POS-PHY™

SATURN COMPATIBLE PACKET OVER SONET INTERFACE SPECIFICATION

FOR PHYSICAL LAYER DEVICES

(Level 2)

Issue 5: December, 1998

CONTENTS

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 was developed with the cooperation of the SATURN Development Group to cover all application bit rates up to and including 622 Mbit/s. It defines the requirements for interoperable single-PHY (one PHY layer device connects to one Link Layer device) and multi-PHY (several PHY layer devices connect 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 and the SCI-PHY Level 2 Specification were used as the basis to 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 2" 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.

2.1 POS-PHY Reference Points

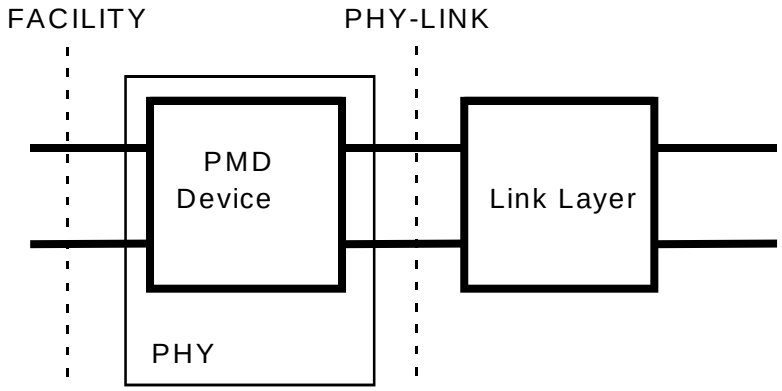

Definitions	
FACILITY:	An optical fiber, twisted pair electrical or coaxial cable electrical transmission facility.
PMD:	Physical Medium Dependant Layer.
PHY:	Physical Layer for Packet over SONET.
PHY-LINK:	Physical Layer to Link Layer electrical interface.

POS-PHY Level 2 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 2 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 table highlights the differences between the Utopia/SCI-PHY and POS-PHY interfaces.

POS-PHY Level 2 differences with ATM SCI-PHY or Utopia Level 2:

- Word based (2 bytes) data format that can accommodate variable size packets.
- Modification to the RSOC/TSOC start of cell signals to form the RSOP/TSOP (receive and transmit start of packet). These signals indicate the first word of a new packet.
- Addition of the REOP/TEOP end of packet signals which indicates the last word of a packet transfer.
- Addition of the RMOD/TMOD modulo signals which indicates if the last word of a packet carries 1 or 2 bytes.
- Addition of the RERR/TERR error signals which during the last word transfer indicates if the transferred packet must be discarded/aborted.
- Modification to the RCA/TCA cell available signals to form the selected and polled STPA, PTPA, and PRPA packet available's. The RVAL receive data valid signal was added to align data transfers with packet boundaries and to indicate FIFO empty conditions. TPA logic values are defined based on the FIFO fill level (in terms of bytes). RPA logic values are defined based on both the FIFO fill level (in terms of bytes) and the presence of an end of packet in the receive FIFO. To cope with the unpredictable nature of packet transfer, both the selected and polled PHY status are provided all the time.
- Allowance for only a 16-bit bus interface running at a maximum speed of 50 MHz.
- Definition of a byte-level and packet-level data transfer.
- In multi-PHY systems only PHY addressing is provided.
- In multi-PHY systems, PHY status indication can be provided using either a polling or a direct status indication scheme.
- Parity generation is an optional requirement.

4. Specification Summary

This section describes the POS-PHY bus operation.

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

SIGNALB Active low

4.2. Bus Widths

POS-PHY compatible devices support a 16-bit data bus width. To accommodate packets with an even or odd number of bytes, the RMOD/TMOD signals are defined and allow a single byte word transfer but only during the last word of a packet transfer.

4.3. Clock Rates

POS-PHY compatible devices can support a transfer clock rate from 25 MHz up to 50 MHz. Some devices may support multiple rates. Generally, devices targeted at single or multi-PHY applications, where the aggregate PHY bit rate is greater than 155 Mbit/s (such as 622 Mbit/s), will support a 50 MHz FIFO clock rate, where as devices targeted at applications where the aggregate PHY bit rate is less than or equal to 155 Mbit/s (such as 155 Mbit/s) will support a 33 MHz or 25 MHz FIFO clock rate. The 50 MHz clock rate is recommended.

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 devices, FIFOs are used. Control signals are provided to both the Link Layer and PHY layer devices to allow either one to exercise flow control. The packet available status granularity is byte-based with the exception of end of packets

which can trigger a transition on the receive packet available (RPA) signals, as explained below.

In the receive direction, when the PHY layer device has stored an end of packet (can be a complete small packet or the end of a larger packet) or some predefined number of bytes in its receive FIFO, it informs the Link Layer device that data is available to be read by asserting a receive packet available signal (RPA). The Link Layer device can then request data on a word by word basis from the PHY layer device by asserting an 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 signal (TPA). The Link Layer device can then write data words 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.

For both transmit and receive interfaces the Link Layer device can at any point suspend the transfer by deasserting its enable signal.

POS-PHY defines both byte-level and packet-level transfer. When doing byte level transfer, direct status indication must be used. In this case, every PHY layer device provides a separate RPA/TPA status. The Link Layer device must process all the incoming RPA/TPA and select the PHY layer devices accordingly using the transmit and receive address signals (TADR[4:0] and RADR[4:0]). With packet level transfer, the Link Layer device is able to do status polling. While TENB or RENB is asserted (data being transferred), the Link Layer device uses the PHY address busses to poll individual PHY layer devices, which all respond onto a common polled PRPA/PTPA signal. The PHY selection for the actual data transfer is performed when TENB and RENB are deasserted. Polling shall still be active when no PHY layer device is selected. 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 separate 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.

When the number of physical layer devices is limited, byte-level transfer provides a simpler implementation at a reasonable pin cost. For a large number of physical layer devices, packet-level transfer provides a lower pin count at the expense of a more

complex protocol. It is left to the system designers to select which approach best suits their needs.

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 definition.

Table 4.1. Interface Bit Rates

Standard Reference	Bit rate (Mbit/s)	Max # of PHYs (200 Mbit/s bus)	Max # of PHYs (800 Mbit/s bus)
SONET STS-1	51.84	4	12
SONET STS-3	155.52	1	4
SDH STM-1	155.52	1	4
SONET STS-12	622.08	-	1
SDH STM-4	622.08	-	1

4.6. PHY and Link Layer Interface Example

Figure 4.1 illustrates a conceptual example of how several multiple-PHY devices may be interfaced to an Link Layer device which uses direct status indication using the DRPA[x]/DTPA[x] signals. Optionally, the Link Layer device can perform multiplexed status polling using the PRPA and PTPA signals. In all cases, the address signals are used to perform PHY selection. In the example the link layer device is connected to a single package four channels PHY layer device.

Figure 4.1. POS-PHY PHY to Link Layer Interface

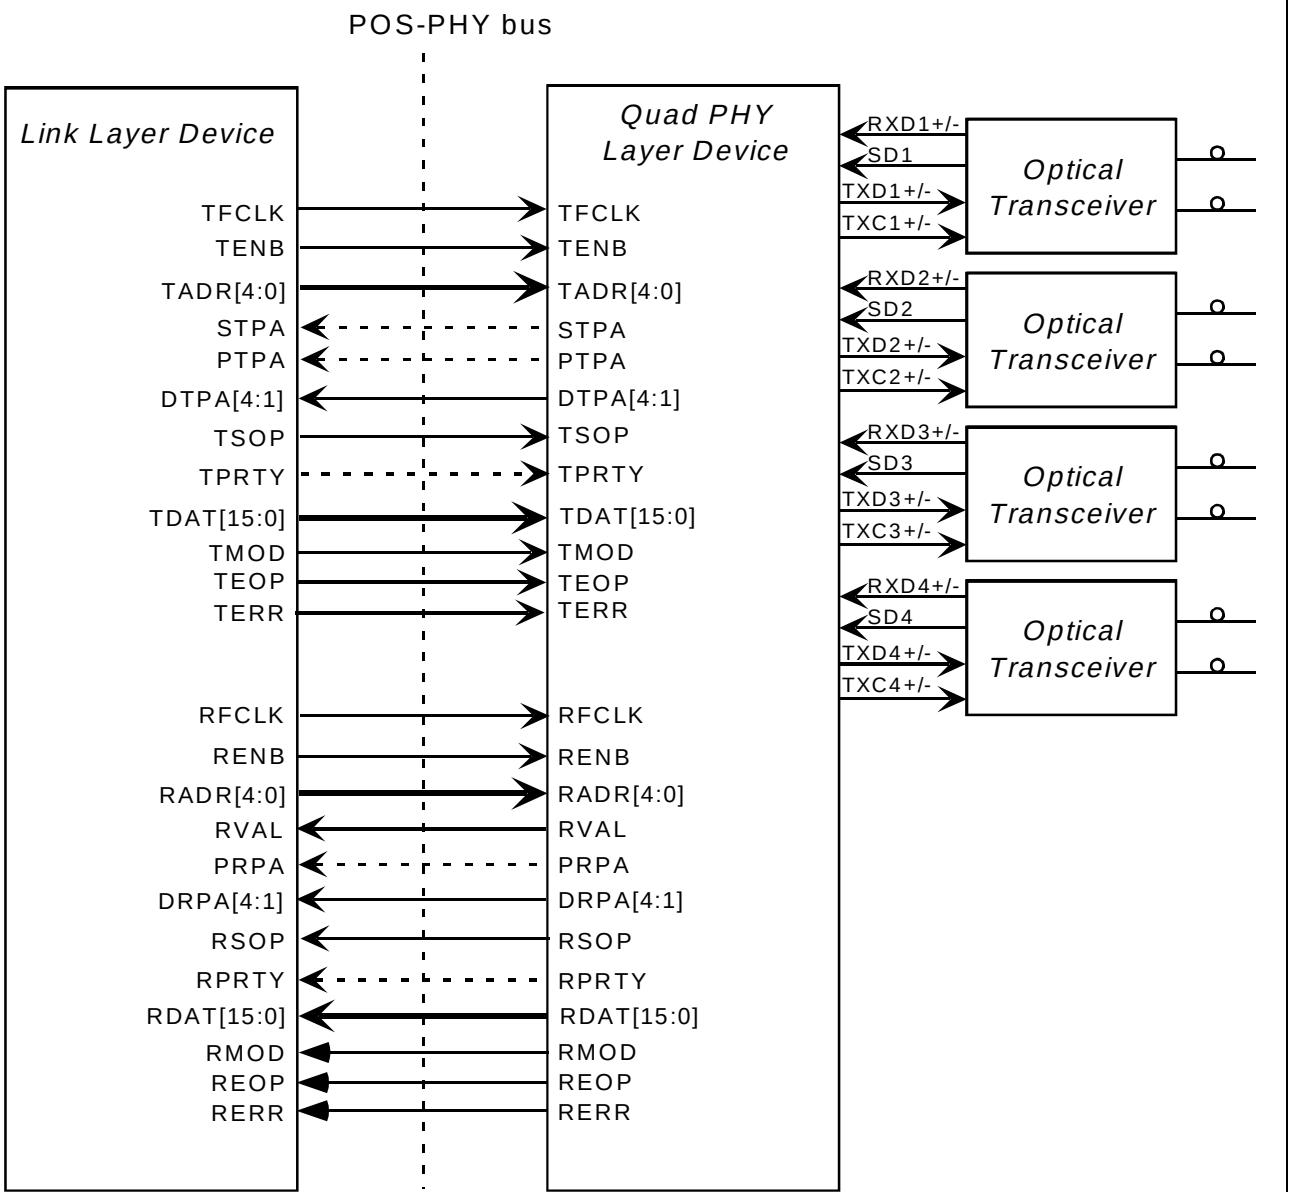

5. Interface Data Structures

Packets shall be written into the transmit FIFO and read from the receive FIFO using one 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. All words are composed of two octets, except the last word of a packet transfer which can have one or two bytes. The POS-PHY specification does not preclude the transfer of 1-byte and 2-byte packets. In this case, both start of packet and end of packet signals shall be asserted simultaneously.

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 however 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. Data Structures

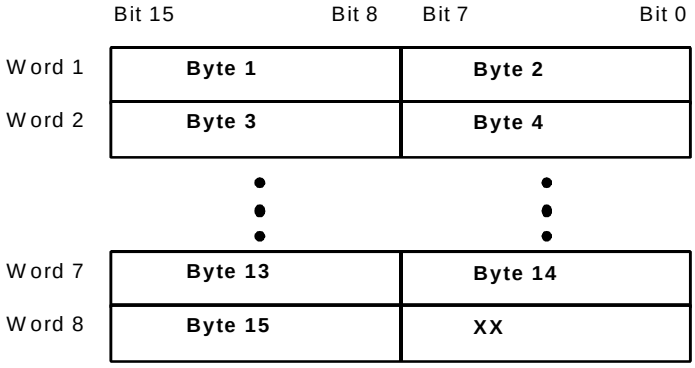

A 15 byte packet

6. Transmit Packet Interface Description

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	DIRECTION	DESCRIPTION
TDAT[15:0]	LINK to PHY	Transmit Packet Data Bus (TDAT[15:0]) bus. This bus carries the packet octets that are written to the selected transmit FIFO. TDAT[15:0] is considered valid only when TENB is simultaneously asserted. Data must be transmitted in big endian order. Given the previously defined data structure, bits are transmitted in the following order: 15, 14 … 8, 7, 6 …1, 0.
TPRTY	LINK to PHY	Transmit bus parity (TPRTY) signal. The transmit parity (TPRTY) signal indicates the parity calculated over the whole TDAT[15:0] bus. When TPRTY is supported, the PHY layer device is required to support both even and odd parity. The PHY layer device is required to report any parity error to higher layers but shall not interfere with the transferred data. TPRTY is considered valid only when TENB is asserted.

SIGNAL	DIRECTION	DESCRIPTION
TMOD	LINK to PHY	The Transmit Word Modulo (TMOD) signal. TMOD indicates the size of the current word. TMOD should always be low, except during the last word transfer of a packet, at the same time TEOP is asserted. During a packet transfer every word must be complete except the last word which can be composed of 1 or 2 bytes. TMOD set high indicates a 1-byte word (present on MSB's, LSB's are discarded), while TMOD set low indicates a 2-byte word.
TSOP	LINK to PHY	Transmit Start of Packet (TSOP) signal. TSOP indicates the first word of a packet. TSOP is required to be present at the beginning of every packet and is considered valid only when TENB is asserted.
TEOP	LINK to PHY	The active high Transmit End of Packet (TEOP) signal. TEOP marks the end of a packet on the TDAT[15:0] bus. When TEOP is high, the last word of the packet is present on the TDAT[15:0] stream and TMOD indicates how many bytes this last word is composed of. It is legal to set TSOP high at the same time TEOP is high. This provides support for one or two bytes packets, as indicated by the value of TMOD.
TERR	LINK to PHY	The transmit error indicator (TERR) signal. TERR is used to indicate that the current packet is aborted and should be discarded. TERR should only be asserted during the last word transfer of a packet.

SIGNAL	DIRECTION	DESCRIPTION
TENB	LINK to PHY	Transmit Multi-PHY Write Enable (TENB) signal. The TENB signal is an active low input which is used along with the TADR[4:0] inputs to initiate writes to the transmit FIFOs. POS-PHY supports both byte-level and packet-level transfer. Packet-level transfer operates with a selection phase when TENB is deasserted and a transfer phase when TENB is asserted. While TENB is asserted, TADR[4:0] is used for polling TPA. Byte level transfer works on a cycle basis. When TENB is asserted data is transferred to the selected PHY. Nothing happens when TENB is deasserted. Polling is not available in byte level transfer mode and direct packet availability is provided by DTPA[x].
TADR[4:0]	LINK to PHY	Transmit PHY Address (TADR[4:0]) bus. The TADR[4:0] bus is used to select the FIFO (and hence port) that is written to using the TENB signal, and the FIFO's whose packet available signal is visible on the TPA output when polling. Note that address 0x1F is the null-PHY address and shall not be identified to any port on the POS-PHY bus.

SIGNAL	DIRECTION	DESCRIPTION
STPA	PHY to LINK Packet Level Mode Only.	Selected-PHY Transmit Packet Available (STPA) signal. STPA transitions high when a predefined (normally user programmable) minimum number of bytes is available in the selected transmit FIFO (the one data is written into). Once high, STPA indicates that the transmit FIFO is not full. When STPA transitions low, it optionally indicates that the transmit FIFO is full or near full (normally user programmable). STPA always provide status indication for the selected PHY in order to avoid FIFO overflows while polling is performed. The PHY Layer device shall tristate STPA when TENB is deasserted. STPA shall also be tristated when either the null-PHY address (0x1F) or an address not matching the PHY layer device address is presented on the TADR[4:0] signals when TENB is sampled high (has been deasserted during the previous clock cycle). STPA is mandatory only if packet-level transfer mode is supported. It shall not be driven in byte-level mode, if available, since the protocol would not prevent bus contention.

SIGNAL	DIRECTION	DESCRIPTION
PTPA	PHY to LINK Packet Level Mode Only.	Polled-PHY Transmit Packet Available (PTPA) signal. PTPA transitions high when a predefined (normally user programmable) minimum number of bytes is 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). PTPA allows to poll the PHY address selected by TADR[4:0] when TENB is asserted. PTPA is driven by a PHY layer device when its address is polled on TADR[4:0]. A PHY layer device shall tristate PTPA when either the null-PHY address (0x1F) or an address not matching available PHY Layer devices is provided on TADR[4:0]. PTPA is mandatory only if packet-level transfer mode is supported. It shall not be driven in byte-level mode, if available, since the protocol would then not prevent bus contention.
DTPA[x]	PHY to LINK Byte Level Mode Only.	Direct Transmit Packet Available (DTPA[x]). The DTPA[x] signals provide direct status indication for the corresponding port (referred to by the index “x”). DTPA[x] transitions high when a predefined (normally user programmable) minimum number of bytes is available in the transmit FIFO. Once high, DTPA[x] indicates that the transmit FIFO is not full. When DTPA[x] transitions low, it optionally indicates that the transmit FIFO is full or near full (normally user programmable). DTPA[x] is mandatory only if byte-level transfer mode is supported. It is optional in packet-level mode.

SIGNAL	DIRECTION	DESCRIPTION
TFCLK	LINK to PHY	Transmit FIFO Write Clock (TFCLK). TFCLK is used to synchronize data transfer transactions from the LINK Layer device to the PHY layer device. TFCLK may cycle at any rate from 25 MHz up to 50 MHz.

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. The logical timing is valid for both single-PHY and multi-PHY configurations.

Byte level transfer can be used for both single-PHY and multi-PHY applications, as shown in Figure 6.1. The PHY layer device indicates that the FIFO is not full by asserting the 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 TENB low and present the word on the TDAT bus. If the presented word is the first word of a packet, the Link Layer device must also assert signal TSOP. At any time, if the Link Layer device does not have a word to write, it can dessert TENB.

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 PHY layer device. In this example the predefined value is one word or two bytes. If the Link Layer writes more than that predefined number of words and DTPA remains deasserted throughout, the PHY layer device will indicate an error condition and ignore additional writes until it asserts DTPA again.

Figure 6.1. Single-PHY and Byte Level-Multi-PHY Transmit Logical Timing

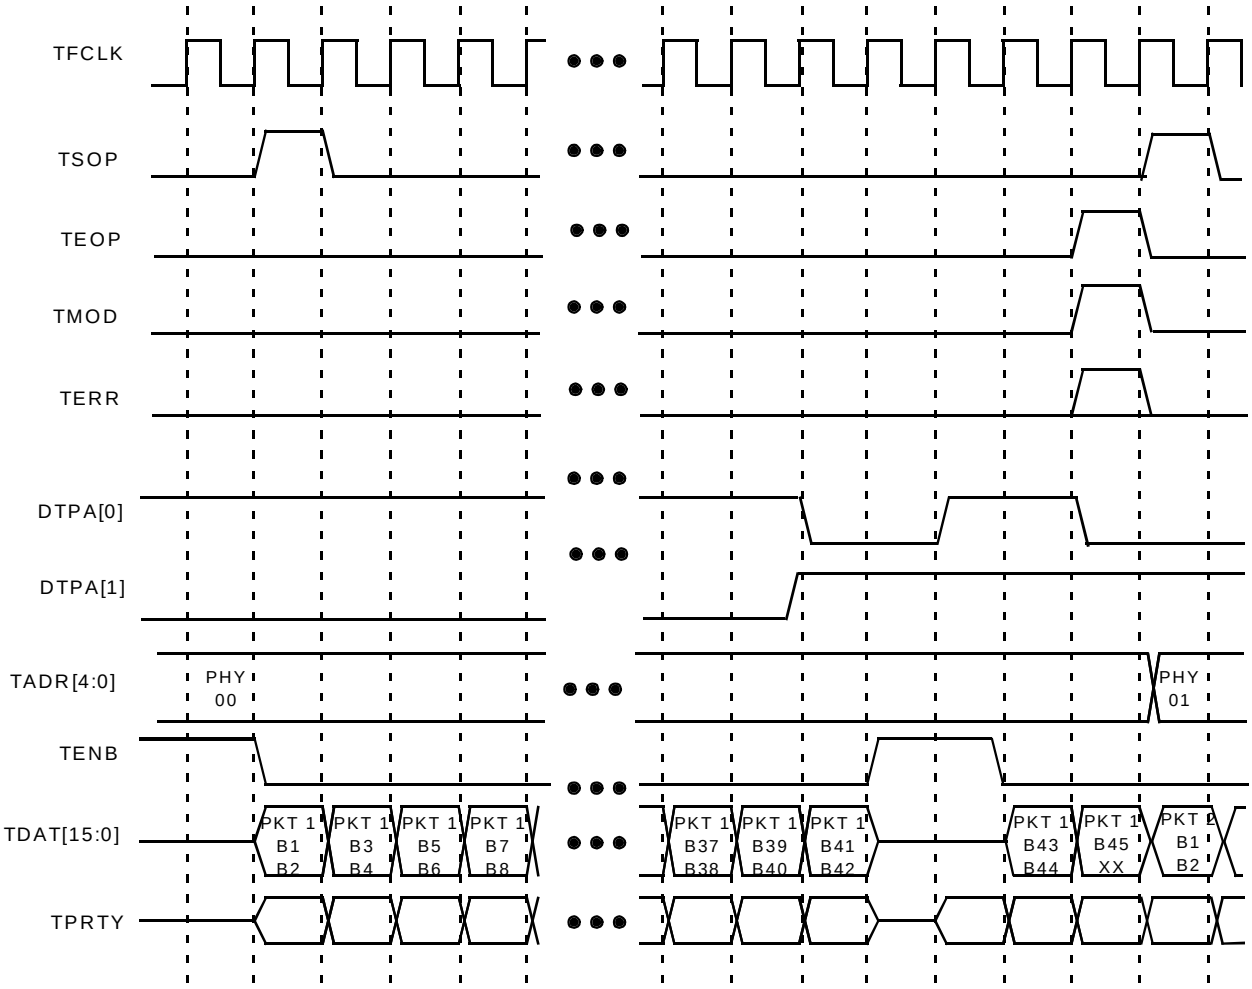

Figure 6.2 is an example of a multi-PHY packet-level transfer, with a polling and a selection sequence. "A", "B" and "C" represent any arbitrary address value; the Link Layer device is not restricted in its polling order. The PHY associated with address "A" indicates it cannot accept data; whereas, PHY "B" indicates its willingness to accept data. As a result, the Link Layer places address "B" on TADR[4:0] the cycle before TENB is asserted to select PHY "B" as the next packet destination. In this example, the PHY "C" status is ignored; the Link Layer device is not constrained to select the latest PHY polled. As soon as the packet transfer is started, the polling process may be recommenced. The 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 as indicated by TENB asserted. In this example, the PHY layer device shown was configured to

deassert STPA two words (4 bytes) before the FIFO is empty. The system could be configured differently.

Figure 6.2. Packet-Level Multi-PHY Addressing Transmit Logical Timing

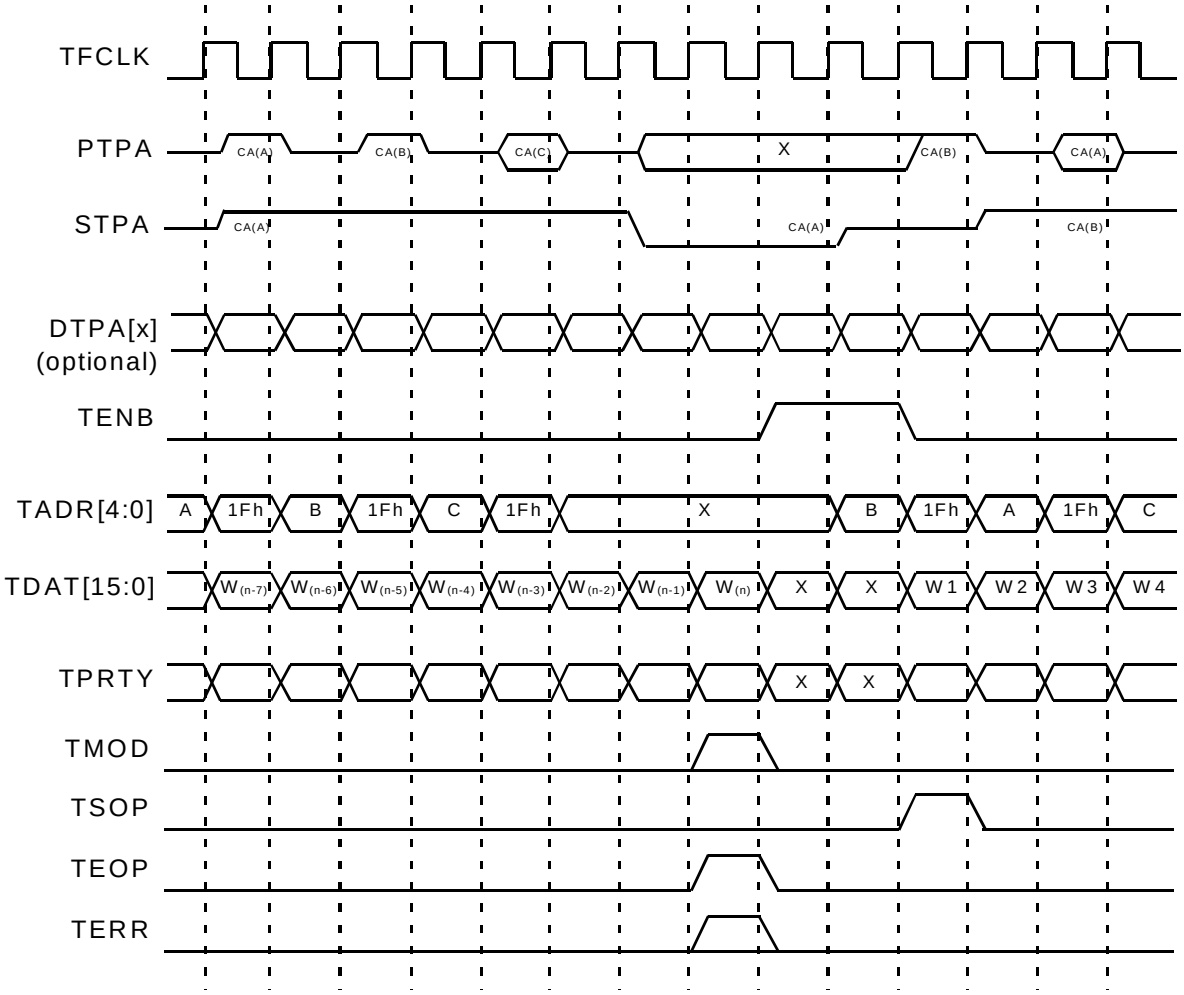

6.3. AC Timing

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

Table 6.2. Transmit Interface Timing, ≤50 MHz

Symbol	Description	Min	Max	Units
	TFCLK Frequency	25	50	MHz
	TFCLK Duty Cycle	40	60	%
tS_{TENB}	TENB Set-up time to TFCLK	4		ns
tH_{TENB}	TENB Hold time to TFCLK	0		ns
tS_{TDAT}	TDAT[15:0] Set-up time to TFCLK	4		ns
tH_{TDAT}	TDAT[15:0] Hold time to TFCLK	0		ns
tS_{TPRTY}	TPRTY Set-up time to TFCLK	4		ns
tH_{TPRTY}	TPRTY Hold time to TFCLK	0		ns
tS_{TSOP}	TSOP Set-up time to TFCLK	4		ns
tH_{TSOP}	TSOP Hold time to TFCLK	0		ns
tS_{TEOP}	TEOP Set-up time to TFCLK	4		ns
tH_{TEOP}	TEOP Hold time to TFCLK	0		ns
tS_{TMOD}	TMOD Set-up time to TFCLK	4		ns
tH_{TMOD}	TMOD Hold time to TFCLK	0		ns
tS_{TERR}	TERR Set-up time to TFCLK	4		ns
tH_{TERR}	TERR Hold time to TFCLK	0		ns
tS_{TADR}	TADR[4:0] Set-up time to TFCLK	4		ns
tH_{TADR}	TADR[4:0] Hold time to TFCLK	0		ns
tP_{STPA}	TFCLK High to STPA Valid	1	12	ns
tZ_{STPA}	TFCLK High to STPA Tristate	1	10	ns
tZB_{SPA}	TFCLK High to STPA Driven	0		ns
tP_{PTPA}	TFCLK High to PTPA Valid	1	12	ns

tZ_{PTPA}	TFCLK High to PTPA Tristate	1	10	ns
tZB_{PTPA}	TFCLK High to PTPA Driven	0		ns
tP_{DTPA}	TFCLK High to DTPA[x] Valid	1	12	ns

* These parameters only apply if the PHY device supports address decoding internally.

Figure 6.3. Transmit Physical Timing

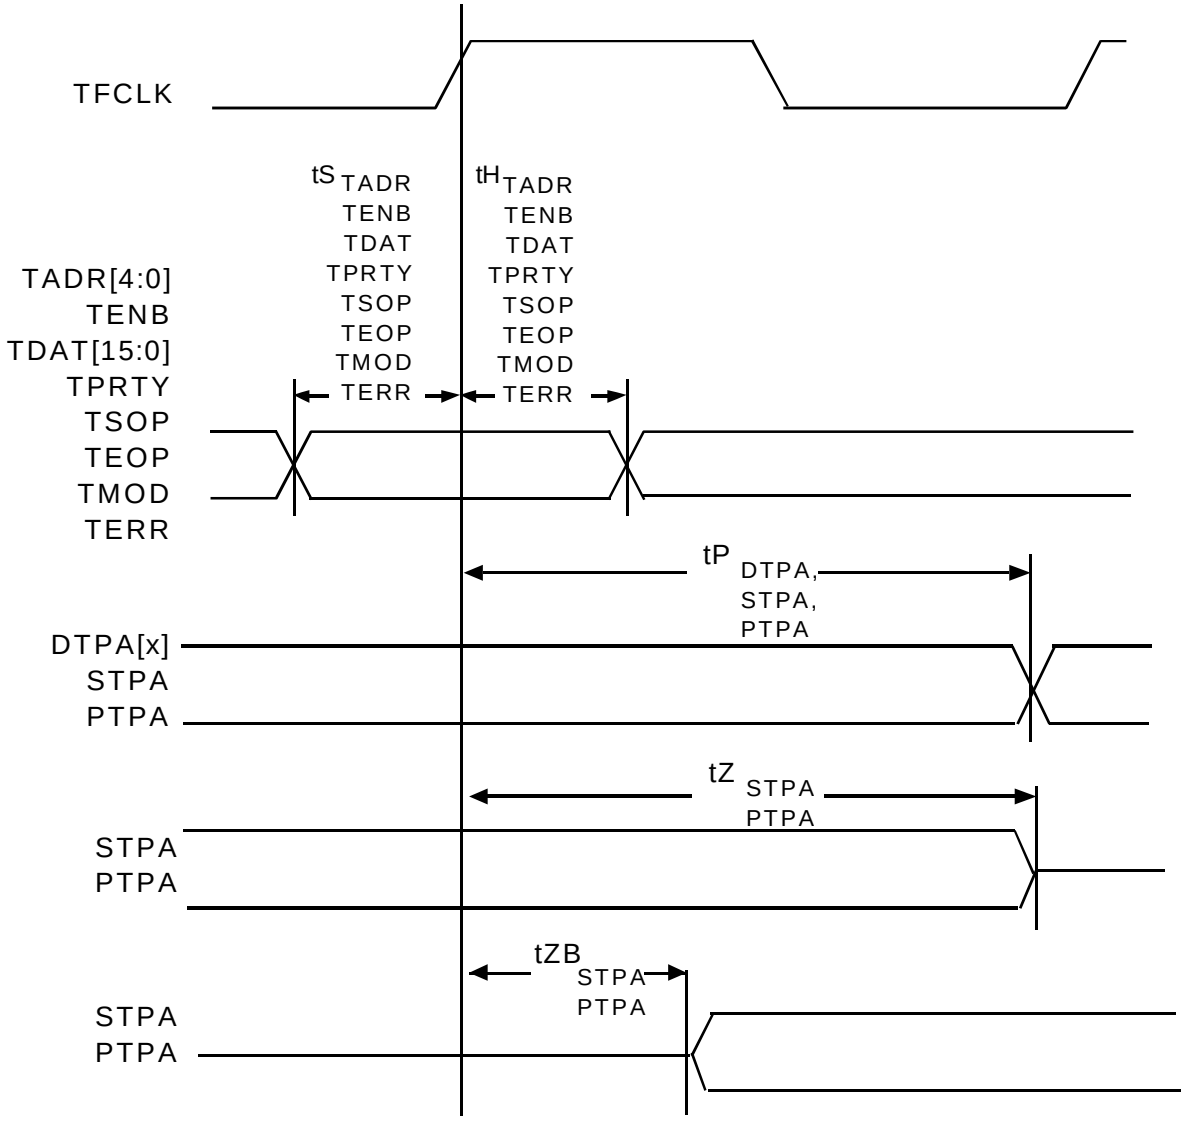

Notes on Transmit I/O Timing:

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. Transmit FIFO Configuration

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 deassertion of the TPA 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.

The PHY layer device shall not, in the transmit direction, 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.

8. Receive Packet Interface Description

8.1. Receive Signals

Table 8.1 lists the receive side POS-PHY specification signals for single-PHY and multi-PHY applications. 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 a 256-byte receive FIFO.

Table 8.1. Receive Signal Descriptions

SIGNAL	DIRECTION	DESCRIPTION
RDAT[15:0]	PHY to LINK	Receive Packet Data Bus (RDAT[15:0]). The RDAT[15:0] bus carries the packet octets that are read from the selected receive FIFO. RDAT[15:0] is considered valid only when RENB is simultaneously asserted and a valid PHY layer device has been selected via the RADR[4:0] signals. Data must be received in big endian order. Given the defined data structure, bits are received in the following order: 15, 14 … 8, 7, 6 …1, 0. The PHY layer device shall tristate RDAT[15:0] when RENB is high. RDAT[15:0] shall also be tristated when either the null-PHY address (0x1F) or an address not matching the PHY layer device address is presented on the RADR[4:0] signals when RENB is sampled high (has been deasserted during the previous clock cycle).

SIGNAL	DIRECTION	DESCRIPTION
RPRTY	PHY to LINK	Receive Parity (RPRTY) signal. The receive parity (RPRTY) signal indicates the parity of the RDAT bus. When RPRTY is supported, the PHY layer device must support both odd and even parity. The PHY Layer device shall tristate RPRTY when RENB is high. RPRTY shall also be tristated when either the null-PHY address (0x1F) or an address not matching the PHY layer device address is presented on the RADR[4:0] signals when RENB is sampled high (has been deasserted during the previous clock cycle).
RMOD	PHY to LINK	The Receive Word Modulo (RMOD) signal. RMOD indicates the size of the current word. RMOD is only used during the last word transfer of a packet, at the same time REOP is asserted. During a packet transfer every word must be complete except the last word which can be composed of 1 or 2 bytes. RMOD set high indicates a 1-byte word (present on MSB's, LSB's are discarded) while RMOD set low indicates a 2-byte word. The PHY Layer device shall tristate RMOD when RENB is high. RMOD shall also be tristated when either the null-PHY address (0x1F) or an address not matching the PHY layer device address is presented on the RADR[4:0] signals when RENB is sampled high (has been deasserted during the previous clock cycle).

SIGNAL	DIRECTION	DESCRIPTION
RSOP	PHY to LINK	Receive Start of Packet (RSOP) signal. RSOP marks the first word of a packet transfer. The PHY layer device must assert RSOP for every packet. The PHY Layer device shall tristate RSOP when RENB is high. RSOP shall also be tristated when either the null-PHY address (0x1F) or an address not matching the PHY layer device address is presented on the RADR[4:0] signals when RENB is sampled high (has been deasserted during the previous clock cycle).
REOP	PHY to LINK	The Receive End Of Packet (REOP) signal. REOP marks the end of packet on the RDAT[15:0] bus. During this same cycle RMOD is used to indicate if the last word has 1 or 2 bytes. It is legal to set RSOP high at the same time REOP is high. This provides support for one or two bytes packets, as indicated by the value of RMOD. The PHY Layer device shall tristate REOP when RENB is high. REOP shall also be tristated when either the null-PHY address (0x1F) or an address not matching the PHY layer device address is presented on the RADR[4:0] signals when RENB is sampled high (has been deasserted during the previous clock cycle).

SIGNAL	DIRECTION	DESCRIPTION
RERR	PHY to LINK	The Receive error indicator (RERR) signal. RERR is used to indicate that the current packet is aborted and should be discarded. RERR can only be asserted during the last word transfer of a packet. Conditions that can cause RERR to be set may be, but are not limited to, FIFO overflow, abort sequence detection, FCS error. The PHY Layer device shall tristate RERR when RENB is high. RERR shall also be tristated when either the null-PHY address (0x1F) or an address not matching the PHY layer device address is presented on the RADR[4:0] signals when RENB is sampled high (has been deasserted during the previous clock cycle).
RENB	LINK to PHY	Receive multi-PHY Read Enable (RENB) signal. The RENB signal is used to initiate reads from the receive FIFO's. The POS-PHY specification supports both byte-level and packet-level transfer. Packet-level transfer operates with a selection phase when RENB is deasserted and a transfer phase when RENB is asserted. While RENB is asserted, RADR[4:0] is used for polling RPA. Byte level transfer works on a cycle basis. When RENB is asserted data is transferred from the selected PHY and RADR[4:0] is used to select the PHY. Nothing happens when RENB is deasserted. In byte-level transfer mode polling is not possible; packet availability is directly indicated by DRPA[x]. RENB must operate in conjunction with RFCLK to access the FIFO's at a high enough rate to prevent FIFO overflows. The system may de-assert RENB at anytime it is unable to accept another byte.

SIGNAL	DIRECTION	DESCRIPTION
RADR[4:0]	LINK to PHY	Receive Read Address (RADR[4:0]) signals. The RADR[4:0] signal is used to select the FIFO (and hence port) that is read from using the RENB signal. For packet-level transfer, RADR[4:0] is also used to determine the FIFO's whose packet available signal is polled on the PRPA output. Address 0x1F is the null-PHY address and shall not be responded to by any PHY layer device.
RVAL	PHY to LINK .	Receive Data Valid (RVAL) signal. RVAL indicates the validity of the receive data signals. When RVAL is high, the Receive signals (RDAT, RSOP, REOP, RMOD, RXPRTY and RERR) are valid. When RVAL is low, all Receive signals are invalid and must be disregarded. RVAL will transition low on a FIFO empty condition or on an end of packet. . No data will be removed from the receive FIFO while RVAL is deasserted. Once deasserted, RVAL will remain deasserted until current PHY has been deselected. RVAL allows to monitor the selected PHY during a data transfer, while monitoring or polling other PHY's is done using PRPA or DRPA[x]. The PHY Layer device shall tristate RVAL when RENB is deasserted. RVAL shall also be tristated when either the null-PHY address (0x1F) or an address not matching the PHY layer device address is presented on the RADR[4:0] signals when RENB is sampled high (has been deasserted during the previous clock cycle)..

SIGNAL	DIRECTION	DESCRIPTION
PRPA	PHY to LINK Packet Level Mode Only.	Receive Polled multi-PHY packet Available (PRPA) signal. PRPA indicates when data is available in the polled receive FIFO. When PRPA is high, the receive FIFO has at least one end of packet or a predefined number of bytes to be read (the number of bytes might be user programmable). PRPA is low when the receive FIFO fill level is below the assertion threshold and the FIFO contains no end of packet. PRPA allows to poll every PHY while transferring data from the selected PHY. PRPA is driven by a PHY layer device when its address is polled on RADR[4:0]. A PHY layer device shall tristate PRPA when either the null-PHY address (0x1F) or an address not matching available PHY Layer devices is provided on RADR[4:0]. PRPA is mandatory only if packet-level transfer mode is supported. It shall not be driven in byte-level mode, if available, since the protocol would then not prevent contention.
DRPA[x]	PHY to LINK Byte Level Mode Only.	Receive Packet Available (DRPA[x]) Direct Status Indication signals. These signals provides direct status indication for the corresponding port (referred to by the index “x”). DRPA[x] indicates when data is available in the receive FIFO. When DRPA[x] is high, the receive FIFO has at least one end of packet or a predefined number of bytes to be read. The number of bytes is usually user programmable. DRPA[x] is low when the receive FIFO fill level is below the assertion thresholdand the FIFO contains no end of packet.. DRPA[x] is mandatory only if byte-level transfer mode is supported. It is optional in packet-level mode.

SIGNAL	DIRECTION	DESCRIPTION
RFCLK	LINK to PHY	Receive FIFO Write Clock (RFCLK). RFCLK is used to synchronize data transfer transactions from the LINK Layer device to the PHY layer device. RFCLK may cycle at a any rate from 25 MHz up to 50 MHz.

8.2. 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, RFCLK. The logical timing is valid for both single-PHY and multi-PHY operating modes.

In general, PHY layer device indicates that a complete packet or some predefined amount of data is available by asserting the receive packet available signal, PRPA or DRPA[x]. 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. Because of the nature of packet transfer, it is not possible to provide a lookahead capability for RPA or RVAL. A consequence is that the Link layer device will sometimes attempt unsuccessful read cycle while the FIFO is empty or following the transfer of an end of packet (REOP). The RVAL signal shall always be used as a data valid indicator. The Link Layer device indicates, by asserting the RENB signal, that the data on the RDAT bus during the next RFCLK cycle will be read from the PHY layer device.

Figure 8.1 illustrates the single-PHY operation. DRPA transitions low when the last word of the last packet is available on the RDAT bus. The RDAT bus, RPRTY, RMOD, RSOP, REOP and RERR are valid in cycles for which RVAL is high and RENB was low is the previous cycle. If the Link Layer device requests a read while RVAL is deasserted, the PHY layer device will ignore the additional reads. Figure 8.1 identifies a dead cycle which occurs when RVAL transitions low. Since the LINK layer device is not guaranteed a fixed number of bytes when it starts reading data, it will not in general be able to deassert RENB after the last data word is read. In general, the link layer will attempt an extra read cycle, for which RVAL will be sampled low, indicating that the data is not valid. After this dead cycle the link layer device should deassert RENB although it is not forced to.

Figure 8.1. Single PHY Receive Logical Timing

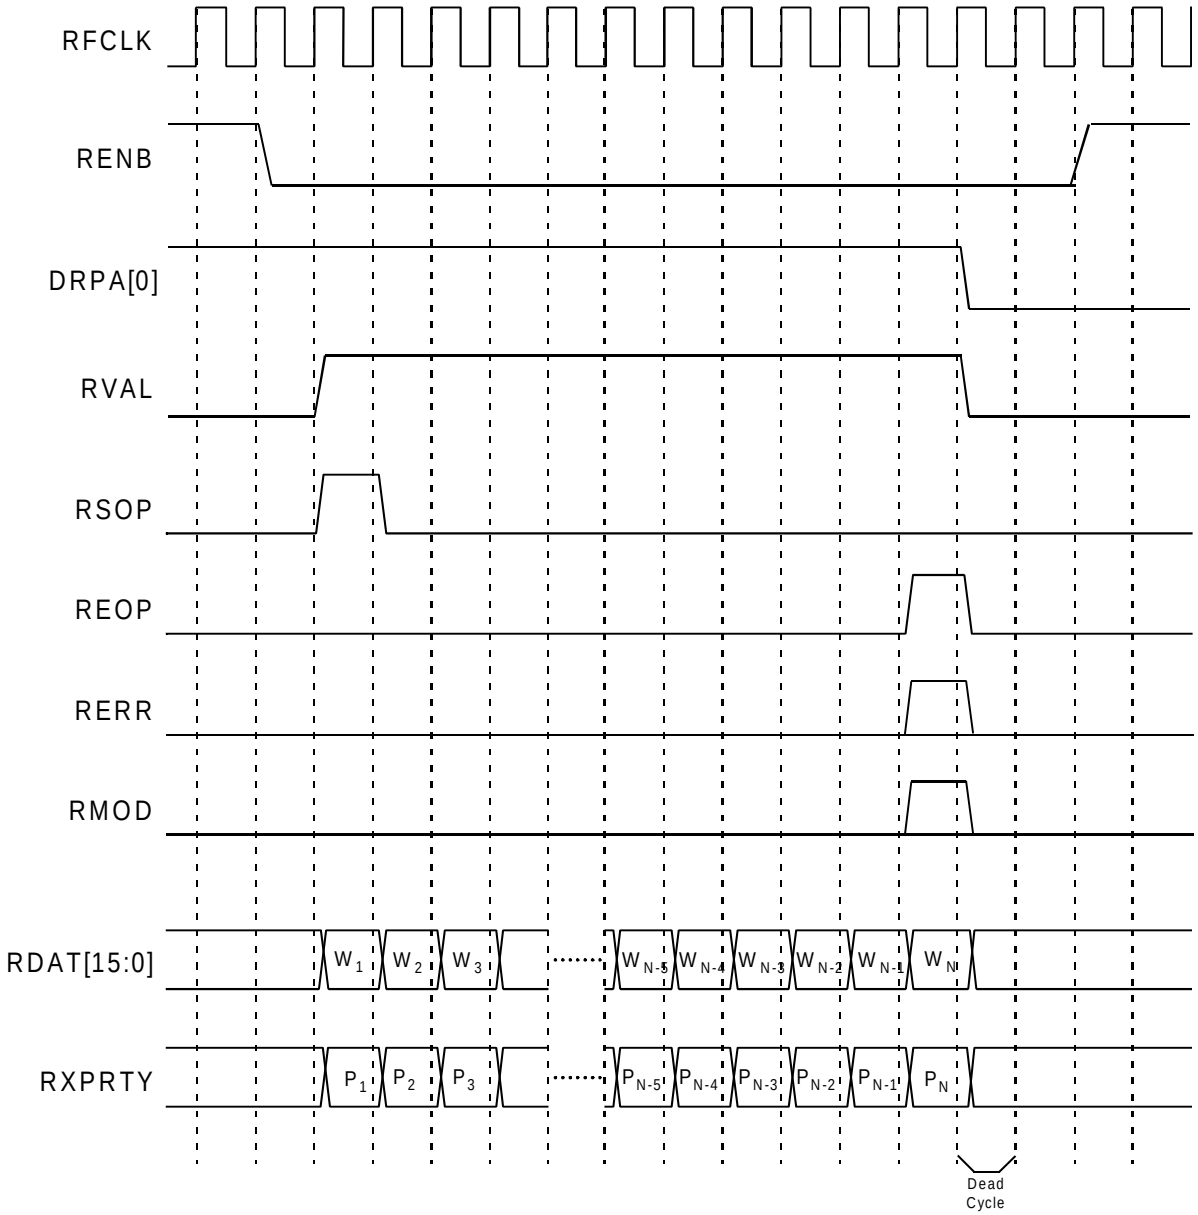

Figure 8.2 illustrates the byte-level multi-PHY operation. Independent DRPA[x] provide status indication for every receive FIFO. In this example we assume that the system consist of 4 logical PHYs implemented using two separate IC's. PHYs #1 and #2 are built within the same IC, as are PHYs #3 and #4. The transition between PHY #1 and PHY #2 occurs without any lost bus cycle. It should be noticed that the addition of a dead cycle would be required when the LINK layer transitions between PHY #2 and #3. This dead cycle is required to avoid contention on the bus when one IC get tristated and the other IC starts driving. It is left to the system implementers to determine if this dead cycle is required or not. Having this dead cycle built into the protocol would provide a more generic interface, at the expense of one lost clock cycle for every RADR[4:0] change.

Figure 8.2. Byte Level Multi-PHY Logical Timing

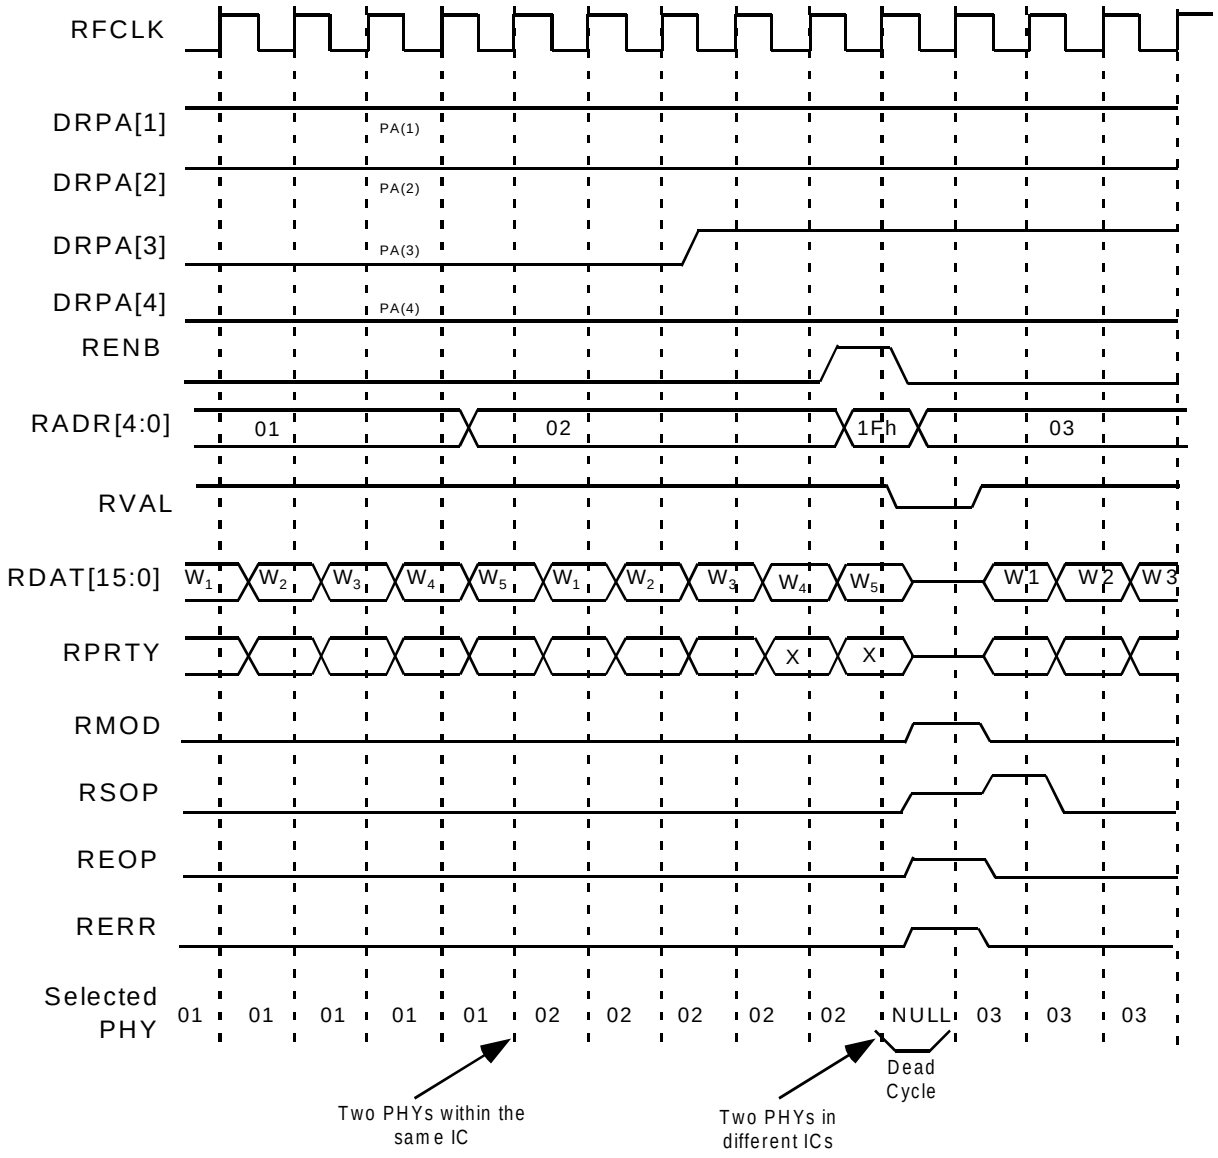

Figure 8.3 is an example of a multi-PHY polling and selection sequence. "A", "B", "C", "D" and "E" represent any arbitrary address values; the Link Layer device is not restricted in its polling order. The PHY associated with address "A" indicates it has a packet available; as well as PHY “B”. As a result, the Link Layer places address "B" on RADR[4:0] the cycle before RENB is asserted to select PHY "B" as the next packet source. In this example, the PHY "C" status is ignored; the Link Layer device is not constrained to select the latest PHY polled. As soon as the packet transfer is started, the polling process may be recommenced.

During multi-PHY operation several PHY layer devices share the RDAT[15:0], RSOP, REOP, RMOD and RPRTY signals. As a result, these signals must be tri-stated in all PHY devices which have not been selected for reading by the Link Layer. Selection of which PHY layer device is being read is made by the value on RADR[4:0] the cycle before RENB is asserted.

Figure 8.3. Packet Level Multi-PHY Logical Timing

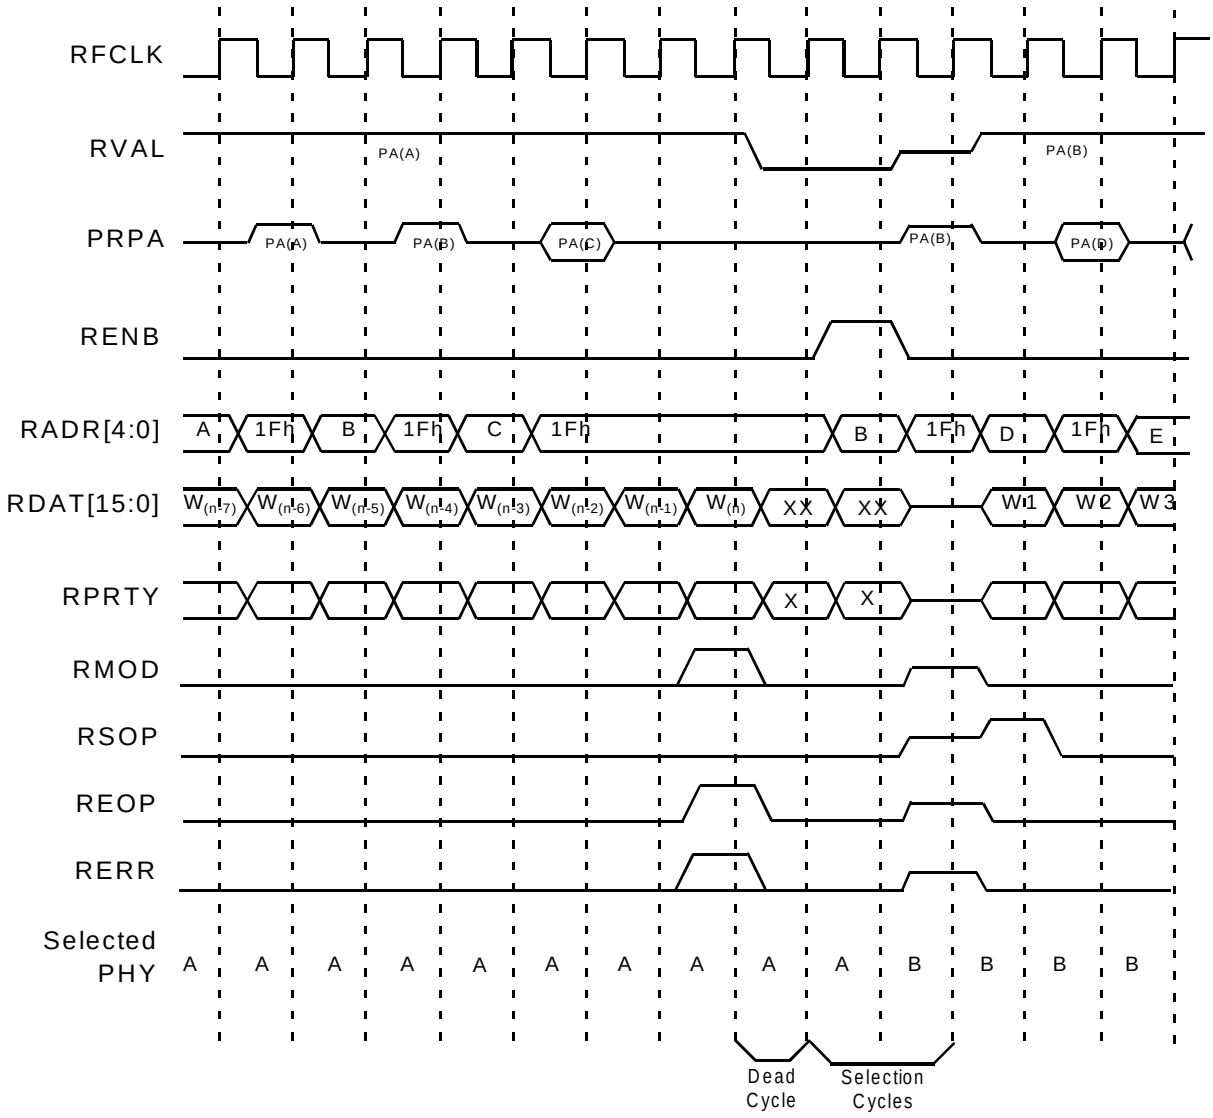

Figure 8.4 illustrates the difference between RPA and RVAL. This illustration presents two transfers. During the first transfer, DRPA[1] get deasserted after the FIFO fill level drop below the programmed 16 bytes threshold. However RVAL remains asserted indicating that there is valid data being transferred. Eventually the

FIFO runs empty and RVAL is deasserted. The LINK Layer device responds by deasserting RENB and waits until DRPA[1] is reasserted, after which it starts the second transfer. In this case, the opposite situation occurs. The FIFO still contains data when an end of packet is encountered and RVAL is deasserted. While RVAL is deasserted, the PHY layer device does not transfer any data. At this point, one option would be that the LINK layer device deasserts RENB and reselects the PHY #1, starting the transfer of a new packet. This example illustrates how RVAL and RPA allow to align data transfers with packet boundaries, at the expense of a few wasted bus cycles.

Figure 8.4. Byte Level Illustration of RVAL and RPA

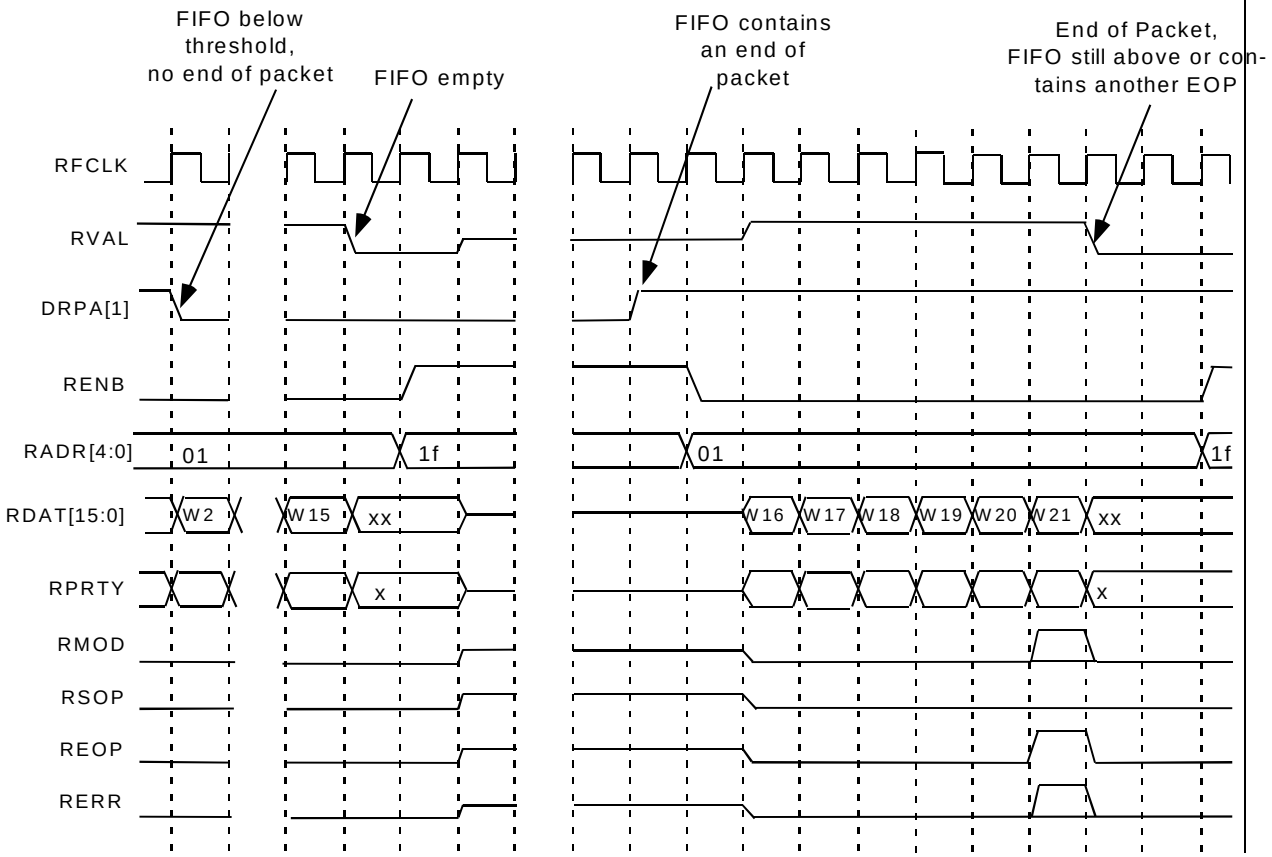

8.3. AC Timing

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

Table 8.3 Receive Interface Timing, ≤50 MHz

Symbol	Description	Min	Max	Units
	RFCLK Frequency	25	50	MHz
	RFCLK Duty Cycle	40	60	%
tS_{RENB}	RENB Set-up time to RFCLK	4		ns
tH_{RENB}	RENB Hold time to RFCLK	0		ns
tS_{RADR}	RADR[4:0] Set-up time to RFCLK	4		ns
tH_{RADR}	RADR[4:0] Hold time to RFCLK	0		ns
tP_{RDAT}	RFCLK High to RDAT Valid	1	12	ns
tZ_{RDAT}	RFCLK High to RDAT Tri-state	1	12	ns
tZB_{RDAT}	RFCLK High to RDAT Driven	0		ns
tP_{RPRTY}	RFCLK High to RPRTY Valid	1	12	ns
tZ_{RPRTY}	RFCLK High to RPRTY Tri-state	1	12	ns
tZB_{RPRTY}	RFCLK High to RPRTY Driven	0		ns
tP_{RSOP}	RFCLK High to RSOP Valid	1	12	ns
tZ_{RSOP}	RFCLK High to RSOP Tri-state	1	12	ns
tZB_{RSOP}	RFCLK High to RSOP Driven	0		ns
tP_{REOP}	RFCLK High to REOP Valid	1	12	ns
tZ_{REOP}	RFCLK High to REOP Tri-state	1	12	ns
tZB_{REOP}	RFCLK High to REOP Driven	0		ns
tP_{RMOD}	RFCLK High to RMOD Valid	1	12	ns
tZ_{RMOD}	RFCLK High to RMOD Tri-state	1	12	ns
tZB_{RMOD}	RFCLK High to RMOD Driven	0		ns

Table 8.3 Receive Interface Timing, ≤50 MHz (Continue)

tP_{RERR}	RFCLK High to RERR Valid	1	12	ns
tZ_{RERR}	RFCLK High to RERR Tri-state	1	12	ns
tZB_{RERR}	RFCLK High to RERR Driven	0		ns
tP_{RVAL}	RFCLK High to RVAL Valid	1	12	ns
tZ_{RVAL}	RFCLK High to RVAL Tri-state	1	12	ns
tZB_{RVAL}	RFCLK High to RVAL Driven	0		ns
tP_{PRPA}	RFCLK High to PRPA Valid	1	12	ns
tZ_{PRPA}	RFCLK High to PRPA Tri-state	1	12	ns
tZB_{PRPA}	RFCLK High to PRPA Driven	0		ns
tP_{DRPA}	RFCLK High to DRPA[x] Valid	1	12	ns

Figure 8.4. Receive Physical Timing

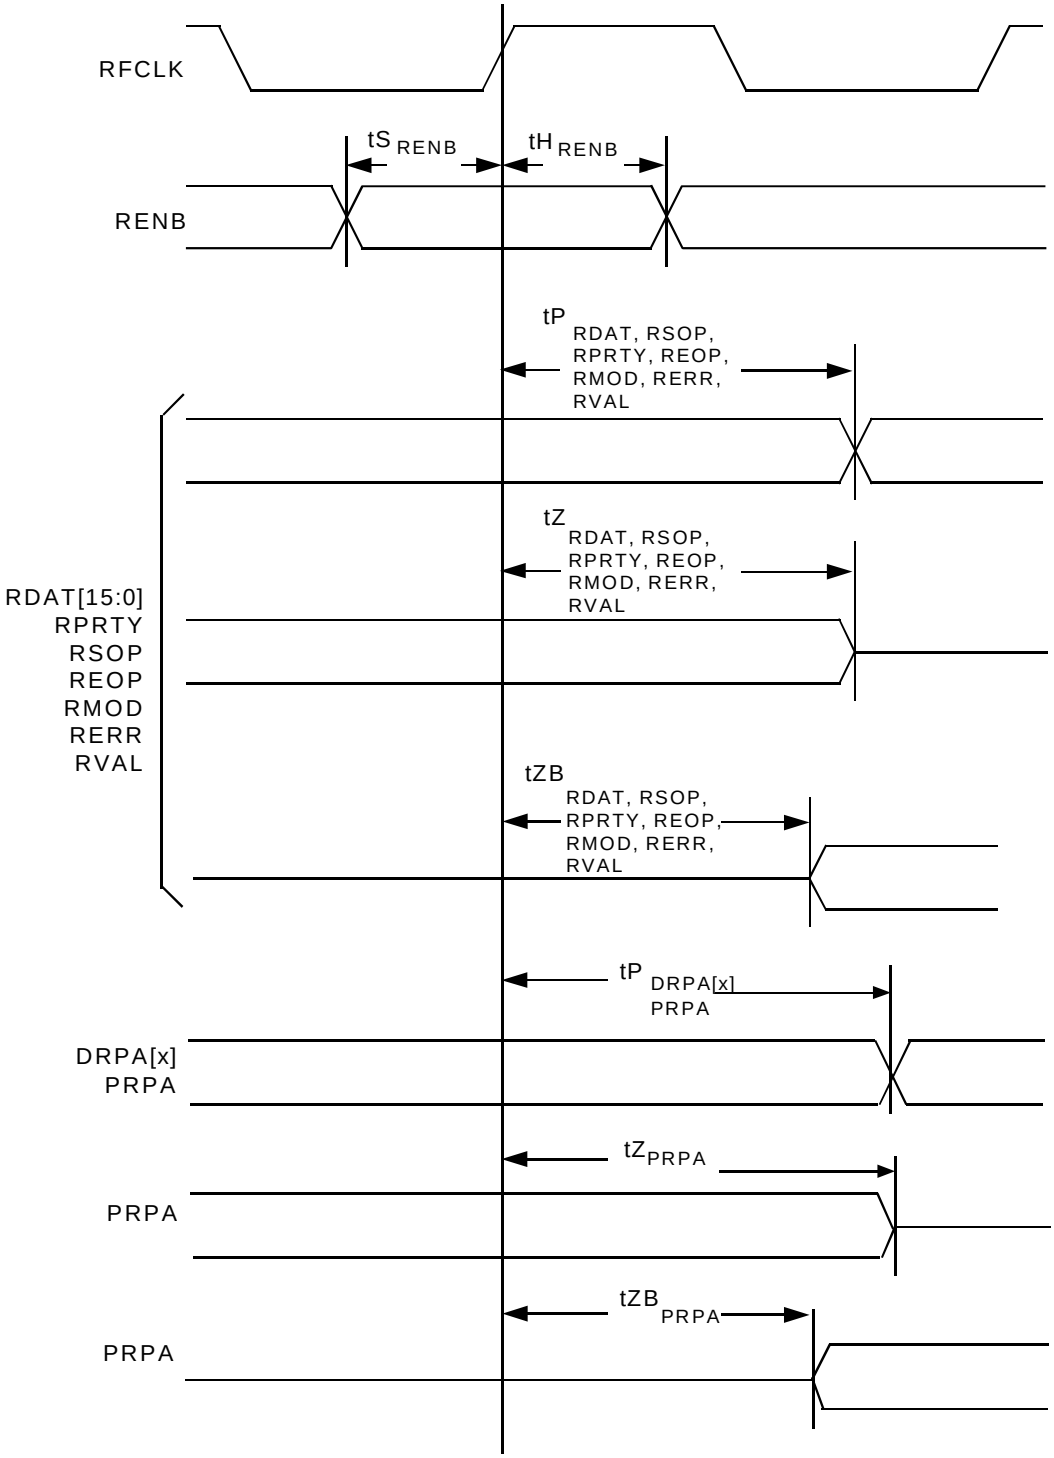

Notes on Receive I/O Timing:

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.

9. Receive FIFO Configuration

The standard FIFO depth for POS-PHY interfaces is 256 octets. The receive buffer shall deassert the RVAL flag once the FIFO is empty or an end of packet is transferred. Due to the variable size nature of packets, it is not practical to provide a lookahead capability on RPA signals. RPA is deasserted when an end of packet is transferred in order to allow the LINK Layer device to align data transfers with packet boundaries, and thus facilitate buffer management. When the PHY layer device provide a programmable RPA assertion fill level, it will be possible to tune the burst transfer size to specific system characteristics.

In addition the PHY layer device shall assert RPA whenever an end of packet is stored in the receive FIFO. This requirement guarantees that small packets or short packet ends, too small to exceed the normal RPA assertion level, will be read from the receive FIFO in a timely manner, providing a better control of the FIFO latency.

10. Remarks About Packet-Level and Byte-Level Transfer Applications

The byte-level transfer mode is intended to simplify the bus protocol and improve throughput by avoiding the PHY selection cycles required in packet-level transfer mode. Skipping the PHY selection cycle will work reliably only if the POS-PHY bus is a point to point bus, that is connecting a single Link Layer device to a single PHY Layer device. That does not preclude multi-PHY applications of the byte-level transfer mode, but it will be usable only if the multiple PHY layer ports are within a single integrated circuit. As an alternative, the system integrator can build the Link Layer device such that it forces the Null PHY address for one cycle whenever TADR[4:0] or RADR[4:0] changes, inserting a single dead cycle during which the bus is tristated. In any case, although more complex, packet-level transfer seems to offer a more reliable mode of operation when multiple PHY's are implemented within several integrated circuits. Also, the packet-level transfer configuration scales nicely when the number of PHY increases, while the byte-level transfer configuration requires additional signals.

From that perspective, byte-level and packet-level transfer are complementary, although an implementer will have to choose which one better suits his needs. For every application both configurations are workable. As a general guideline, byte-level transfer is preferable when few PHYs are provided within a single IC, while packet level transfer is better suited for applications where the number of PHY is large and they are implemented using separate IC's.

NOTES:

PMC-971147(R5) Issue date: December, 1998.

Printed in Canada