

To: T10 Technical Committee

From: Bill Bissonette (Bill.Bissonette@intel.com)

Date: February 23, 2005

Subject: 05-077r1 SAS-1.1 SAS-SATA OOB Algorithm Summary

Revision history

Revision 0 (February 16, 2005) First revision

Revision 1 (February 24, 2005) Corrected hot plug time-out to 10ms – 500ms, added transient discussion.

Related documents

SAS-1.1r8

SATA 1.0a

SATA II Electrical spec

Overview

This paper is an expository on the realities of the SAS OOB algorithm described in SAS-1.1r8 section 5.3.7.5. In particular, it describes the scenarios in which a SATA drive will be exposed to SAS-level OOB bursts, and then describes the electrical realities of such. Finally, it calls for a response from the SATA disk drive suppliers to help the SAS Phy TWG determine whether this is indeed a problem that needs to be addressed with changes to the SAS-1.1 specification.

SAS-SATA OOB 'Toggling' Algorithm

SAS-1.1 working draft currently specifies an OOB algorithm as follows:

5.3.7.5 Transmitter device signal output levels for OOB signals

Transmitter devices 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 transmitter device shall increase its transmit levels to the SAS voltage levels specified in table 32 (see 5.3.7) and table 33 (see 5.3.7.3). If a COMINIT is not received within a hot-plug timeout at SATA 1.0 signal levels, the transmitter device 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 transmitter device shall initiate another OOB sequence using SATA 1.0 signal levels. The transmitter device 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 transmitter device shall switch to SATA 1.0 voltage levels and repeat the OOB sequence.

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

Figure 1: SAS-1.1r8 SAS-SATA OOB text

OOB scenarios for SATA drives in a SAS system

If a SATA drive is already installed in a system when the power is applied or when the system experiences a reset sequence, there is little probability that the SATA drive will be exposed to a SAS amplitude OOB sequence since it will respond to the first COMINIT (interpreted by the SATA drive as a COMRESET) with a COMINIT which will, in turn, cause the expander or initiator to maintain SATA OOB levels from there on out (until the next hard reset or power cycle). See Figure 2.

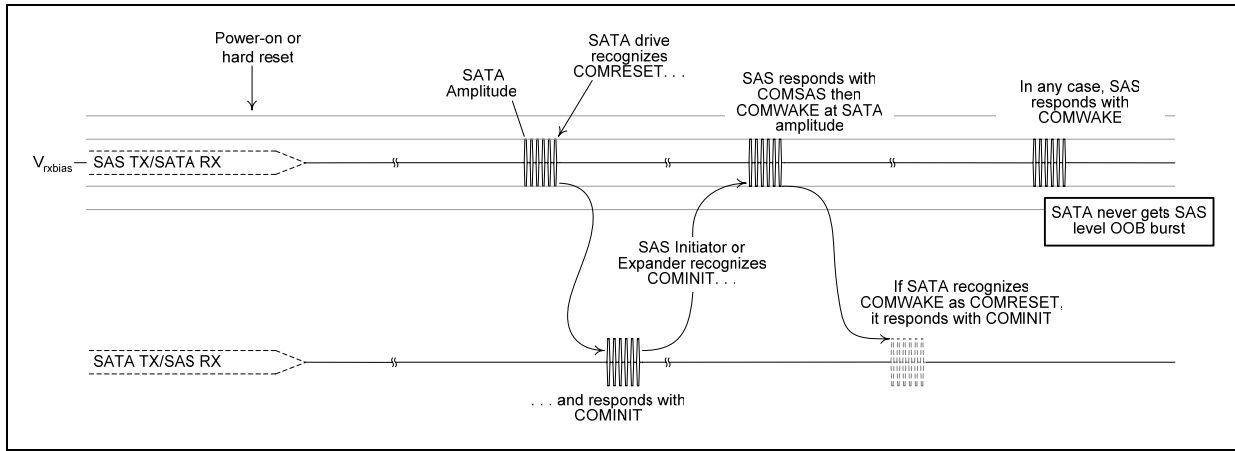


Figure 2: The "normal" scenario

After a power-on or hard reset event, if the SATA drive is not initialized before the SAS Expander or Initiator has issued the first COMINIT at SATA levels, the SATA drive will be exposed to a SAS-level COMINIT as shown in Figures 3 & 4. The SATA drive may or may not recognize the SAS-level COMINIT but the subsequent sequence is the same for either case. It is also possible in these cases that the RX termination is turned off for some time after the reset. A high impedance termination **could** cause an amplification of the received signal level. **Is this valid?**

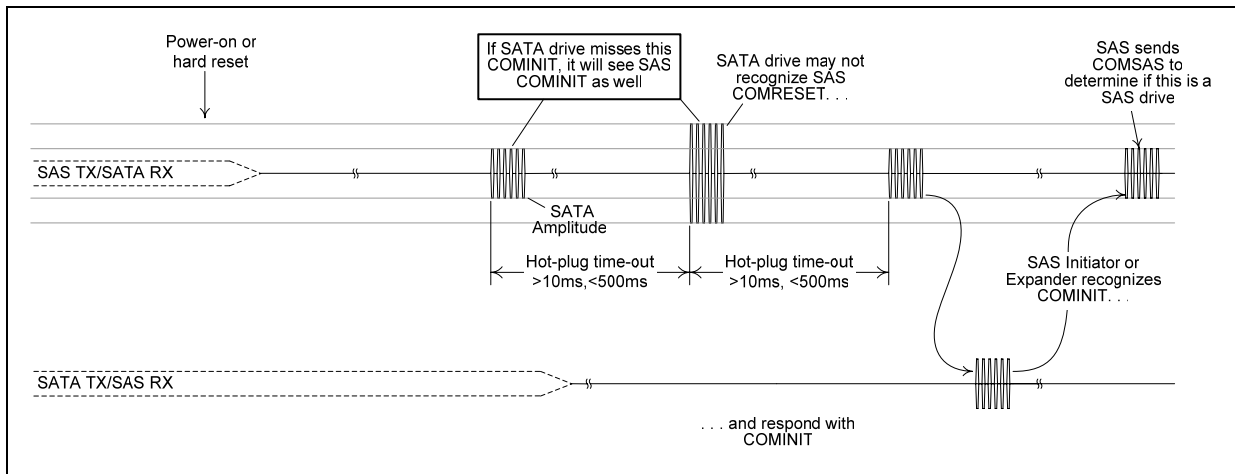


Figure 3: SATA drive wakes up too late, doesn't recognize SAS-level COMINIT

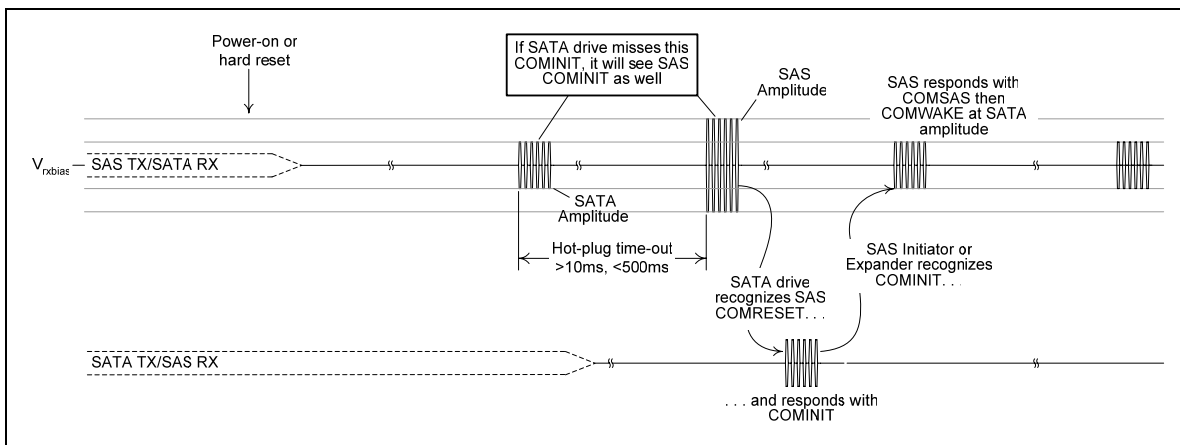


Figure 4: SATA drive recognizes SAS-level COMINIT

As shown in figure 5, a SATA drive that is 'hot-plugged' into a backplane has a significant likelihood (on the order of 50%) of exposure to a SAS-amplitude COMINIT.

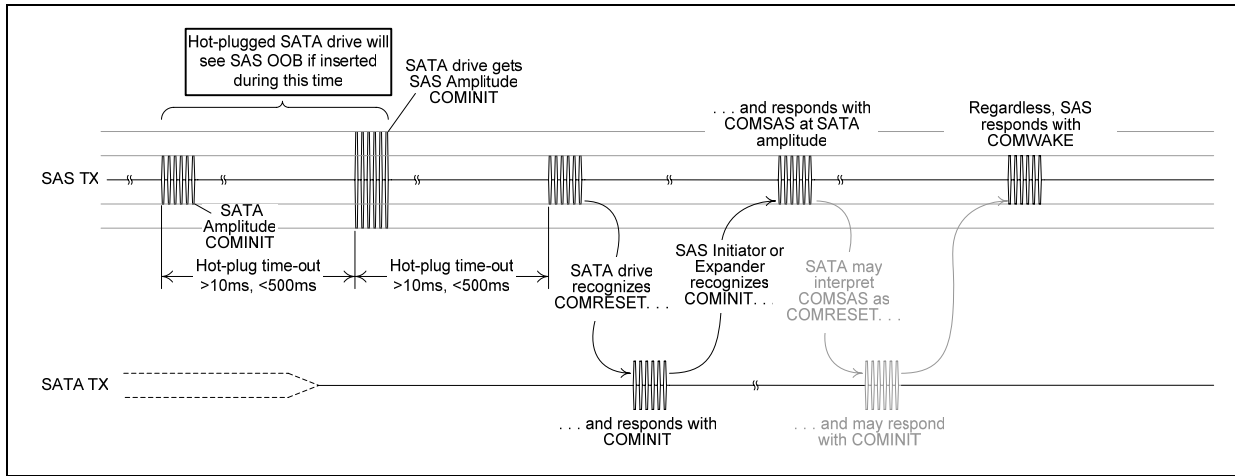


Figure 5: SATA drive hot-plugged

Voltage Considerations at the Receiver

Proposal T10 05-019 proposes that all OOB sequences from initiators and expanders that support SATA be transmitted at SATA amplitudes. There is concern from SAS drive suppliers that SAS devices may fail to detect SATA-level OOBs. And there are existing, shipping SAS-1.0 and SAS-1.1 initiator and expander devices that have implemented the OOB toggling algorithm and implementation of such a change to the OOB algorithm would cause design changes. The outstanding question has become: **Is there a risk that SAS level OOB bursts will damage a SATA drive or cause the drive to fail an OOB initialization sequence?**

The SATA II electrical specification states that the maximum differential amplitude at the receiver is 600mV maximum for Gen1i and 750mV maximum for Gen2i. These are operational specifications and the 'absolute maximum' amplitudes are not specified (*are they?*). It has been noted by SAS Phy TWG members that some SAS expanders or initiators DO transmit at nearly 1600mV.

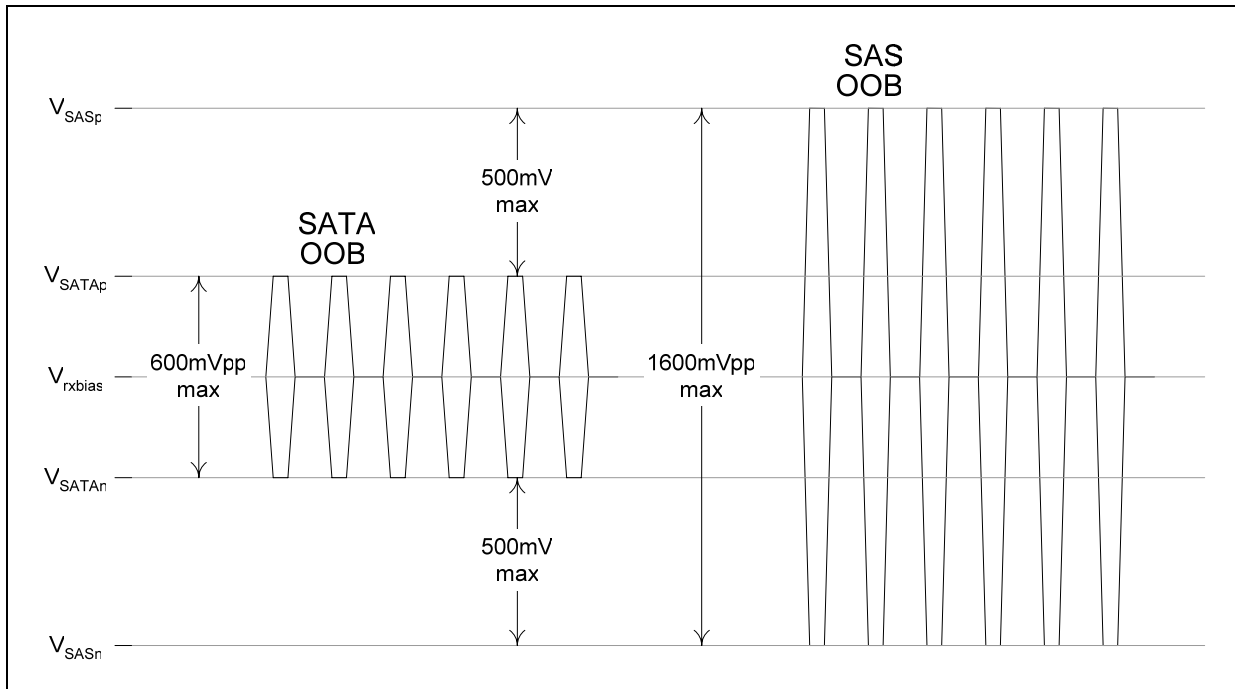


Figure 6: Differential Amplitude Case

Figure 6 illustrates worst-case SATA Gen1i and SAS level OOB bursts as seen by a SATA drive receiver. The individual voltages within the burst are not shown, but rather the burst envelopes for a COMINIT burst sequence. The differential amplitude beyond the maximum Gen1i specification is 1000mV. **Will this cause any damage or cause the SATA OOB detection circuitry to miss a subsequent (as little as 10ms later) SATA level OOB burst?**

Concern: The receiver input circuitry will saturate, debias or suffer transient nonfunctionality.

Discussion: SATA spec requires 3us max recovery time from “transient” event. SAS spec requires at least 10ms (hot plug time-out) before next COMINIT which is >3,000 times longer that required. Receiver has plenty of time to recover before next COMINIT. **Does this apply?**

SAS transmitter IS AC coupled and RX bias will settle well before the next OOB burst (when SAS TX launch amplitudes are adjusted within the 10 - 500ms hot-plug time-out).

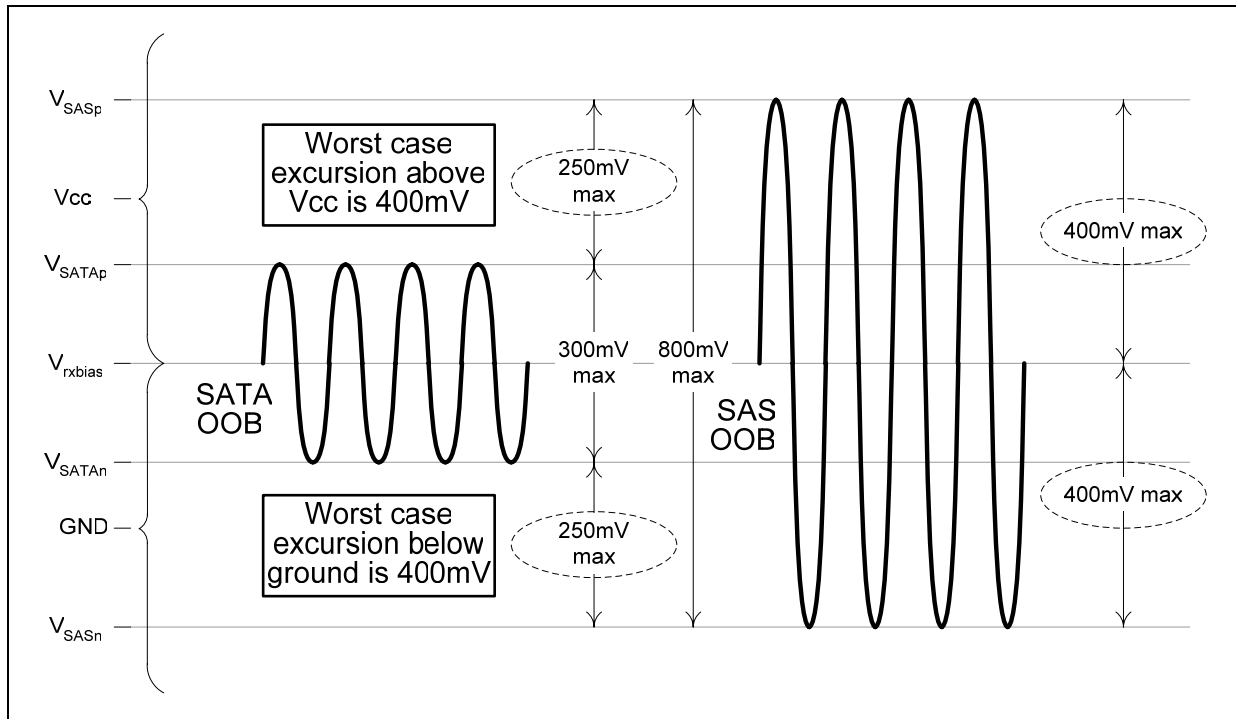


Figure 7: Single-ended Amplitude Case

Figure 7 illustrates the single-ended voltages as seen by either the RXn or RXp inputs on a SATA drive for the SATA and SAS level OOBs. The actual amplitude excursion beyond the maximum SATA Gen1i specification is 250mV maximum, with a 400mV maximum excursion above and below the RX bias voltage. Note that V_{RXbias} can be any voltage between GND and V_{cc} on the drive, depending on the design of the receiver circuitry and, in particular, the voltage the receiver termination resistors are tied to.

Concern: Excursions below ground or above the power supply rail may cause ESD diodes to conduct and, over time, this would damage the receiver. **Is this a realistic scenario?**

Discussion: V_{RXbias} can be anywhere between V_{cc} and GND. Excursions are no more than 400mV above power rail or 400mV below ground. ESD forward bias voltages are typically greater than 400mV, so this scenario is not likely.

SATA ‘Sequencing Transient Voltage’

Table 2 in section 6.2.1 of the SATA II electrical document specifies a +2V to -2V maximum sequencing transient voltage. This limits single-ended transients in a system to less than $\pm 2V$ but also requires attached SATA devices to withstand transients up to that limit and is stated as: “The common mode transient requirements defined in Table 2 were determined sufficient to limit stresses on the attached components under transient conditions. . . .” If a SATA receiver complies with the intent of this SATA requirement, does that insure that same receiver will be able to withstand the comparatively weaker 400mV_{S-E} SAS OOB without damage or operational failure? What are mechanisms in the receiver that insure compliance to this requirement (e.g. ESD structures)?

The explanatory text from the SATA II Electrical Specification is included in Figure 8 for handy reference.

From “6.2.2.1.8 AC Coupled Common Mode Voltage”

The SATA interface defined as Gen1i may be AC or DC coupled as shown in Figure 13. The SATA interfaces defined as Gen1x, Gen2i, Gen2x shall be AC-coupled. Figure 14 shows an example of a fully AC-coupled system.

Compliance points for SATA are defined at the connector. The AC coupled common mode voltage in Table 2 defines the open circuit DC voltage level of each single-ended signal at the IC side of the coupling capacitor in an AC coupled PHY and it shall be met during all possible power and electrical conditions of the PHY including power off and power ramping. Since the Gen1x, Gen2i, Gen2x specification defines only the signal characteristics as observable at the connector, this value is not applicable to those specifications. The common mode transient requirements defined in Table 2 were determined sufficient to limit stresses on the attached components under transient conditions which was the sole intent of the AC coupled common mode voltage requirement. Due to this, the following is true even for Gen1i where $V_{cm,ac\ coupled}$ applies: AC coupled common mode voltage levels outside the specified range may be used provided that the transient voltage requirements of Table 2 are met.

Figure 8: SATA II text for AC coupled transient voltage

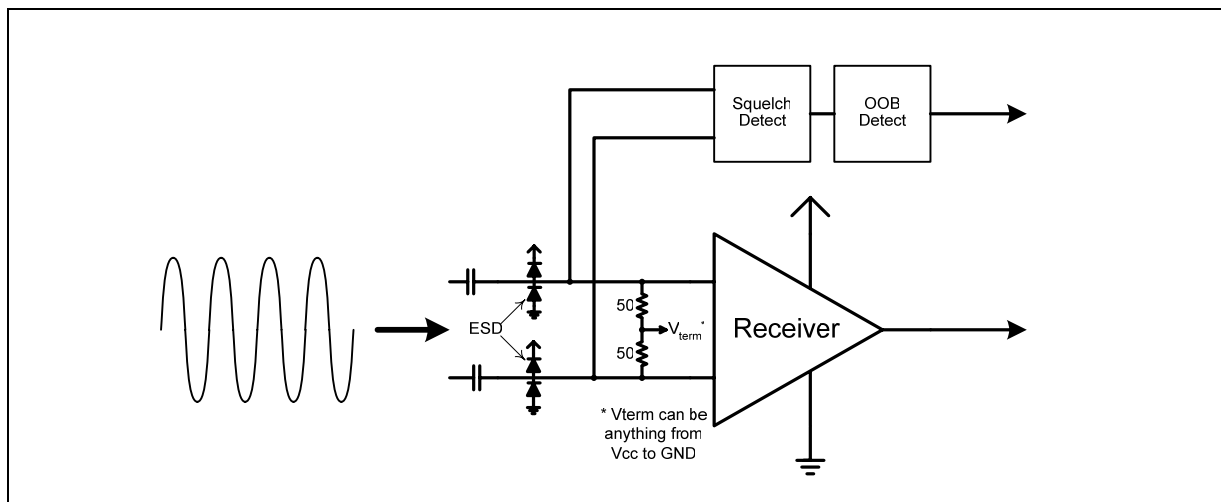


Figure 9: SATA receiver diagram

Figure 9 has been added to this document only as a reference diagram for possible discussion around this subject.

Other considerations not treated here

In a hot-plug scenario, there may be the additional factor of transients that can act as additive to the OOB burst amplitudes – particularly in a single-ended sense. The whole transient concern is a hot topic in both the SAS and SATA phy working groups. Can the additive effects of the SAS OOB signaling be dealt with in those transient discussions instead of over-complicating this SAS OOB algorithm issue?

Conclusion and call for assistance from SATA Phy TF members

With the SAS specified SAS-SATA OOB algorithm, SATA drives, particularly in hot-plug scenarios, will be exposed to SAS level OOB bursts. SAS system suppliers depend in part on the very attractive SATA support feature. It is a feature that is now demonstrating market acceptance and will cause market pull for SATA drives in the enterprise and SMB markets. The T10 technical committee intends to release SAS-1.1 to letter ballot shortly after the March 7-11 plenary week and humbly requests the assistance of the SATA phy WG and SATA drive supplier representatives to help resolve this issue in the interest of a robust SAS and SATA-in-the-enterprise deployment this year.

Thanks! -- T10 Phy Technical Committee