

To: T10 Technical Committee
 From: Rob Elliott (elliott@hp.com) and Thomas Grieff (thomas.grieff@hp.com), HP
 Date: 11 July 2002
 Subject: T10/02-198r5 SAS OOB timing

Revision History

Revision 0 (20 May 2002) first revision

Revision 1 (27 May 2002) incorporated feedback from the 21 May 2002 conference call.

Revision 2 (2 June 2002) added power on to phy reset sequence and hard reset to phy reset sequence times.

Revision 3 (19 June 2002) incorporated feedback from 6 June 2002 SAS WG. Use ALIGN and ALIGN(1) rather than D10.2 and ALIGN for speed negotiation. Clarified how COMSAS negation time is used in the state machine. Added SMP report bit to indicate problems with speed negotiation, so applications can turn off high speeds if problems are being reported.

Revision 4 (28 June 2002) incorporated feedback from 24 June 2002 SAS WG. Set a minimum lock time that ensures time for the ALIGN(1) to be detected by the other side. Changed the OOB signal detection signal to trigger after 4 idles, not 4 bursts.

Revision 5 (11 July 2002) incorporated feedback from 9 July 2002 SAS conference call. Add clarification of terms and definition of the sequences. Specify RCD in UI(OOB) rather than ns.

Related Documents

sas-r00 - Serial Attached SCSI revision 0

Overview

The OOB timing requirements need to be more precise, listing the transmit burst time and specifying receiver tolerances. Various other minor changes are also included.

Also, the May T10 SAS study group chose to make the COMSAS burst time match the other OOB signals, with its idle time increased to 3x the COMINIT idle time.

Suggested Changes

6.3 Out of band (OOB) signals

~~SATA defines "out of band" (OOB) signals used for phy resets. OOB signals are low-speed signal patterns detected by the phy that do not appear in normal data streams. They consist of 160 UI (106.7 ns) bursts of ALIGNs followed by defined amounts of idle time. The signals are differentiated by the length of idle time between the bursts of ALIGNs.~~

~~SATA defines two OOB signals: COMINIT/COMRESET and COMWAKE. SAS devices identify themselves with a SAS-specific OOB signal called COMSAS.~~

[Editor's note: COMSAS timing is under debate again.]

~~Table 21 describes the SATA OOB signals. Each signal shall be sent six times. Each signal shall be received four times to be detected.~~

Table 21. OOB signal transmitter requirements

Signal	Idle time minimum	Idle time nominal	Idle time maximum	Notes
COMWAKE	55 ns	106.7 ns (160 UIs)	175 ns	
COMINIT/COMRESET	175 ns	320 ns (480 UIs)	525 ns	Named COMRESET when sent by an initiator, and named COMINIT when sent by a target.

~~SATA defines "out of band" (OOB) signals used for phy resets. OOB signals are low-speed signal patterns detected by the phy that do not appear in normal data streams. They consist of bursts (repeated sequences) of ALIGNs followed by defined amounts of idle time (driving with no transitions). During the burst time, ALIGN or ALIGN(1) primitives are sent repeatedly. During the idle time, the transmitter shall be enabled and transmit either~~

a low or high level but shall not toggle between the levels. The signals are differentiated by the length of idle time between the bursts of ALIGNs.

SATA defines two OOB signals: COMINIT/COMRESET and COMWAKE. COMINIT and COMRESET are used in this standard interchangeably. SAS devices identify themselves with an additional SAS-specific OOB signal called COMSAS.

Table 1 defines the timing specifications for OOB signals.

Table 1. OOB signal timing specifications

Parameter	Nominal	Minimum	Maximum	Comments
COMSAS detector off threshold		175 ns	525 ns	Detector shall reject all bursts with spacings outside of this range.
COMSAS detector on threshold	320 ns	304 ns	336 ns	Detector shall detect all burst with spacings meeting this period.
COMSAS transmit spacing	320 ns	310,4 ns	329,6 ns	As measured from receiver thresholds of last and first differential crosspoints of the burst.
<u>Unit interval during OOB (UI(OOB))</u>	666,667 ps	666,600 ps	666,734 ps	<u>The time basis for ALIGN bursts and idle times used to create OOB signals.</u> <u>Based on 1,5 Gbps clock tolerance.</u>
Hot-Plug Timeout	100 ms	10 ms	500 ms	The hot plug retry timer. <u>How often a device should resend COMINIT to detect if a device has been attached.</u>
Rate change detect (RCD)	480 UI(OOB)	0 ns	320 ns	Transmitter shall change speeds within this time. The time the transmitter shall transmit idle (no transitions) between rates during speed negotiation. Used by the transmitter and receiver to calculate the width of the Speed Negotiation Window.
Speed Negotiation Window Speed Negotiation Transmit Time (SNTT)	109,232 µs	109,220 µs	109,24 µs	The speed negotiation window is- <u>The time during which ALIGN or ALIGN(1) is transmitted at each rate during negotiation. Derived from:</u> <u>UI(OOB) * 4_096 * 40.</u> <u>Used by the transmitter and receiver to calculate the width of the Speed Negotiation Window.</u>
Speed Negotiation Lock Time (SNLT)	104,475 µs			Receiver's time during the speed negotiation window to detect ALIGN or ALIGN(1) and reply with ALIGN(1). Derived from: SNTT - (3 * (COMSAS Negation Time))
COMSAS detect timeout period	3,5 µs <u>10,88 µs</u>			Timer- <u>The time a receiver shall allow to detect COMSAS after sending COMSAS.</u>
ALIGN Detect Timeout			880 µs	<u>The time during SATA speed negotiation that an initiator phy shall allow for an ALIGN to be received after detecting the COMWAKE negation</u>

				<u>period.</u> <u>Allows over 32 768 G1</u> <u>words.</u>
--	--	--	--	---

Table 1 describes the OOB signal transmitter requirements for the burst time and idle time that comprise each OOB signal. When the link is synchronized, before sending an OOB signal, the transmitter shall transmit idle (drive with no transitions) for 1 ms to ensure the receiver drops link synchronization and detects the OOB signal.

Table 2. OOB signal transmitter requirements

<u>Signal</u>	<u>Burst time</u>	<u>Idle time</u>
<u>COMWAKE</u>	<u>160 UI(OOB)</u>	<u>160 UI(OOB)</u>
<u>COMINIT/RESET</u>	<u>160 UI(OOB)</u>	<u>480 UI(OOB)</u>
<u>COMSAS</u>	<u>160 UI(OOB)</u>	<u>1 440 UI(OOB)</u>

Table 3 describes the OOB signal receiver requirements for the idle times that distinguish each OOB signal.

Table 3. OOB signal receiver requirements

<u>Signal</u>	<u>Idle time</u>		
	<u>may detect</u>	<u>shall detect</u>	<u>shall not detect</u>
<u>COMWAKE</u>	<u>55 ns to 175 ns</u>	<u>101,3 ns to 112 ns</u>	<u>< 55 ns or > 175 ns</u>
<u>COMINIT/COMRESET</u>	<u>175 ns to 525 ns</u>	<u>304 ns to 336 ns</u>	<u>< 175 ns or > 525 ns</u>
<u>COMSAS</u>	<u>525 ns to 1 575 ns</u>	<u>911,7 ns to 1008 ns</u>	<u>< 525 ns or > 1 575 ns</u>

[Editor's note: SATA table 11 "transmit spacing" times correspond to COMWAKE and COMINIT minimum, nominal, and maximum idle times in the next table.]

[Editor's note: Receiver burst times: SATA table 11 does not specify any receiver burst times. The sample circuit in SATA figure 31 might ignore bursts longer than 114,4 ns and definitely ignores bursts longer than 171,6 ns. Since the sample circuit is not normative, this behavior cannot be depended upon.]

[Editor's note: Receiver idle times: SATA "detector on threshold" corresponds to the "shall detect" region and "detector off threshold" corresponds to the "shall not detect" region. A "may detect region is assumed to sit between the two.]

[Editor's note: Systems must guarantee that the transmitters meeting both the minimum and maximum transmit times result in a signal at the receiver in the shall detect region.]

[Editor's note: SATA numbers are inconsistent. The COMINIT numbers are all 3x the COMWAKE numbers, except the minimum detect time of 55*3=165 not 175 (58.3*3 is 175). Left the 55 as is since that's what SATA 1.0 specifies.]

Figure 1 illustrates the timing required for the generation and detection of all the SAS OOB signals. The signals labeled ComXxxDet and ComXxxSent (where Xxx is Wake, Init, or Sas), are inputs to the SAS phy Initialization layer state machine. ComXxxDet is an output from the OOB detection circuitry that indicates the detection of an incoming COMINIT, COMWAKE, or COMSAS sequencesignal. ComXxxSent is an output from the OOB generation circuitry that indicates that a burst of six COMWAKE, COMINIT, or COMSAS sequencesignal has been output onto the link.

COMINIT as defined in the SATA specification is the same as COMRESET, and they are used in this standard interchangeably.

The ALIGNs used in OOB sequencesignallingsignals are not required to be at generation 1 (G1) rates (1,5 Gbps), as this rate might not be supported in future generations of SAS devices. The ALIGNs are only required to generate an envelope for the detection circuitry, as required for any signaling that may be AC coupled. If G2 ALIGNs are used, the number of ALIGNs doubles compared with G1 ALIGNs. Likewise tTheThe idle time for

~~between the burst of ALIGNs. So, the these bursts and spaces IDLEs is important used to distinguish between the OOB signals, not the absolute number of ALIGNs for the proper detection of the signal in each burst.~~

A SAS transmitter should ~~send OOB sequences with use~~ ALIGNs at the G1 rate to create the burst portion of the OOB signal, but may ~~send them use~~ ALIGNs at its slowest supported rate if it does not support the G1 rate and shall not send them at a rate faster than its slowest supported rates.

A SAS receiver shall detect OOB ~~sequences signals~~ comprised of ALIGNs transmitted at any rate up to its highest supported rate. This includes rates below its lowest supported rates.

~~To send an OOB signal, a transmitter shall send the OOBALIGN Bburst preceded and followed by IDLE idle time signal at least six times. The last burst of IDLEs idle time is serves as the OOB signal negation time and will be different than the inter ALIGN IDLE time is longer than the normal idle time for that OOB signal. A receiver shall detect an OOB signal after receiving at four consecutive bursts of transitions and idle times. It shall not detect the same signal again until it has detected a different OOB signal (e.g., if the idle time changes) or has detected lack of transitions for the OOB signal negation time.~~

~~Table 3 shows the OOB signal negation times IDLE for each OOB signal, defining the negation period. and after detecting an OOB signal before entering the speed negotiation section.~~

Table 3. OOB signal negation times

Parameter	Time
<u>COMWAKE negation time</u>	<u>175 ns</u>
<u>COMINIT/COMRESET negation time</u>	<u>525 ns</u>
<u>COMSAS negation time</u>	<u>1 575 ns</u>

[Editor's note: this time is used between the last OOB signal and the start of speed negotiation.]

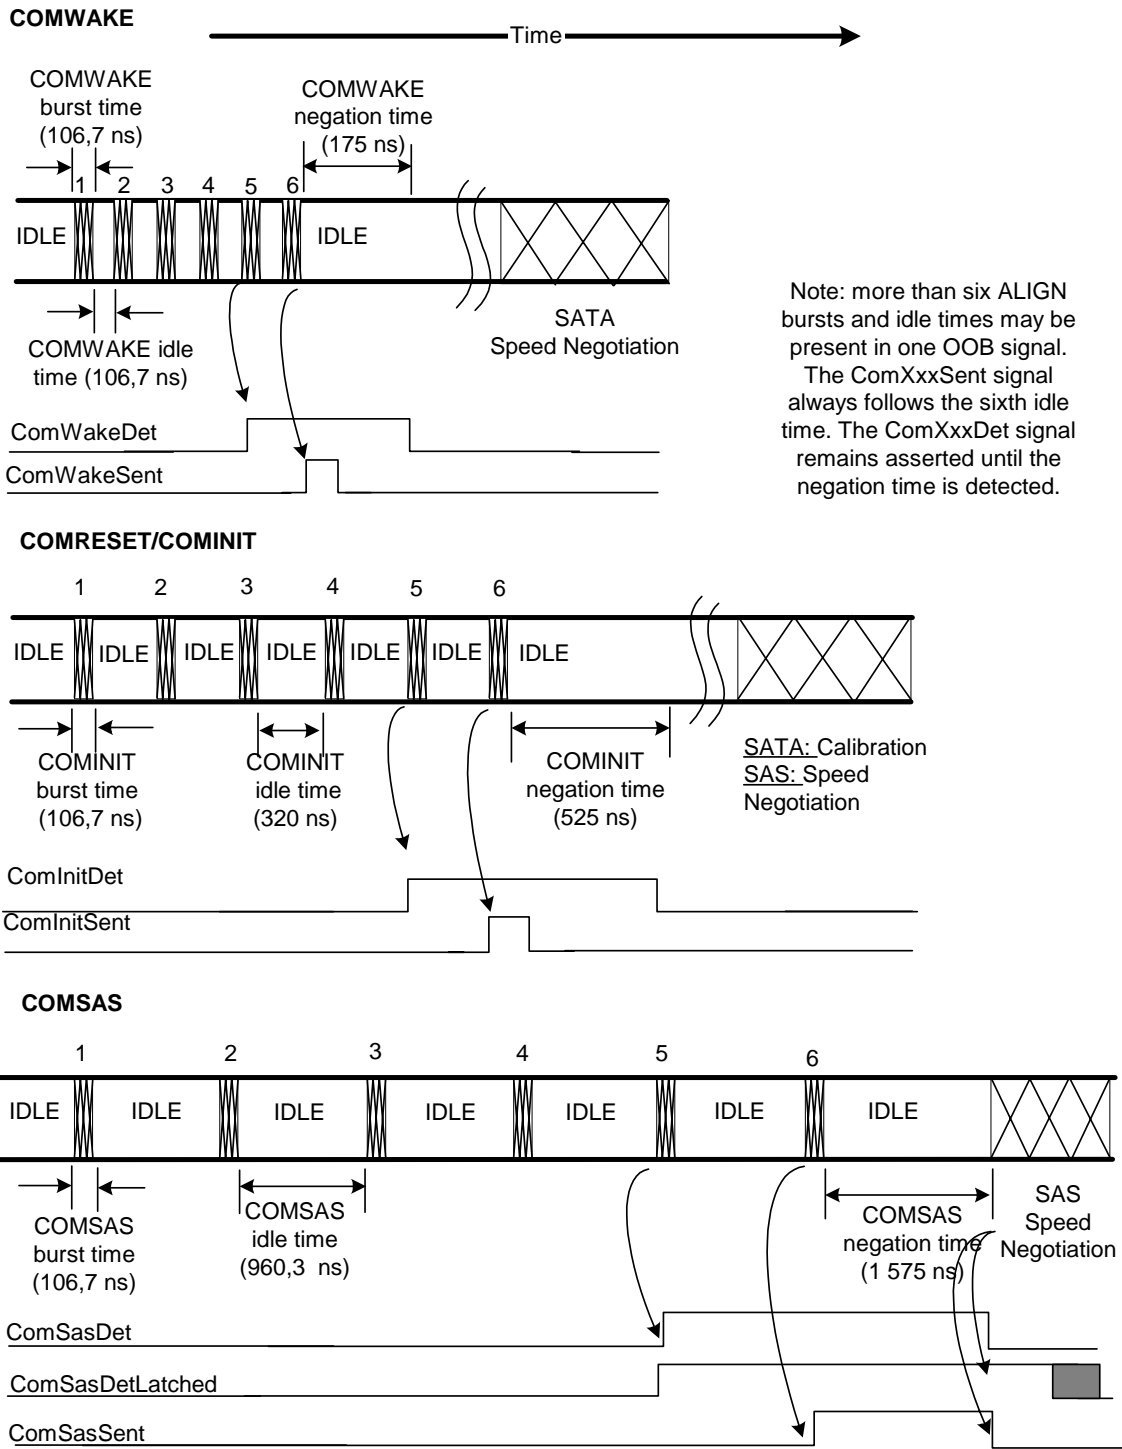


Figure 1. Phy reset sequence OOB signal generation transmission and detection

Expanders shall not pass through OOB signals. An expander device shall run the link reset sequence independently on each physical link (initiator phy to expander phy, expander phy to expander phy, or expander phy to target phy).

The phy reset sequence shall only affect the phy, not the port or device containing the phy or other phys in the same port or device.

6.8 Phy reset sequences

6.8.x Phy reset sequence overview

After a HARD_RESET, a device should start the phy reset sequence within 250 ms.

6.8.1 SATA phy reset sequence

Figure 34 depicts the link reset sequence between a SATA initiator device and target device. See SATA for detailed requirements.




Figure 34. SATA phy reset sequence

6.8.2 SAS to SATA phy reset sequence

SAS phys shall send COMSAS during the SATA calibration phase to identify themselves as SAS phys rather than SATA phys.

Figure 35 shows a reset sequence between a SAS phy and a SATA phy (i.e., between an expander device and a SATA target device, or between a SAS initiator device and a SATA target device). The two possible cases are presented. The first case is that a legacy SATA phy misinterprets the COMSAS to be a COMRESET and responds with a COMINIT. The second case is that the SATA phy ignores the COMSAS and provides no response within the COMSAS detect timeout of 3,5 μs. The SAS phy state machine treats these two cases the same, and determines that a SATA phy is attached. The normal SATA reset sequence shall be used thereafter.




Figure 35. SAS device to SATA device phy reset sequence

6.8.3 SAS to SAS phy reset sequence

A SAS device shall distinguish between COMINIT and COMSAS and continue with a SAS speed negotiation sequence rather than the SAS reset sequence.

Figure 36 illustrates several different ~~power-up~~ sequences based on whether ~~one a~~ SAS phy powers up before, after, or at the same time as the SAS phy to which it is attached.




Figure 36. SAS-to-SAS phy reset sequence

The SAS speed negotiation sequence is a peer-to-peer negotiation technique that does not assume initiator device and target device roles like the SATA speed negotiation sequence. The sequence consists of a set of Speed Negotiation Windows for each link rate, starting with 1,5 Gbps, then 3,0 Gbps, then the next rate. Each Speed Negotiation Window lasts a Rate Change Delay plus a Speed Negotiation Transfer Time. The sequence may be shortened if a device does not support faster rates.

First, each device shall transmit IDLEidle (drive with no transitions) stop transmitting for a Rate Change Delay. Then, each device shall start transmitting ALIGN1s at that link rate if it supports the link rate or shall transmit nothing IDLEidle for the entire Speed Negotiation Window if it does not support the link rate.

In order to avoid mis-communication, where one device receives ALIGN1s and then starts transmitting them, but it is too late for the original device to recognize, a Speed Negotiation Lock Time is specified. This provides a safeguard against this occurrence. Since the Speed Negotiation Window is only as accurate as the detection of the COMSAS negation, there will always be some discrepancy on when SAS devices start and finish the speed negotiation process. Dword lock and the ALIGN1 handshake completion must be obtained sufficiently before the

end of the negotiation window to avoid an incomplete handshake. It is possible that the device being communicated with may have already transitioned to the next window.

If the device supports the link rate, it shall attempt to synchronize on an incoming series of dwords at the same rate for the Speed Negotiation Lock Time. The incoming dwords may be ALIGNs or ALIGN(1)s. If the device achieves dword synchronization within the Speed Negotiation Lock Time, it shall change from transmitting ALIGNs to transmitting ALIGN(1)s for the remainder of the Speed Negotiation Transmit Time (i.e., the remainder of the Speed Negotiation Window). If the device does not achieve dword synchronization within the Speed Negotiation Lock Time, it shall continue transmitting ALIGNs for the remainder of the Speed Negotiation Transmit Time. Figure xx illustrates the Speed Negotiation Window.

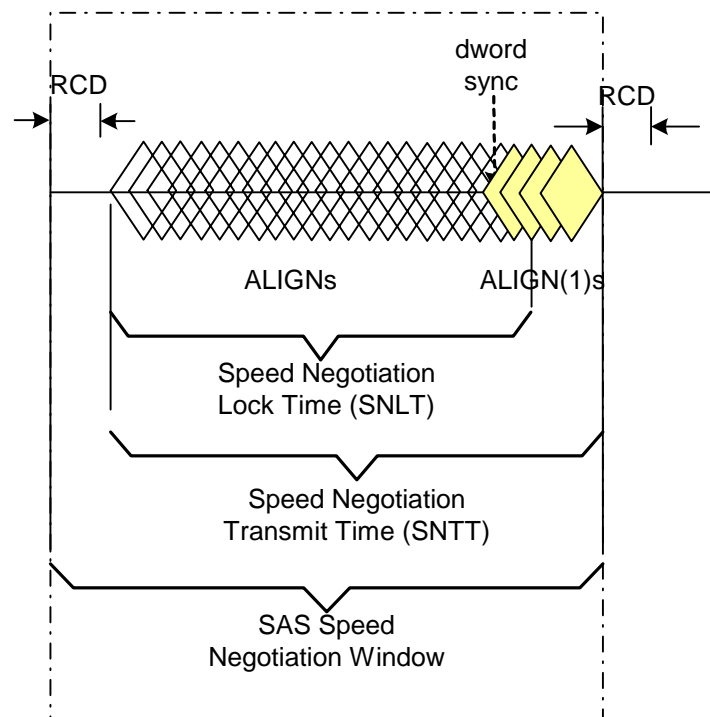


Figure 2. Speed Negotiation Window

At the end of the Speed Negotiation Transmit Time, if a device is both transmitting and receiving ALIGN(1)s, it shall consider that link rate valid for later use. It shall then proceed to the next ~~time~~ Speed Negotiation window.

As a Window. A device shall participate in all Speed Negotiation Windows up to its highest supported link rate plus one, or until it has detected a valid link rate in the most recent Speed Negotiation Window and runs a Speed Negotiation Window that does not detect a valid link rate. At this point, if it has detected a valid link rate in the previous Speed Negotiation Window, it shall enter the final Speed Negotiation Window using the highest previously successful link rate.

Figure xx shows a speed negotiation between a SAS device A that supports only the G1 link rate attached to a SAS device B that only supports the G1 link rate. They both run the G1 Speed Negotiation Window (valid) and the G2 Speed Negotiation Window (invalid), then select G1 for the final Speed Negotiation Window used to establish the negotiated rate.

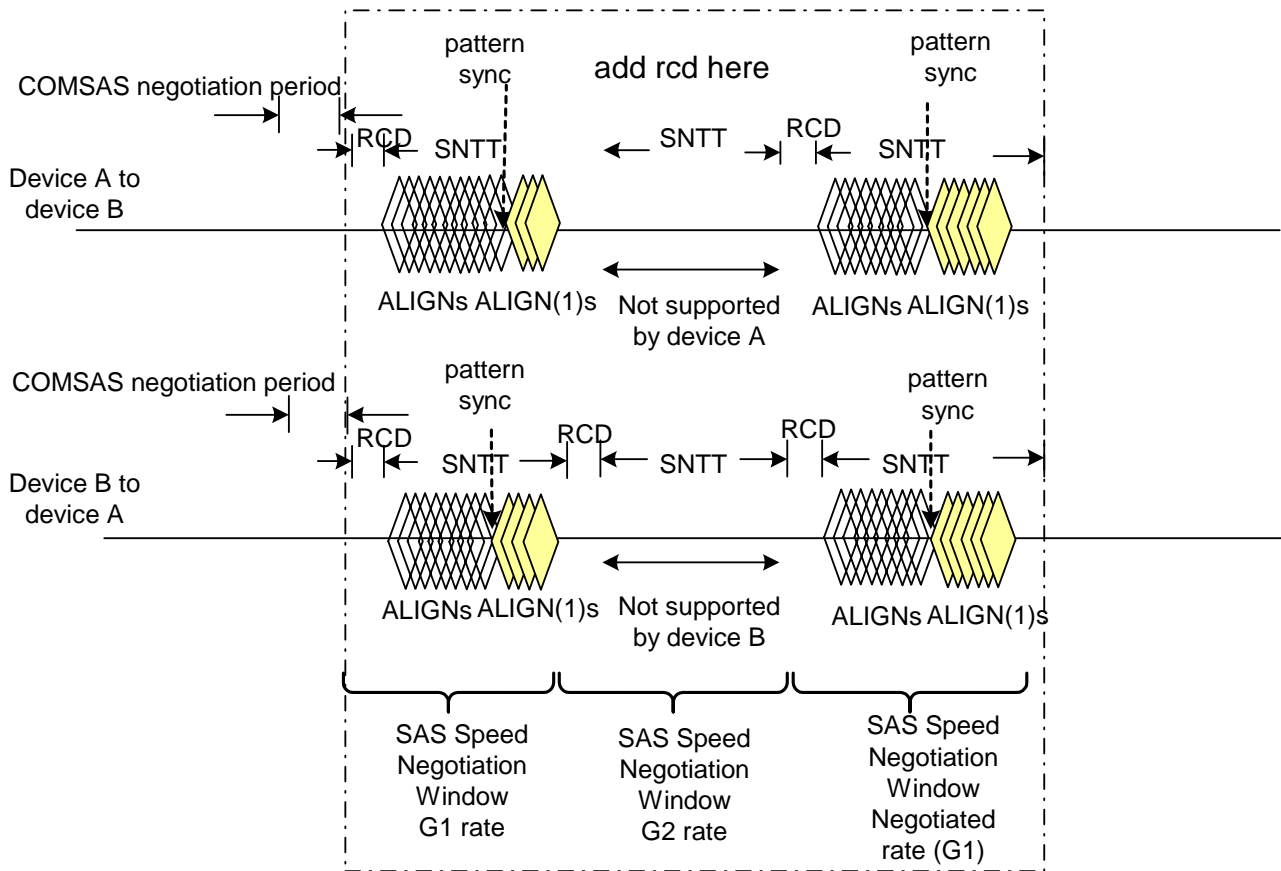


Figure xx. SAS speed negotiation sequence (G1 only with G1 only)

Figure 37 shows a speed negotiation between a SAS phy-device A that supports speeds-G1 thru G3 link rates attached to a SAS phy-device B that only supports speeds-G1 and G2 link rates. They both run the G1 Speed Negotiation Window (valid), the G2 Speed Negotiation Window (valid), and the G3 Speed Negotiation Window (supported by device A but not by device B, so invalid), then select G2 for the final Speed Negotiation Window used to establish the negotiated rate.

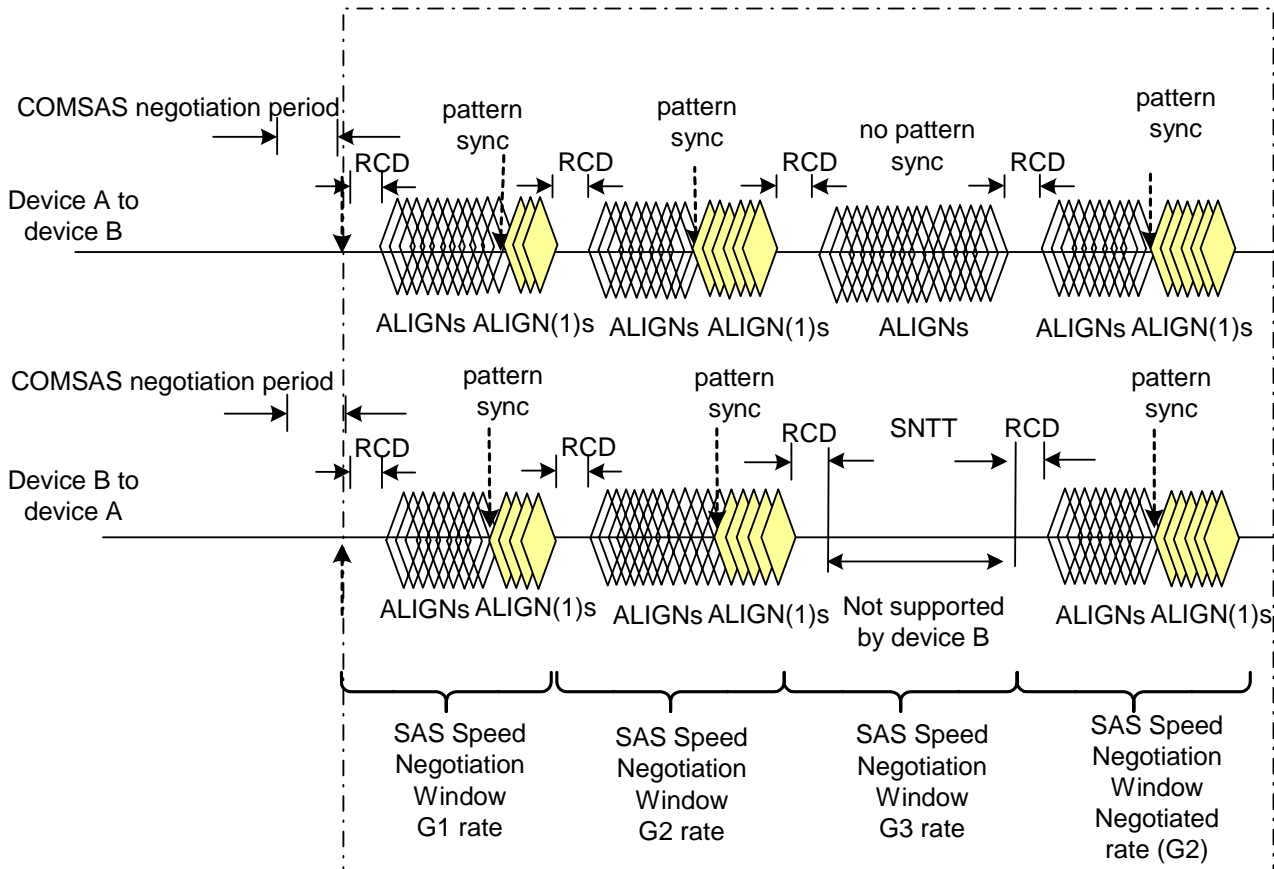


Figure 37. SAS speed negotiation sequence (example 4-G1, G2, G3 with G1, G2)

Figure 38 shows speed negotiation between a SAS device A that supports G1 thru G3 link rates and a SAS device B that only supports the G2 link rate. They both run the G1 Speed Negotiation Window (supported by device A but not by device B, so invalid), the G2 Speed Negotiation Window (valid), and the G3 Speed Negotiation Window (supported by device A but not by device B, so invalid), then select G2 for final Speed Negotiation Window to establish the negotiated rate.

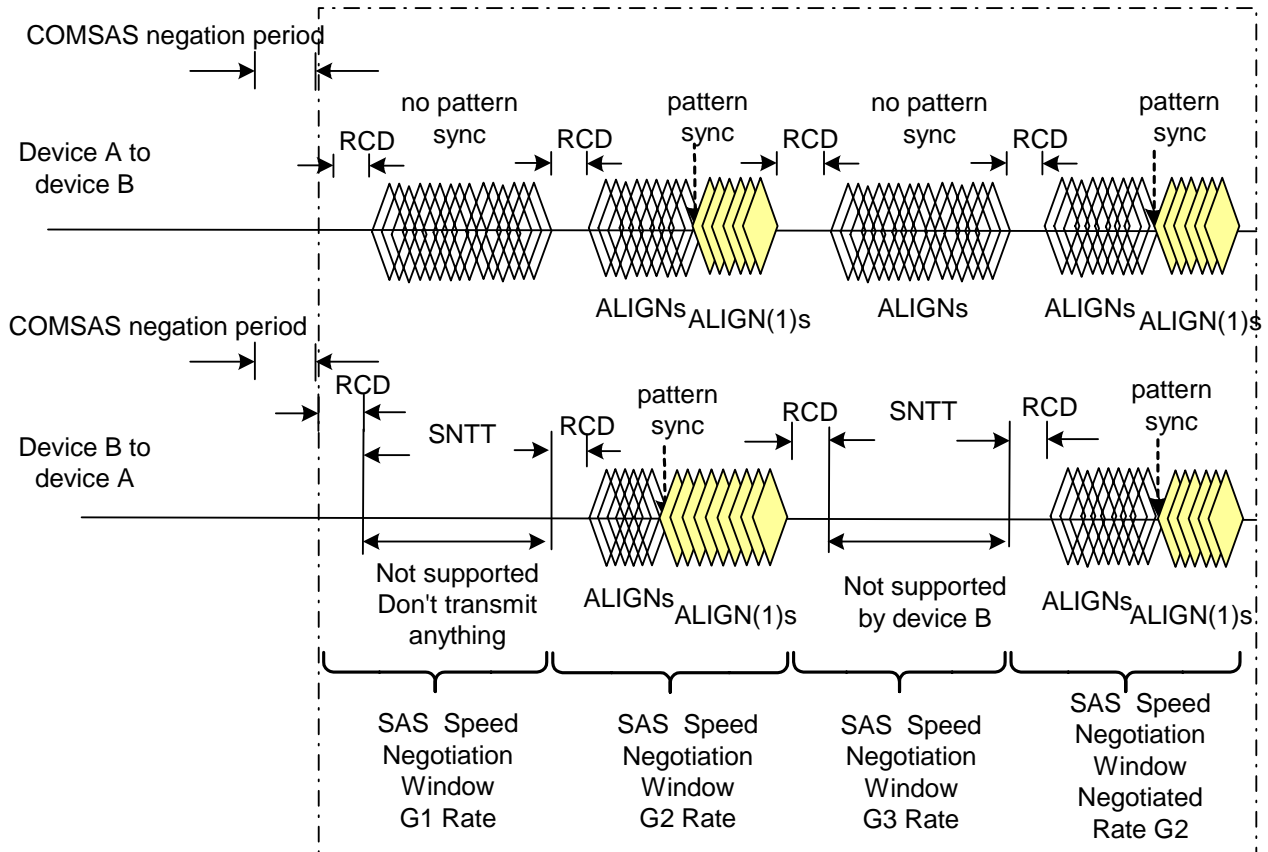


Figure 38. SAS speed negotiation sequence (example-2G1, G2, G3 with G2 only)

Figure 39 shows the same speed negotiation sequence as in Figure 37, except for some reason, one of the SAS device's PLLs does not obtain ~~pattern_dword~~ synchronization and detect ALIGNs during the final speed negotiation window. If this occurs, the handshake is not complete and the OOB sequence shall be retried starting with COMINIT, forcing the ~~link-phy~~ to retry the whole reset sequence.

Any time a SAS phy fails speed negotiation, it shall wait the Hot Plug Time-out Delay before attempting a retry.




Figure 39. SAS speed negotiation sequence (example 3G1,G2,G3 with G1,G2 and failure)

6.9 Phy reset sequence after signal cable insertion

Since SATA and SAS signal cable connectors do not include power lines, it is not possible to detect the physical insertion of the signal cable connector into a receptacle. In order to get around this limitation, SAS phys should periodically transmit a COMINIT sequence if they have not detected a COMINIT sequence within a Hot-Plug Timeout period.

Figure 40 illustrates how two SAS phys complete the phy reset sequence if the phys are not attached at power on. In this scenario, SAS phy B is connected to SAS phy A via a cable some time before SAS phy B's second Hot-Plug timeout occurs. SAS phy B's OOB detection circuitry detects a COMINIT after the connection is made, and therefore SAS phy B transmits the COMSAS sequence, since it has both sent and detected a COMINIT sequence. Upon detection of COMSAS, SAS phy A transmits its own COMSAS sequence bypassing the sent/received COMINIT requirement. The normal SAS speed negotiation process follows.




Figure 40. Hot-plug and the phy reset sequence

...

9.4.4.6 REPORT PHY function

The REPORT PHY function returns information about the specified phy. This function may implemented by any type of device.

The REPORT PHY function returns information about the specified phy. This function may implemented by any type of device.

Table 4 defines the request format.

Table 4. REPORT PHY request

Byte	7	6	5	4	3	2	1	0
0	FUNCTION (10h)							
1	Reserved							
2								
3	PHY IDENTIFIER							

The PHY IDENTIFIER field indicates the phy ([see xx](#)) for which information shall be reported.

Table 5 defines the response format.

Table 5. REPORT PHY response

Byte	7	6	5	4	3	2	1	0
0	FUNCTION RESULT							
1	REPORT PHY RESULT							

2	Rsvd	Rsvd	Rsvd	Rsvd	PHY RESET PROBLEM	SAS PROTOCOL	SATA PENDING	PHYSICAL LINK ACTIVE
3	MAXIMUM PHYSICAL LINK RATE				HARDWARE MAXIMUM PHYSICAL LINK RATE			
4	MINIMUM PHYSICAL LINK RATE				HARDWARE MINIMUM PHYSICAL LINK RATE			
5	Reserved				CURRENT PHYSICAL LINK RATE			
6	ATTACHED PHY IDENTIFIER							
7	ATTACHED DEVICE TYPE	ATTACHED STP INITIATOR	ATTACHED STP TARGET	ATTACHED SSP INITIATOR	ATTACHED SSP TARGET	ATTACHED SMP INITIATOR	ATTACHED SMP TARGET	
8	(MSB)	ATTACHED DEVICE NAME						(LSB)
15								
16								
31	Reserved							

The FUNCTION RESULT field is defined in table xx.

The REPORT PHY RESULT field is defined in Table 6.

Table 6. Report PHY result

REPORT PHY RESULT	Description
00h	Phy exists; rest of data is valid.
01h	Phy does not exist; rest of data is invalid.
All others	Reserved.

The PHYSICAL LINK ACTIVE bit is set to one if the phy has completed the link reset sequence and is active.

The SATA PENDING bit is set to one if the physical link is inactive because the beginning of the link reset sequence indicated SATA capability, and the expander device did not complete the link reset sequence.

The SAS PROTOCOL bit is set to one if the link reset sequence indicated SAS capability and is set to zero if the link reset sequence indicated SATA capability. If the SAS PROTOCOL bit is one, the ATTACHED fields indicate the information received during the link initialization sequence. If the SAS PROTOCOL is zero, the attached fields are all set to zero.

The PHY RESET PROBLEM bit is set to one if the phy state machine obtained dword synchronization for at least one rate during speed negotiation (either SAS or SATA), but the speed negotiation sequence failed. The phy may have retried the negotiation and succeeded, but the bit remains set to one. The bit remains set to one until the phy is reset without any failures.

The HARDWARE MAXIMUM PHYSICAL LINK RATE field indicates the maximum physical link rate supported by the phy.

The MAXIMUM PHYSICAL LINK RATE field indicates the maximum physical link rate set by the PHY CONTROL function.

The HARDWARE MINIMUM PHYSICAL LINK RATE field indicates the minimum physical link rate supported by the phy.

The MINIMUM PHYSICAL LINK RATE field indicates the minimum physical link rate set by the PHY CONTROL function.

The CURRENT PHYSICAL LINK RATE field indicates the physical link rate negotiated during the link reset sequence. The current physical link rate may be outside the minimum physical link rate and maximum physical link rate if they have been changed since the link reset sequence.

The ATTACHED PHY IDENTIFIER field indicates the PHY IDENTIFIER value received during the link reset sequence.

The ATTACHED DEVICE TYPE field indicates the DEVICE TYPE value received during the link reset sequence.

The ATTACHED STP INITIATOR bit indicates the SSP INITIATOR value received during the link reset sequence.

The ATTACHED STP TARGET bit indicates the SSP TARGET value received during the link reset sequence.

The ATTACHED SSP INITIATOR bit indicates the SSP INITIATOR value received during the link reset sequence.

The ATTACHED SSP TARGET bit indicates the SSP TARGET value received during the link reset sequence.

The ATTACHED SMP INITIATOR bit indicates the SSP INITIATOR value received during the link reset sequence.

The ATTACHED SMP TARGET bit indicates the SSP TARGET value received during the link reset sequence.

If SAS PROTOCOL is set to one, the ATTACHED DEVICE NAME field indicates the device name received during the link reset sequence. If SAS PROTOCOL is set to zero, this field contains the unique device name assigned by the expander device.

REPORT PHY ERROR LOG function

The REPORT PHY ERROR LOG returns error logging information about the specified phy. This function may be implemented by any type of device.

Table 7 defines the request format.

Table 7. REPORT PHY ERROR LOG request

Byte	7	6	5	4	3	2	1	0
0	FUNCTION (11h)							
1	Reserved							
2	Reserved							
3	PHY IDENTIFIER							

The PHY IDENTIFIER field indicates the phy ([see table XX](#)) for which information shall be reported.

Table 8 defines the response format.

Table 8. REPORT PHY ERROR LOG response

Byte	7	6	5	4	3	2	1	0
0	FUNCTION RESULT							
1	REPORT PHY ERROR LOG RESULT							
2	Reserved							
3	Reserved							
4	(MSB)	INVALID CHARACTER COUNT						(LSB)
7	Reserved							
8	(MSB)	DISPARITY ERROR COUNT						(LSB)
11	Reserved							
12	(MSB)	LOSS OF BIT SYNCHRONIZATION COUNT						(LSB)
15	Reserved							
16	PHY RESET PROBLEM COUNT							
19	Reserved							
20	Reserved							
31	Reserved							

The FUNCTION RESULT field is defined in table XX.

The REPORT PHY ERROR LOG RESULT field is defined in Table 9.

Table 9. Report phy error log result

REPORT PHY ERROR LOG RESULT	Description
00h	Phy exists; rest of data is valid.
01h	Phy does not exist; rest of data is invalid.
All others	Reserved.

The INVALID CHARACTER COUNT field indicates the number of invalid 8b10b characters that have been received.

The DISPARITY ERROR COUNT field indicates the number of disparity errors that have been detected.

The LOSS OF BIT SYNCHRONIZATION COUNT field indicates the number of times bit synchronization has been lost.

| The PHY RESET COUNT field indicates the number of times the phy reset sequence has failed.