

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 4 November 2007
 Subject: 07-119r4 SAS-2 Far-end retimed loopback phy test function

Revision history

Revision 0 (9 March 2007) First revision

Revision 1 (12 April 2007) Incorporated comments from March 2007 SAS physical WG: removed analog loopback mode. Per a request from Mark Evans (Western Digital), added an optional PHY TEST FUNCTION address frame to invoke phy test functions in a simpler manner than the SEND DIAGNOSTIC command or the SMP PHY TEST FUNCTION function; this could be used for initiators.

Revision 2 (28 April 2007) Incorporated comments from April 2007 SAS physical WG

Revision 3 (4 July 2007) Incorporated comments from May 2007 SAS physical WG

Revision 3 (4 November 2007) Incorporated comments from July 2007 SAS protocol WG. Moved TRAIN and TRAIN_DONE patterns into 07-480.

Related documents

sas2r08 - Serial Attached SCSI - 2 (SAS-2) revision 8

02-218r0 SAS Remove far-end retimed loopback (Jim Coomes, Seagate) - accepted into sas-r00b

07-091r2 or later - SAS-2 SMP function support for SNW-3 phy capabilities (Rob Elliott, HP)

07-479 SAS-2 Phy test pattern transmitter controls (Rob Elliott, HP)

07-480 SAS-2 More phy test patterns - TRAIN, TRAIN_DONE, and idle dwords (Rob Elliott, HP)

Serial ATA 2.6 (from the Serial ATA International Organization - <http://www.sata-io.org>)

PCI Express Base Specification Revision 2.0 (from the PCI SIG - <http://www.pcisig.org>)

Overview

Many phy designs implement far-end loopback (receiver -> transmitter) phy test modes, where the phy retransmits everything it receives. As with test patterns, there should be a common mechanism for entering these test modes.

Other high-speed serial interface standards define standard ways to enter loopback modes.

The *Serial ATA* specification defines two far-end loopback modes and defines enabling them via the BIST Activate FIS: far-end retimed (required) and far-end analog (optional). In far-end retimed mode:

- a) the tester must send valid 8b10b characters
- b) the tester must send two ALIGNs every 256 dwords (normal SATA rule)
- c) the phy may consume up to two ALIGNs every 256 dwords (preventing elasticity buffer overflow)
- d) the phy may add a single ALIGN (on elasticity buffer underflow) as long as it is directly preceded or followed by a tester ALIGN(s) that was not consumed
- e) the phy may add two ALIGNs back-to-back at any time
- f) the tester may send a burst of four ALIGNs to guarantee it gets at least two ALIGNs back, regardless of physical link rate tolerance differences.
- g) the phy may change disparity
- h) exit on receiving COMRESET/COMINIT

The *PCI Express* specification defines a Loopback mode "for test and fault isolation use":

- a) entry is via two consecutive TS1 Ordered Sets with the Enable Loopback bit set to one.
- b) phy's link rate is selected via TS1/TS2 bits
- c) phy's deemphasis level is selected via TS1/TS2 bits
- d) tester must go to electrical idle for 1 ms if changing speeds
- e) phy must go to electrical idle for 2 ms if changing speeds
- f) tester must send enough TS1s for the phy to lock before sending test patterns
- g) tester must send valid 8b10b symbols (characters)
- h) the phy must preserve disparity
- i) the phy may drop SKP symbols (the equivalent to SAS deletable primitives) if needed
- j) the phy may add SKP symbols adjacent to tester SKP symbols if needed (same disparity)
- k) the phy must preserve 10-bit data (even invalid 10 bit patterns)

- l) exit on receipt of four consecutive EIOS (Electrical Idle Ordered Set) patterns
- m) at 2.5 GT/sec, exit on electrical idle for 128 usec (this is all analogous to the SAS/SATA COMINIT signal).

SAS-1 originally included far-end retimed loopback, but it was removed in SAS revision 1b - it didn't seem worth the effort to resolve the details at the time. Many SAS phy designs do implement this feature, however, with vendor-specific mechanisms; a standard mechanism and standardized operational behavior would increase its usefulness.

Proposal

1. Add standardized controls using the Protocol-Specific diagnostic page and the SMP PHY TEST function to place a phy into far-end retimed loopback mode.

2. The SCSI Protocol-Specific diagnostic page and the SMP PHY TEST FUNCTION function are not well-suited for putting initiators into test modes, since initiators don't often implement SSP and/or SMP target ports. Define a new PHY TEST FUNCTION address frame for this purpose, similar to the SATA BIST Activate frame. Its contents are nearly identical to the SMP PHY TEST FUNCTION request frame, just lacking the ability to specify the phy identifier. It can be supported by any phy, not just those in initiators.

The address frame sends the attached phy it into the test mode; the diagnostic page and SMP function cannot send that phy into the test mode, they can only send other phys into the test mode. The frame is not forwarded through expanders (there's no destination SAS address like in an OPEN address frame).

In a typical design, hardware may interrupt software upon reception of a valid PHY TEST FUNCTION address frame. Software may take longer than 1 ms to parse the frame and respond, so an AF_ACK is proposed to acknowledge the PHY TEST FUNCTION address frame (within 1 ms of reception). After processing the request, software chooses whether to transmit REJECT_PHY_TEST or START_PHY_TEST.

3. Define that START_PHY_TEST is also used by the existing Transmit Pattern phy test mode.

4. Define that a phy beginning any test function transmits BREAK before START_PHY_TEST. If it was in the middle of a connection, this helps get the phy isolated without disturbing the rest of the fabric (with the address frame, the phy won't be in a connection; it could be for the SMP and mode page methods).

Editor's Note 1: Could also involve the SL and XL state machines in this process, and if the phy is already running, make it wait until between connections before directing the SP state machine to start the phy test function. For a test mode, though, BREAK seems adequate.

Suggested changes to SAS-2

4.10 Phy test functions

4.10.1 Phy test functions overview

~~The optional Protocol-Specific diagnostic page for SAS (see 10.2.9.1) provides methods for an application-client to enable and disable a phy test function (e.g., transmission of the CJTPAT) for a selected phy in a SAS target device with an SSP target port. The optional SMP PHY TEST FUNCTION function (see 10.4.3.25) provides similar methods for expander devices and SAS target devices with SMP target ports.~~

Phy test functions (e.g., transmission of test patterns) are used for phy and interconnect characterization and diagnosis. The phy may be attached to test equipment while performing a phy test function. The following optional mechanisms are defined for invoking phy test functions:

- a) the PHY TEST FUNCTION address frame (see 7.9.4) invokes a phy test function in a phy that receives the address frame. The phy test function starts some time after the phy transmits AF_ACK;
- b) the Protocol-Specific diagnostic page for SAS (see 10.2.9.1) invokes a phy test function in a selected phy other than the phy that receives the diagnostic page in a SAS target device with an SSP target port. The SEND DIAGNOSTIC command may be sent through any SSP target port to any logical unit in the SAS target device that contains the phy that is to perform the phy test function. The phy test

function starts some time after the SSP target port receives an ACK for the RESPONSE frame transmitted in response to the SEND DIAGNOSTIC command; and

- c) the SMP PHY TEST FUNCTION function (see 10.4.3.25) invokes a phy test function in a phy controlled by a management device server other than the phy that receives the function. The phy test function starts some time after the SMP target port transmits the SMP response frame.

Each phy test function is optional.

~~The application client sends a SEND DIAGNOSTIC command with the Protocol-Specific diagnostic page or an SMP PHY TEST FUNCTION function specifying the phy in the SAS target device that is to perform the phy test function and the phy test function to be performed. If the phy test function requires a specific phy test pattern and/or phy test pattern physical link rate, then it also specifies the phy test pattern and phy test pattern physical link rate are also specified.~~

~~The SEND DIAGNOSTIC command may be sent through any SSP target port to any logical unit in the SAS target device that contains the phy that is to perform the phy test function.~~

~~For the SEND DIAGNOSTIC command, the phy shall begin the specified phy test function after the SSP target port receives an ACK for the RESPONSE frame transmitted in response to the SEND DIAGNOSTIC command that requested the phy test function.~~

~~For the SMP PHY TEST FUNCTION function, the phy shall begin the specified phy test function after the SMP target port transmits the SMP response frame.~~

Upon starting the phy test function, the phy:

- 1) transmits BREAK;
- 2) transmits START_PHY_TEST; and
- 3) sets its physical link rate to the specified physical link rate.

~~Once a SAS phy has begin While a phy is performing a phy test function, it shall ignore its receiver the link layer receivers (i.e., the SL_IR receiver, SL receiver, SSP receiver, STP receiver, and SMP receiver) shall ignore all incoming dwords and the OOB signal detector shall detect COMINIT. The phy shall ignore any other OOB signals (i.e., COMSAS and COMWAKE).~~

~~To stop a SAS phy from performing a phy test function, an application client sends a SEND DIAGNOSTIC command or an SMP PHY TEST FUNCTION function to a SAS phy in the SAS target device that is not performing a phy test function requesting a phy test function of 00h (i.e., STOP). If no such phy is available, the phy test function only stops on power loss.~~

A phy stops performing a phy test function:

- a) after the SCSI device server, if any, processes a Protocol-Specific diagnostic page specifying the phy and specifying a phy test function of 00h (i.e., STOP);
- b) after the management device serve, if any, processes an SMP PHY TEST FUNCTION request specifying the phy and specifying a phy test function of 00h (i.e., STOP);
- c) after the phy receives COMINIT; or
- d) upon power off.

It is vendor-specific how long a phy takes to stop performing the phy test function. After a phy stops performing a phy test function, it performs a link reset sequence.

4.10.2 Transmit pattern phy test function

While a phy is performing the transmit pattern phy test function, the test equipment attached to that phy:

- a) shall not transmit COMSAS or COMWAKE; and
- b) shall not transmit COMINIT except to stop the phy test function.

When performing the transmit pattern phy test function, a phy:

- a) shall ignore all dwords received; and
- b) shall repeatedly transmit the specified pattern at the specified physical link rate.

4.10.3 Transmit pattern phy test function

A phy performing the transmit pattern phy test function is considered to be in the transmit pattern phy test mode.

Editor's Note 2: It might be handy to have an optional bit to have the phy transmit TRAIN for 20 ms before beginning the specified pattern. If the phy is being used to feed patterns into a receiver under test, the receiver may need that to train its DFE first.

4.10.4 Far-end retimed loopback phy test function

A phy performing the far-end retimed loopback phy test function is considered to be in the far-end retimed loopback phy test mode.

While a phy is performing the far-end retimed loopback phy test function, the test equipment attached to that phy:

- a) shall transmit D.C. idle or valid 8b10b characters at the specified physical link rate;
- b) after transmitting D.C. idle, shall transmit the following before resuming transmission of 8b10b characters:
 - 1) TRAIN patterns for MTWT (see 6.7.4.2.2); and
 - 2) ALIGN (0)s for SNTT (see 6.7.4.2.2);

NOTE 1 - If the test equipment is detached and exchanged while a phy is performing the far-end retimed loopback phy test function, the receiver in the phy may not be able to train and reattain dword synchronization. Sending the TRAIN and ALIGN (0) patterns provides it an opportunity to do so. Since the phy does not respond with TRAIN DONE patterns or ALIGN (1)s, the patterns are transmitted for their maximum times.

- c) while transmitting 8b10b characters, shall transmit deletable primitives according to 7.3;
- d) shall not transmit COMSAS or COMWAKE; and
- e) shall not transmit COMINIT except to stop the phy test function.

While performing the far-end retimed loopback phy test function, the phy:

- a) shall maintain dword synchronization and receive dwords into loopback buffer;
- b) may delete deletable primitives as necessary to avoid overflowing an elasticity buffer (see 7.3);
- c) shall either delete invalid dwords, convert them into ERROR primitives, or preserve their 10 bit patterns. The phy should not delete invalid dwords;
- d) shall retransmit, using the transmit clock, each dword that was not deleted;
- e) shall transmit a deletable primitive whenever the transmitter underflows (i.e., whenever the transmitter is ready to transmit a dword but the receiver has not provided another dword to transmit); and
- f) shall transmit dwords using valid disparity based on the transmitted dword stream. The disparity for a dword after an invalid dword is vendor-specific.

While the phy is not receiving valid 8b10b characters (e.g., while it receives D.C. idle, or after receiving D.C. idle while retraining and/or trying to attain dword synchronization), it should not try to repeat characters and may transmit anything (e.g., D.C. idle, rotating ALIGNs, or COMINIT).

Figure 1 shows the far-end retimed loopback phy test mode.

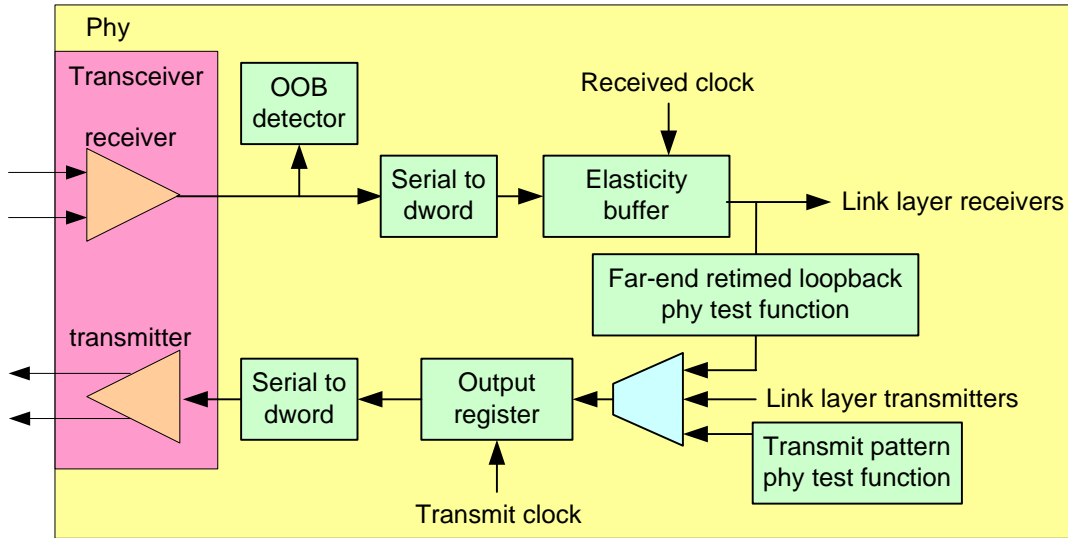


Figure 1 — Far-end retimed loopback phy test mode

Figure 1 shows usage of the far-end retimed loopback test mode.

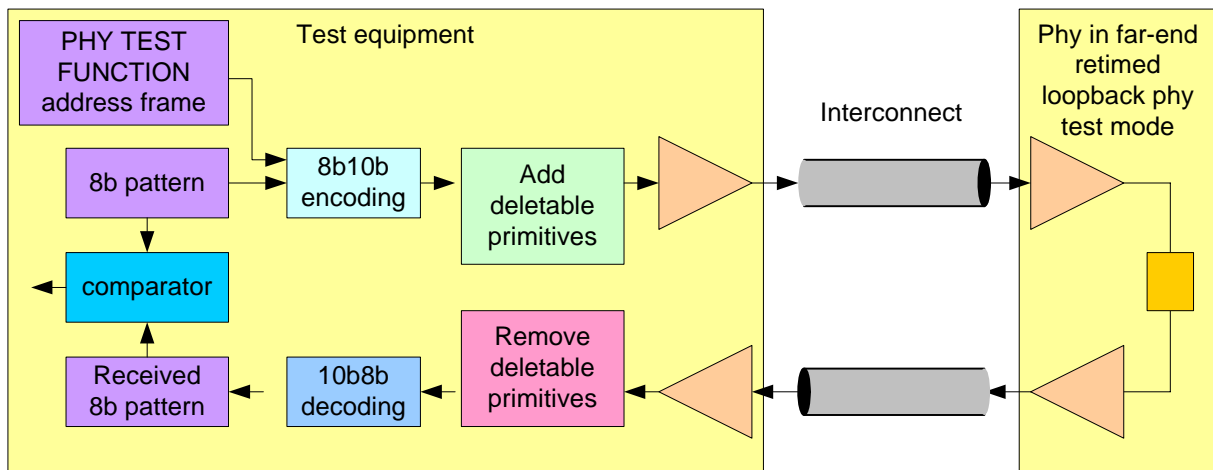


Figure 2 — Far-end retimed loopback phy test mode usage

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:

- a) OOB sequence (OOB) states;
- b) SAS speed negotiation (SAS) states; and
- c) SATA host emulation (SATA) states.

This state machine consists of the following states:

- a) SP0:OOB_COMINIT (see 6.8.3.2)(initial state);
- b) SP1:OOB_AwaitCOMX (see 6.8.3.3);
- c) SP2:OOB_NoCOMSASTimeout (see 6.8.3.4);
- d) SP3:OOB_AwaitCOMINIT_Sent (see 6.8.3.5);
- e) SP4:OOB_COMSAS (see 6.8.3.6);
- f) SP5:OOB_AwaitCOMSAS_Sent (see 6.8.3.7);

- g) SP6:OOB_AwaitNoCOMSAS (see 6.8.3.8);
- h) SP7:OOB_AwaitCOMSAS (see 6.8.3.9);
- i) SP8:SAS_Start (see 6.8.4.2);
- j) SP9:SAS_WindowNotSupported (see 6.8.4.3);
- k) SP10:SAS_AwaitALIGN (see 6.8.4.4);
- l) SP11:SAS_AwaitALIGN1 (see 6.8.4.5);
- m) SP12:SAS_AwaitSNW (see 6.8.4.6);
- n) SP13:SAS_Pass (see 6.8.4.7);
- o) SP14 SAS_Fail (see 6.8.4.8);
- p) SP15:SAS_PHY_Ready (see 6.8.4.9);
- q) SP16:SATA_COMWAKE (see 6.8.5.2);
- r) SP17:SATA_AwaitCOMWAKE (see 6.8.5.3);
- s) SP18:SATA_AwaitNoCOMWAKE (see 6.8.5.4);
- t) SP19:SATA_AwaitALIGN (see 6.8.5.5);
- u) SP20:SATA_AdjustSpeed (see 6.8.5.6);
- v) SP21:SATA_Transmit_ALIGN (see 6.8.5.7);
- w) SP22:SATA_PHY_Ready (see 6.8.5.8);
- x) SP23:SATA_PM_Partial (see 6.8.5.9);
- y) SP24:SATA_PM_Slumber (see 6.8.5.10);
- z) SP25:SATA_PortSel (see 6.8.6.2);
- aa) SP26:SATA_SpinupHold (see 6.8.7.2);
- ab) SP27:SAS_Settings (see 6.8.4.10);
- ac) SP28:SAS_TrainSetup (see 6.8.4.11);
- ad) SP29:SAS_Train (see 6.8.4.12); **and**
- ae) SP30:SAS_TrainingDone (see 6.8.4.13);
- [af\) SPxx:Transmit Pattern \(see 6.8.4.xx\); and](#)
- [ag\) SPxy:Far End Retimed Loopback \(see 6.8.4.xy\).](#)

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.~~

The SP state machine shall maintain a Current SNW state machine variable to determine the current speed negotiation window (e.g., SNW-1, SNW-2, SNW-3, Final-SNW, Train-SNW, or Unsupported Phy Attached).

If the phy status is available through any of the following:

- a) the SMP DISCOVER response (see 10.4.3.9);
- b) the SMP DISCOVER LIST response (see 10.4.3.15);
- c) the Phy Control And Discover mode page (see 10.2.7.5); or
- d) the Protocol-Specific Port log page (see 10.2.8.1),

then the SP state machine shall maintain a ResetStatus state machine variable to determine the NEGOTIATED PHYSICAL LINK RATE field and/or the NEGOTIATED LOGICAL LINK RATE field.

If the phy supports SNW-3, then the SP state machine shall maintain a Common Supported Settings state machine variable that contains supported settings common between the transmitted supported settings bits and the received supported settings bits.

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 92.

Table 1 — SP state machine timers

Timer	Initial value
COMSAS Detect Timeout timer	COMSAS detect timeout (see table 77 in 6.6.1)
Await ALIGN Timeout timer	Await ALIGN timeout (see table 84 in 6.7.2.2)
Hot-Plug Timeout timer	Hot plug timeout (see table 83 in 6.7.1)
RCDT timer	RCDT (see table 85 in 6.7.4.2)
SNLT timer	SNLT (see table 85 in 6.7.4.2)
SNTT timer	SNTT (see table 85 in 6.7.4.2)
TLT timer	TLT (see table 85 in 6.7.4.2)
MTT timer	MTT (see table 85 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 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 Settings Bits;
- i) Transmit TRAIN Pattern;
- j) Transmit TRAIN_DONE Pattern; ~~and~~
- k) Transmit MUX Sequence; ~~;~~
- l) [Transmit BREAK;](#)
- m) [Transmit START_PHY_TEST;](#) and
- n) [Transmit REJECT_PHY_TEST.](#)

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 81 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) Settings 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) Settings Bits Completed 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.

[6.8.3 State transitions applicable to all states](#)

Overview

shows the transitions that apply to all states in the SP state machine.

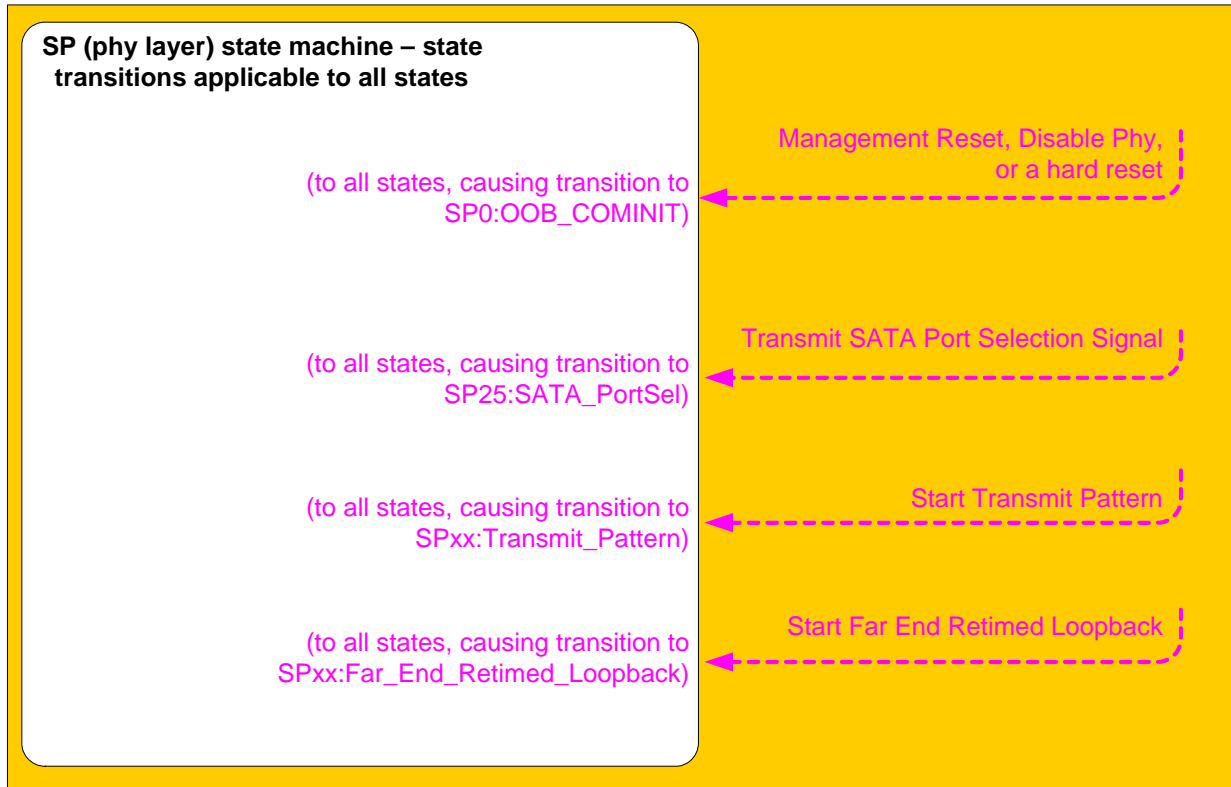


Figure 3 — SP (phy layer) state machine - transitions applicable to all states

6.8.3.1 Transition to SP0:OOB_COMINIT

If any state receives a Management Reset request, it shall set the MgmtReset state machine variable to one and perform this transition.

If any state receives a Disable Phy request or detects a hard reset, it shall set the MgmtReset state machine variable to zero and perform this transition.

6.8.3.2 Transition to SP25:SATA_PortSel

If the phy supports SATA port selectors, this transition shall occur after receiving a Transmit SATA Port Selection Signal request.

6.8.3.3 Transition to SPxx:Transmit_Pattern

This transition shall occur after receiving a Start Transmit Pattern request (e.g., from the SMP PHY TEST FUNCTION function requesting a phy test function of TRANSMIT_PATTERN in an expander device).

6.8.3.4 Transition to SPxy:Far_End_Retimed_Loopback

This transition shall occur after receiving a Start Far End Retimed Loopback request (e.g., from the SMP PHY TEST FUNCTION function requesting a phy test function of FAR_END_RETIMED_LOOPBACK in an expander device).

6.8.4 OOB sequence states

6.8.4.1 OOB sequence states overview

Figure 138 shows the OOB sequence states. These states are indicated by state names with a prefix of OOB.

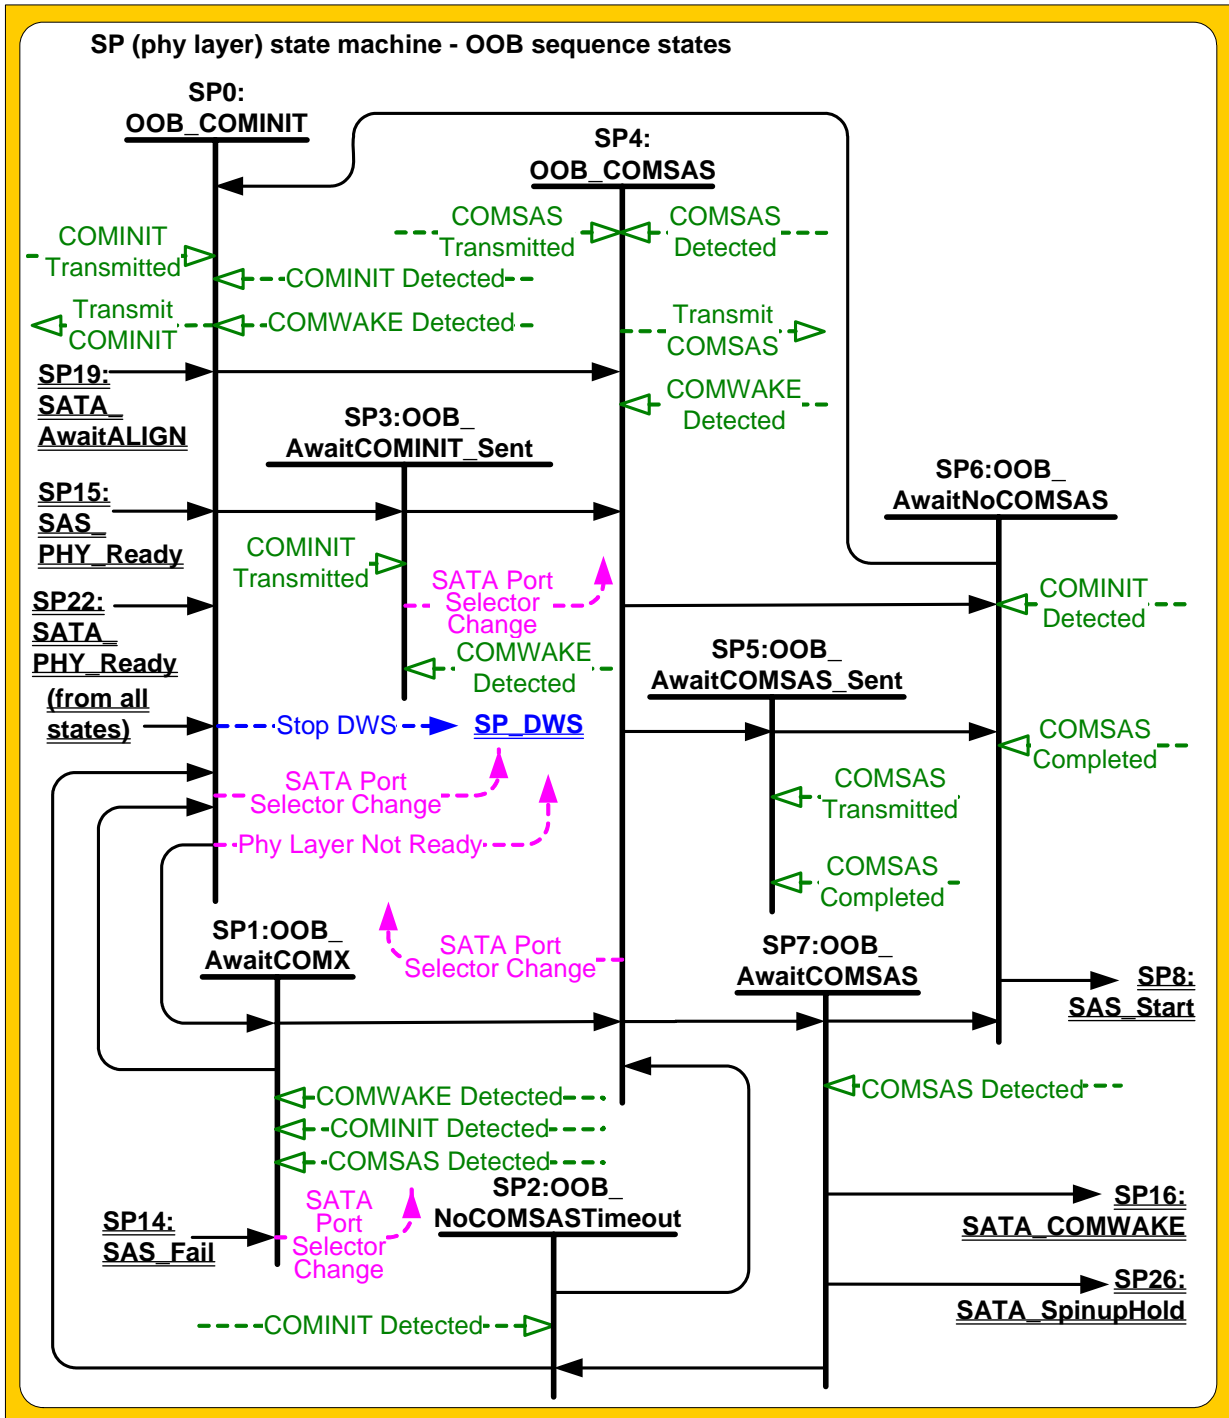


Figure 138 — SP (phy layer) state machine - OOB sequence states [modified]

6.8.5 SAS speed negotiation states

6.8.5.1 SAS speed negotiation states overview

Figure 139 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.

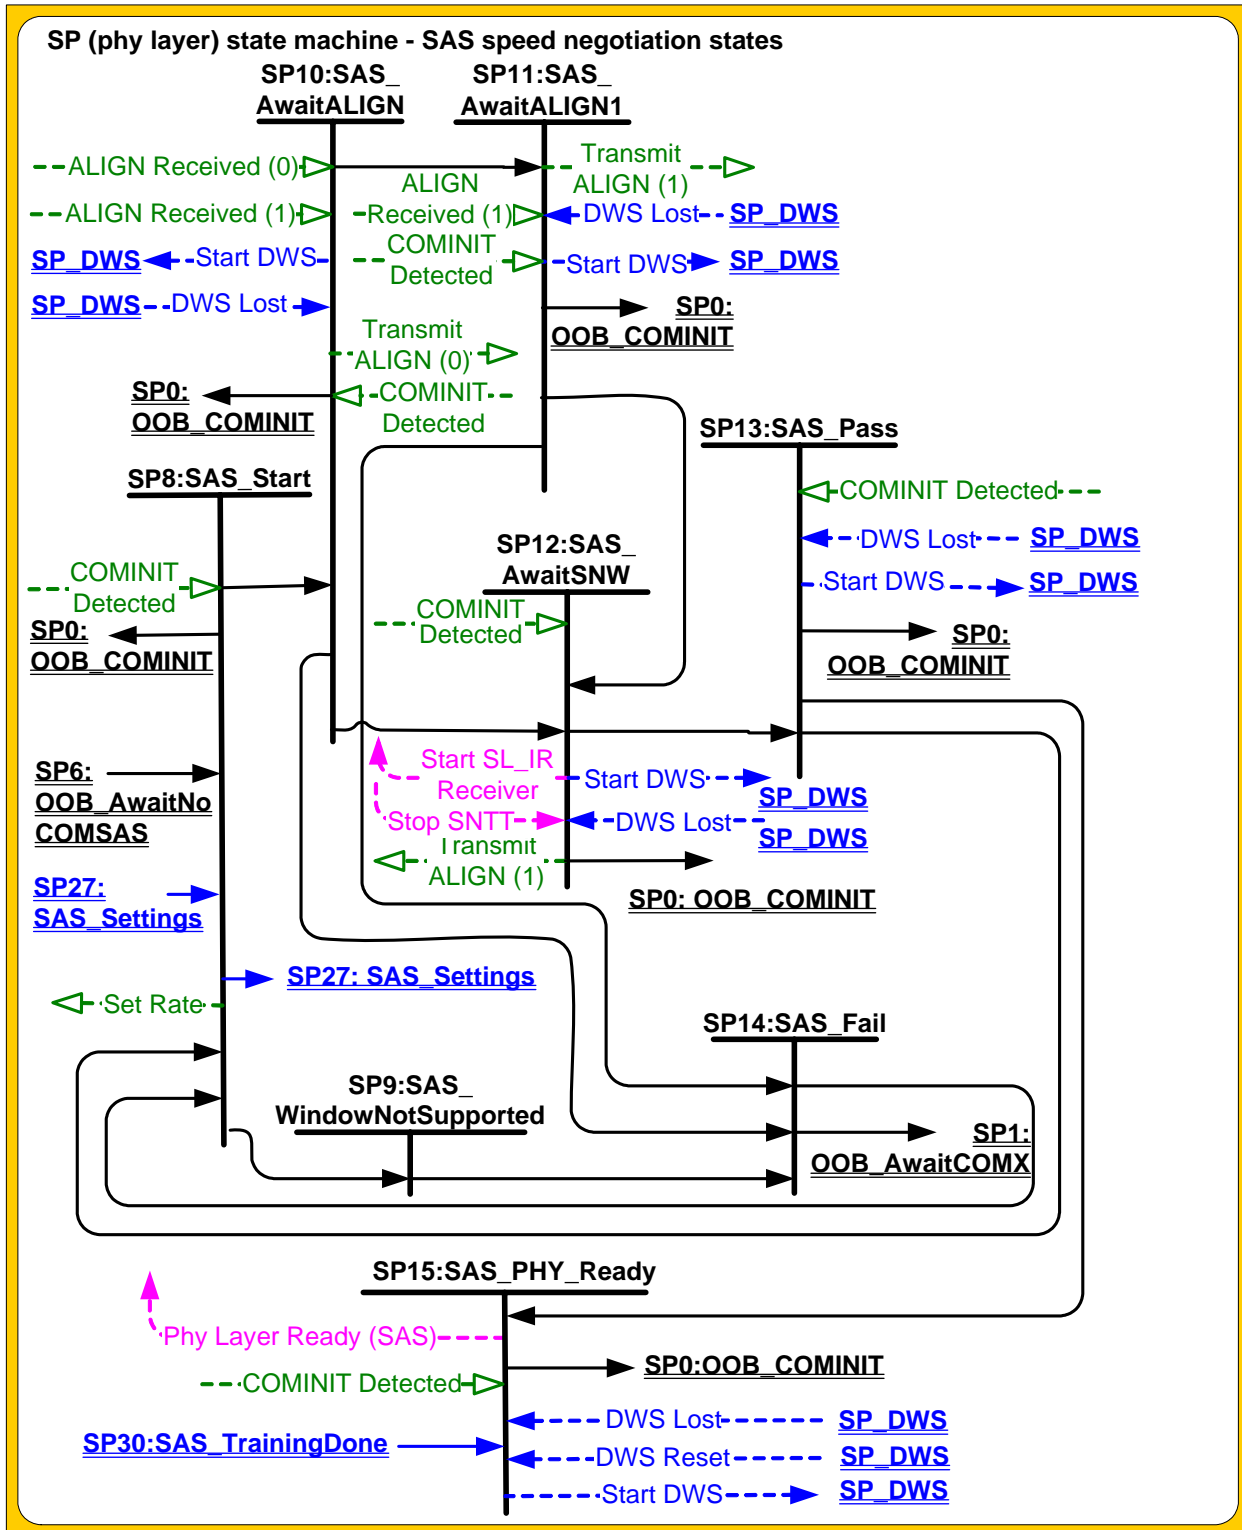


Figure 139 — SP (phy layer) state machine - SAS speed negotiation states [modified]

Figure 140 shows the SAS speed negotiation states related to SNW-3 and Train-SNW.

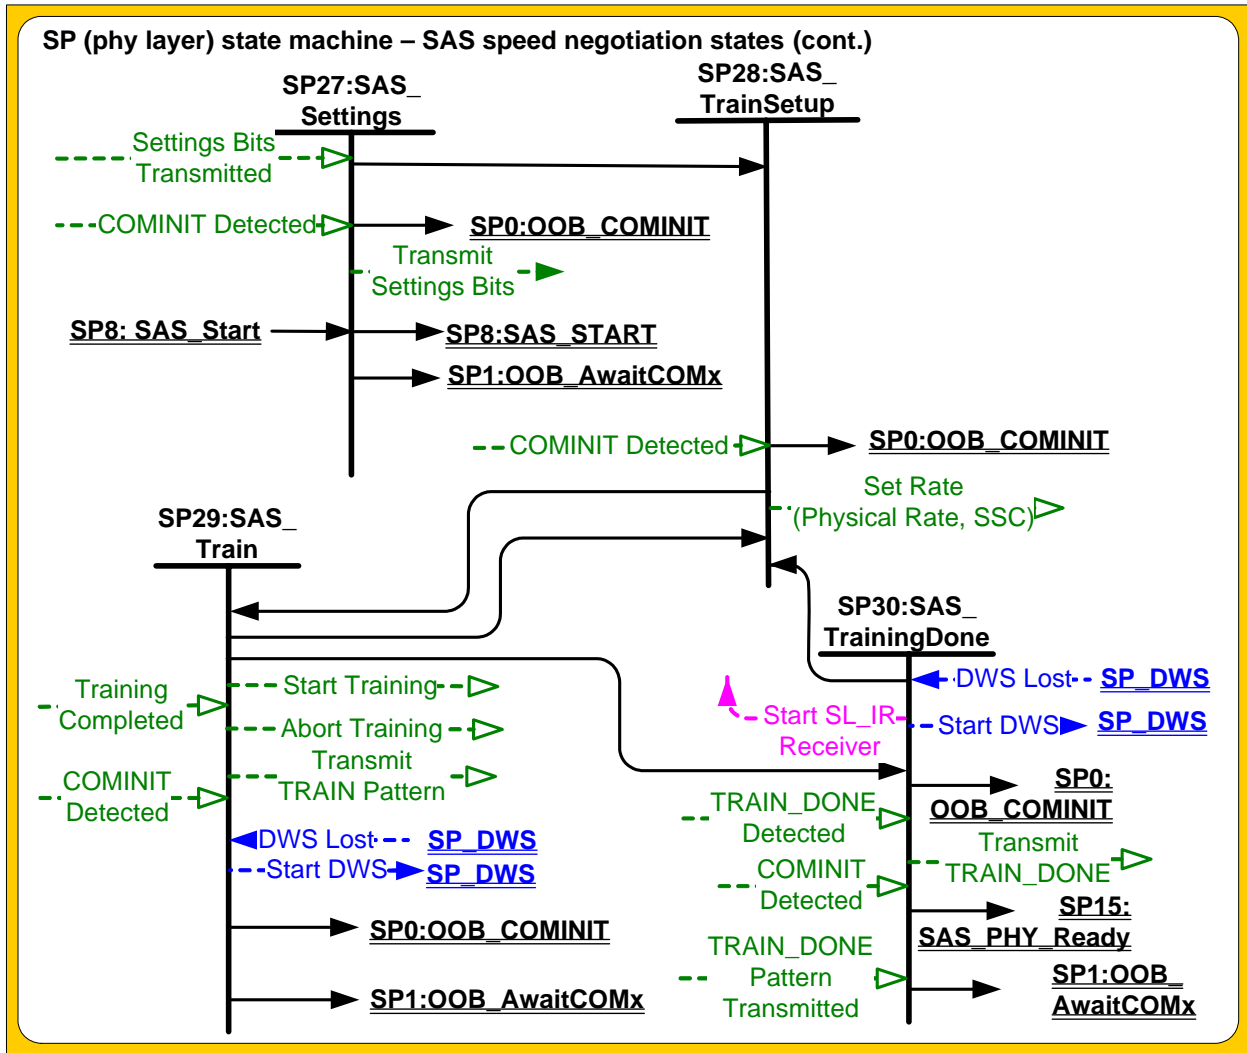


Figure 140 — SP (phy layer) state machine - SAS speed negotiation states for SNW-3 and Train-SNW [modified]

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.

Upon entry into this state, the phy shall:

- if multiplexing is enabled (see table 86 in 6.7.4.2.3.3), send a Transmit MUX Sequence message to the SP transmitter;
- 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;
- if the SP transmitter is transmitting at 1.5 Gbps, set the ResetStatus state machine variable to G1;
- if the SP transmitter is transmitting at 3 Gbps, set the ResetStatus state machine variable to G2;
- if the SP transmitter is transmitting at 6 Gbps, set the ResetStatus state machine variable to G3.

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.

If multiplexing is disabled, 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.

NOTE 2 - If multiplexing is enabled and this state receives a DWS Lost message, this state does not send a Start DWS message and the state machine transitions to SP0:OOB_COMINIT.

[If this state receives a Reject Phy Test request, this state shall send a Transmit REJECT_PHY_TEST message to the SP transmitter.](#)

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).

Editor's Note 3: Add the following two new states to the SAS-2 state diagram

[6.8.4.9 SPxx:Transmit Pattern state](#)

[6.8.4.9.1 State description](#)

[Upon entry into this state, the phy shall:](#)

- 1) [set the ResetStatus state machine variable to UNKNOWN;](#)
- 2) [send a Phy Layer Not Ready confirmation to the link layer;](#)
- 3) [if this state was entered from SP15:SAS_PHY_Ready, send a Transmit BREAK message to the SP transmitter;](#)
- 4) [if this state was entered from SP15:SAS_PHY_Ready or SP22:SATA_PHY_Ready, send a Transmit START_PHY_TEST message to the SP transmitter; and](#)
- 5) [send a Set Rate message to the SP transmitter with the Physical Link Rate argument, SSC On argument, and SSC Off argument set to the values received in the Start Transmit Pattern request;](#)
- 6) [perform the transmit pattern phy test function \(see 4.10\).](#)

[This state waits for a Stop Phy Test Function request.](#)

Editor's Note 4: Referring to a value received with a request while in the previous state generally requires passing the value "with the transition" in our horrid state machine conventions. Set Rate conceptually takes a long time (up to an RCDT), but there is no Set Rate Completed message coming back; should that interlock be added or ignored?

Editor's Note 5: Could define a new ResetStatus called PHY_TEST_FUNCTION. However, there is only one reserved value remaining (7h) for the Negotiated Physical Link Rate field. If a new value is chosen, 7h will have to be assigned to mean "look at another field" and another field in DISCOVER, et al. will need to be assigned where more values can be returned.

6.8.4.9.2 Transition SPxx:Transmit Pattern to SP0:OOB COMINIT

This transition shall occur after:

- a) receiving a COMINIT Detected message; or
- b) receiving a Stop Phy Test Function request.

Before the transition, this state shall set the ResetStatus state machine variable to UNKNOWN.

Editor's Note 6: Not documenting how the application layer decides to send a Stop Phy Test Function request to SP. SL/XL could be intermediaries.

6.8.4.9 SPxy:Far End Retimed Loopback state**6.8.4.9.1 State description**

Upon entry into this state, the phy shall:

- 1) send a Phy Layer Not Ready confirmation to the link layer;
- 2) if this state was entered from SP15:SAS PHY Ready, send a Transmit BREAK message to the SP transmitter;
- 3) if this state was entered from SP15:SAS PHY Ready or SP22:SATA PHY Ready, send a Transmit START_PHY_TEST message to the SP transmitter; and
- 4) send a Set Rate message to the SP transmitter with the Physical Link Rate argument, SSC On argument, and SSC Off argument set to the values received in the Start Far End Retimed Loopback request;
- 5) perform the far end retimed loopback phy test function (see 4.10).

This state waits for a COMINIT Detected message or a Stop Phy Test Function request.

Every time this state receives a DWS Lost message, this state shall send a Start Training message to the SP receiver and a Start DWS message to the SP DWS state machine.

Editor's Note 7: Not documenting the loopback process itself with any kind of link layer state machine.

Editor's Note 8: The SP receiver doesn't report a D.C. Idle Received message, so exactly how training is restarted is not precise (just based on DWS Lost).

Editor's Note 9: The PHY READY states only honor COMINIT Detected if it is preceded by DWS Lost (to avoid problems if the OOB detector falsely triggers in the middle of low amplitude data patterns that look like a COMINIT). If this state needs to make that distinction, then it might be better split into two states - one with dword sync, one without.

6.8.4.9.2 Transition SPxy:Far End Retimed Loopback to SP0:OOB COMINIT

This transition shall occur after:

- a) receiving a COMINIT Detected message; or
- b) receiving a Stop Phy Test Function request.

Before the transition, this state shall set the ResetStatus state machine variable to UNKNOWN.

Editor's Note 10: Not documenting how the application layer decides to send a Stop Phy Test Function request to SP. SL/XL could be intermediaries.

6.8.5 SATA host emulation states

6.8.5.1 SATA host emulation states overview

Figure 141 shows the SATA host emulation states, in which the phy has detected that it is attached to a SATA phy and behaves as if it were a SATA host phy, initiating the SATA speed negotiation sequence. These states are indicated by state names with a prefix of SATA.

The power management states defined in this standard are for SAS initiator phys that support SATA; expander devices that support SATA do not support power management in this standard.

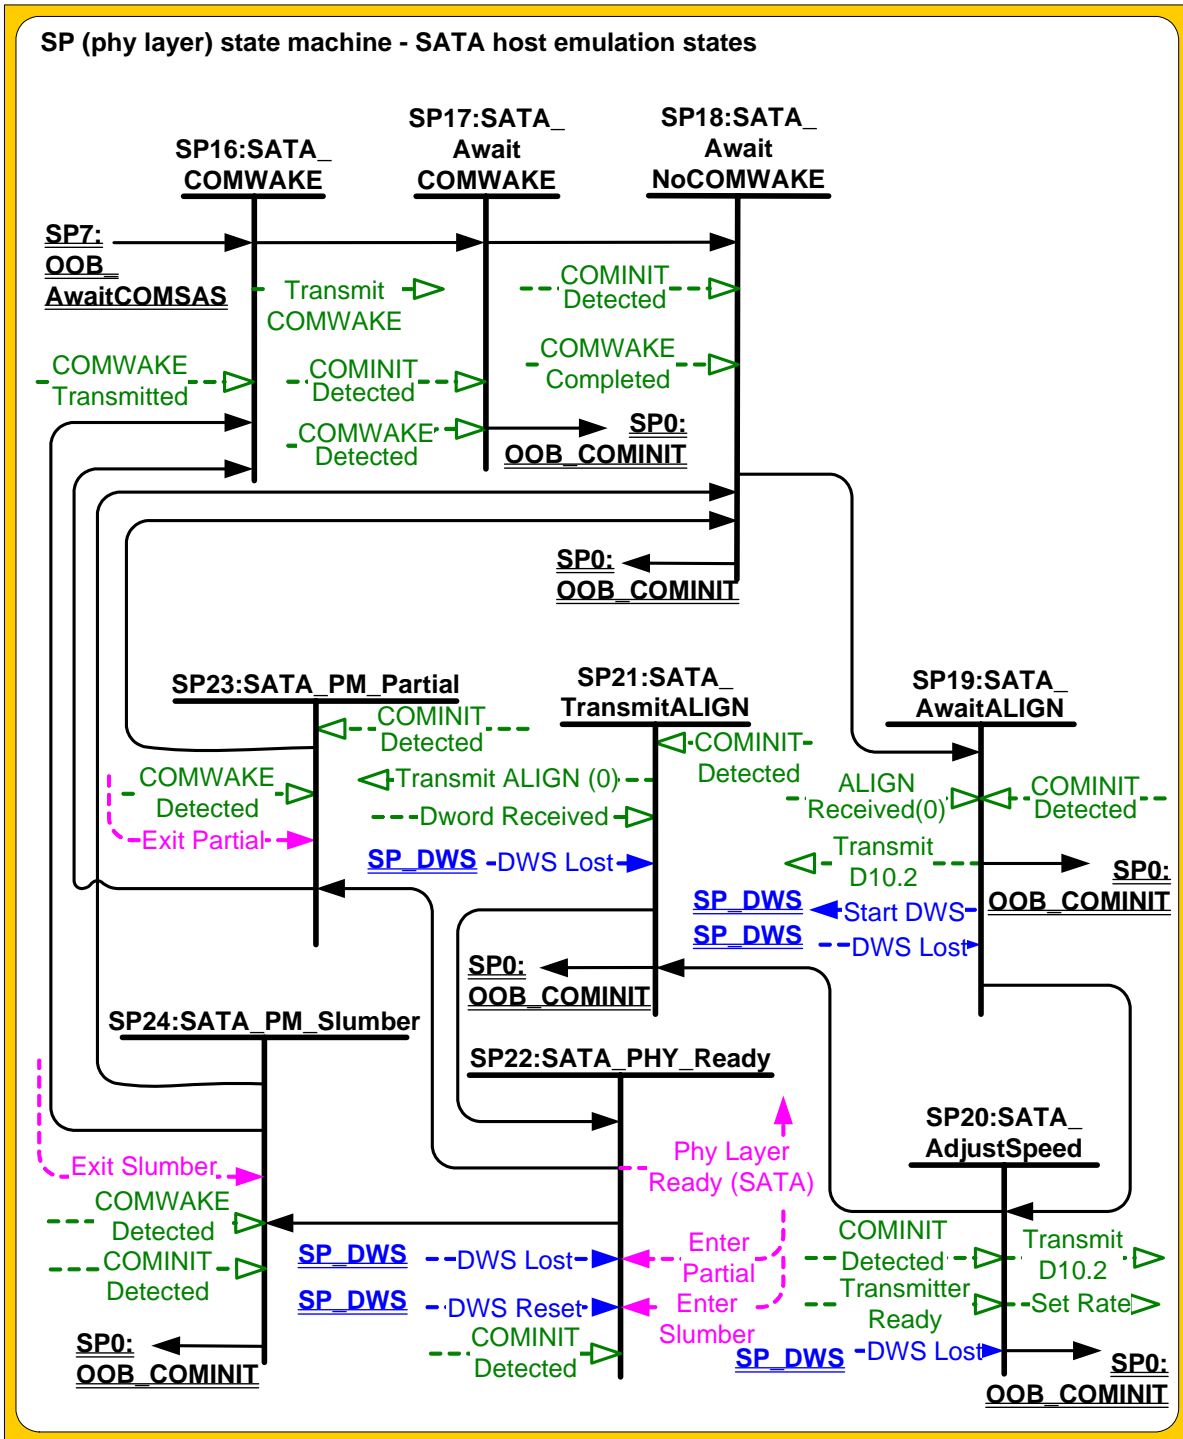


Figure 141 — SP (phy layer) state machine - SATA host emulation states [modified]

6.8.5.8 SP22:SATA_PHY_Ready state

6.8.5.8.1 State description

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

Upon entry into this state, this state shall:

- a) if the SP transmitter is transmitting at 1.5 Gbps, set the ResetStatus state machine variable to G1;
- b) if the SP transmitter is transmitting at 3 Gbps, set the ResetStatus state machine variable to G2; or
- c) if the SP transmitter is transmitting at 6 Gbps, set the ResetStatus state machine variable to G3.

This state shall send a Phy Layer Ready (SATA) confirmation to the link layer to indicate that the physical link has been brought up successfully in SATA mode.

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

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.

6.8.5.8.2 Transition SP22:SATA_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;
- a) receiving a DWS Lost message followed by a COMINIT Detected message, if this state does not send a Start DWS message; or
- b) 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).

6.8.5.8.3 Transition SP22:SATA_PHY_Ready to SP23:SATA_PM_Partial

This transition shall occur after receiving an Enter Partial request.

6.8.5.8.4 Transition SP22:SATA_PHY_Ready to SP24:SATA_PM_Slumber

This transition shall occur after receiving an Enter Slumber request.

6.8.6 SATA port selector state SP25:SATA_PortSel

6.8.6.1 State description

Figure 142 shows the SP25:SATA_PortSel state. This state controls transmission of the SATA port selection signal when a specified phy processes a Transmit SATA Port Selection Signal request.

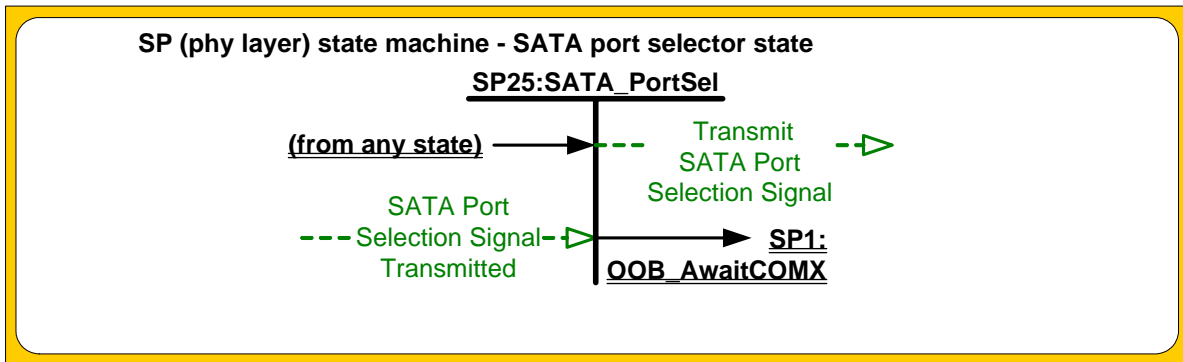


Figure 142 — SP (phy layer) state machine – SATA port selector state [\[modified\]](#)

Upon entry into this state, the phy shall:

- a) set the ResetStatus state machine variable to UNKNOWN;
- b) send a Transmit SATA Port Selection Signal message to the SP transmitter;

- c) set the ATTACHED SATA PORT SELECTOR bit to zero in the SMP DISCOVER response (see 10.4.3.9); and
- d) set the ATTACHED SATA DEVICE bit to zero in the SMP DISCOVER response.

6.8.6.2 Transition SP25:SATA_PortSel to SP1:OOB_AwaitCOMX

This transition shall occur after receiving a SATA Port Selection Signal Transmitted message.

6.8.7 SATA spinup hold state SP26:SATA_SpinupHold

6.8.7.1 State description

Figure 143 shows the SP26:SATA_SpinupHold state.

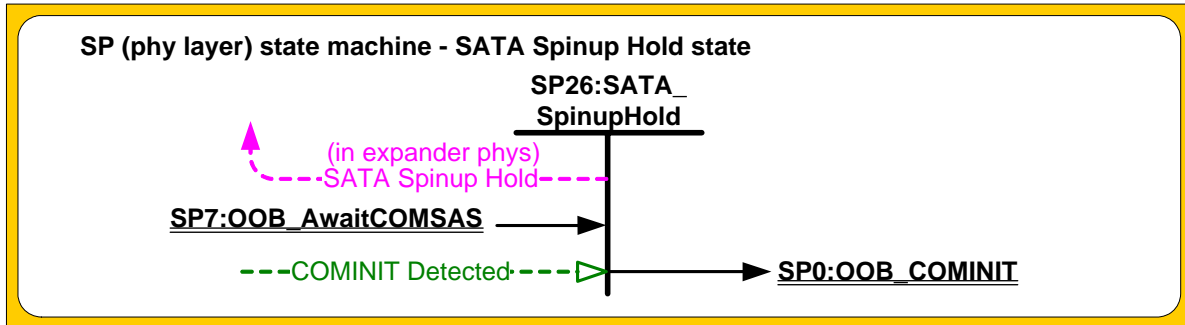


Figure 143 — SP (phy layer) state machine - SATA spinup hold state [\[modified\]](#)

Upon entry into this state, this state shall set the ResetStatus state machine variable to SPINUP_HOLD.

If this state machine is in an expander phy, this state shall send a SATA Spinup Hold confirmation to the link layer.

6.8.7.2 Transition SP26:SATA_SpinupHold to SP0:OOB_COMINIT

This transition shall occur if this state receives a COMINIT Detected message.

7.2 Primitives

7.2.1 Primitives overview

...

7.2.2 Primitive summary

Editor's Note 11: Add "PhyTest" footnote to all the primitive tables like this.

Table 91 — Deletable primitives

Primitive	Use ^a	From ^b			To ^b			Primitive sequence type ^c
		I	E	T	I	E	T	
^a The Use column indicates when the primitive is used: a) NoConn: SAS logical links, outside connections; b) Conn: SAS logical links, inside connections; c) All: SAS logical links, both outside connections or inside any type of connection; d) STP: SAS logical links, inside STP connections; of e) SpNeg: SAS physical links, during speed negotiation; or f) PhyTest: SAS physical links, during phy test functions. ^b ... ^c ...								

Table 92 defines the primitives not specific to the type of connection.

Table 92 — Primitives not specific to type of connection

Primitive	Use ^a	From ^b			To ^b			Primitive sequence type ^c
		I	E	T	I	E	T	
...								
AF_ACK	NoConn	I	E	I	I	E	I	Redundant
START_PHY_TEST	PhyTest	I	E	I	I	E	I	Redundant
REJECT_PHY_TEST	NoConn	I	E	I	I	E	I	Redundant
^a The Use column indicates when the primitive is used: a) NoConn: SAS logical links, outside connections; b) Conn: SAS logical links, inside connections; c) All: SAS logical links, both outside connections or inside any type of connection; d) STP: SAS logical links, inside STP connections; of e) SpNeg: SAS physical links, during speed negotiation; or f) PhyTest: SAS physical links, during phy test functions. ^b ...								

...

7.2.3 Primitive encodings

...

Table 96 defines the primitive encoding for primitives not specific to type of connection.

Table 96 — Primitive encoding for primitives not specific to type of connection

Primitive	Character				Hexadecimal
	1 st	2 nd	3 rd	4 th (last)	
...	K28.5	D02.0	D04.7	D01.4	BC02E481h
AF ACK	K28.5	D24.0	D27.4	D02.0	BC189B02h
START PHY TEST	K28.5	D24.0	D24.0	D24.0	BC181818h
REJECT PHY TEST	K28.5	D24.0	D29.7	D04.7	BC18FDE4h

7.2.3.1 AF ACK

[AF ACK](#) is used to positively acknowledge certain address frames (e.g., the [PHY TEST FUNCTION](#) address frame).

[See 7.8 for details on address frames.](#)

7.2.3.2 START PHY TEST

[START PHY TEST](#) is used to indicate the start of a phy test function.

[See 4.x for details on phy test functions.](#)

7.2.3.3 REJECT PHY TEST

[REJECT PHY TEST](#) is used to indicate that a phy test function requested by the the [PHY TEST FUNCTION](#) address frame is not supported.

[See 4.x for details on phy test functions.](#)

7.8 Address frames

7.8.1 Address frames overview

Address frames are used for the identification sequence and for connection requests.

Address frames are preceded by SOAF and followed by EOAF as shows in figure 158.




Figure 158 — Address frame transmission

Address frames shall only be sent outside connections. Address frames shall not be terminated early. All data dwords in an address frame shall be scrambled.

Table 114 defines the address frame format.

Table 114 — Address frame format

Byte\Bit	7	6	5	4	3	2	1	0
0					ADDRESS FRAME TYPE			
1	Frame type dependent bytes							
27								
28	(MSB)	CRC						
31								(LSB)

The ADDRESS FRAME TYPE field indicates the type of address frame and is defined in table 115. This field determines the definition of the frame type dependent bytes.

Table 115 — ADDRESS FRAME TYPE field

Code	Address frame type	Description	Acknowledgement method(s)
0h	IDENTIFY	Identification sequence	See 7.9
1h	OPEN	Connection request	See 7.12.2.2
Fh	PHY TEST FUNCTION	Phy test function	AF ACK
All others	Reserved		

The CRC field contains a CRC value (see 7.5) that is computed over the entire address frame prior to the CRC field.

Address frames with unknown address frame types, incorrect lengths, or CRC errors shall be ignored by the recipient.

7.8.2 IDENTIFY address frame

...

7.8.3 OPEN address frame

...

7.8.4 PHY TEST FUNCTION address frame

Table 116 defines the PHY TEST FUNCTION address frame, which is used to request that a phy perform a phy test function (see 4.x).

Table 116 — PHY TEST FUNCTION address frame format [all new]

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved				ADDRESS FRAME TYPE (Fh)			
1	Reserved							
9	Reserved							
10	PHY TEST FUNCTION							
11	PHY TEST PATTERN							
12	Reserved							
14	Reserved							
15	Reserved				PHY TEST FUNCTION PHYSICAL LINK RATE			
16	Reserved							
18	Reserved							
19	PHY TEST PATTERN DWORDS CONTROL							
20	PHY TEST PATTERN DWORDS							
27	PHY TEST PATTERN DWORDS							
28	(MSB)	CRC						(LSB)
31	(LSB)							

The ADDRESS FRAME TYPE field shall be set to Fh.

All fields not defined in this subclause are defined in the SMP PHY TEST FUNCTION (see 10.x.x.x), with the selected phy being the phy that received the PHY TEST FUNCTION address frame.

The CRC field is defined in 7.8.1.

7.13 Phy test function sequence [all new]

A phy may use the PHY TEST FUNCTION address frame to request the phy to which it is attached perform a phy test function.

After a phy transmits a PHY TEST FUNCTION address frame, it shall:

- a) not transmit an OPEN address frame;
- b) start a Phy Test Function Request timer;

- c) expect one of the results listed in table 117.

Table 117 — Results of a PHY TEST FUNCTION address frame

Result	Description
Receive AF_ACK	The attached phy received the PHY TEST FUNCTION address frame and is processing it.
Receive OPEN address frame	The attached phy is making a connection request. The phy shall ignore the connection request. This leads to the attached phy's Open Timeout timer expiring.
Receive PHY TEST FUNCTION address frame	The attached phy is concurrently requesting a phy test function. The phy shall ignore the PHY TEST FUNCTION address frame and expect another result in this table. This leads to the attached phy's Phy Test Function Request timer expiring.
Phy Test Function Request timer expires	The attached phy did not receive the PHY TEST FUNCTION address frame (e.g., due to a CRC error), or does not support it

After receiving AF_ACK, the phy shall start a Phy Test Function Decision timer and expect one of the results listed in table 118.

Table 118 — Results of a PHY TEST FUNCTION address frame after AF_ACK

Result	Description
START_PHY_TEST	The attached phy has accepted the request and is starting the phy test function
REJECT_PHY_TEST	The attached phy has rejected the request and is resuming normal operation. The phy may resume using the physical link (e.g., transmitting OPEN address frames).
Phy Test Function Decision timer expires	The attached phy did not respond. The phy shall perform a link reset sequence.

[\[end of all-new section\]](#)

7.14 SL (link layer for SAS phys) state machine

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 AF_ACK](#);
- c) Transmit SOAF/Data Dwords/EOAF;
- d) Transmit OPEN_ACCEPT;
- e) Transmit OPEN_REJECT with an argument indicating the specific type (e.g., Transmit OPEN_REJECT (Retry));
- f) Transmit BREAK;
- g) Transmit BREAK_REPLY;
- h) Transmit BROADCAST; and
- i) 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.

...

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 [or a PHY TEST FUNCTION 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 **Open1h (i.e., OPEN) or Fh (i.e., PHY TEST FUNCTION)**;
- 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 [address frame](#) is ~~not discarded~~ [an OPEN address frame](#) 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.

[If the address frame is a PHY TEST FUNCTION address frame then this state machine shall:](#)

- a) [send a Transmit AF ACK message to the SL transmitter; and](#)
- b) [send a Phy Test Function Received confirmation to the management application layer including the contents of the address frame.](#)

Editor's Note 12: In the "SL state machines (part 2)" figure, add a Phy Test Function Received confirmation from SL_RA

Editor's Note 13: Not documenting how the application layer parses the address frame contents and decides to send a Start Far End Retimed Loopback request or Start Transmit Pattern request to SP.

Editor's Note 14: No other SL changes proposed for now. Could have SL_CC0:Idle receive the Send Phy Test Function request (from the application client parsing the diagnostic page or otherwise deciding it wants to send it) and transmit the address frame, then transition to a new SL_CC state to wait for AF_ACK. After AF_ACK is received or a timeout, transition to a new SL_CC state to wait for START_PHY_TEST or REJECT_PHY_TEST. This would keep SL_CC from parsing new address frames after the test has begun. Also, not sending SL_CC into a new state during the test function (incoming requests).

Editor's Note 15: Not documenting transmitting the PHY TEST FUNCTION address frame based on processing the SMP function. Not sending XL into a new state during the test function.

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_HOLD 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 3 - The XL transmitter may always insert a deletable primitive before transmitting a BREAK, BREAK_REPLY, CLOSE, or SATA_HOLD 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_HOLD 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 4 - 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.

7.15.3 XL0:Idle state

7.15.3.1 State description

This state is the initial state and is the state that is used when there is no connection pending or established.

If a Phy Layer Not Ready confirmation is received, this state shall send a Broadcast Event Notify (Phy Not Ready) request to the BPP.

If a SATA Spinup Hold confirmation is received, this state shall send a Broadcast Event Notify (SATA Spinup Hold) request to the BPP.

If an Enable Disable SAS Link (Enable) message is received, this state shall send a Broadcast Event Notify (Identification Sequence Complete) request to the BPP.

If a SATA Port Selector Change confirmation is received, this state shall send a Broadcast Event Notify (SATA Port Selector Change) request to the BPP.

If a PHY TEST FUNCTION Address Frame Received message is received, this state shall send a Phy Test Function Received confirmation to the management application layer including the contents of the address frame.

If a BROADCAST Received message is received, this state shall send a Broadcast Event Notify request to the BPP with the argument indicating the specific BROADCAST primitive received (e.g., Change Received).

If a Transmit Broadcast indication is received, this state shall send a Transmit BROADCAST message to the XL transmitter with an argument specifying the specific type from the Transmit Broadcast indication. Otherwise, this state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the XL transmitter.

If a BREAK Received message is received and the BREAK_REPLY method of responding to received BREAK primitive sequences is enabled (see 7.12.5), then this state shall send a Transmit BREAK_REPLY message to the XL transmitter.

7.15.3.2 Transition XL0:Idle to XL1:Request_Path

This transition shall occur if:

- a) an Enable Disable SAS Link (Enable) message has been received;
- b) a Forward Open indication is not being received; and
- c) an OPEN Address Frame Received message is received.

10.2 SCSI application layer

10.2.9 SCSI diagnostic parameters

10.2.9.1 Protocol-Specific diagnostic page

The Protocol-Specific diagnostic page for SAS provides a method for an application client to enable and disable phy test functions (see 4.10) for selected phys. The diagnostic page format is specified in SPC-4.

The Protocol-Specific diagnostic page is transmitted using the SEND DIAGNOSTIC command. If the device server receives a RECEIVE DIAGNOSTIC RESULTS command with the PAGE CODE field set to 3Fh, it shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and

the additional sense code set to INVALID FIELD IN PARAMETER LIST. Table 203 defines the Protocol-Specific diagnostic page for SAS.

Table 203 — Protocol-Specific diagnostic page for SAS

Byte\Bit	7	6	5	4	3	2	1	0
0	PAGE CODE (3Fh)							
1	Reserved				PROTOCOL IDENTIFIER (6h)			
2	(MSB)	PAGE LENGTH (001Ch)						(LSB)
3								
4	PHY IDENTIFIER							
5	PHY TEST FUNCTION							
6	PHY TEST PATTERN							
7	Reserved				PHY TEST PATTERN <u>FUNCTION</u> PHYSICAL LINK RATE			
8	Reserved							
10								
11	PHY TEST PATTERN DWORDS CONTROL							
12								
19	PHY TEST PATTERN DWORDS							
20	Reserved							
31								

The PAGE CODE field shall be set to 3Fh.

The PROTOCOL IDENTIFIER field shall be set to 6h indicating this is a SAS SSP specific diagnostic page.

The PAGE LENGTH field shall be set to the number of bytes in the diagnostic page after the PAGE LENGTH field (i.e., 001Ch)

The PHY IDENTIFIER field specifies the phy identifier (see 4.2.7) of the phy that is to perform or to stop performing a phy test function (i.e., the selected phy). If the PHY IDENTIFIER field specifies a phy that does not exist, then the device server shall terminate the SEND DIAGNOSTIC command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The PHY TEST FUNCTION field specifies the phy test function to be performed and is defined in table 204. If the PHY TEST FUNCTION field specifies a phy test function that is not supported, then the device server shall terminate the SEND DIAGNOSTIC command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

Table 204 — PHY TEST FUNCTION field

Code	Name	Description
00h	STOP	<p>If the selected phy is performing a phy test function, then the selected phy shall stop performing the phy test function and originate a link reset sequence.</p> <p>If the selected phy is not performing a phy test function, then this function has no effect on the selected phy. ^a</p>
01h	TRANSMIT_PATTERN	<p>If the selected phy is not performing a phy test function, the selected phy shall <u>perform the transmit pattern phy test function (see 4.10.2) be set to transmit the using the</u> phy test pattern specified by the PHY TEST PATTERN field at the physical link rate specified by the PHY TEST PATTERN PHYSICAL LINK RATE field and set to ignore its receiver. If the selected phy receives data while transmitting the pattern, then the selected phy shall ignore the received data.</p> <p>If the selected phy is performing a phy test function, the device server shall terminate the SEND DIAGNOSTIC command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to PHY TEST FUNCTION IN PROGRESS. ^a</p>
<u>02h</u>	<u>FAR_END_RETIMED_LOOPBACK</u>	<p><u>If the selected phy is not performing a phy test function, the selected phy shall perform the far-end retimed loopback phy test function (see 4.10.3) at the physical link rate specified by the PHY TEST PATTERN PHYSICAL LINK RATE field.</u></p> <p><u>If the selected phy is performing a phy test function, the device server shall terminate the SEND DIAGNOSTIC command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to PHY TEST FUNCTION IN PROGRESS.</u> ^a</p>
02h <u>03h</u> - EFh	Reserved	
F0h - FFh	Vendor specific	
<p>^a If there is no SSP target port available to receive a SEND DIAGNOSTIC command to stop a phy from performing a phy test function, then a power on may be required to cause the phy to stop performing the function and originate a phy reset sequence.</p>		

If the PHY TEST FUNCTION field is set to TRANSMIT_PATTERN (i.e., 01h), then the PHY TEST PATTERN field specifies the phy test pattern to be transmitted as defined by table 205. If the PHY TEST PATTERN field specifies a phy test pattern that is not supported by the specified SAS phy, then the device server shall terminate the

SEND DIAGNOSTIC command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

Table 205 — PHY TEST PATTERN field

Code	Name	Description
00h	Reserved	
01h	JTPAT	The selected phy shall continuously transmit the JTPAT for RD+ and RD- (see A.1).
02h	CJTPAT	The selected phy shall continuously transmit the CJTPAT (see A.2).
03h - 3Fh	Reserved	
40h	TWO_DWORDS	The selected phy shall continuously transmit the dwords specified by the PHY TEST PATTERN DWORDS CONTROL field and the PHY TEST PATTERN DWORDS field without scrambling. This pattern is only for use for characterization of the transmitter device and the passive interconnect. Phys are not required to support all patterns that may be specified.
41h - EFh	Reserved	
F0h - FFh	Vendor specific	

The PHY TEST PATTERN PHYSICAL LINK RATE field specifies the physical link rate at which the phy test pattern shall be transmitted and is defined in table 206. If the physical link rate specified by the PHY TEST PATTERN PHYSICAL LINK RATE field is less than the hardware minimum physical link rate or greater than the hardware maximum physical link rate, then the device server shall terminate the SEND DIAGNOSTIC command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

Table 206 — PHY TEST PATTERN PHYSICAL LINK RATE field

Code	Description
0h - 7h	Reserved
8h	1,5 Gbps
9h	3 Gbps
Ah	6 Gbps
Bh - Fh	Reserved for future physical link rates

The PHY TEST PATTERN DWORDS CONTROL field and PHY TEST PATTERN DWORDS field are only used if the PHY TEST PATTERN field is set to TWO_DWORDS (see table 205).

The PHY TEST PATTERN DWORDS CONTROL field defined in table 207 controls whether the bytes in the PHY TEST PATTERN DWORDS field are sent as control characters or data characters.

Table 207 — PHY TEST PATTERN DWORDS CONTROL field

Code	Description
00h	Each byte in the PHY TEST PATTERN DWORDS field shall be sent as a data character (i.e., Dxx.y)(see 6.3.3) without scrambling.
08h	The fifth byte in the PHY TEST PATTERN DWORDS field shall be sent as a control character (i.e., Kxx.y)(see 6.3.3); each other byte shall be sent as a data character without scrambling.
80h	The first byte in the PHY TEST PATTERN DWORDS field shall be sent as a control character; each other byte shall be sent as a data character without scrambling.
88h	The first and fifth bytes in the PHY TEST PATTERN DWORDS field shall each be sent as a control character; each other byte shall be sent as a data character without scrambling.
All others	Reserved

The PHY TEST PATTERN DWORDS field contains the two dwords that are sent during a TWO_DWORDS test pattern. Whether each byte in the dwords is sent as a control character or a data character is specified by the PHY TEST PATTERN DWORDS CONTROL field. A byte specifying a control character shall only specify a control character which is used in this standard (see table 75 in 6.3) and is supported by the phy (i.e., all phys support K28.5 (i.e., BCh), but only phys supporting STP support K28.3 (i.e., 7Ch) or K28.6 (i.e., DCh)).

The device server shall terminate a SEND DIAGNOSTIC command specifying any unsupported combination with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

Table 208 lists some examples of TWO_DWORDS phy test patterns.

Table 208 — TWO_DWORDS phy test pattern examples

PHY TEST PATTERN DWORDS CONTROL field	PHY TEST PATTERN DWORDS field	Description
00h	4A4A4A4A 4A4A4A4Ah	D10.2 characters (see table 73 in 6.3.3). This high-frequency pattern contains 01b repeating and has the highest possible frequency. This pattern may be used for measuring skew and rise/fall times (see table 59 in 5.3.6.2, table 60 in 5.3.6.3, and table 62 in 5.3.7.2).
00h	B5B5B5B5 B5B5B5B5h	D21.5 characters (see table 73 in 6.3.3). This high-frequency pattern contains 10b repeating and has the highest possible frequency. This pattern may be used for measuring skew and rise/fall times (see table 59 in 5.3.6.2, table 60 in 5.3.6.3, and table 62 in 5.3.7.2).
00h	78787878 78787878h	D24.3 characters (see table 73 in 6.3.3). This mid-frequency pattern contains 0011b or 1100b repeating (depending on starting disparity) and has half the highest possible frequency.
00h	D926D926 D926D926h	Pairs of D25.6 and D6.1 characters (see table 73 in 6.3.3). This mid-frequency pattern contains 1001b repeating and has half the highest possible frequency.
00h	7E7E7E7E 7E7E7E7Eh	D30.3 characters (see table 73 in 6.3.3). This low-frequency pattern contains four bits of one polarity, three bits of the other polarity, and three bits of the first polarity (e.g., 1111000111b), followed by the inverse (e.g., 0000111000b).
00h	EBF4EBF4 EBF4EBF4h	Pairs of these D11.7 and D20.7 characters (see table 73 in 6.3.3). This pattern contains a single bit of one polarity after five bits of the other polarity (i.e., 0000010b and 1111101b). This pattern may be used as a lone bit pattern for measuring jitter in external cables (see table 52 in 5.2.6).
88h	BC4A4A7B BC4A4A7Bh	ALIGN (0) primitives (see table 102 in 7.2.3). This pattern appears during OOB bursts (see 6.6), the SATA speed negotiation sequence (see 6.7.2.2), and the SAS speed negotiation sequence (see 6.7.4.2).
88h	BC070707 BC070707h	ALIGN (1) primitives (see table 102 in 7.2.3). This pattern appears during the SAS speed negotiation sequences (see 6.7.4.2).
80h	BC4A4A7B 4A787E7Eh	Pairs of an ALIGN (0) primitive (see table 102 in 7.2.3) and a dword containing D10.2, D24.3, D30.3, and D30.3 characters (see table 73 in 6.3.3).

10.4 Management application layer

10.4.3.27 PHY TEST FUNCTION function

The PHY TEST FUNCTION function requests actions by the specified phy. This SMP function may be implemented by any management device server. In zoning expander devices, if zoning is enabled then this function shall only be processed from SMP initiator ports that have access to zone group 2 or the zone group of the specified phy (see 4.9.3.2).

Table 306 defines the request format.

Table 306 — PHY TEST FUNCTION request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (92h)							
2	Reserved							
3	REQUEST LENGTH (09h)							
4	(MSB)	EXPECTED EXPANDER CHANGE COUNT						(LSB)
5								
6	Reserved							
8								
9	PHY IDENTIFIER							
10	PHY TEST FUNCTION							
11	PHY TEST PATTERN							
12								
14	Reserved							
15	Reserved				PHY TEST PATTERN FUNCTION PHYSICAL LINK RATE			
16	Reserved							
18								
19	PHY TEST PATTERN DWORDS CONTROL							
20								
27	PHY TEST PATTERN DWORDS							
28								
39	Reserved							
40	(MSB)	CRC						(LSB)
43								

Editor’s Note 16: Whether the phy chooses to enable SSC to perform its transmit test at the specified physical link rate is vendor-specific. If a field is added to PHY CONTROL to control this, then a similar field should be added here. The field should be able to specify none, center-spreading, and down-spreading; expanders would support all, while end devices would not support center-spreading.

The SMP FRAME TYPE field shall be set to 40h.

The FUNCTION field shall be set to 92h.

The REQUEST LENGTH field shall be set to 09h. For compatibility with previous versions of this standard, a REQUEST LENGTH field set to 00h specifies that there are 9 dwords before the CRC field.

The EXPECTED EXPANDER CHANGE COUNT field is defined in the SMP CONFIGURE GENERAL request (see 10.4.3.15).

The PHY IDENTIFIER field specifies the phy (see 4.2.7) to which the SMP PHY TEST PATTERN request applies.

If the PHY IDENTIFIER field specifies the phy which is being used for the SMP connection, the management device server shall not perform the requested operation and shall return a function result of SMP FUNCTION FAILED in the response frame.

The PHY TEST FUNCTION field specifies the phy test function to be performed, and is defined in table 307. If the PHY TEST FUNCTION field specifies a phy test function that is not supported by the phy, the management device server shall return a function result of UNKNOWN PHY TEST FUNCTION in the response frame.

Table 307 — PHY TEST FUNCTION field

Code	Name	Description
00h	STOP	<p>If the selected phy is performing a phy test function, then the selected phy shall stop performing the phy test function and originate a link reset sequence.</p> <p>If the selected phy is not performing a phy test function, then this function has no effect on the selected phy.</p>
01h	TRANSMIT_PATTER	<p>If the selected phy is not performing a phy test function, the selected phy shall perform the transmit pattern phy test function (see 4.10.2) be set to transmit the using the phy test pattern specified by the PHY TEST PATTERN field at and the physical link rate specified by the PHY TEST PATTERN PHYSICAL LINK RATE field and set to ignore its receiver. If the selected phy receives data while transmitting the pattern, then the selected phy shall ignore the received data.</p> <p>If the selected phy is performing a phy test function, the management device server shall return a function result of PHY TEST FUNCTION IN PROGRESS in the response frame.</p>
02h	FAR_END_RETIMED_LOOPBACK	<p>If the selected phy is not performing a phy test function, the selected phy shall perform the far-end retimed loopback phy test function (see 4.10.3) at the physical link rate specified by the PHY TEST PATTERN PHYSICAL LINK RATE field.</p> <p>If the selected phy is performing a phy test function, the device server shall return a function result of PHY TEST FUNCTION IN PROGRESS in the response frame.</p>
02h 03h - EFh	Reserved	
F0h - FFh	Vendor specific	

If the PHY TEST FUNCTION field is set to 01h ([i.e., TRANSMIT_PATTER](#)), the PHY TEST PATTERN field specifies the phy test pattern to be performed, and is the same as that defined in table 205 for the Protocol-Specific diagnostic page (see 10.2.9.1). The phy test pattern shall be sent at the physical link rate specified by the PHY TEST PATTERN PHYSICAL LINK RATE field.

The PHY TEST **PATTERNFUNCTION** PHYSICAL LINK RATE field specifies the physical link rate at which the phy test function, if any, shall be performed. Table 308 defines the values for this field.

Table 308 — PHY TEST **PATTERNFUNCTION PHYSICAL LINK RATE field**

Code	Description
0h - 7h	Reserved
8h	1,5 Gbps
9h	3 Gbps
Ah	6 Gbps
Bh - Fh	Reserved for future physical link rates

The PHY TEST PATTERN DWORDS CONTROL field and the PHY TEST PATTERN DWORDS field are as defined in table 205 for the Protocol-Specific diagnostic page (see 10.2.9.1).

The CRC field is defined in 10.4.3.1.

Table 309 defines the response format.

Table 309 — PHY TEST FUNCTION response

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (92h)								
2	FUNCTION RESULT								
3	RESPONSE LENGTH (00h)								
4	(MSB)	CRC							
7								(LSB)	

The SMP FRAME TYPE field shall be set to 41h.

The FUNCTION field shall be set to 92h.

The FUNCTION RESULT field is defined in 10.4.3.2.

The RESPONSE LENGTH field shall be set to 00h.

The CRC field is defined in 10.4.3.2.

Annex K (informative)

Primitive encoding

Table K.1 describes a set of the K28.5-based primitive encodings whose 40-bit values (after 8b10b encoding with either starting running disparity) have a Hamming distance (i.e., the number of bits different in two

patterns) of at least 8. All the primitive encodings in 7.2 except for TRAIN and TRAIN_DONE were selected from this list. Unassigned encodings may be used by future versions of this standard.

Table K.1 — Primitives with Hamming distance of 8

Character				Assignment
1 st	2 nd	3 rd	4 th	
				...
K28.5	D24.0	D24.0	D24.0	START PHY TEST
K28.5	D24.0	D27.4	D02.0	AF ACK
K28.5	D24.0	D29.7	D04.7	REJECT PHY TEST