

POS-PHY™ Level 3

SATURN COMPATIBLE PACKET OVER SONET INTERFACE SPECIFICATION FOR PHYSICAL AND LINK LAYER DEVICES



Issue 4: June 2000

CONTENTS

1 Introduction 1

2 POS-PHY Interface Reference Definition.....2

3 Compatibility Options3

4 Specification Summary.....5

 4.1 Signal Naming Conventions5

 4.2 Bus Widths.....5

 4.3 Clock Rates.....5

 4.4 Packet Interface Synchronization.....5

 4.5 Application Line Rates7

 4.6 PHY and Link Layer Interface Example8

5 Interface Data Structures 11

6 Transmit Packet Interface Description 13

 6.1 Transmit Signals 13

 6.2 Examples 18

 6.3 AC Timing21

7 Receive Packet Interface Description24

 7.1 Receive Signals.....24

 7.2 Examples29

 7.3 AC Timing33

1 Introduction

This document specifies PMC-Sierra's recommended interface for the interconnection of Physical Layer (PHY) devices to Link Layer devices implementing Packet over SONET (POS). POS-PHY fulfills the need for system designers to target a standard POS Physical Layer interface. Although targeted at implementing POS, the POS-PHY specification is not restricted to this application. It provides a versatile bus interface for exchanging packets within a communication system.

POS-PHY Level 3 was developed with the cooperation of the SATURN Development Group to cover all application bit rates up to and including 2.4 Gbit/s. It defines the requirements for interoperable single-PHY (one PHY layer device connected to one Link Layer device) and multi-PHY (multiple PHY layer devices connected to one Link Layer device) applications. It stresses simplicity of operation to allow forward migration to more elaborate PHY and Link Layer devices.

The ATM Forum Utopia Level 2 Specification, the SCI-PHY Level 2 Specification, the SATURN POS-PHY Level 2 Specification and ATM Forum proposals for Utopia Level 3 were used as the basis for this POS-PHY specification, with several adaptations to support variable packet sizes. However, the POS-PHY specification does not intend to be compatible with the above mentioned specifications.

This specification defines 1-the physical implementation of the POS-PHY bus, 2-the signaling protocol used to communicate data and 3-the data structure used to store the data into holding FIFO's.

This information is provided as a reference to independent developers of integrated circuits or system-level circuits who wish to interoperate with SATURN Compatible components. Going forward, references to "POS-PHY" shall be taken to indicate "POS-PHY Level 3" unless otherwise noted .

2 POS-PHY Interface Reference Definition

The POS-PHY interface defines the interface between SONET/SDH Physical layer devices and Link Layer devices, which can be used to implement several packet-based protocols like HDLC and PPP.

Figure 2.1: POS-PHY Reference Points

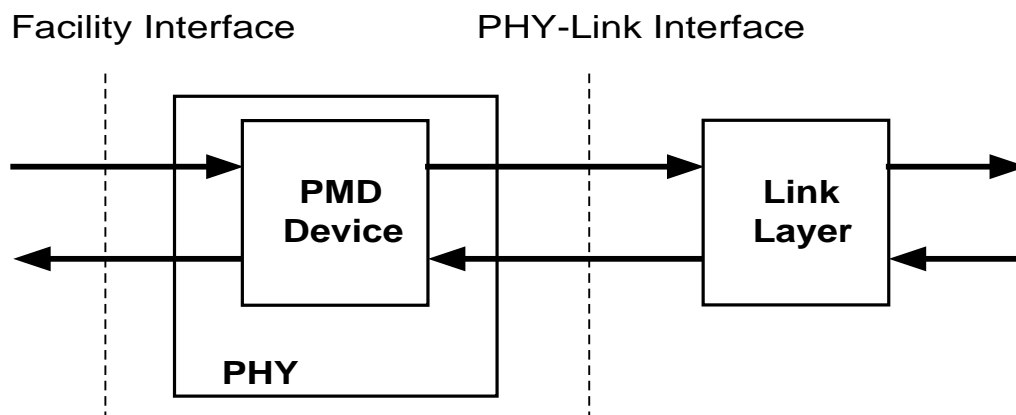


Diagram Definitions

- Facility:** An optical fiber, twisted pair electrical or coaxial cable electrical transmission facility.
- Link Layer:** Switching Function Layer.
- PHY:** Physical Layer for Packet over SONET.
- PHY-LINK:** Physical Layer to Link Layer electrical interface.
- PMD:** Physical Medium Dependent Layer.

POS-PHY Level 3 specifies the PHY-LINK interface. The Facility Interface (such as SONET OC-3) is defined by several National and International standards organizations including Bellcore and ITU.

3 Compatibility Options

The POS-PHY Level 3 specification does not attempt to be compatible to any existing standard. There is no existing equivalent standard. Specifically, POS-PHY does not intend to be compatible with similar ATM specifications like Utopia and SCI-PHY. Although this information is not critical to any implementation, the following bullets highlight the differences between the Utopia/SCI-PHY and POS-PHY interfaces.

- Allowance for an 8-bit bus or a 32-bit bus interface running at a maximum speed of 104 MHz. The bus interface is point-to-point (one output driving only one input load).
- Byte or double-word (4 bytes) data format that can accommodate variable size packets.
- Modification to the RSOC/TSOC start of cell signals to identify the start of packets being transferred over the interface. Renamed the signals to RSOP/TSOP.
- Addition of the REOP/TEOP end of packet signals which delineate the end of packets being transferred over the interface.
- Addition of the RMOD[1:0]/TMOD[1:0] modulo signals which indicate if the last double-word of the packet transfer contains 1, 2, 3 or 4 valid bytes of data.
- Addition of the RERR/TERR error signals which, during the end of the packet, indicates if the transferred packet must be discarded/aborted.
- Deletion of the RCA signal. Receive interface of the PHY pushes packet data to the Layer device. Multi-port PHY devices are responsible for performing round-robin servicing of their ports. PHY address is inserted in-band with the packet data.
- Transmit interface of the PHY is selected using an in-band address that is provided on the same bus transferring the packet data.
- Addition of the RSX/TSX start of transfer signals which identify when the in-band port address of the PHY is on the RDAT/TDAT bus.
- Modification of the TCA cell available signals to form the TPA packet available signals. TPA logic values are defined based on the FIFO fill level (in terms of bytes). In multi-port PHY devices, PHY status indication can be provided using either a polling or a direct status indication scheme. Polled PHY address is provided by a separate address bus and has pipelined timing.

- Interface FIFO fill level granularity is byte-based. For the Transmit Interface FIFO, the packet available status and start of transmission FIFO fill levels are programmable. For the Receive Interface, the maximum burst transfer size is programmable.

4 Specification Summary

4.1 Signal Naming Conventions

The interface where data flows from the Link Layer device to the Physical layer device will be labeled the Transmit Interface. The interface where data flows from the Physical Layer device to the Link Layer device will be labeled the Receive Interface. All signals are active high unless denoted by a trailing "B".

SIGNAL	Active high signaling.
SIGNALB	Active low signaling.

4.2 Bus Widths

POS-PHY compatible devices support an 8-bit and/or a 32-bit data bus structure. The bus interface is point-to-point (one output driving only one input load) and thus a 32-bit data bus would support only one device. To support multiple lower rate devices with point-to-point connections, an 8-bit data bus structure is defined. Thus, each PHY device would use an 8-bit interface reducing the total number of pins required.

To support variable length packets, the RMOD[1:0]/TMOD[1:0] signals are defined to specify valid bytes in the 32-bit data bus structure. Each double-word must contain four valid bytes of packet data until the last double-word of the packet transfer which is marked with the end of packet REOP/TEOP signal. This last double-word of the transfer will contain up to four valid bytes specified by the RMOD[1:0]/TMOD[1:0] signals.

4.3 Clock Rates

POS-PHY compatible devices can support a transfer clock rate up to 104 MHz. Some devices may support multiple rates. Generally, devices targeted at single or multi-PHY applications, where the aggregate PHY bit rate approaches 622 Mbit/s will use the 8-bit data bus structure with a 104 MHz FIFO clock rate. Devices targeted at applications where the aggregate PHY bit rate approaches 2.4 Gbit/s will use the 32-bit data bus structure with a 104 MHz FIFO clock rate.

4.4 Packet Interface Synchronization

The POS-PHY packet interface supports transmit and receive data transfers at clock rates independent of the line bit rate. As a result, PHY layer devices must support packet rate decoupling using FIFOs.

To ease the interface between the Link Layer and PHY layer devices and to support multiple PHY layer interfaces, FIFOs are used. Control signals are provided to both the Link Layer and PHY layer devices to allow either one to exercise flow control. Since the bus interface is point-to-point, the receive interface of the PHY device pushes data to the Link Layer device. For the transmit interface, the packet available status granularity is byte-based.

In the receive direction, when the PHY layer device has stored an end-of-packet (a complete small packet or the end of a larger packet) or some predefined number of bytes in its receive FIFO, it sends the in-band address followed by FIFO data to the Link Layer device. The data on the interface bus is marked with the valid signal (RVAL) asserted. A multi-port PHY device with multiple FIFOs would service each port in a round-robin fashion when sufficient data is available in its FIFO. The Link Layer device can pause the data flow by deasserting the enable signal (RENB).

In the transmit direction, when the PHY layer device has space for some predefined number of bytes in its transmit FIFO, it informs the Link Layer device by asserting a transmit packet available (TPA). The Link Layer device can then write the in-band address followed by packet data to the PHY layer device using an enable signal (TENB). The Link Layer device shall monitor TPA for a high to low transition, which would indicate that the transmit FIFO is near full (the number of bytes left in the FIFO can be user selectable, but must be predefined), and suspend data transfer to avoid an overflow. The Link Layer device can pause the data flow by deasserting the enable signal (TENB).

POS-PHY defines both byte-level and packet-level transfer control in the transmit direction. In byte-level transfer, FIFO status information is presented on a cycle-by-cycle basis. With packet-level transfer, the FIFO status information applies to segments of data. When using byte level transfer, direct status indication must be used. In this case, the PHY layer device provides the transmit packet available status of the selected port (STPA) in the PHY device. As well, the PHY layer device may provide direct access to the transmit packet available status of all ports (DTPA[]) in the PHY device if the number of ports is small. With packet level transfer, the Link Layer device is able to do status polling on the transmit direction. The Link Layer device can use the transmit port address TADR[] to poll individual ports of the PHY device, which all respond on a common polled (PTPA) signal.

Since the variable size nature of packets does not allow any guarantee as to the number of bytes available, in both transmit and receive directions, a selected PHY transmit packet available is provided on signal STPA and a receive data valid on signal RVAL. STPA and RVAL always reflect the status of the selected PHY to or from which data is being transferred. RVAL indicates if valid data is available on the receive data bus and is defined such that data transfers can be aligned with packet boundaries.

Physical layer port selection is performed using in-band addressing. In the transmit direction, the Layer device selects a PHY port by sending the address on the TDAT[] bus marked with the TSX signal active and TENB signal inactive. All subsequent TDAT[] bus operations marked with the TSX signal inactive and the TENB active will be packet data for the specified port. In the receive direction, the PHY device will specify the selected port by sending the address on the RDAT[] bus marked with the RSX signal active and RVAL signal inactive. All subsequent RDAT[] bus operations marked with RSX inactive and RVAL active will be packet data from the specified port.

Both byte-level and packet-level modes are specified in this standard in order to support the current low density multi-port physical layer devices and future higher density multi-port devices. When the number of ports in the physical layer device is limited, byte-level transfer using DTPA[] signals provides a simpler implementation and reduces the need for addressing pins. In this case, direct access will start to become unreasonable as the number of ports increase. Packet-level transfer provides a lower pin count solution using the TADR[] bus when the number of ports is large. In-band addressing ensures the protocol remains consistent between the two approaches. However, the final choice left to the system designers and physical layer device manufacturers to select which approach best suits their desired applications.

4.5 Application Line Rates

The numerous combinations of clock rates and bus widths allow the Packet over SONET Interface for PHY layer devices (POS-PHY) to support a wide range of line rates. Table 4.1 gives examples of line rates supported by POS-PHY interfaces and the maximum number of channels supported by the interface definitions.

Table 4.1: Interface Bit Rates

Standard Reference	Bit Rate (Mbit/s)	Number of PHYs (800 Mbit/s bus)	Number of PHYs (3.2 Gbit/s bus)
SONET STS-1	51.84	12	48
SONET STS-3 SDH STM-1	155.52	4	16
SONET STS-12 SDH STM-4	622.08	1	4
SONET STS-48 SDH STM-16	2488.32	N/A	1

4.6 PHY and Link Layer Interface Example

Figure 4.1 illustrates a conceptual example of how a single multi-port PHY device may be interfaced to a Link Layer device. In the example, the Link Layer device is connected to a single package four channel PHY layer device using the 32-bit interface. Figure 4.2 illustrates a conceptual example of how multi-port PHY devices may be interfaced to a single Link Layer device. The Link Layer device is connected to two four-channel PHY layer devices using 8-bit interfaces.

In both examples, the PHY devices are using the direct status indication signals DTPA[]. Optionally, the Link Layer device can perform multiplexed status polling using the PTPA signals.

Figure 4.1: POS-PHY PHY to Link Layer 32-bit Interface

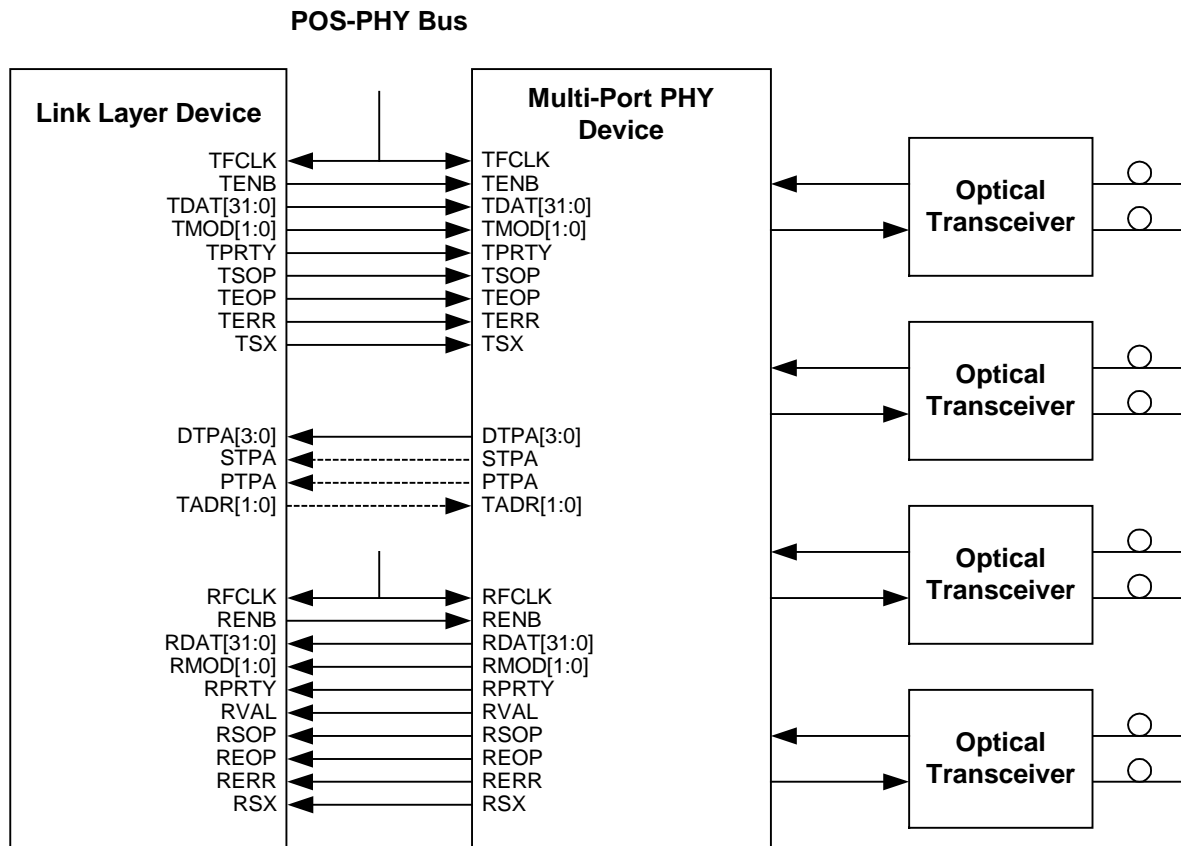
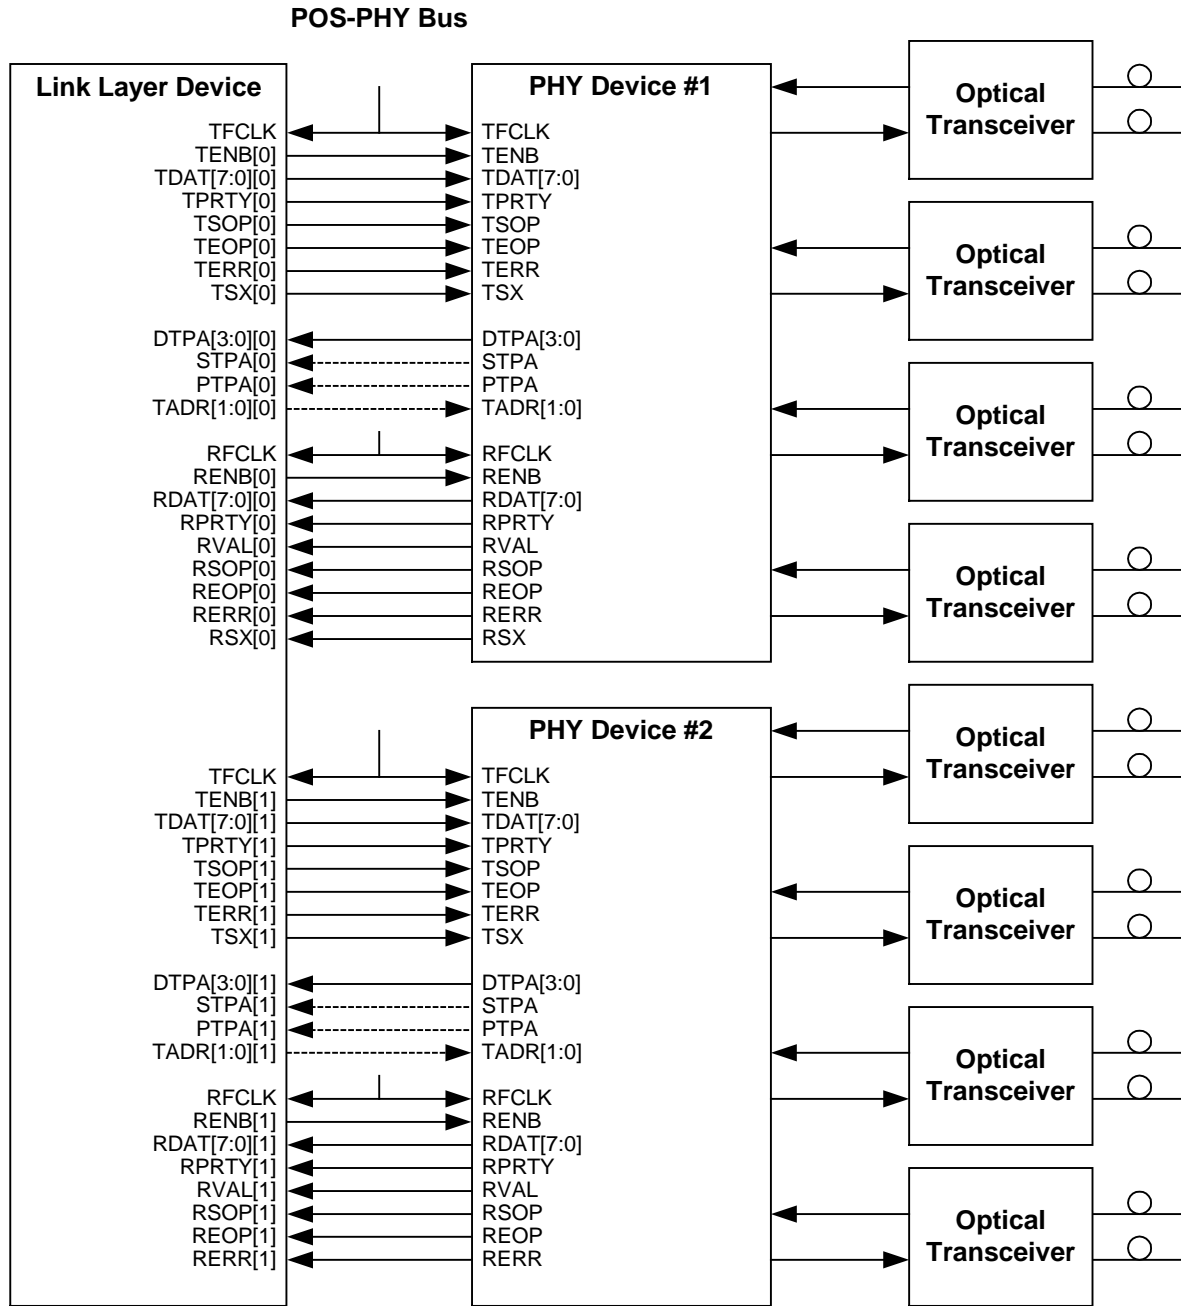


Figure 4.2: POS-PHY PHY to Link Layer 8-Bit Interface



5 Interface Data Structures

Packets shall be written into the transmit FIFO and read from the receive FIFO using a defined data structure. Octets are written in the same order they are to be transmitted or they were received on the SONET line. Within an octet, the MSB (bit 7) is the first bit to be transmitted. The POS-PHY specification does not preclude the transfer of 1-byte packets. In this case, both start of packet and end of packet signals shall be asserted simultaneously.

For packets longer than the PHY device FIFO, the packet must be transferred over the bus interface in sections. The number of bytes of packet data in each section may be fixed or variable depending on the application. In general, the Receive Interface will round-robin between receive FIFOs with fill levels exceeding a programmable high water mark or with at least one end of packet stored in the FIFO. The Receive Interface would end the transfer of data when an end of a packet is transferred or when a programmable number of bytes have been transferred. The Link Layer device may send fixed size sections of packets on the Transmit Interface or use the TPA signal to determine when the FIFO reaches a full level.

Figure 5.1 illustrates the data structure for the 32-bit bus interface. The double-word with the last byte of the packet is marked with TEOP asserted and TMOD[1:0] specifying the number of valid bytes. Figure 5.2 illustrates the data structure for the 8-bit bus interface. The first byte of the packet is marked with TSOP asserted. The last byte of the packet is marked with TEOP asserted. In all cases, the PHY address is marked with TSX asserted.

In both illustrations, the in-band port address for multi-port PHY devices is not shown. The Transmit Interface would send the PHY port address, on the same bus as the data, marked with the TSX signal active and the TENB signal inactive. Subsequent data transfers on the Transmit Interface would use the transmit FIFO selected by the in-band address. On the Receive Interface, the PHY device reports the receive FIFO address in-band with the RSX signal active and the RVAL signal inactive before transferring packet data. For both cases, large packets which exceed the FIFO size will be transferred over the POS-PHY interface in sections with appropriate in-band addressing prefixing each section.

The in-band address is specified in a single clock cycle operation marked with the RSX/TSX signals. The port address is specified by the TDAT[7:0]/RDAT[7:0] signals. The address is the numeric value of the TDAT[7:0]/RDAT[7:0] signals where bit 0 is the least significant bit and bit 7 is the most significant bit. Thus, up to 256 ports may be supported by a single interface. With a 32-bit interface, the upper 24 bits shall be ignored.

The POS-PHY specification does not define the usage of any packet data. In particular, POS-PHY does not define any field for error correction. Notice that if the Link Layer device uses the PPP protocol, a Frame Check Sequence (FCS) must be processed. If the Physical Layer device does not insert the FCS field before transmission, these bytes should be included at the end of the packet. If the Physical Layer device does not strip the FCS field in the receive direction, these bytes will be included at the end of the packet.

Figure 5.1: 32-bit Interface Data Structures

	Bit 31	Bit 24	Bit 23	Bit 16	Bit 15	Bit 8	Bit 7	Bit 0
Dword 1	Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6	Byte 7	Byte 8
Dword 2	Byte 9	Byte 10	Byte 11	Byte 12	Byte 13	Byte 14	Byte 15	Byte 16
Dword 3	Byte 17	Byte 18	Byte 19	Byte 20	Byte 21	Byte 22	Byte 23	Byte 24
Dword 4	Byte 25	Byte 26	Byte 27	Byte 28	Byte 29	Byte 30	Byte 31	Byte 32
Dword 5	Byte 33	Byte 34	Byte 35	Byte 36	Byte 37	Byte 38	Byte 39	Byte 40
Dword 6	Byte 41	Byte 42	Byte 43	Byte 44	Byte 45	Byte 46	Byte 47	Byte 48
	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Dword 28	Byte 109	XX	XX	XX	XX	XX	XX	XX

A 109 Byte Packet

Figure 5.2: 8-bit Interface Data Structures

	Bit 7	Bit 0
Byte 1	Byte 1	
Byte 2	Byte 2	
Byte 3	Byte 3	
Byte 4	Byte 4	
Byte 5	Byte 5	
Byte 6	Byte 6	
	⋮	
Byte 62	Byte 62	

A 62 Byte Packet

6 Transmit Packet Interface Description

The standard FIFO depth for POS-PHY interfaces is 256 octets. The transmit buffer shall have a programmable threshold defined in terms of the number of bytes available in the FIFO for the assertion and deassertion of the transmit packet available flags.

In this fashion, transmit latency can be managed, and advance TPA lookahead can be achieved. This will allow a Link Layer device to continue to burst data in, without overflowing the transmit buffer, after TPA has been deasserted.

In the transmit direction, the PHY layer device shall not initiate data transmission before a predefined number of bytes or an end of packet flag has been stored in the transmit FIFO. This capability does not affect the POS-PHY bus protocol, but is required to avoid transmit FIFO underflow and frequent data retransmission by the higher layers.

6.1 Transmit Signals

Table 6.1 lists the transmit side POS-PHY specification signals. All signals are expected to be updated and sampled using the rising edge of the transmit FIFO clock TFCLK. A fully compatible POS-PHY Physical Layer device requires at least a 256 byte deep FIFO.

Table 6.1: Transmit Signal Descriptions

Signal Name	Direction	Function
TFCLK	Clock Source to Link and PHY	Transmit FIFO Write Clock (TFCLK). TFCLK is used to synchronize data transfer transactions between the LINK Layer device and the PHY layer device. TFCLK may cycle at a rate up to 104 MHz.
TERR	LINK to PHY	Transmit Error Indicator (TERR) signal. TERR is used to indicate that there is an error in the current packet. TERR should only be asserted when TEOP is asserted; it is considered valid only when TENB is simultaneously asserted.

Signal Name	Direction	Function
TENB	LINK to PHY	<p>Transmit Write Enable (TENB) signal.</p> <p>The TENB signal is used to control the flow of data to the transmit FIFOs. When TENB is high, the TDAT, TMOD, TSOP, TEOP and TERR signals are invalid and are ignored by the PHY. The TSX signal is valid and is processed by the PHY when TENB is high.</p> <p>When TENB is low, the TDAT, TMOD, TSOP, TEOP and TERR signals are valid and are processed by the PHY. Also, the TSX signal is ignored by the PHY when TENB is low.</p>
TDAT[31:0]	LINK to PHY	<p>Transmit Packet Data Bus (TDAT[]) bus.</p> <p>This bus carries the packet octets that are written to the selected transmit FIFO and the in-band port address to select the desired transmit FIFO. The TDAT bus is considered valid only when TENB is simultaneously asserted.</p> <p>When a 32-bit interface is used, data must be transmitted in big endian order on TDAT[31:0]. Given the define data structure, bit 31 is transmitted first and bit 0 is transmitted last.</p> <p>When an 8-bit interface is used, the PHY supports only TDAT[7:0].</p>
TPRTY	LINK to PHY	<p>Transmit bus parity (TPRTY) signal.</p> <p>The transmit parity (TPRTY) signal indicates the parity calculated over the TDAT bus. When an 8-bit interface is used, the PHY only supports TPRTY calculated over TDAT[7:0]. TPRTY is considered valid only when TENB or TSX is asserted.</p> <p>When TPRTY is supported, the PHY layer device is required to support odd parity. The PHY layer device is required to report any parity error to higher layers, but shall not interfere with the transferred data.</p>

Signal Name	Direction	Function
TMOD[1:0]	LINK to PHY	<p>Transmit Word Modulo (TMOD[1:0]) signal.</p> <p>TMOD[1:0] indicates the number of valid bytes of data in TDAT[31:0]. The TMOD bus should always be all zero, except during the last double-word transfer of a packet on TDAT[31:0]. When TEEP is asserted, the number of valid packet data bytes on TDAT[31:0] is specified by TMOD[1:0].</p> <p style="padding-left: 40px;"> TMOD[1:0] = "00" TDAT[31:0] valid TMOD[1:0] = "01" TDAT[31:8] valid TMOD[1:0] = "10" TDAT[31:16] valid TMOD[1:0] = "11" TDAT[31:24] valid </p> <p>When an 8-bit interface is used, the TMOD[1:0] bus is not required.</p> <p>TMOD is considered valid only when TENB is simultaneously asserted.</p>
TSX	LINK to PHY	<p>Transmit Start of Transfer (TSX) signal.</p> <p>TSX indicates when the in-band port address is present on the TDAT bus. When TSX is high and TENB is high, the value of TDAT[7:0] is the address of the transmit FIFO to be selected. Subsequent data transfers on the TDAT bus will fill the FIFO specified by this in-band address.</p> <p>For single port PHY devices, the TSX signal is optional as the PHY device will ignore in-band addresses when TENB is high.</p> <p>TSX is considered valid only when TENB is not asserted.</p>
TSOP	LINK to PHY	<p>Transmit Start of Packet (TSOP) signal.</p> <p>TSOP is used to delineate the packet boundaries on the TDAT bus. When TSOP is high, the start of the packet is present on the TDAT bus.</p> <p>TSOP is required to be present at the beginning of every packet and is considered valid only when TENB is asserted.</p>

Signal Name	Direction	Function
TEOP	LINK to PHY	<p>Transmit End of Packet (TEOP) signal.</p> <p>TEOP is used to delineate the packet boundaries on the TDAT bus. When TEOP is high, the end of the packet is present on the TDAT bus.</p> <p>When a 32-bit interface is used, TMOD[1:0] indicates the number of valid bytes the last double-word is composed of when TEOP is asserted. When an 8-bit interface is used, the last byte of the packet is on TDAT[7:0] when TEOP is asserted.</p> <p>TEOP is required to be present at the end of every packet and is considered valid only when TENB is asserted.</p>
TADR[]	LINK to PHY Packet-Level Mode	<p>Transmit PHY Address (TADR[]) bus.</p> <p>The TADR bus is used with the PTPA signal to poll the transmit FIFO's packet available status.</p> <p>When TADR is sampled on the rising edge of TFCLK by the PHY, the polled packet available indication PTPA signal is updated with the status of the port specified by the TADR address on the following rising edge of TFCLK.</p>
DTPA[]	PHY to LINK Byte-Level Mode	<p>Direct Transmit Packet Available (DTPA[]).</p> <p>The DTPA bus provides direct status indication for the corresponding ports in the PHY device.</p> <p>DTPA transitions high when a predefined (normally user programmable) minimum number of bytes is available in its transmit FIFO. Once high, the DTPA signal indicates that its corresponding transmit FIFO is not full. When DTPA transitions low, it optionally indicates that its transmit FIFO is full or near full (normally user programmable).</p> <p>DTPA is required if byte-level transfer mode is supported. DTPA is updated on the rising edge of TFCLK.</p>

Signal Name	Direction	Function
STPA	PHY to LINK Byte-Level Mode	<p>Selected-PHY Transmit Packet Available (STPA) signal.</p> <p>STPA transitions high when a predefined (normally user programmable) minimum number of bytes are available in the transmit FIFO specified by the in-band address on TDAT. Once high, STPA indicates the transmit FIFO is not full. When STPA transitions low, it indicates that the transmit FIFO is full or near full (normally user programmable).</p> <p>STPA always provides status indication for the selected port of the PHY device in order to avoid FIFO overflows while polling is performed. The port which STPA reports is updated on the following rising edge of TFCLK after the PHY address on TDAT is sampled by the PHY device.</p> <p>STPA is required if byte-level transfer mode is supported. STPA is updated on the rising edge of TFCLK.</p>
PTPA	PHY to LINK Packet-Level Mode	<p>Polled-PHY Transmit Packet Available (PTPA) signal.</p> <p>PTPA transitions high when a predefined (normally user programmable) minimum number of bytes are available in the polled transmit FIFO. Once high, PTPA indicates that the transmit FIFO is not full. When PTPA transitions low, it optionally indicates that the transmit FIFO is full or near full (normally user programmable).</p> <p>PTPA allows the polling of the PHY selected by the TADR address bus. The port which PTPA reports is updated on the following rising edge of TFCLK after the PHY address on TADR is sampled by the PHY device.</p> <p>PTPA is required if packet-level transfer mode is supported. PTPA is updated on the rising edge of TFCLK.</p>

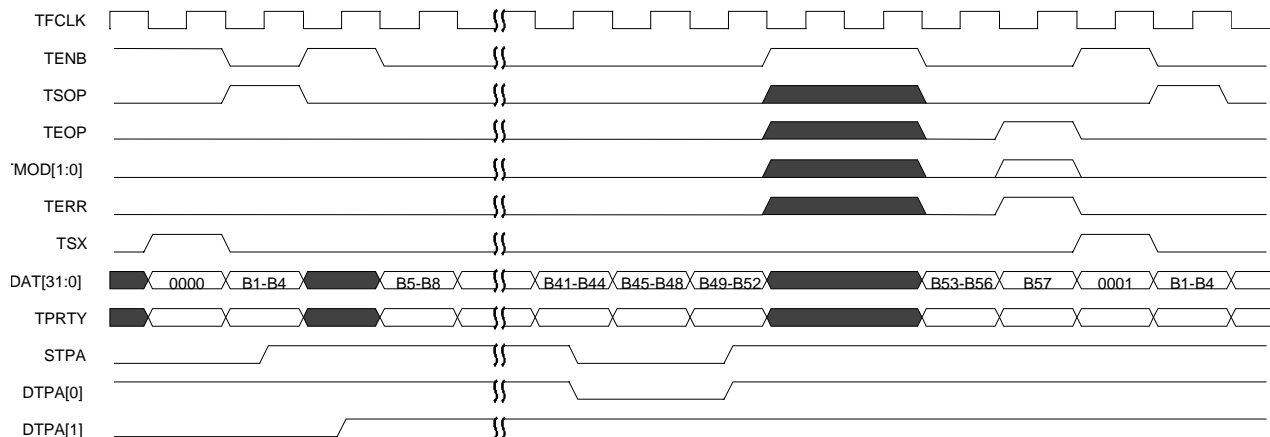
6.2 Examples

The following examples are not part of the requirements definition of the POS-PHY compatibility specification. They are only informative and provide an aid in the visualization of the interface operation. The examples only present a limited set of scenarios; they are not intended to imply restrictions beyond that presented in the text of the specification. If any apparent discrepancies exist between the examples and the text, the text shall take precedence.

The POS-PHY transmit interface is controlled by the Link Layer device using the TENB signal. All signals must be updated and sampled using the rising edge of the transmit FIFO clock, TFCLK. Figure 6.1 is an example of a multi-port PHY device with two channels. The PHY layer device indicates that a FIFO is not full by asserting the appropriate transmit packet available signal DTPA. DTPA remains asserted until the transmit FIFO is almost full. Almost full implies that the PHY layer device can accept at most a predefined number of writes after the current write.

If DTPA is asserted and the Link Layer device is ready to write a word, it should assert TSX, deassert TENB and present the port address on the TDAT bus if required. Subsequent data transfers with TENB low are treated as packet data which is written to the selected FIFO. At any time, if the Link Layer device does not have data to write, it can deassert TENB. The TSOP and TEOP signals must be appropriately marked at the start and end of packets on the TDAT bus.

Figure 6.1: Transmit Logical Timing



When DTPA transitions low and it has been sampled, the Link Layer device can write no more than a predefined number of bytes to the selected FIFO. In this example, the

predefined value is two double-words or eight bytes. If the Link Layer writes more than that predefined number of words and DTPA remains deasserted throughout, the PHY layer device should indicate an error condition and ignore additional writes until it asserts DTPA again.

Figure 6.2 is an example of the Link Layer device using the polling feature of the Transmit Interface. For comparison purposes, the direct transmit packet available signals for the example ports are provided in the diagram. The status of a given PHY port may be determined by setting the polling address TADR bus to the port address. The polled transmit packet available signal PTPA is updated with the transmit FIFO status in a pipelined manner. The Link Layer device is not restricted in its polling order. The selected transmit packet available STPA signal allows monitoring the selected PHY status and halting data transfer once the FIFO is full. The PTPA signal allows polling other PHY's at any time, including while a data transfer is in progress. The system could be configured differently.

Figure 6.2: Packet-Level Transmit Polling Logical Timing


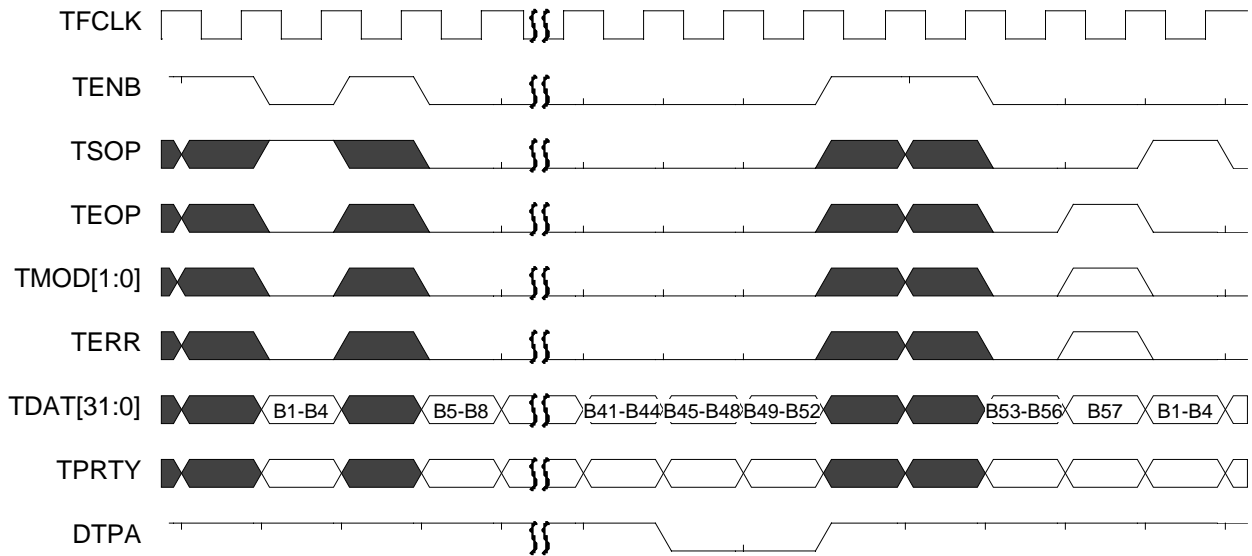


Figure 6.3 shows an example of data transfer to a single-port PHY device without the use of in-band addressing, which is optional for single-port applications. During packet transfer, the Link Layer device pauses for its own internal reasons. As in Figure 6.1, the PHY device indicates its FIFO status by means of DTPA.

Figure 6.3: Transmit Logical Timing
 (single-port example without in-band addressing)



6.3 AC Timing

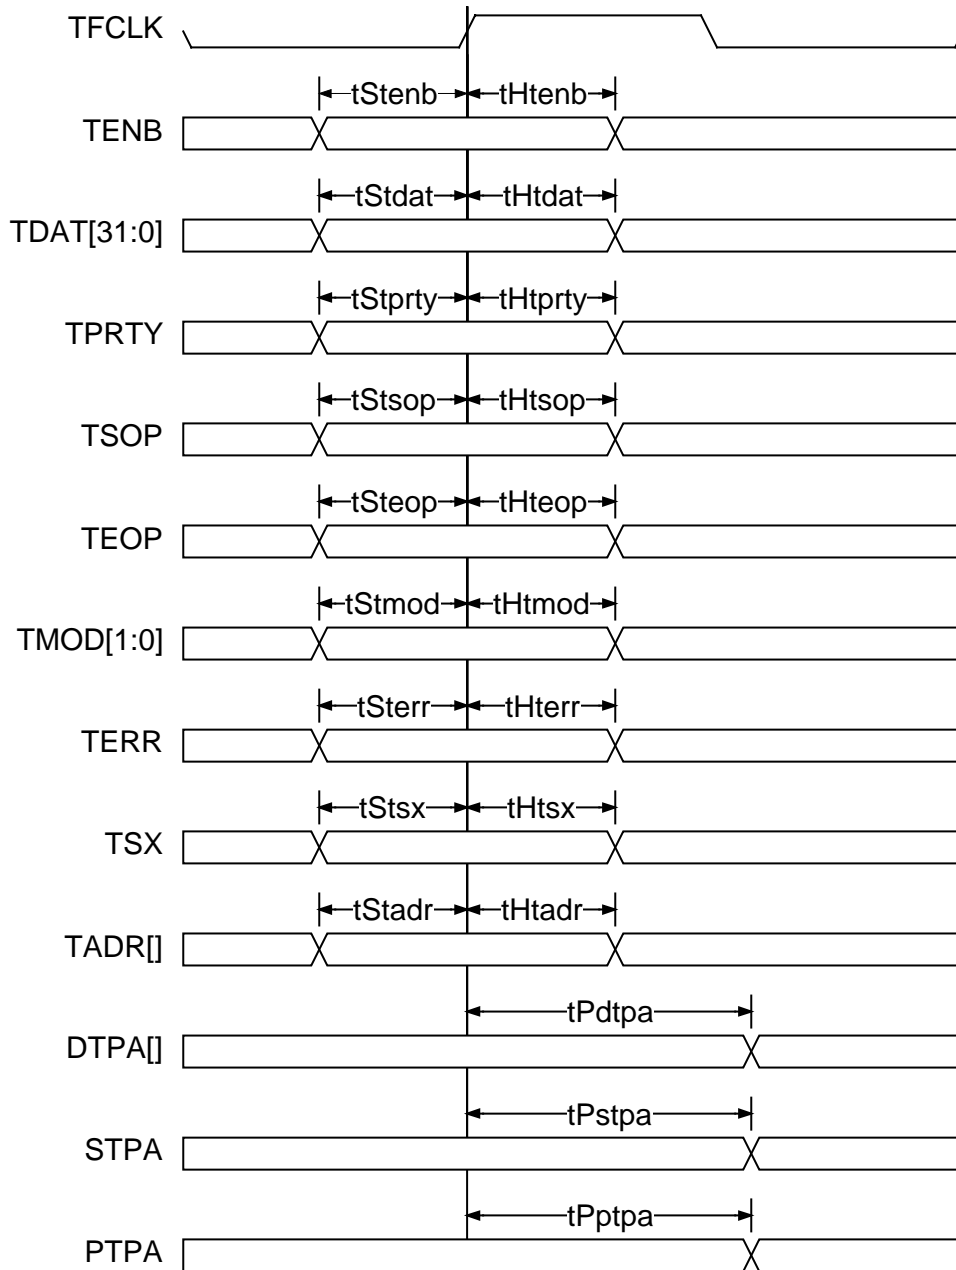
All AC Timing is from the perspective of the PHY layer device in a PHY-LINK interface.

Table 6.2: Transmit Interface Timing

Symbol	Description	Min	Max	Units
	TFCLK Frequency		104	MHz
	TFCLK Duty Cycle	40	60	%
$t_{S_{tenb}}$	TENB Set-up time to TFCLK	2.0		ns
$t_{H_{tenb}}$	TENB Hold time to TFCLK	0.5		ns
$t_{S_{tdat}}$	TDAT[15:0] Set-up time to TFCLK	2.0		ns
$t_{H_{tdat}}$	TDAT[15:0] Hold time to TFCLK	0.5		ns
$t_{S_{tprty}}$	TPRTY Set-up time to TFCLK	2.0		ns
$t_{H_{tprty}}$	TPRTY Hold time to TFCLK	0.5		ns
$t_{S_{tsop}}$	TSOP Set-up time to TFCLK	2.0		ns
$t_{H_{tsop}}$	TSOP Hold time to TFCLK	0.5		ns
$t_{S_{teop}}$	TEOP Set-up time to TFCLK	2.0		ns
$t_{H_{teop}}$	TEOP Hold time to TFCLK	0.5		ns
$t_{S_{tmod}}$	TMOD Set-up time to TFCLK	2.0		ns
$t_{H_{tmod}}$	TMOD Hold time to TFCLK	0.5		ns
$t_{S_{terr}}$	TERR Set-up time to TFCLK	2.0		ns
$t_{H_{terr}}$	TERR Hold time to TFCLK	0.5		ns
$t_{S_{tsx}}$	TSX Set-up time to TFCLK	2.0		ns
$t_{H_{tsx}}$	TSX Hold time to TFCLK	0.5		ns
$t_{S_{tadr}}$	TADR[4:0] Set-up time to TFCLK	2.0		ns
$t_{H_{tadr}}$	TADR[4:0] Hold time to TFCLK	0.5		ns
$t_{P_{dtpa}}$	TFCLK High to DTPA Valid	1.5	6	ns

Symbol	Description	Min	Max	Units
tP_{stpa}	TFCLK High to STPA Valid	1.5	6	ns
tP_{ptpa}	TFCLK High to PTPA Valid	1.5	6	ns

Figure 6.3. Transmit Physical Timing



Notes on Transmit I/O Timing:

- Note 1: When a set-up time is specified between an input and a clock, the set-up time is the time in nanoseconds from the 1.4 Volt point of the input to the 1.4 Volt point of the clock.
- Note 2: When a hold time is specified between an input and a clock, the hold time is the time in nanoseconds from the 1.4 Volt point of the clock to the 1.4 Volt point of the input.
- Note 3: Output propagation delay time is the time in nanoseconds from the 1.4 Volt point of the reference signal to the 1.4 Volt point of the output.
- Note 4: Maximum output propagation delays are measured with a 30 pF load on the outputs.

7 Receive Packet Interface Description

The standard FIFO depth for POS-PHY interfaces is 256 octets. As the interface is point-to-point, the PHY device is required to push receive packet data to the Link Layer device. This arrangement simplifies the interface between the PHY device and the Link Layer device. Traditional polling schemes for the receive side are not required, saving a significant number of pins.

The receive FIFO shall have a programmable threshold defined in terms of the number of bytes of packet data stored in the FIFO. A multi-port PHY device must service each receive FIFO with sufficient packet data to exceed the threshold or with an end of packet. The PHY should service the required FIFOs in a round-robin fashion. The type of round-robin algorithm will depend on the various data rates supported by the PHY device and is outside this specification.

The amount of packet data transferred, when servicing the receive FIFO, is bounded by the FIFO's programmable threshold. Thus, a transfer is limited to a maximum of 256 bytes of data (64 cycles for a 32-bit interface or 256 cycles for an 8-bit interface) or until an end of packet is transferred to the Layer device. At the end of a transfer, the PHY device will round-robin to the next receive FIFO.

The PHY device should support a programmable minimum pause of 0 or 2 clock cycles between transfers. A pause of 0 clock cycles maximizes the throughput of the interface. A pause of 2 clock cycles allows the Layer device to pause between transfers.

7.1 Receive Signals

Table 7.1 lists the receive side POS-PHY specification signals. All signals are expected to be updated and sampled using the rising edge of the receive FIFO clock, RFCLK. A fully compatible POS-PHY Physical Layer device requires at least a 256-byte receive FIFO.

Table 7.1: Receive Signal Descriptions

Signal Name	Direction	Function
RFCLK	Clock Source to Link and PHY	Receive FIFO Write Clock (RFCLK). RFCLK is used to synchronize data transfer transactions between the Link Layer device and the PHY layer device. RFCLK may cycle at a rate up to 104 MHz.

Signal Name	Direction	Function
RVAL	PHY to LINK	<p>Receive Data Valid (RVAL) signal.</p> <p>RVAL indicates the validity of the receive data signals. RVAL is low between transfers and when RSX is asserted; it is also low when the PHY pauses a transfer due to an empty receive FIFO. When a transfer is paused by holding RENB high, RVAL will hold its value unchanged, although no new data will be present on RDAT[31:0] until the transfer resumes.</p> <p>When RVAL is high, the RDAT[31:0], RMOD[1:0], RSOP, REOP and RERR signals are valid. When RVAL is low, the RDAT[31:0], RMOD[1:0], RSOP, REOP and RERR signals are invalid and must be disregarded.</p> <p>The RSX signal is valid when RVAL is low.</p>
RENB	LINK to PHY	<p>Receive Read Enable (RENB) signal.</p> <p>The RENB signal is used to control the flow of data from the receive FIFO's. During data transfer, RVAL must be monitored as it will indicate if the RDAT[31:0], RPRTY, RMOD[1:0], RSOP, REOP, RERR and RSX are valid. The system may deassert RENB at anytime if it is unable to accept data from the PHY device.</p> <p>When RENB is sampled low by the PHY device, a read is performed from the receive FIFO and the RDAT[31:0], RPRTY, RMOD[1:0], RSOP, REOP, RERR, RSX and RVAL signals are updated on the following rising edge of RFCLK.</p> <p>When RENB is sampled high by the PHY device, a read is not performed and the RDAT[31:0], RPRTY, RMOD[1:0], RSOP, REOP, RERR, RSX and RVAL signals will remain unchanged on the following rising edge of RFCLK.</p>

Signal Name	Direction	Function								
RDAT[31:0]	PHY to LINK	<p>Receive Packet Data Bus (RDAT[31:0]).</p> <p>The RDAT[31:0] bus carries the packet octets that are read from the receive FIFO and the in-band port address of the selected receive FIFO. RDAT[31:0] is considered valid only when RVAL is asserted.</p> <p>When a 32-bit interface is used, data must be received in big endian order on RDAT[31:0]. Given the defined data structure, bit 31 is received first and bit 0 is received last.</p> <p>When an 8-bit interface is used, the PHY supports only RDAT[7:0].</p>								
RPRTY	PHY to LINK	<p>Receive Parity (RPRTY) signal.</p> <p>The receive parity (RPRTY) signal indicates the parity calculated over the RDAT bus. When an 8-bit interface is used, the PHY only supports RPRTY calculated over RDAT[7:0].</p> <p>When RPRTY is supported, the PHY layer device must support odd parity. RPRTY is considered valid only when RVAL or RSX is asserted.</p>								
RMOD[1:0]	PHY to LINK	<p>Receive Word Modulo (RMOD) signal.</p> <p>RMOD[1:0] indicates the number of valid bytes of data in RDAT[31:0]. The RMOD bus should always be all zero, except during the last double-word transfer of a packet on RDAT[31:0]. When REOP is asserted, the number of valid packet data bytes on RDAT[31:0] is specified by RMOD[1:0].</p> <table data-bbox="743 1453 1360 1591"> <tr> <td>RMOD[1:0] = "00"</td> <td>RDAT[31:0] valid</td> </tr> <tr> <td>RMOD[1:0] = "01"</td> <td>RDAT[31:8] valid</td> </tr> <tr> <td>RMOD[1:0] = "10"</td> <td>RDAT[31:16] valid</td> </tr> <tr> <td>RMOD[1:0] = "11"</td> <td>RDAT[31:24] valid</td> </tr> </table> <p>When an 8-bit interface is used, the RMOD bus is not required. RMOD[1:0] is considered valid only when RVAL is asserted.</p>	RMOD[1:0] = "00"	RDAT[31:0] valid	RMOD[1:0] = "01"	RDAT[31:8] valid	RMOD[1:0] = "10"	RDAT[31:16] valid	RMOD[1:0] = "11"	RDAT[31:24] valid
RMOD[1:0] = "00"	RDAT[31:0] valid									
RMOD[1:0] = "01"	RDAT[31:8] valid									
RMOD[1:0] = "10"	RDAT[31:16] valid									
RMOD[1:0] = "11"	RDAT[31:24] valid									

Signal Name	Direction	Function
RSOP	PHY to LINK	<p>Receive Start of Packet (RSOP) signal.</p> <p>RSOP is used to delineate the packet boundaries on the RDAT bus. When RSOP is high, the start of the packet is present on the RDAT bus.</p> <p>RSOP is required to be present at the start of every packet and is considered valid when RVAL is asserted.</p>
REOP	PHY to LINK	<p>Receive End Of Packet (REOP) signal.</p> <p>REOP is used to delineate the packet boundaries on the RDAT bus. When REOP is high, the end of the packet is present on the RDAT bus.</p> <p>When a 32-bit interface is used, RMOD[1:0] indicates the number of valid bytes the last double-word is composed of when REOP is asserted. When an 8-bit interface is used, the last byte of the packet is on RDAT[7:0] when REOP is asserted.</p> <p>REOP is required to be present at the end of every packet and is considered valid only when RVAL is asserted.</p>
RERR	PHY to LINK	<p>Receive error indicator (RERR) signal.</p> <p>RERR is used to indicate that the current packet is in error. RERR shall only be asserted when REOP is asserted.</p> <p>Conditions that can cause RERR to be set may be, but are not limited to, FIFO overflow, abort sequence detection and FCS error.</p> <p>RERR is considered valid only when RVAL is asserted.</p>

Signal Name	Direction	Function
RSX	PHY to LINK	<p>Receive Start of Transfer (RSX) signal.</p> <p>RSX indicates when the in-band port address is present on the RDAT bus. When RSX is high, the value of RDAT[7:0] is the address of the receive FIFO to be selected by the PHY. Subsequent data transfers on the RDAT bus will be from the FIFO specified by this in-band address.</p> <p>For single port PHY devices, the RSX signal is optional as the PHY device will not need to generate in-band addresses.</p> <p>For multi-port PHY devices, RSX must be asserted at the beginning of each transfer.</p> <p>When RSX is high, RVAL must be low.</p>

7.2 Examples

The following examples are not part of the requirement definition of the POS-PHY compatibility specification. They are only informative and provide an aid in the visualization of the interface operation. The examples only present a limited set of scenarios; they are not intended to imply restrictions beyond that presented in the text of the specification. If any apparent discrepancies exist between the examples and the text, the text shall take precedence.

The POS-PHY Receive Interface is controlled by the Link Layer device using the RENB signal. All signals must be updated and sampled using the rising edge of the receive FIFO clock. The RDAT bus, RPRTY, RMOD, RSOP, REOP and RERR signals are valid in cycles for which RVAL is high and RENB was low in the previous cycle. When transferring data, RVAL is asserted and remains high until the internal FIFO of the PHY layer device is empty or an end of packet is transferred. The RSX signal is valid in the cycle for which RVAL is low and RENB was low in the previous cycle.

Figure 7.1 is an example of a multi-port PHY device with at least two channels. The PHY informs the Link Layer device of the port address of the selected FIFO by asserting RSX with the port address on the RDAT bus. The Link Layer may pause the Receive Interface at any time by deasserting the RENB signal. When the selected FIFO is empty, RVAL is deasserted. In this example, the RVAL is re-asserted, without changing the selected FIFO, transferring the last section of the packet. The end of the packet is indicated with the REOP signal. Thus, the next subsequent FIFO transfer for this port would be the start of the next packet. If an error occurred during the reception of the packet, the RERR would be asserted with REOP. Since another port's FIFO has sufficient data to initiate a bus transfer, RSX is again asserted with the port address. In this case, an intermediate section of the packet is being transferred.

Figure 7.1: Receive Logical Timing




Figure 7.2 is an example of a multi-port PHY configured to gap transfers for two clock cycles. The first transfer is a complete 3-byte packet and the second transfer is the end of a 36-byte packet. The pause allows the Link Layer device to halt the transfer of data between transfers. In order to handle an end of packet, the Link Layer device may deassert the RENB signal when it samples REOP active. As shown in the diagram, the Link Layer device pauses the PHY device on the in-band address for two clock cycles.

Figure 7.2: Receive Logical Timing with Pausing





Figure 7.3 shows an example of data transfer from a single-port PHY device, without the use of the RSX signal and in-band addressing, which are optional for single-port applications. Just before reaching end-of-packet, the PHY device happens to pause for two cycles (for internal reasons).

**Figure 7.3: Receive Logical Timing
(single-port without in-band addressing)**



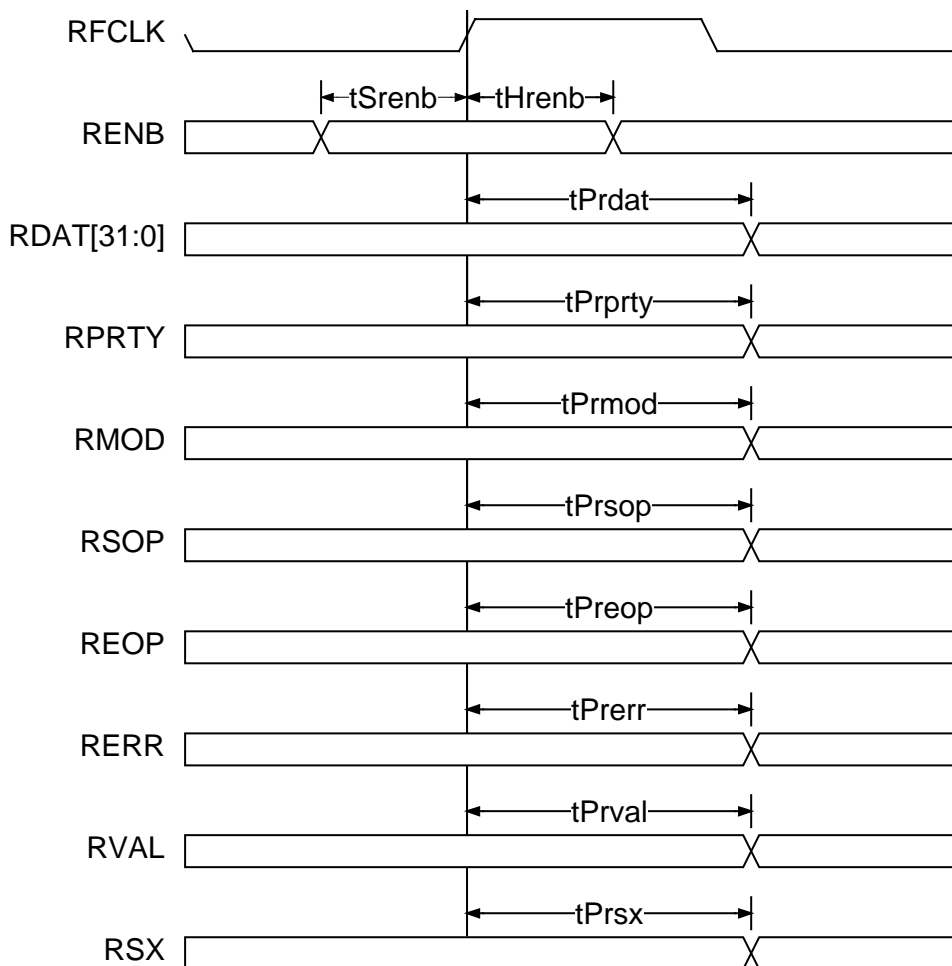
7.3 AC Timing

All AC Timing is from the perspective of the PHY layer device in a PHY-LINK interface.

Table 7.3: Receive Interface Timing

Symbol	Description	Min	Max	Units
	RFCLK Frequency		104	MHz
	RFCLK Duty Cycle	40	60	%
$t_{S_{renb}}$	RENB Set-up time to RFCLK	2.0		ns
$t_{H_{renb}}$	RENB Hold time to RFCLK	0.5		ns
t_{Prdat}	RFCLK High to RDAT Valid	1.5	6	ns
t_{Prppty}	RFCLK High to RPRTY Valid	1.5	6	ns
t_{Prsop}	RFCLK High to RSOP Valid	1.5	6	ns
t_{Preop}	RFCLK High to REOP Valid	1.5	6	ns
t_{Prmod}	RFCLK High to RMOD Valid	1.5	6	ns
t_{Prerr}	RFCLK High to RERR Valid	1.5	6	ns
t_{Prval}	RFCLK High to RVAL Valid	1.5	6	ns
t_{Prsx}	RFCLK High to RSX Valid	1.5	6	ns

Figure 7.4: Receive Physical Timing



Notes on Receive I/O Timing:

Note 1: When a set-up time is specified between an input and a clock, the set-up time is the time in nanoseconds from the 1.4 Volt point of the input to the 1.4 Volt point of the clock.

Note 2: When a hold time is specified between an input and a clock, the hold time is the time in nanoseconds from the 1.4 Volt point of the clock to the 1.4 Volt point of the input.

- Note 3: Output propagation delay time is the time in nanoseconds from the 1.4 Volt point of the reference signal to the 1.4 Volt point of the output.
- Note 4: Maximum output propagation delays are measured with a 30 pF load on the outputs.

NOTES:

None of the information contained in this document constitutes an express or implied warranty by PMC-Sierra, Inc. as to the sufficiency, fitness or suitability for a particular purpose of any such information or the fitness, or suitability for a particular purpose, merchantability, performance, compatibility with other parts or systems, of any of the products of PMC-Sierra, Inc., or any portion thereof, referred to in this document. PMC-Sierra, Inc. expressly disclaims all representations and warranties of any kind regarding the contents or use of the information, including, but not limited to, express and implied warranties of accuracy, completeness, merchantability, fitness for a particular use, or non-infringement.

In no event will PMC-Sierra, Inc. be liable for any direct, indirect, special, incidental or consequential damages, including, but not limited to, lost profits, lost business or lost data resulting from any use of or reliance upon the information, whether or not PMC-Sierra, Inc. has been advised of the possibility of such damage.

© 1997, 1998, 1999, 2000 PMC-Sierra, Inc.

PMC-1980495 Issue date: June 2000.

Printed in Canada