To: T10 Technical Committee
From: Don Johnson, HP (don.johnson@hp.com)
Date: 8 July 2004
Subject: SATA Enhanced Signal Modes

This document is informational only. HP will submit the proposal that follows to the SATA II committee for consideration. The content of this proposal is intended to be complementary with the T10/04-220r0 proposal.

As SATA devices become more commonplace in business environments, all the way to the data center, they will be integrated into new server and storage subsystems. Inherent with many of these server configurations will be more complex backplane interconnects which exhibit greater signal loss between the system and the SATA device. The increased signal losses will be especially troublesome with 3Gb and higher signaling rates and will assuredly erode signal margins and adversely effect signal integrity across the link. A hint to the SATA device, about additional losses through the link, would allow adjustment of the device transmit characteristics, thereby, optimizing signal margin. When lossy interconnects are encountered, adjustments to certain transmit driver characteristics will significantly improve signal margins and thus, overall signal integrity across the link.

This proposal assumes that the link does not introduce catastrophic losses such that the device cannot complete initialization at nominal signal levels, but may introduce sufficient losses to affect link error rates if allowed to operate without adjustment.

As stated above, as SATA devices find their way into server configurations, along side of, and plug compatible with SAS devices, a common signaling solution is necessary for both types of device. The “Class” concept and methods of setting the class, as well as other signal characteristics, are consistent with a similar proposal being presented for SAS device signal enhancement. The proposed use of SET FEATURES and IDENTIFY DEVICE will allow a common programming model between SATA and the SAS SCSI mode page structures.

The suggested approach is to define three classes of interconnect characterizations shown here. Each class has some basic assumptions about the physical link.

Class 0	**Unknown Interconnect**
	• The system is not providing any hint of the transport characteristics
Class 1 Class is supported but signaling is at nominal (default) values	**Direct Cable**
	• Recommended no more than 2 interconnects (controller – cable – drive) • Maximum of 2.3 dB loss @ 1.5 GHz • Maximum of 4.5 dB loss @ 3.0 GHz
Class 2	**Typical Backplane and Cable**
	• Recommended no more than 3 interconnects (controller – cable – backplane – drive) • Maximum of 3.8 dB loss @ 1.5 GHz • Maximum of 7.5 dB loss @ 3.0 GHz
Class 3	**Lossy Backplane and Media**
	• 3 or more interconnects • Maximum of 5.4 dB loss @ 1.5 GHz • Maximum of 10.9 dB loss @ 3.0 GHz • Additional transitions through lossy interconnects

The method to allow adjustment to the data signal characteristics of a SATA device is through the established ATA/6, SET FEATURES command. Two new Sector Count Values will need to be defined under Serial ATA Features.

Prior to using the SET FEATURES command to effect any change in the signaling characteristics, it is proposed that the IDENTIFY DEVICE command be used to determine enhanced signaling capabilities and/or current signaling mode. Two bits of word 76 (fixed portion) of the IDENTIFY DEVICE command are required to indicate the level of support along with four concurrent words in the 0-255 word range to indicate the current signal class and programmable step settings.

To complement the ability to set and adjust signal characteristics at any given signaling rate, it is necessary to be able to determine the negotiated link signaling rate. Three bits of word 79 (variable portion) of the IDENTIFY DEVICE command are required to indicate current signaling rate. It is helpful in the programming model for the capable and negotiated link rate to be available from this register rather than a separate access to the SStatus register.

To facilitate the system designer in the characterization process of various interconnects and to better determine signal margins during qualification procedures of systems and devices, this proposal further defines a method to allow incremental adjustments of certain transmit signal characteristics.

Scope of Proposal

It is the intent of this proposal to limit the mechanism, in which to set the *Enhanced Signaling Modes*, to a SATA target device (i.e. disk drive). It is the SATA device that must be informed, through the SET FEATURES command, that it is to operate across a lossy interconnect. The command will most appropriately come from a system level component (BIOS or driver) that comprehends it own SATA interconnect topology. Further, it is assumed that the system (i.e. controller and any expanders) will comprehend its own interconnect topology and provide appropriate transmit signal characteristics back to the device.

SET FEATURES Command Specifics

Serial ATA features are controlled using:

Features Value	Description
10h	Enable use of Serial ATA feature
90h	Disable use of Serial ATA feature

The Sector Count register contains the specific Serial ATA feature to enable, disable or set value. The specific Serial ATA feature values for Enhanced Signaling Mode are:

Sector Count Value	Description
yyh	Enhanced Signaling Mode – Signal Class
zzh	Enhanced Signaling Mode – Signal Step Value

yyh and zzh = Two new requested Sector Count Values

Register assignment for Signal Class (Sector Count yyh)

Register	Function	Bits (7:0)
Features	Serial ATA Feature	10h
Sector Count	Signal Class	yyh
LBA Low	Signal Class for Link Rate 1.5 Gb/s	0xh
LBA Mid	Signal Class for Link Rate 3.0 Gb/s	0xh
LBA High	Signal Class for future link rate	0xh

Register assignment for Signal Step Values (Sector Count zzh)

Register	Function	Bits (7:0)

Features	Serial ATA Feature	10h
Sector Count	Signal Step Value	zzh
LBA Low	Use nominal or default value for Tx signal level	00h
Continued	Increase Tx signal level by programmable steps	01h-7Fh
Continued	Decrease Tx signal level by programmable steps	FFh-80h
LBA Mid	Use nominal or default value for Tx pre-emphasis	00h
Continued	Increase Tx pre-emphasis by programmable steps	01h-7Fh
Continued	Decrease Tx pre-emphasis by programmable steps	FFh-80h
LBA High	Use nominal or default value for Tx Slew Rate	00h
Continued	Increase Tx slew rate by programmable steps	01h-7Fh
Continued	Decrease Tx slew rate by programmable steps	FFh-80h

The value for the Signal Class in the respective signaling rate field is a 0 to 3 value to directly indicate the Signal Class mode. If the Signal Class requested in invalid, the device should return command aborted.

The value in the programmable Signal Step control field contains the step offset from the default transmitter signal characteristics for the current Signal Class. The field value should be treated as a 2's-complement signed value that can range from –128 to +127. If the device does not support a programmable stepping for the transmitter characteristic, or if the step requested is out of range of the particular transmitter characteristic, the device should return command aborted.

The SET FEATURES command for any of the enhances signaling modes should be issued to the device after the OOB sequence and early in the device initialization prior to any data transfer commands (i.e. along with the transfer mode settings). Immediately after successful command and status completion, the device will initiate a unique OOB sequence allowing the actual physical link signaling changes to occur on an inactive link. See the **SATA Device Specifics** section for more detail on device behavior.

IDENTIFY DEVICE Command Specifics

Serial ATA device capabilities and current settings are determined using words 75-79 as follows:

Word	O/M	F/V	Description
0-75			As defined in the ATA reference
76	O		Serial ATA capabilities
		F	15-12 Reserved
		F	11 1 = Supports transmit signal programmable stepping
		F	10 1 = Supports Signal Class
		F	9 Supports receipt of host-initiated interface power management requests
			8 Supports native command queuing
		F	7-4 Reserved
		F	3 Reserved for future Serial ATA
		F	signaling speed
			2 1 = Supports Serial ATA Gen-2 signaling
		F	speed (3.0 Gbps)
			1 1 = Supports Serial ATA Gen-1 signaling
		F	speed (1.5Gbps)
			0 Reserved (set to 0)
		F	
77			Reserved for future Serial ATA definition
78	O		Serial ATA features supported
			15-4 Reserved

		F	3	1 = device supports initiating interface power management
		F	2	1 = supports DMA Setup Auto-Activate optimization
		F	1	1 = supports non-zero buffer offsets in DMA Setup FIS
		F	0	Reserved (set to 0)

79	O		Serial ATA features enabled
			15-7 Reserved
		V	6 Reserved for future signaling speed
			5 1 = Signaling speed 3.0 Gbps
		V	4 1 = Signaling speed 1.5Gbps
		V	3 1 = device initiating interface power
		V	management enabled
			2 1 = DMA Setup Auto-Activate optimization
		V	enabled
			1 1 = non-zero buffer offsets in DMA Setup
		V	FIS enabled
			0 Reserved (set to 0)
		V	
80-255			As defined in the ATA reference
xxx	O	V	Current Enhanced Signaling Class
xxx		V	Current programmable stepping value – Tx Level
xxx		V	Current programmable stepping value – Tx Pre-emphasis
xxx		V	Current programmable stepping value – Tx Slew Rate

Key:
O/M = Mandatory/optional requirement.
M = Support of the word is mandatory.
O = Support of the word is optional.
F/V = Fixed/variable content
F = the content of the word is fixed and does not change. For removable media devices, these values may change when media is removed or changed.
V = the contents of the word is variable and may change depending on the state of the device or the commands executed by the device.
X = the content of the word may be fixed or variable.

SATA Device Specifics

Initial Device Mode

The target device should power-up with Signal Class 0 characteristics. The target is not expected to use the Signal Class mechanisms to compensate for the inability to complete a Phy Reset Sequences at a specific link rate. The target is not expected to retain the Signal Class information across a catastrophic “hard” reset or loss of power. Such events will reset the Signal Class and programmable stepping registers back to power-up defaults.

Changing the Signal Class

The target device Signal Class mode and signaling link rate are independent. Any change of the signal class using the SET FEATURES command will use the appropriate register bits corresponding to the current link rate. Changing the signal class should only be done after the current negotiated link rate is determined using the IDENTIFY DEVICE command.

Important note:
Changing the Signal Class (or making programmable step changes to any of the transmit characteristics) will modify the physical link transmit signal and should only be done on an idle link. Command status, for the SET FEATURES command, that initiated the change in signal characteristic, must be returned to the system before the device can negotiate a unique COMRESET when the transmit characteristics are actually modified. The unique COMRESET will

use the appropriate SET FEATURES register contents to effect the signaling characteristic change and issue sufficient align primitive cycles to allow the link to settle. Additionally, all previous device firmware settings shall remain unchanged. On completion of the COMRESET, the device should confirm the new negotiated link rate and requested signal class are consistent.

Link Error Behavior

During system test and qualification cycles, and for the purpose of determining signal margin limits, it is expected to generate link level errors and potentially a link reset condition. Since only the target device transmit characteristics are being adjusted, any link error conditions will be detected by the system/controller side of the link and target device will be unaware of such errors. This is true to the extent of a link reset condition generated by the system/controller, in which case, the target device will reset Signal Class and programmable stepping registers back to power-up defaults.