9:SAS_RateNotSupported to SP14:SAS_Fail

This transition shall occur after the SNTT timer expires.

6.8.4.4 SP10:SAS_AwaitALIGN state

6.8.4.4.1 State description

Upon entry into this state, the SNTT timer and SNLT timer shall be initialized and started and this state shall repeatedly send Transmit ALIGN (0) messages to the SP transmitter.

Each time this state receives a DWS Lost message, this state may send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization without running a new link reset sequence.

Upon entry into this state, this state shall send a Start DWS message to the SP_DWS state machine.

6.8.4.4.2 Transition SP10:SAS_AwaitALIGN to SP0:OOB_COMINIT

This transition shall occur after receiving a DWS Lost message if this state does not send a Start DWS message, or after receiving a COMINIT Detected message.

6.8.4.4.3 Transition SP10:SAS_AwaitALIGN to SP11:SAS_AwaitALIGN1

This transition shall occur if this state receives an ALIGN Received (0) message before the SNLT timer expires.

6.8.4.4.4 Transition SP10:SAS_AwaitALIGN to SP12:SAS_AwaitSNW

This transition shall occur if this state receives an ALIGN Received (1) message before the SNLT timer expires.

6.8.4.4.5 Transition SP10:SAS_AwaitALIGN to SP14:SAS_Fail

This transition shall occur if the SNTT timer expires.

6.8.4.5 SP11:SAS_AwaitALIGN1 state

6.8.4.5.1 State description

This state shall repeatedly send Transmit ALIGN (1) messages to the SP transmitter.

Each time this state receives a DWS Lost message, this state may send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization without running a new link reset sequence.

6.8.4.5.2 Transition SP11:SAS_AwaitALIGN1 to SP0:OOB_COMINIT

This transition shall occur after receiving a DWS Lost message if this state does not send a Start DWS message, or after receiving a COMINIT Detected message.

6.8.4.5.3 Transition SP11:SAS_AwaitALIGN1 to SP12:SAS_AwaitSNW

This transition shall occur if this state receives an ALIGN Received (1) message before the SNLT timer expires and the current speed negotiation window is either SNW-1 or SNW-2. This indicates that the other phy has been able to achieve dword synchronization in the current speed negotiation window.

6.8.4.5.4 Transition SP11:SAS_AwaitALIGN1 to SP14:SAS_Fail

This transition shall occur if the SNTT timer expires. This indicates that the other phy has not been able to achieve dword synchronization in the current speed negotiation window.

6.8.4.6 SP12:SAS_AwaitSNW state

6.8.4.6.1 State description

This state shall repeatedly send Transmit ALIGN (1) messages to the SP transmitter. If this is the final speed negotiation window, this state shall send a Start SL_IR Receiver confirmation to the link layer.

Each time this state receives a DWS Lost message, this state may send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization without running a new link reset sequence.

This state waits for the SNTT timer to expire or for a Stop SNTT request.

6.8.4.6.2 Transition SP12:SAS_AwaitSNW to SP0:OOB_COMINIT

This transition shall occur after receiving a DWS Lost message if this state does not send a Start DWS message, or after receiving a COMINIT Detected message.

6.8.4.6.3 Transition SP12:SAS_AwaitSNW to SP13:SAS_Pass

This transition shall occur in the SNW-1 or SNW-2 windows after the SNTT timer expires or after receiving a Stop SNTT request.

6.8.4.7 SP13:SAS_Pass state

6.8.4.7.1 State description

This state determines if:

- a) another SAS speed negotiation window is required; or
- b) the SAS speed negotiation sequence is complete.

Each time this state receives a DWS Lost message, this state may send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization without running a new link reset sequence.

6.8.4.7.2 Transition SP13:SAS_Pass to SP0:OOB_COMINIT

This transition shall occur after receiving a DWS Lost message if this state does not send a Start DWS message, or after receiving a COMINIT Detected message.

6.8.4.7.3 Transition SP13:SAS_Pass to SP8:SAS_Start

This transition shall occur if this is not the final speed negotiation window.

This transition shall include a SAS Speed Negotiation Window Rate argument with the transition set to the next higher speed negotiation window rate or to 1,5 Gbps if the next speed negotiation window is SNW-3.

6.8.4.7.4 Transition SP13:SAS_Pass to SP15:SAS_PHY_Ready

This transition shall occur if this is the final speed negotiation window.

6.8.4.8 SP14:SAS_Fail state**6.8.4.8.1 State description**

This state determines if:

- a) another SAS speed negotiation window is required; or
- b) the SAS speed negotiation sequence is complete.

6.8.4.8.2 Transition SP14:SAS_Fail to SP1:OOB_AwaitCOMX

This transition shall occur if the current speed negotiation window is:

- a) the maximum SAS speed negotiation window; or
- b) the final SAS speed negotiation window.

6.8.4.8.3 Transition SP14:SAS_Fail to SP8:SAS_Start

If the previous SAS speed negotiation window (SNW-1 or SNW-2, but not SNW-3) was successful, this transition shall occur and shall include:

- a) a SAS Speed Negotiation Window Rate argument set to the previous speed negotiation window rate; and
- b) a Final SAS Speed Negotiation Window argument.

If the previous SAS speed negotiation window failed and the current speed negotiation window is not the maximum SAS speed negotiation window, this transition shall occur and shall include a SAS Speed Negotiation Window Rate argument set to the next higher speed negotiation window rate.

6.8.4.9 SP15:SAS_PHY_Ready state**6.8.4.9.1 State description**

This state waits for a COMINIT Detected message, a DWS Lost message, or a DWS Reset message.

While in this state dwords from the link layer are transmitted at the negotiated physical link rate at the rate established in the previous speed negotiation window.

Upon entry into this state, this state shall send a Phy Layer Ready (SAS) confirmation to the link layer to indicate that the physical link has been brought up successfully in SAS mode.

Each time this state receives a DWS Lost message, this state may send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization without running a new link reset sequence.

6.8.4.9.2 Transition SP15:SAS_PHY_Ready to SP0:OOB_COMINIT

This transition shall occur after:

- a) receiving a DWS Lost message, if this state does not send a Start DWS message;
- b) receiving a DWS Lost message followed by a COMINIT Detected message, if this state does not send a Start DWS message; or
- c) receiving a DWS Reset message.

This transition may but should not occur after receiving a COMINIT Detected message before receiving a DWS Lost message, or after receiving a COMINIT Detected message after sending a Start DWS message (i.e., the SP state machine should ignore COMINIT Detected messages unless the SP_DWS state machine has indicated loss of dword synchronization).

<insertion of new paragraphs – exact paragraph number to be determined at time of incorporation into the SAS-2 draft>

6.8.4.10 SP27:SAS2_Settings state

6.8.4.10.1 State description

This state shall send Transmit Settings Bits message to the SP transmitter. This state shall transmit and receive the settings bits and negotiate the common settings according to table 74. This state waits for the SNTT timer to expire or for a Stop SNTT request.

6.8.4.10.2 Transition SP27:SAS2_Settings to SP0:OOB_COMINIT

This transition shall occur after receiving a COMINIT Detected message.

6.8.4.10.3 Transition SP27:SAS2_Settings to SP28:SAS2_FinalSNW

This transition shall occur after the SNTT timer expires if a Settings Bits Detected message has been received and the settings bits are valid.

6.8.4.10.4 Transition SP27:SAS2_Settings to SP14:SAS_Fail

This transition shall occur if the SNTT timer expires without receiving valid setting bits.

6.8.4.11 SP28:SAS2_FinalSNW state

6.8.4.11.1 State description

Upon entry into this state, this state shall:

- a) initialize and start the RCDT timer;
- b) send a Set Rate message to the SP transmitter with the arguments set to the negotiated value of the physical link rate common to both the transmitted and received settings bits according to the priority given in Table 74.
- c) send a Set SSC message to the SP transmitter if needed.
- d) send a Set SSC message to the SP receiver if needed.

During this state D.C. idle shall be transmitted.

6.8.4.11.2 Transition SP28:SAS2_FinalSNW to SP0:OOB_COMINIT

This transition shall occur after receiving a COMINIT Detected message.

6.8.4.11.3 Transition SP28:SAS2_FinalSNW to SP29:SAS2_Training

This transition shall occur after the RCDT timer expires.

6.8.4.12 SP29:SAS2_Training state**6.8.4.12.1 State description**

Upon entry into this state, the MTT timer shall be initialized and started and this state shall send Transmit TRAIN sequence message to the SP transmitter and a Start Training message to the SP receiver.

Each time this state receives a DWS Lost message, this state may send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization without running a new link reset sequence.

This state shall send a Start SL_IR Receiver confirmation to the link layer when a TRAIN Completed message is received and dword synchronization is acquired.

Upon entry into this state, this state shall send a Start DWS message to the SP_DWS state machine.

6.8.4.12.2 Transition SP29:SAS2_Training to SP0:OOB_COMINIT

This transition shall occur after receiving a DWS Lost message if this state does not send a Start DWS message, or after receiving a COMINIT Detected message.

6.8.4.12.3 Transition SP29:SAS2_Training to SP30:SAS2_TrainingDone

This transition shall occur if this state receives a TRAIN Completed message before the MTT timer expires and dword synchronization is acquired.

6.8.4.12.4 Transition SP29:SAS2_Training to SP28:SAS2_FinalSNW

This transition shall occur if the MTT timer expires and the common settings exchanged in SNW-3 have not been exhausted. This indicates that this phy has not been able to complete training and achieve dword synchronization in the final speed negotiation window. This transition shall include an argument to indicate the next highest common settings of the received and transmitted settings in SNW-3 according to table 74.

6.8.4.12.5 Transition SP29:SAS2_Training to SP14:SAS_Fail

This transition shall occur if the MTT timer expires and the common settings exchanged in SNW-3 have been exhausted. This indicates that this phy has not been able to complete training and achieve dword synchronization in the final speed negotiation window.

6.8.4.13 SP30:SAS2_TrainingDone state**6.8.4.13.1 State description**

This state shall send Transmit TRAIN_DONE sequence message to the SP transmitter. A minimum of 4 training segments with a TRAIN_DONE redundant primitive sequence must be transmitted.

Each time this state receives a DWS Lost message, this state may send a Start DWS message to the SP_DWS state machine to re-acquire dword synchronization without running a new link reset sequence.

This state waits for the MTT timer to expire or for a Stop MTT request.

6.8.4.13.2 Transition SP30:SAS2_TrainingDone to SP0:OOB_COMINIT

This transition shall occur after receiving a DWS Lost message if this state does not send a Start DWS message, or after receiving a COMINIT Detected message.

6.8.4.13.3 Transition SP30:SAS2_TrainingDone to SP13:SAS_Pass

This transition shall occur if this state receives a TRAIN_DONE Detected message before the MTT timer expires, or after receiving a Stop MTT request. This indicates that the other phy has been able to complete training and achieve dword synchronization in the final speed negotiation window.

6.8.4.13.4 Transition SP30:SAS2_TrainingDone to SP28:SAS2_FinalSNW

This transition shall occur if the MTT timer expires and the common settings exchanged in SNW-3 have not been exhausted. This indicates that the peer phy has not been able to complete training and achieve dword synchronization in the final speed negotiation window. This transition shall include an argument to indicate the next highest common settings of the received and transmitted settings in SNW-3 according to table 74.

6.8.4.13.5 Transition SP30:SAS2_TrainingDone to SP14:SAS_Fail

This transition shall occur if the MTT timer expires. This indicates that the other phy has not been able to complete training and achieve dword synchronization in the final speed negotiation window.