Maxtor Corporation
500 McCarthy Boulevard
Milpitas, CA 95035 USA

To: T10 SAS Phy Working Group
Contact: Dick Uber
Email: richard_uber@maxtor.com
Date: 05 November 2004

Subject: SAS 1.1, Phy hot plug and transients on SAS and SAS/SATA environments

1 Introduction

Previous discussions about transients focused on a circuit such as shown in Figure 1. An aggressor (either transmitter or receiver) is subject to a power up or power-down step function which results in a significant voltage surge on the victim device (either transmitter or receiver).

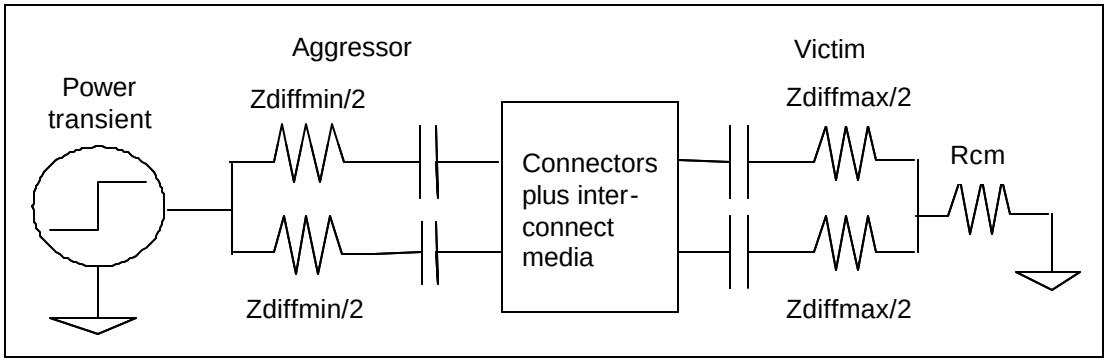

Figure 1 — Previous transient circuit

Similar transients can result from multiple causes:

a) System power applied to the aggressor PCB;
b) Secondary voltage regulator power up / down on the aggressor PCB;
c) Driver enable / disable on the aggressor;
d) Driver analog shift (i.e., driver strength, pre-emphasis strength, or enable/disable); and
e) Physical mating within the bus interconnect with either or both devices powered, such as when:
   A) An HDD is connected to a cable, and the cable is mated to host;
   B) An HBA/Expander is connected to a cable, and the cable is mated to HDD; and
   C) An HDD is mated to a back plane connector.

Note: Neither the aggressor silicon or the victim silicon actually operates on either of the power supplies passed over the SAS/SATA connectors. This assures that during a system level power up / down, the power up/down of the host silicon and target silicon will be uncorrelated events. As mentioned above, transmitter and receiver ports can both be aggressors, and can both be victims.

To encourage the community of SAS component vendors to make relatively uniform trade-offs between “robustness during transients” and “device cost”, we need to provide bounds on the strength of the aggressor designs and the robustness of the victim designs. Of course, all devices are at times both aggressors and victims. A restatement of the transient specifications will be proposed. An attempt will be made to include hot insertion and hot removal cases.

2 Discussion of the issues

The current draft specification is based on the circuit shown in Figure 2.

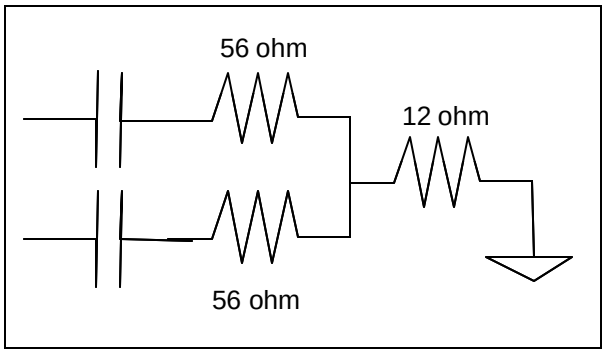

Figure 2 — Circuit in current draft

Note: the ground leg in the Y circuit is shown as resistive only. In practice, this leg is often very reactive. Also, there have been anecdotal reports that the real portion of the ground leg impedance can be as large as 10 Kohms.

Values specified for SAS were chosen to sum up to the 40 ohm minimum common mode impedance specification. The problem with studying or specifying transients with this circuit is that it is a “small signal AC model.” It does not include the additional (vendor unique) circuitry that will be added to insure that the SAS/ SATA component will meet reliability and performance goals.

One possible vendor specific complication may be the inclusion of significant capacitance on the common mode point as shown in Figure 3. Given the frequencies of interest here, small capacitances are unavoidable.

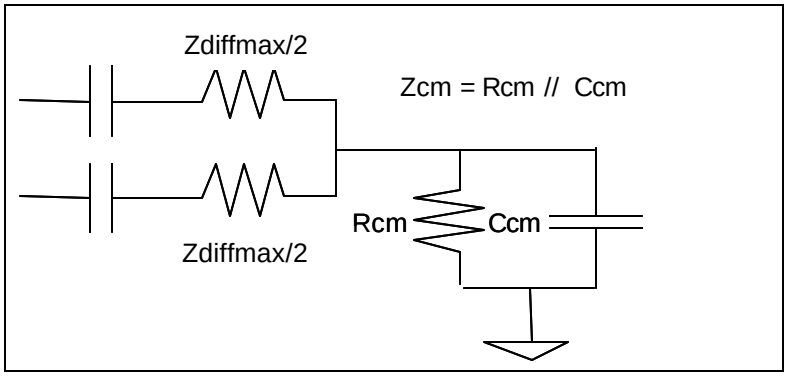

Figure 3 — Capacitance on the common mode point

The values of the capacitor can vary widely. On the low end, with an on-chip regulator, the total might be 10 pF, about 1000 times smaller than the coupling capacitor. On the high end, the common mode point could be a power supply bypass capacitor on the PCB. Bypass values and could be 10 uF, about 1000 times larger than the coupling capacitors. Having six orders of magnitude for the range of common mode capacitance is regrettable in the context of a standards effort.

Another feature absent in the small signal model is ESD structures as shown in Figure 4. Diodes are sometimes used to shunt large signals to a safe return path.

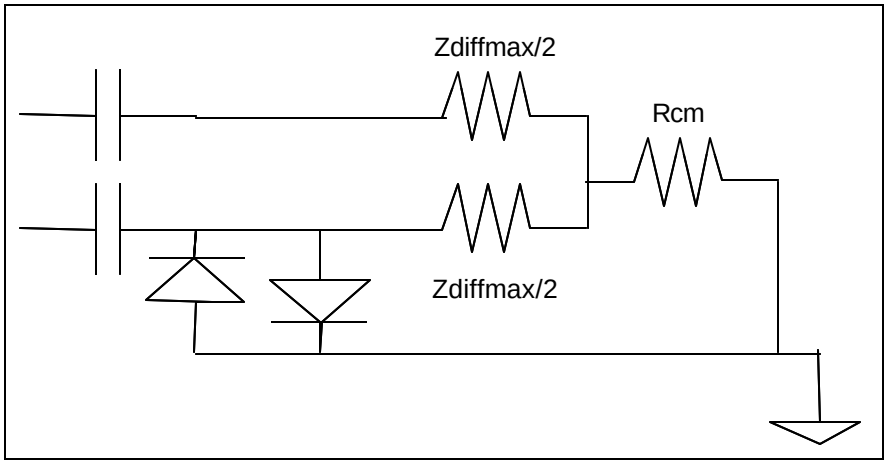

Figure 4 — Shunt diodes added to the circuit

3 Proposed changes to the standard

Because of the vendor specific large signal components such as in the examples above, it is difficult to cleanly specify limitations for aggressors and robustness for victims. The two examples given are examples of the many variations used on SAS and SATA devices. It would be even harder to predict what circuits may be implemented with future deep sub-micron semiconductor processes. Since it seems to be too late to force more common design details, I propose a path to follow:

a) Put bounds on aggressor impedances during transients;
b) Specify a test circuit for defining the maximum voltage that can be induced on the pins of a victim device; and
c) Place the burden of characterizing victim tolerance levels on the vendor.

Among the specification styles possible would be:

a) Vendor would specify four voltages for the SAS interface:
   A) Transmitter can induce no more than Vtx_ag on the standard load
   B) Receiver can induce no more than Vrx_ag on the standard load
   C) Receiver can tolerate transients which induce at least Vrx_vic on the standard load
   D) Transmitter can tolerate transients which induce at least Vtx_vic on the standard load
b) Ratings could be established tiers which correspond to ranges of voltages. Remember P-SCSI Hot Swap? Hot swap tolerance was broken into several classes of robustness. The same could be applied here:
   A) Class 2 transient tolerance might be 1.4v (margin on 1.2v supply);
   B) Class 3 transient tolerance might be 2.0v (margin on 1.8v supply); and
   C) A superb product might only generate class 2 transients, but would tolerate class 5 transients.

The rating system chosen should permit a non-scientist to select a safe mix of host/expander and target devices for his SAS system.

4 Details of the proposal

Bounds on aggressors would need to constrain both voltage swing and source impedance:

a) Current common mode impedance limits are based on small signal models and only appear to cover active bus operation;
b) Minimum common mode impedance needs to cover ALL of the following situations:
   A) Operational for data transfers;
   B) Power on / power off;
   C) Disable / enable active circuitry;
   D) Adjusting drive levels on terminators; and
   E) Enable / disable Pre-Emphasis.

c) For rise time / fall time:
   A) This is seldom an issue with power cycles;
   B) It is possibly a concern for enable / disable of active circuits;
   C) It is always a concern for hot plug; and
   D) It can be addressed by inclusion in standard load circuit.

There needs to be a standard load circuit definition where the transient test load includes resistors and 4ns transmission lines as shown in Figure 5.

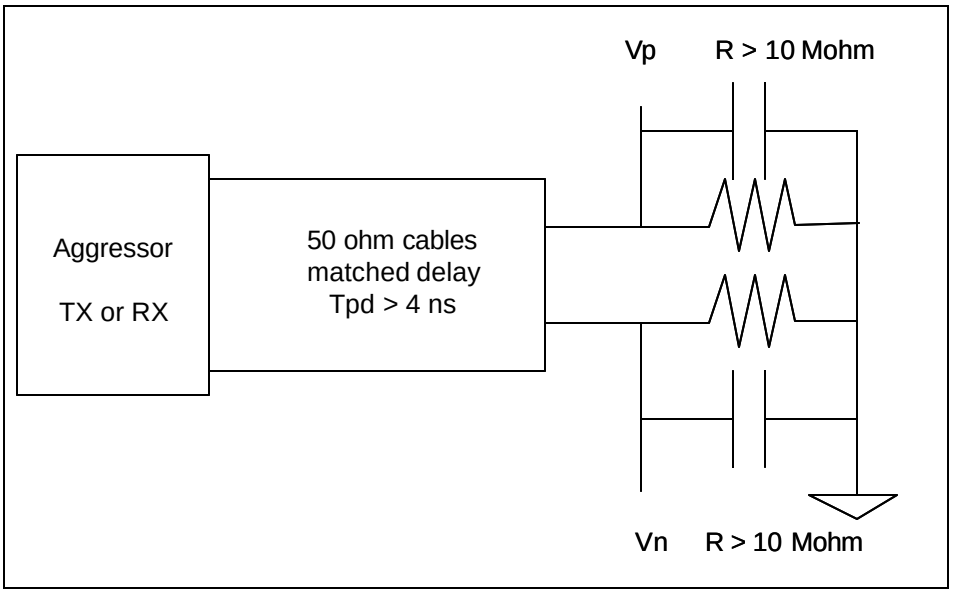

Figure 5 — Standard load circuit

Capacitors shown are for minimal fixture loads and for scope probes. Size should be bounded at a maximum set in the range 8 - 15 pF.

Note: This diagram differs from the SATA-II diagram in the requirement that the cable delays be greater than 4 ns.

For measurements on the standard transient load circuit:

a) We need two values to distinguish between brief over voltage case from transmission line reflection doubling, and the slower transient related to aggressor impedance and coupling capacitor;
b) Voltage amplitude should be measured at a level where the pulse width = 5 ns (this captures transmission line effects during hot insertion), and the higher of Vp and Vn is reported;
c) Voltage amplitude should be measured at a level where the pulse width = 20 ns (this measurement will better reflect the transient induced by charging the AC coupling capacitors) and the higher of Vp and Vn is reported;
d) Report largest voltages from ALL possible transients; and
e) Resistance in the load circuit should be specified to include scope probes and other test fixtures or instruments.

The above clarifications and additions to the PHY specification will provide vendors with sufficient information about the aggressor (voltage swing, source impedance, and slew rate effects). The vendor should be able to verify device tolerance to transients, including hot plug, through both simulation and test even though:

a) Specific implementations may or may not be sensitive to the transmission line overshoots; and
b) Backplanes which power disk drives and contact their signals through the same connector will not need to test for hot drive plugged into a cold backplane.

5 Additional comments on hot plugging

Hot plug has been treated here as a step function. Some inductive overshoot is inevitable on the leading edge of the step, The choice of looking at the voltage where the pulse width is 5 ns wide is an attempt to fairly capture the voltage doubling, but not any ringing.

Ringing from a contact bounce has not been considered.

The receiver and transmitter/interconnect form a network in the common mode. A receiver circuit can be highly reactive, and thus have high Q. The transmitter/interconnect may be mismatched in both termination and line impedance. The transmitter/interconnect is likely to limit the overall Q of the attached circuit. But, when the hot plug event occurs, the common mode circuit is excited by a step function. The response will ring according to the Q, and peak voltages and currents could exceed 2*Vinc. Common mode resonance effects merit further investigation.

Your feedback is requested!