Proposed Draft

Serial ATA II Workgroup

Revision 0.75
10-March 2003

Serial ATA II Specification
Port Selector

Serial ATA II Port Selector Technical Editor:

Amber Huffman
Intel Corporation
2111 NE 25th Ave M/S JF2-53
Hillsboro, OR 97124 USA
Tel: (503) 264-7929
Email: amber.huffman@intel.com

Revision History

Revision	Date	Comments
0.10	September 19, 2002	Initial skeleton document. ADH
0.11	September 26, 2002	Clarified BIST support, out-of-band selection, COMINIT is only propagated and not generated by SFX, etc. ADH
0.12	October 4, 2002	Fixed some minor language. ADH
0.13	October 10, 2002	Added interlocked OOB and in-band default port handling. ADH
0.14	October 31, 2002	Incorporated summit feedback. ADH
0.15	November 14, 2002	Incorporated more summit feedback and feedback from 11/11 meeting. ADH
0.16	November 28, 2002	Added overview, example applications, and power-on reset material. ADH
0.50	December 3, 2002	Moved revision to 0.5 as requested by group, accepted red-lined changes. ADH
0.51	December 12, 2002	Added in comments from spec page turner. ADH
0.52	December 21, 2002	Highlighted wakeup budget impact out of partial/slumber, added section on orderly switch, added state machine for protocol-based port selection. ADH
0.53	January 6, 2003	Added protocol-based port selection state machine. ADH
0.70	January 13, 2003	Updated revision level for summit version of spec. ADH
0.71	January 30, 2003	Changes from summit and new marketing name. ADH
0.72	February 6, 2003	Instead of power "states" used power "conditions". Changed reset terminilogy in some sections to "initialized". ADH
0.73	February 26, 2003	When no active host port, new active host port is selected when any OOB signal is received over a host port. ADH
0.74	March 3, 2003	Updated "OOB signal" from last revision to "COMRESET or COMWAKE". ADH
0.75	March 10, 2003	Summit version of the spec. ADH

Table of Contents

1. Introduction

Port Selector is a mechanism that allows two different host ports to connect to the same device in order to create a redundant path to that device. In combination with RAID, the Port Selector allows system providers to build fully redundant solutions. A Port Selector can be thought of as a simple multiplexer as shown in Figure 1.

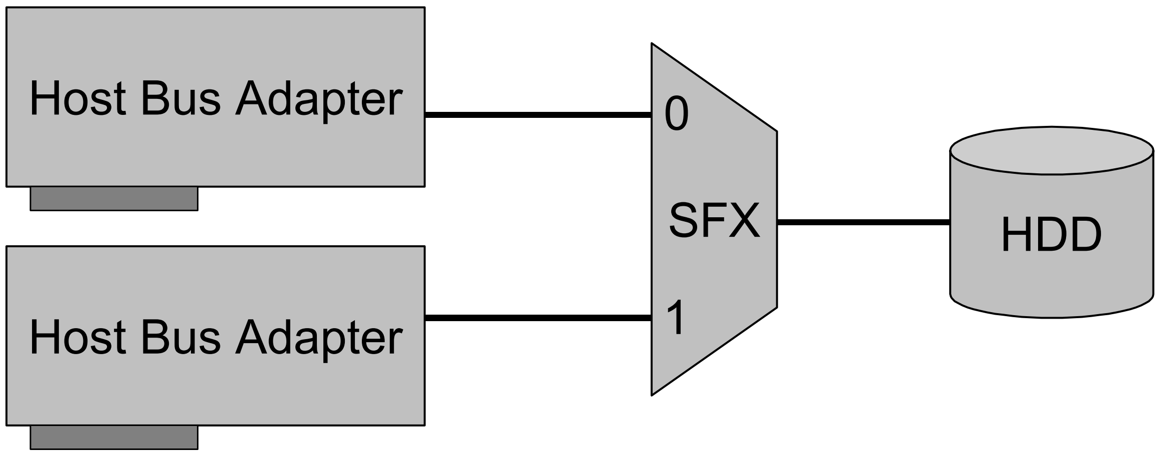

Figure 1 Port Selector Overview

1.1. Goals, Objectives, & Constraints

This specification defines an ingredient, Port Selector, that creates redundant paths to a device. The Port Selector allows system providers to build fully redundant solutions in combination with RAID.

Some of the goals and requirements for the Port Selector definition include:

- Serial ATA 1.0 devices may be attached without modification
- Serial ATA 1.0 host bus adapters may be attached with only software modifications
- Transport, Link, and Phy layer compatibility with Serial ATA 1.0 must be maintained for both hosts and devices
- No new primitives may be added as part of the definition
- No new FIS types may be added as part of the definition
- A Port Selector should not require a full Link layer or Transport layer

Some of the constraints include:

- Only two host connections are provided by the Port Selector
- Only one of the two host ports is active at a time (no active/active)

1.2. Example Applications

One example application of a Port Selector, as shown in Figure 2, is to provide a means for redundant access to a device. This ingredient, along with RAID, allows a system with no single point of failure to be built. Typically the Port Selector would be packaged in the hard drive carrier to create a single serviceable unit in case the hard drive failed. The total system would consist of two hosts each connected to a RAID array where each drive in the system had a Port Selector attached that was connected to each host. One host could be considered the live host and the other host may be the spare. In this configuration, the live host would maintain access to all of the devices and the spare host would only take over access to the devices if the live host had a failure.

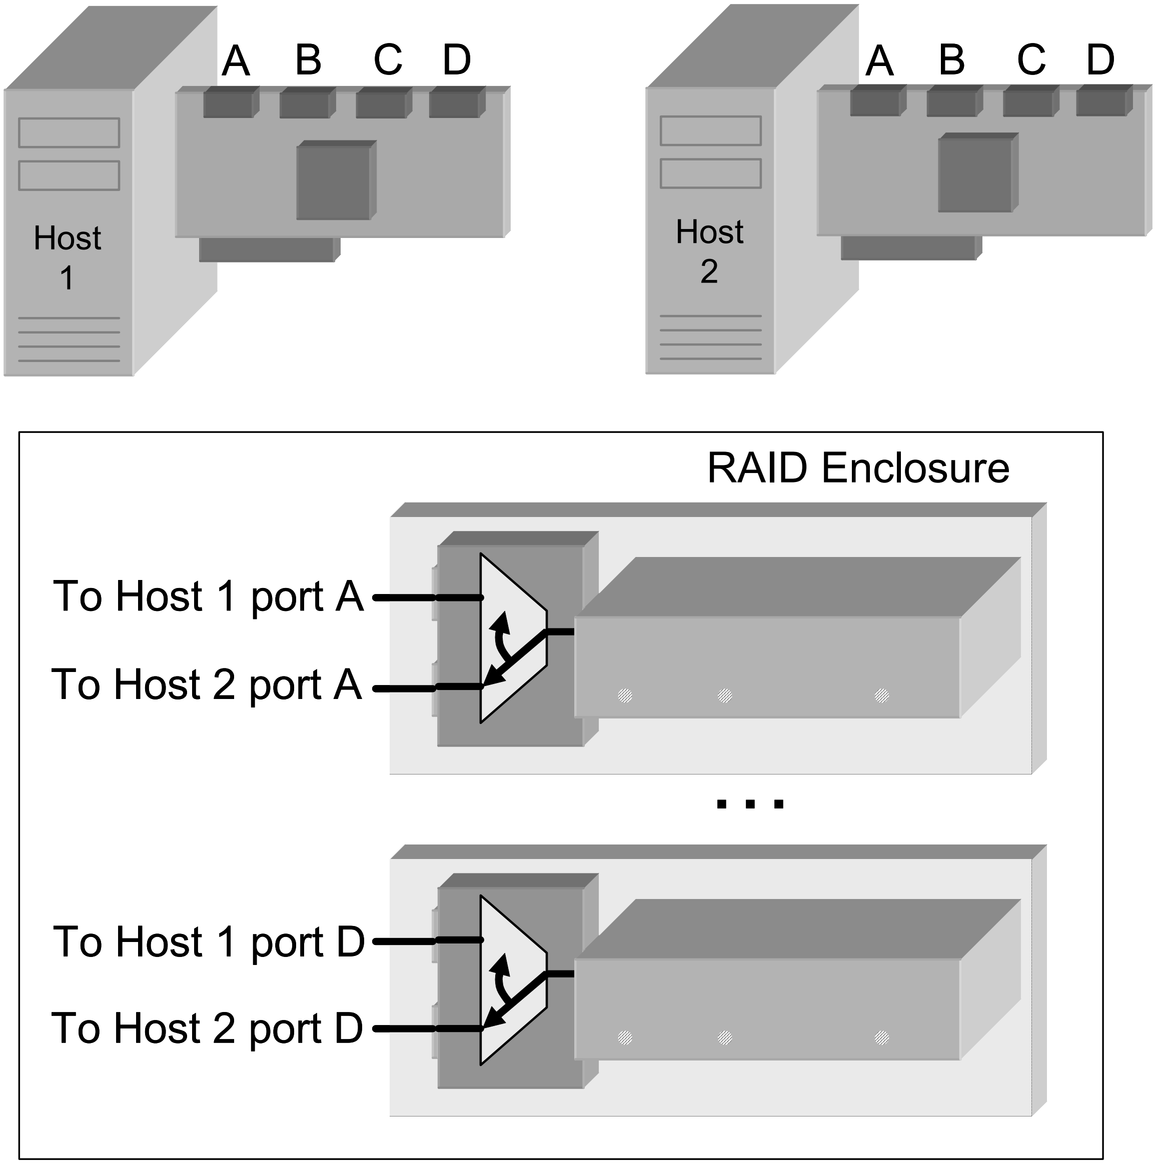

Figure 2 Example Failover Application with Two Hosts

1.3. Definitions, abbreviations, and conventions

1.3.1. Definitions and Abbreviations

The terminology used in this specification is consistent with the terminology used in the Serial ATA 1.0 and Serial ATA II specifications, and all definitions and abbreviations defined in those specifications are used consistently in this document. Additional terms and abbreviations introduced in this specification are defined in the following sections.

1.3.1.1. Active port

The active port is the currently selected host port on the Port Selector.

1.3.1.2. Device port

The device port is the port on the Port Selector that is connected to the device.

1.3.1.3. Host port

A port on the Port Selector that is connected to a host. There are two host ports on a Port Selector.

1.3.1.4. Inactive port

The inactive port is the host port that is not currently selected on the Port Selector.

1.3.1.5. Protocol-based port selection

Protocol-based port selection is a method that may be used by a host to select the host port that is active. Protocol-based port selection uses a sequence of Serial ATA out-of-band Phy signals to select the active host port.

1.3.1.6. Quiescent power condition

Entering a quiescent power condition for a particular Phy is defined as the Phy entering the idle bus condition as defined in section 6.7.5 of the Serial ATA 1.0 specification.

1.3.1.7. Side-band port selection

Side-band port selection is a method that may be used by a host to select the host port that is active. Side-band port selection uses a mechanism that is outside of the Serial ATA protocol for determining which host port is active. The mechanism used in implementations that support side-band port selection is outside the scope of this specification.

1.4. References

This specification is an extension to the Serial ATA 1.0 specification. The Serial ATA 1.0 specification is presumed as the underlying baseline for this specification. This specification makes reference to the following specifications:

Serial ATA: High Speed Serialized AT Attachment revision 1.0.

Serial ATA II Specification: Extensions to Serial ATA 1.0.

The specifications are available for download at www.serialata.org.

2. Overview

Port Selector is a mechanism that allows two different host ports to connect to the same device in order to create a redundant path to that device. Only one host connection to the device is active at a time. Effective use of a Port Selector requires coordinated access to the device between the two host ports. The host(s) must coordinate to determine which host port should be in control of the device at any given point in time. Definition of the coordination mechanism or protocol is beyond the scope of this specification.

Once the host(s) determine the host port that should be in control of the device, the host that contains the host port to be made active will take control of the device by selecting that host port to be active. The active host selects a port to be active by using either a protocol-based or side-band port selection mechanism. A side-band port selection mechanism can be as simple as a hardware select line that is pulled high to activate one host port and low to activate the other. The side-band port selection mechanism is outside the scope of this specification. A protocol-based port selection mechanism uses the Serial ATA protocol to cause a switch of active port. This specification defines a protocol-based port selection mechanism that uses a particular Morse coding of COMRESET signals to cause a switch of active host port. A Port Selector shall only support one selection mechanism at any point in time. The externally visible behavior of a Port

Selector is the same regardless of whether a protocol-based or side-band port selection mechanism is used.

Systems must be specifically configured to utilize a Port Selector.

3. Active Port Selection

The Port Selector has a single active host port at a time. The Port Selector shall support one of two mechanisms for determining which of the two host ports is active. The first mechanism is called side-band port selection. Side-band port selection uses a mechanism outside of the Serial ATA protocol for determining which host port is active. The second mechanism is called protocol-based port selection. Protocol-based port selection uses a sequence of Serial ATA out-of-band Phy signals to select the active host port.

Whether a protocol-based or side-band port selection mechanism is used, the Port Selector shall exhibit the behavior defined within this specification.

After selection of a new active port, the device and its connection speed to the Port Selector is in an unknown state. The device may have active commands outstanding from the previous active host that need to be flushed. The speed negotiated between the Port Selector and the device may also be in an indeterminate state. After an active port switch has been performed it is strongly suggested that the active host issue a COMRESET to the device to ensure that the device and device connection speed to the Port Selector is in a known state.

3.1. Protocol-based Port Selection

Protocol-based port selection is an active port selection mechanism that uses a sequence of Serial ATA out-of-band Phy signals to select the active host port. A Port Selector that supports protocol-based port selection shall have no active host port selected upon power-up. The first COMRESET or COMWAKE received over a host port shall select that host port as active. The host may then issue explicit switch signals to change the active host port.

Reception of the protocol-based port selection signal on the inactive host port causes the Port Selector to deselect the currently active host port and select the host port over which the selection signal is received. The protocol-based port selection signal is defined such that it can be generated using the Serial ATA 1.0 superset Status and Control registers and such that it can be received and decoded without the need for the Port Selector to include a full Link or Transport layer (i.e. direct Phy detection of the signal).

3.1.1. Port Selection Signal Definition

The port selection signal is based on a pattern of COMRESET out-of-band signals transmitted from the host to the Port Selector. As illustrated in Figure 3, the Port Selector shall qualify only the timing from the assertion of a COMRESET signal to the following assertion of the COMRESET signal in detecting the port selection signal.

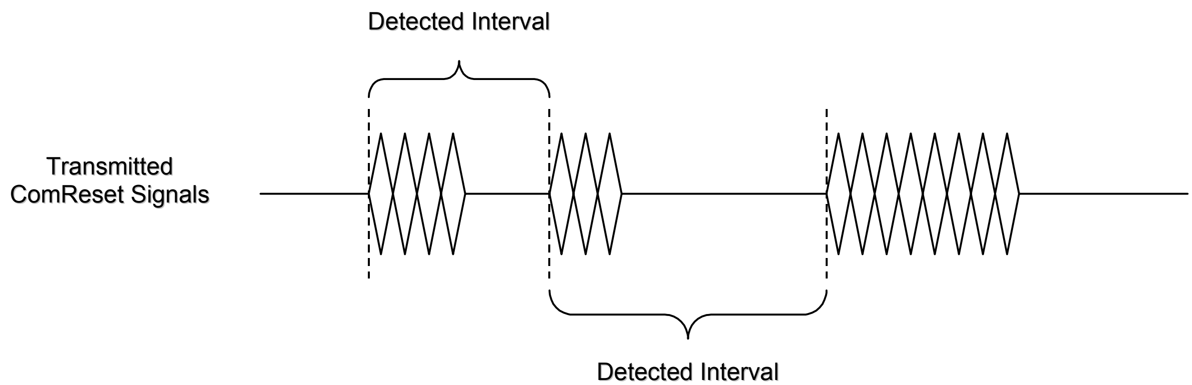

Figure 3 Port selection signal based on assertion of COMRESET to assertion of following COMRESET

The port selection signal is defined as a series of COMRESET signals with the timing from the assertion of one COMRESET signal to the assertion of the next as defined in Table 1 and illustrated in Figure 4. The Port Selector shall select the port, if inactive, on the de-assertion of COMRESET after receiving two complete back-to-back sequences with specified inter-burst spacing over that port (i.e. two sequences of two COMRESET intervals comprising a total of five COMRESET bursts with four inter-burst delays). Specifically, after receiving a valid port selection signal, the Port Selector shall not select that port to be active until the entire fifth COMRESET burst has been de-asserted (i.e. there has been an invalid inter-burst spacing for the fifth COMRESET signal). The Port Selector is only required to recognize the port selection signal over an inactive port. Reception of COMRESET signals over an active port is propagated to the device without any action taken by the Port Selector, even if the COMRESET signals constitute a port selection signal. This may result in multiple device resets.

The timings detailed in Table 1 shall be independent of the signaling speed used on the link. For example, the inter-reset timings are the same for links using Gen1 or Gen2 speeds.

	Nom	Min	Max	Units	Comments
T1	2.0	1.6	2.4	ms	Inter-reset assertion delay for first event of the selection sequence
T2	8.0	7.6	8.4	ms	Inter-reset assertion delay for the second event of the selection sequence

Table 1 Port selection signal inter-reset timing requirements

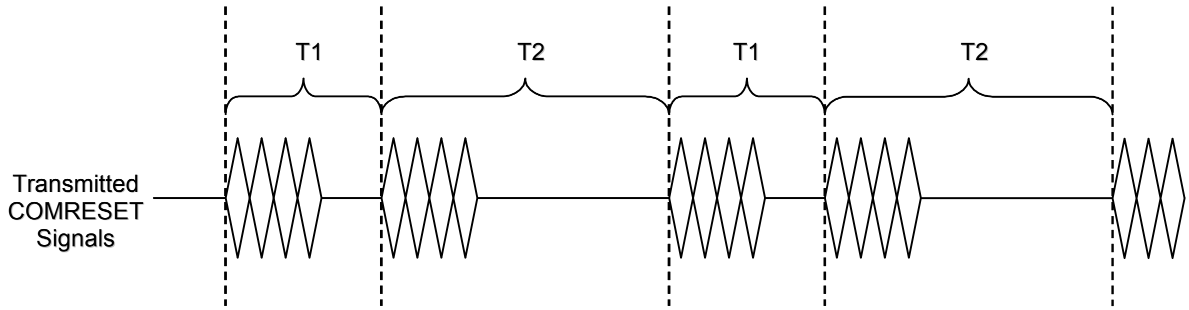

Figure 4 Complete port selection signal consisting of two sequences with requisite inter-reset spacings

The interpretation and detection of the COMRESET signal by the Port Selector is in accordance with the Serial ATA 1.0 definition, i.e. the COMRESET signal is detected upon receipt of the fourth burst that complies with the COMRESET signal timing definition. The inter-reset timings referred to here for the port selection signal are from the detection of a valid COMRESET signal to the next detection of such a signal, and are not related to the bursts that comprise the COMRESET signal itself.

3.1.2. Protocol-based Port Selection State Machine

The following state machine describes how the active host port is selected when using protocol-based port selection.

PBPS1: Reset[1]	Set Port Selector to initial state as specified in section 5.1. The Port Selector shall then issue COMINIT over each host port.		
1. Unconditional		→	NoActiveHostPort
NOTE: 1. This state is entered upon power-on reset			

PBPS2: NoActiveHostPort			
1. (COMRESET or COMWAKE) not received over host port A or host port B		→	NoActiveHostPort
2. (COMRESET or COMWAKE) received over host port A		→	SelectHostPortA
3. (COMRESET or COMWAKE) signal received over host port B and (COMRESET or COMWAKE) not received over host port A			SelectHostPortB

PBPS3: SelectHostPortA	The Port Selector performs the actions listed in the following order: 1. Host port B is selected as the inactive host port. 2. Host port B Phy is placed in the quiescent power condition. 3. Internal state of the Port Selector is initialized. 4. Host port A is placed in the active power condition. 5. Host port A is selected as the active host port.		
1. Unconditional		→	HostPortA_Active

PBPS4: HostPortA_Active		
1. Protocol-based port selection signal received over host port B	→	SelectHostPortB
2. No protocol-based port selection signal received over host port B	→	HostPortA_Active

PBPS5: SelectHostPortB	The Port Selector performs the actions listed in the following order: 1. Host port A is selected as the inactive host port. 2. Host port A Phy is placed in the quiescent power condition. 3. Internal state of the Port Selector is reset. 4. Host port B is placed in the active power condition. 5. Host port B is selected as the active host port.	
1. Unconditional	→	HostPortB_Active

PBPS6: HostPortB_Active		
1. Protocol-based port selection signal received over host port A	→	SelectHostPortA
2. No protocol-based port selection signal received over host port A	→	HostPortB_Active

3.1.3. Host Transmission Considerations (Informative)

In order to ensure the port selection signal is reliably conveyed to the Port Selector, the host should account for any other interface activity that may interfere with the transmitted COMRESET port selection sequence. For example, if the host periodically issues a COMRESET signal as part of a hardware-polled device presence detection mechanism, a periodic COMRESET signal could occur during the port selection signaling sequence, thereby corrupting the port selection sequence. In order to avoid such interactions, the host may elect to continually transmit the port selection sequence while monitoring the Phy status in the associated superset Status and Control register. When the port selection signal is recognized by the Port Selector and has taken effect, the host can detect a change in the PhyRdy status since the associated port will be activated and communications with it will be established. It is recommended that the host check the PhyRdy signal immediately before issuing each COMRESET burst in the protocol-based selection signal and only issue the next COMRESET burst if PhyRdy is not present.

3.2. Side-band Port Selection

The active host port may be selected by a side-band mechanism. Side-band port selection uses a mechanism outside of the Serial ATA protocol for determining which host port is active. One example of a side-band port selection mechanism is a hardware select line. The side-band selection mechanism used is outside the scope of this specification. A Port Selector that supports side-band port selection shall exhibit the behavior defined within this specification.

3.3. Behavior during a change of active port

During a change of active port, the previous host connection is broken and all state is flushed before the connection with the new active host is made. When a new active host port is selected, the Port Selector shall perform the following procedure:

1. The Port Selector shall immediately stop transmitting and enter the quiescent power condition on the previously active host port Phy (now the inactive host port).
2. The Port Selector shall initialize all internal state other than the active host port.
3. The Port Selector shall enter the active power condition on the new active host port.
4. The Port Selector shall allow out-of-band and in-band traffic to proceed between the new active host port and the device.

A Port Selector may support an orderly switch to a new active host port. A Port Selector that supports an orderly switch ensures that primitive alignment with the device Phy is maintained during the switch to the new active host port. Maintaining primitive alignment ensures that PhyRdy remains present between the Port Selector and the device throughout the switch to the new active host port.

After selection of a new active port, the device and its connection speed to the Port Selector is in an unknown state. The device may have active commands outstanding from the previous active host that need to be flushed. The speed negotiated between the Port Selector and the device may also be in an indeterminate state. After an active port switch has been performed it is strongly suggested that the active host issue a COMRESET to the device to ensure that the device and connection speed is in a known state.

4. Behavior and Policies

4.1. BIST support

A Port Selector is not required to respond to a BIST Activate FIS. The resultant behavior of sending a BIST Activate FIS through a Port Selector is undefined.

4.2. Flow control signaling latency

The Port Selector must satisfy the flow control signaling latency specified in section 7.4.7 of the Serial ATA 1.0 specification. The Port Selector shall ensure that the flow control signaling latency is met on a per link basis. Specifically, the Port Selector shall ensure that the flow control signaling latency is met between:

1. The Port Selector active host port and the host it is connected to
2. The Port Selector device port and the device it is connected to

The Port Selector shall not reduce the flow control signaling latency budget of the active host it is connected to or the device it is connected to.

4.3. Power Management

The Port Selector must maintain the active host port across power management events and only allow an active host port change after receiving a valid active port switch signal.

The Phy on the inactive host port shall be in the quiescent power condition. Upon detecting that the PhyRdy signal is not present for the active host port or device port, the Port Selector shall place that Phy in a quiescent power condition.

If the PhyRdy signal is not present between the device and the Port Selector, the Phy connected to the active host port shall enter the quiescent power condition and immediately squelch the Phy transmitter. If the PhyRdy signal is not present between the active host and the Port Selector, the Phy connected to the device shall enter the quiescent power condition and immediately squelch the Phy transmitter. During these periods while PhyRdy is not present, OOB signals shall still be propagated between the active host and the device to ensure that communication can be established.

If the Port Selector is able to determine that the active host and device negotiated a SLUMBER power management transition, the Port Selector may recover from the quiescent power condition in the time defined by the SLUMBER power state. If the Port Selector does not have this ability, the Port Selector shall recover from the quiescent power condition in the time defined by the PARTIAL power state.

4.3.1. Wakeup Budget (Informative)

The wakeup budget out of PARTIAL or SLUMBER may increase when a Port Selector is connected to a device. When the active host Phy comes out of low power condition, the Port Selector active host Phy may wakeup before causing the Port Selector device Phy to wakeup which in turn will wakeup the device. This may add up to twice the latency to come out of a low power condition. System designers should take this extra latency into account in their designs.

4.4. Out of Band (OOB) Phy signals

Out of Band communication shall be end-to-end. The Port Selector shall only propagate OOB signals received from the active host to the device and shall not generate OOB signals to the device on its own. The Port Selector shall not propagate OOB signals received from the inactive port to the device. The Port Selector shall only propagate OOB signals received from the device to the host over the active host port. If no active host port is selected, the Port Selector shall propagate OOB signals received from the device over both host ports. The Port Selector is allowed to delay delivery of OOB signals to be propagated but must maintain the relative timing between OOB signals.

The Port Selector shall not respond to COMRESET signals received over the inactive host port. The inactive host port Phy shall remain in the quiescent power condition when COMRESET is received over the inactive host port.

4.5. Hot Plug

The Port Selector support for hot plug is completely defined by the interlocked handling of out of band Phy signals and the power management behavior.

The Port Selector shall only generate a COMINIT over a host port when a COMINIT signal is received from the device as detailed in section 4.4. If a drive connected to a Port Selector is hot plugged, the drive will issue a COMINIT sequence as part of its normal power-up sequence in accordance with Serial ATA 1.0. The Port Selector will propagate the COMINIT over the active host port (or both host ports if both host ports are inactive). The host will see the COMINIT signal and then interrogate the port to determine whether a drive is attached.

If a drive connected to a Port Selector is hot unplugged, the Port Selector shall squelch the transmitter for the active host port as detailed in section 4.3. The active host will determine that the PhyRdy signal is no longer present and will determine that there is no longer a drive present.

4.6. Spread spectrum clocking

The Port Selector shall support spread spectrum clocking receive on all of its ports. The Port Selector may support spread spectrum transmit. It is recommended that a configuration jumper be used to enable/disable spread spectrum clocking if it is dynamically settable. There is no in-band means provided to enable/disable spread spectrum clocking if it is statically configurable.

If spread spectrum clocking is used, the spreading domain between the host and the Port Selector is not required to be the same as the spreading domain between device and the Port Selector. The signals passing through a Port Selector may be re-spread.

5. Power-up and Resets

5.1. Power-up

Upon power-up, the Port Selector shall reset all internal state, including the active host port. This will cause no active host port to be selected when protocol-based port selection is used.

5.1.1. Presence detection of Port Selector

There is no mechanism using the Serial ATA protocol to detect the presence of a Port Selector. Presence detect of a Port Selector may be done via a mechanism outside the scope of this specification or by user configuration that is outside the scope of this specification.

5.2. Resets

5.2.1. COMRESET

When COMRESET is received over the active host port the Port Selector shall reset all internal state, the active host port shall remain unchanged. The Port Selector shall take no reset action upon receiving a COMRESET signal over the inactive host port except as specified in section 3.1 when there is no active host port selected after power-up.

5.2.2. Software reset and DEVICE RESET

The Port Selector shall not reset in response to receiving a Software reset or the DEVICE RESET command.