

To: T10 Technical Committee
From: Bill Lye, PMC-Sierra (lye@pmc-sierra.com)
Yuriy Greshishchev, PMC-Sierra (greshish@pmc-sierra.com)
Date: 4 May 2004
Subject: T10/04-128r1 SAS-1.1 OOB Signal Rate @ 1,5G only

Revision History

Revision 0 (26 April 2004) First Revision
[Revision 1 \(4 May 2004\) Editorial changes](#)

Related Documents

sas1r04 – Serial Attached SCSI 1.1 Revision 4

Overview

As currently written, the SAS 1.1 specification states (Sections 5.3.4 and 6.5) that SAS initiator and target PHY's transmit OOB signals at the lowest supported transfer rate (i.e. 1,5 Gbps for a G1 or G1/G2 Transmitter or 3,0 Gbps for a G2-only Transmitter). This assumes that the squelch detection circuitry the RX OOB circuit can always detect the peak of any 1,5 Gbps, 3,0 Gbps, or higher rate (6,0 Gbps) transmitter. However, due to finite bandwidth effects, detecting the peak of a 3,0 Gbps or 6,0 Gbps OOB signal has a significant cost in complexity, area and/or power, which may be disadvantageous for high port count devices.

As a result, the following changes are proposed in sections 5.3.4 and 6.5 to ensure that the OOB burst signal rate remains at G1 (1,5 Gbps) rates.

Revised text follows:

[Section 5.3.4 \(pages 88-89\)](#)

The OOB sequence shall be performed at signal voltage levels corresponding to the lowest supported transfer rate. Expander phys supporting being attached to SATA devices shall use SATA 1.0 signal levels (see ATA/ATAPI-7 V3) during the first OOB sequence after a power on or hard reset ~~if the 1,5 Gbps transfer rate is supported~~. As soon as COMSAS has been exchanged, the expander phy shall increase its transmit levels to the SAS voltage levels specified in table 27. If a COMINIT is not received within a hot-plug timeout at SATA 1.0 signal levels, the expander phy shall increase its transmit levels to the SAS voltage levels and perform the OOB sequence again. If no COMINIT is received within a hot-plug timeout of the second OOB sequence the expander phy shall initiate another OOB sequence using SATA 1.0 signal levels. The expander phy shall continue alternating between sending COMINIT at SATA 1.0 signal levels and SAS signal levels until a COMINIT is received.

If the OOB sequence is completed at the SAS voltage level and a SATA device is detected rather than a SAS target device, the expander phy shall switch to SATA 1.0 voltage levels and repeat the OOB sequence.

NOTE 9 - SAS initiator phys supporting being attached to SATA devices may use the same algorithm as expander phys.

SAS initiator phys and SAS target phys shall transmit OOB signals at generation 1 (G1) physical link rates (i.e., 1,5Gbps) ~~the lowest supported transfer rate~~ using SAS signal levels.

Section 6.5 (Transmitter – page 108)

The ALIGNs used in OOB signals ~~are not required to~~ shall be at generation 1 (G1) physical link rates (i.e., 1,5Gbps), ~~as although this data rate may not be supported in phys compliant with future generations of this standard. The ALIGNs are only required to generate an envelope for the detection circuitry, as required for any signaling that may be A.C. coupled. If G2 ALIGNs are used, the number of ALIGNs doubles compared with G1 ALIGNs.~~


~~A SAS transmitter should transmit ALIGNs at the G1 physical link rate to create the burst portion of the OOB signal, but may transmit ALIGNs at its lowest supported physical link rate if it does not support the G1 physical link rate and shall not transmit them at a physical link rate faster than its lowest supported physical link rate.~~

Section 6.5 (Receiver – page 110)

A SAS receiver shall detect OOB signals comprised of ALIGNs transmitted at generation 1 (G1) physical link rates (i.e., 1,5Gbps) ~~any rate up to its highest supported physical link rate~~. This includes physical link rates below its lowest supported physical link rate (e.g., a SAS receiver supporting only 3,0 Gbps needs to detect 1,5 Gbps based ALIGNs to interoperate with a SAS transmitter supporting both 1,5 Gbps and 3,0 Gbps).

Background

A general block diagram of a SATA/SAS device with embedded SERDES is shown below.



Bandwidth requirements for the signal path are well known from other protocols; in order to pass data through the signal path (including the channel, connectors, package, and any buffering) a minimum raw bandwidth of about 70% of the bit rate is required. In other words, a 3,0 Gbps system requires a signal path bandwidth of 2,1 GHz or higher.

The receiver is also only has to detect zero crossings of the incoming data, and is therefore insensitive to amplitude errors in the system caused by channel dispersion and impedance mismatches. The squelch detector, on the other hand, has to detect absolute amplitudes and is therefore sensitive to such amplitude errors, with the result that the squelch detector threshold actually ends up dictating system signal level requirements

For the system to operate properly, the squelch detector requires the overall system to have a bandwidth greater than what would otherwise be required for passing NRZ data.


For example, a 3,0 Gbps system with a bandwidth of 2,1 GHz will have a squelch detector amplitude detection error of as much as 30% due to signal frequency (i.e. pattern) variation. This must then be accounted for in the squelch detector circuit design and in the already-tight OOB threshold specification. For example, if the intended OOB signal level range is 2:1 (as it is currently – 240mV to 120mV), a 30% reduction on the high end reduces the spec to 1.4:1, almost a 70% reduction in design target.

As an aside, the requirement for increased bandwidth is actually implied by Note c on Table 27 of sas1r04 (Page 91), which implies that the system bandwidth (including the squelch detector) would have to be 4,5 GHz for a 3,0 Gbps link, or over 2 times as large as would otherwise be required for an NRZ-only system.

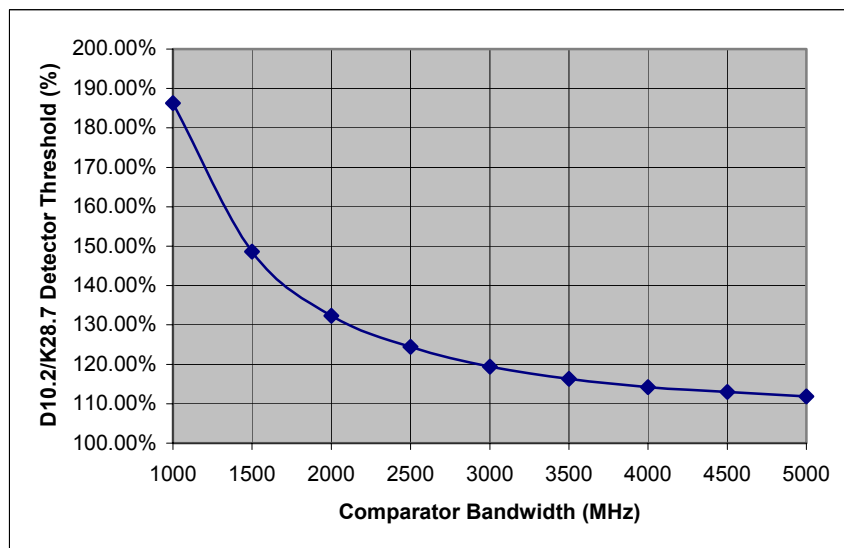
In a number of other standards (for example SONET and FibreChannel) the squelch detector on the receive side of the device is used only for Loss of Signal (LOS) detect. As a result, frequency limitations of the squelch detector can be somewhat mitigated by using long detection time constants, thereby increasing the likelihood that the squelch detector will see a run of low-frequency data in the incoming stream and decreasing the likelihood that the squelch detector will indicate a false LOS.

However, in SAS and SATA, the squelch detector is used for OOB, and not for LOS. OOB requires the detector to have a short detection time constant, which therefore reduces the likelihood that it will see low-frequency data in the incoming stream (increasing the likelihood of asserting a false OOB), which increases the importance of having a wide bandwidth in the squelch detector.

In CMOS, where diode-based full-wave rectifiers are uncommon, squelch detection is commonly done using peak detectors. In order to understand the frequency requirement of a peak detector, a simple linear model of a peak detector (shown below) is constructed. Comparators are modeled as saturating transconductance elements with a bandwidth-limiting load and nonlinear switches are modeled as clipping current sources driving into the peak detector time constant. In simulations, the comparator bandwidth, signal frequency, and signal amplitude are varied, while the peak detector time constant is held constant.



Simulation results from this model are shown below. Because we are concerned with the amplitude error seen when comparing data with high frequency components and low frequency components, we have run two simulations, one with a 1,5 GHz square wave (D10.2 at 3,0 Gbps) and the second with a 300 Mhz square wave (the lowest frequency component of 3,0 Gbps 8B/10B coded data, also known as K28.7 in XAUI). The ratio of detector threshold between the two patterns is shown as a function of comparator bandwidth.



As can be seen, in order to minimize the frequency-dependent effects, the comparator bandwidth must be maximized, and in fact even with a bandwidth of 4,5 GHz there is still an amplitude error of 13%. In practice, additional second order effects (not modeled in these simulations) tend to accumulate, increasing the frequency-dependent amplitude error even more.

This wider bandwidth requirement, combined with the relatively tight absolute specification on OOB thresholds, can easily cause the cost (in area, complexity, power, process choice, test time, or yield fallout) of the squelch detect circuit to exceed that of the signal path, and this increased cost may become prohibitive for devices with large port counts.

Our proposed modification would change the data rate of the ALIGNs in the OOB burst to be 1,5 Gbps only. This ensures that, in G2 and future generations, the cost of the squelch detect circuits does not increase beyond what it is in G1 designs.